Issue N. 27 - 2024

ARGO

New Frontiers in Practical Risk Management

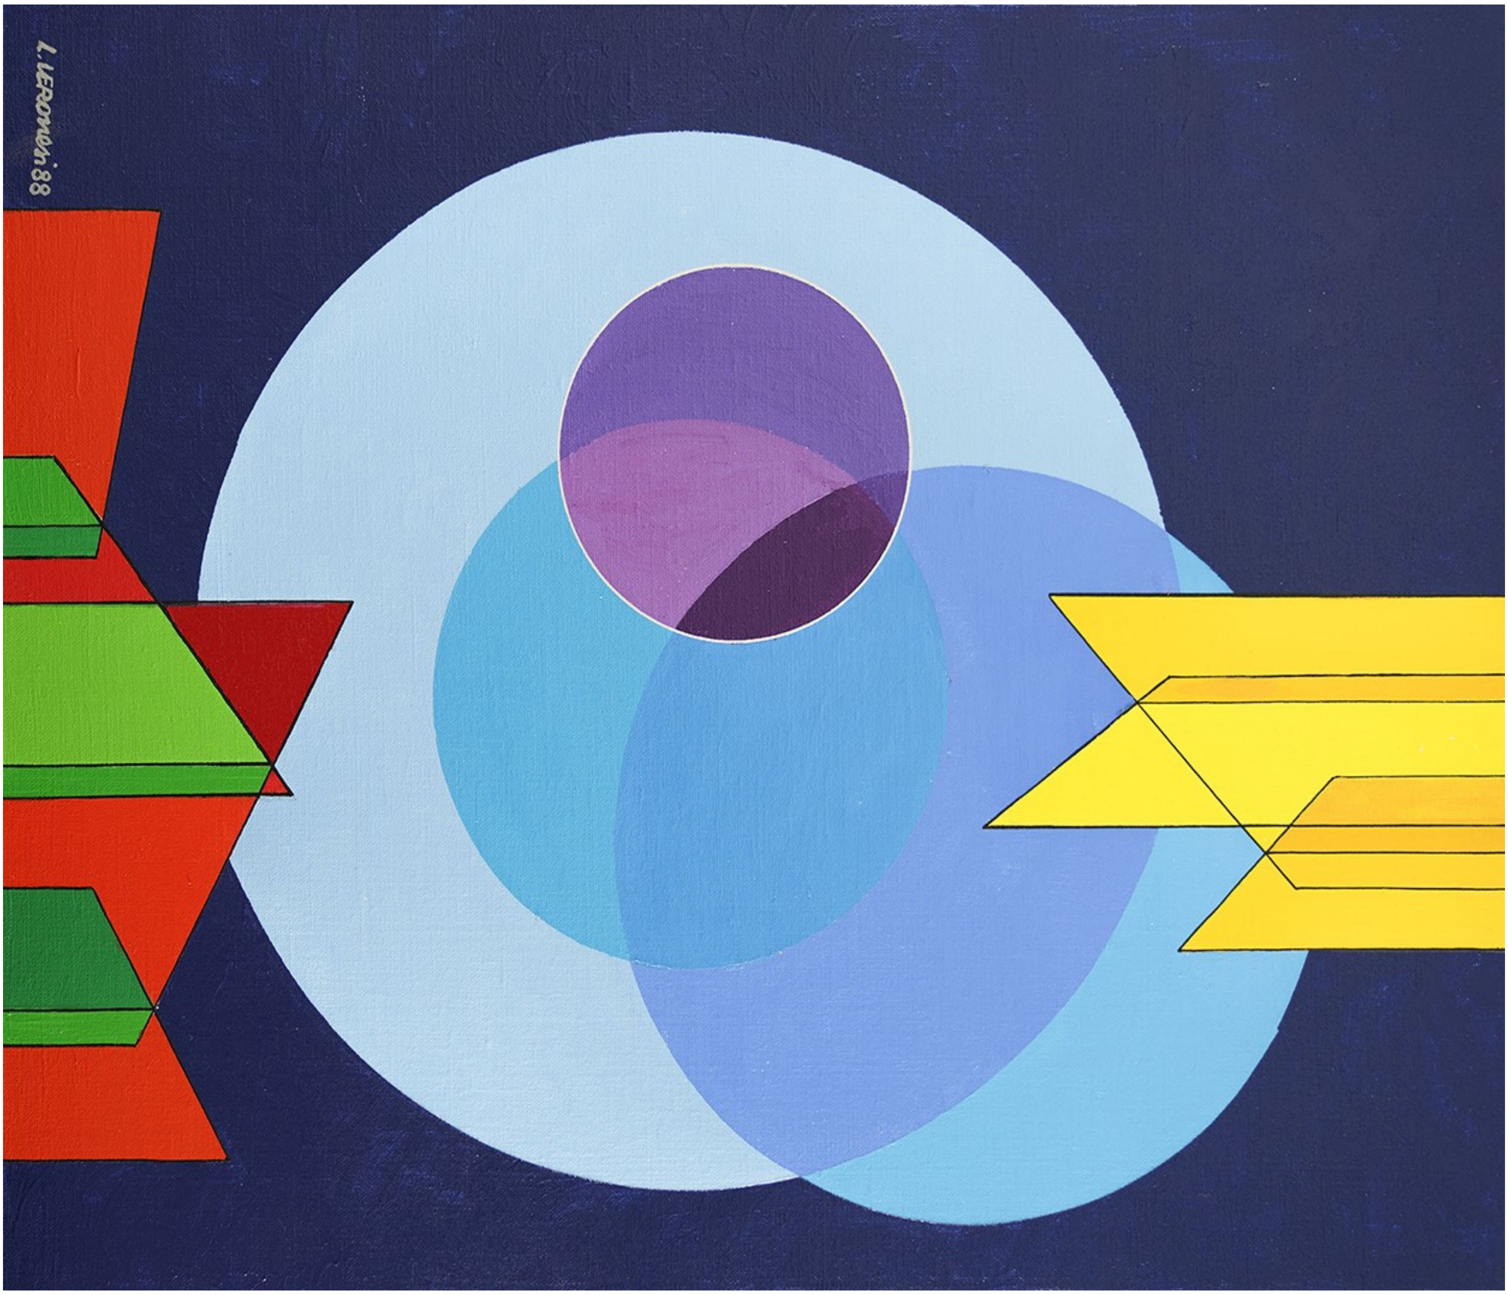

iason

ESSENTIAL SERVICES FOR
FINANCIAL INSTITUTIONS

iason

Argo magazine

Year 2024 - Issue Number 27

Published in December 2024
First published in October 2013

Last published issues are available online:
http://www.iasonltd.com

Front Cover: **Luigi Veronesi**, *Composizione PH3*, 1988.

ESSENTIAL SERVICES FOR
FINANCIAL INSTITUTIONS

Milan Headquarter:
Corso Europa, 15
20122 Milan
Italy

London Headquarter:
3rd Floor, 120 Baker Street
W1U 6TU London
United Kingdom

Madrid Headquarter:
Calle Miguel Ángel, 16
28010 Madrid
Spain

Contact Information:
info@iasonltd.eu
www.iasonltd.com

Articles submission guidelines
Argo welcomes the submission of articles on topical subjects related to the risk management. The articles can be indicatively, but not exhaustively, related to models and methodologies for market, credit, liquidity risk management, valuation of derivatives, asset management, trading strategies, statistical analysis of market data and technology in the financial industry. All articles should contain references to previous literature. The primary criteria for publishing a paper are its quality and importance to the field of finance, without undue regard to its technical difficulty. Argo is a single blind refereed magazine: articles are sent with author details to the Scientific Committee for peer review. The first editorial decision is rendered at the latest within 60 days after receipt of the submission. The author(s) may be requested to revise the article. The editors decide to reject or accept the submitted article. Submissions should be sent to the technical team (**info@iasonltd.eu**). LaTeX or Word are the preferred format, but PDFs are accepted if submitted with LaTeX code or a Word file of the text. There is no maximum limit, but recommended length is about 4,000 words. If needed, for editing considerations, the technical team may ask the author(s) to cut the article.

Table of Contents

Editorial p.6

Just in Time - iason Notes p.8

Asset Management

Comparative Analysis of Portfolio Performance: a CVaR-Based Approach with and without Cryptocurrency Allocation
G. Colombo, S. Donadoni, A. Prati and L. Ratibondi

About the Authors p.15
Literature Review p.18
Crypto Overview p.18
Research Method p.21
Empirical Results p.25
Conclusion p.31
Reference p.32

Portfolio Optimization: a New Frontier through VAE-LSTM-Based Reinforcement Learning
G. Colombo, G. Crupi, G. Giannoni and R. Paradiso

About the Authors p.35
Literature Review p.38
Theoretical Models p.38
Methodologies p.43
Results p.66
Conclusion and Future Works p.71
Reference p.73

Systematic Machine Learning Asset Allocation Benchmarking
E. Veksin

About the Authors p.76
Introduction: Background and Client Requests p.77
Database Setup p.79
Feature Engineering and Unsupervised Learning p.79
Feature Evaluation p.82
Model Selection p.82
Model Backtesting and Portfolio Implementation p.84
Summary p.89
Reference p.91
Technical Appendix p.92

Credit Risk

Computation of RWAs for Securitisation Exposures: the iSEC Calculation Engine
G. Morisani, M. Bainotti and C. Papetti

About the Authors p.102
Risk Weighted Assets Calculation p.105
Seniority Reclassification p.111
iSEC Calculation Engine p.111
References p.113
Annex p.114

Insurance Risk

Ensuring the Future: the Potential Evolution of the Insurance Market

V. Ciminelli and J. Figuriello

About the Authors	p.118
Insurance Foundations	p.119
Digital Ledger Technology	p.124
Decentralized Finance (DeFi)	p.127
Open Insurance	p.129
Market Applications	p.131
Integrating Insurance DeFis and Open Insurance	p.137
Conclusions	p.139
Reference	p.141

Dear Readers,

Welcome to the Winter 2024 edition of Argo Magazine. This edition features a selection of cutting-edge research and innovative approaches that are shaping the future of the banking and insurance sectors.

This edition brings together a series of compelling papers that explore some of the most significant and innovative topics in finance today. These include advanced risk management techniques in asset allocation and the transformative integration of machine learning, blockchain, and other digital innovations in financial and insurance practices.

*In the Asset Management section, we feature a paper entitled "**Comparative Analysis of Portfolio Performance: a CVaR-Based Approach with and without Cryptocurrency Allocation**" by Giacomo Colombo et al. This paper explores the role of cryptocurrencies in modern portfolio management. By examining the impact of digital assets on portfolio performance through Conditional Value at Risk (CVaR), they assess how cryptocurrencies affect the risk-return balance in diversified portfolios. The paper concludes that, while cryptocurrencies can enhance returns, their volatile nature necessitates careful consideration and strategic asset selection to manage potential risks effectively.*

*The following article, entitled "**Portfolio Optimization: a New Frontier through VAE-LSTM-Based Reinforcement Learning**", authored by G. Colombo et al., presents a pioneering approach to portfolio management.*
The authors present a dynamic method that combines Variational Autoencoders (VAE) with Long Short-Term Memory (LSTM) networks and reinforcement learning techniques. This innovative model offers a more adaptive and effective strategy for optimising asset allocation, with superior performance compared to traditional portfolio models.
By leveraging the capabilities of machine learning, this paper underscores the potential for AI-driven methodologies to transform investment management practices.

*The chapter concludes with "**Systematic Machine Learning Asset Allocation Benchmarking**" by P. Bortolotti and E. Veksin, which examines the implementation of a systematic investment process powered by machine learning signals. The paper demonstrates how this framework can be applied not only to direct investments but also as a benchmark for evaluating the performance of traditional investment portfolios. By combining unsupervised learning with supervised regression algorithms, the authors illustrate how systematic approaches can deliver* *superior risk-adjusted returns, even when faced with diverse investment constraints.*

*In the Credit Risk section, we present "**Computation of RWAs for Securitisation Exposures: the iSEC Calculation Engine**" by Giuseppe Morisani, Mattia Bainotti, and Caterina Papetti. This paper addresses the complexities of calculating Risk-Weighted Assets (RWAs) for securitisation exposures within the current regulatory framework. The authors introduce iSEC, a Python-based tool developed by iason, which provides an efficient and automated solution for accurately computing capital requirements related to securitisation exposures. By analysing existing methodologies, the paper demonstrates how iSEC streamlines the computation process for the SEC-SA, SEC-ERBA, and SEC-IRBA approaches, offering risk managers a flexible and practical tool for navigating today's regulatory landscape.*

*Finally, in the Insurance section, we present "**Ensuring the Future: the Potential Evolution of the Insurance Market**", a paper by V. Ciminelli and J. Figuriello which explores the transformative potential of digital innovations such as asset tokenisation, decentralised finance (DeFi) and open insurance. The paper assesses how these emerging technologies are set to transform the insurance industry, increasing efficiency, improving accessibility and creating new opportunities for risk management.*

By examining the convergence of digital innovation and traditional insurance practices, the authors present a strategic vision for the future of the sector in an increasingly digital and decentralised world.

We hope you find these articles both inspiring and informative, providing valuable insights into the evolving landscape of finance and insurance.

We invite you to engage with these topics, share your thoughts, and continue the discussion with us. As always, we look forward to bringing you the latest in advanced research and practice in the fields of finance and technology.

Antonio Castagna
Luca Olivo
Giulia Perfetti

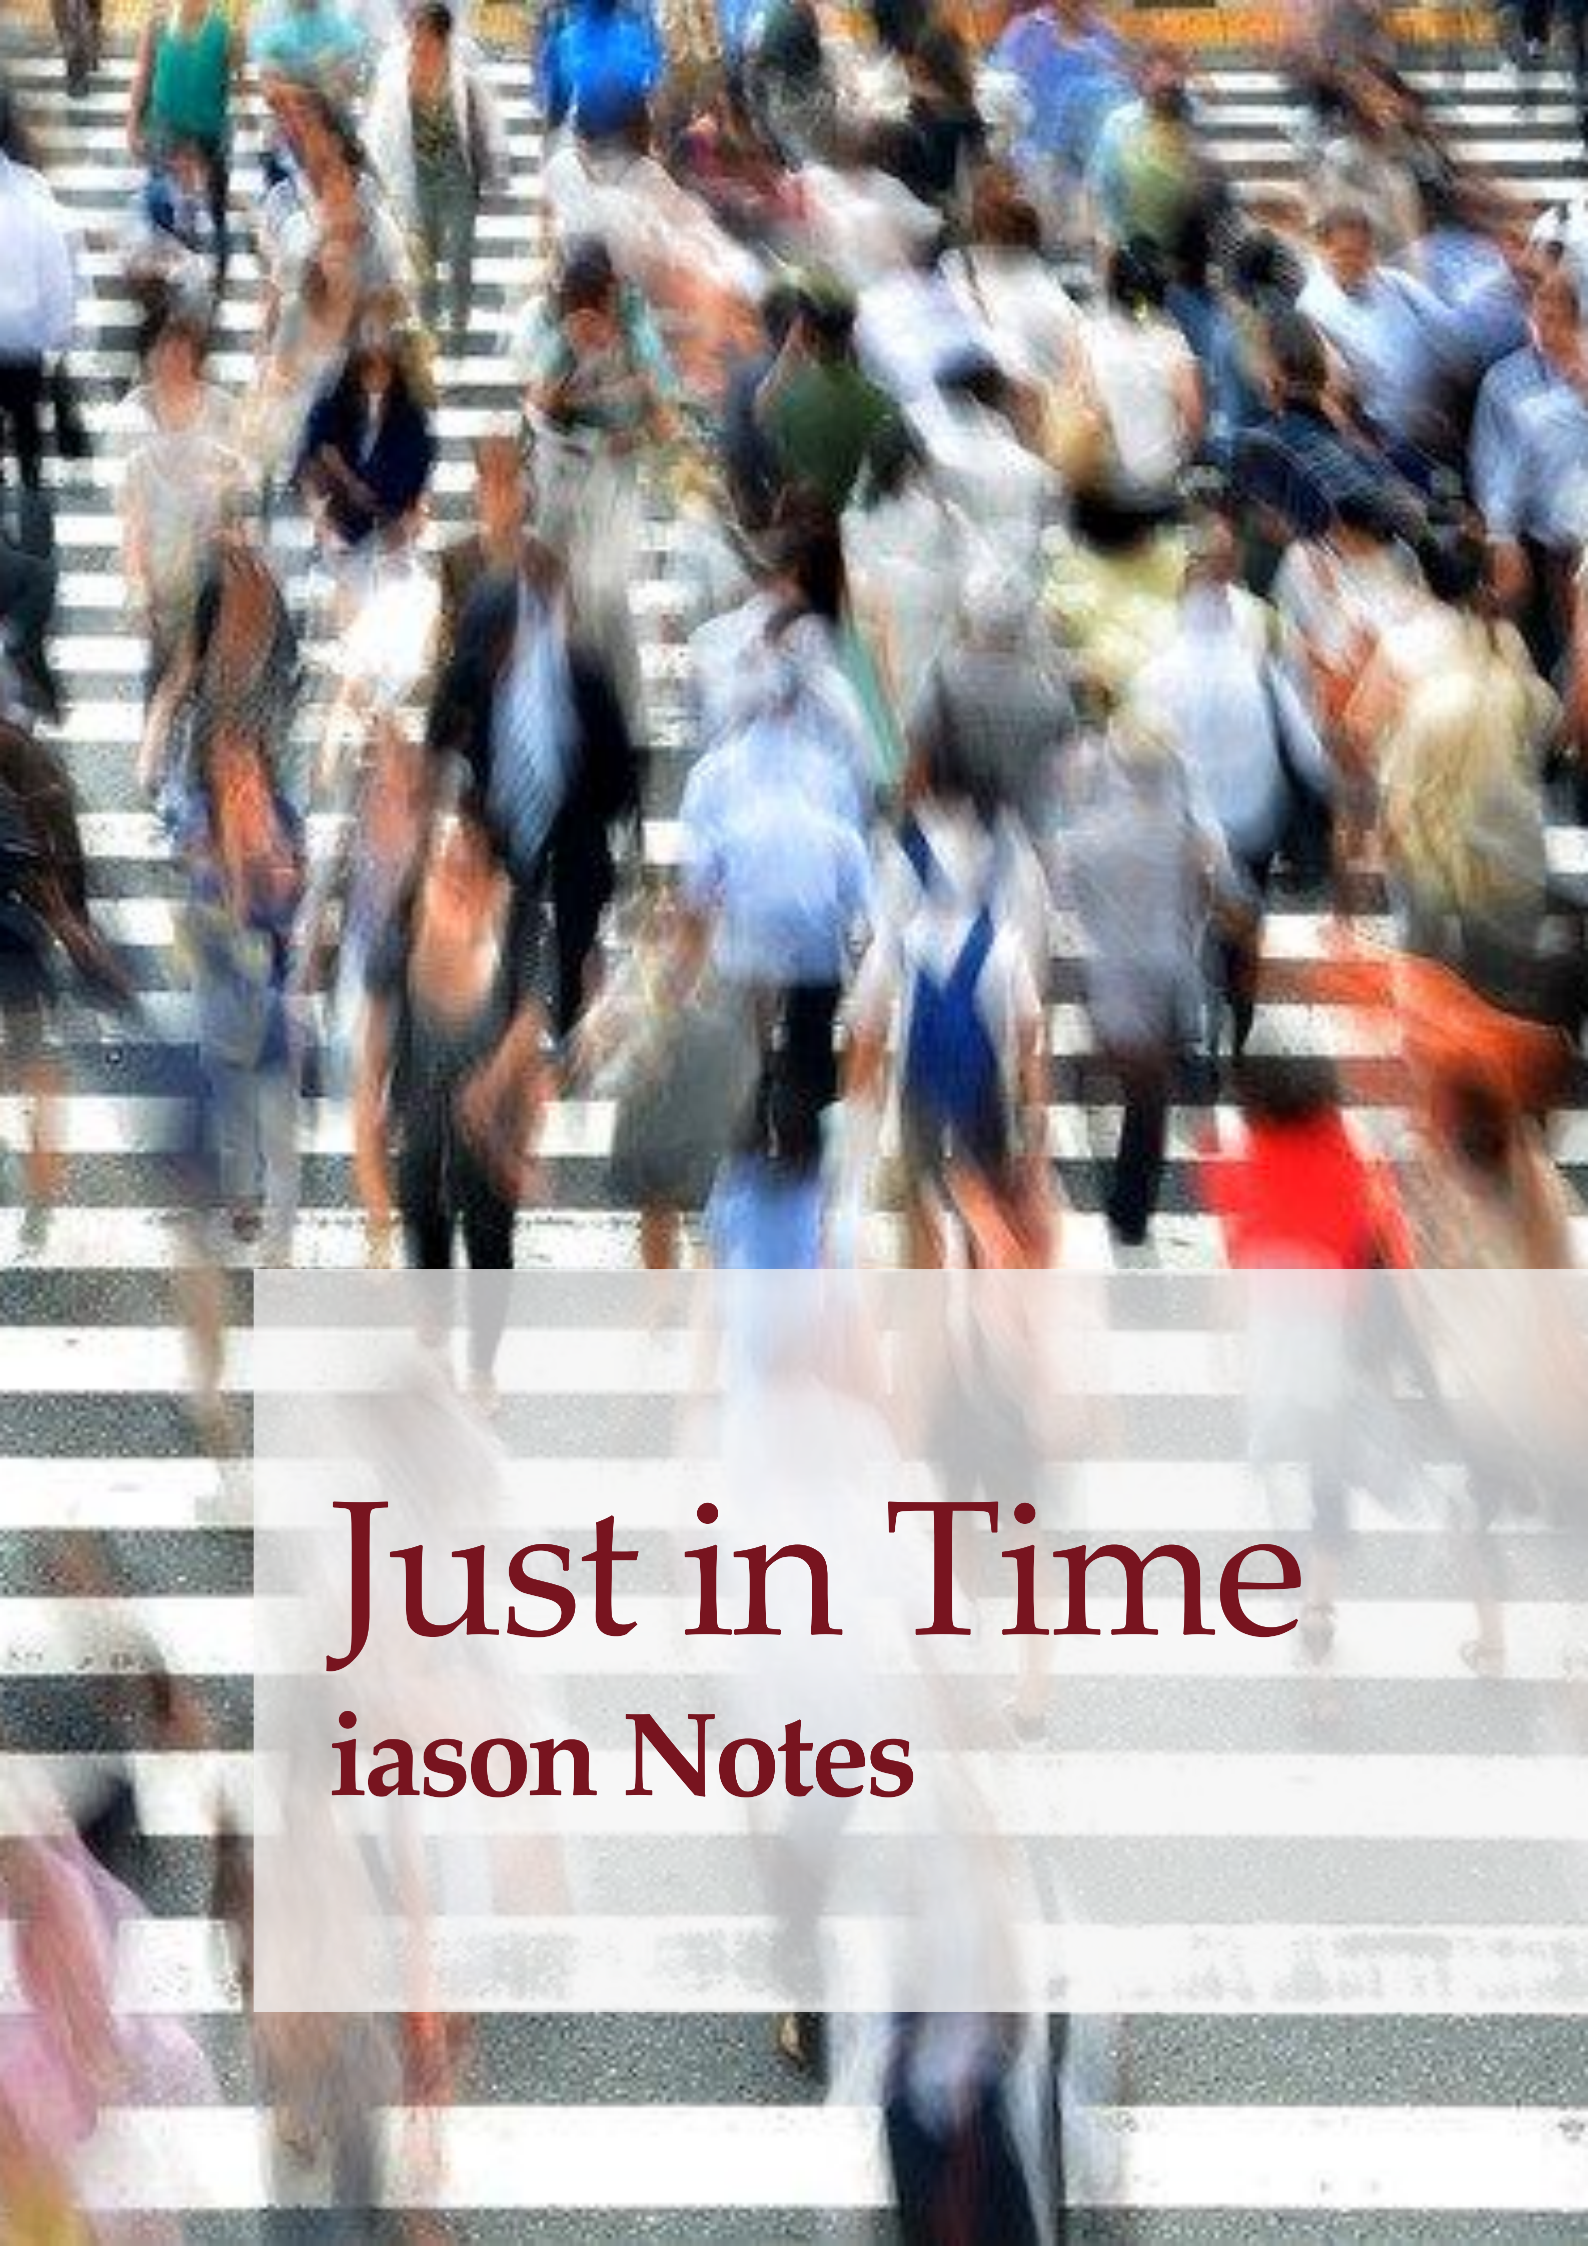

Just in Time

iason Notes

Nobel Prize 2024 in Physics Assigned for AI and ML Contributions

The Nobel Prize in Physics 2024 was awarded jointly to John J. Hopfield and Geoffrey E. Hinton "for foundational discoveries and inventions that enable machine learning with artificial neural networks". This assignation triggered a large debate in the social media (both specialized and general purpose) about this choice. Do ML, AI, Artificial Neural Networks actually match what we usually mean by "physics"? In the following, some details of the technical topics and a review of the different perspectives. We remember that G. Hinton left Google in 2023 after 10 years: he said he quit to speak freely about the dangers of AI.

read more

Date November 2024

The ESAs: Work Programme for 2025

The 2025 Work Programme of the Joint Committee of European Supervisory Authorities (ESAs) focuses on coordination initiatives between the main EU financial regulators.

The document identifies the main priorities of the year, aimed at improving financial stability and consumer protection in a context of growing economic and geopolitical risks.

A deep dive has been given to the EBA's work programme. EBA adopted some priorities for a three-year horizon: EU Single Rulebook, financial stability, data, DORA oversight and MiCAR supervision, as well as conduct and AML/CFT.

read more

Date November 2024

Stress Test 2025

The European Banking Authority (EBA) published its draft methodology, templates, and guidance for the 2025 EU-wide stress test. This phase initiates discussions with the banking sector, enhancing the 2023 exercise's methodology with updates based on recent insights and regulatory developments. Significant changes are being introduced, particularly the incorporation of the forthcoming Capital Requirements Regulation (CRR3), which will take effect on January 1, 2025. In this document we analyze the main changes with respect to the methodology of the 2023 exercise, with a breakdown based on risk type and highlighting the major challenges that the institutions should face in the upcoming exercise.

read more

Date October 2024

Accelerating Settlement: Potential Impacts of T+1 on European Settlement Cycle

Research on post-trade settlement cycles shows the crucial role that efficient and timely processes have in reducing systemic risks and creating resilient financial markets. Reduction of counterparty risk, increase in liquidity, and decrease in failure rates can all be achieved with shorter settlement cycles. Considering this, in May 2024 US and Canada adopted a T+1 post-trade cycle, shifting from the classic T+2 framework, with several benefits in terms of reduction of failure rates and margin requirements. Following these examples, other countries, for instance, the UK and Switzerland, have started to consider the beneficial impacts of adopting a shorter post-trade cycle that also aligns with the US cycle

read more

Date December 2024

Loss-Given-Default and Macroeconomic Conditions

In July 2024, the European Central Bank (ECB) published a working paper analyzing the sensitivity of the realized loss given default (LGD) to macroeconomic conditions. The study utilized Global Credit Data's confidential dataset on cash flows from defaulted loans to address three key research questions: i) does LGD increase during adverse macroeconomic conditions?; ii) does the timing of cash flows influence the relationship between LGD and macroeconomic conditions?; iii) is the sensitivity of LGD to macroeconomic conditions more significant for secured loans than for unsecure loans?.

read more

Date October 2024

Peer Review Report Guidelines on the Application of the Definition of Default

In July 2024, the European Banking Authority (EBA) published a Peer Review on its Guidelines on the application of the definition of default, outlining the findings of six Competent Authorities' supervision of credit risk, focusing on application of the definition of default and the EBA Guidelines across three major areas: implementation of EBA/GL/2016/07 in the supervisory framework; effectiveness of the procedure for the submission of the application; effectiveness of the assessment for checking compliance with the definition of default.

read more

Date September 2024

EIOPA: Comparative Study on Market and Credit Risk Modelling

Market and credit risk significantly impact the solvency capital requirement (SCR) of insurance companies and are crucial for most internal model undertakings. As a result, in early 2018, the EIOPA Board of Supervisors decided to conduct annual European-wide comparative studies on market and credit risk modeling. These studies are managed by a joint project group comprising National Competent Authorities (NCAs) and EIOPA. Companies with substantial exposure to Euro-denominated assets and an approved internal model for market and credit risk are required to participate in this annual study.

read more

Date September 2024

The Impact of ECB Banking Supervision on Climate Risk and Sustainable Finance

"The paper explores the effects of the ECB's supervisory actions on climate risk management and sustainable finance among banks, in particular, the introduction of the climate-risk-related supervisory efforts since 2020 to enhance banks' awareness and preparedness for managing climate-related risks. The first analysis are included in the «Guide on C&E Risks» and in the first climate risk stress test in 2022 where aims to assess the impact of these supervisory efforts on banks' climate risk exposure and management, as well as their green finance activities. The data highlights that these improvements are robust across various control variables and fixed effects, confirming the effectiveness of ECB's supervisory measures.

read more

Date September 2024

Principles for the Sound Management of Third-party Risk

Banks have long relied on third-party service provider (TPSP) arrangements for a variety of reasons. The "Principles for Sound Third-Party Risk Management" is the consultative (end of consultation by 9th October 2024) document published by the Basel Committee on Banking Supervision (BCBS) provides a common guidance for effectively managing the risk associated with third-party service. The consultative document sets out 12 principles aimed at improving banks' operational resilience, reducing operational fragmentation, and promoting international collaboration.

read more

Date September 2024

iason Weekly Insights

Regulatory/Supervisory Pills

Among iason's various publications we also find the iason Pills.

With these daily Pills, iason aims to offer a summary on information, mostly, of the main regulatory and supervisory news in the banking and finance sector on both Pillar I and Pillar II risks of the Basel framework. The main purpose of these publications is to give the reader an effective, timely and brief overview of the main topics of the moment.

The authors of the Iason Pills are Dario Esposito and Cecchin Matteo.
read more

Market View

Among iason's weekly insight you can also find the iason Market View, a weekly update on financial market by Sergio Grasso.

The author, with almost three decades of investment experience, presents an accurate analysis of market fluctuations of the week, giving a critical view of observed phenomenos and suggesting interesting correlations with the main world events.
read more

Comparative Analysis of Portfolio Performance: a CVaR-Based Approach with and without Cryptocurrency Allocation

About the Authors

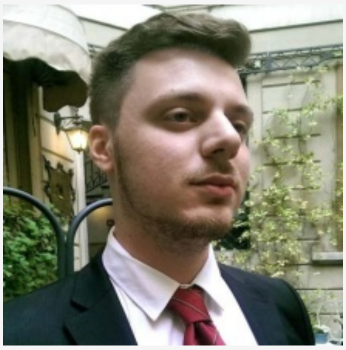

Giacomo Colombo:
Senior Manager
He holds a master's degree in Economics and Finance, with a specialisation in Corporate Finance. After two years in the M&A sector, he moved on to Banking field, with particular focus on Market contents. He is highly specialised in data modelling, methodology and automated computations. In his current role as Senior Manager, he oversees several tasks at the largest Italian Banks and follows numerous Asset Management streams, also collaborating with an important University.

Silvio Donatoni:
Developer
He obtained a Bachelor's degree in Computer Science and a Master's degree in Data Science at Verona University. He carried out various work experiences in the field of computer science, acquiring cross-disciplinary skills from front-end to back-end development, as well as in machine learning and data analysis. He joined iason in March 2023, driven by his growing interest in the financial-economic sector in addition to computer science and data science. At Iason, he presented his Master's thesis titled "Expected Shortfall Forecasting Models & Backtesting Methods".

Alessandro Prati:
Senior Business Analyst
He holds a MSc in Economic and Finance obtained at the University of Milano-Bicocca. He joined iason in 2021 working on projects mainly focused on Market Risk. In the last year, together with his work team he has mainly carried out activities relating to the development and improvement of data models to support regulatory exercises.

Luigi Ratibondi:
Business Analyst
He holds a MSc in Quantitative Finance & Insurance obtained at the University of Torino. He joined Iason in 2023 working on projects mainly focused on Financial & Market Risk. Since late 2023, he has worked with the team in the data collection and testing process on several projects.

Comparative Analysis of Portfolio Performance: a CVaR-Based Approach with and without Cryptocurrency Allocation

Giacomo Colombo Silvio Donatoni Alessandro Prati
Luigi Ratibondi

This study aims to examine the utility of cryptocurrencies as a portfolio enhancement tool, thus offering valuable analysis for those wishing to introduce cryptocurrencies into their investments. The analysis takes as a starting point an already optimized portfolio, to which different cryptocurrencies are added and tested using key metrics for evaluating a portfolio.
The main metric considered is Conditional Value at Risk (CVaR), which, due to the volatile nature of cryptocurrencies, is a particularly important metric for this thesis. The use of CVaR makes it possible to measure the risk of extreme losses, providing a more comprehensive assessment of risk than other traditional metrics such as simple standard deviation or Value at Risk (VaR).
The objective is to determine whether and how the introduction of cryptocurrencies may affect the overall performance of an optimized portfolio. Various investment scenarios, including different combinations of cryptocurrencies, will be analyzed to assess the impact on parameters such as Average Return, Annualized Average Return, Standard Deviation, Annualized Standard Deviation, VaR, CVaR, Sharpe Ratio, Sortino Ratio, Maximum Drawdown and Probability of Loss. In this way, the research aims to provide practical and theoretical guidance for investors interested in diversifying their portfolios with digital assets.
In addition, the study will explore the differences among major cryptocurrencies, assessing which one offers the best risk-return ratio. Through in-depth analysis, the distinctive characteristics of the most promising cryptocurrencies will be identified, helping investors make informed decisions.
This research is a further contribution to the investment field, as it provides a rigorous and detailed assessment of the role of cryptocurrencies in modern portfolios, using CVaR as the main risk measurement tool.
The results show that the integration of cryptocurrencies into diversified portfolios represents a promising opportunity which allows for benefits, particularly in returns, but requires careful attention to asset selection, time horizon, and market context, thereby balancing potential returns with risk management.

The cryptocurrency market has exceeded a capitalization of 2 trillion dollars, becoming a new asset class that can no longer be disregarded in financial markets. Despite this, many institutional investors remain cautious due to the inherent risks associated with this market. This study aims to evaluate how the inclusion of cryptocurrencies impacts a benchmark institutional portfolio, particularly in terms of risk-return dynamics and market diversification.
The hypothesis underlying this research is that, by maintaining portfolio risk at or slightly above the benchmark level with the inclusion of cryptocurrencies, it is possible to achieve higher returns, provided that investments are made during an appropriate phase of the cryptocurrency market, given its cyclical nature.
To accomplish this, five cryptocurrencies were selected from the top ten by market capitalization: Bitcoin (BTC), Ethereum (ETH), Binance Coin (BNB), Ripple (XRP), and Cardano (ADA). These cryptocurrencies were chosen based on their higher market capitalization and lower volatility, while recently introduced cryptocurrencies such as Solana, Avalanche, and Toncoin were excluded due to insufficient historical data.
Starting from a professionally constructed benchmark portfolio, two additional portfolios were generated. The first portfolio is composed by equally weighted instruments, without any optimization logic. The second one is designed in order to minimize the 95% Conditional Value at Risk (CVaR), allowing for an unconstrained reallocation of asset weights. Each of these three portfolios was evaluated both with and without the inclusion of cryptocurrencies, and several performance, risk, and risk-return metrics were computed. These metrics include daily and annual returns, standard deviation, Value at Risk (VaR) and CVaR at 95% and 99%, maximum drawdown, probability of loss, and risk-return ratios such as the Sharpe and Sortino ratios.
The historical data available for cryptocurrencies begins in early 2018, with the exception of Bitcoin, for which data is available from late 2014. As such, the analysis covers the period from June 2018 to June 2024. Furthermore, the study segments the analysis into three well-known cryptocurrency market phases-bull markets, bear markets, and recovery markets-related to Bitcoin's halving cycles. Two instances of each market phase are considered.

The scenario analysis is crucial for understanding returns. Empirical evidence suggests that while portfolio risk remains relatively stable across different market phases, the inclusion of cryptocurrencies during a bull market significantly enhances returns. In contrast, returns decline sharply during bear markets, while recovery markets yield moderate return increases. Over the entire period, portfolio returns rise, reflecting the strong growth experienced by the cryptocurrency market in recent years.
Risk-return indicators such as the Sharpe and Sortino ratios improve with the inclusion of cryptocurrencies, confirming the benefits of adding this asset class to institutional portfolios. However, the portfolio optimized to minimize CVaR, tends to either exclude cryptocurrencies or allocate them in minimal proportions due to their high volatility, which often results in fat-tailed distributions in profit and loss outcomes. Nevertheless, Bitcoin, being the most capitalized and least volatile cryptocurrency, is assigned a higher proportion in these portfolios, reflecting the expectation that, as its capitalization continues to grow, its volatility will stabilize. A correlation analysis with the Eurostoxx 50 index[1] reveals that the inclusion of cryptocurrencies reduces the portfolio's correlation with traditional markets. The higher the proportion of cryptocurrencies in the portfolio, the lower its correlation with the index.
This document is structured as follows. The literature review provides an overview of existing research on asset management, with a particular emphasis on the inclusion of cryptocurrencies in institutional investment portfolios. The second section presents a detailed analysis of the cryptocurrency market, focusing on key factors such as market capitalization, trading volumes, volatility, correlations, and the significance of Bitcoin's halving cycle. The methodology section outlines the risk-return metrics used to evaluate the three portfolios. The empirical results section presents the outcomes of the backtesting performed on the three portfolios over the entire period and during different cryptocurrency market phases, analyzing various metrics commonly used in the financial sector to assess the impact of cryptocurrencies on portfolio risk, return, and diversification. Finally, the conclusion section summarizes the findings and provides recommendations for investors, emphasizing the appropriate timeframes and factors to consider when including cryptocurrencies in an institutional investment portfolio.

Literature Review

Over the past decade, cryptocurrencies have garnered significant attention not only as speculative investment assets but also as potential components of a well-diversified portfolio. A growing body of literature has investigated the distinct risk-return profiles of cryptocurrencies and their correlation with traditional asset classes, yielding insights into their potential role in portfolio diversification strategies. Notably, studies by Corbet et al. (201) [7] and Baur et al. (2018) [4] have highlighted that cryptocurrencies tend to exhibit low to negative correlations with traditional financial assets such as equities and bonds, which enhances their attractiveness as diversification tools. The findings of Huang et al. (2022) [13] further support this view, demonstrating that Bitcoin, despite its pronounced volatility, provides considerable diversification benefits, especially during periods of economic uncertainty.
In a more focused analysis, Będowska-Sójka, B., & Kliber, A. (2021) [5] examined the quantile coherency between the S&P 500 Index and two major cryptocurrencies, Bitcoin and Ether, to assess their suitability as hedge or safe-haven assets. Kliber's findings suggest that both Bitcoin and Ether may indeed serve as safe-haven assets, though this characteristic fluctuates over time and is influenced by the investment horizon in question. These results indicate that while cryptocurrencies may offer hedging potential, this quality is neither consistent nor universal.
On the other hand, a study by Goodell and Goutte (2021) [11] employed a range of econometric methodologies, including neural network analyses, to evaluate the co-movements of four major cryptocurrencies with seven global equity indices, both before and during the COVID-19 pandemic. Their findings revealed that correlations between cryptocurrencies and equity indices gradually increased over time, particularly during the pandemic, thus challenging the notion that cryptocurrencies can reliably act as safe havens. Ji et al. (2020) [15] got similar conclusions, further questioning the consistency of cryptocurrencies as protective assets during periods of financial stress.
Another critical consideration in the realm of risk management is the behavior of cryptocurrencies during episodes of market stress. Dyhrberg (2016) [9] explored this issue by comparing Bitcoin to traditional safe-haven assets such as gold, positing that Bitcoin's decentralized nature could render it an effective hedge against systemic risk. However, the context-specific nature of this finding is emphasized by Fang et al. (2019) [10], who demonstrated that the risk-hedging properties of Bitcoin and other cryptocurrencies vary significantly depending on broader global economic conditions and policy-related uncertainty.
Empirical research on the performance of investment portfolios that include allocations to cryptocurrencies has produced mixed outcomes. Guesmi et al. (2019) [12] found that portfolios with small allocations to Bitcoin (typically between 1% and 5% of the total asset mix) outperformed those without such allocations in terms of risk-adjusted returns. This enhanced performance is attributed to the ability of small cryptocurrency exposures to capture potential upside while mitigating the adverse effects of high volatility. In a similar vein, Klein et al. (2018) argue that Bitcoin's low correlation with traditional assets, such as stocks and bonds, enhances overall portfolio performance, particularly during periods of economic stress, as reflected in improved Sharpe ratio. However, other studies, such as Almeida and Gonçalves (2023) [2], underscored the extreme volatility of cryptocurrencies, noting that while such volatility can bolster portfolio performance under favorable conditions, it can also severely detract from returns during market downturns.

In summary, while the inclusion of cryptocurrencies in investment portfolios can potentially enhances diversification and improves risk-adjusted returns, particularly during economic turmoil, the high volatility and fluctuating correlation with traditional assets make their role in portfolio management both complex and context-dependent. The suitability of cryptocurrencies as hedges or safe-haven assets remains a contentious issue, with their effectiveness varying widely based on market conditions, time horizons, and broader economic dynamics.

Crypto Overview

Cryptocurrency market has experienced exponential growth over the past decade, transforming from a small niche market of 500 Million dollars in 2016 into a significant asset class of 2.5 Trillion dollars in the global financial ecosystem today. Since the bull market of 2020-2021, many institutional

[1] The Eurostoxx 50 index was chosen since most of the assets in the benchmark portfolio are European.

investors have started to approach this emerging market and also the recent introduction of spot ETFs on Bitcoin and Ethereum allows them to enter it more easily and safely. This section provides a generic overview of the cryptocurrency market, including key metrics and drivers influencing the inclusion of cryptocurrencies in investment portfolios.

Market Capitalization & Trading Volumes

Market capitalization (market cap) is a crucial metric used to gauge the size and stability of a cryptocurrency. It is calculated by multiplying the current price of the cryptocurrency by its total circulating supply. Trade volume is another critical metric indicating the liquidity and market activity of a cryptocurrency. High trade volumes suggest a more liquid market, enabling investors to enter and exit positions with minimal price slippage. Liquidity is mandatory for institutional investors who require substantial trade capacities without significantly impacting the market price.

As of 2024, Bitcoin (BTC) remains the dominant cryptocurrency, considered the father of the entire market, with a dominance around 55% of the overall market, while Ethereum (ETH) is the second biggest crypto covering around 17% of the market dominance.

Volatility

Cryptocurrencies are known for their high volatility compared to traditional financial assets. Volatility measures the degree of variation in a cryptocurrency's price over time.

This inherent volatility presents both opportunities and risks for investors. High volatility can lead to substantial gains but also significant losses, making risk management strategies like Conditional Value at Risk (CVaR) essential when incorporating cryptocurrencies into a portfolio.
Periods of high volatility in the crypto market have historically occurred and are known as 'bull' markets (2017, 2021) and 'bear' markets (2018, 2022) where there was significant financial speculation. However, due to the continuous growth in market capitalization, cryptocurrencies (or at least the largest ones) will tend over time to flatten this high volatility.
In figure 1 is compared historical 1-year volatility of the top ten cryptos and the ones with a more stable market are choosen (Bitcoin, Ethereum, Ripple, Binance and Cardano). Since Avalanche, Toncoin and Solana was born recently (2020-2021) they have been excluded from this research, finally Doge experienced a high volatility due to its highly speculative nature during bull and bear markets (see 2021-22).

Assets Correlation

Let's consider now the correlation between the Benchmark portfolio with the chosen 5 cryptocurrencies in order to check a possible improvement of the portfolio risk diversification.
From the correlation matrix, it can be observed that the least correlated cryptocurrencies with other assets are Cardano (ADA: 0.39) and Ripple (XRP: 0.41). Bitcoin, on the other hand, is in the middle (BTC: 0.47), while the most correlated are Ethereum (ETH: 0.51) and Binance (BNB: 0.49).

Bull & Bear Markets: The Halving Cycle

As shown in Table 1, Bitcoin is an inflationary cryptocurrency, as it has a maximum supply of 21 million. Additionally, its consensus algorithm is programmed for decreasing inflation, which means that approximately every 4 years (every 210,000 transaction blocks), a 'halving' of the Bitcoin supply occurs. This has happened four times so far [3]:

- **2012**: from 50 to 25 BTC per block reward;
- **2016**: from 25 to 12.5 BTC per block reward;
- **2020**: from 12.5 to 6.25 BTC per block reward;
- **2024**: from 6.25 to 3.125 BTC per block reward.

This algorithm halves the Bitcoins supply production about every four years, causing a shortage of Bitcoins in circulation in a programmed manner that is immutable over time, so anyone can know how many Bitcoins will be created at a certain time in the future, a property that makes the asset unique.
The empirical evidences shows as the halving algorithm has dictated a long-term price cycle (from one halving to the next, approximately 4 years) structured as follows:

- **Bull Market:** in the time period close to halving, investors, being aware of its arrival, accumulate Bitcoin by increasing demand against a halved supply. The price at this point begins to rise starting the bull market. Here the behavioural economics mechanism called fear-of-missing-out (FOMO) is triggered which continues to increase demand pushing the price to new all time highs. This situation leads to an overbought market in the final part of the bull market. Historically, bull markets have lasted about a year and a half.
- **Bear Market:** at the top of the bull market, the market trend is no longer sustainable and investors who have been in the market for longer begin to monetize the gains made during the bull market. Thus the selling pressure begins to increase and causes strong price corrections that lead to a bear market. During this period a situation of fear, uncertainty and

Name	Price	MarketCap	Volume(24h)	Circ.Supply	Max.Supply
Bitcoin(BTC)	$66.429,90	$1.310Bln	$27.7Bln	19.7MlnBTC	21MlnBTC
Ethereum(ETH)	$3.329,98	$400.4Bln	$13.5Bln	120MlnETH	INFETH
Solana(SOL)	$164,55	$87.2Bln	$2.5Bln	466MlnSOL	INFSOL
Binance(BNB)	$586,90	$85.6Bln	$1.8Bln	146MlnBNB	INFBNB
Ripple(XRP)	$0,65	$36.5Bln	$2.4Bln	56BlnXRP	100BlnXRP
Toncoin(TON)	$6.95	$17.5Bln	$700Mln	2.5BlnTON	INFTON
Dogecoin(DOGE)	$0.11	$15.6Bln	$1Bln	145BlnDOGE	INFDOGE
Cardano(ADA)	$0,41	$14.5Bln	$270Mln	36BlnADA	45MlnADA
Avalanche(AVAX)	$27.2	$11Bln	$240Mln	395MlnAVAX	715MlnAVAX
Tron(TRX)	$0.12	$10Bln	$220Mln	87BlnTRX	INFTRON
TotalMarket	N.D.	$2.450Bln	$78Bln	N.D.	

TABLE 1: *Market Overview: Top 10 Cryptocurrencies by Market Cap (06/2024).*

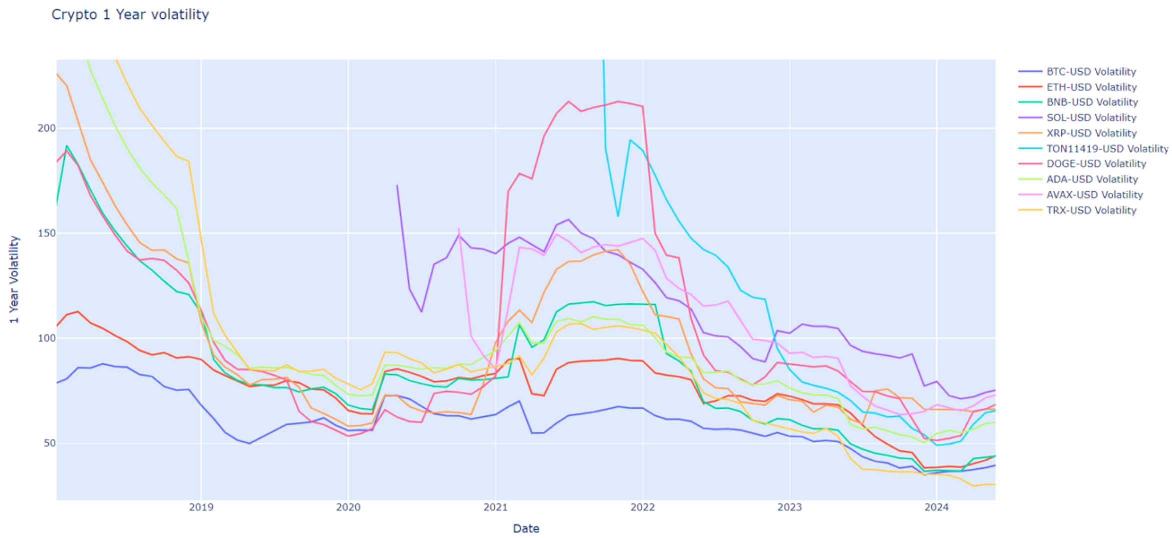

FIGURE 1: Top 10 cryptos yearly volatility compared from 2017 to 2024.

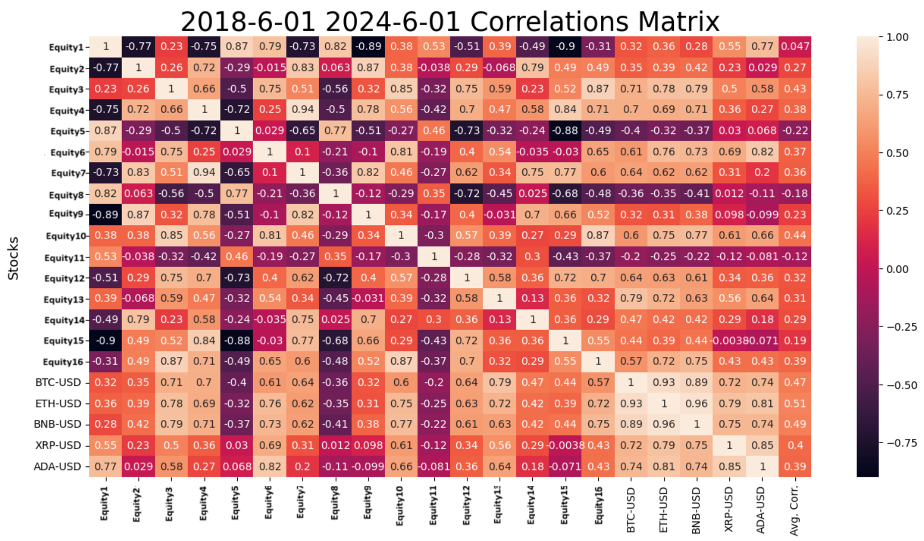

FIGURE 2: 6 years asset correlation between benchmark portfolio and crypto assets.

doubt (FUD) is unleashed dragging the market into an oversold situation and the price to new relative lows. Historically, bear markets have lasted about a year.

- **Recovering Market:** starting from the bottom of the bear market, forward-looking investors see big investment opportunities at very low prices. So in addition to buying back assets, they begin to refinance projects, industry and crypto-related start-ups. This recovering market brings demand and price to a more stable situation. As the new halving approaches, the phase of crypto accumulation and increased demand restart, setting the ideal conditions for a new bull market.

The same cycle pattern was repeated also in 2012-2016, but, since the market was very small and the price very low, it is not clearly visible from this plot.
Accordingly, Bitcoin's price cycle influences the entire crypto market, and the inclusion of cryptocurrencies in the investment portfolio of a financial institution yields risk and return outcomes that are highly dependent on the period and time horizon for which they are held. Therefore, this paper will also distinguish the results based on the three market scenarios listed above.

Crypto & Traditional Markets Correlation

Another phenomenon to be taken into account is the correlation between the traditional stock market and the crypto market. Figure 4 shows a similar breakdown of scenarios for the traditional stock market and Figure 5 shows the correlation between Bitcoin (as the highest representative of the crypto market) and Nasdaq Composite (as the highest representative of the stock market).
From the scenario and rolling correlation analysis, a growing trend emerges in which the crypto market increasingly aligns with the traditional market over time, especially from 2021 onwards, where scenarios become overlapping. This occurs in conjunction with the persistent entry of institutional investors, such as banks and investment funds, into cryptocurrencies and the resulting increase in market capitalization. However, as seen in the Chapter "Assets Correlation", this correlation tends to decrease when considering less capitalized cryptocurrencies (see Ripple and Cardano).

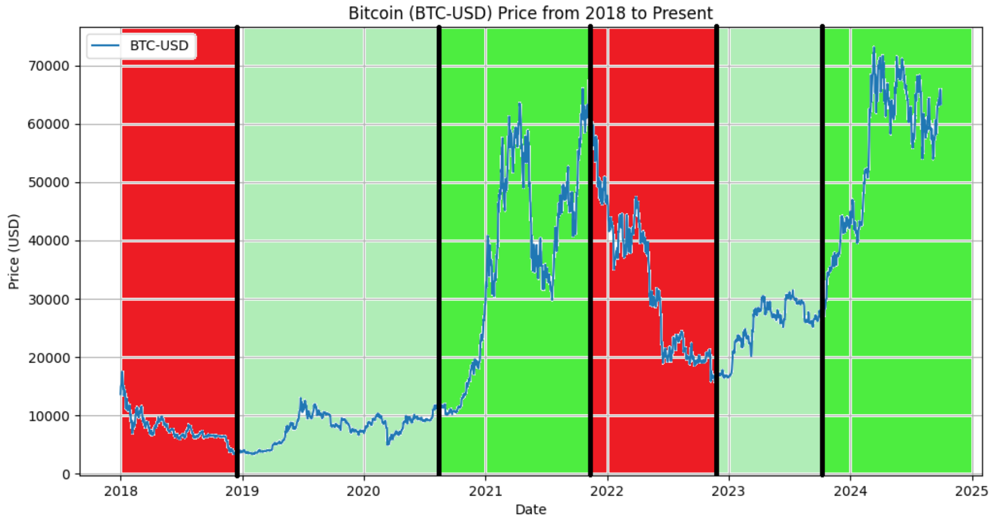

FIGURE 3: Bull and Bear market periods shown from Bitcoin empirical data.

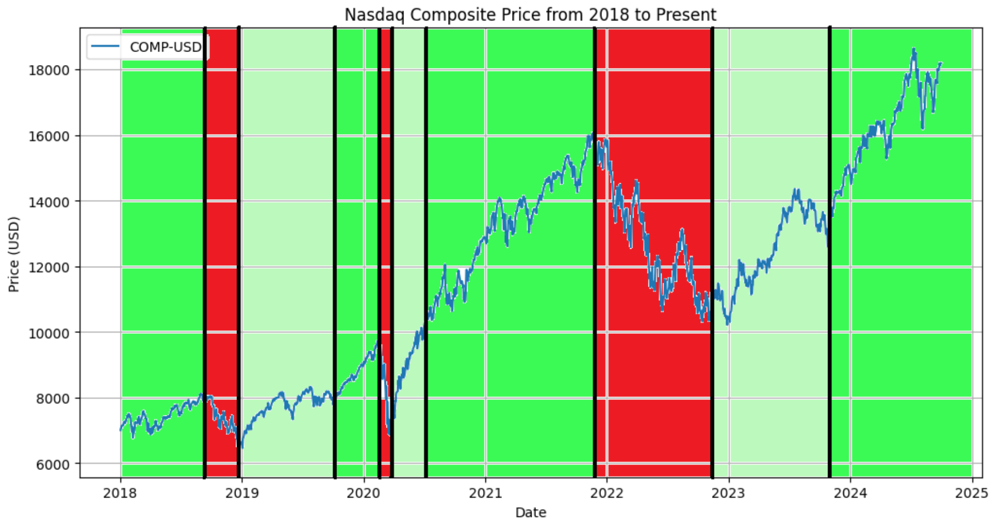

FIGURE 4: Bull and Bear market periods shown from Nasdaq Composite empirical data.

Research Method

To evaluate the inclusion of cryptocurrencies in a traditional investment portfolio, a time-series data collection was conducted on the target portfolio's assets and the cryptocurrencies to be included. Specifically, the window from June 2018 to June 2024 was selected as the time frame of interest for data collection in order to ensure the most comprehensive data set possible and consistent with the availability of time series for cryptocurrencies. Furthermore, taking into consideration the importance of the time horizon and related market conditions of an investment, the analysis was extended to intermediate time windows characterized by significant market conditions for the crypto world.
In fact, as discussed in the previous chapter, the crypto market has been characterized by a cyclical pattern that alternates among bull, bear, and recovering market periods with each halving cycle. Therefore, in order to evaluate the integrated portfolio with cryptocurrencies with the addition of the time horizon as a driver, the following time windows were identified and data extracted:

- **2 Bull Markets**: the first from May 11, 2020 to November 11, 2021 corresponding to the third cycle of Bitcoin halving, and the second from Oct 12, 2023 to June 1, 2024 corresponding to the fourth cycle of Bitcoin halving.
- **2 Bear Markets**: the first from January 1, 2018 to December 31, 2018 corresponding to the second cycle of Bitcoin halving, while the second from November 12, 2021 to November 14, 2022 corresponding to the third cycle of Bitcoin halving.
- **2 Recovering Markets**: the first from January 1, 2019 to May 10, 2020 corresponding to the second cycle of Bitcoin halving, while the second from November 15,

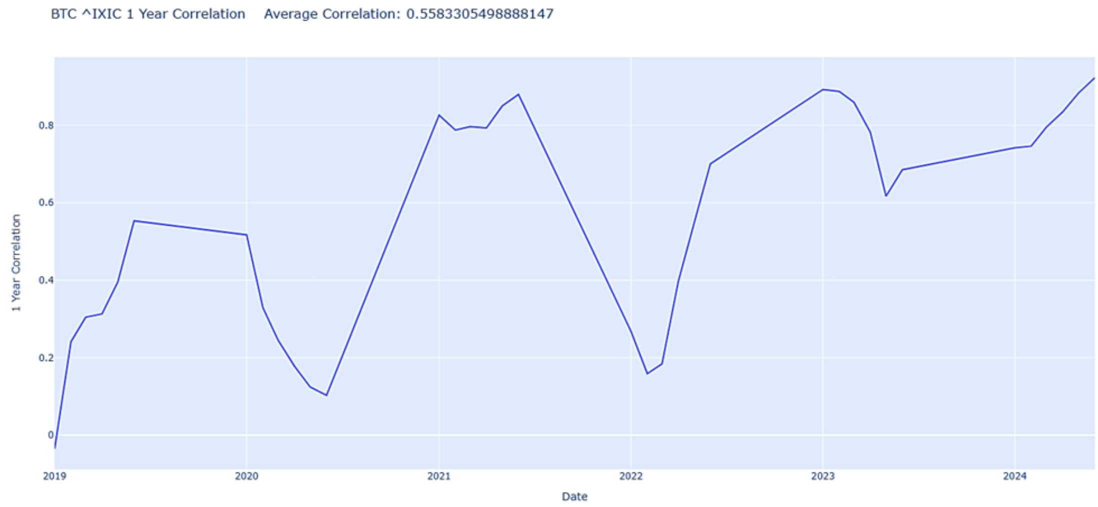

FIGURE 5: *Bitcoin & Nasdaq Composite 1 year rolling correlation.*

2022 to October 11, 2023 corresponding to the third cycle of Bitcoin halving.

In this regard, two additional portfolios were constructed starting from the benchmark one, and these evaluated with and without the inclusion of cryptos for each time horizon considered. This allows for a more detailed comparison and a deeper understanding of how cryptocurrencies can affect portfolio performance. The three portfolios construction methods used in this study are:

- **Portfolio A (Benchmark)**: this portfolio has been created and optimized by industry professionals. It serves as the reference point for the study, representing a well-balanced and professionally managed portfolio.
- **Portfolio B (Equally Weighted)**: in this portfolio, all stocks of the benchmark portfolio are assigned equal weights. The primary utility of this portfolio is to provide a baseline comparison to highlight the quality and performance of the benchmark portfolio. Additionally, it allows for the study of the inclusion of cryptocurrencies in a scenario where no specific optimization has been applied, thereby showing their raw impact.
- **Portfolio C (Minimizing CVaR)**: this portfolio is optimized such that the weights assigned to the stocks minimize the Conditional Value at Risk at 95%(CVaR). The goal is to determine whether or not crypto should be included, and in what percentage, within a portfolio that brings the CVaR value closest to zero, thus focusing on risk management and the reduction of extreme losses.

By comparing these three portfolios, the study aims to provide a thorough evaluation of how the inclusion of cryptocurrencies can enhance or impact portfolio performance. Portfolio A, as the benchmark, offers a standard approach with respect to the other methods. Portfolio B helps to understand the performance of cryptocurrencies when equal weighting is applied to all stocks, giving insight into their impact without any optimization biases. Portfolio C specifically addresses risk management by focusing on minimizing CVaR, offering a perspective on how cryptocurrencies might be utilized to mitigate extreme risks in a portfolio.
The framework is a comprehensive tool for financial analysis, particularly centered on assessing and optimizing portfolio risk metrics. The primary focus is about understanding and managing the risks associated with a portfolio, which is composed by different instruments, including traditional stocks and cryptocurrencies.
To begin with, the analysis involves calculating several important risk measures commonly used in portfolio management. These include the Value at Risk (VaR), a metric that estimates the potential loss in the portfolio's value over a given time period at a specified confidence level. In addition, the Conditional Value at Risk (CVaR) is also computed, representing the average loss that occurs beyond the VaR threshold, which provides a deeper insight into tail risk. Additional measures include the Sharpe ratio, which assesses risk-adjusted returns by comparing excess returns relative to their volatility, and the Sortino ratio, a variant that focuses solely on downside risk. Moreover, the framework evaluates the maximum drawdown, a measure of the largest peak-to-trough decline in the portfolio's value, and the probability of loss, which quantifies how often the portfolio experiences negative returns.
The assets eligible to be included into the portfolio can be summarized in a mix of stocks and currency pairs. Historical price data for these assets is collected over a specified date range, extending from June 2018, to June 2024. After downloading and organizing the adjusted closing prices, the daily returns for each asset are calculated and prepared in a format suitable for analysis. The returns data is then reindexed to cover the entire date range, and any missing values are filled to provide a complete dataset for the entire period under analysis.
The next step in the analysis involves constructing a portfolio with equal weights assigned to each asset. This initial portfolio serves as a baseline for comparison. Using the daily returns data, portfolio-level returns are computed based on these equal weights. Several key metrics are then calculated for the portfolio, including average daily returns, annualized returns, and the standard deviation of returns, both on a daily and annualized basis. To evaluate the risk profile of the portfolio, the VaR and CVaR are computed at 95% and 99% confidence levels. Furthermore, the framework calculates other performance indicators such as the Sharpe ratio, Sortino ratio, maximum drawdown, and the probability of loss.
Afterward, the framework adjusts the asset weights to ac-

count for cryptocurrency holdings. This adjustment introduces official weights for various assets, ensuring that the portfolio's total weight is always equals to 100% when it is combined with the introduction of cryptocurrencies. Portfolio returns are recalculated using these revised weights, and the associated risk and performance metrics are computed again, allowing for a new comparative evaluation between the initial and adjusted portfolios.
Finally, an optimization procedure is employed to improve the risk profile of the portfolio, with a specific focus on minimizing the CVaR. The optimization process determines the optimized weights allocation among these assets, with constraints ensuring that the total weight remains flat and that short selling is not allowed. Upon solving this optimization problem, the portfolio returns are recalculated using the optimized weights. The same set of risk measures is then computed once again, allowing for a comprehensive evaluation of how the optimization improves the portfolio's risk-adjusted performance.
This framework thus provides a consistent methodology for analyzing and optimizing portfolios, offering valuable insights into risk management and performance evaluation across a wide array of financial assets.
This approach allows for a comprehensive analysis of cryptocurrencies as a portfolio enhancement tool. By considering various investment strategies and risk metrics, the study provides a detailed comparison that helps investors understand the potential benefits and risks associated with integrating cryptocurrencies into their investment portfolios.

Benchmark Portfolio

The study begins with a benchmark portfolio that has already been optimized by industry professionals[2].
The portfolio optimization is achieved applying a two-steps methodology: initially, the strategic asset allocation is defined by an expert based approach (qualitative step), where the single stocks are chosen evaluating the financial market conditions, the well-positioned sectors and geographical areas and the most attractive stocks in terms of upside forecasting. Once the single stocks have been selected, the weights of each stock are settled using mean-variance and resampled optimization techniques (note that the weight of some stocks may be zero).
Full disclosure about the single components is not given due to privacy reasons. However, the main portfolio's drivers are briefly summarized as follows: The portfolio is composed of 16 instruments, mainly focused on Energy (40.76%), Defense (19.76%) and Transportation (11.50%) sectors. For what concerns the geographical area distribution, the West Europe is the most picked (66.23%), followed by the North America (23.90%) and the Emerging Countries (9.87%). Furthermore, the portfolio shows a strong component of growth stocks (46.72%).
This diversified sector and geographical allocation allows the portfolio to capture potential growth opportunities across various industries and regions, while maintaining a balance between risk and return. However, to further enhance diversification and explore new potential sources of return, the composition of the study portfolio was extended in the following ways:

- By adding one cryptocurrency at a time with a weight equal to 1/N, where N is the number of constituents in the portfolio, and rescaling the weight of the other assets proportionally. By doing this, the impact of the individual cryptocurrency on portfolio performance could be assessed.
- Adding all possible combinations of the 5 cryptocurrencies by assigning each of them a percentage weight equal to a value between 0%, 1%, 2%, 3%, 4%, and 5%, and proportionally rescaling the weight of the other assets. This is in order to assess the joint impact of adding the cryptocurrencies.

Equally Weighted Portfolio

The Equally Weighted (EW) portfolio, commonly known as the "1/N" or "naive" portfolio, represents an investment strategy where equal weights are assigned to each asset in the investment universe. This approach is formulated as follows:

$$w_i = \frac{1}{N}, \quad \forall i = 1, \ldots, N.$$

Here, w_i denotes the weight allocated to the i-th asset, and N signifies the total number of assets within the portfolio. By distributing an identical fraction of the total investment across all N assets, the EW portfolio achieves the minimum concentration in terms of portfolio weights, embodying a straightforward and intuitive method for diversification.
Extending the composition of the study portfolio was done in two ways:

- By adding one cryptocurrency at a time with a weight of 1/N, where N is the number of constituents in the portfolio, with the constituents also

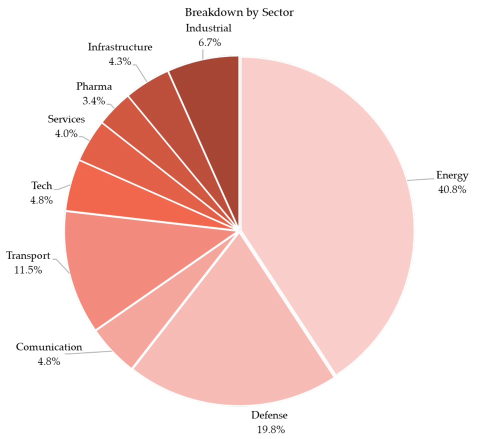

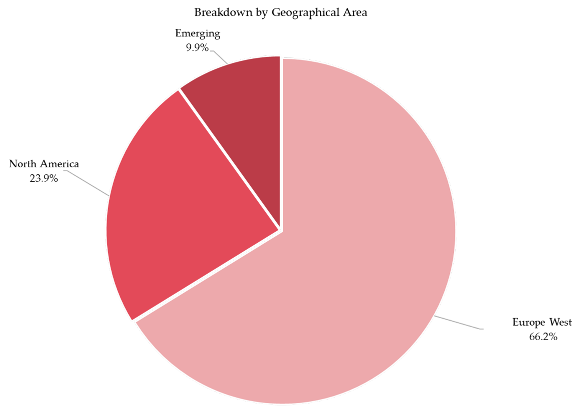

FIGURE 6: Breakdown by Sector and Geographical Area.

[2]For simplicity, the portfolio consists solely of stocks.

weighted 1/N. By doing this, one can compare a portfolio with generic weights with the (optimized) Benchmark portfolio.

- Adding all possible combinations of the 5 cryptocurrencies by assigning each of them a percentage weight equal to 1/N.

Despite its simplicity, this naive heuristic of equal allocation is not without its shortcomings. As highlighted by Braga (2015) [6], the assumption that the simple $\frac{1}{N}$ rule inherently guarantees effective diversification can be misleading. This is especially true in scenarios where individual asset risks exhibit significant variability. In such cases, an equally weighted portfolio might inadvertently result in concentrated risk exposures, contrary to the diversification goals.
Moreover, from equation (1), it is evident that the $\frac{1}{N}$ portfolio only aligns with the ex-ante Markowitz mean-variance efficient portfolio under stringent conditions: identical expected returns, equal volatilities, and uniform correlations among all assets. These conditions are rarely met in practical investment environments, thus questioning the optimality of the EW portfolio from a theoretical standpoint.
Nonetheless, extensive empirical research suggests that the EW strategy frequently outperforms more sophisticated strategies when evaluated ex-post. For instance, the seminal study by DeMiguel et al. (2007) [8] demonstrated that no portfolio weighting strategy consistently surpasses the $\frac{1}{N}$ rule in terms of the Sharpe ratio, certainty-equivalent return, or turnover. Their findings underscore a significant gap between the theoretical promises of optimal portfolio choice and its practical, out-of-sample performance. They concluded that substantial progress is still required to realize the benefits of optimal portfolio strategies in real-world scenarios.
Strictly speaking, the EW portfolio is not a reliable risk-based strategy, as it does not necessitate any parameter estimates for risks or correlations, nor does it involve the optimization of an objective function. However, within the context of investment literature, the EW strategy is often categorized as risk-based for a fundamental reason. Investors typically perceive the EW portfolio as a defensive strategy designed to allocate wealth equally across all assets, thereby aiming to enhance diversification and mitigate risk. This objective aligns closely with the general goals of risk-based portfolio optimization strategies, making the EW strategy a relevant benchmark in this domain.
In summary, while the Equally Weighted portfolio provides a simple and intuitive method for diversification, its limitations must be acknowledged, particularly in environments with diverse risk profiles among assets. Nevertheless, its consistent ex-post performance and alignment with the objectives of risk-based strategies ensure its continued relevance and utility in the field of portfolio management.
The introduction of this model into the analysis proves crucial for several reasons. First, it is crucial to have a portfolio that can be compared with a reference portfolio, which is as neutral as possible and has some robustness. This approach not only facilitates a more accurate analysis, but also allows a clear and objective evaluation criterion to be established. Moreover, adopting such a model allows for a study of cryptocurrencies that is distinguished by its impartiality. Each component of the portfolio is treated fairly, without favoritism toward their nature or past performance. This means that all cryptocurrencies are considered with equal weight, thus offering a more balanced and unbiased analysis.
Finally, the implementation of this model not only improves comparability with existing benchmarks, but also promotes a deeper understanding of market dynamics, fostering more informed and strategic decisions.

Minimizing CVaR Portfolio

The Minimizing Conditional Value at Risk (CVaR) model represents a sophisticated approach to portfolio optimization that focuses on mitigating potential losses in adverse market conditions. CVaR, also referred to as Expected Shortfall, measures the average losses that occur beyond a specified Value at Risk (VaR) threshold. This model is formulated as follows:

$$\mathrm{CVaR}_\alpha = \frac{1}{1-\alpha}\int_{-\infty}^{-\mathrm{VaR}_\alpha} x f(x)\,dx.$$

Here, α denotes the confidence level, while $f(x)$ represents the probability density function of portfolio returns. By concentrating on the tail of the loss distribution, the CVaR model provides a more comprehensive assessment of risk compared to traditional metrics like VaR, which only captures the maximum loss at a specified confidence level.
Extending the composition of the study portfolio was done in two ways:

- By adding one cryptocurrency at a time leaving freedom to the optimizer to determine the weights of all components, with no particular restrictions.
- Adding all 5 cryptocurrencies together and leaving freedom to the optimizer to determine the weights of all components, without special restrictions.

The significance of the CVaR model lies in its ability to enhance risk management practices, particularly for investors who prioritize capital preservation. According to Rockafellar and Uryasev (2000) [17], minimizing CVaR leads to portfolios that are less susceptible to extreme negative returns, thereby aligning more closely with the risk tolerance of many investors. This robustness against tail risks is especially pertinent in volatile markets, such as those observed in cryptocurrency trading, where price fluctuations can lead to substantial losses.
Despite its advantages, the CVaR model is not devoid of challenges. For instance, the optimization process can become computationally intensive, particularly in high-dimensional portfolios with numerous assets. Additionally, the effectiveness of CVaR as a risk metric relies heavily on accurate estimations of the return distribution. If the underlying assumptions about the asset returns are violated, the model may not perform as expected, leading to potentially misleading risk assessments.
Empirical research supports the utility of the CVaR model in achieving better risk-adjusted returns. For instance, studies have demonstrated that portfolios optimized using CVaR often outperform those constructed with traditional mean-variance techniques, especially in environments characterized by non-normal return distributions (Acerbi & Tasche, 2002 [1]; Rockafellar et al., 2006 [18]). These findings highlight the importance of considering tail risk when constructing investment portfolios.
In other words, the implementation of a CVaR-based optimization strategy enables investors to balance risk and return more effectively. By focusing on minimizing potential losses in extreme scenarios, the CVaR model aligns well with the objectives of risk-averse investors. Furthermore, this approach provides a framework for evaluating cryptocurrencies impartially, treating each asset equally regardless of historical performance or volatility.
In summary, while the Minimizing Conditional Value at Risk model offers a robust method for enhancing portfolio resilience against extreme losses, its practical application must consider the challenges associated with estimation and computation. Nevertheless, its effectiveness in improving risk-adjusted returns and fostering informed decision-making underscores its relevance in contemporary portfolio management.

In the context of the case study, the introduction of the Conditional Value at Risk (CVaR) model is particularly relevant due to the inherently volatile nature of cryptocurrencies. Investors often experience heightened anxiety when engaging with these digital assets, primarily driven by concerns related to potential losses and their magnitude. This volatility makes it imperative to employ robust risk management strategies.
The CVaR model, which focuses on minimizing the expected losses in the worst-case scenarios, offers a compelling framework for evaluating how cryptocurrencies can be integrated into an investment portfolio. By specifically targeting extreme losses, the CVaR approach aligns well with the risk aversion exhibited by many investors in this space. It allows for a more nuanced understanding of the risks associated with cryptocurrency investments and facilitates a strategic allocation that seeks to minimize potential downturns.
Given the unpredictable price movements typical of cryptocurrencies, assessing their role within a portfolio that prioritizes CVaR optimization is of significant interest. This analysis not only helps in understanding the compatibility of cryptocurrencies with traditional assets but also provides insights into how a CVaR-based portfolio can be constructed to enhance overall resilience against market fluctuations.
Ultimately, the focus on minimizing CVaR within the context of cryptocurrency investments aims to foster a more secure and informed approach for investors, allowing them to navigate the complexities of the digital asset landscape with greater confidence.

Results Markers

In this section, the indicators used to evaluate the models are presented. These indicators provide both quantitative measures of return and risk. Each one is defined mathematically to ensure clarity and rigor in its interpretation.
Average Return. The average return, denoted as $\bar{r}$, is the arithmetic mean of the returns r_t over the evaluation period, and is computed as follows:

$$\bar{r} = \frac{1}{T}\sum_{t=1}^{T} r_t,$$

where T is the number of time periods.
Annualized Average Return. The annualized average return, denoted as $\bar{r}_{\text{ann}}$, is the geometric mean of the returns over the evaluation period, annualized by the number of periods per year N:

$$\bar{r}_{\text{ann}} = \left(\prod_{t=1}^{T}(1+r_t)\right)^{\frac{N}{T}} - 1.$$

This allows for comparison across different time horizons.
Standard Deviation. The standard deviation σ measures the volatility of returns, defined as the square root of the variance:

$$\sigma = \sqrt{\frac{1}{T-1}\sum_{t=1}^{T}(r_t - \bar{r})^2}.$$

It provides a measure of the dispersion of returns from their mean.
Annualized Standard Deviation. The annualized standard deviation, σ_{ann}, adjusts the standard deviation for annual periods:

$$\sigma_{\text{ann}} = \sigma\sqrt{N},$$

where N is the number of periods in a year.
Value at Risk (VaR) at 95.0% Confidence Levels. The Value at Risk (VaR) quantifies the potential loss in portfolio value over a specified time period, with a given confidence level α. For a confidence level α (e.g., 95.0%), the VaR is defined as the threshold loss L such that the probability of a loss exceeding L is $1-\alpha$. Mathematically, for returns r_t, VaR is given by:

$$\text{VaR}_\alpha = \inf\{L \in \mathbb{R} : \mathbb{P}(r_t < L) \leq 1-\alpha\}.$$

Conditional Value at Risk (CVaR) at 95.0% Confidence Levels. The Conditional Value at Risk (CVaR), also known as Expected Shortfall, measures the expected loss given that the loss has exceeded the VaR level. For a confidence level α, the CVaR is computed as:

$$\text{CVaR}_\alpha = \mathbb{E}[r_t \mid r_t < \text{VaR}_\alpha].$$

This indicator provides a more conservative risk assessment by focusing on the average loss in extreme scenarios.
Sharpe Ratio. The Sharpe Ratio measures the excess return per unit of risk, defined as:

$$\text{Sharpe Ratio} = \frac{\bar{r} - r_f}{\sigma},$$

where r_f is the risk-free rate. A higher Sharpe ratio indicates better risk-adjusted performance.
Sortino Ratio. The Sortino Ratio refines the Sharpe Ratio by considering only downside risk. It is calculated by dividing the excess return by the downside deviation, which measures only the standard deviation of negative returns. Formally:

$$\text{Sortino Ratio} = \frac{\bar{r} - r_f}{\sigma_d},$$

where σ_d is the downside deviation, defined as:

$$\sigma_d = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\min(r_t - \bar{r}, 0)^2}.$$

This measure is useful for investors primarily concerned with downside risk, rather than overall volatility.
Maximum Drawdown. The maximum drawdown (MDD) quantifies the largest peak-to-trough decline observed during the evaluation period. It is defined as:

$$MDD = \max_{1\leq t\leq T}\left(\frac{\max_{0\leq u\leq t} P_u - P_t}{\max_{0\leq u\leq t} P_u}\right),$$

where P_t is the portfolio value at time t. This indicator is particularly relevant for evaluating a portfolio's susceptibility to large losses during market downturns.
Probability of Loss. The probability of loss (P_L) measures the likelihood that the portfolio experiences a negative return during the evaluation period. It is mathematically expressed as the proportion of time periods where returns are negative:

$$P_L = \frac{1}{T}\sum_{t=1}^{T}\mathbb{I}(r_t < 0),$$

where $\mathbb{I}$ is the indicator function, which equals 1 if the condition is met and 0 otherwise.
In summary, these indicators offer a comprehensive view of the portfolio's performance, allowing both return and risk to be analyzed quantitatively. The combination of statistical risk measures (such as standard deviation and VaR) with performance ratios (like the Sharpe and Sortino Ratios) ensures a balanced evaluation of the model's effectiveness.

Empirical Results

This chapter presents and analyzes the empirical results obtained from back-testing the three portfolios across the various time horizons outlined earlier, illustrating how the risk-return metrics evolve over time.
As explained in the previous chapter, for both the equally

weighted and benchmark portfolio models, the allocation to each cryptocurrency was set at 1/N, where N represents the number of assets in the portfolio. On the other hand, the minimum Conditional Value at Risk (CVAR) portfolio was allowed to reallocate asset percentages without such constraints.
For investors with a long-term horizon (exceeding four years), particular attention should be given to the results for the overall period from 2018 to 2024. Meanwhile, those with a medium-term focus should consider the portfolio performance across distinct market phases, such as bull, bear, and recovery periods, to better align their investment timing with market conditions.
The primary objective is to capitalize on periods of strong upward trends, i.e., bull markets, while avoiding prolonged downturns, i.e., bear markets, which could negatively impact portfolio returns.

Overall Period

First of all, the longest available period across all times series of the cryptocurrencies considered has been analyzed, so the time horizon between 2018 and 2024. The empirical results are shown in the table 2
Across all three portfolio models, the inclusion of cryptocurrencies has the following broad impacts:

Pros of Adding Cryptos:

- **Higher Returns:** the inclusion of cryptocurrencies enhances portfolio yearly returns, rising from 14% to as much as 20%, with assets such as Binance and Ethereum driving much of this growth. Over a time horizon exceeding a halving cycle (approximately four years), the positive performance of the crypto market results in notable gains.
- **Value at Risk (VaR) and Conditional Value at Risk (CVaR):** the inclusion of cryptocurrencies, particularly Bitcoin, results in minimal to no increase in VaR and CVaR measures. This is a critical factor, as these risk metrics are widely adopted in the regulatory framework of financial institutions, such as the Basel Accords.
- **Improved Risk-Adjusted Returns:** cryptocurrencies positively impact risk-adjusted return measures, with the Sharpe and Sortino Ratios increasing from 0.6 and 0.66 to 0.8 and 1, respectively. This indicates that the returns relative to the risk undertaken have improved.

Cons of Adding Cryptos:

- **Higher Volatility and Drawdowns:** the inclusion of cryptocurrencies marginally raises the portfolio's standard deviation (from 20% to 21%) and drawdowns (from 40% to 43%), exposing investors to slightly higher short-term risks and potential losses.
- **Higher Probability of Loss:** adding cryptocurrencies increases the probability of incurring a loss, particularly in the Equally Weighted and Benchmark portfolios, where this likelihood exceeds 46%. This result highlights the high volatility of cryptocurrencies in the short term. Due to their significant fluctuations, the probability of experiencing a loss on a single day is higher. However, this measure does not account for the magnitude of gains and losses, which are ultimately considered in the yearly return ratio instead.

For **risk-averse investors**, incorporating cryptocurrencies into a **Minimum CVaR portfolio** provides a balance between enhancing returns through a higher allocation to cryptocurrencies and controlling risk by minimizing CVaR through linear programming optimization. The results indicate that Bitcoin is the most stable option, as expected, given its large market capitalization, lower volatility compared to other cryptocurrencies, and thinner tail loss distribution. On the other hand, **risk-tolerant investors** may prefer the **Equally Weighted** or **Benchmark** portfolios, which offer higher returns, particularly through assets like Ethereum or Binance that can outperform Bitcoin.
However, they must account for the increased volatility inherent in both the portfolio structure and the more volatile nature of these alternative cryptocurrencies. A practical alternative could be to combine these assets with Bitcoin but in lower proportions to achieve a more balanced risk-return trade-off. This strategy is further examined in the following section.
Multiple Cryptos' Combination. This section explores the inclusion of various combinations of the five cryptocurrencies within the portfolio, with allocations ranging from 0% to 5% for each, resulting in a total crypto allocation of between 0% and 25%. The findings reveal a direct correlation between the percentage of crypto allocation and both risk and return. Specifically, a portfolio with 5% allocated to each cryptocurrency exhibits the highest annual return (up to 31%) and the highest CVaR (up to 3.30%). Conversely, portfolios without cryptocurrency display the lowest annual return (14%) and rank among those with the lowest CVaR (2.53%), although portfolios with Bitcoin allocations share the same CVaR.
Moreover, portfolios with the highest Sharpe and Sortino Ratios (0.98 and 1.19, respectively) consistently allocate 5% to Bitcoin, Ethereum, and Binance, while assigning more moderate percentages to Cardano and Ripple. This analysis confirms that a significant allocation to cryptocurrencies within a traditional portfolio leads to a substantial increase in risk; however, when appropriately balanced, it can maintain risk levels comparable to those without cryptocurrency.

Bull Market

Then the same analysis was performed on 2 bull market periods (2020-2021, 2023-2024) which usually starts near the halving and continue for about a year and a half, driven by the euphoria of fear-of-missing-out phenomena(FOMO) created by the halving of Bitcoin supply. This period ends with an all-time-high price and overbought market conditions. The empirical results are shown in the table 3.
Again, the inclusion of cryptocurrencies has the following broad impacts: **Pros of Adding Cryptos:**

- **Strongly Higher Returns:** cryptocurrencies improve strongly the portfolio's returns, moving from 73% up to over 100% in the first bull market, with Binance and Cardano as top performers. While for the second returns moves from 24-28% to 36-41%, with Binance but also Bitcoin and Ethereum as top performers. This result shows that during bull periods, returns increase dramatically in periods of only about a year and a half.
- **Value at Risk (VaR) and Conditional Value at Risk (CVaR):** again VaR and CVaR metrics experienced little to no increase with the inclusion of cryptocurrencies, even if min CVaR model includes only Bitcoin in a small proportion(1-2%), it can be seen that all Bitcoin portfolios don't increase CVaR, while other cryptos increased CVaR slightly.
- **Better Risk-Adjusted Returns:** cryptos enhance the Sharpe and Sortino Ratios, from 3 and 4 up to 3.6 and 5 in the first bull market, from 1.5 and 2 to 2.1 and 2.8 and in the second bull market indicating

Portfolio	Crypto	%	Return	S.D.	V.95%	CV.95%	Sh.R	So.R.	M.D.	P.L.
minCVaR	None	0.00%	14.00%	15.98%	-1.22%	-1.91%	0.76	0.84	-29.69%	33.83%
minCVaR	Bitcoin	3.51%	16.39%	15.94%	-1.22%	-1.90%	0.89	1.05	-30.11%	46.97%
minCVaR	Ethereum	0.09%	14.00%	15.98%	-1.21%	-1.91%	0.76	0.89	-29.76%	47.24%
minCVaR	Binance	1.55%	15.67%	15.99%	-1.22%	-1.90%	0.85	1.00	-30.47%	47.24%
minCVaR	Ripple	1.26%	14.70%	15.97%	-1.21%	-1.91%	0.80	0.95	-30.08%	48.65%
minCVaR	Cardano	0.78%	14.56%	15.99%	-1.22%	-1.91%	0.79	0.93	-30.17%	47.83%
Eq.We.	None	0.00%	13.85%	19.62%	-1.54%	-2.42%	0.61	0.66	-39.69%	33.56%
Eq.We.	Bitcoin	5.88%	16.95%	19.67%	-1.50%	-2.40%	0.75	0.86	-39.18%	46.56%
Eq.We.	Ethereum	5.88%	17.62%	20.16%	-1.56%	-2.49%	0.76	0.88	-40.61%	46.47%
Eq.We.	Binance	5.88%	20.09%	20.19%	-1.56%	-2.45%	0.86	1.01	-39.67%	46.83%
Eq.We.	Ripple	5.88%	16.35%	20.37%	-1.55%	-2.44%	0.69	0.84	-40.33%	48.29%
Eq.We.	Cardano	5.88%	17.10%	20.50%	-1.60%	-2.50%	0.72	0.86	-41.68%	46.74%
Benchmark	None	0.00%	14.26%	20.56%	-1.63%	-2.53%	0.60	0.66	-42.72%	33.74%
Benchmark	Bitcoin	5.88%	17.34%	20.52%	-1.61%	-2.52%	0.73	0.85	-40.13%	46.42%
Benchmark	Ethereum	5.88%	18.01%	21.00%	-1.68%	-2.61%	0.74	0.87	-43.47%	46.69%
Benchmark	Binance	5.88%	20.49%	21.04%	-1.63%	-2.56%	0.84	1.00	-40.57%	47.15%
Benchmark	Ripple	5.88%	16.74%	21.19%	-1.65%	-2.55%	0.68	0.84	-43.21%	48.34%
Benchmark	Cardano	5.88%	17.49%	21.34%	-1.68%	-2.62%	0.71	0.85	-44.48%	47.47%

***TABLE 2:** Empirical Results of Backtesting Overall Period (2018-2024).*

strongly improved returns relative to the amount of risk taken.

Cons of Adding Cryptos:

- **Higher Volatility and Drawdowns:** cryptos slightly increase the portfolios' standard deviation (+ 0.4% in the first bull market, +1% in the second except for Bitcoin) exposing investors to a slightly higher short-term risks and potential losses.
- **Higher Probability of Loss:** including cryptos increases the likelihood of experiencing a loss, particularly in the Equally Weighted and Benchmark portfolios, it increases of about 10%, similarly to the overall period.

The bull market has been the most profitable period in the crypto market, if the intention is to invest in time frames of a few years, it is crucial to try to arrive at the beginning of the bull market having already added the crypto component within your portfolio. Again min CVaR portfolio has a more risk-adverse setting keeping risk lower but losing a good part of returns, while Equally Weighted and Benchmark portfolios are more risk-tolerant and benefit from all the high yields of this period.

Bear Market

Following bull market periods, the analysis move on to the bear market periods that arise at the end of bull markets from an overbought and overpriced crypto market situation. In this period the phenomenon of fear-uncertainty-doubt(FUD) in investors developed, producing strong selling pressure and a sharp drop in prices for about a year moving to an oversold market situation. Analysing the 2 bear periods (2018 and 2022) we obtain the results on Table 4.

In this market scenario the inclusion of cryptocurrencies produces substantial differences from the previous case:

Pros of Adding Cryptos:

- **No Advantages:** generally adding cryptos during bear markets worsens portfolio performance in both risk and return, except for some outliers cases where in the first bear market min CVaR portfolio keep a small percentage of Bitcoin and Binance keeping CVaR basically the same, Binance also slightly improves returns.

Cons of Adding Cryptos:

- **Higher Risks:** substantially all risk metrics from standard deviation (+2% in the first, +1% in the second) to CVaR (+0.2%) and the others worsen with the addition of cryptos without exception.
- **Lower Returns:** the inclusion of cryptos also leads to lower returns due to the large loss of value in the market, with the only exception of Binance in the first bear market.
- **Lower Risk-Adjusted Returns:** even Sharpe and Sortino Ratio worsen, again with the exception of Binance in the first bear market.

In contrast to the bull market, the bear market is the market's worst-performing period, so it should be avoided especially in medium-term investments as it risks to remove much of the returns achieved in the previous bull market. the only portfolio that might be resilient to this negative market period is the min CVaR, which tries to limit risk by introducing cryptos only in small percentages (2018 case) or not introducing them (2022 case).

Recovering Market

Finally, the last market scenario is the recovering market, this period, from the perspective of fundamental analysis, can be considered the most important, as investors begin to re-fund the industry, blockchain application developers and related start-ups. The market restarts to gain value from the relative low of the previous bear market, returning closer to the all time high of the previous bull market as we move towards the new halving. Here the 2 recovering market phases are exposed (2019-2020, 2023):

Pros of Adding Cryptos:

- **Higher Returns:** cryptocurrencies improve the portfolio's returns, moving from -2 and 0% up to over 4 and 6% in the first recovering market, with Bitcoin

Start	End	Portfolio	Crypto	%	Return	S.D.	V.95%	CV.95%	Sh.R.	So.R.	M.D.	P.L.
11/5/20	11/11/21	minCVaR	None	0.00%	50.49%	14.67%	-1.13%	-1.48%	2.72	4.20	-7.68%	31.45%
11/5/20	11/11/21	minCVaR	Bitcoin	0.82%	52.24%	14.71%	-1.09%	-1.48%	2.79	4.27	-7.69%	44.91%
11/5/20	11/11/21	minCVaR	Ethereum	0.00%	50.49%	14.67%	-1.13%	-1.48%	2.72	4.10	-7.68%	42.36%
11/5/20	11/11/21	minCVaR	Binance	0.30%	50.76%	14.64%	-1.10%	-1.48%	2.74	4.14	-7.46%	43.64%
11/5/20	11/11/21	minCVaR	Ripple	0.00%	50.49%	14.67%	-1.13%	-1.48%	2.72	4.13	-7.68%	45.27%
11/5/20	11/11/21	minCVaR	Cardano	0.00%	50.49%	14.67%	-1.13%	-1.48%	2.72	4.12	-7.68%	44.36%
11/5/20	11/11/21	Eq.We.	None	0.00%	73.33%	17.98%	-1.33%	-1.98%	3.01	4.04	-12.31%	30.55%
11/5/20	11/11/21	Eq.We.	Bitcoin	5.88%	84.34%	17.96%	-1.37%	-1.98%	3.35	4.71	-10.70%	41.82%
11/5/20	11/11/21	Eq.We.	Ethereum	5.88%	95.63%	18.49%	-1.29%	-2.08%	3.58	4.94	-11.64%	38.55%
11/5/20	11/11/21	Eq.We.	Binance	5.88%	102.81%	19.27%	-1.28%	-2.04%	3.62	5.24	-9.99%	40.55%
11/5/20	11/11/21	Eq.We.	Ripple	5.88%	91.35%	19.47%	-1.35%	-2.04%	3.29	4.96	-12.37%	43.82%
11/5/20	11/11/21	Eq.We.	Cardano	5.88%	103.06%	19.21%	-1.38%	-2.08%	3.64	5.26	-13.21%	41.64%
11/5/20	11/11/21	Benchmark	None	0.00%	73.74%	19.42%	-1.49%	-2.16%	2.79	3.80	-13.10%	29.45%
11/5/20	11/11/21	Benchmark	Bitcoin	5.88%	84.75%	19.25%	-1.54%	-2.16%	3.14	4.38	-11.50%	41.64%
11/5/20	11/11/21	Benchmark	Ethereum	5.88%	96.06%	19.77%	-1.64%	-2.26%	3.36	4.64	-12.54%	40.00%
11/5/20	11/11/21	Benchmark	Binance	5.88%	103.26%	20.58%	-1.48%	-2.22%	3.40	4.92	-10.98%	41.27%
11/5/20	11/11/21	Benchmark	Ripple	5.88%	91.77%	20.62%	-1.58%	-2.22%	3.11	4.64	-13.08%	44.00%
11/5/20	11/11/21	Benchmark	Cardano	5.88%	103.50%	20.51%	-1.55%	-2.26%	3.42	4.92	-13.91%	42.36%
10/12/23	6/1/24	minCVaR	None	0.00%	36.91%	12.34%	-0.90%	-0.99%	2.47	4.71	-5.69%	32.91%
10/12/23	6/1/24	minCVaR	Bitcoin	1.73%	41.21%	12.38%	-0.89%	-0.99%	2.71	4.88	-5.33%	43.16%
10/12/23	6/1/24	minCVaR	Ethereum	0.00%	36.91%	12.34%	-0.90%	-0.99%	2.47	4.43	-5.69%	44.02%
10/12/23	6/1/24	minCVaR	Binance	0.00%	36.91%	12.34%	-0.90%	-0.99%	2.47	4.43	-5.69%	44.02%
10/12/23	6/1/24	minCVaR	Ripple	0.00%	36.91%	12.34%	-0.90%	-0.99%	2.47	4.43	-5.69%	47.01%
10/12/23	6/1/24	minCVaR	Cardano	0.00%	36.91%	12.34%	-0.90%	-0.99%	2.47	4.43	-5.69%	45.73%
10/12/23	6/1/24	Eq.We.	None	0.00%	23.92%	14.95%	-1.28%	-1.79%	1.37	1.81	-8.51%	32.48%
10/12/23	6/1/24	Eq.We.	Bitcoin	5.88%	34.36%	14.58%	-1.29%	-1.74%	1.96	2.57	-6.84%	41.45%
10/12/23	6/1/24	Eq.We.	Ethereum	5.88%	34.29%	14.74%	-1.31%	-1.78%	1.93	2.54	-6.85%	40.60%
10/12/23	6/1/24	Eq.We.	Binance	5.88%	36.27%	14.82%	-1.39%	-1.77%	2.02	2.71	-7.47%	41.45%
10/12/23	6/1/24	Eq.We.	Ripple	5.88%	24.44%	14.94%	-1.35%	-1.78%	1.40	1.91	-8.61%	45.30%
10/12/23	6/1/24	Eq.We.	Cardano	5.88%	31.59%	15.28%	-1.28%	-1.77%	1.73	2.43	-7.87%	41.88%
10/12/23	6/1/24	Benchmark	None	0.00%	28.76%	16.04%	-1.33%	-1.89%	1.51	2.05	-8.96%	30.77%
10/12/23	6/1/24	Benchmark	Bitcoin	5.88%	39.28%	15.55%	-1.20%	-1.83%	2.07	2.74	-7.27%	41.03%
10/12/23	6/1/24	Benchmark	Ethereum	5.88%	39.21%	15.76%	-1.24%	-1.86%	2.04	2.74	-7.35%	41.88%
10/12/23	6/1/24	Benchmark	Binance	5.88%	41.26%	15.71%	-1.38%	-1.86%	2.14	2.88	-7.89%	41.45%
10/12/23	6/1/24	Benchmark	Ripple	5.88%	29.01%	15.84%	-1.29%	-1.86%	1.55	2.13	-9.05%	45.30%
10/12/23	6/1/24	Benchmark	Cardano	5.88%	36.42%	16.13%	-1.19%	-1.86%	1.86	2.65	-8.31%	43.59%

TABLE 3: *Empirical Results of Backtesting Bull Markets.*

Start	End	Portfolio	Crypto	%	Return	S.D.	V.95%	CV.95%	Sh.R.	So.R.	M.D.	P.L.
1/1/18	31/12/18	minCVaR	None	0.00%	-2.97%	12.74%	-1.26%	-1.39%	-0.32	-0.46	-13.93%	35.89%
1/1/18	31/12/18	minCVaR	Bitcoin	2.64%	-4.33%	12.52%	-1.22%	-1.37%	-0.43	-0.62	-14.35%	46.30%
1/1/18	31/12/18	minCVaR	Ethereum	0.00%	-2.97%	12.74%	-1.26%	-1.39%	-0.32	-0.43	-13.93%	48.22%
1/1/18	31/12/18	minCVaR	Binance	1.12%	-1.68%	12.65%	-1.23%	-1.38%	-0.21	-0.30	-13.75%	49.59%
1/1/18	31/12/18	minCVaR	Ripple	1.95%	-2.02%	12.73%	-1.21%	-1.39%	-0.24	-0.33	-12.04%	47.67%
1/1/18	31/12/18	minCVaR	Cardano	0.00%	-2.97%	12.74%	-1.26%	-1.39%	-0.32	-0.43	-13.93%	49.32%
1/1/18	31/12/18	Eq.We.	None	0.00%	-7.08%	14.73%	-1.38%	-1.71%	-0.57	-0.77	-18.76%	34.79%
1/1/18	31/12/18	Eq.We.	Bitcoin	5.88%	-11.65%	14.88%	-1.39%	-1.78%	-0.90	-1.22	-20.37%	48.77%
1/1/18	31/12/18	Eq.We.	Ethereum	5.88%	-12.68%	15.64%	-1.53%	-1.96%	-0.93	-1.24	-23.92%	48.77%
1/1/18	31/12/18	Eq.We.	Binance	5.88%	-3.37%	16.67%	-1.34%	-1.88%	-0.27	-0.39	-21.47%	49.59%
1/1/18	31/12/18	Eq.We.	Ripple	5.88%	-12.17%	16.20%	-1.55%	-1.91%	-0.86	-1.22	-18.76%	49.86%
1/1/18	31/12/18	Eq.We.	Cardano	5.88%	-16.72%	16.44%	-1.56%	-1.97%	-1.17	-1.67	-24.54%	49.86%
1/1/18	31/12/18	Benchmark	None	0.00%	-6.56%	14.83%	-1.47%	-1.75%	-0.52	-0.71	-18.06%	35.34%
1/1/18	31/12/18	Benchmark	Bitcoin	5.88%	-11.17%	14.97%	-1.45%	-1.81%	-0.86	-1.18	-20.04%	46.58%
1/1/18	31/12/18	Benchmark	Ethereum	5.88%	-12.21%	15.79%	-1.48%	-1.98%	-0.89	-1.19	-23.78%	47.95%
1/1/18	31/12/18	Benchmark	Binance	5.88%	-2.85%	16.81%	-1.42%	-1.91%	-0.23	-0.34	-21.01%	50.68%
1/1/18	31/12/18	Benchmark	Ripple	5.88%	-11.70%	16.42%	-1.45%	-1.95%	-0.82	-1.15	-18.58%	50.41%
1/1/18	31/12/18	Benchmark	Cardano	5.88%	-16.27%	16.61%	-1.58%	-2.00%	-1.13	-1.61	-24.34%	50.14%
12/11/21	14/11/22	minCVaR	None	0.00%	19.33%	16.90%	-1.42%	-1.93%	0.99	1.37	-13.67%	35.60%
12/11/21	14/11/22	minCVaR	Bitcoin	0.00%	19.33%	16.90%	-1.42%	-1.93%	0.99	1.39	-13.67%	50.00%
12/11/21	14/11/22	minCVaR	Ethereum	0.00%	19.33%	16.90%	-1.42%	-1.93%	0.99	1.38	-13.67%	49.73%
12/11/21	14/11/22	minCVaR	Binance	0.00%	19.33%	16.90%	-1.42%	-1.93%	0.99	1.38	-13.67%	49.46%
12/11/21	14/11/22	minCVaR	Ripple	0.00%	19.33%	16.90%	-1.42%	-1.93%	0.99	1.39	-13.67%	50.54%
12/11/21	14/11/22	minCVaR	Cardano	0.00%	19.33%	16.90%	-1.42%	-1.93%	0.99	1.38	-13.67%	49.18%
12/11/21	14/11/22	Eq.We.	None	0.00%	-7.01%	22.65%	-2.01%	-2.89%	-0.36	-0.44	-23.93%	36.68%
12/11/21	14/11/22	Eq.We.	Bitcoin	5.88%	-12.63%	23.22%	-1.99%	-2.98%	-0.62	-0.79	-27.05%	51.90%
12/11/21	14/11/22	Eq.We.	Ethereum	5.88%	-11.61%	23.84%	-2.06%	-3.07%	-0.56	-0.72	-26.79%	51.63%
12/11/21	14/11/22	Eq.We.	Binance	5.88%	-9.51%	23.17%	-2.00%	-2.96%	-0.47	-0.61	-25.33%	50.82%
12/11/21	14/11/22	Eq.We.	Ripple	5.88%	-11.23%	23.19%	-1.99%	-2.97%	-0.56	-0.72	-25.08%	50.82%
12/11/21	14/11/22	Eq.We.	Cardano	5.88%	-13.72%	23.90%	-2.05%	-3.04%	-0.66	-0.85	-28.19%	48.64%
12/11/21	14/11/22	Benchmark	Nessuna	0.00%	-7.46%	24.27%	-2.13%	-3.06%	-0.36	-0.45	-24.97%	36.41%
12/11/21	14/11/22	Benchmark	Bitcoin	5.88%	-13.02%	24.70%	-2.08%	-3.13%	-0.61	-0.77	-27.58%	50.54%
12/11/21	14/11/22	Benchmark	Ethereum	5.88%	-12.01%	25.29%	-2.18%	-3.22%	-0.55	-0.71	-27.30%	50.27%
12/11/21	14/11/22	Benchmark	Binance	5.88%	-9.91%	24.63%	-2.09%	-3.12%	-0.46	-0.60	-25.87%	50.54%
12/11/21	14/11/22	Benchmark	Ripple	5.88%	-11.63%	24.64%	-2.08%	-3.12%	-0.54	-0.71	-25.60%	50.54%
12/11/21	14/11/22	Benchmark	Cardano	5.88%	-14.10%	25.35%	-2.22%	-3.21%	-0.64	-0.83	-28.70%	49.73%

TABLE 4: *Empirical Results of Backtesting Bear Markets.*

and Binance as top performers. While for the second returns moves from 0-3% to 4-6%, with Bitcoin and Ripple as top performers. This result shows that also during recovery periods, returns increases.

- **Value at Risk (VaR) and Conditional Value at Risk (CVaR):** again VaR and CVaR metrics experienced little to no increase with the inclusion of cryptocurrencies in the second recovering market, even if min CVaR model, not includes cryptos it can be seen that all Bitcoin portfolios don't increase CVaR, while other cryptos increased CVaR slightly. In the first scenario, it can be observed that the benchmark and equally weighted portfolios slightly reduce the CVaR, while the minimum CVaR portfolio allocates a higher percentage to cryptocurrencies than in all other cases (ranging from 4% to 6%). This allocation increase is due to the period of increased volatility caused by the COVID pandemic, which affected all portfolio assets by increasing their overall risk. As a result, the addition of cryptocurrencies no longer has a significant impact on the total portfolio risk as in the other cases.

Cons of Adding Cryptos:

- **Slightly Higher Standard Deviation:** Standard deviation slightly increase, in particular for smaller cap cryptos, such as Ripple and Cardano.
- **Higher Probability of Loss:** this metric basically changes in the same way as overall and bull market cases, decreasing of around 10% with cryptocurrencies.

Portfolio-Indexes Correlation

This chapter analyzes the correlation between the previously discussed portfolios and the benchmark indices, Eurostoxx 50 and Nasdaq Composite in the overall period. A step-by-step addition of cryptocurrencies was made to each portfolio, introducing one cryptocurrency at a time from the five selected, to observe how the correlation coefficient changes as the percentage of cryptocurrencies in the portfolio increases and to see if any benefits in market

Start	End	Portfolio	Crypto	%	Return	S.D.	V.95%	CV.95%	Sh.R.	So.R.	M.D.	P.L.
1/1/19	5/10/20	minCVaR	None	0.00%	18.78%	18.97%	-1.15%	-2.33%	0.85	0.78	-27.67%	33.00%
1/1/19	5/10/20	minCVaR	Bitcoin	6.02%	21.78%	19.27%	-1.10%	-2.23%	0.97	0.97	-28.21%	44.67%
1/1/19	5/10/20	minCVaR	Ethereum	5.08%	22.17%	19.77%	-1.05%	-2.26%	0.96	0.98	-28.60%	46.68%
1/1/19	5/10/20	minCVaR	Binance	4.69%	23.24%	19.52%	-1.09%	-2.27%	1.02	1.04	-28.69%	47.48%
1/1/19	5/10/20	minCVaR	Ripple	5.56%	18.06%	19.58%	-1.16%	-2.28%	0.80	0.82	-27.52%	47.08%
1/1/19	5/10/20	minCVaR	Cardano	4.59%	19.95%	19.81%	-1.07%	-2.27%	0.87	0.88	-28.38%	45.67%
1/1/19	5/10/20	Eq.We.	None	0.00%	-0.29%	24.43%	-1.56%	-3.29%	-0.05	-0.04	-39.09%	31.79%
1/1/19	5/10/20	Eq.We.	Bitcoin	5.88%	5.22%	24.38%	-1.49%	-3.16%	0.17	0.16	-39.18%	4.27%
1/1/19	5/10/20	Eq.We.	Ethereum	5.88%	3.70%	24.80%	-1.55%	-3.23%	0.11	0.10	-39.72%	46.48%
1/1/19	5/10/20	Eq.We.	Binance	5.88%	6.58%	24.85%	-1.61%	-3.23%	0.22	0.21	-39.67%	46.48%
1/1/19	5/10/20	Eq.We.	Ripple	5.88%	-1.02%	24.47%	-1.56%	-3.21%	-0.08	-0.08	-39.37%	46.48%
1/1/19	5/10/20	Eq.We.	Cardano	5.88%	3.19%	24.98%	-1.60%	-3.28%	0.09	0.08	-39.82%	45.88%
1/1/19	5/10/20	Benchmark	None	0.00%	-2.27%	24.60%	-1.62%	-3.29%	-0.13	-0.12	-40.00%	34.21%
1/1/19	5/10/20	Benchmark	Bitcoin	5.88%	3.25%	24.54%	-1.49%	-3.19%	0.09	0.09	-40.09%	46.28%
1/1/19	5/10/20	Benchmark	Ethereum	5.88%	1.75%	24.93%	-1.51%	-3.26%	0.03	0.03	-40.62%	48.09%
1/1/19	5/10/20	Benchmark	Binance	5.88%	4.59%	25.00%	-1.63%	-3.24%	0.14	0.14	-40.57%	47.69%
1/1/19	5/10/20	Benchmark	Ripple	5.88%	-2.87%	24.59%	-1.62%	-3.24%	-0.16	-0.16	-40.28%	47.08%
1/1/19	5/10/20	Benchmark	Cardano	5.88%	1.25%	25.12%	-1.58%	-3.29%	0.01	0.01	-40.71%	47.69%
11/15/22	10/11/23	minCVaR	None	0.00%	22.48%	11.39%	-0.84%	-1.08%	1.69	2.81	-7.46%	36.25%
11/15/22	10/11/23	minCVaR	Bitcoin	0.00%	22.48%	11.39%	-0.84%	-1.08%	1.69	2.71	-7.46%	51.66%
11/15/22	10/11/23	minCVaR	Ethereum	0.00%	22.48%	11.39%	-0.84%	-1.08%	1.69	2.71	-7.46%	51.96%
11/15/22	10/11/23	minCVaR	Binance	0.00%	22.48%	11.39%	-0.84%	-1.08%	1.69	2.71	-7.46%	50.76%
11/15/22	10/11/23	minCVaR	Ripple	0.00%	22.48%	11.39%	-0.84%	-1.08%	1.69	2.71	-7.46%	50.76%
11/15/22	10/11/23	minCVaR	Cardano	0.00%	22.48%	11.39%	-0.84%	-1.08%	1.69	2.71	-7.46%	50.76%
11/15/22	10/11/23	Eq.We.	None	0.00%	0.47%	15.12%	-1.21%	-1.76%	-0.04	-0.05	-17.05%	35.65%
11/15/22	10/11/23	Eq.We.	Bitcoin	5.88%	4.17%	14.94%	-1.16%	-1.73%	0.21	0.29	-16.45%	52.27%
11/15/22	10/11/23	Eq.We.	Ethereum	5.88%	2.56%	15.17%	-1.17%	-1.78%	0.10	0.14	-16.76%	53.47%
11/15/22	10/11/23	Eq.We.	Binance	5.88%	-0.79%	14.95%	-1.15%	-1.76%	-0.12	-0.17	-16.70%	52.57%
11/15/22	10/11/23	Eq.We.	Ripple	5.88%	4.30%	16.46%	-1.23%	-1.75%	0.19	0.31	-18.11%	53.78%
11/15/22	10/11/23	Eq.We.	Cardano	5.88%	-0.45%	15.59%	-1.25%	-1.80%	-0.09	-0.14	-17.16%	53.78%
11/15/22	10/11/23	Benchmark	None	0.00%	2.86%	16.09%	-1.36%	-1.91%	0.11	0.15	-17.47%	35.95%
11/15/22	10/11/23	Benchmark	Bitcoin	5.88%	6.50%	15.83%	-1.34%	-1.89%	0.33	0.45	-16.84%	51.06%
11/15/22	10/11/23	Benchmark	Ethereum	5.88%	4.85%	16.09%	-1.35%	-1.93%	0.23	0.32	-17.16%	50.76%
11/15/22	10/11/23	Benchmark	Binance	5.88%	1.43%	15.86%	-1.31%	-1.92%	0.03	0.04	-17.09%	50.45%
11/15/22	10/11/23	Benchmark	Ripple	5.88%	6.63%	17.36%	-1.38%	-1.90%	0.31	0.47	-18.50%	52.57%
11/15/22	10/11/23	Benchmark	Cardano	5.88%	1.78%	16.48%	-1.41%	-1.95%	0.05	0.06	-17.55%	51.36%

TABLE 5: Empirical Results of Backtesting Recovering Markets.

diversification are achieved. Results are shown in Table 6.

The results indicate that the correlation with the Eurostoxx 50 decreases significantly with the addition of crypto assets, while a larger proportion is required to reduce correlation with the Nasdaq Composite. This occurs because the benchmark and equally weighted portfolios are predominantly composed of European stocks, with a smaller allocation to American stocks. In conclusion, the inclusion of cryptocurrencies provides notable benefits in terms of market risk diversification, particularly for European markets, which constitute the larger portion of the portfolio.

StartDate	EndDate	Portfolio	Crypto	%	EurostocksCorr.	NasdaqCorr.
6/1/2018	6/1/2024	MinCVAR	None	0%	0.6217	0.5681
6/1/2018	6/1/2024	MinCVAR	Bitcoin	3.51%	0.6333	0.5820
6/1/2018	6/1/2024	MinCVAR	Ethereum	0.09%	0.6205	0.5669
6/1/2018	6/1/2024	MinCVAR	Binance	1.55%	0.6292	0.5750
6/1/2018	6/1/2024	MinCVAR	Ripple	1.26%	0.6212	0.5715
6/1/2018	6/1/2024	MinCVAR	Cardano	0.78%	0.6210	0.5702
6/1/2018	6/1/2024	MinCVAR	BTC+ETH	3.51%	0.6333	0.5820
6/1/2018	6/1/2024	MinCVAR	BTC+ETH+BNB	3.51%	0.6333	0.5820
6/1/2018	6/1/2024	MinCVAR	BTC+ETH+BNB+XRP	3.51%	0.6333	0.5820
6/1/2018	6/1/2024	MinCVAR	BTC+ETH+BNB+XRP+ADA	3.51%	0.6333	0.5820
6/1/2018	6/1/2024	Eq.We.	None	0%	0.7502	0.6061
6/1/2018	6/1/2024	Eq.We.	Bitcoin	5.88%	0.7502	0.6325
6/1/2018	6/1/2024	Eq.We.	Ethereum	5.88%	0.7475	0.6366
6/1/2018	6/1/2024	Eq.We.	Binance	5.88%	0.7440	0.6271
6/1/2018	6/1/2024	Eq.We.	Ripple	5.88%	0.7344	0.6237
6/1/2018	6/1/2024	Eq.We.	Cardano	5.88%	0.7389	0.6353
6/1/2018	6/1/2024	Eq.We.	BTC+ETH	11.72%	0.7292	0.6434
6/1/2018	6/1/2024	Eq.We.	BTC+ETH+BNB	17.64%	0.6982	0.6345
6/1/2018	6/1/2024	Eq.We.	BTC+ETH+BNB+XRP	23.52%	0.6582	0.6176
6/1/2018	6/1/2024	Eq.We.	BTC+ETH+BNB+XRP+ADA	29.40%	0.6204	0.6027
6/1/2018	6/1/2024	Benchmark	None	0%	0.7405	0.5654
6/1/2018	6/1/2024	Benchmark	Bitcoin	5.88%	0.7419	0.5930
6/1/2018	6/1/2024	Benchmark	Ethereum	5.88%	0.7395	0.5975
6/1/2018	6/1/2024	Benchmark	Binance	5.88%	0.7359	0.5889
6/1/2018	6/1/2024	Benchmark	Ripple	5.88%	0.7283	0.5868
6/1/2018	6/1/2024	Benchmark	Cardano	5.88%	0.7315	0.5971
6/1/2018	6/1/2024	Benchmark	BTC+ETH	11.72%	0.7205	0.6074
6/1/2018	6/1/2024	Benchmark	BTC+ETH+BNB	17.64%	0.6821	0.5985
6/1/2018	6/1/2024	Benchmark	BTC+ETH+BNB+XRP	23.52%	0.6288	0.5768
6/1/2018	6/1/2024	Benchmark	BTC+ETH+BNB+XRP+ADA	29.40%	0.5747	0.5543

TABLE 6: *Correlations between Portfolios and Indices from 6/1/2018 to 6/1/2024.*

Conclusion

The inclusion of cryptocurrencies in an optimized portfolio offers both potential benefits and risks that warrant careful evaluation. Surely, cryptocurrencies can provide higher returns in specific market contexts, such as bull markets, without increasing or slightly increasing CVaR, enhance the risk-return ratio, and expand diversification opportunities. On the other hand, cryptocurrencies also entail significant volatility, which can increase the overall portfolio risk, particularly during bear market phases.
An analysis of major crypto assets - Bitcoin, Ethereum, and others such as Binance Coin, Ripple, and Cardano - reveals distinctive risk-return profiles. Bitcoin and Ethereum stand out for having lower volatility compared to the altcoins, thereby reducing overall portfolio risk relative to emerging and less capitalized cryptocurrencies. However, even these primary assets require careful entry and exit timing, especially for short-to-medium term investments.
Looking forward, while past performance is no guarantee of future returns, the cyclical bull-bear nature of cryptocurrencies has consistently been observed through the past Bitcoin halvings, which trigger supply changes and subsequent market fluctuations. Explained why, how and when, these long-term price cycles happen, and periodically verifying the cycle's recurrence through price action, this paper suggests two primary investment strategies based on time horizon:

1. **Long-term (over 4-5 years)**: holding cryptocurrencies over longer periods (more than an halving cycle) helps mitigate the effects of bull and bear market recurrence, allowing investors to focus on overall market growth.
2. **Short-to-medium-term (less than 4 years)**: here, the bull-bear cycle has a significant impact, making it crucial to time entries, the optimal entry period is at the end of the bear market or/and during recovery, and exits during or/and at the end of the bull market to maximize returns and minimize risks.

Additionally, this article highlights an increasing correlation between the crypto market and traditional financial markets, driven by the entry of institutional investors and the expansion of cryptocurrency market capitalization. This alignment, along with the decreasing volatility of the largest cryptocurrencies, suggests that future bull-bear cycles could be less pronounced, especially for the top capitalized assets.
In summary, the integration of cryptocurrencies into diversified portfolios represents a promising opportunity but requires careful attention to asset selection, time horizon, and market context, thereby balancing potential returns with risk management.

References

[1] **Acerbi, C. and Tasche, D.** *On the use of the Conditional Value-at-Risk for the measurement of risk*. Journal of Banking & Finance, 26(7), pp. 1467-1484, 2002.

[2] **Almeida, J. and Gonçalves, T. C.** *A systematic literature review of investor behavior in the cryptocurrency markets. Journal of Behavioral and Experimental Finance*, 37, 100785, 2023.

[3] **Ammous, S.** *The Bitcoin Standard: The Decentralized Alternative to Central Banking*. Hoboken, NJ: Wiley, 2018.

[4] **Baur, D. G., Hong, K. H. and Lee, A. D.** *Bitcoin: Medium of exchange or speculative assets?*. Journal of International Financial Markets, Institutions and Money, Vol. 54, 2018.

[5] **Będowska-Sójka, B. and Kliber, A.** *Is there one safe-haven for various turbulences? The evidence from gold, Bitcoin and Ether. North American Journal of Economics and Finance*, Vol. 56, pp. 101390, 2021.

[6] **Braga, M. D.** *Risk-based approaches to asset allocation: Concepts and practical applications*. Springer, 2015.

[7] **Corbet, S., Lucey, B., Urquhart, A. and Yarovaya, L.** *Cryptocurrencies as a financial asset: A systematic analysis*. International Review of Financial Analysis, Vol. 62, pp. 182-199, 2019.

[8] **De Miguel, V., Garlappi, L. and Uppal, R.** *Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?* The Review of Financial Studies, Vol. 22(5), pp. 1915-1953, 2007.

[9] **Dyhrberg, A. H.** *Bitcoin, gold and the dollar - A GARCH volatility analysis*. Finance Research Letters, Vol. 16, pp. 85-92, 2016.

[10] **Fang, L., Bouri, E., Gupta, R. and Roubaud, D.** *Does global economic uncertainty matter for the volatility and hedging effectiveness of Bitcoin?* International Review of Financial Analysis, Vol. 61, pp. 29-36, 2019.

[11] **Goodell, J. W. and Goutte, S.** *Diversifying equity with cryptocurrencies during COVID-19*. International Review of Financial Analysis, 76, 2021.

[12] **Guesmi, K., Saadi, S., Abid, I. and Ftiti, Z.** *Portfolio diversification with virtual currency: Evidence from bitcoin*. International Review of Financial Analysis, Vol. 63, pp. 431-437, 2019.

[13] **Huang, X., Han, W., Newton, D., Platanakis, E., Stafylas, D. and Sutcliffe, C.** *The diversification benefits of cryptocurrency asset categories and estimation risk: Pre and post Covid-19*. The European Journal of Finance, Vol. 29(7), pp. 800-825, 2022.

[14] **Ji, Q., Bouri, E., Lau, C. K. M. and Roubaud, D.** *Dynamic connectedness and integration in cryptocurrency markets*. International Review of Financial Analysis, Vol. 63, pp. 257-272, 2019.

[15] **Ji, Q., Zhang, D. and Zhao, Y.** *Searching for safe-haven assets during the COVID-19 pandemic*. International Review of Financial Analysis, Vol. 71, pp. 101526, 2020.

[16] **Klein, T., Pham Thu, H. and Walther, T.** *Bitcoin is not the New Gold - A comparison of volatility, correlation, and portfolio performance*.

International Review of Financial Analysis, Vol. 59, pp. 105-116, 2018.

[17] **Rockafellar, R. T. and Uryasev, S.** *Optimization of Conditional Value-at-Risk*. Journal of Risk, Vol. 2(3), pp. 21-41, 2000.

[18] **Rockafellar, R. T., Uryasev, S. and Zabarankin, M.** *Generalized Metrics of Risk for Portfolio Optimization*. Risk Theory and Decision Theory, Vol. 12(1), pp. 1-34, 2006.

**Portfolio Optimization:
a New Frontier through VAE-LSTM-Based
Reinforcement Learning**

About the Authors

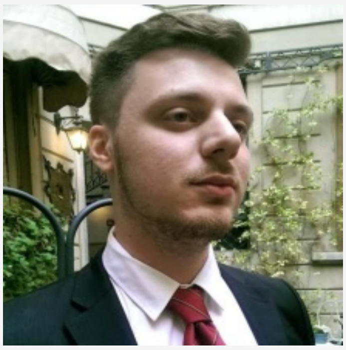

Giacomo Colombo:
Senior Manager
He holds a master's degree in Economics and Finance, with a specialisation in Corporate Finance. After two years in the M&A sector, he moved on to Banking field, with particular focus on Market contents. He is highly specialised in data modelling, methodology and automated computations. In his current role as Senior Manager, he oversees several tasks at the largest Italian Banks and follows numerous Asset Management streams, also collaborating with an important University.

Giuseppe Crupi:
Quant Analyst
He graduated in Data Physics from the Università Statale di Milano. His thesis, "Development of a Machine Learning framework for Anomaly Detection in Financial Time Series", demonstrates how machine learning and finance can coexist to drive future advancements. In 2023, he joined Iason working on projects focused on Market Risk with one of the major Italian banks.

Giacomo Giannoni:
Senior Financial Engineer
He holds a master's degree in Mathematical Engineering - Quantitative Finance from Politecnico di Milano. He joined iason in 2020 working on projects mainly focused on Market Risk, as part of the Risk Management of one of the major Italian banks. He is currently involved in projects regarding the development and implementation of innovative solutions and methodologies for the computation and analysis of multiple risk measures.

Raffaele Paradiso:
Financial Engineer
He holds a Bachelor's degree in Physics and is currently pursuing a Master's degree in the Physics of Complex Systems. As a quantitative analyst, he has spent the past year working in market risk with a major banking player, focusing on analyzing and managing complex financial risk models.

This article was written in collaboration with Andrea Cannizzo, who at the time was working for Iason Consulting.

Portfolio Optimization: a New Frontier through VAE-LSTM-Based Reinforcement Learning

Giacomo Colombo Andrea Cannizzo Giuseppe Crupi
Giacomo Giannoni Raffaele Paradiso

Portfolio allocation is a critical aspect of investment management, as it involves the strategic distribution of assets within a portfolio to maximize returns and minimize risk. Traders and portfolio managers commonly use various financial features and metrics to inform their portfolio allocation decisions. These features help evaluate the risk, return, and correlation between different assets, enabling more informed investment strategies. The integration of Machine Learning (ML) into portfolio allocation is transforming investment management, providing innovative approaches to optimize risk and return. With its capacity to handle large datasets, identify complex patterns, and adapt to evolving information, ML offers a powerful, dynamic alternative for portfolio strategy. In this work, a novel method that combines Neural Networks (NN) and Reinforcement Learning (RF) is proposed to enhance portfolio allocation. The approach involves introducing various financial features typically used by traders to determine whether to buy or sell an asset and then utilizing a VAE-LSTM model to condense this information into a compact set of features. The Variational Autoencoder (VAE) component captures the underlying structure of the data, while the Long Short-Term Memory (LSTM) identifies temporal patterns. In the second step, we concatenate these features with the original dataset to predict optimal portfolio weights over a specified period using fundamental RF techniques, such as Actor to Critic (A2C), Proximal Policy Optimization (PPO) and Deep Deterministic Policy Gradient (DDPG). The results demonstrated a clear advantage of dynamic allocation strategies with RF techniques over traditional static weight portfolios, improving both returns and risk management. The algorithms employed in this study consistently outperformed classic methods, providing superior drawdown control. However, they exhibited a more conservative approach, which made them less effective compared to the maximum Sharpe ratio method.

Portfolio management plays a vital role in financial services, involving the allocation of funds across diverse assets to generate uncorrelated returns, minimize risk, and control operational costs. Portfolios may span multiple asset classes, like cash, bonds, and equities, or focus within a single asset class, such as optimizing a mix of stocks in an equity portfolio. Investors typically optimize based on performance metrics, aiming to maximize returns relative to risk. Since Markowitz's (1952) introduction of Modern Portfolio Theory [22], substantial advancements have been made in portfolio optimization, including refining optimization methods and introducing constraints aligned with rational investment principles. In recent years, ML and RL have significantly impacted portfolio optimization. More specifically, ML has enhanced feature selection, forecasting, and the estimation of asset means and covariances, while gradient-based methods have become integral to optimization. Concurrently, RL and, later, Deep Reinforcement Learning (DRL) have gained traction in fields like gaming, robotics, and natural language processing, leading to financial applications in automated stock trading, risk management (Buehler et al., 2019, [8]), and portfolio optimization. Especially in the portfolio optimization framework, while DRL-based strategies have shown impactful improvements over conventional approaches, they also embed some structural limitations that need to be addressed and solved (e.g. producing discrete trading signals that restrict broader applicability).

For the sake of argument, Markowitz's mean-variance framework established a new standard for efficiency while also showing a notable limitation where the performance of the optimization strategies sometimes yielded non-optimal portfolios. Even though many other alternative models have been proposed during the last decade, they are still generally challenging to implement in practice due to computational demands and scalability issues. While feasible for small asset universes, these models become intractable with larger portfolios, requiring considerable computational resources and expertise. Machine learning faced similar challenges historically, with early models constrained by limited optimization capabilities. However, advancements in the late 1990s provided the foundation for handling large-scale, high-dimensional data, transforming ML applications. This paper examines how modern large-scale optimization algorithms can improve portfolio allocation.

This document is structured as follows: The literature review provides an overview of existing research on asset allocation with ML techniques, in the Theoretical Models section, the models used in this paper are introduced, which include the VAE, the LSTM, and Deep Reinforcement Learning. This section outlines the core frameworks that underpin the analysis, highlighting their relevance and applications. Each model is briefly discussed in order to providing the necessary insights. In the Methodological section, both the dataset and the methodologies used are introduced along with the novel approach to the portfolio optimization strategy. Finally, in the last two chapters, the results of the optimization and the possible future works are shown.

Literature Review

The goal of our research is to develop a novel methodology for portfolio allocation, which is a critical aspect of investment management that focuses on the strategic distribution of assets to optimize returns while minimizing risk.
The incorporation of ML into portfolio allocation is evolving investment practices by offering advanced techniques to enhance risk-return optimization. Having the ability to process vast datasets, identify complex patterns, and adapt to rapidly changing market conditions, ML offers a robust and dynamic approach for designing alternative efficient portfolio strategies.
As a result, in recent years, there has been a significant increase in the use of deep learning applications for portfolio asset allocation. Starting from NN, the work conducted by Chi-Ming Lin, Jih-Jeng Huang, Mitsuo Gen and Gwo-Hshiung Tzeng (2006) [10], provides an example of a dynamic portfolio selection model, incorporating Recurrent Neural Network (RNN) and cross-covariance matrices. Other deep neural networks models can be found in the works of Fabio D. Freitas, Alberto F. De Souza, Ailson R. de Almeida (2009) [11], where they propose a prediction-based portfolio optimization model, using NN predictors model, or in the research conducted by Seyed Taghi Akhavan Niaki and Saeid Hoseinzade (2013) [23] where they used Artificial Neural Network (ANN) to forecast the daily direction of Standard Poor's 500 (SP 500) index, as well as in the research of J. B. Heaton, N. G. Polson and J. H. Witte (2016) [13], where the authors explored the use of DL hierarchical models for prediction and classification in financial problems.
Following the introduction of RL based approaches, in the work conducted by Zhengyao Jiang and Jinjun Liang (2017) [15], they explored the usage of Convolutional Neural Network (CNN) in the optimization of the weights of a cryptocurrency portfolio. Other works on RL approaches include the research of Saud Almahdi and Steve Y. Yang (2017) [1] where the authors extended the existing work in Recurrent Reinforcement Learning (RRL) building an optimal variable weight portfolio allocation under a coherent downside risk measure, as well as the work conducted by S. Obeidat, D. Shapiro, Mathieu Lemay, M. K. MacPherson and M. Bolic (2018) [24] where they try to predict returns using LSTM neural networks. Following their work on CNN (2017) [15], Zhengyao Jiang, Dixing Xu and Jinjun Liang (2017) [16] investigated the impact, in portfolio management, of different approaches, such as CNN, LSTM and RNN.
Finally, covering the state-of-the-art reinforcement learning algorithm, such as, A2C, PPO, Policy Gradient (PG) and DDPG, the work conducted by Zhipeng Liang, Hao Chen, Junhao Zhu, Kangkang Jiang and Yanran Li (2018) [19] provides a practical example for these algorithms, where they implemented and compared DDPG, PPO and PG in portfolio management, while in their research, Hongyang Yang, Xiao-Yang Liu, Shan Zhong and Anwar Walid (2020) [29], combined PPO, A2C and DPPG to create a new trading strategy by maximizing investment return.

Theoretical Models

This chapter discusses key concepts and techniques in modern artificial intelligence and machine learning. It begins with an exploration of Artificial Neural Networks (ANNs), which serve as the foundation for many AI models allowing machines to learn from data through structures that mimic the human brain. Next, it covers VAEs, a type of generative model that represents complex data distributions, and LSTM networks, designed to handle sequential data and capture long-term dependencies. The discussion then shifts to RL, where agents learn to make decisions through trial and error by interacting with their environments. Within RL, advanced techniques such as A2C methods are examined, which combine the advantages of value-based and policy-based learning. Finally, the chapter reviews cutting-edge algorithms like PPO and DDPG, both of which enhance the stability and efficiency of training for complex tasks in RL.

Artificial Neural Network

ANNs are nonlinear models that draw inspiration from the structure and functioning of the human brain. By combining simple linear and nonlinear functions, ANNs can tackle complex problems effectively. The structure of an ANN consists of interconnected nodes, or neurons, organized in layers. The input layer receives the data, which is then processed through one or more hidden layers before reaching the output layer. Each neuron in a layer is connected to neurons in the subsequent layer, and these connections are assigned weights that are adjusted during the training process. ANNs learn from data using learning algorithms, such as backpropagation, which enable them to identify patterns and relationships within the data itself. Through this training process, ANNs can function as "universal approximators", capable of modeling complex data without requiring formal assumptions, unlike traditional statistical methods [25]. One of the key advantages of ANNs is their ability to outperform traditional linear and nonlinear models, particularly in pattern recognition tasks, even when dealing with noisy data. They are well-suited for predictive tasks, such as stock market forecasting, where they have consistently outperformed regression models in accuracy [5, 9]. An example of a NN is illustrated in Figure 7, where: in image *a*, an activation function is depicted, and in image *b*, a feed-forward fully connected neural network is shown. In this work, two specific types of ANNs are employed: VAE and LSTM. VAEs are a type of generative model that learns a compressed representation of the input data, while LSTMs are a variant of RNNs designed to capture long-term dependencies in sequential data.

Variational Autoencoder

A VAE is a type of generative model that combines neural networks with probabilistic approaches to learn complex data distributions. Introduced by Kingma and Welling [17], the VAE extends the traditional autoencoder by using a probabilistic encoder-decoder framework.
Instead of mapping an input directly to a compressed latent space, VAEs represent the latent variables as probability distributions, usually Gaussian and then sample from them. This allows for the generation of new data by sampling from these latent distributions, which makes VAEs highly useful in applications like image generation, anomaly detection, and dimensionality reduction. One key feature of VAEs is the use of the reparameterization trick, which enables backpropagation through stochastic nodes, a challenge in earlier generative models. The model optimizes a loss function that includes both reconstruction loss (how well the input is reconstructed) and a regularization term, usually the Kullback-Leibler (KL) divergence [18], which ensures the latent space distribution is close to a known prior, typically a standard normal distribution. Because of their ability to

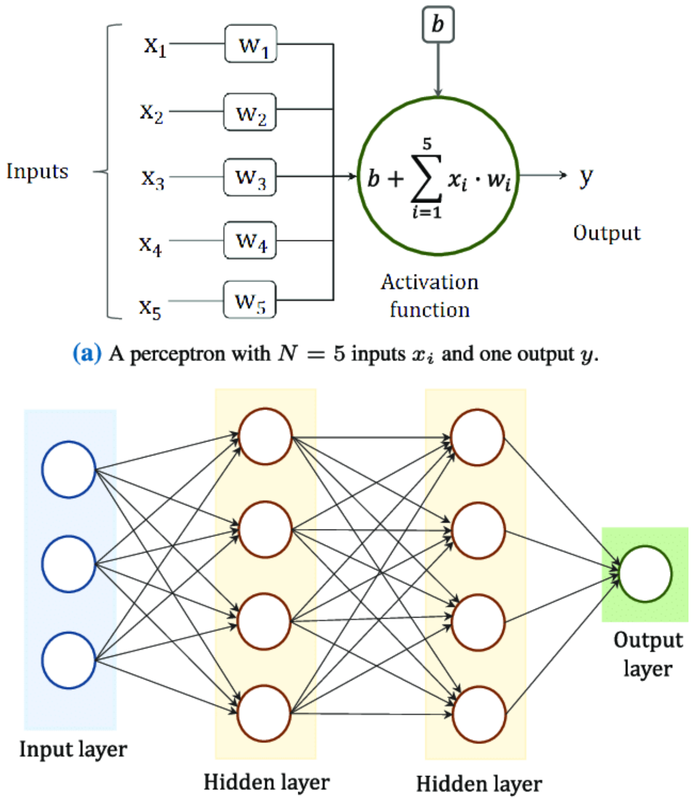

(a) A perceptron with $N = 5$ inputs x_i and one output y.

(b) Feed-forward fully connected neural network.

FIGURE 7: In image a is depicted an example of activation function which takes as input 5 variable, x_i and, after elaborating it with an activation function, it gives as output y. In image b an example of NN structure is shown.

generate diverse and realistic data, VAEs have been widely adopted in tasks such as image generation and natural language processing.
Stated more formally, for $x \in X$ and $z \in Z$ with z a vector of latent variables generally with values taken from $\mathbb{R}$, the likelihood $p(x)$ sometimes called the evidence or the marginal likelihood (Bayes' Rules) can be defined as the marginalization over the joint probability w.r.t. latent variable:

$$p(\boldsymbol{x}) = \int_z p(\boldsymbol{x}, \mathbf{z}) d\mathbf{z} = \int_z p(\boldsymbol{x}|\mathbf{z}) p(\mathbf{z}) d\mathbf{z},$$

or, using the chain rule of probability, it can be written as:

$$p(\boldsymbol{x}) = \frac{p(\boldsymbol{x}, \mathbf{z})}{p(\mathbf{z} \mid \boldsymbol{x})}.$$

Computing the probability of the evidence is frequently intractable because it needs to evaluate this integral over all latent variables Z.
To obtain a tractable $p(\boldsymbol{x})$ a tractable $p(\mathbf{z} \mid \boldsymbol{x})$ is needed, and to have a tractable $p(\mathbf{z} \mid \boldsymbol{x})$ a tractable $p(\boldsymbol{x})$ is required.

$$p(\mathbf{z} \mid \boldsymbol{x}) = \frac{p(\boldsymbol{x}, \mathbf{z})}{p(\boldsymbol{x})}.$$

An attempt can be made to approximate the true posterior parametrized by θ with an approximate posterior parametrized by φ.

$$p_\theta(\mathbf{z} \mid \boldsymbol{x}) \approx q_\varphi(\mathbf{z} \mid \boldsymbol{x}).$$

Following the above approximations, the core of VAEs lies in the optimization of the Evidence Lower Bound (ELBO).

$$\log p_\theta(\boldsymbol{x}) = \underbrace{E_{q_\varphi(\mathbf{z}|\boldsymbol{x})}\left[\log \frac{p_\theta(\boldsymbol{x}, \mathbf{z})}{q_\varphi(\mathbf{z} \mid \boldsymbol{x})}\right]}_{\text{ELBO}} + \underbrace{D_{KL}\left(q_\varphi(\mathbf{z} \mid \boldsymbol{x}) \| p_\theta(\mathbf{z} \mid \boldsymbol{x})\right)}_{\geq 0}.$$

Since the KL divergence is non-negative, the ELBO forms a lower bound on $\log p_\theta(x)$. The KL divergence determines two distances:

- The distance of the approximate posterior from the true posterior;
- The gap between the ELBO and the marginal likelihood $\log p_\theta(\mathbf{x})$.

Splitting the expectation in the ELBO:

$$L(\theta, \varphi, x) = \underbrace{E_{q_\varphi(\mathbf{z}|\mathbf{x})}\left[\log p_\theta(\boldsymbol{x} \mid \mathbf{z})\right]}_{\text{Reconstruction Term}} \underbrace{- D_{KL}\left(q_\varphi(\mathbf{z} \mid \boldsymbol{x}) \| p_\theta(\mathbf{z})\right)}_{\text{Regularization Term}}.$$

The goal is differentiate and optimize the lower bound $L(\theta, \varphi, x)$ w.r.t. variational parameters φ and generative parameters θ. However, the gradient of the lower bound w.r.t. φ is a bit problematic ($\nabla_\varphi E_{q_\varphi(\mathbf{z}|\mathbf{x})}\left[\log p_\theta(\mathbf{x} \mid \mathbf{z})\right]$).
This estimator is unbiased, on average it will converge to the true expectation, but it is stochastic and it has a high variance and it's not possible to run a back propagation through it.
To make the ELBO differentiable with respect to φ, VAEs use the reparameterization trick, instead of sampling directly

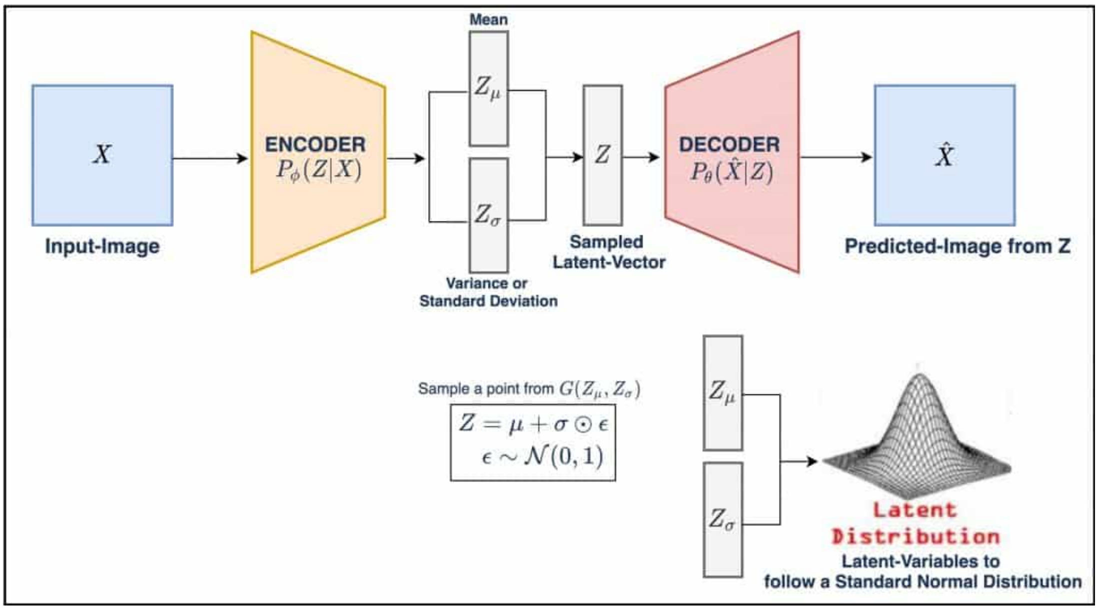

FIGURE 8: Example of VAE structure: the NN takes as input the dataset X and with the encoder it extrapolate the mean Z_μ and the standard deviation Z_σ, which are used to sample the latent vector Z. After that the decoder give as output the predicted dataset $\hat{X}$.

from $q_\varphi(\mathbf{z} \mid \boldsymbol{x})$, samples are drawn from a fixed distribution $\epsilon \sim p(\epsilon)$ and then transformed:

$$z = g_\varphi(\epsilon, x).$$

For a Gaussian approximate posterior, this might look like:

$$z = \mu_\varphi(x) + \sigma_\varphi(x) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I).$$

Long Short-Term Memory

LSTM is a type of RNN architecture specifically designed to address the problem of long-term dependencies in sequence data, which standard RNNs struggle with due to vanishing or exploding gradients. LSTMs were introduced by Hochreiter and Schmidhuber [14] and have since become widely used in tasks such as natural language processing, speech recognition, and time series prediction.
LSTMs are structured with special units called memory cells that maintain information over long periods. These cells are controlled by three key gates: the input gate, which decides what information is updated in the cell; the forget gate, which controls what information to discard; and the output gate, which determines the next output from the cell, the overall structure is showed in Figure 9. This gating mechanism allows LSTMs to effectively retain and update relevant information over time, making them superior to standard RNNs for handling long sequences.
The gates are controlled by specific mathematical functions:

- The forget gate determines which information in the memory cell should be discarded. It takes in the previous hidden state h_{t-1} and the current input x_t, applies a sigmoid activation function, and outputs a value between 0 and 1:

  $$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f), \tag{1}$$

  where W_f is the weight matrix and b_f is the bias. A value closer to 0 indicates forgetting information, while a value closer to 1 retains it.

- The input gate decides what new information should be added to the memory cell. It also uses a sigmoid function to filter the incoming information:

  $$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i). \tag{2}$$

  Simultaneously, a candidate memory $\tilde{C}_t$ is created using the tanh function to ensure values are between -1 and 1:

  $$\tilde{C}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_C). \tag{3}$$

  The memory cell C_t is then updated by combining the previous cell state C_{t-1} with the new information:

  $$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t. \tag{4}$$

- The output gate determines the output of the cell for the current time step. Like the other gates, it uses a sigmoid function to produce a scaling factor:

  $$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o). \tag{5}$$

  The final hidden state h_t is computed by applying the output gate to the updated memory cell state C_t:

  $$h_t = o_t * \tanh C_t. \tag{6}$$

Reinforcement Learning

RL has its roots in the early days of cybernetics, as well as in developments across statistics, psychology, neuroscience, and computer science. Over the past decade, it has gained significant attention in the fields of machine learning and artificial intelligence. Its appeal lies in the potential to train agents through rewards and punishments, without explicitly defining the steps required to complete the task [3, 4, 7].
Specifically, RL is a branch of ML where an agent learns to make decisions by interacting with its environment. Unlike supervised learning, where models learn from a fixed dataset of labeled examples, RL operates in a dynamic setting where the agent's actions influence the future states it will encounter. The core objective of the RL agent is, then, to maximize a cumulative reward over time by learning an optimal policy, which is a mapping from states to actions. This decision-making process is modeled through

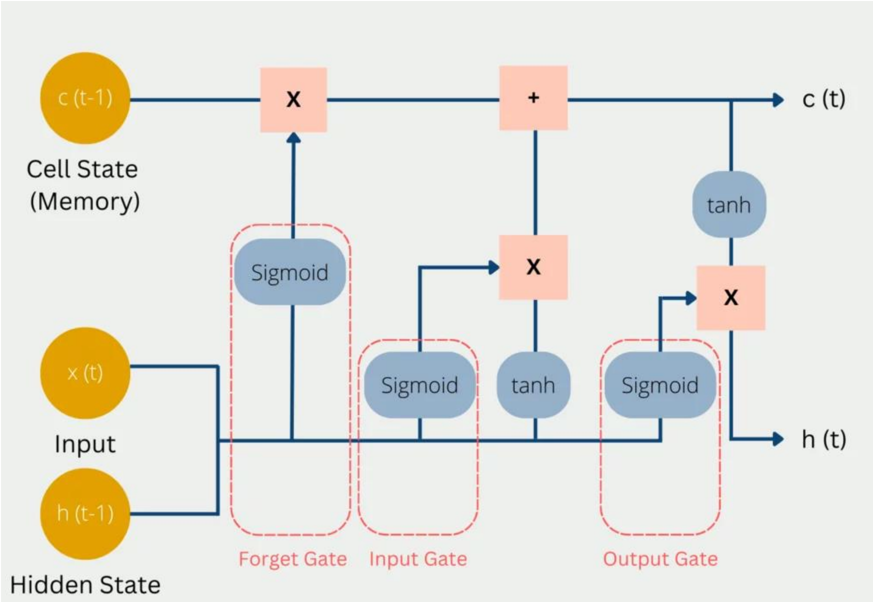

FIGURE 9: *Example of LSTM structure.*

trial and error, with the agent improving its behavior based on the feedback received in the form of rewards or penalties [26, 28].

In particular, in RL the agent interacts with an environment $\mathcal{E}$, which provides feedback in the form of a reward signal r_t, indicating how good the taken action a_t is at a particular state s_t, where t stands for a specific time step.

A classic model for RL problems is the Markov Decision Process (MDP) which can be summarized as follows. Consider a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ where $\mathcal{S}$ represents the set of states, $\mathcal{A}$ the set of actions, $\mathcal{P} = \mathcal{P}(s', s, a)$ is the transition probability function: the probability of moving from state s to s' by taking the action a, $\mathcal{R} = \mathcal{R}(s, a)$ is the reward function: the immediate reward received after taking action a in state s and $\gamma \in [0, 1]$ is the discount factor, determining how much future rewards are weighted. At each time step t, the agent observes a state $s_t \in \mathcal{S}$, selects an action $a_t \in \mathcal{A}$, transitions to a new state s_{t+1} based on $\mathcal{P}(s_{t+1}, s_t, a_t)$ and receives a reward r_t. The sequence of states, actions, and rewards is known as the agent's trajectory and the goal of RL is to learn a policy $\pi(a|s)$ that defines the agent's behavior by specifying the probability of taking action a in state s.

The return G_t is the total cumulative reward from time step t, defined as:

$$G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}, \tag{7}$$

where r_{t+k} is the reward at time step $t + k$ and γ is the discount factory that determines the importance of future rewards.

To evaluate policies, RL relies on value functions, which estimate the expected return starting from a state or a state-action pair. The two key value functions are the state value function $V^\pi(s)$ and the action value function $Q^\pi(s, a)$ that are respectively:

$$V^\pi(s) = \mathbb{E}_\pi[G_t | s_t = s] = \mathbb{E}_\pi\left[\sum_{k=0}^{\infty} \gamma^k r_{t+k} | s_t = s\right],$$

$$\begin{aligned} Q^\pi(s, a) &= \mathbb{E}_\pi[G_t | s_t = s, a_t = a] \\ &= \mathbb{E}_\pi\left[\sum_{k=0}^{\infty} \gamma^k r_{t+k} | s_t = s, a_t = a\right]. \end{aligned}$$

The goal in RL is to find the optimal policy $\pi*$ that maximizes the expected return from all states. The optimal value function $V*(s)$ and the optimal action value function $Q*(s, a)$, then, are those that satisfy the Bellman optimality equations.

On a historical note, RL has its foundations in behavioral psychology, where animals learn from the consequences of their actions. In RL, in fact, the agent faces a similar situation, where it explores various strategies, gradually learning which actions yield the highest rewards in the long run. This framework has gained in the past decade a lot of attention due to its ability to solve complex sequential decision-making problems, where explicit instructions for optimal behavior are not available, and the agent must learn from experience.

Actor to Critic

A2C is a specific RL algorithm that combines the strengths of two key approaches in RL: policy-based methods and value-based methods. The goal of A2C is to optimize an agent's behavior by learning both a policy (how to act) and a value function (how good a state or action is), leveraging the synergy between these two components to improve learning efficiency and performance. A2C is part of a broader class of Actor-Critic algorithms, which form the foundation for some of the most advanced RL techniques used in complex, high-dimensional environments.

As mentioned before, in reinforcement learning, the agent interacts with an environment to maximize cumulative rewards over time. Traditionally, value-based methods, such as Q-learning, estimate the expected future rewards (value) for each state or state-action pair, guiding the agent toward rewarding actions. On the other hand, policy-based methods directly optimize the policy by selecting actions that maximize rewards, without explicitly estimating the value of states or actions. Actor-to-Critic algorithms bridge these two approaches by employing two key components: the actor, which selects actions based on a policy, and the critic, which evaluates the quality of the chosen actions using a value function.

A2C, specifically, is a synchronous variant of the Actor-Critic framework and improves upon its predecessor, Asyn-

chronous Advantage Actor-Critic (A3C), by simplifying parallelization and ensuring more stable updates. This approach has been widely adopted due to its ability to balance exploration and exploitation while efficiently learning from large, continuous state spaces.
More specifically, Actor-Critic methods combine the following approaches:

- The Actor is responsible for selecting actions based on a policy $\pi_\theta(a|s)$, which is parametrized by θ and to adjust the policy to maximize long-term rewards.
- The Critic evaluates the actor's actions by estimating the value function, which predicts the expected return from each state-action pair. Furthermore, the Critic can either estimate the state value function $V(s)$ or the action-value function $Q(s,a)$.

By having the actor optimize the policy and the critic evaluate it, Actor-Critic methods form a loop where the actor improves using the feedback from the critic, and the critic, in turn, refines its evaluation based on the actor's performance.
Among some of the advantages of A2C methods, unlike value-based methods like Q-learning, which struggle with continuous action spaces, A2C methods naturally handle continuous actions since the actor directly outputs actions based on the policy. Moreover, A2C methods are more sample-efficient because they utilize both policy and value function information. The actor improves its policy based on the critic's evaluation, while the critic continually updates based on the actor's exploration of the environment.

Proximal Policy Optimization

PPO is a state-of-the-art reinforcement learning algorithm that addresses key challenges in training agents to interact effectively with even more complex environments. Proposed by OpenAI in 2017, PPO builds upon policy gradient methods and aims to improve both stability and performance in policy-based reinforcement learning while maintaining simplicity and computational efficiency. Due to these features, it is widely recognized for its ability to achieve high performance in a variety of environments, ranging from game-playing tasks like Atari and Go to robotics and autonomous control systems.
In this framework, the agent also maximizes the cumulative reward it receives from interactions with its environment and this is typically done by learning a policy, a function that maps states to actions. Policy gradient methods directly optimize the policy by adjusting it in the direction that increases expected rewards. However, traditional policy gradient algorithms suffer from instability due to large, uncontrolled updates during training, leading to poor performance and convergence issues. PPO addresses these challenges by incorporating mechanisms that constrain policy updates, preventing the agent from making overly large or destructive changes to its behavior.
Furthermore, PPO strikes a balance between flexibility and stability, using a surrogate objective function that penalizes drastic deviations from the current policy. This approach allows this new model to achieve robust results across a wide range of reinforcement learning tasks, making it one of the most popular algorithms in contemporary RL research and applications.
The core idea behind PPO is to maximize a clipped objective that keeps the new policy close to the old policy, ensuring more stable updates. To do so, PPO method aims to optimize a policy $\pi_\theta(a|s)$ to maximize the expected reward. This is done by updating the policy in the direction of the gradient of the expected reward:

$$J(\theta) = \mathbb{E}_{s_t, a_t \sim \pi_\theta}\left[\sum_{t=0}^{T} \gamma^t r_t\right], \tag{8}$$

where r_t is the reward at time step t and $\gamma \in [0,1]$ is the discount factor. Furthermore, PPO is based on tracking how the new policy π_θ differs from the old policy $\pi_{\theta\text{old}}$ and, to do so, the following ratio that measures how much the policy has changed at state s_t for action a_t is introduced

$$r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta\text{old}}(a_t|s_t)}. \tag{9}$$

In summary, as a reinforcement learning method, PPO method offers an effective way to optimize policies while maintaining stability in updates and, by using a clipped objective function, it prevents the policy from changing too drastically in a single update, which makes it more robust compared to earlier policy gradient methods.
It's also worth observing that, at the same time, this method encourages the agent to explore new states of the environment by adding an entropy bonus.

Deep Deterministic Policy Gradient

DDPG is a powerful algorithm designed to tackle environments with continuous action spaces. Developed by Lillicrap et al. in 2015 [20], DDPG merges the strengths of policy-based and value-based approaches to overcome the challenges of continuous control tasks, where traditional discrete action methods such as Q-learning are inadequate. By combining the deterministic policy gradient method with deep neural networks, DDPG also provides an efficient way to learn optimal policies for high-dimensional control problems, such as robotic manipulation, autonomous driving, and other real-world applications.
While traditional algorithms like Deep Q-Networks (DQN) have demonstrated significant success in discrete action spaces, their applicability is limited in tasks where the action space is continuous. For example, controlling the steering angle of a car or adjusting the joints of a robot requires precise and continuous values for each action dimension. DDPG addresses this gap by using an actor-critic architecture, where the actor directly learns a deterministic policy, and the critic evaluates this policy using a Q-function.
DDPG has become a foundational algorithm for continuous control in reinforcement learning, combining the stability of Q-learning with the flexibility of policy gradient methods. It is widely used in environments that require fine-grained control and has inspired a range of subsequent innovations in deep reinforcement learning.
Unlike stochastic policy gradient methods (such as PPO) where the policy outputs a probability distribution over actions, DDPG uses a deterministic policy. This means that given a state s the policy outputs a single specific action a, instead of a probability distribution over possible actions.
The actor network is parameterized by θ^μ and the policy is defined as $a = \mu_\theta(s)$, where $\mu_\theta(s)$ is the action determined by the actor for state s. The critic in DDPG is responsible for evaluating the actor's actions by estimating the Q-value, which represents the expected return from taking an action a in a state s and following the policy thereafter.
The critic network is parameterized by θ^Q and the Q-value is given by:

$$Q(s, a|\theta^Q). \tag{10}$$

The critic learns to minimize the difference between its predicted Q-values and the target Q-values using the Bellman equation for continuous control defined as

$$y_t = r_t + \gamma Q'(s_{t+1}, \mu'(s_{t+1})), \tag{11}$$

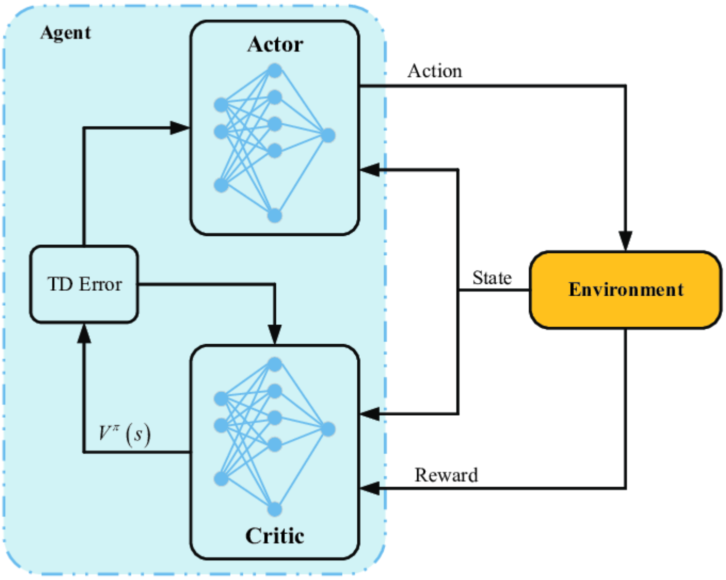

FIGURE 10: *Scheme of deep Actor to Critic reinforcement learning.*

where y_t is the target Q-value, r_t the reward, γ the discount factor and Q' and μ' refer to the so-called target networks. The main objective of a DDPG algorithm is to use an experience replay to break correlations in the training data and target networks to make the learning process more stable. Moreover, to ensure exploration, the algorithm adds noise to the actor's actions during training. The use of deterministic policy gradients, combined with the critic's guidance and the stability provided by the target networks, allows a DDPG algorithm to be an effective and powerful approach for solving complex continuous control tasks.

Methodologies

This chapter delves into the various methodologies employed in this research. It provides a comprehensive overview of the techniques and approaches that were implemented.
It begins by introducing the datasets used in this study, which consist of financial data for two distinct portfolios. The first represents a traditional optimized equity portfolio that mimics strategies employed by asset managers, while the second is based on the NASDAQ index, serving as a benchmark for comparative analysis. The datasets include a range of features derived from historical stock prices, enabling a thorough examination of the efficacy of our approach in both active and passive portfolio management contexts. Feature engineering plays a pivotal role in improving model performance, especially in financial applications where raw data might not fully capture underlying market dynamics. This chapter details the process of transforming raw market data into a set of meaningful technical indicators, which are essential for understanding trends, momentum, and other key attributes of financial assets. As the volume of features increases, it becomes crucial to reduce dimensionality while retaining the most informative characteristics. A VAE combined with a LSTM network is introduced to achieve this. This hybrid architecture enables the extraction of compressed, yet insightful, features that are subsequently used for further analysis.
The chapter also covers the design and training of reinforcement learning algorithms, focusing on A2C, DDPG, and PPO. These algorithms aim to learn the optimal portfolio allocation by simulating multiple scenarios and adjusting the asset weights dynamically. Each of these models employs custom loss functions tailored to the specific requirements of portfolio management, such as risk control, diversification, and minimization of trading costs.
The integration of RL techniques allows for the continuous adaptation of portfolio weights based on the evolving market environment, making these methods suitable for both short-term tactical adjustments and long-term strategic allocation. This chapter provides an in-depth explanation of the architectures, training procedures, and evaluation methods for each algorithm, illustrating how machine learning can enhance decision-making in the context of financial portfolio management.

Dataset Description

In this work, two distinct portfolios are used as starting points. To simplify the analysis, the portfolios are limited to equity stocks, although further extensions could be explored. The first portfolio is constructed using a combination of classical and innovative approaches employed by practitioners, representing an optimized equity portfolio commonly used by asset managers. The second portfolio represents the NASDAQ stock index, serving as a benchmark to compare performance against assets with updated weights. This choice allows for assessing the behavior of our methodology in two different contexts: active and passive portfolio management. A key objective of this research is to determine whether the developed algorithm performs more effectively in one or both approaches.
In the first approach, portfolio optimization follows a two-step methodology. First, strategic asset allocation is determined using an expert-based approach (qualitative step), where individual stocks are selected by evaluating financial market conditions, well-performing sectors, geographical regions, and stocks with the most attractive upside potential. After the stock selection, the weights of each stock are assigned using mean-variance optimization and resampling techniques to refine the allocation.
Due to privacy reasons, full disclosure of individual components is not provided. However, the key drivers of the portfolio are summarized as follows: The portfolio con-

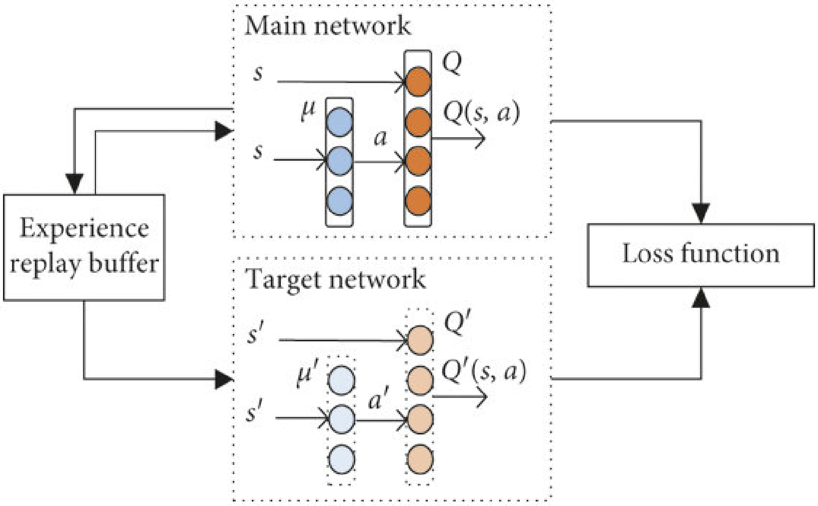

FIGURE 11: *Scheme of a Deep Deterministic Policy Gradient algorithm.*

sists of 15 instruments, with a strong focus on the Energy sector (40.76%), followed by Defense (19.76%) and Transportation (11.50%), as showed in Figure 12. In terms of geographical distribution, Western Europe dominates the allocation (66.23%), followed by North America (23.90%) and Emerging Markets (9.87%), as in Figure 13. Additionally, the portfolio has a significant emphasis on growth stocks, accounting for 46.72% of its composition, like Figure 14. Finally, some key performance indicators are reported:

- Sharpe Ratio: 2.385;
- Candidate Mean: 25%;
- Standard Deviation: 0.74%;
- Max Weight: 12.50%;
- Min Weight: 3.41%.

Each dataset includes information on assets identified by their tickers. For each day, the dataset records the following: closing price, high price, low price, opening price and volume, that spans from 2017 to 2024. As suggested by the name:

- Close price is the last price at which a security is traded during the regular market hours on a given trading day;
- High price is the highest price at which a security traded during a particular trading session;
- Low price is the lowest price at which a security traded during a specific trading session;
- Opening price is the first price at which a security is traded when the market opens for a given trading day;
- Volume represents the total number of shares, contracts, or units of a security traded during a day.

Since not all assets are quoted every day, one of the initial data manipulation steps involved extracting a specific set of business dates and carrying forward data from the previous day when a particular asset's data was missing for a given date. This ensures that each asset has consistent data entries for every business day.

As an additional step in data manipulation, since the prices are in different currencies and have varying magnitudes, percentage variations are used in this work. This can be expressed using the following formula:

$$r_{\%}(x_i, x_{i-1}) = \frac{x_i - x_{i-1}}{x_{i-1}}. \tag{12}$$

A result on how the final dataset would like is presented in Table 7.

As a further example, Figure 15 illustrates the trend of Closing Prices for four different assets. Conversely, Figure 16 displays the corresponding percentage returns. It is evident that, although the magnitude of *Stock1*'s closing prices is significantly higher than that of *Stock4*, in the tens, their returns are easily comparable when analyzed in percentage terms.

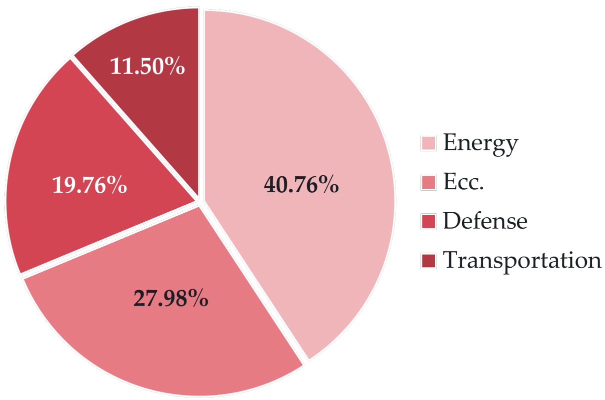

FIGURE 12: Pie chart representing the sector allocation of the first portfolio. Energy dominates with over 40% of the portfolio, followed by the defense sector at nearly 20%, and the transportation sector accounting for 11.50%.

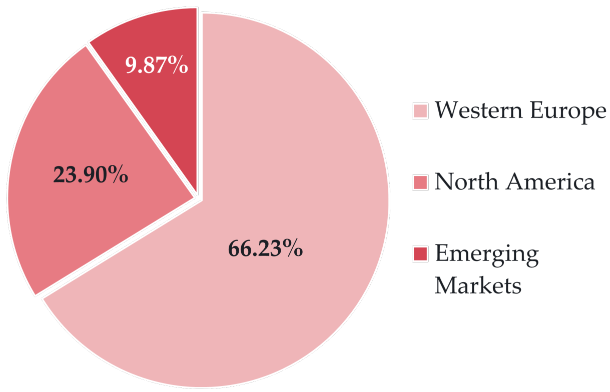

FIGURE 13: Pie chart representing the geographical distribution of the first portfolio. Western Europe leads with 66.23% of the allocation, followed by North America at 23.90%, and Emerging Markets at 9.87%.

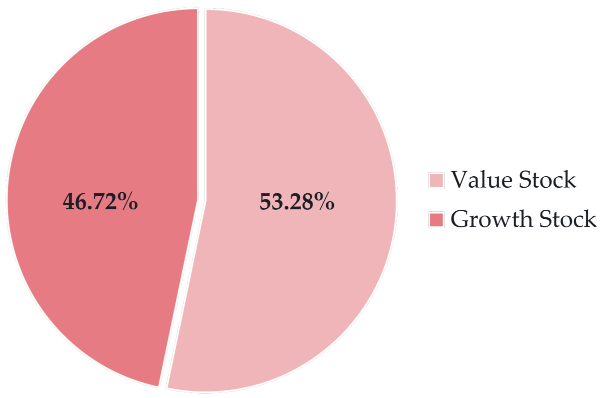

FIGURE 14: Pie chart representing the emphasis of our portfolio on growth stocks, as they take the 46.72% of the first portfolio.

Date	Ticker	Close	High	Low	Open	Volume
03/01/2017	Stock1	$r_{Close}^{0-Stock1}$	$r_{High}^{0-Stock1}$	$r_{Low}^{0-Stock1}$	$r_{Open}^{0-Stock1}$	$r_{Volume}^{0-Stock1}$
03/01/2017	Stock2	$r_{Close}^{0-Stock2}$	$r_{High}^{0-Stock2}$	$r_{Low}^{0-Stock2}$	$r_{Open}^{0-Stock2}$	$r_{Volume}^{0-Stock2}$
...	...	...	...	...	...	...
03/01/2017	Stock15	$r_{Close}^{0-Stock15}$	$r_{High}^{0-Stock15}$	$r_{Low}^{0-Stock15}$	$r_{Open}^{0-Stock15}$	$r_{Volume}^{0-Stock15}$
04/01/2017	Stock1	$r_{Close}^{1-Stock1}$	$r_{High}^{1-Stock1}$	$r_{Low}^{1-Stock1}$	$r_{Open}^{1-Stock1}$	$r_{Volume}^{1-Stock1}$
...	...	...	...	...	...	...
02/01/2021	Stock1	$r_{Close}^{i-Stock1}$	$r_{High}^{i-Stock1}$	$r_{Low}^{i-Stock1}$	$r_{Open}^{i-Stock1}$	$r_{Volume}^{i-Stock1}$
...	...	...	...	...	...	...
23/09/2024	Stock1	$r_{Close}^{N-Stock1}$	$r_{High}^{N-Stock1}$	$r_{Low}^{N-Stock1}$	$r_{Open}^{N-Stock1}$	$r_{Volume}^{N-Stock1}$
...	...	...	...	...	...	...
23/09/2024	Stock15	$r_{Close}^{N-Stock15}$	$r_{High}^{N-Stock15}$	$r_{Low}^{N-Stock15}$	$r_{Open}^{N-Stock15}$	$r_{Volume}^{N-Stock15}$

TABLE 7: Table with an example of the structure of the dataset used in this paper. It's composed of seven columns, the first one containing the business date, which goes from 03/01/2017 to 23/09/2024, rather the second one contains the list of the assets in the portfolio, named as Stock n. From the third to the seventh columns the percentage returns of Close, High, Low, Open and Volume are presented.

FIGURE 15: An illustration of the historical close prices for four equity stocks analyzed in our dataset.

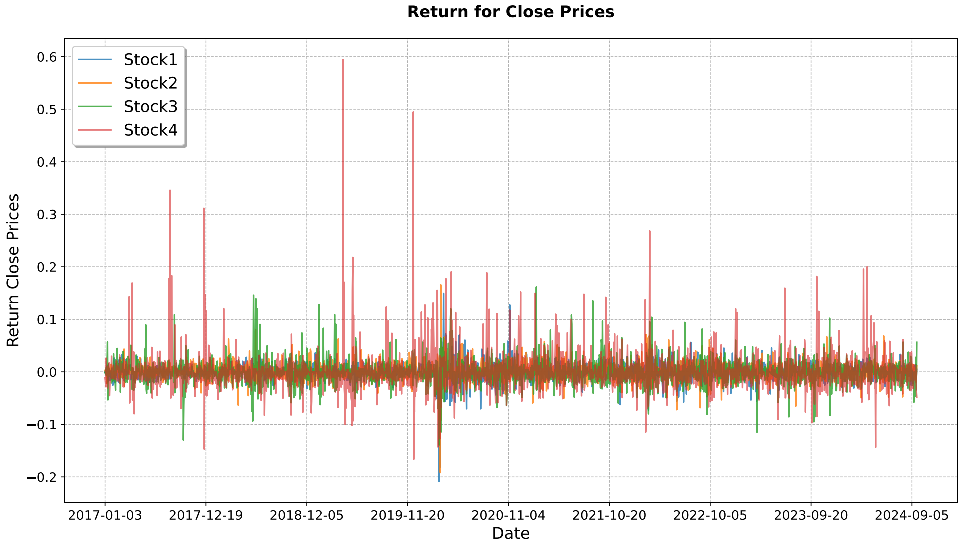

FIGURE 16: An illustration of the historical returns of close prices for four equity stocks analyzed in our portfolio.

Feature Engineering

Feature engineering is a critical step in the ML pipeline that involves transforming raw data into meaningful features that enhance the performance of predictive models. The process includes selecting, modifying and creating new features to capture the underlying patterns in the data more effectively.
In portfolio management, feature engineering is essential for creating quantitative metrics that can evaluate risk and optimize allocation decisions. By transforming raw market data into meaningful and relevant features, feature engineering enables financial models, including ML and DL algorithms, to better capture relationships between different financial variables and improve overall portfolio performance.
To enhance the model's portfolio allocation capabilities, several technical indicators are incorporated into the dataset. These indicators are widely used by traders and portfolio managers to this day, as they provide signals based on historical price and volume data, helping identify trends, reversals, momentum, and potential entry and exit points.
The features added include:

1. Volatility Avarage True Range (ATR), a technical indicator used to measure market volatility. It specifically captures the degree of price movement for an asset over a specified period of time. It is particularly useful for understanding how much an asset's price typically fluctuates on a day-to-day basis, regardless of the direction (up or down) [27]. It is derived from the True Range (TR), which captures the range of price movement over a given period. The TR is defined as the greatest of the following three values:

    - The difference between the current high and low (current trading range);
    - The absolute difference between the current high and the previous close;
    - The absolute difference between the current low and the previous close.

    This method ensures that the TR accounts for any gaps or sudden spikes in pricem making it more accurate measures of volatility. The ATR is then computed as the moving average of the TR over a specified number of periods:

$$ATR = \frac{1}{n}\sum_{i=1}^{n} TR_i. \tag{13}$$

    A plot of the ATR for one of the stock in our portfolio is shown in Figure 17.

2. Bollinger Band Width (BBW) measures the relative width between the upper and lower Bollinger Bands, providing insights into the volatility of an asset's price over time [6].

It consists of three main components:

- *Middle Band*: The Simple Moving Average (SMA) of the asset's price;
- *Upper Band*: This is calculated as the Middle Band plus a multiple (typically 2) of the standard deviation of the price;
- *Lower Band*: This is calculated as the Middle Band minus the same multiple of the standard deviation.

Then the BBW is defined as the difference between the upper and lower bands, normalized by dividing by the Middle Band. The formula is:

$$BBW = \frac{UpperBand - LowerBand}{MiddleBand} \times 100. \tag{14}$$

An example is proposed in Figure 18.

3. On-Balance Volume (OBV) is a technical analysis indicator that uses volume flow to predict changes in stock prices. It measures the cumulative buying and selling pressure by summing up the volume on up days and subtracting it on down days. The idea is that volume precedes price movement, meaning that significant changes in volume can signal a potential price movement in the same direction [12]. The formula for OBV is straightforward:

    - If the current closing price is higher than the previous period's closing price: $OBV = PreviousOBV + CurrentPeriodVolume$;
    - If the current closing price is lower than the previous period's closing price: $OBV = PreviousOBV - CurrentPreiodVolume$;

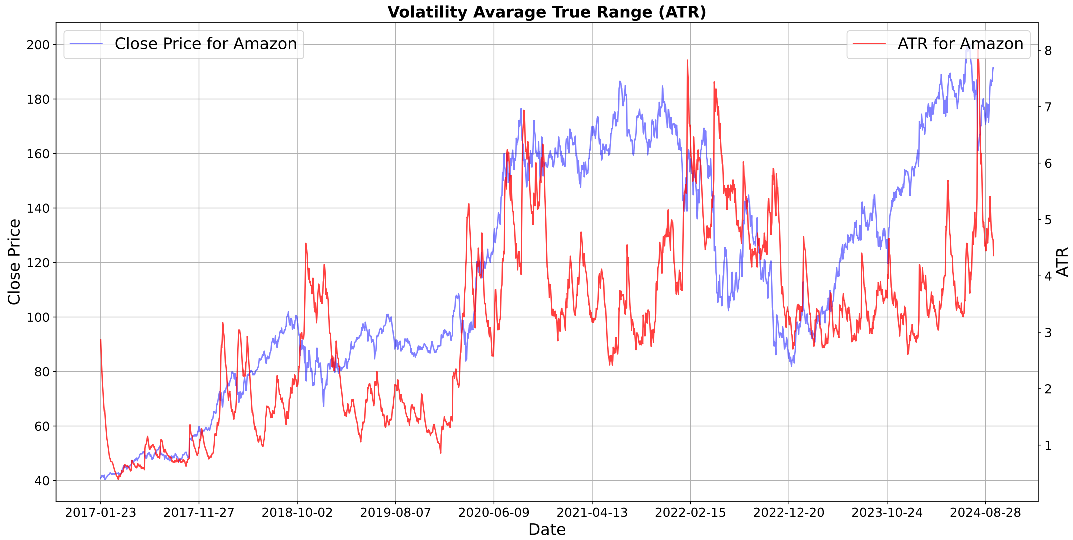

FIGURE 17: In the figure is presented a graph with the Close Price (blue line) and the correspondent Volatility Avarage True Range (red line). The left y-axis represents the scale for the Close Price, while the right y-axis corresponds to ATR values. High ATR indicates high market volatility. This can occur during strong price movements, either up or down, or when the market experiences sharp price swings. While low ATR suggests lower volatility, which is typical during consolidation periods when the price tends to move within a narrow range.

- If the current closing price is equal to the previous period's closing price: $OBV = PreviousOBV$.

An example is shown in Figure 19.

4. Chaikin Money Flow (CMF) measures the accumulation and distribution of money flow over a specified period by evaluating both price movement and volume. CMF helps traders identify the buying and selling pressure on an asset, which can be used to predict potential trends or reversals [21]. The calculation of CMF involves three main steps:

   - Calculate the Money Flow Multiplier (MFM) for each period:

   $$MFM = \frac{(Close - Low) - (High - Close)}{High - Low}. \tag{15}$$

   The MFM ranges from -1 to 1 and reflects where the close is relative to the high and low of the period. If the close is near the

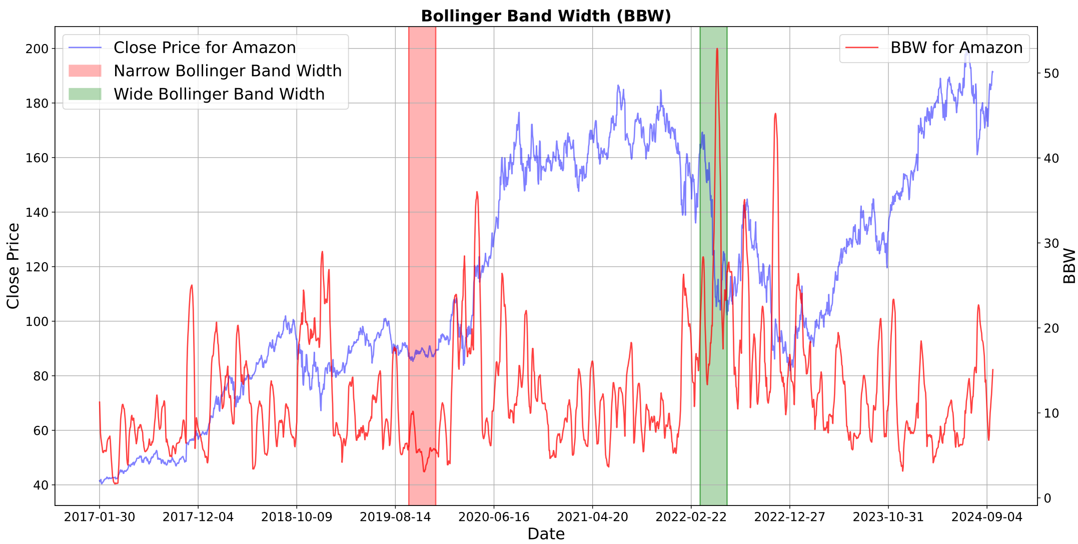

FIGURE 18: In the figure is presented a graph with the Close Price (blue line) and the Bollinger Band Width (red line). The left y-axis represents the scale for the Close Price, while the right y-axis corresponds to BBW prices. A narrowing of the Bollinger Band Width indicates low volatility and a period of consolidation, like the one identified by the red area in the plot. This is often a precursor to a significant price move or breakout, as low volatility periods tend to be followed by high volatility. Rather, a widening of the Bollinger Band Width suggests high volatility and usually indicates a strong price trend, either bullish or bearish. This can also signal that the current trend is approaching exhaustion, and a reversal or retracement may be imminent.

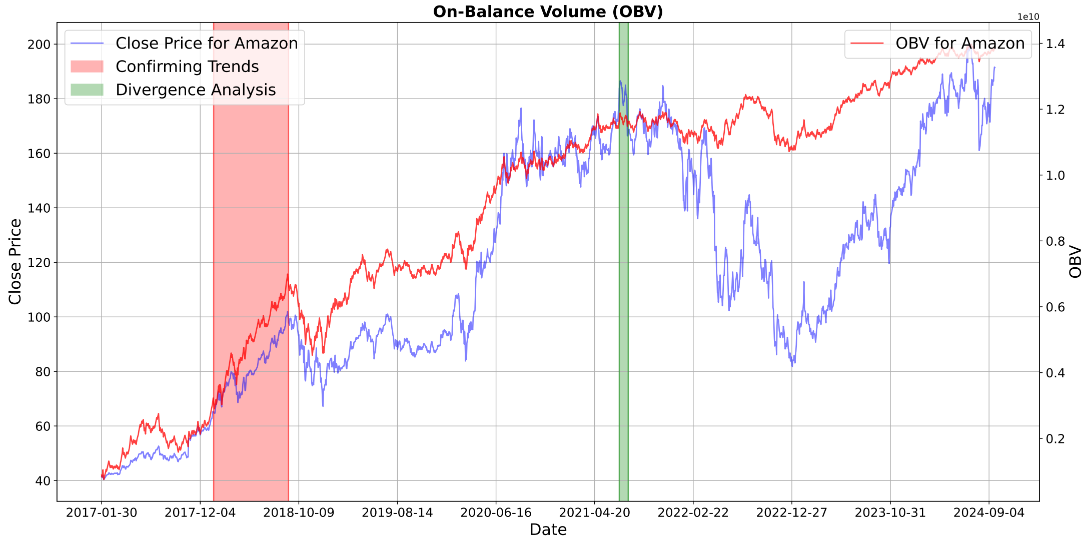

FIGURE 19: *The figure displays a graph with the Close Price (blue line) and the On-Balance Volume (OBV) (red line). The left y-axis shows the scale for the Close Price, while the right y-axis corresponds to OBV values. OBV can be used to confirm the direction of a trend. For instance, if the price is rising and the OBV is also increasing, as seen in the red-highlighted period, it suggests that the uptrend is supported by strong buying volume, making it more likely to continue. Conversely, a divergence between the price and OBV, as shown in the green area, may indicate a potential trend reversal. If the price is making new lows while the OBV remains flat or increases, it suggests weakening buying pressure, signaling that the uptrend could be losing momentum.*

high, the multiplier will be positive, indicating buying pressure. If the close is near the low, the multiplier will be negative, indicating selling pressure;

- Calculate the Money Flow Volume (MFV):

$$MFV = MFM \times Volume. \tag{16}$$

- Calculate the CMF:

$$CMF = \frac{\sum_{i=1}^{n} MFV_i}{\sum_{i=1}^{n} Volume_i}. \tag{17}$$

An example of this statistic is proposed in Figure 20.

5. The Moving Average Convergence Divergence (MACD) is a popular trend-following momentum indicator in technical analysis, used to identify potential changes in the strength, direction, momentum, and duration of a trend in a stock's price. It does this by comparing two moving averages of a security's price [2].

MACD consists of three main elements:

- The MACD Line, which is calculated as the difference between a short-term Exponential Moving Average (EMA) and a long-term EMA. $MACDLine = EMA_{12} - EMA_{26}$;
- The Signal Line, which is a 9-period EMA of the MACD Line. It smooths out the MACD line, making it easier to spot potential buy or sell signals. $SignalLine = EMA_9(MACDLine)$;
- The MACD Histogram, which is the difference between the MACD Line and the Signal Line. It visually represents the convergence or divergence between the two lines: $MACDHistogram = MACDLine - SignalLine$.

An example plot is shown in Figure 21.

6. The Average Directional Index (ADX) is a technical indicator used to measure the strength of a trend in a market, regardless of its direction. The ADX helps traders determine whether a market is trending or moving sideways, and how strong that trend is [27]. The ADX calculation involves several steps and uses the relationship between the price highs and lows over a specified period (usually 14 periods). The steps are as follows:

- Calculate the Directional Movement (+DM and -DM), where +DM (Positive Directional Movement) measures upward movement when the current high minus the previous high is greater than the current low minus the previous low, rather -DM (Negative Directional Movement) measures downward movement when the previous low minus the current low is greater than the current high minus the previous high;
- Calculate the True Range (TR);
- Calculate the Directional Index (DI):

$$+DI = \frac{\sum +DM}{\sum TR} \times 100, \tag{18}$$

$$-DI = \frac{\sum -DM}{\sum TR} \times 100. \tag{19}$$

- Calculate the Directional Movement Index (DX, which measures the absolute difference between +DI and -DI, divided by their sum, and multiplied by 100:

$$DX = \frac{|+DI - -DI|}{+DI + -DI} \times 100. \tag{20}$$

- Calculate the ADX as a smoothed moving average of the DX, typically over a 14-period timeframe:

$$ADX = EMA\ of\ DX\ over\ 14\ period. \tag{21}$$

An example is proposed in Figure 22.

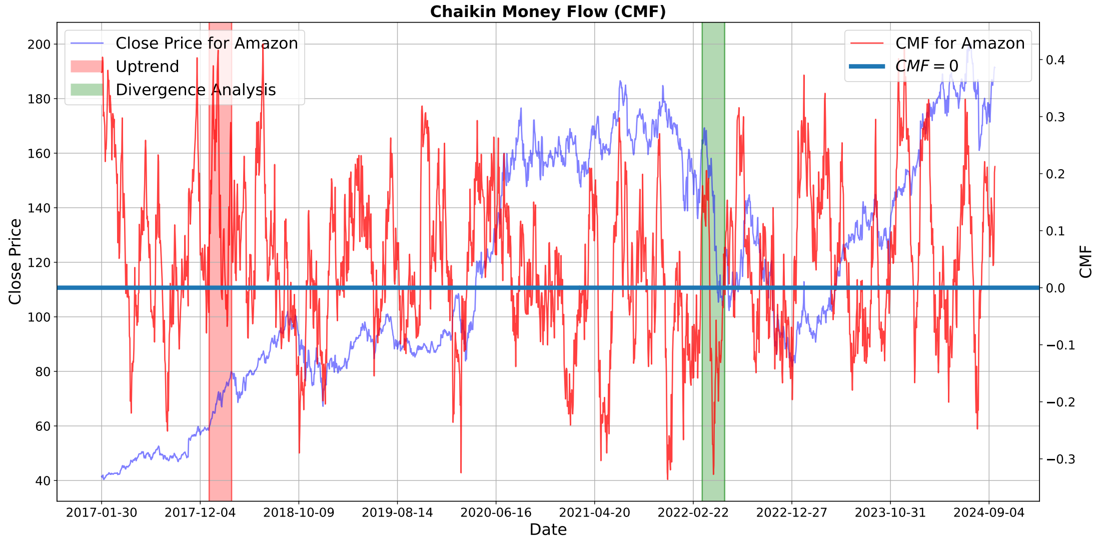

FIGURE 20: *The figure displays a graph with the Close Price (blue line) and the Chaikin Money Flow (red line). The left y-axis shows the scale for the Close Price, while the right y-axis corresponds to CMF values. When CMF is above the blue zero line, it indicates that the asset is being accumulated, suggesting that buyers are stronger and pushing the price higher. Sustained positive CMF values often confirm an uptrend, like the red shaded area. Rather, when CMF is below zero, it signifies that the asset is being distributed, indicating selling pressure. Sustained negative CMF values confirm a downtrend, such as the green area.*

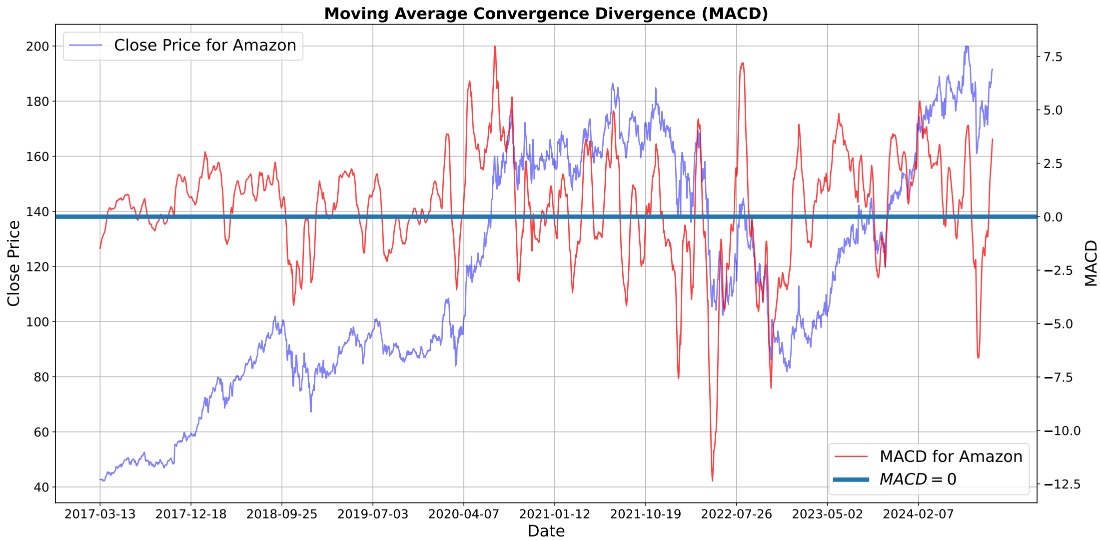

FIGURE 21: *The figure displays a graph with the Close Price (blue line) and the Moving Average Convergence Divergence (red line). The left y-axis shows the scale for the Close Price, while the right y-axis corresponds to MACD values. When the MACD Line crosses above the zero line (the horizontal blue line), it indicates that the short-term momentum is stronger than the long-term momentum, suggesting bullish strength. Rather, when MACD Line crosses below the zero line, it suggests bearish strength, as the short-term momentum is weaker than the long-term momentum.*

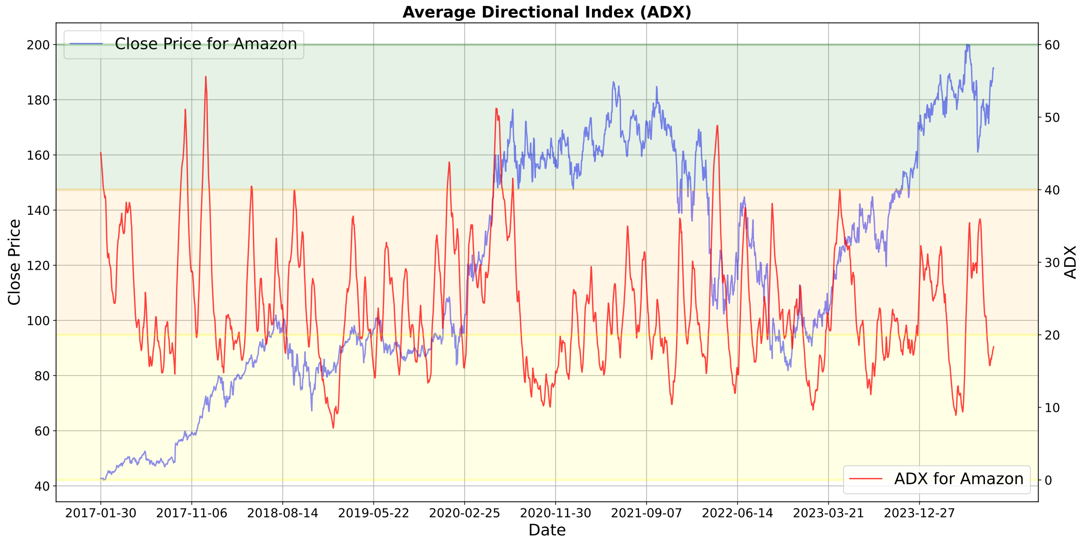

FIGURE 22: *The figure displays a graph with the Close Price (blue line) and the Average Directional Index (red line). The left y-axis shows the scale for the Close Price, while the right y-axis corresponds to ADX values. When ADX is in range 0-20, it indicates a weak trend or sideways market, like highlighted by the yellow area. Traders often avoid trend-based strategies during this period. Range 20-40 indicates a moderate or strengthening trend, as in the orange faded area. This range is often seen at the early stages of a trend. Range 40-60 suggests a strong trend, such as the green area in the plot. Trend-following strategies are usually more effective during this range. Range 60-100 reflects a very strong trend. While a high ADX indicates trend strength, it can also suggest that the trend is overextended and may be nearing exhaustion.*

7. The Fast Simple Moving Average (Fast SMA) is a type of Simple Moving Average (SMA) that uses a shorter time period to calculate the average price of a security. It is called ”fast" because it responds more quickly to recent price changes than an SMA with a longer time period, making it useful for identifying short-term trends [32]. It is calculated by summing up the closing prices of an asset over a specific number of periods and then dividing by that number of periods. The shorter the period used in the calculation, the faster the SMA responds to price changes. This is why a Fast SMA typically uses a period of 5, 10, or 20, as opposed to longer-term SMAs that might use 50, 100, or 200 periods.

$$SMA = \frac{Sum\ of\ closing\ prices\ over\ n\ periods}{n}. \tag{22}$$

In Figure 23 an example of this feature is presented.

8. The Fast Exponential Moving Average (Fast EMA) is a type of Exponential Moving Average (EMA) that uses a shorter time period to calculate the moving average. Like the Fast Simple Moving Average (SMA), it is called "fast" because it responds more quickly to recent price changes compared to a longer-period EMA, making it useful for identifying short-term trends. The formula for calculating an EMA is as follows:

$$\begin{aligned} EMA = {} & Price_{today} \times K \\ & + EMA_{yesterday} \times (1-K), \end{aligned} \tag{23}$$

where K is the smoothing constant, calculated as $K = \frac{2}{n+1}$. An example is proposed in Figure 24.

9. The Commodity Channel Index (CCI) is a momentum-based technical indicator used to measure the deviation of a security's price from its historical average price. Despite its name, the CCI can be applied not only to commodities but also to stocks, indices and other financial instruments [31]. The CCI quantifies the relationship between the current price, a moving average of the price and the mean deviation from that moving average. This measurement helps determine if the asset is trading significantly above or below its average value, signaling potential overbought or oversold conditions. The formula for calculating the CCI is:

$$CCI = \frac{TypicalPrice\ \text{-}\ SMA(TypicalPrice)}{0.015 \times MeanDeviation}, \tag{24}$$

where $TypicalPrice = \frac{High+Low+Close}{3}$ and $MeanDeviation$ is the average of the absolute deviations of the Typical Price from its SMA over the same period.

10. The Momentum Relative Strength Index (RSI) is a widely used momentum oscillator in technical analysis that measures the speed and change of price movements. It is designed to identify overbought or oversold conditions in a market, as well as potential reversal points [30]. In particular, it calculates the ratio of recent upward price movements to recent downward price movements over a given period, typically set at 14 days or periods by default. The formula for RSI is as follows:

$$RSI = 100 - \frac{100}{1+RS}, \tag{25}$$

where $RS = \frac{AveregeGain}{AveregeLoss}$.

An example is provided in Figure 26.

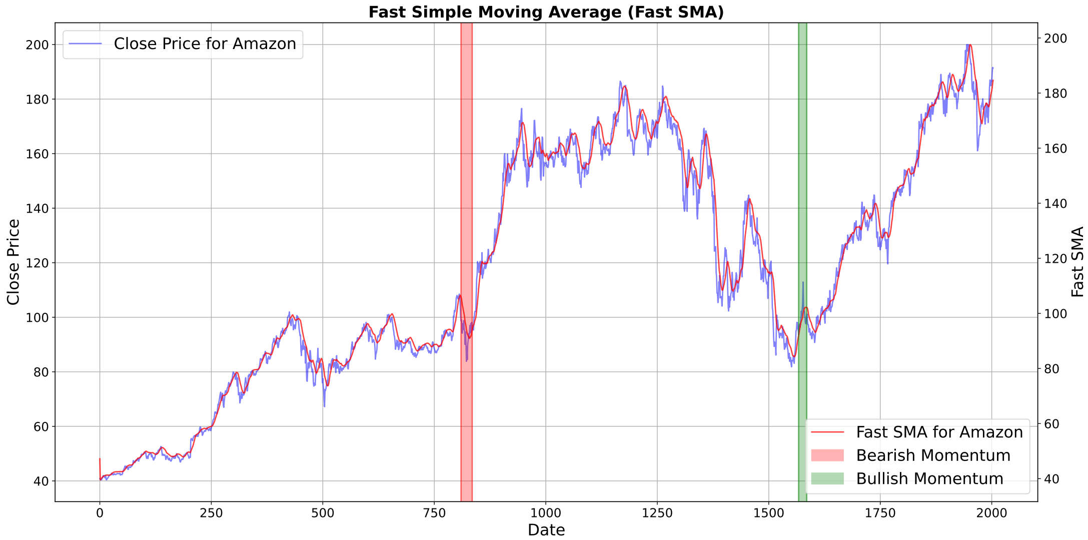

FIGURE 23: The figure displays a graph with the Close Price (blue line) and the Fast Simple Moving Average (red line). The left y-axis shows the scale for the Close Price, while the right y-axis corresponds to Fast SMA values. A Fast SMA can be used independently to identify price crossovers. When the price moves above the Fast SMA, it indicates bullish momentum, like the green area, and when it falls below, it signals bearish momentum, like the red area.

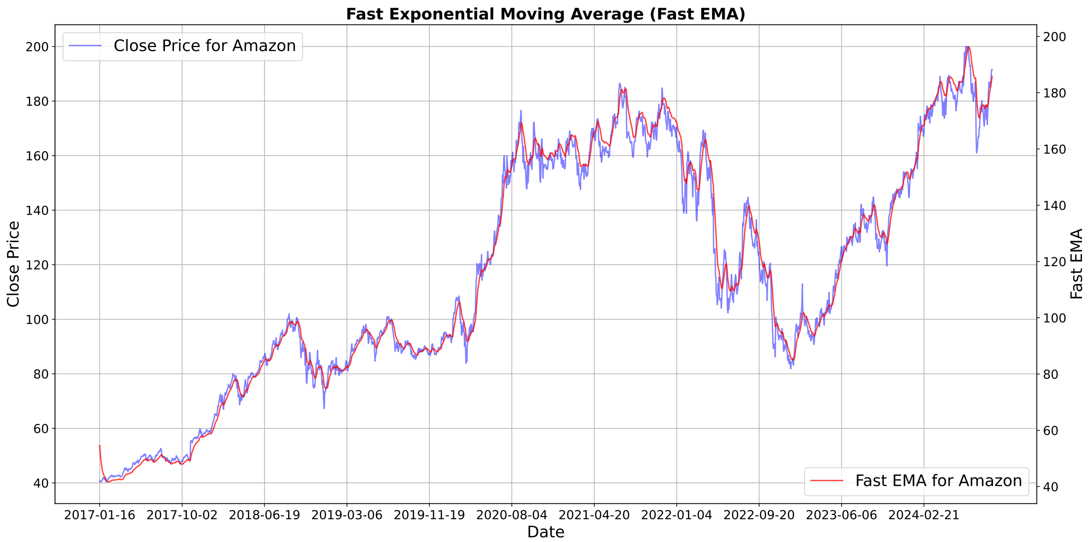

FIGURE 24: The figure displays a graph with the Close Price (blue line) and the Fast Exponential Moving Average (red line). The left y-axis shows the scale for the Close Price, while the right y-axis corresponds to Fast EMA values. As the Fast SMA, traders often use price crossovers to identify buy and sell signals. When the price crosses above the Fast EMA, it indicates bullish momentum, while a move below the Fast EMA signals bearish momentum.

Feature Reduction with VAE-LSTM

Certainly, adding new features to the database enhances the information available for statistical analysis. However, with the current implementation, each individual stock has fifteen features, making it difficult to interpret them collectively. To create more readable features that condense the information from these fifteen inputs, we designed a neural network structure composed of a Variational Autoencoder (VAE) and a Long Short-Term Memory (LSTM) network, a scatch of this architecture is shown in Figure 27. The input to the VAE is a 3D tensor of dimensions $(n_{assets} \times n_{feeature} \times n_{days})$. However, this can't be fed directly into the VAE, so the first layer of the neural network is a Flatten layer. This type of layer reshapes the multi-dimensional tensor into a one-dimensional vector, making it suitable for further processing by fully connected layers.

The resulting output is a 2D array of shape $[(n_{assets} \times n_{feeature}) \times n_{days}]$, which is then passed through a MinMaxScaler layer. This layer normalizes the feature values to a fixed range, in our case $[0, 1]$. Such scaling is crucial since many machine learning algorithms are sensitive to the magnitude of input features. The transformation formula used is:

$$X_{scaled} = \frac{X - X_{min}}{X_{max} - X_{min}}. \tag{26}$$

The normalized input dataset is then fed through two Dense layers, also known as fully connected layers, where each input node is connected to every output node. The first Dense layer has 64 nodes, while the second has 32 nodes, serving as the initial layers of the encoder. These Dense layers gradually reduce the dimensionality of the data, allowing the network to extract complex features that effectively repre-

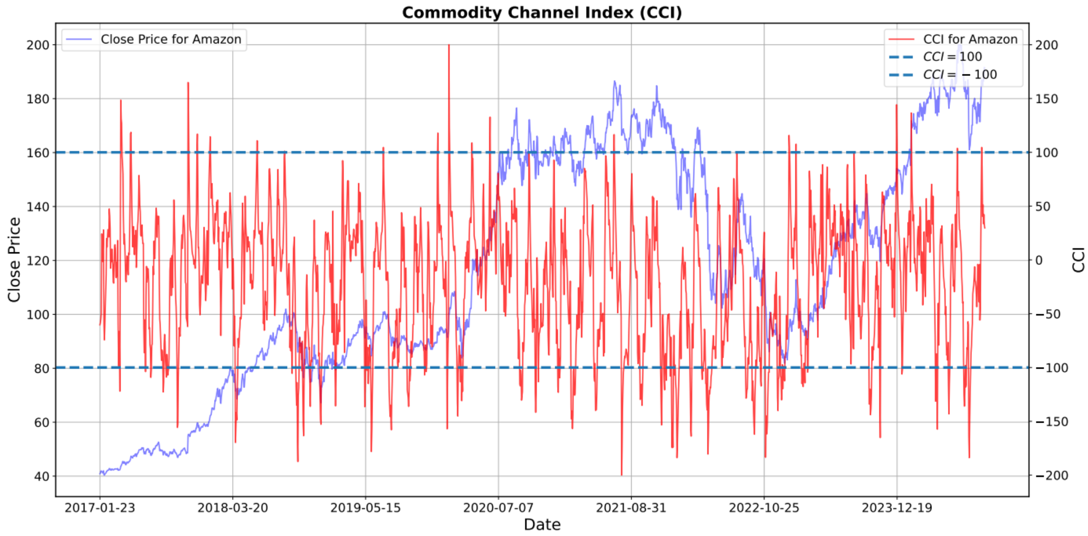

FIGURE 25: The figure shows how the CCI (red line) changes in relation to the Close Price (blue line) for Amazon from 2017 to 2024. The left y-axis represents the scale for the Close Price, while the right y-axis corresponds to CCI values. The graph includes two horizontal lines at $CCI = -100$ and $CCI = 100$, which are typical thresholds used for the CCI indicator. When $CCI > 100$, it signals that the asset is overbought and may experience a price correction or pullback. Conversely, $CCI < -100$ suggests that the asset is oversold and could be set for a rebound or trend reversal.

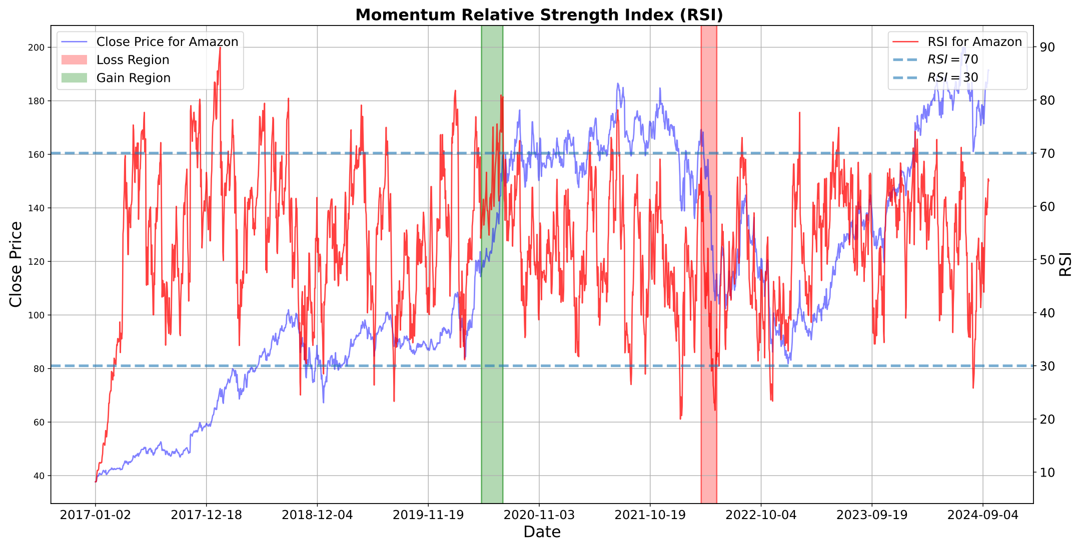

FIGURE 26: The plot displays Amazon's closing price (blue line) and its Relative Strength Index (RSI) (red line) from 2017 to 2024. Green-shaded area highlights a period of price gain, while red-shaded area represents a period of loss. When the RSI is above 70 during a green-shaded region, it signals that the stock is overbought. This could prompt the trader to take profits or consider a short position in anticipation of a potential price pullback. Conversely, the red-shaded area coincides with the RSI dropping below 30, indicating that the stock is oversold. This condition serves as a potential buy signal for the trader, expecting a price rebound.

sent the input.
Next, the compressed data is passed through two separate layers: the Mean Layer and the Log Standard Deviation Layer, each with two nodes. The outputs from these layers are the mean values and log standard deviations that define a multivariate Gaussian distribution, forming the latent space representation of the input.
In Figure 28 a plot of the latent space of the mean array is shown. It is evident from the plot that each year maintains a distinctive pattern, forming clusters of points, which are highlighted with different colors in the figure. These patterns can be explained by the fact that data from adjacent time periods are encoded into similar latent variables due to the comparable behavior of the financial features. This becomes even more apparent when observing the plot with the gradient scale in Figure 29, where the color transitions are smoother and more continuous compared to the previous graph.
After evaluating the mean and log standard deviation arrays, the next step for the neural network is to define each latent variable, Z. As described in Section "Variational Autoencode" each latent variable is derived from a bivariate Gaussian distribution, using statistics computed from the mean and log standard deviation arrays.

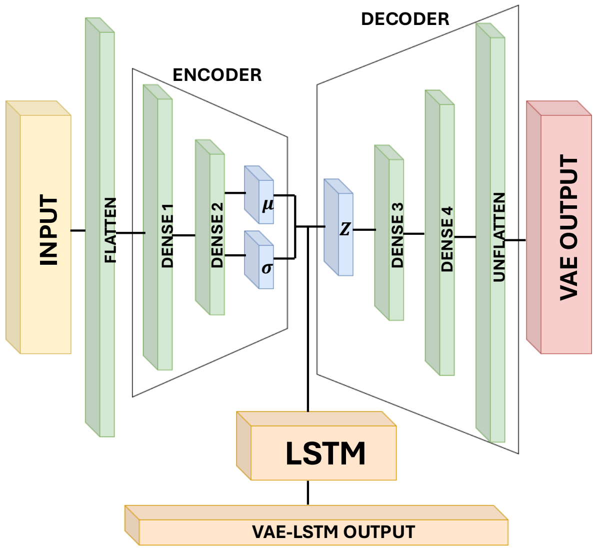

FIGURE 27: The figure shows the architecture of the VAE-LSTM model built for this paper. The yellow layer represents the input, which is first passed through a Flatten layer to adjust its structure, followed by two Dense layers. Afterward, the input dimension is reduced into two layers: μ and σ, representing the mean and log standard deviation of the latent variable Z.This latent variable is then used to reconstruct the original input, represented by the VAE output layer, highlighted in red in the figure. Once the latent variables μ and σ are obtained, they are fed into the LSTM layer, completing the architecture and producing the final VAE-LSTM output.

For instance, consider the first data point in our dataset, hence the first mean pair of the array just mentioned, which is given by $[-1.286438, 0.0488874]$, and the corresponding log standard deviation pair is $[-2.7194648, 0.14427687]$. The resulting multivariate distribution with the variable Z casually picked are illustrated both as a 2D graph in Figure 30 and as a 3D plot in Figure 31.
The final layers of the neural network consist of three Dense layers with 32, 64 and $(n_{assets} \times n_{feeature})$ nodes. These layers are designed to process the output of the VAE, which is only used to minimize the network's Loss Function and determine the optimal weights for each layer, resulting in the lowest possible loss. As discussed in Section "Variational Autoencode" the Loss Function of a VAE is defined as the sum of the Reconstruction Loss and the Kullback-Leibler loss, as shown in the formula below:

$$L = L_{Reconstruction} + L_{Kullback-Leibler}. \tag{27}$$

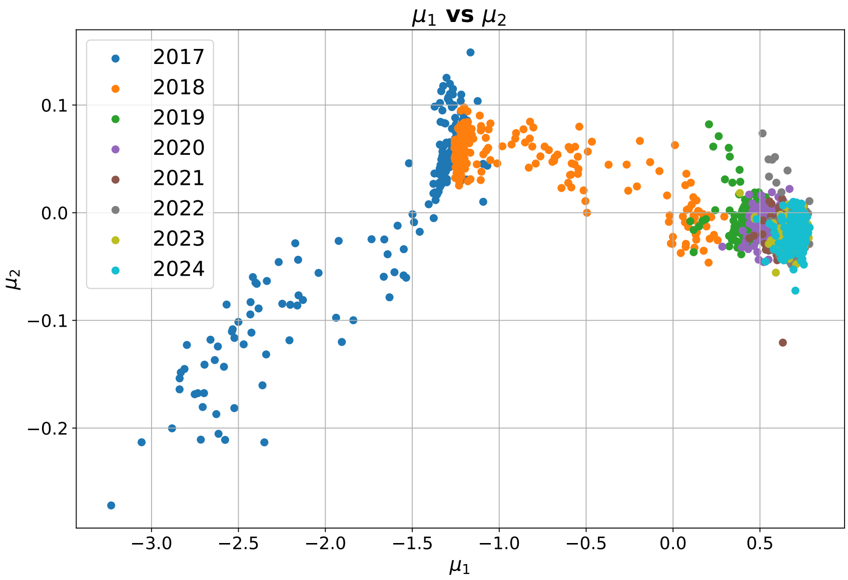

FIGURE 28: This plot presents a scatter plot of the mean array produced by the VAE encoder, μ_1, in the x axis, versus μ_2, in the y axis. Different colors are used to represent distinct time periods, as shown in the legend of the graph, making it easier to identify patterns across the graph.

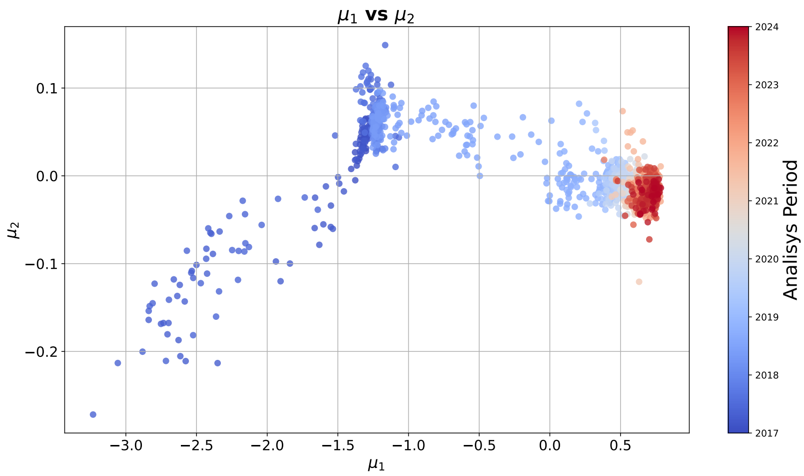

FIGURE 29: This plot presents a scatter plot of the mean array produced by the VAE encoder, μ_1, in the x axis, versus μ_2, in the y axis. Here the color scale is used to identify the different years, the dark blue refers to data from 2017, rather dark red points refers to 2024 data.

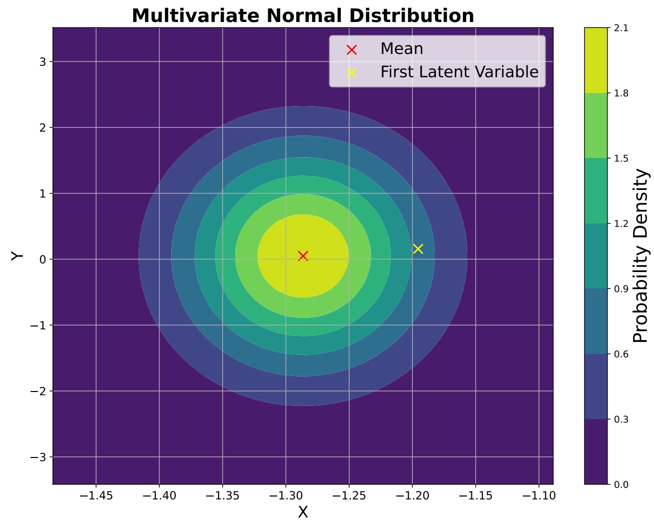

FIGURE 30: Plot of the multivariate distribution generated using the statistics encoded for the first day, with parameters $\mu_1 = -1.3063028$, $\mu_2 = 0.06409347$, $\sigma_1 = -2.8236599$, and $\sigma_2 = 0.1387112$. The surface uses a color gradient ranging from dark blue (indicating points with lower probability density) to yellow (indicating points with the highest probability density). The red cross marks the mean of the distribution, while the yellow cross highlights the latent variable extracted. As observed, this point does not necessarily coincide with the most probable combination.

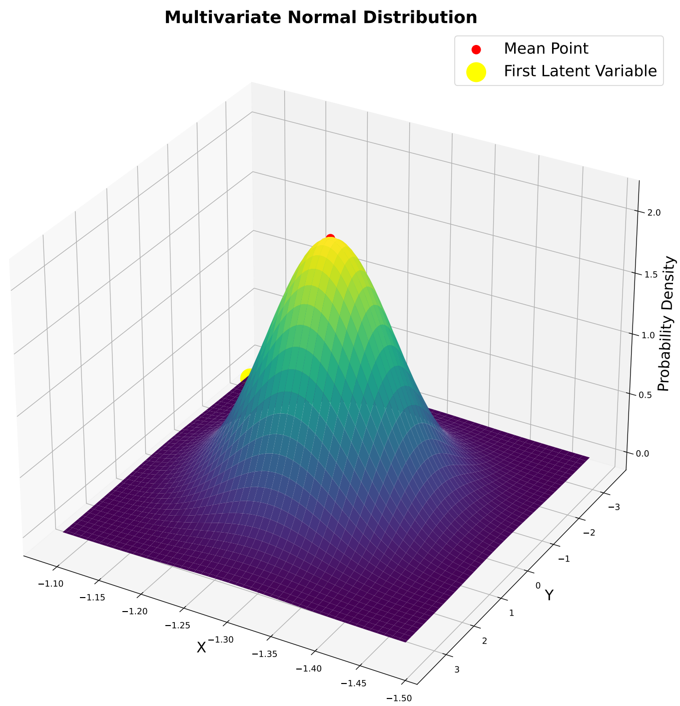

FIGURE 31: Tridimensional plot of the multivariate distribution generated using the statistics encoded for the first day, with parameters $\mu_1 = -1.3063028$, $\mu_2 = 0.06409347$, $\sigma_1 = -2.8236599$, and $\sigma_2 = 0.1387112$. The surface uses a color gradient ranging from dark blue (indicating points with lower probability density) to yellow (indicating points with the highest probability density). The red cross marks the mean of the distribution, while the yellow cross highlights the latent variable extracted.

To fit VAE's weights, we used the following parameter: $n_{epochs} = 100$ and $n_{batch} = 32$, where an epoch is one complete pass through the entire training dataset. During training, the model iterates over the data multiple times to learn and improve its weights and biases. Each time the model has seen every sample in the training set once, it com-

pletes one epoch. Rather, batch size refers to the number of samples the model processes before updating its internal parameters. Instead of updating weights after each individual sample (which can be slow) or after all samples (which might not fit into memory), batch size controls how many samples are processed before a weight update is applied.
Another parameter set was the validation split, which is a parameter used when training machine learning models to specify the proportion of the training data that should be set aside for validation, in our case $split = 0.2$, indicating that the 20% of the dataset is used as validation set.
Figure 32 shows the trend of the VAE Loss with respect to the number of epochs. It can be observed that the loss asymptotically approaches its minimum value over time.
Specifically, due to the relatively small size of the dataset, the total VAE loss (green line), along with its components, the Reconstruction Loss (blue line) and KL Loss (yellow line), nearly reaches its minimum after around 20 epochs.
Once the latent space variables (μ, σ) are extracted using the VAE, they are passed through the LSTM layer to capture temporal patterns in the latent space, which can then guide the model's weight adjustments. The goal of the LSTM in this context is not to generate new data, but rather to forecast existing data points.
Specifically, we aim to use a sequence of data points, x_i with i in $0, ..., n$ to predict x_{n+1}. As an initial step, it's crucial to determine the length of the sequence n that will be fed into the LSTM. Considering the temporal nature of our dataset, a reasonable choice is to use five consecutive data points to predict the sixth, simulating a weekly data pattern for forecasting the subsequent value.
Going deeper in the problem, let's define z_i the array consisting of μ_1^i, μ_2^i, σ_1^i and σ_2^i. Since we are working with windows of dimension $n = 5$, it's necessary to reshape the dataset to be fed to the LSTM as follows:

$$\begin{bmatrix} [z_0, z_1, z_2, z_3, z_4] \\ [z_1, z_2, z_3, z_4, z_5] \\ \dots \\ [z_i, z_{i+1}, z_{i+2}, z_{i+3}, z_{i+4}] \\ \dots \\ [z_{N-6}, z_{N-5}, z_{N-4}, z_{N-3}, z_{N-2}] \\ [z_{N-5}, z_{N-4}, z_{N-3}, z_{N-2}, z_{N-1}] \end{bmatrix}.$$

So the goal of the LSTM is to predict z_5 based on the patterns and trends observed in the sequence $[z_0, z_1, z_2, z_3, z_4]$. More generally, it aims to forecast z_i using the preceding array $[z_{i-5}, z_{i-4}, z_1, z_{i-3}, z_{i-2}, z_{i-1}]$. Figure 33 provides a visual representation of this process to aid comprehension.
As in previous cases, the weights and biases of the neural network are optimized by minimizing the loss function. To fit LSTM's weights and biases, we used the following parameter: $n_{epochs} = 50$, $n_{batch} = 32$ and $split = 0.2$. The trend of the loss function respect the number of epochs is shown in Figure 34.
The output generated by the LSTM has a shape of $[(n_{days} - 4) \times n_{feature}]$, as determined by the network's architecture. Figures 35 and 36 illustrate how two components of the outputs, $\hat{\mu}_1$ and $\hat{\mu}_2$ (yellow lines), compare to the original μ_1 and μ_2 (red lines). As shown, the yellow lines have a smoother shape compared to the red lines, which exhibit more abrupt variations. This behavior highlights the tendency of the LSTM to smooth out fluctuations while still capturing the overall trend, which is particularly evident in the second plot, where the y-axis scale enhances this visualization. A result on how the output would like is presented in Table 8.
As the final step of this section, Table 9 presents the dataset that will serve as input for the neural networks in the upcoming sections. This dataset is constructed by concatenating the close returns with the VAE-LSTM outputs. As shown, each date corresponds to a single set of outputs, meaning that each implemented feature has the same value across all assets and only varies from one day to the next.

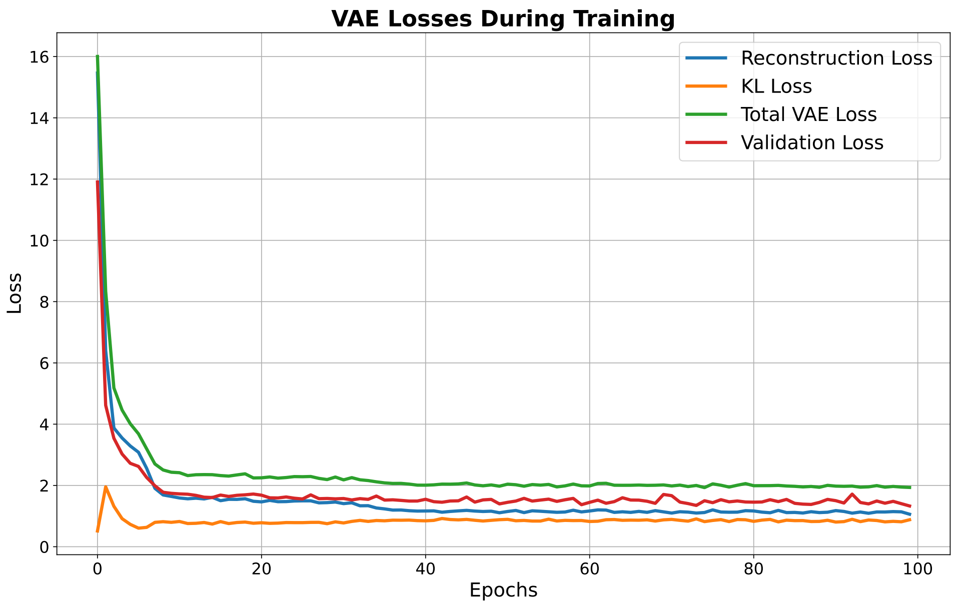

FIGURE 32: The plot illustrates the trends of various losses over the number of epochs. The blue line represents the Reconstruction Loss, while the yellow line corresponds to the KL Loss. Their sum forms the total VAE Loss, shown by the green line. The red line depicts the Validation Loss, which is evaluated on the validation dataset. All losses reach their asymptotes at around 20 epochs and show no further improvement afterward, indicating that the VAE model can be effectively trained without requiring many epochs.

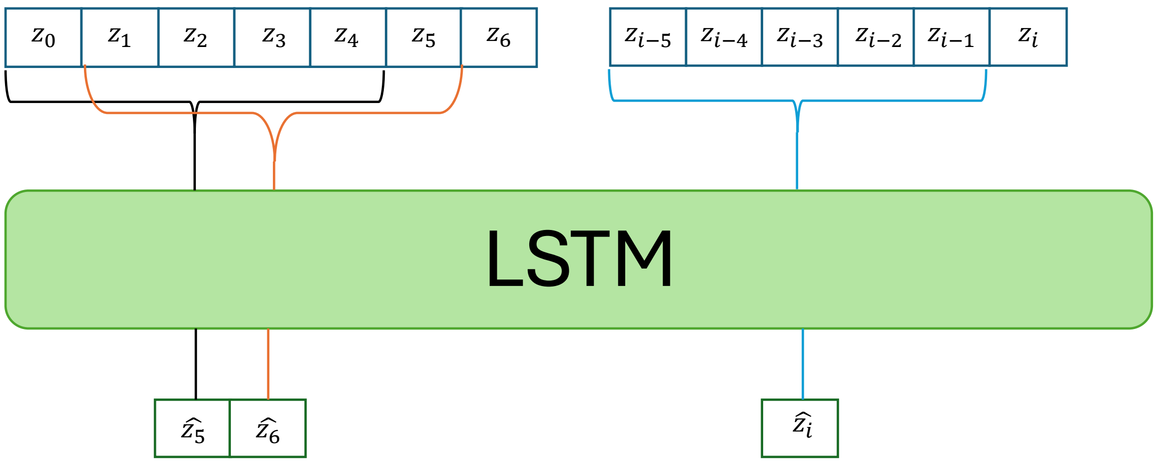

FIGURE 33: The plot illustrates an example structure of the LSTM model. The upper section represents the input, which consists of the individual components of each z_i. At each step, a sequence of five variables is fed into the LSTM to generate the predicted output $\hat{z}$ output.

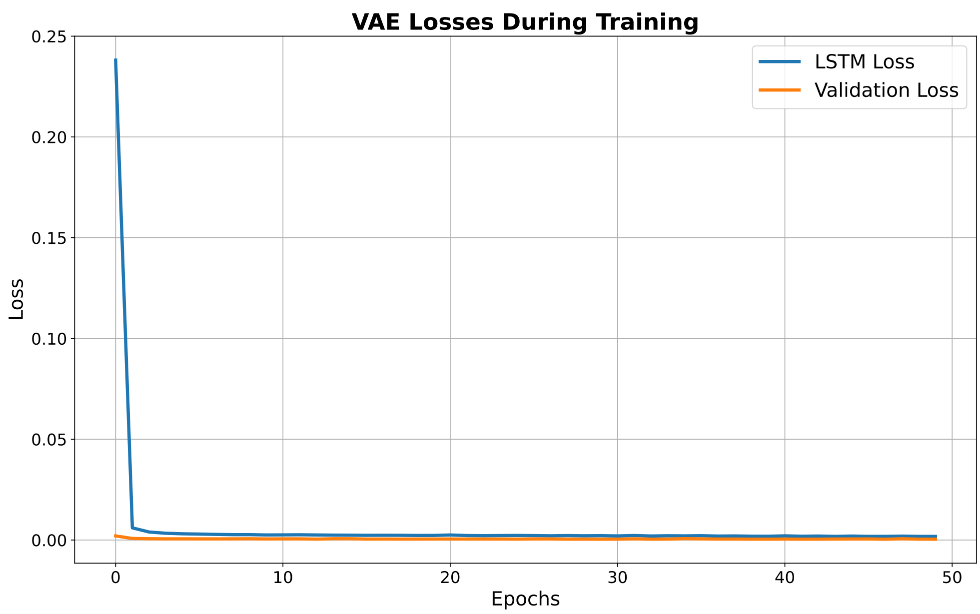

FIGURE 34: The plot illustrates the trends losses over number of epochs. The blue line represent the reconstruction loss, while the orange one shows the validation loss.

Neural Network for Weights Selection

Since the RL algorithm simulates the portfolio multiple times to identify the optimal weight configuration for each day, this section discusses a single simulation to clarify how weights are selected during each iteration and outline the assumptions made.
The core of the simulation involves determining the asset weights for each individual day under analysis. This is accomplished using a Dense layer in the neural network. The Dense layer takes as input the current day's returns along with the four statistical features introduced earlier and outputs a set of n_{asset} weights, one for each asset. Incorporating these additional statistical features as inputs provides the neural network with extra information, which helps reduce the number of simulations required to find the best portfolio configuration.
To ensure that the sum of the weights for each day equals 1 (or 100% in percentage terms), we employ the Softmax activation function for the Dense layer. The Softmax function transforms a vector of raw values into a probability distribution, ensuring that the sum of all elements is 1 and that all values are non-negative. This property is crucial for portfolio allocation, where weights must represent proportions of total investment. In mathematical terms, the Softmax function converts a vector of raw scores $\vec{z} = [z_1, z_2, ..., z_n]$ into a vector of probabilities $\vec{p} = [p_1, p_2, ..., p_n]$, where:

$$p_i = \frac{e^{z_i}}{\sum_{j=1}^{n} e^{z_j}}. \tag{28}$$

Date	$\hat{\mu}_1$	$\hat{\mu}_2$	$\hat{\sigma}_1$	$\hat{\sigma}_2$
06/01/2017	$\hat{\mu}_1^0$	$\hat{\mu}_2^0$	$\hat{\sigma}_1^0$	$\hat{\sigma}_2^0$
07/01/2017	$\hat{\mu}_1^1$	$\hat{\mu}_2^1$	$\hat{\sigma}_1^1$	$\hat{\sigma}_2^1$
...	...	...	...	...
02/01/2021	$\hat{\mu}_1^i$	$\hat{\mu}_2^i$	$\hat{\sigma}_1^i$	$\hat{\sigma}_2^i$
...	...	...	...	...
02/01/2023	$\hat{\mu}_1^j$	$\hat{\mu}_2^j$	$\hat{\sigma}_1^j$	$\hat{\sigma}_2^j$
...	...	...	...	...
23/09/2024	$\hat{\mu}_1^N$	$\hat{\mu}_2^N$	$\hat{\sigma}_1^N$	$\hat{\sigma}_2^N$

TABLE 8: Table with the resulting output of the VAE-LSTM architecture.

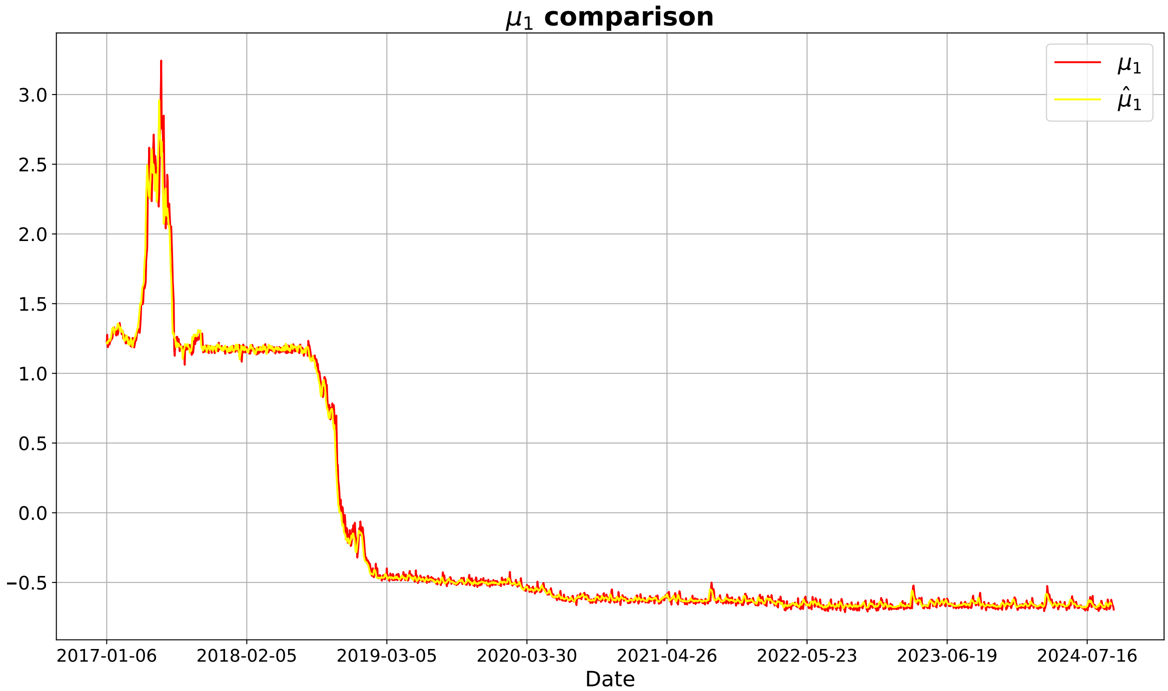

FIGURE 35: Comparison of the latent variable μ_1 (red line) and the output of the LSTM, $\hat{\mu}_1$ (yellow line).

This ensure that $\sum_{i=1}^{n} p_i = 1$.
Additionally, as this work does not consider short selling or borrowing, all weights must be positive. Softmax naturally enforces non-negative weights, making it well-suited for traditional portfolios that allow only positive allocations.
Another constraint imposed is to prevent any single weight from exceeding a maximum exposure limit. Setting maximum exposure limits helps maintain portfolio diversification and prevents over-concentration in a few assets, thereby reducing overall risk. To enforce this, a customized version of Softmax, known as Softmax with Max Clipping, is used. This method combines the Softmax transformation with a mechanism that caps the maximum value of each weight or probability. Specifically, after applying Softmax to obtain the probabilities p_i, it's ensured that no weight exceeds a pre-defined threshold, such as $max\_clipping = 20\% = 0.2$. Mathematically, max clipping modifies each p_i as follows:

$$p_i = \min(p_i, max\_clipping). \tag{29}$$

After clipping, the sum of the probabilities might not equal 1. Therefore, an additional normalization step is applied:

$$p_i = \frac{p_i}{\sum_{j=1}^{n} p_j}. \tag{30}$$

This ensures that the final set of weights meets the desired sum constraint while respecting the maximum exposure limit.
Figure 37 illustrates the neural network architecture as described above.

Custom Loss Functions Design

In RL, the loss function is a crucial component that guides the learning process of an agent by providing feedback on its performance. Unlike in supervised learning, where loss functions minimize the difference between predicted outputs and true labels, RL loss functions are designed to maximize cumulative rewards over time. This difference makes the design of RL loss functions more complex and highly dependent on the specific algorithm being used, such

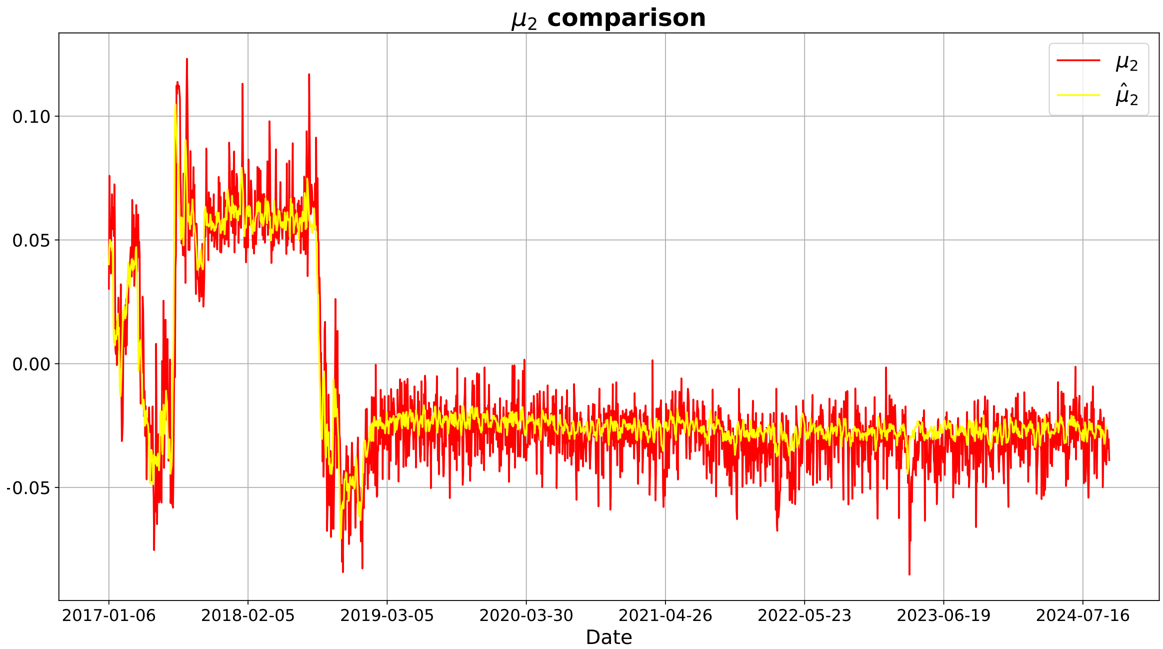

FIGURE 36: *Comparison of the latent variable μ_2 (red line) and the output of the LSTM, $\hat{\mu}_2$ (yellow line).*

Date	Ticker	Return Close Price	$\hat{\mu}_1$	$\hat{\mu}_2$	$\hat{\sigma}_1$	$\hat{\sigma}_2$
06/01/2017	Stock1	r^0_{Stock1}	$\hat{\mu}^0_1$	$\hat{\mu}^0_2$	$\hat{\sigma}^0_1$	$\hat{\sigma}^0_2$
06/01/2017	Stock2	r^0_{Stock2}	$\hat{\mu}^0_1$	$\hat{\mu}^0_2$	$\hat{\sigma}^0_1$	$\hat{\sigma}^0_2$
...	...	...	...	...	...	...
06/01/2017	Stock15	$r^0_{Stock15}$	$\hat{\mu}^0_1$	$\hat{\mu}^0_2$	$\hat{\sigma}^0_1$	$\hat{\sigma}^0_2$
07/01/2017	Stock1	r^1_{Stock1}	$\hat{\mu}^1_1$	$\hat{\mu}^1_2$	$\hat{\sigma}^1_1$	$\hat{\sigma}^1_2$
...	...	...	...	...	...	...
02/01/2021	Stock1	r^i_{Stock1}	$\hat{\mu}^i_1$	$\hat{\mu}^i_2$	$\hat{\sigma}^i_1$	$\hat{\sigma}^i_2$
...	...	...	...	...	...	...
02/01/2021	Stock15	$r^i_{Stock15}$	$\hat{\mu}^i_1$	$\hat{\mu}^i_2$	$\hat{\sigma}^i_1$	$\hat{\sigma}^i_2$
...	...	...	...	...	...	...
23/09/2024	Stock1	r^N_{Stock1}	$\hat{\mu}^N_1$	$\hat{\mu}^N_2$	$\hat{\sigma}^N_1$	$\hat{\sigma}^N_2$
...	...	...	...	...	...	...
23/09/2024	Stock15	$r^N_{Stock15}$	$\hat{\mu}^N_1$	$\hat{\mu}^N_2$	$\hat{\sigma}^N_1$	$\hat{\sigma}^N_2$

TABLE 9: *Table with the input of the Reinforcement Learning.*

as policy gradient methods, Q-learning or actor-critic models.
For our portfolio allocation with A2C algorithm, the goal of the loss function is to minimize the negative factors that detract from portfolio performance, thereby indirectly maximizing total returns over the observation period. Given this context, we have developed a custom loss function tailored to address the unique requirements of portfolio management. The function combines several key components to reduce portfolio risk, balance diversification and minimize trading costs and is showed below:

$$\begin{aligned}\mathcal{L}_{TOT} = -\gamma(&\mathcal{L}_{expected\_returns} - \mathcal{L}_{risk\_penalty}\\ &- \mathcal{L}_{portfolio\_volatility\_penalty} + \mathcal{L}_{IDR}\\ &- \mathcal{L}_{entropy} - \mathcal{L}_{turnover\_penalty}\\ &- \mathcal{L}_{market\_impact\_penalty}).\end{aligned}$$

Each of these components will be discussed in detail below. To highlight how each term of the loss function varies across different portfolios, we present examples based on two distinct portfolios: the first with uniform weights and the second with divergent weights. These examples, illustrated in Figure 38, provide a clear comparison of how the loss function's components behave under different weight allocations.

- $\mathcal{L}_{expected\_returns}$ computes the expected return of the portfolio, which is the sum of the weights allocated to each asset multiplied by their respective returns. Maximizing this value corresponds to increasing the overall return of the portfolio. In formula:

$$\mathcal{L}_{expected\_returns} = \sum_{i=1}^{N} w_i \cdot r_i. \quad (31)$$

  To compare the two portfolios illustrated in Figure 38, let's consider a set of returns as shown in Figure 39. Using these returns, we can calculate two loss values based on the scalar multiplication with the respective weights: $\mathcal{L}^1_{expected\_returns} = 0.0107$ for the uniform portfolio and $\mathcal{L}^2_{expected\_returns} = 0.0092$ for the divergent portfolio. In this example, the uniform portfolio is a better choice since it yields a higher expected return.

- $\mathcal{L}_{risk\_penalty}$ represents the portfolio risk and is calculated as the weighted sum of the standard deviations of each asset. This penalizes portfolios with higher volatility, encouraging a more stable return profile. This term in the final loss function controls the degree of risk aversion. In formulas:

$$\mathcal{L}_{risk\_penalty} = \lambda_{risk\_rate} \sum_{i=1}^{N} w_i \cdot \sigma_i, \quad (32)$$

  where σ_i is the array of standard deviation of each asset, an example is shown in Figure 40 and λ_{risk_rate} is a constant used to balance the weight of this component in the overall loss. Using the portfolios in Figure 38, the two losses results $\mathcal{L}^1_{risk\_penalty} = 0.01827$ and $\mathcal{L}^1_{risk\_penalty} = 0.01861$, hence the first one is less risky than the second.

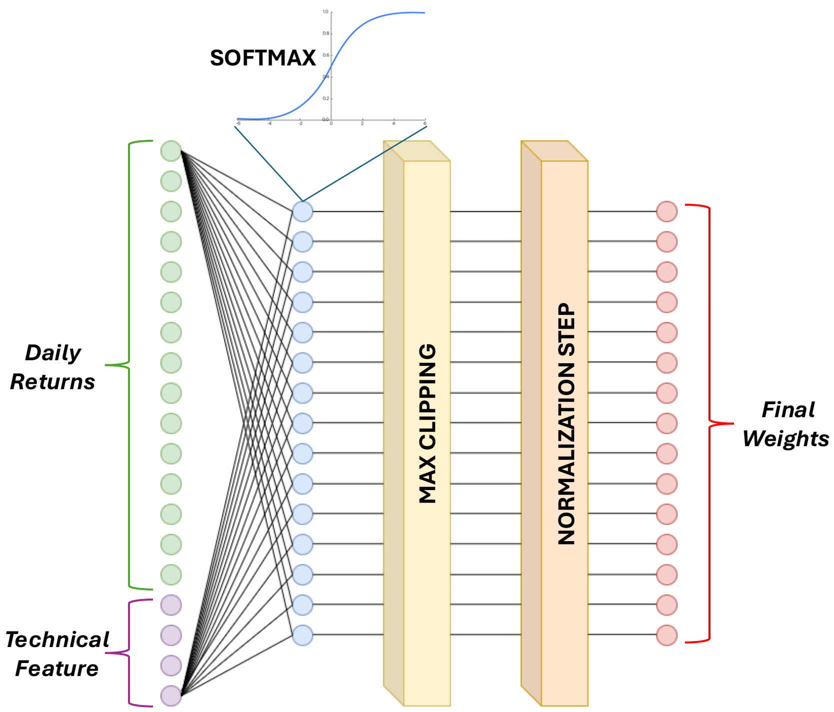

FIGURE 37: The figure shows a schematic of the neural network used to predict the weights for each day. The network takes as inputs the daily returns r_i (represented by green nodes) along with technical features (μ_1,μ_2,σ_1 and σ_2, represented by violet nodes). These inputs are fed into a Dense layer, which is a fully connected layer (only illustrated for the first and last nodes of the input). The activation function used for this layer is Softmax, as depicted in the function plot over the network. The Softmax function converts the raw input values into a probability distribution, ensuring that all elements are non-negative and their sum equals 1. The output of the Dense layer is then passed to a Max Clipping layer, which imposes a post-processing step to ensure no individual weight or probability exceeds a specified threshold. Finally, the output undergoes a normalization step, ensuring that the sum of the final weights is 1.

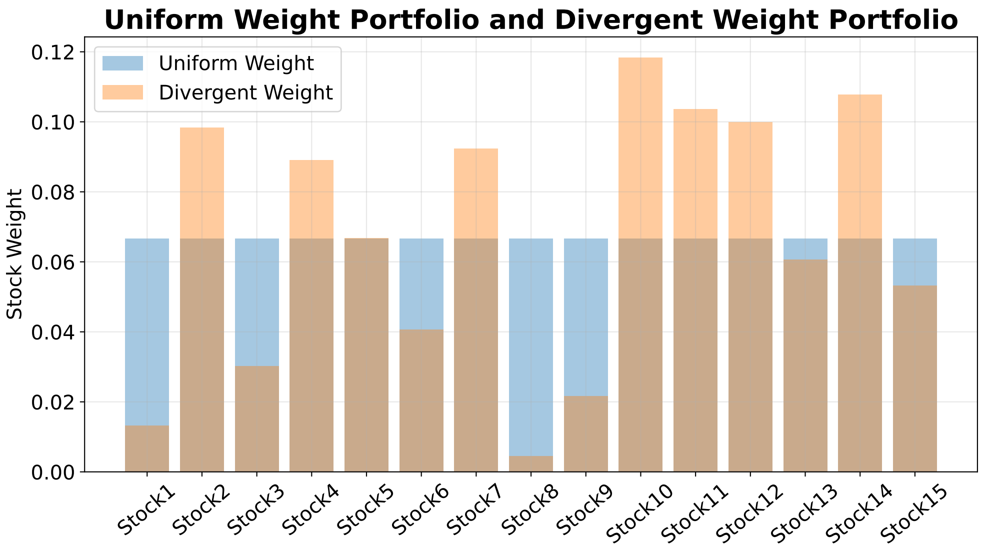

FIGURE 38: Example of two kind of portfolios: the blue one is for an equal weight portfolio, rather the orange one represents a divergent weight portfolio.

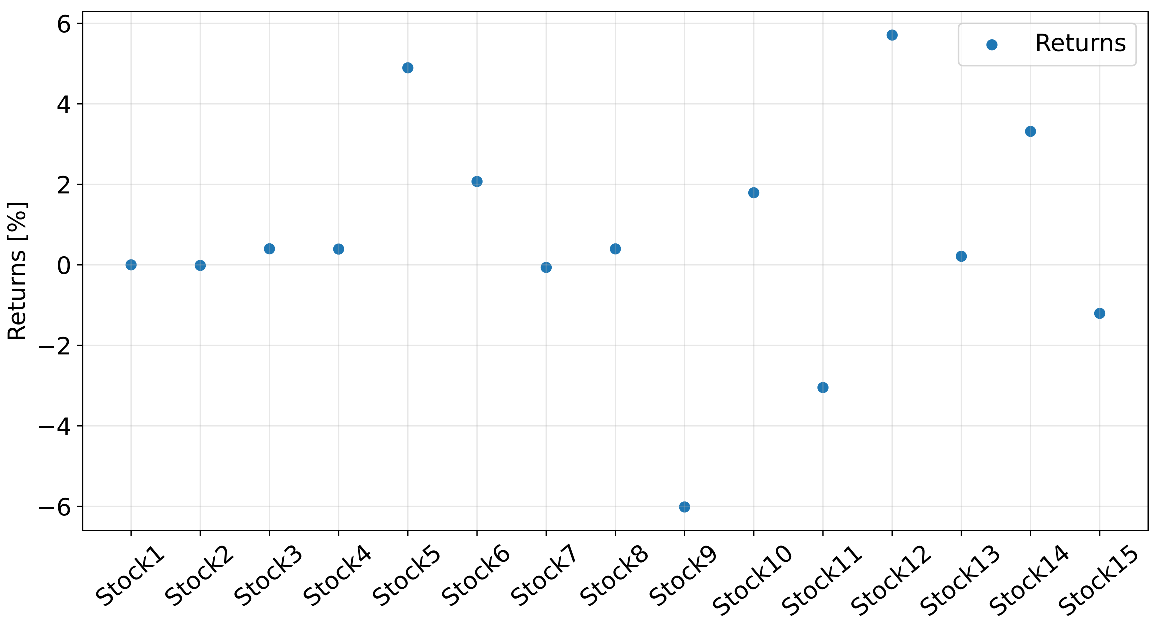

FIGURE 39: Example of percentage returns for each asset of the portfolio.

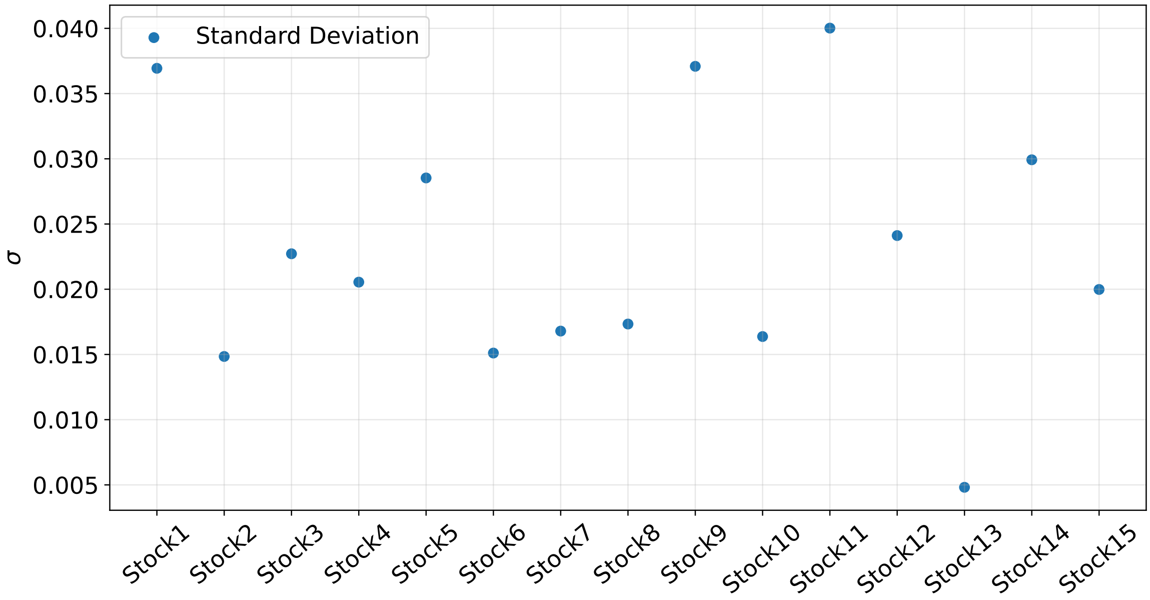

FIGURE 40: Example of Standard Deviation of each asset over a period of 250 days. Greater σ stands for more variable stock, in this case Stock11 has the greatest σ, meaning that its price fluctuates significantly over time.

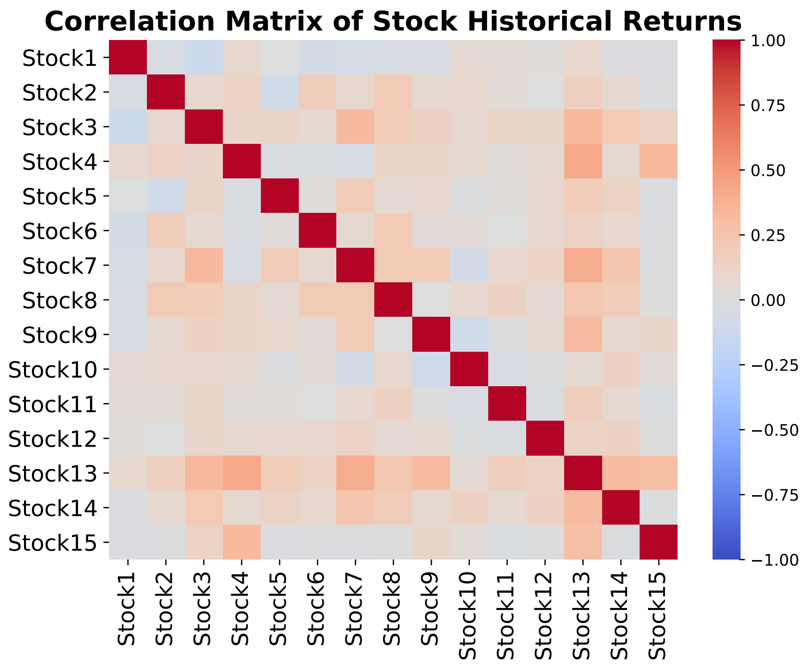

FIGURE 41: Heatmap of the correlation between historical series of returns of close prices. As can be seen the correlation between each stock is nearly zero, conferming that our portfolio is diversificated in term of asset.

- $\mathcal{L}_{portfolio\_volatility\_penalty}$ is necessary because, in addition to penalizing individual asset risk, we want to penalize the overall portfolio volatility. This can be done using the covariance matrix of returns. For a portfolio with weights w_i and asset returns covariance matrix Σ, the portfolio volatility loss can be calculated as:

$$\mathcal{L}_{portfolio\_volatility\_penalty} = \lambda_{portfolio\_volatility} \sum_{i=1}^{N} \sqrt{w_i^T \Sigma w_i}, \tag{33}$$

where $\lambda_{portfolio\_volatility}$ is constant used to balance the weight of this component in the overall loss. Given the covariance matrix Σ in Figure 41, which shows how the historical series of returns are correlated, the resulting losses for the two portfolios in Figure 38 are: $\mathcal{L}^1_{portfolio\_volatility\_penalty} = 1.254 \cdot 10^{-4}$ and $\mathcal{L}^2_{portfolio\_volatility\_penalty} = 1.526 \cdot 10^{-4}$, confirming that the second portfolio is riskier than the first one.

- $\mathcal{L}_{IDR}$ which stay for Inverse Diversification Ratio and is a measure used in portfolio optimization to evalu-

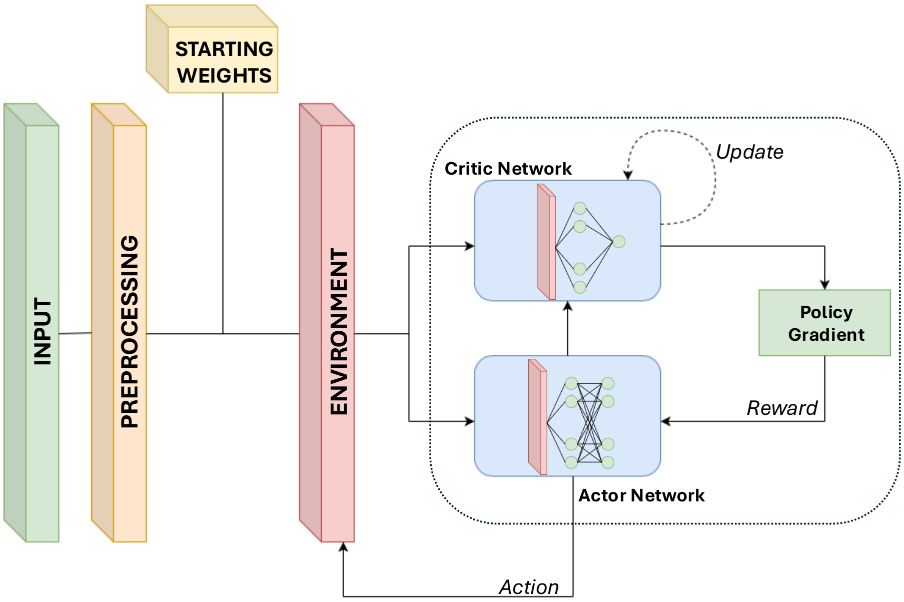

FIGURE 42: The figure illustrates the architecture used in the PPO algorithm. The input (shown as the green layer) after undergoing a preprocessing stage (orange layer), combined with the initial weights of the network represents the environment (red layer). This input is simultaneously processed by both the actor and critic networks. The critic updates its value estimate and provides feedback to the actor, guiding it on which actions to prioritize. Finally, the actor sends the selected action to the environment.

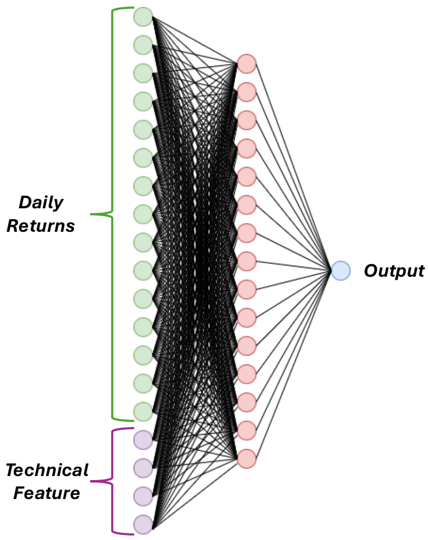

FIGURE 43: In the figure, the Critic Network used by the A2C algorithm is illustrated. This neural network evaluates the value of a given state, which is represented by the output, highlighted as the blue node.

ate and control the level of diversification within a portfolio. It aims to discourage overly concentrated portfolios by penalizing the allocation of too much weight to a few assets. In reinforcement learning or optimization contexts, this penalty can be used to ensure that the resulting portfolio is well-diversified. In formula it can be written as:

$$\mathcal{L}_{IDR} = \lambda_{IDR} \sum_{i=1}^{N} w_i^2, \tag{34}$$

where λ_{IDR} is constant used to balance the weight of this component in the overall loss. For the example portfolio of Figure 38, $\mathcal{L}^1_{IDR} = 0.0666$ and $\mathcal{L}^2_{IDR} = 0.0840$, since the first one has a lower value it suggests a more diversified portfolio, where the risk is more evenly spread.

- $\mathcal{L}_{entropy}$ is to encourage a more diversified and balanced allocation of assets. It is derived from the concept of entropy in information theory, which quantifies the level of uncertainty or randomness in a distribution. In the context of portfolio allocation, a higher entropy value indicates a more evenly distributed allocation, promoting diversification and reducing concentration risk. The formula is:

$$\mathcal{L}_{entropy} = -\lambda_{entropy} \sum_{i=1}^{N} w_i \log(w_i), \quad (35)$$

where $\lambda_{entropy}$ is constant used to balance the weight of this component in the overall loss. For the example portfolios in Figure 38, $\mathcal{L}^1_{entropy} = 2.7081$ and $\mathcal{L}^2_{entropy} = 2.5663$, hence the first portfolio would be preferred in term of entropy loss since it's more balance respect the second one.

- $\mathcal{L}_{turnover\_penalty}$ is a mechanism used in portfolio optimization and reinforcement learning to limit the amount of trading activity in a portfolio. It aims to minimize transaction costs, reduce excessive rebalancing and promote stability in portfolio weights over time. This is particularly useful when transaction costs or slippage can significantly impact portfolio performance. The relative formula is:

$$\mathcal{L}_{turnover\_penalty} = \lambda_{turnover} \sum_{i=1}^{N} |w_i - w_{i-1}| \times fee\_rate, \quad (36)$$

  where $\lambda_{turnover}$ is constant used to balance the weight of this component in the overall loss.

- $\mathcal{L}_{market\_impact\_penalty}$ which is a concept used in trading and portfolio management to account for the costs associated with executing large orders in the market. When an investor or trader executes a sizable trade, it can affect the asset's price, causing it to move unfavorably. This adverse price movement is known as market impact. A common way to calculate it is to assume that the market impact cost is proportional to the square of the trading volume. In formula:

$$\mathcal{L}_{market\_impact\_penalty} = \lambda_{market\_impact} \sum_{i=1}^{N} (w_i - w_{i-1})^2, \quad (37)$$

  where λ_{market_impact} is constant used to balance the weight of this component in the overall loss.

The term γ in the $\mathcal{L}_{TOT}$ formula is known as advantage and in this work is determinated each i-th day as:

$$\gamma_i = cumulative\_return_i + critic\_output_{i+1} - critic\_output_i. \quad (38)$$

In the A2C algorithm the other Loss function, which is used to optimize the Critic Network, is the so called Critic Loss, which is defined as:

$$\mathcal{L} = \sum_{i=1}^{N} (\hat{x}_i - \hat{x}_{i+1})^2. \quad (39)$$

Since DDPG is an actor-critic algorithm, it involves two separate loss functions, similar to A2C: one for the actor and one for the critic. The structure of each loss function follows a format comparable to that used in the A2C algorithm.
Rather, the actor loss function in Proximal Policy Optimization (PPO) is designed to address some of the instability issues present in traditional policy gradient methods. Typically, policy gradient algorithms update the policy by directly using the gradients of the expected reward. However, these updates can sometimes cause the policy to change drastically, leading to instability. This is especially problematic in environments where even small adjustments to the policy can result in significant changes in performance. To mitigate this, PPO introduces a "clipped objective function," which stabilizes the policy updates.
In PPO, the action probability under the previous policy, denoted as *old_prob*, corresponds to the total loss evaluated for the previous set of weights, $\mathcal{L}^{i-1}_{TOT}$. The action probability under the new policy, *new_prob*, is associated with the total loss evaluated for the new set of weights, $\mathcal{L}^{i}_{TOT}$. Additionally, the advantage of an action, *advantage*, and a hyperparameter known as the clip ratio, are used. The clip ratio sets a threshold for how much deviation from the old policy is permitted. The probability ratio is calculated as follows:

$$ratio = \frac{new\_prob}{old\_prob + 1e^{-8}}. \quad (40)$$

This ratio measures the change in how the new policy assigns probabilities to actions compared to the old policy. When *ratio* is close to 1, it indicates that the policies are similar. A *ratio* greater than 1 means the new policy assigns a higher probability to the action than the old one, whereas a *ratio* less than 1 indicates a lower probability.
The objective of PPO is to maximize a surrogate objective, defined as:

$$\mathcal{L}_{PPO} = \mathbb{E}\left[\min(ratio \cdot advantage, clipped\_ratio \cdot advantage)\right], \quad (41)$$

where *clipped_ratio* is obtained by clamping the *ratio* within the range $[1 - clip\_ratio, 1 + clip\_ratio]$.This clipping mechanism prevents the *ratio* from reaching extreme values, which would otherwise lead to large, destabilizing changes in the policy. By constraining how much the policy can change in each update, PPO ensures more stable and reliable learning.

Reinforcement Learning Algorithm Description

This section describes the training process for each of the previously mentioned algorithms and explain how their hyperparameters are evaluated.
As a starting point, let's consider the dataset divided into training set, validation set and test set. The training set goes from 09/01/2017 to 23/08/2021, the validation one goes from 24/08/2021 to 07/03/2023 and the test one goes from 08/03/2023 to 20/09/2024. The first one covers the 60% of the dataset, the second the 20% and the third one covers the 20%.
After that, since each of the algorithm are based on an actor-critic structure, we will illustrate all the process using the A2C algorithm, as the structure is nearly identical for the other two.
To better explain the algorithm, Figure 42 provides a helpful illustration. The A2C algorithm begins by taking the training set of returns as input, along with the number of episodes, which refers to how many times the agent interacts with the environment to complete a full trajectory (or episode) from start to finish. It also takes the various rates, such as the λ_{risk_rate} and the λ_{fee_rate} , discussed in the previous section, which are used to compute the loss function. These inputs are passed through a preprocessing layer (shown in orange in Figure 42), where the rolling standard deviation and covariance matrix are calculated, these are essential features for the loss computation. The output from the previously mentioned layer, along with the initial weights (representing the asset allocation on day zero), are used to create the environment. This environment is then fed into the Agent Network, which has the same structure as shown in Figure 37. The Agent Network's task is to determine a set of weights for each day. Once these weights are determined, a Loss Actor function is evaluated for each set, following the structure discussed in the previous chapter.
Simultaneously, the environment feeds a state into the Critic Network, a neural network with a structure like that shown

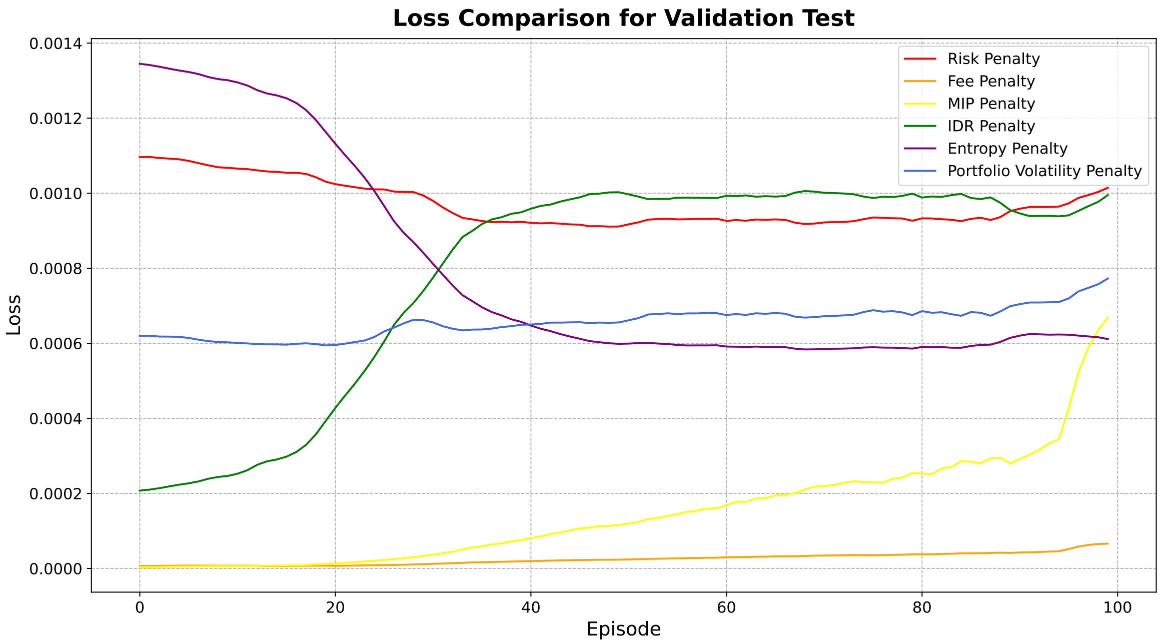

FIGURE 44: In figure the trend of each component of the loss function which is used to penalize the return in order to have a more balanced and conservative portfolio.

in Figure 43. This network is essential for producing an output value that assesses the quality of the decisions made by the actor. Like the Actor Network, the Critic Network must also be trained, but in this case, the loss function to minimize is:

$$\mathcal{L}_{critic} = \sum_{i=0}^{N} [\hat{x}_i - (\hat{x}_{i+1} + cumulative\_returns)]^2, \quad (42)$$

where $\hat{x}_i$ represents the output of the Critic Network for the i-th day, $\hat{x}_{i+1}$ is the output from the same network for the $(i+1)$-th day, and *cumulative_returns* refers to the cumulative returns generated from the initial state (the initial random asset allocation). The output of the Critic Network is used by the Actor to update its policy, utilizing policy gradient methods. The Actor then sends the selected action back to the environment.

The training step is the same for the other two algorithm since the PPO is a more stable and robust version of A2C which uses a policy optimization with clipping update, that prevent drastic policy changes, but follows the same structure illustrated in Figure 42, and DDPG, that, besides it is a off-policy algorithm, uses two separate network along with a replay buffer to store past experiences, which helps in training stability and efficient use of collected samples.

After completing the training step, the validation set is used to fine-tune the algorithm's hyperparameters. The key hyperparameters are the λ rates, which control the weighting of each component in the total loss, $\mathcal{L}_{TOT}$. These values are manually adjusted to maintain a balanced contribution from each component. The resulting trends for each loss component, observed during the validation process of the A2C algorithm, are illustrated in Figure 44.

Initially, the starting weights are considered riskier, as indicated by the high loss they generate. Over the course of the epochs, the weights are adjusted to achieve a more balanced and less risky allocation. However, towards the final epochs, adopting a riskier portfolio leads to higher returns. The fee penalty, represented by the orange line, appears to have minimal influence, as the algorithm consistently chooses small adjustments, keeping transaction costs low. Similarly, the market impact rate behaves in a comparable manner, though its quadratic formulation makes it more impactful than the fee penalty.

The green line, representing the inverse diversification ratio, starts lower due to an evenly balanced portfolio. As the episodes progress, the portfolio reallocates weights, favoring assets with higher potential returns, leading to a more concentrated and unbalanced portfolio.

This shift is also reflected in the entropy loss (violet line), which decreases as the portfolio focuses more on a smaller set of assets. Lastly, the portfolio volatility penalty starts lower due to a less risky allocation. It increases after several episodes but does not reach levels that significantly disrupt the overall portfolio strategy.

The best combination of hyperparameters, hence the one used for the final model is:

$$\begin{aligned}
\lambda_{risk_rate} &= 0.05 = 5\%, \\
\lambda_{portfolio_volatility} &= 0.05 = 5\%, \\
\lambda_{IDR} &= 0.003 = 0.3\%, \\
\lambda_{entropy} &= 0.0005 = 0.05\%, \\
\lambda_{turnover} &= 0.02 = 2\%, \\
\lambda_{market_impact} &= 0.05 = 5\%.
\end{aligned}$$

Results

The results of the performances of each method described so far are showed for the test set, so for data going from March 2023 to September 2024, which covers the 20% of the entire dataset.

In this chapter the algorithm explained before is compared to other conservative techniques which mantain the same weights over time, without any kind of change. *Portfolio 1* is the uniform weight portfolio, a type of investment strategy where each asset in the portfolio is allocated the same proportion of the total investment. *Portfolio 2* is a portfolio which uses the weights decided at the beginning but keeping it uniform for the rest of the time, rather *Portfolio 3* is the max sharpe portfolio an investment strategy that seeks to allocate assets in a way that maximizes the Sharpe Ratio of the portfolio. The Sharpe Ratio measures the risk-adjusted return, which helps investors understand how much excess return they are receiving for the extra risk they take on.

The other portfolio strategies, are the three dynamic models that adjust weights daily and have benne explained in the previous chapters: *Portfolio 4* which uses the Actor to Critic method, *Portfolio 5*, which employs Proximal Policy Optimization and *Portfolio 6*, which utilizes Deep Deterministic Policy Gradient.

Figure 45 illustrates the cumulative returns of each portfolio from March 2023 to September 2024, calculated using the following formula:

$$CR_i = \prod_{j=0}^{i} (r_i + 1), \tag{43}$$

where r_i represents the daily return.

From the graph, it's evident that both strategies with fixed uniform weights and with pre-determined weights underperform, resulting in cumulative returns of $CR_{uniform\_weights} = 0.999$ (blue line) and $CR_{our\_weight} = 1.003$ (orange line) by the end of the observation period.

In contrast, *Portfolio 3*, which aims to maximize the Sharpe Ratio, shows strong performance with a final cumulative return of $CR_{max\_sharpe} = 1.581$ (green line). However, as seen in Figure 46, this approach tends to concentrate over 80% of the portfolio in just four assets, leading to higher risk and reduced diversification.

On the other hand, the dynamically managed portfolios (*Portfolio 4,5,6*) developed in this study achieved cumulative returns of $CR_{A2C} = 1.318$, $CR_{PPO} = 1.445$ and $CR_{DDPG} = 1.409$. Although lower than Portfolio 3, these portfolios provide a more balanced risk distribution across the year.

Figure 47 illustrates the daily fluctuations in the weights assigned to the stock AIR.PA. It is evident that the weight never exceeds 6%, and there are instances when the stock is not included in the portfolio at all, resulting in a weight close to 0%.

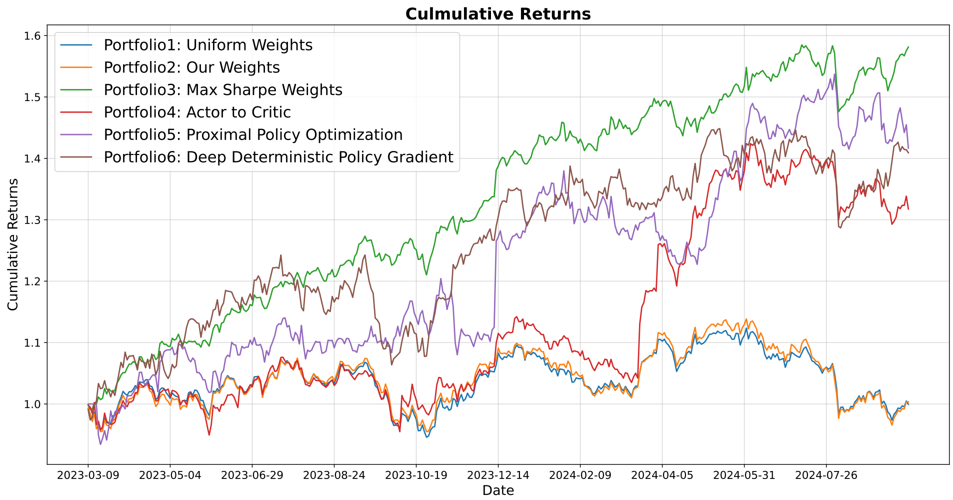

FIGURE 45: The plot displays the cumulative returns of our portfolio using various techniques. The blue, orange, and green lines represent three static portfolios: the first with uniform weights, the second with preselected weights, and the third with weights optimized to maximize the Sharpe ratio. In contrast, the red, purple and brown lines correspond to dynamic portfolios, each managed using different reinforcement learning techniques introduced in previous chapters—A2C, PPO, and DDPG. Over the one-year analysis period, the static portfolios show varying cumulative returns, with the first two yielding minimal gains. Meanwhile, the dynamic portfolios demonstrate significantly higher returns. Although Portfolio 3 (optimized for the Sharpe ratio) achieves the highest return, it tends to be less balanced, heavily favoring certain assets. In comparison, the reinforcement learning-based portfolios deliver strong returns while maintaining better balance due to constraints applied during the weight allocation process. A significant reward spike is observed for the A2C and PPO algorithms, indicated by the large vertical lines in December 2023 and March 2024. These spikes occur because, during training, the network identified an allocation strategy that maximized returns. This approach allocated nearly 70% of the portfolio across six assets, resulting in a high-risk profile similar to that of a Sharpe ratio-based strategy.

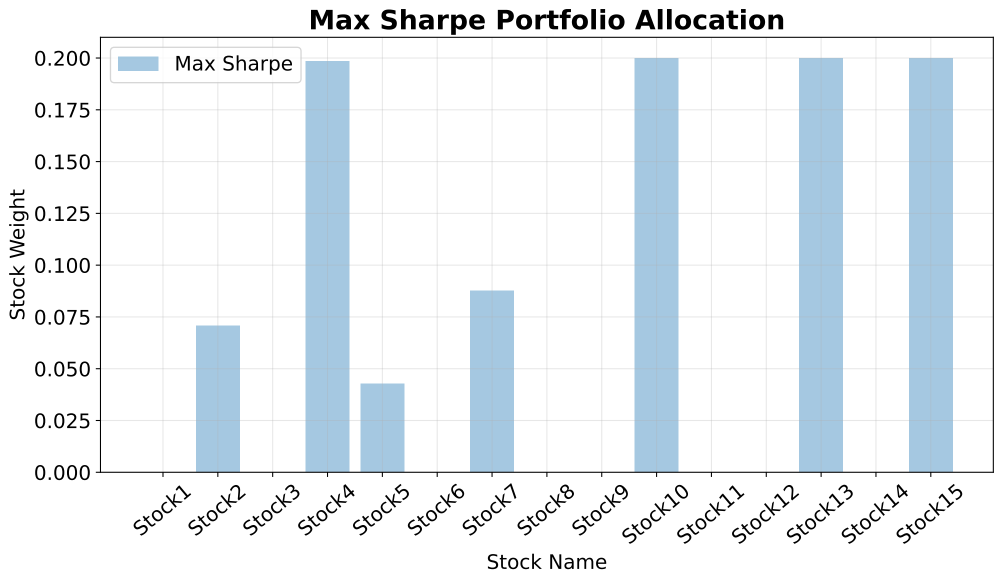

FIGURE 46: The figure shows the asset allocation aimed at maximizing the Sharpe Ratio, it is evident that the portfolio primarily concentrates on seven out of the fifteen available assets. Notably, four of these assets account for 80% of the portfolio, leading to reduced diversification and, consequently, increased risk.

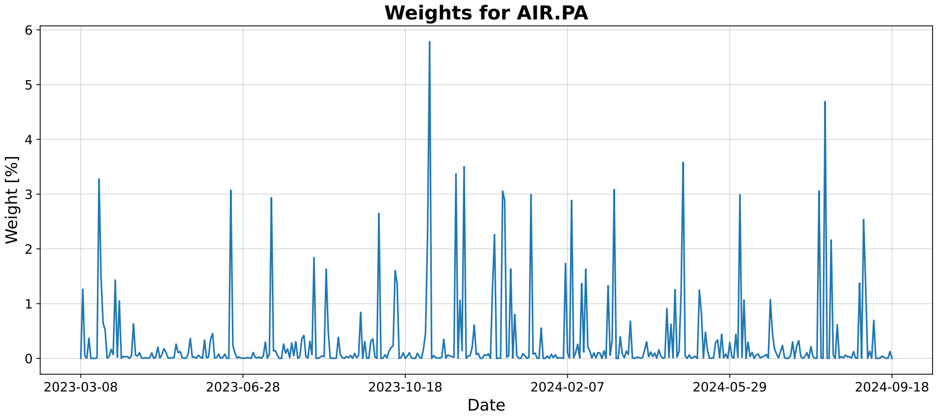

FIGURE 47: Example of weight allocation for AIR.PA with the A2C method.

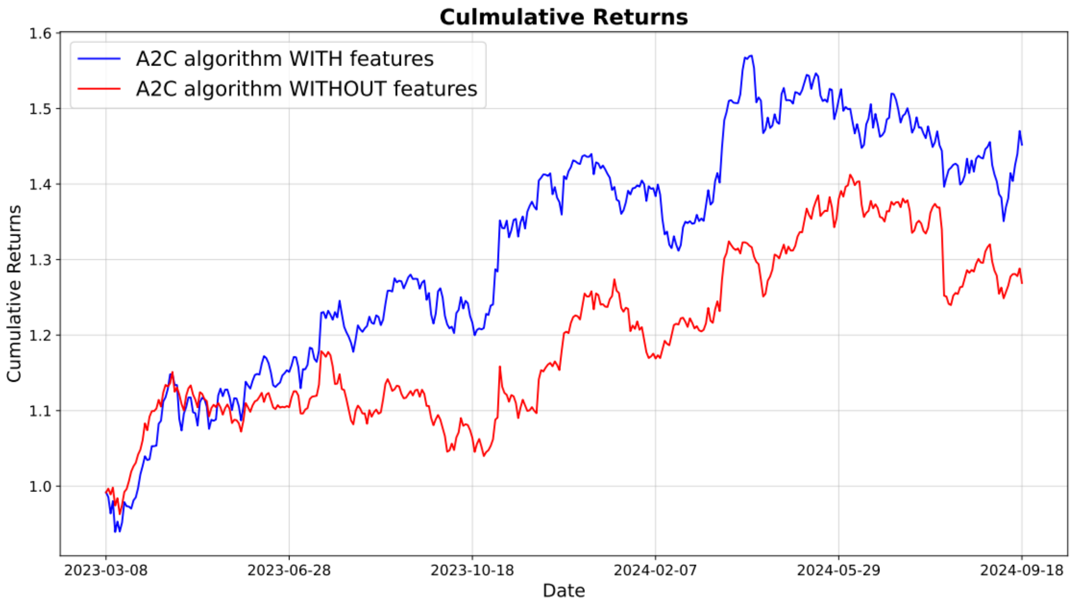

FIGURE 48: Comparison of the A2C algorithm with features (blue line) and without features (red line) shows that, by the end of the analysis period, the blue line achieves a higher cumulative return than the red line. This indicates that the model utilizing synthetic financial features delivers superior results.

To highlight the effectiveness of using this approach with four synthetic features derived from financial metrics commonly used by traders, we compared this approach with the same model using only returns as input. Figure 48 presents the results, showing the A2C algorithm with features (blue line) and without features (red line). The model incorporating financial features consistently outperforms the model without features, with the performance gap becoming more pronounced over longer periods.
To demonstrate the effectiveness of these algorithms on a different type of portfolio, each was applied to a benchmark dataset—the Dow Jones index. As illustrated in Figure 49, the original index with its existing weights achieved a cumulative return of $CR_{Dow\_Jones} = 1.118$. In contrast, the new methods produced higher returns, as shown by the orange line (A2C), the green line (PPO), and the red line (DDPG). Specifically, the cumulative returns are $CR_{A2C} = 1.490$, $CR_{PPO} = 1.427$ and $CR_{DDPG} = 1.634$, indicating significantly better performance compared to the actual returns of the index.

Therefore, it is demonstrated that these algorithms perform effectively even with different types of equity-based portfolios or random indices, such as the Dow Jones, which was used as an example in this study.
In the loss definition implemented in this work, the risk associated with the portfolio is included. However, to further evaluate the risk of the portfolio proposed by our method, an additional risk measurement is introduced.
As an alternative measure of risk, we use the drawdown, which assesses the decline in the value of an investment or portfolio from its peak to its lowest point over a specific period. It indicates the potential loss an investor would incur if they bought at the peak and sold at the trough.
The drawdown is calculated as the difference between the peak value of the portfolio and the subsequent lowest value (trough), expressed as a percentage of the peak value. The formula is:

$$Drawdown = \frac{Peak\_Value - Trough\_Value}{Peak\_Value}. \tag{44}$$

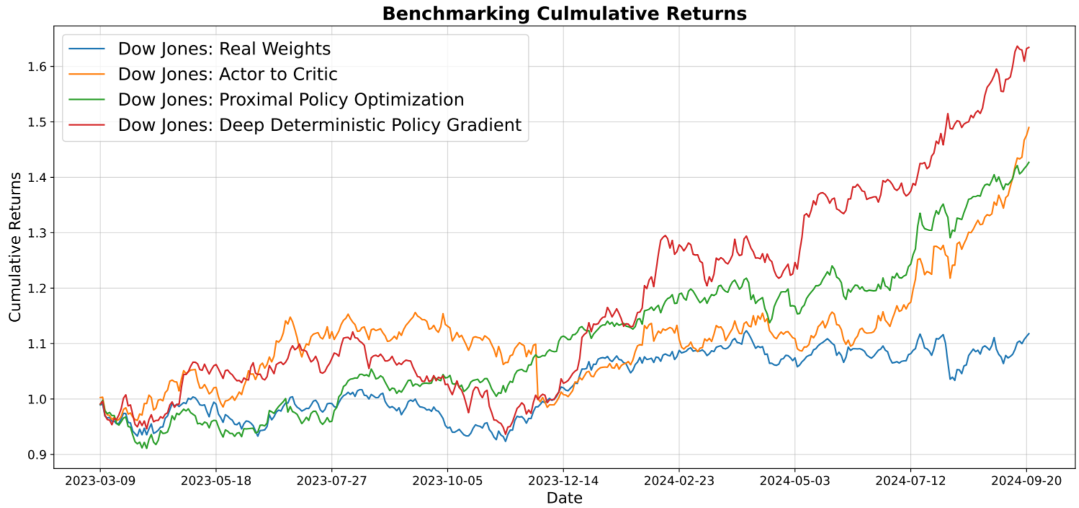

FIGURE 49: *The plot presents the cumulative return of the benchmark portfolio, the Dow Jones index. The blue line represents the actual performance of the index, showing a return of approximately 10% at the end of the one-year evaluation period. In contrast, each of the algorithms introduced in this study delivers significantly higher returns, all exceeding 40%, as illustrated by the orange, green, and red lines in the figure.*

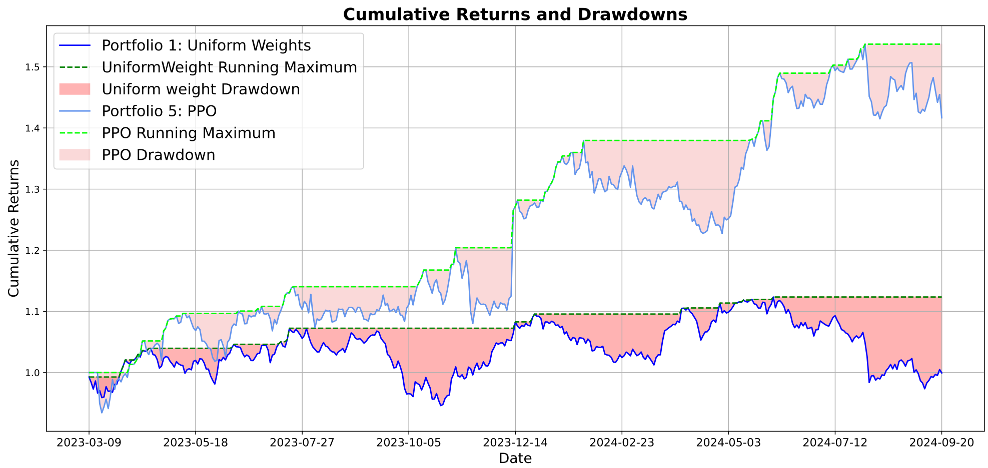

FIGURE 50: *Drawdown comparison for Portfolio 1 and Portfolio 5: The dark blue line represents the uniform weights portfolio, which exhibits lower cumulative returns and experiences longer periods of loss, indicated by the darker red shaded areas. In contrast, the PPO portfolio shows higher cumulative returns with shorter drawdown periods, highlighted by the lighter red areas.*

Figure 50 presents the drawdown comparison between*Portfolio 1* and *Portfolio 5*, illustrating the risk associated with a uniform weight strategy versus a dynamic allocation approach. These two portfolios are selected to compare a static allocation strategy with a dynamic one.
As shown, *Portfolio 1* experiences longer drawdown periods. For a clearer comparison of the percentage drawdowns, Figure 51 provides a more detailed visualization. It highlights that, during certain periods, the uniform weight portfolio exhibits higher drawdowns than the dynamically allocated PPO portfolio. Notably, the maximum drawdown for *Portfolio 1* (marked by a dark blue dot) is significantly greater than that of *Portfolio 5* (light blue dot).
Similarly, in Figure 52, the maximum drawdown is illustrated for two Dow Jones portfolios. The dark blue line represents the cumulative returns of the real weights index, with the darker red areas indicating periods of drawdown. In contrast, the light blue line shows the cumulative returns achieved using the PPO algorithm, with the lighter red areas highlighting its drawdown periods.
The difference in drawdown becomes even clearer in Figure 53, where the maximum drawdown for both portfolios is plotted. The red line represents the maximum drawdown for the real weights portfolio, while the green line shows the maximum drawdown for the PPO portfolio. As seen, the real weights portfolio experiences higher maximum percentage drawdowns, suggesting a higher risk level compared to the PPO portfolio, which appears to be less risky due to its construction.

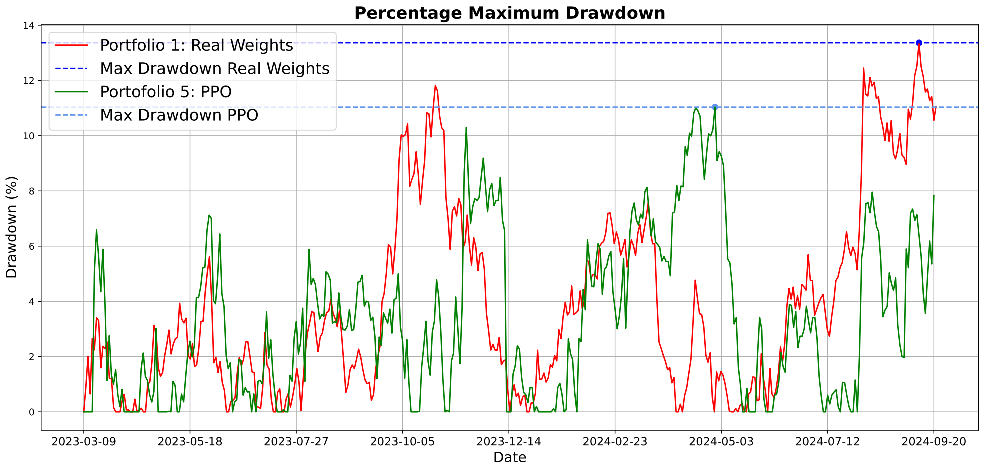

FIGURE 51: *The percentage maximum drawdown for two portfolio analyzed in previous chapter is shown, with the red line representing the Portfolio 1 (uniform weights portfolio). Its maximum drawdown, marked by the dark dotted blue line, exceeds 13.5%. In contrast, the green line, which corresponds to Portfolio 5 (PPO portfolio), demonstrates a lower drawdown over time, peaking at around 11%.*

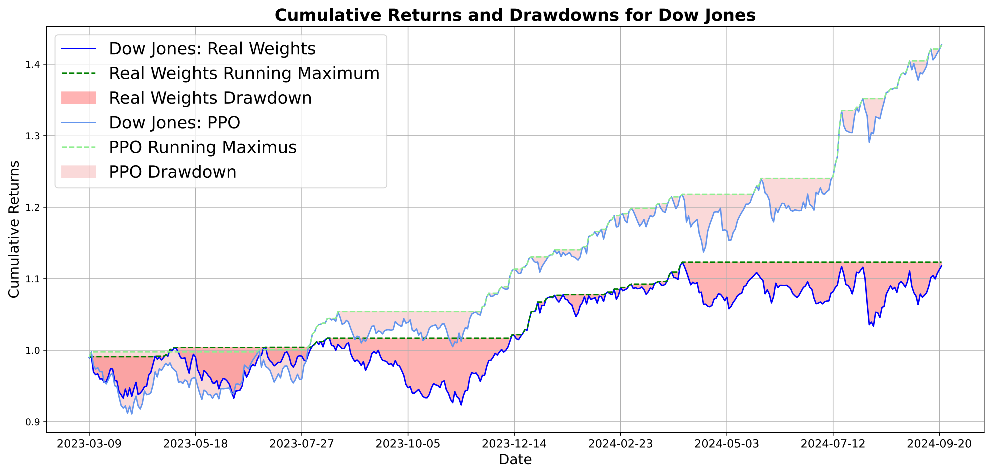

FIGURE 52: *Drawdown comparison for two Dow Jones portfolios: The dark blue line represents the real weights portfolio, which exhibits lower cumulative returns but experiences longer periods of loss, indicated by the darker red shaded areas. In contrast, the PPO portfolio shows higher cumulative returns with shorter drawdown periods, highlighted by the lighter red areas.*

Conclusion and Future Works

This work introduced a novel approach to portfolio allocation by integrating NN and RL, demonstrating how advanced machine learning techniques can enhance investment strategies. By incorporating various financial features traditionally used by traders, the proposed method allows for a more comprehensive analysis of market dynamics. The use of a VAE-LSTM model effectively condensed complex financial data, with the VAE capturing the underlying structure and the LSTM identifying critical temporal patterns. Subsequently, these features were combined with the original dataset to predict optimal portfolio weights, utilizing RL techniques such as A2C, PPO and DDPG. This dynamic approach to portfolio allocation allows for continuous adaptation to evolving market conditions, thus providing a robust solution for maximizing returns while managing risk. The results highlight a significant advantage of dynamic allocation strategies over traditional static weight portfolios, enhancing both returns and risk management. Each algorithm applied in this study outperformed classic methods, offering improved drawdown control; however, they tended to be slightly more conservative and less effective when compared to the maximum Sharpe ratio approach. Additionally, our findings show that portfolio allocation using financial features outperforms methods that base weights solely on return data. These advantages are consistently seen in benchmark equity portfolios, such as the Dow Jones index, underscoring the versatility of these algorithms across various equity portfolios as they depend mainly on return data. Future research could expand these strategies to a wider range of financial instruments, including bonds, futures, and additional asset classes. Furthermore, applying these

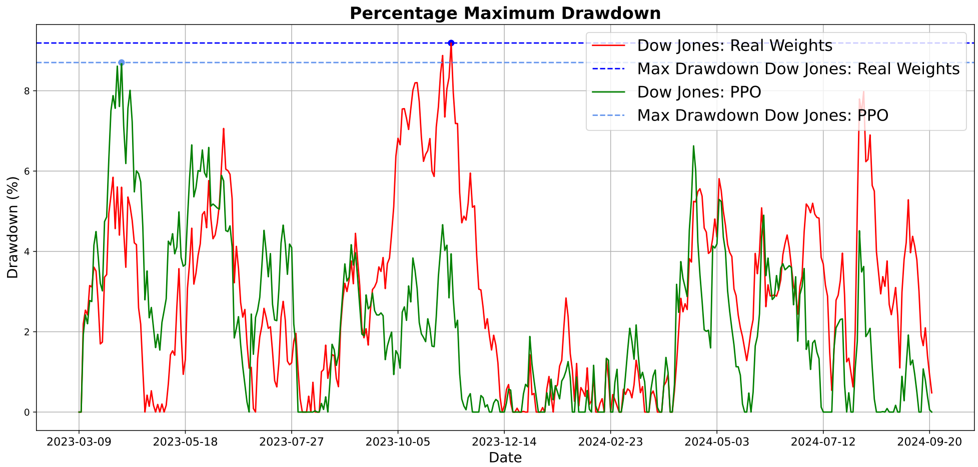

FIGURE 53: The percentage maximum drawdown for two Dow Jones portfolios is shown, with the red line representing the real weights portfolio. Its maximum drawdown, marked by the dark dotted blue line, exceeds 9%. In contrast, the green line, which corresponds to the PPO portfolio, demonstrates a lower drawdown over time, peaking at around 8.5%.

methods to other indices could broaden their applicability. Incorporating correlations between asset classes as an added feature could enhance weight allocation strategies. Further insights might also be gained by exploring alternative risk measures in loss evaluation, such as Conditional Value at Risk (CVaR), or by examining additional reinforcement learning techniques to compare with the observed performance in this study.

References

[1] **Almahdi, S. and Yang, S. Y.** *An adaptive portfolio trading system: A risk-return portfolio optimization using recurrent reinforcement learning with expected maximum drawdown*. Expert Systems with Applications, Vol. 87, pp. 267-279, 2017.

[2] **Appel, G.** *Technical Analysis: Power Tools for Active Investors*. Financial Times and Prentice Hall, 2005.

[3] **Baird, L.** *Residual Algorithms: Reinforcement Learning with Function Approximation*. Morgan-Kaufmann, 1995.

[4] **Bertsekas, D. P.** *Dynamic programming: deterministic and stochastic models*. 1987.

[5] **Bo, Q. and Khaled, R.** *Stock market prediction with multiple classifiers*. Applied Intelligence, 2007.

[6] **Bollinger, J.** *Bollinger on Bollinger Bands*. McGraw Hill, 2001.

[7] **Bradtke, S.** *Reinforcement Learning Applied to Linear Quadratic Regulation*. Morgan-Kaufmann, 1992.

[8] **Buehler, H., Gonon, L., Teichmann, J. and Wood, B.** *Deep hedging*. Quantitative Finance, Vol. 9, N. 8, pp. 1271-1291, 2019.

[9] **Chiang, W. C., Urban, T. L. and Baldridge, G. W.** *A neural network approach to mutual fund net asset value forecasting*. Omega, Vol. 24, N. 2, pp. 205-215, 1996.

[10] **Chi-Ming, L., Jih-Jeng, H., Mitsuo, G. and Gwo-Hshiung, T.** *Recurrent neural network for dynamic portfolio selection*. Applied Mathematics and Computation, Vol. 175, N. 2, pp. 1139-1146, 2006.

[11] **Freitas, F. D., De Souza, A. F. and De Almeida, A. R.** *Prediction-based portfolio optimization model using neural networks*. Neurocomputing, Vol. 72, N. 10, pp. 2155-2170, 2009.

[12] **Granville, J.** *New Key to Stock Market Profits*. 2010.

[13] **Heaton, J. B., Polson, N. G. and Witte, J. H.** *Deep Learning for Finance: Deep Portfolios*. Applied Stochastic Models in Business and Industry, Vol. 33, N. 1, pp. 3-12, 2016.

[14] **Hochreiter, S. and Schmidhuber, J.** *Long Short-Term Memory*. Neural Computation, Vol. 9, N. 8, pp. 1735-1780, 1997.

[15] **Jiang, Z. and Liang, J.** *Cryptocurrency Portfolio Management with Deep Reinforcement Learning*. 2017.

[16] **Jiang, Z., Xu, D. and Liang, J.** *A Deep Reinforcement Learning Framework for the Financial Portfolio Management Problem*. 2017.

[17] **Kingma, D. P. and Welling, M.** *Auto-Encoding Variational Bayes*. 2022.

[18] **Kullback, S. and Leibler, R. A.** *On Information and Sufficiency*. The Annals of Mathematical Statistics, Vol. 22, N.1, pp. 79-86, 1951.

[19] **Liang, Z., Chen, H., Zhu, J., Jiang, K. and Li, Y.** *Adversarial Deep Reinforcement Learning in Portfolio Management*. 2018.

[20] **Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D. and Wierstra, D.** *Continuous control with deep reinforcement learning*. 2015.

[21] **Lund, B.** *Unlocking the power of the Chaikin money flow indicator*. Encyclopedia Britannica, 2023.

[22] **Markowitz, H.** *Portfolio Selection.* The Journal of Finance, 1952.

[23] **Niaki, S. T. A. and Hoseinzade, S.** *Forecasting S&P 500 index using artificial neural networks and design of experiments*. Journal of Industrial Engineering International, Vol. 9, 2013.

[24] **Obeidat, S. F., Shapiro, D., Lemay, M., MacPherson, M. K. and Bolic, M.** *Adaptive Portfolio Asset Allocation Optimization with Deep Learning*. 2018.

[25] **Skapura, D. M.** *Building Neural Networks*. 1996.

[26] **Watkins, C. J. C. H.** *Learning from delayed reward*. 1989.

[27] **Welles Wilder, J.** *New Concepts in Technical Trading Systems*. 1978.

[28] **Williams, R. J. and Baird, L. C.** *Tight performance bounds on greedy policies based on imperfect value functions*. 1993.

[29] **Yang, H., Liu, X. Y., Zhong, S. and Walid, A.** *Deep Reinforcement Learning for Automated Stock Trading: An Ensemble Strategy*. 2020.

Sitography

[30] **Fidelity Investments.** *Website*.

[31] **IG Bank.** *Website.*

[32] **Trading View.** *Website*.

Systematic Machine Learning Asset Allocation Benchmarking

About the Authors

Evgenii Veksin:
Analyst
He holds bachelor's degree in economics at Higher School of Economics (Moscow). Secondly, he holds a master's degree in Management Engineering at Politecnico di Milano. In Iason, he was involved on the Credit Risk topics: Stress Test project, internal GRisk models. Currently, he is involved in Asset Allocation project.

This document was prepared in collaboration with Paolo Bortolotti who at the time was working for Iason Consulting.

Systematic Machine Learning Asset Allocation Benchmarking

Paolo Bortolotti Evgenii Veksin

We present a case study in which a systematic investment process is implemented based on signals derived from a machine learning model. This case arises from a specific experience conducted by Iason Ltd in response to various clients requests for support. The focus is on describing the process and emphasizing the potential use of such a framework for both direct investment purposes and as a benchmark for the internal performance evaluations of "traditional" management portfolios. This concept emerged from a client's need for a cost-effective and easily implementable tool that can be used to evaluate discretionary investment teams by comparing their performance to that of systematic portfolios operating under the same constraints and "advantages" afforded to the investment team. The modeling framework use an integrated approach where Unsupervised Learning techniques are integrated in a Supervised Learning Regression algorithms. Four distinct portfolio implementations, based on investment signals from the developed model have been presented. All portfolios demonstrate risk-adjusted performance that exceeds that of the benchmark established in the investment mandate. The different implementations vary in their levels of compliance with the investment risk framework agreed upon with the final investor, with only the last of the four achieving full compliance with all investment limits while still maintaining more than adequate outperformance.

This paper presents a case study synthesizing recent developments and modeling activities undertaken by the Iason team during consulting engagements in the Asset Management field. The processes and models presented are not specific to any one client; instead, they represent a simplified and stylized exercise[3]. Nonetheless, they serve as a fair example of the types of activities we have conducted for our clients. The work is presented as an informative case study, with the technical and mathematical content minimized to facilitate reading and comprehension. Additionally, it is accompanied by a series of appendices that explain the subjects presented in a more technical manner.

Introduction: Background and Client Requests

The client had the following needs[4]:

- **Development of Quantitative Tools to Support Investment Processes:** The Investment Committee requested the creation of a fully systematic asset allocation model for direct use that would comply with the investment guidelines agreed upon with investors and adhere to the internal risk management framework.
- **Innovation in Investment Processes:** Senior management and the Investment team sought concrete support in adopting machine learning techniques within investment processes, aiming to innovate corporate practices and enhance skill sets.
- **Automation:** The Investment and Risk teams expressed a need to improve automation processes. We were specifically tasked with implementing processes based on Python scripts and professional data management using relational databases.
- **Performance Evaluation:** The board required the development of a systematic model portfolio to serve as an operational benchmark, in addition to the actual mandate benchmark. The goal was to utilize this systematic model as a minimally viable, realistic alternative for comparing corporate investment processes.
- **Passive versus Active Management:** One client requested the development of a systematic active management tool to be used tactically alongside a purely passive replicating portfolio (ETF) to enhance the fund's performance within the agreed-upon risk framework.

The model framework developed can be used directly or indirectly, where a particularly skilled and professional investment team is present. In the latter scenario, the systematic model serves as a comparative tool for the discretionary choices made by the investment team, enabling verification of deviations at the end of each month. The client established an internal performance fee structure based on the outperformance achieved by the investment team compared to what the model generated, all at the same and comparable level of risk. The rationale was to compensate (sometimes significantly) for all "human" discretionary investment decisions that led to outperformance over the systematic investment decisions achieved by the model. In other words, the board of the company, having a concrete, realistic, low-cost, and efficient investment alternative at

[3]The model settings are intentionally suboptimal and unoptimized. This approach allows the reader to focus on the process and the potential value added, rather than on achieving the best risk-adjusted performance.

[4]The needs are presented as a single list from one client; however, in reality, they represent a collection of requests from various clients across different projects. Not all clients had all the needs listed.

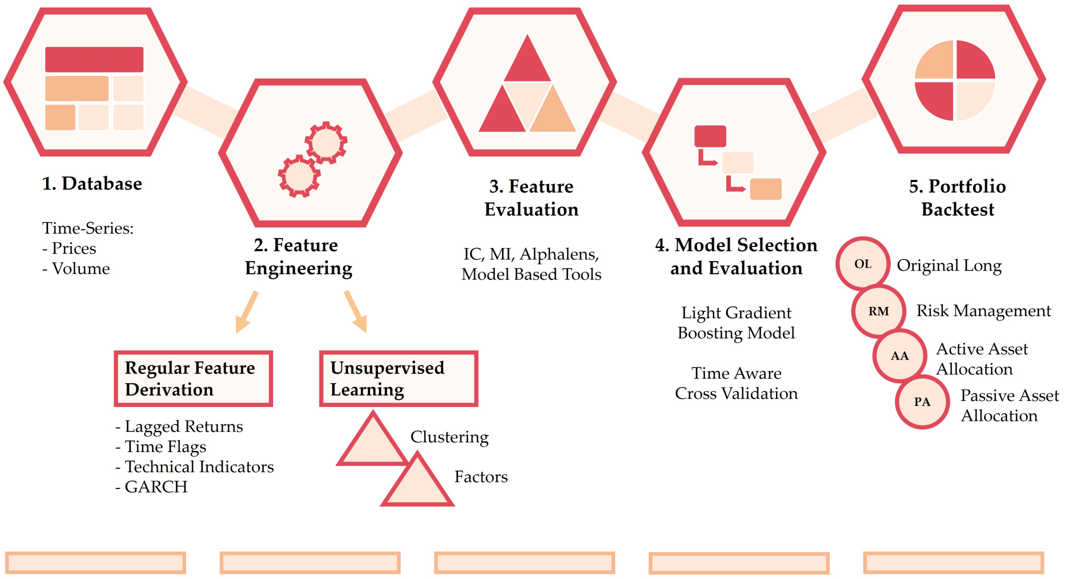

FIGURE 54: Process Workflow.

their disposal, decided to reward only those portfolio managers who genuinely added value compared to systematic approach.
Supporting this view from the client is the fact that, unlike a traditional benchmark index, the model operates on the same playing field as the portfolio manager, sharing the same inefficiencies inherent in the manager's investment activities - such as trade optimization, transaction cost reduction, investment timing, and irregular rebalancing. Simultaneously, the model offers the same advantages (albeit with increased risk) provided to portfolio managers by the investment mandate. This includes the ability to quickly adjust asset allocations and concentrations of securities in the portfolio, as well as the flexibility, often afforded by the investment mandate, to allocate portions of the portfolio to securities that fall outside the benchmark constituents (known as "off-benchmark" investments or allocations that deviate from a specified benchmark index).
Consequently, the board of the client considers outperforming the systematic model portfolio to be the minimum acceptable performance, with the systematic model serving as a managerial tool for evaluating and comparing the performance of managed portfolios.

Portfolio Risk Management Risk and Mandate

The portfolio constructed based on the model had to comply with at least the following minimum risk management framework and mandate requirements:

1. Daily VaR limits of 2%;
2. TEV (Tracking Error Volatility) limits of 8%;
3. Long positions only; short selling is not permitted;
4. Intraday trading not allowed; positions are daily;
5. Investments permitted only within the investment universe defined by the constituents of the index benchmark[5].

Model Proposal

Given the aforementioned requirements, risk framework constraints, and the typology of clients (small to medium firms), we proposed the simplest internally developed model. This model, in fact, has the advantage of being easily implementable even in contexts with a low level of quantitative sophistication, such as those found in small-sized investment firms. It relies on minimal and nearly free financial data, specifically time series of prices (OHLC)[6] prices and volumes of the benchmark's constituents and it is independent of external or proprietary data sources, such as costly sentiment analysis databases or private text data repositories.

Integrated Approach of Unsupervised and Supervised Learning

The **key aspect of the model** developed is an **integrated use of unsupervised and supervised machine learning methods to identify and select investment decisions.** This combined approach leverages the strengths of both paradigms to enhance overall predictive power and generate actionable insights from the data. Let's first briefly describe the approach:

- We start from applying Unsupervised Machine Learning tools to the set of explanatory variables.
- This extended data representation is an input for the Supervised Machine Learning model where label variable is introduced.

By using unsupervised learning to generate new features, models can capture important relationships and underlying structure from complex data structures. This should lead to a more accurate and generalizable performance of the final model.

[5]This limit can be relaxed to allow the model to have the same investment universe as the portfolio manager specified in the client's investment mandate. For simplicity, we assume that both the model and the portfolio manager can invest only in the constituents of the benchmark.
[6]The time series of prices taken are the so called OHLC - Open, High, Low, Close.

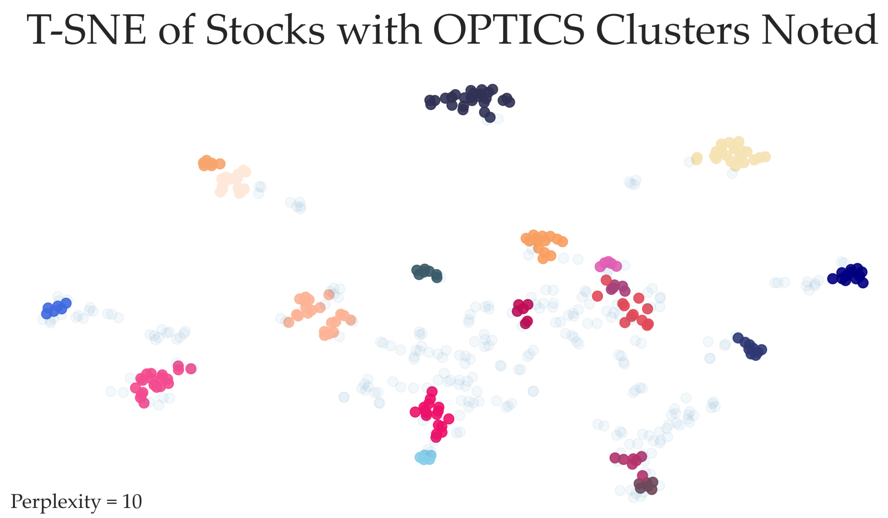

FIGURE 55: T-SNE visualization of clustering.

Model Framework and Development Process

This case study outlines the steps for developing and managing a machine learning-based systematic equity portfolio, using the S&P 500 index as the benchmark. Although this is a simplified and illustrative example, the same concepts have been applied to strategies benchmarked against combinations of equity indices or blends of indices that include assets from different asset classes (commonly referred to as balanced portfolios).
The steps followed in the portfolio construction process are as follows: Database Setup, Feature Engineering, Feature Evaluation, Model Selection, Model Evaluation, Portfolio contraction and Backtesting.
The figure 54 outlines the workflow of the process undertaken to develop, test, and select the model/portfolios.

Database Setup

The first step in the process involves populating a relational database with data obtained from information providers, including historical prices and volumes for the constituents of the benchmark. To comply with the investment processes, we need to archive the daily historical series for all tickers that make up the S&P 500. This is essential because the model can only invest in stocks within the benchmark, accounting for new entrants and divesting from stocks that exit the benchmark each day.
Additionally, to facilitate analysis, calibration, and backtesting, we must archive the historical series of prices and volumes for all stocks that have entered and exited the S&P 500 over time. Therefore, the database tables will need to contain information on significantly more tickers than just the 500 constituents of the benchmark.

Feature Engineering and Unsupervised Learning

Feature engineering is a fundamental step in the process of creating Machine Learning models. It involves the creation, transformation, and selection of features (input variables) from raw data to enhance the performance of predictive models. For our analysis we divide arbitrarily this phase in two distinct sub-process:

Regular Features Derivation

The features calculated are very standard and well known, for the sake of brevity, it is not possible to provide detailed information on each individual feature derived, instead, we will categorize them as follows:

- **General Variables:** lagged returns and time indicators;
- **Momentum Technical Analysis Features:** Features derived from momentum indicators that help capture the strength and direction of price trends;
- **Volume Technical Analysis indicators:** Features based on trading volume data, which help to assess the strength of price movements and potential reversals;
- **Volatility Metrics and Estimates Features:** Features that measure price and returns volatility and incorporate estimates to provide insights into potential price movements and risks.

Applying Unsupervised Learning

This phase aims to create features by integrating results from an unsupervised method into the supervised model. We have generated two types of features: one related to risk factors and the second related to clustering. Both features are derived from a reduced form of the data representation produced via PCA[7]. The steps are as follows:

[7]Principal Component Analysis (PCA) is a statistical technique used for dimensionality reduction while preserving as much variability as possible in a dataset. It transforms the original correlated variables into a new set of uncorrelated variables known as principal components.

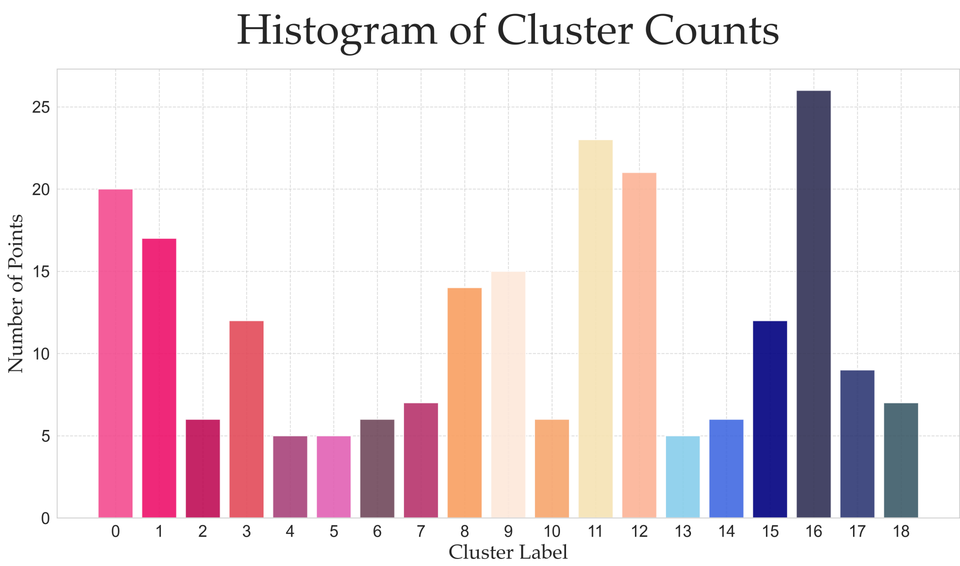

FIGURE 56: *Structure of Clustering.*

1. Reduce the dimensionality of the data at hand;
2. From the reduced data-set, generate clusters:
   - Derivate from PCA the risk premiums (or the eigenvectors) for each ticker;
   - Apply unsupervised algorithms (like OPTICS) to the tickers/risk premiums to generate clusters or groups of stocks that share common behavior in terms of variance;
   - Use these clusters as categorical features.
3. From the reduced data-set, generate a Principal Components Factors Time Series:
   - From this, calculate the betas of each stock based on these factors' data series;
   - Use these betas as predictive features.

The following sections explain each individual step described above[8]:

1. **Reduced Data Dimensionality**
   This step is instrumental for accurately estimating both clusters and risk factors. By reducing the dimensionality of the data, we simplify the complex dataset, minimize noise that can obscure patterns, mitigate the risk of overfitting, and enhance computational efficiency. To achieve this compact representation for each asset, we apply Principal Component Analysis (PCA) to the dataset, resulting in a reduced data representation. In particular from a matrix of price time series with dimensions 1260 (dates) x 500 (tickers/stocks) to a matrix of principal components with dimensions 1260 (dates) x 15 (principal factors).
2. **Cluster Determination Features:**
   This application infer meaningful clusters of assets. Here, the clusters are stocks that share common patterns in variance. The sequential steps to arrive at the identification of the clusters starts from dimensionality reduction with Principal Component Analysis (PCA) described above. On this reduced form of the data set we apply a clustering algorithm. We choose OPTICS algorithm, which belongs to the family of density-based clustering algorithms, is chosen because it can form clusters of arbitrary shapes, eliminating the need for Gaussian assumptions. It is very flexible because it does not require all stocks to receive a cluster label. The clustering application is dynamic: each 3 months the clustering structure changes[9] in order to capture the latest information and for coherence with current S&P 500 constituents. The investment universe of S&P 500 Index constituents is grouped to 15-30 clusters with 5-35 stocks within. The approximate structure of the clusters can be visualized using the t-SNE visualization tool, which is a manifold dimensionality reduction technique used for visualizing high-dimensional data. In this bidimensional representation, 19 clusters are displayed. Stocks that are not grouped are considered not to exhibit common behaviors. Figure 55 illustrates one example of the 19 clusters generated from the stocks representing the constituents of the S&P 500. Figure 56 illustrates the number of stocks in each cluster generated by the algorithm.
3. **Data Driven Factor Decomposition:**
   Factor analysis, in portfolio risk management, is a statistical method used to identify and understand the underlying factors that influence the returns and risks of a portfolio. By grouping the variability of asset returns into a smaller number of common factors, investors can gain insights into the sources of risk and performance in their portfolios (Avellaneda & Lee [2]), (Lustig et al. [9]). Factors can be derived from various sources, including financial theory, market indicators, and data-driven techniques like Principal Component Analysis (PCA). We opted to use the PCA reduced data set in point 1 to generate factors time series.

[8]For a detailed methods applied please refer to the appendix called "Unsupervised Learning Application", where the Concept of PCA, Data Implied Factors and Clustering Algorithms OPTICS is explained in more details.
[9]This also is consistent with the calibration of the model as we will see later.

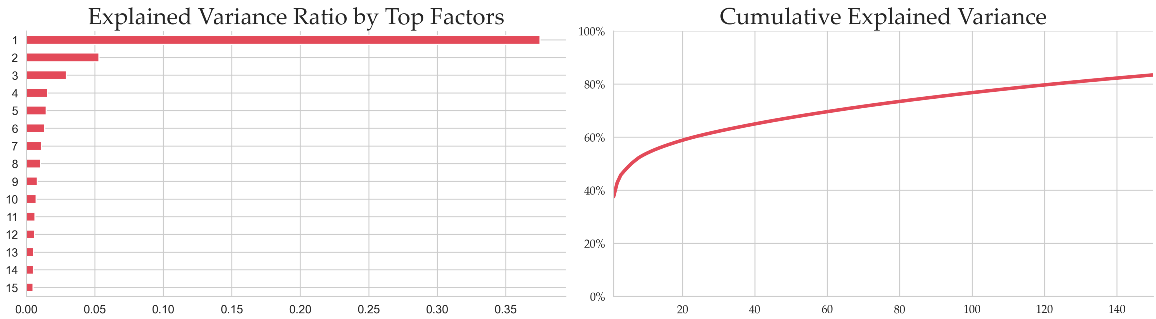

FIGURE 57: Variance structure and convergence.

Group	Variable	Selected
General Variables	Lagged returns family	Yes
	Time indicators: weekday, month, year	Yes
Momentum Indicators	Hilbert Transform	Yes
	Stop And Reverse Indicator	Yes
	Average Directional Movement Index	Yes
	Average Directional Movement Index Rating	No
	Percentage Price Oscillator	Yes
	Aroon Oscilator	No
	Balance Of Power	No
	Commodity Channel Index	No
	Ultimate Oscillator	Yes
	Exponential Regression, Clenow momentum	Yes
	MACD: Original, Signal, Hist	Yes
	Relative Strength Index	No
	Stochastic Relative Strength Index	Yes
	Stochastic Oscillator	No
	Williams %R Indicator	No
	Momentum	Yes
Volume Indicators	Chaikin Advance/Decline	Yes
	Chaikin Advance/Decline Oscillator	No
	On Balance Volume	Yes
	Money Flow Index	No
	Dollar volume	No
	Dollar volume rank	No
Volatility Indicators	GARCH	Yes
	Average True Range	Yes
	Bollinger Bands: Up, Low, Squeeze	Yes
Unsupervised Learning Features	Returns top-15 Principal Components Betas	Yes
	OPTICS Cluster	Yes

TABLE 10: List of Variables.

A risk factor model utilizes a selection of principal components as features, which are then employed in the supervised learning algorithm to forecast future returns. The following graphs illustrate the distribution of variance among the principal components and the dynamic convergence toward a maximum of 100% explanation as the number of principal components increases. Since PCA is a dimensionality reduction technique and we aim for a parsimonious model, we have selected the top 15 principal components, which collectively explain approximately 55% of the variance, as presented in figure 57.

Feature Evaluation

Feature evaluation is a component of the machine learning pipeline, aiming to determine the relevance and importance of input variables in predictive modelling. This process helps identify which features significantly contribute to the model's accuracy, reducing complexity and improving performance. By performing a preliminary feature evaluation, the end result (often) is in model performance improvement, model interpretability, and in an efficient computation.
In data science, various methods are available to evaluate features before using them in predictive models. The following methods have been applied: Information Coefficient, Mutual Information, Preliminary Model-Based Evaluation (Feature Importance and Shapley Additive Explanation), and Alphalens[10].
The metrics applied are based on different concepts, and sometimes the results can be controversial. Table 10 presents the features originally collected for the analysis alongside those that were ultimately selected for use in the model.

Model Selection

Our task is to replicate systematically the investment process of a retail asset manager. Typically, this business does not engage in daily trading; instead, portfolio positions are held for the medium to long term. Generally, portfolios are adjusted on a weekly/monthly basis, depending on market conditions and asset class. For the implementation of the case study this translates into a model that must predict the five-day forward return for each asset within the investment universe, as defined by the investment mandate, using the set of financial features we have previously selected. This type of task falls under regression in supervised machine learning.

Model Selected

The model selected was the light GBM of the family of Gradient Boosting Models[11]. This is an ensemble learning technique that combines the predictions from multiple weak learners, typically decision trees, to produce a strong predictive model.
The model is applied to the entire investment universe, which includes each of the 500 individual constituents of the S&P 500[12]. It uses the predictive relationships of the features selected in the previous process (features engineering) to estimate (regress) a five-day forward return as the label (or target) variable.
To verify the model's performance, we use the cross-validation technique. This involves separating the data into training and testing samples. The test data is then evaluated using models that were estimated on the training data. In the context of stock prices, cross-validation must adhere to the time-series structure of the data. The test period must always follow the training data, avoiding random sampling. In figure 58, is an illustration of cross-validation over the course of one year.
Predictions for each three-month trading interval are generated using models calibrated on data from the previous five years. The hyperparameters tuned for these models include variations in feature fraction and pruning-related parameters, both aimed at preventing overfitting. Finally, the best models from all combinations are selected to generate predictions for the three-month test data. For portfolio selection and backtesting, we established a three-year trading (test) period that consists of estimates for each three-month interval by repeatedly applying this cross-validation technique sequentially for 12 times (quarters).

Model Evaluation

Model evaluation is a process that assesses how well a model performs on unseen data. Effective evaluation ensures that the model generalizes well to new, unseen sample data rather than just fitting the training data (a phenomenon known as overfitting). The aim is to estimate how accurately the model will perform in real-world scenarios.
Depending on the type of machine learning task (classification, regression, etc.), various metrics can be used to evaluate model performance. For the portfolio modeling it often use IC (Information Coefficient)[13] and the python library Alphalens.
The Information Coefficient reflects how well the model's predictions align with actual outcomes. While the IC, presented in figure 59, varies across trading periods, it generally remains positive.
Alphalens, a Python library previously mentioned in the context of feature evaluation, can also be employed for model evaluation. It is particularly useful for assessing the predictive power of the model. Specifically, we utilize the "Mean Period Wise Returns by Factors Quantile Analysis" provided by the library, which examines the average returns across different periods for assets grouped by their model prediction scores. This analysis aids in assessing the effectiveness of the model in predicting returns.
The following analysis figure 60 reports the average return per quantile for the S&P 500 constituents, as determined by the rankings derived from the model. If higher quantiles are consistently associated with higher average returns, this suggests that the model in question may possess predictive power.
Alphalens factor decomposition shows the mean return for 1, 2 and 5 day holding period return of groups of stocks ag-

[10] The explanation of the metrics used for model evaluations is outside the scope of this article. However, a brief description has been reported in the appendix .

[11] A detailed explanation of the model Light GBM used in this case study is outside the scope of this article. However, a brief discussion of the model, including its calibration and its pros and cons, is provided in the appendix .

[12] The model is applied exclusively to the 500 constituents of the S&P 500 index; however, it can be extended to include additional stocks if "off-benchmark" allocations are permitted by the investment mandate.

[13] Information coefficient quantifies the relationship between predicted returns of the model and actual returns, helping investors and asset managers understand how effectively their models capitalize on expected returns. It is the Spearman correlation coefficient between the predicted returns generated by an investment model and the actual observed returns of the assets over a specified period. It reflects how well the model's predictions align with the real outcomes.

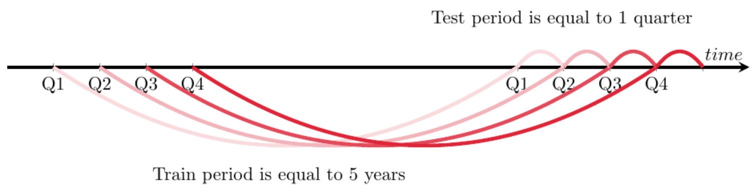

FIGURE 58: Time Aware Cross validation.

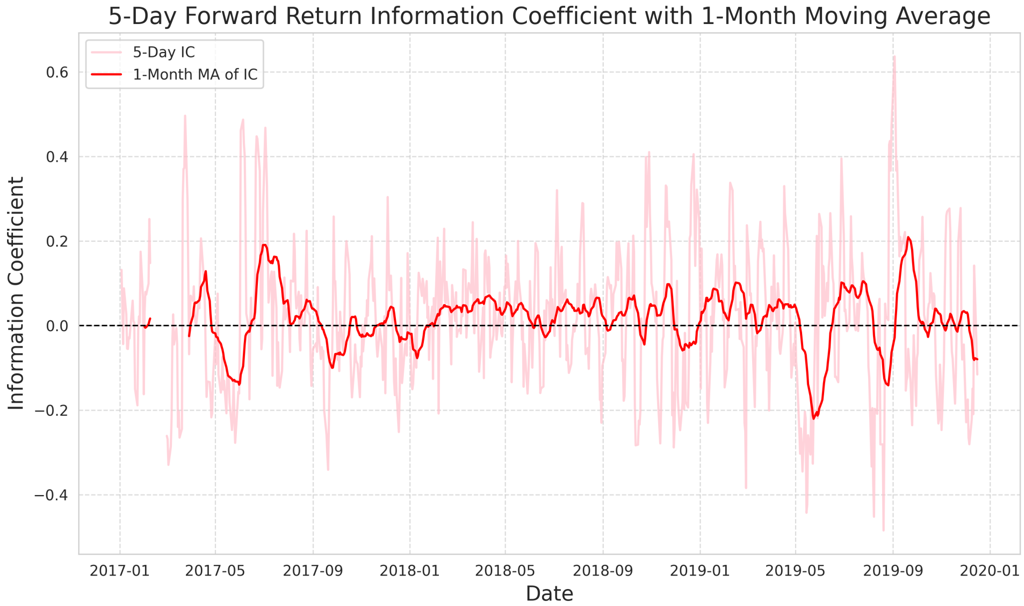

FIGURE 59: Rolling Information Coefficient.

gregated by model score quintile. The results indicate that the model exhibits potential predictive power for the top 10 quintile (the top 50 performing stocks according to the model), achieving an average return of nearly 9 basis points over a five-day holding period. Interestingly, although the model was developed for clients with a long-only mandate, it performs well at identifying underperforming stocks. Specifically, the first quintile (the bottom 50 stocks based on the model ranking) shows the poorest returns, with an average negative return of nearly -7 basis points for a five-day holding period. This suggests that the model can also be also effectively utilized for long/short strategies[14].

[14]A long/short portfolio using this model has not been developed and tested, as it was outside the client mandate. However, it could be easily implemented and backtested.

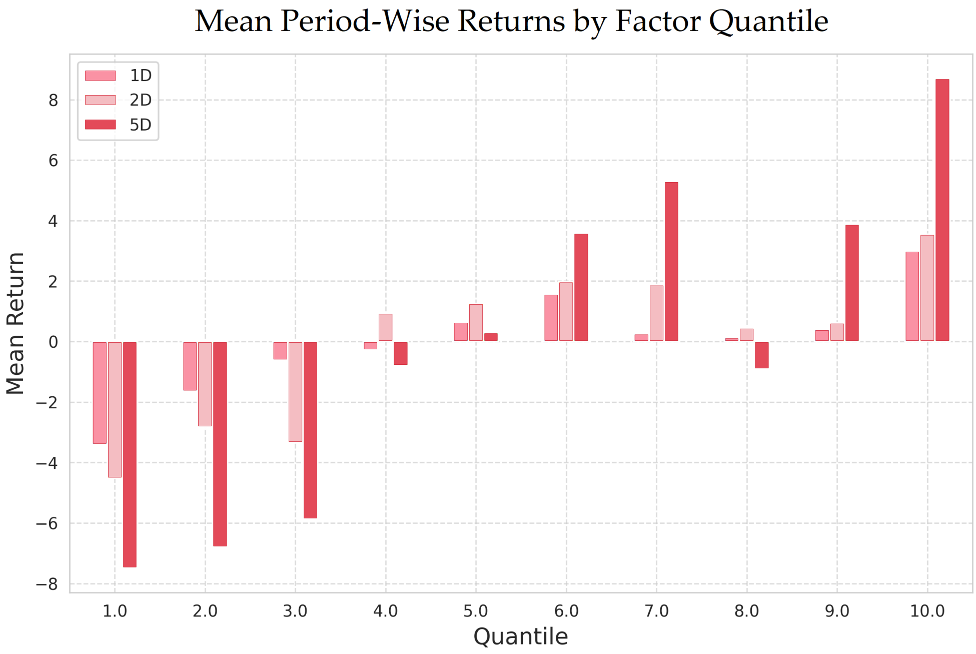

FIGURE 60: Alphalens of the Model.

Model Backtesting and Portfolio Implementation

Backtesting portfolio construction serves as the final attempt to replicate real market and trading conditions and verify the performance. The model backtest is conducted using Python libraries[15] that specialize in simulating realistic market conditions, the impact of market dynamics on trade executions, and the effects of trading commissions. We present four portfolio specifications and their performance based on the predictive power of signals from the model:

- **i.** Long portfolio (Original Long) - this includes outlined backtest performance based on a portfolio with only long positions.
- **ii.** Long Strategy with Basic Risk Management - mirror of point (i) but incorporates simple stop-loss techniques to mitigate risk.
- **iii.** Active Asset Allocation - model is used as an active portfolio allocation tool within a passive ETF strategy. This approach combines a percentage allocation to an S&P 500 ETF (SPY)[16] with the allocations from portfolio (ii). The allocation is dynamic and is a function of the market volatility forecast.
- **iv.** Passive Asset Allocation - this is a passive asset allocation based on a fixed percent allocation to ETF (SPY) and on the Model portfolio developed in point (ii).

The period selected for the backtest covers three years, from 2017 to 2019. This timeframe is deemed interesting as it reflects an overall upward trend, typical of the stock market, while also featuring notable drawdowns and increases in return volatility.
Overall, from the first to the fourth portfolio, there is a trend of reduced absolute performance, accompanied by decreased overall absolute and relative risk, while all four portfolios outperform the benchmark. The core portfolio settings, common to all the presented models, are as follows:

- **20 Long Positions:** Derived from the best predictions of the models.
- **Weekly Rebalancing:** New positions are added to the portfolio and lower-ranked stocks are removed.
- **Trading Commission:** A fee of 2 basis points is applied to the transaction value.

Below we present a summary of the risk adjusted performance on each portfolio implementation. Please refer to the Annex for a definition of the performance and risk metrics calculated.

Long Only Original Strategy

The following we report the performance of a portfolio construction that takes direct input from the model signals, without any correction from risk management techniques. This portfolio is long only.
Figure 61 and table 11 are the risk-adjusted performance indicators and accompanying graph.
The portfolio demonstrates absolute over-performance with a cumulative return of over 90% and outperforms the benchmark by approximately 40 percentage points at the end of the backtest period. Even the Sharpe Ratio is superior for the portfolio model, which implies a better performance over the benchmark per unit of risk. However, it requires

[15] Zipline, Backtrader
[16] SPY is an ETF managed by State Street Global Advisors. The fund's portfolio consists of a basket of stocks that are designed to replicate the performance of the S&P 500 Index.
[17] The Tracking Error is quite high, although it is largely influenced by the strong overperformance over the benchmark in the second part of the test period.

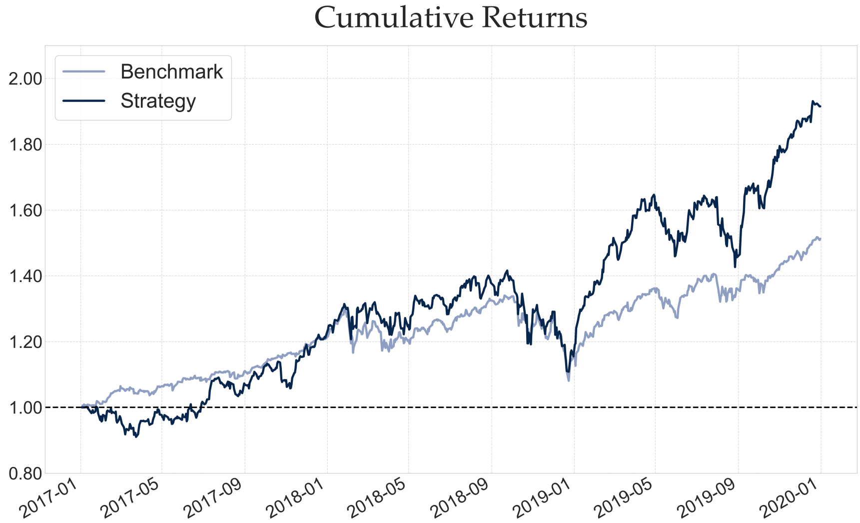

FIGURE 61: *Cumulative Returns: Original Long Model.*

Metric	Original Model	Benchmark
Annual return	24.25%	14.83%
Cumulative returns	91.50%	51.25%
Annual volatility	19.06%	12.82%
Average Tracking Error Volatility	17.36%	-
Average Beta	0.73	-
Sharpe Ratio	1.23	1.14
Treynor Ratio	0.30	0.13
Calmar Ratio	1.12	0.77
Information Ratio	0.54	-
Information Ratio (Sortino method)	0.88	-
Stability	0.91	0.87
Max drawdown	-21.72%	-19.35%
Omega ratio	1.24	1.24
Sortino ratio	1.81	1.58
Skew	-0.23	-
Kurtosis	3.17	-
Tail ratio	1.12	0.86
Daily value at risk	-2.31%	-1.56%

TABLE 11: *Indicators of the Original Long Model.*

adjustments, as the risk is deemed unacceptable for the institutional client. Specifically, the VaR is above the 2% daily limit, here is a significant maximum drawdown of 22%, and a high tracking error volatility of 17% compared to a limit of 8%[17].

Long Strategy with Basic Risk-Management

This portfolio is a modification of original strategy from previous point. Here we apply disposal of assets that have significantly fallen in one day and are likely to cause losses further. The stop loss scheme is implemented at both single stock and portfolio levels. The first deletes fallen positions individually while portfolio level stoploss implies closing of all positions till the next trade iteration. For explanation and demonstration purposes, we chose the following settings for the stop loss: 15% at the single stock level and 3% at the portfolio level.
Figure 62 and table 12 are the risk-adjusted performance indicators and accompanying graph.

The application of the presented framework delivers an improved risk adjusted performance. Specifically, it achieves a cumulative return of over 83%, outperforming the benchmark by approximately 32 percentage points. The portfolio has a lower systematic risk (beta), the Sharpe Ratio is higher, indicating better performance compared to the benchmark per unit of risk. Moreover, both absolute and relative risk metrics demonstrate improved performance in the model portfolio compared to the benchmark. The average daily Value-at-Risk (VaR) is 1.93%, which is well below the limit of 2%. The maximum drawdown of 18.5% is also lower than that of the benchmark. While the average tracking error volatility is still relatively high at 16%, exceeding the limit of 8%, this metric is significantly influenced by the portfolio's positive outperformance compared to the benchmark.
The two portfolio implementation tests developed and reported above have demonstrated interesting performance; however, it is still not acceptable for an institutional client that manages portfolios for retail investors. The asset manager must comply with a strict risk management framework

Metric	Original Model	Benchmark
Annual return	22.32%	14.83%
Cumulative returns	82.74%	51.25%
Annual volatility	15.99%	12.82%
Average Tracking Error Volatility	16.00%	-
Average Beta	0.54	-
Sharpe Ratio	1.34	1.14
Treynor Ratio	0.37	0.13
Calmar Ratio	1.20	0.77
Information Ratio	0.46	-
Information Ratio (Sortino method)	0.79	-
Stability	0.90	0.87
Max drawdown	-18.65%	-19.35%
Omega ratio	1.26	1.24
Sortino ratio	1.99	1.58
Skew	-0.16	-
Kurtosis	2.04	-
Tail ratio	1.05	0.86
Daily value at risk	-1.93%	-1.56%

TABLE 12: Indicators of the Risk Management Model.

agreed upon with investors. The two main metrics that are typically tracked very rigorously are:

- **Tracking Error Volatility:** This aims at maintaining the portfolio's performance alignment with the client-chosen benchmark under all circumstances.
- **VaR:** This focuses on avoiding or limiting excessive absolute losses for the client, regardless of the benchmark[18].

The last portfolio implementation process (ii) presented is compliant with the second metric (VaR) but has significant drawbacks regarding Tracking Error Volatility. For this reason, two additional implementations have been developed with the specific goal to improve the process and better support the mandates and requests of retail clients.

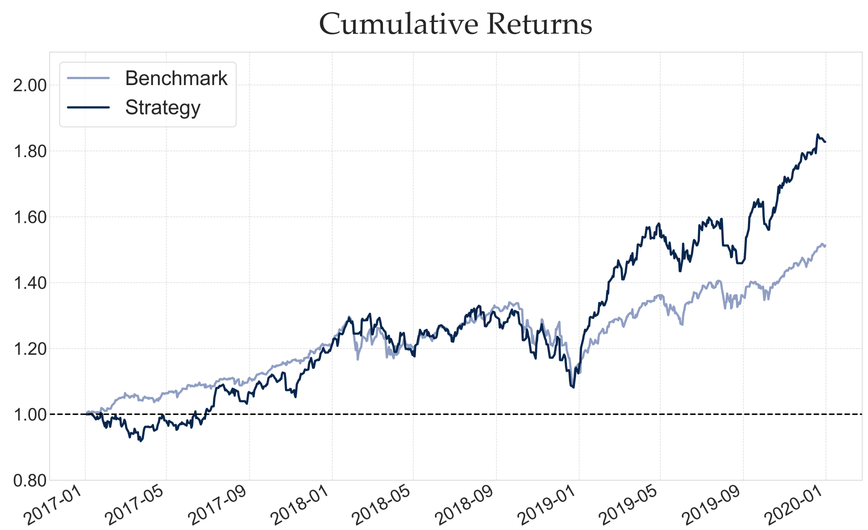

FIGURE 62: Cumulative Returns: Risk Management Model.

[18]Under certain market conditions, these two metrics may conflict when limits are established, as it may not always be feasible to adhere to both simultaneously. For investors capital protection is often given a higher priority limits on Value at Risk (VaR).

Metric	Original Model	Benchmark
Annual return	22.56%	14.83%
Cumulative returns	83.80%	51.25%
Annual volatility	14.83%	12.82%
Average Tracking Error Volatility	13.32%	-
Average Beta	0.65	-
Sharpe Ratio	1.45	1.14
Treynor Ratio	0.31	0.13
Calmar Ratio	1.32	0.77
Information Ratio	0.58	-
Information Ratio (Sortino method)	0.95	-
Stability	0.93	0.87
Max drawdown	-17.12%	-19.35%
Omega ratio	1.28	1.24
Sortino ratio	2.10	1.58
Skew	-0.39	-
Kurtosis	2.52	-
Tail ratio	1.01	0.86
Daily value at risk	-1.78%	-1.56%

TABLE 13: *Indicators of the Active Asset Allocation Model.*

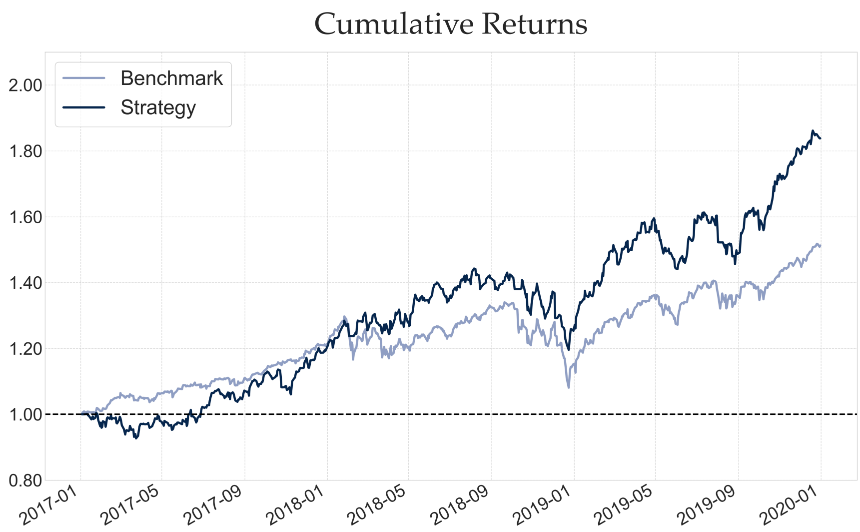

FIGURE 63: *Cumulative Returns: Active Asset Allocation Model.*

Active Asset Allocation Management

This implementation combines a passive allocation to the ETF SPY[19] with an active allocation to the portfolio developed in point (ii). The allocation percentage is dynamically managed by signals coming from the volatility forecast on the VIX index[20].
We have empirically observed that the model tends to overreact to subtle changes in market regimes, particularly when the overall volatility of the benchmark increases, resulting in wide drawdowns in price trends. To mitigate this negative behavior, we implement a portfolio construction strategy that allocates a portion of the portfolio value to the SPY ETF if forecasted market volatility is high. This approach allows the portfolio dynamics to partially replicate those of the benchmark during periods of high market volatility. Meanwhile, during trending markets, the model is utilized with full stock allocation.
As a reallocation trigger, the framework incorporates a volatility forecast obtained using the Asymmetric GARCH model applied to the percentage change on VIX Index time series data. In summary, the settings for this specification are as follows[21]:

- The model includes stop-loss settings from the pre-

[19] The SPDR S&P 500 ETF Trust is an exchange-traded fund which trades on the NYSE Arca under the symbol SPY. The ETF is designed to track the S&P 500 index by holding a portfolio comprising all 500 companies on the index.
[20] The VIX index, often referred to as the "fear gauge," is a popular measure of the stock market's expectation of volatility, derived from the prices of S&P 500 index options. It reflects the market's expectation of volatility over the next 30 days.
[21] The asset allocation to the passive ETF SPTY is intentionally suboptimal and unoptimized. The approach aim to explain the process, rather than on achieving the best risk-adjusted performance.

vious point (ii).

- If expected volatility is moderate, 30% of the portfolio value is allocated to the SPY ETF.
- If expected volatility is high, 45% of the portfolio value is allocated to the SPY ETF.

The following figure 63 and table 13 is the performance of the portfolio construction described above.
The resulting portfolio has demonstrated outperformance of the benchmark throughout most of the considered period, achieving an overall cumulative return of 84%. Additionally, it has significantly enhanced risk-adjusted metrics: a Sharpe Ratio of 1.45 Information Ratio at 0.58 (higher than the previous 2 portfolio constructions) and a maximum drawdown smoothed to -17.2%, which is even lower than that of the benchmark, along with a low daily VaR of -1.78%. The tracking error volatility decreased to 13%. Despite the model being more aligned with the benchmark, this value still exceeds the 8% limit[22]. The Tracking Error Volatility is largely influenced by the strong overperformance over the benchmark in the second part of the test period. On this regard note that Information Ratio (Sortino implementation) is higher, as it uses the downside Tracking Error Volatility This specification offers good absolute and risk-adjusted performance, making it highly suitable for maximizing returns while maintaining a lower level of risk. However, it remains somewhat unsuitable from a client's perspective. Aligning the strategy with the benchmark and adjusting the tracking error volatility requires an additional step.

[22]The Tracking Error is further improved although still high. However as explained for the previous 2 portfolio implementation is largely influenced by the strong overperformance over the benchmark in the second part of the test period. On this regard note that Information Ratio (Sortino implementation) is higher, as it use the downside Tracking Error Volatility.

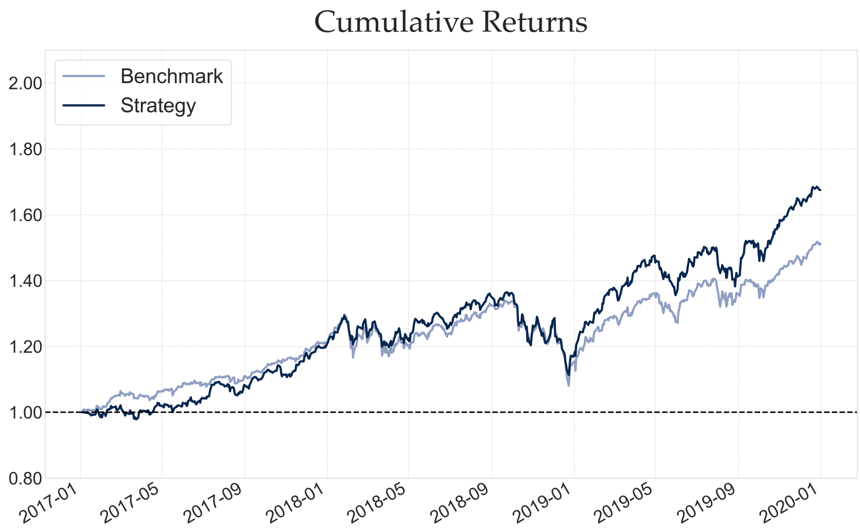

FIGURE 64: Cumulative Returns: Passive Asset Allocation Model.

Metric	Original Model	Benchmark
Annual return	18.82%	14.83%
Cumulative returns	67.50%	51.25%
Annual volatility	12.25%	12.82%
Average Tracking Error Volatility	7.98%	-
Average Beta	0.77	-
Sharpe Ratio	1.47	1.14
Treynor Ratio	0.21	0.13
Calmar Ratio	1.02	0.77
Information Ratio	0.50	-
Information Ratio (Sortino method)	0.84	-
Stability	0.91	0.87
Max drawdown	-18.42%	-19.35%
Omega ratio	1.29	1.24
Sortino ratio	2.13	1.58
Skew	-0.41	-
Kurtosis	3.15	-
Tail ratio	1.02	0.86
Daily value at risk	-1.47%	-1.56%

TABLE 14: Indicators of the Passive Asset Allocation Model.

Passive Asset Allocation Management

The final specification is designed to address the issue with tracking error volatility. Instead of making adjustments only during market downturns, we incorporate a constant allocation to the SPY ETF. This significantly impacts the tracking error, as the constant allocation reduces the mismatch even during trending markets. Here we report an example that invest 50% of the portfolio value in the SPY ETF. Thus, the portfolio utilizes the core settings and risk management tools from Portfolio (ii), while adding a constant 50% allocation to the SPY ETF.
The following figure 64 and table 14 is the performance of the portfolio construction described above.
Overall, the resulting portfolio maintains outperformance while reducing tracking error volatility to approximately the required level of 8%. Although cumulative returns are limited to around 68%, the overperformance over the benchmark remains at a more than acceptable level of 17 percentage points. The Sharpe Ratio of this specification is the highest, the drawdown is less than the benchmark, the overall absolute volatility is less than the benchmark and the daily VaR is at the lowest level at 1.47% particularly in comparison to the benchmark at 1.56%. The 3-rd portfolio remain, however, superior in terms of performance adjusted by relative risk (the Information ratio is superior) and performance adjusted by systematic risk (Treynor Ratio).

Summary

We have presented a case study in which a systematic investment process is executed based on signals derived from a machine learning model. This case focuses on describing the process, highlighting the potential use of such a framework

for both direct investment purposes and as a benchmark for internal performance evaluations of "traditional" management portfolios. The latter idea, born from a client's need to have an economical and easily implementable tool that can be effectively used to reward discretionary investment teams by comparing them with the performances of systematic portfolios. The modeling framework use an integrated approach where Unsupervised Learning techniques are integrated in a Supervised Learning Regression algorithms. The end-to-end process for the development of the model and the backtest has been presented.
The results are encouraging; four levels of portfolio implementations based on investment signals from the developed machine learning model have been presented. All portfolios exhibit risk-adjusted performances superior to those of the benchmark agreed upon in the investment mandate. The different implementations are characterized by varying levels of compliance with the investment risk framework agreed upon with the final investor, with only the last of the four achieving full compliance with all investment limits while maintaining adequate outperformance.
The best portfolio construction is the 3rd (Active Asset Allocation Management) which achieve an overperformance over 3 years of 33% points, a Sharpe Ratio of 1.45, An Information Ratio of 0.58, a Treynor ratio of 0.31 and a daily VaR of 1.78%. The one that fully comply with client Risk Management mandate is the 4-th (Passive Asset Allocation Management) which achieve an overperformance over 3 years of 17% points, a Sharpe Ratio of 1.47 (highest of all the portfolio implementations).
The case study conveys recent field-acquired experiences in a divulgative, informative, and simplified manner. Research is underway to further develop the model to:

- Improve risk-adjusted performance by testing and implementing new features, also using the integration of exogenous data beyond price and trading volume data.
- Enhance risk-adjusted performance by experimenting with new machine learning models or new hyperparameter calibration settings.
- Improve risk-adjusted performance by employing different investment frameworks, with varying holding periods and training and testing periods of the model.

References

[1] **Ankerst, M., Breunig, M. M., Kriegel, H. P. and Sander J.** *Optics: ordering points to identify the clustering structure*. ACM Sigmod Rec, June 1999.

[2] **Avellaneda, M. and Lee, J.** *Statistical Arbitrage in the U.S. Equities Market*. Quantitative Finance, August 2008.

[3] **Ester, M., Kriegel, H. P., Sander, J. and Xu, X.** *A density-based algorithm for discovering clusters in large spatial databases with noise*. Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, June 1998.

[4] **Fadda, S.** *Comparing symmetric and asymmetric volatility estimates for S&P index prices*. Periodicals of Engineering and Natural Sciences, February 2019.

[5] **Jansen, S.** *Machine Learning for Algorithmic Trading*. Second Edition, Packt Publishing, 2020.

[6] **Jolliffe, I.** *Principal component analysis.* Springer, January 2014.

[7] **Krauss, C., Do, X. A. and Huck, N.** *Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500*. European Journal of Operational Research, 2016.

[8] **Lundberg, S. M. and Lee, S.** *A Unified Approach to Interpreting Model Predictions*. NIPS, 2017.

[9] **Lustig, H., Roussanov, N. and Verdelhan, A.***Common Risk Factors in Currency Markets*. The Review of Financial Studies, June 2008.

[10] **Prasan, N. H.** *OPTICS: Hierarchical Density-Based Clustering*.

Technical Appendix

Unsupervised Learning Application

As described in the article, we have applied an integrated approach by incorporating features from unsupervised learning into supervised learning models. We have generated two types of features: one related to Risk factors and the second related to Clustering. Both features are generated from a reduced form of the data representation produced via PCA. The steps are as follows:

1. Reduce the dimensionality of the data at hand via PCA;
2. From the reduced dataset, generate clusters;
3. From the reduced dataset, generate a Principal Components Risk Factors Time Series.

1. Reduce the dimensionality of the data at hand via PCA
PCA is a statistical procedure that uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of linearly uncorrelated variables, the principal components. The transformation is defined such that the first principal component accounts for as much of the variability in the data as possible. Each succeeding component, in turn, has the highest variance possible under the constraint that it must be orthogonal to the preceding components.
The application of PCA is as follows: First, the return series for security (i) at time (t), denoted as $R_{i,t}$, is derived from the price series P_i of the security.

$$R_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}}. \tag{45}$$

Next, the return series needs to be normalized because PCA is sensitive to the relative scaling of the original variables. This normalization involves subtracting the mean, R_i, and then dividing by the standard deviation, σ_i. Once the return series of all assets is normalized, the correlation matrix is computed.

$$Y_i = \frac{R_i - \bar{R}_i}{\sigma_i}, \quad \rho_{i,j} = \frac{1}{T-1}\sum_{t=T}^{M} Y_{i,t}Y_{j,t}. \tag{46}$$

The next step involves extracting the eigenvectors and eigenvalues from the covariance matrix to create the principal components. The eigenvectors indicate the directions of maximum variance, while the eigenvalues represent the amount of variance in those corresponding directions. Both eigenvalues and eigenvectors can be calculated using singular value decomposition (SVD) on the stacked matrix of normalized returns (A).

$$A = USV^T. \tag{47}$$

Here U is an orthogonal matrix of singular vectors, S is the matrix of singular values, and V is the transposed matrix of singular vectors. The result of multiplying A^T by A is:

$$A^T A = VS^2V^T. \tag{48}$$

Here, $A^T A$ forms the correlation matrix, and the matrix V now represents the eigenvectors of $A^T A$. To reduce dimensionality, we select the k eigenvectors that correspond to the k directions of maximum variance, where k indicates the number of features we wish to use to describe the transformed data. The inclusion of more eigenvectors allows for a more accurate representation of the data.
We apply the principal component analysis to the returns time series of the S&P index constituents with a training period of five years. As a result, we have reduced the dataset from 1260 (dates) x 500 (tickers) to 1260 (dates) x 15 (principal components). The example for the starting date of the training period is in the table 15. This resulting reduced dataset then serves as the foundation for Clustering and Risk Factor Decomposition analyses.

	PC_1	PC_2	…	PC_15
03.01.2012	-18.10	13.85	…	-0.02
04.01.2012	3.24	7.48	…	-0.22
…	…	…	…	…
03.01.2017	-11.86	4.37	…	-1.69

TABLE 15: Example of reduced dataset for factors.

	PC_1	PC_2	…	PC_15
Ticker 1	0.57	-0.38	…	-0.21
Ticker 2	1.72	0.71	…	-3.94
…	…	…	…	…
Ticker 500	0.37	0.24	…	-0.34

TABLE 16: Example of reduced dataset for clustering.

2. Generate Clusters
To conduct the clustering analysis, we need to transform the reduced dataset from stock-specific data to data specific to individual tickers. We utilize the eigenvectors from the previously conducted PCA to obtain the loadings (unnormalized weights) for each ticker in relation to the principal components. This can be interpreted as risk premium since the principal components reflect the underlying risks. In summary, the input dataset for clustering appears as in table 16.
In this manner, we have created a reduced dataset that captures the main patterns in the variance of stocks within the S&P Index. Clustering enables us to group stocks that exhibit a common variance structure over time, thereby highlighting the heterogeneity among the stocks. Specifically, we utilize the OPTICS algorithm, which is an extension of the DBSCAN algorithm.
Density-based clustering algorithms offer several advantages in this context. First, they allow for clusters with arbitrary shapes, eliminating the need for Gaussianity assumptions about the data's shape. Second, these algorithms are naturally robust to outliers since they do not require every point in the dataset to be grouped. Lastly, they do not necessitate prior specification of the number of clusters. Therefore, we will explore how density-based clustering can be applied in this analysis.
A density-based clustering algorithm views the clustering space as an open set in Euclidean space that can be divided into a collection of connected components. The DBSCAN (Density-Based Spatial Clustering of Applications with Noise) algorithm (Ester et al., 1996) is the most prominent algorithm in this category. According to the authors, clusters of points can be readily detected because the density of points within each cluster is significantly higher than that outside the cluster.
To formalize the notion of clusters being dependent on density, the key concept relies on understanding density and connectivity. Definitions 3.1 to 3.5 pertain to the DBSCAN algorithm and include two user-defined parameters: (ε) and (minPts).
Definition 3.1: (ε)-neighborhood: The (ε)-neighborhood of a point (q) is defined as $N\varepsilon(q) = p \in X | d(q, p) \leq \varepsilon$, where $(d(q, p))$ represents the distance between points (q) and (p), and X is the set of all points.
Definition 3.2: Core Point: A point (q) is considered a core point if it satisfies $|N\varepsilon(q)| \geq minPts$, where $|N\varepsilon(q)|$ denotes the number of points within the ε-neighborhood of (q). It is important to note that $minPts$ includes the point itself.
Definition 3.3: Directly Density-Reachable: A point (p) is directly density-reachable from point (q) within a set of points X if $p \in N\varepsilon(q)$ and (q) is a core point.
Definition 3.4: Density-Reachable: A point (p) is density-reachable from point (q) if there exists a chain of objects $(p_1, ..., p_n)$ such that $(p_1 = q)$ and $(p_n = p)$, with (p_{i+1}) being directly density-reachable from (p_i).
Definition 3.5: Density-Connected: A point (p) is density-connected to point (q) if both points are density-reachable from a common core point.
Figure 65 illustrates the concepts just defined, with the parameter $minPts$ set to 5. Figure 65 a) depicts a core point. Point (q) is identified as a core point because it contains at least $minPts$ within its ε-neighborhood, which is represented by a circle with a radius of ε. Point (p_1) is part of the ε-neighborhood of (q). Similarly, a point that is not a core point but still falls within the ε-neighborhood is known as a border point. In contrast, a point like (p_2), which does not belong to any neighborhood, is classified as an outlier.
Figure 65 b) illustrates the concepts of directly density-reachable and density-reachable points. Point (p_1) is directly density-reachable from (q). Furthermore, point (p_2) is density-reachable from (q), as both (q) and (p_1) are core points. However, (q) is not density-reachable from (p_2) because (p_2) is not a core point.
Finally, 65 c) demonstrates the idea of density connection. Points (p_2) and (p_4) are density-connected since both points are density-reachable from (q).
These definitions are established for the DBSCAN algorithm, which requires ε to be a fixed predefined parameter. The OPTICS extension introduces additional calculations:

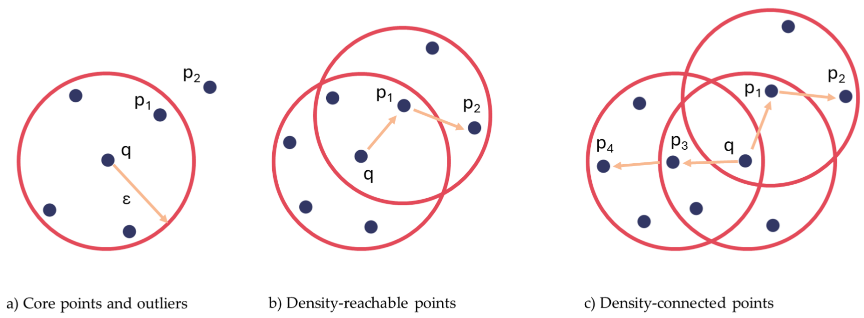

FIGURE 65: *Main Concepts of Density-Based Clustering.*

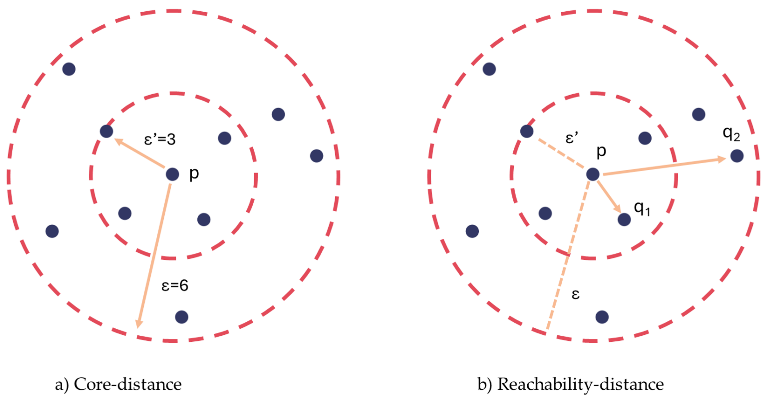

FIGURE 66: *Core of OPTICS algorithm.*

Definition 3.6: Core-Distance: Let $(minPts)$-distance(p) be the distance from a point (p) to its $minPts$ neighbors. Thus, the core-distance of (p) is defined as:

$$\text{core\_dist}_{\varepsilon,min\_pts}(p) = \begin{cases} \text{Undefined, if } |N_\varepsilon(q) < \text{min\_pts}| \\ \text{min\_pts-distance, otherwise} \end{cases}. \tag{49}$$

The core-distance of a point (p) is, in simple terms, the smallest distance ε between (p) and any point in its ε-neighborhood such that (p) is considered a core point with respect to (ε). If (p) is not a core point to begin with, then the core-distance is undefined.

3.7: Reachability-Distance: The reachability-distance of point (p) with respect to point (o) is defined as follows:

$$\text{reach\_dist}_{\varepsilon,min\_pts}(p,o) = \begin{cases} \text{Undefined, if } |N_\varepsilon(o) < \text{min\_pts}| \\ \text{max((core-distance(o)), (distance(o,p))), otherwise} \end{cases}. \tag{50}$$

The reachability-distance of a point (p) with respect to a point (o) can be understood as the minimum distance required for (p) to be directly density-reachable from (o). This condition necessitates that (o) is a core point; therefore, the reachability-distance cannot be less than the core-distance. If it were, (o) would not qualify as a core point.
Figure 67 illustrates the concepts of core-distance and reachability-distance as described. The scenario depicted defines $(minPts = 5)$ and $(\varepsilon = 6)$ mm. In the left image, the core-distance (represented as ε) indicates the minimum distance at which (p) qualifies as a core point. For a radius of $(\varepsilon' < \varepsilon)$, there are already five points $(minPts)$ within the circle. The right image illustrates the concept of reachability-distance.
The implementation of OPTICS is fundamentally similar to that of DBSCAN; however, instead of keeping a set of known but unprocessed cluster members, OPTICS utilizes a priority queue.
The algorithm outputs the points along with their corresponding reachability-distances. These points are ordered in such a manner that spatially closest points are neighbors in this ordering. This information allows for the construction of a reachability plot by arranging the ordered points. Similar to a dendrogram, the reachability plot represents the hierarchical structure of clusters. Because points within the same cluster are closer to one another, they exhibit a low reachability-distance to their nearest neighbor.
In summary, OPTICS enhances the foundational principles of DBSCAN by addressing one of its key limitations: the ability to utilize a range of epsilon ε values to identify clusters with varying densities. Unlike DBSCAN, which relies on a fixed ε, OPTICS investigates multiple ε radius values. This capability allows for the identification of clusters that exhibit diverse shapes, sizes, and densities based on the reachability graph. This flexibility positions OPTICS as a robust tool for analyzing

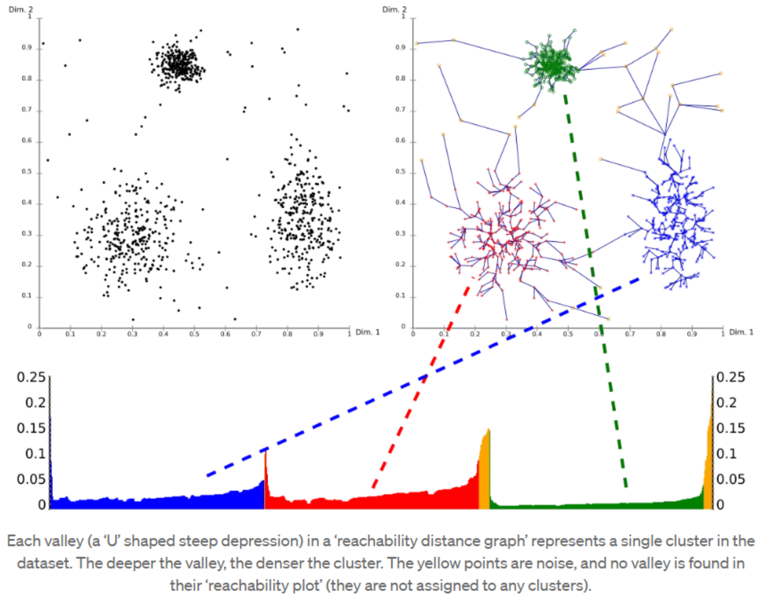

FIGURE 67: Reachability Plot [10].

datasets with complex structures.

3. Principal Components Risk Factors Time Series
Factor analysis in portfolio risk is a statistical technique employed to identify and understand the underlying factors influencing the returns and risks of a portfolio. By consolidating the variability of asset returns into a smaller number of common factors, investors can gain insights into the sources of risk and performance within their portfolios.
One of the most established applications of this approach is the Fama-French factors, which can be interpreted as portfolios based on specific characteristics of firms, such as size, growth, and robustness. However, these direct data points are primarily used for academic purposes rather than for online trading algorithms, and generally, factors require additional data for implementation. Instead, it is possible to derive factors directly from the available dataset by using PCA.
To extract the common underlying risk factors, we apply PCA to the return series, specifically utilizing the reduced dataset generated in point 1. It is particularly noteworthy that each principal component (PC) of the PCA can be viewed as representing a distinct risk factor. These components can serve as:

- **Latent Factors:** The principal components can represent underlying factors that drive asset returns based on patterns identified in multidimensional data. For example, in a dataset containing various stock returns, PCA might unveil a component that captures overall market movements, while other components may highlight sector-specific trends or risk exposures (Avellaneda & Lee [2]).
- **Risk Factors:** Principal components derived from financial metrics can encompass common risks affecting multiple assets. For instance, one component may capture market volatility, while another may reflect sensitivity to interest rates (Lustig et al., [9]).

The calculation of these factors is dynamic to capture the latest market information and avoid lookahead bias-the usage of information that is not available for the specified day. In the context of a machine learning model for predicting returns, the PCA time series risk factors are utilized to generate betas or sensitivities to these factors for each stock. To achieve this, we apply rolling regression with a one-year window. The regression model for one iteration can be represented as follows:

$$R_i = \alpha + \sum_{k=1}^{15} \beta_k PC_k + \epsilon_i. \tag{51}$$

Here, we measure the sensitivity for each individual stock, where R_i represents the returns for a specific stock and PC_k denotes the time series of the three risk factors. In summary, we obtain 16 variables: one constant (also referred to as Alpha) and 15 coefficients (Betas) that capture the effects of changes in the underlying factors.

Evaluation Metrics

i. Information Coefficient
The first metric is Information coefficient that measures correlation between two variables. In the scope of the feature evaluation we measure the correlation between target variable and explanatory variable. In the scope of model evaluation we measure the correlation between actual and predicted returns. The correlation coefficient can be based on Pearson or Spearman correlation.
Pearson correlation coefficient is based on the covariance between features and then apply normalization with standard deviations. It calculates linear relationships:

$$\rho = \frac{\text{Cov}(X,Y)}{\sigma_X \sigma_Y}. \tag{52}$$

The second option is Spearman's rank correlation. Instead of comparing absolute values, it ranks the variables (by sorting them by value and assigning positions) and then applies covariance to these ranks. This method assesses the monotonicity of the relationship between variables, allowing for the detection of non-linear relationships:

$$r_s = 1 - \frac{6\sum_{i=1}^{n} d_i^2}{n(n^2-1)}. \tag{53}$$

Where $d_i = R[X_i] - R[Y_i]$, difference of rankings.
We use Information Coefficient in feature evaluation and model evaluation, Spearman rank correlation is applied. Particularly, it can be considered as more suitable for non-linear financial data.

ii. Mutual Information Regression

This metric is based on the entropy framework. It measures the reduction in uncertainty for one variable given the value of the other variable. Primarily, this tool is used to assess the preliminary importance of the feature. In this case, mutual information is calculated for target and explanatory variables.
Usually, this metric is univariate: We measure the associations between the target variable and one explanatory variable. However, it can be extended to calculate multivariate mutual information.
Conceptually, coefficient of mutual information can be defined as follows:

$$I(X;Y) = H(X) - H(X|Y). \tag{54}$$

Where H(x) - entropy for X and H(X | Y) - conditional entropy for X given Y.
Since the entropy for a continuous variable is defined as:

$$H(X) = -\int p(x)\log p(x)\,dx, \tag{55}$$

and conditional entropy is defined as:

$$H(X,Y) = -\int\int p(x,y)\log p(x,y)\,dx\,dy. \tag{56}$$

The mutual information coefficient can be rewritten with Kullback-Leibler divergence between the joint probability distribution and the product of the marginal probabilities for each variable. Here we present the process for continuous variables. Process for discrete variables is almost the same besides usage of discrete summation instead of integrals.

$$I(X;Y) = D_{KL}\left(P_{(X,Y)}||P_X \otimes P_Y\right). \tag{57}$$

Where $P_{(X,Y)}$ is a joint distribution and the marginal distributions are P_X and P_Y.
So, process of calculation is defined as follows:

$$I(X,Y) = \int_Y\int_X P_{(X,Y)}(x,y)\log\left(\frac{P_{(X,Y)}(x,y)}{P_X(x)P_Y(y)}\right)dx\,dy, \text{ for continuous variables.} \tag{58}$$

If the variables are completely independent, the mutual information coefficient is equal to 0, which represents the lower bound. The value of the coefficient depends on the base of the logarithm used (typically 2, e, or 10), which reflects the level of dependency between the variables. Selection of the features is based on the comparison of coefficients for all variables. Since the calculation is logarithmic, the increase of coefficient is slow and high values (compared to base of logarithm) are uncommon.

iii. Features Importance

First model-based metric in our case study is the feature importance. This metric is derived from a preliminary gradient boosting model that includes all the variables being considered.. The output of the gradient boosting model consists of a sequence of decision trees, where each decision tree is built by splitting the data based on the threshold of a variable. The variable and split value are selected based on their efficiency, which, in the case of a regression task, is measured by variance reduction.

$$\text{Var}(Y) = \frac{1}{n}\sum_{i=1}^{n}(y_i - \bar{y})^2, \tag{59}$$

$$\Delta\text{Var} = \text{Var}_{\text{parent}} - \left(\frac{n_{\text{left}}}{n_{\text{total}}}\cdot\text{Var}_{\text{left}} + \frac{n_{\text{right}}}{n_{\text{total}}}\cdot\text{Var}_{\text{right}}\right). \tag{60}$$

The best split separates the data in a two subsets that have lowest possible variance within. It is the most important part that is used in feature evaluation. The feature importance algorithm evaluates which variables were selected for the splits and how much gain (variance reduction) each variable contributed to the splits within the tree. So, for variable j and all t nodes in one tree, importance of the feature j is calculated as follows:

$$I_j = \sum_t \Delta\text{Impurity}_{j,t}. \tag{61}$$

Once the importance scores for each feature are calculated, they are averaged across all trees in the model to produce the final feature importance score. The metric is univariate and does not give direct information about interaction effect. However, set of other variable affects construction of the trees and has effect on feature importance. Additionally, Selection of the features

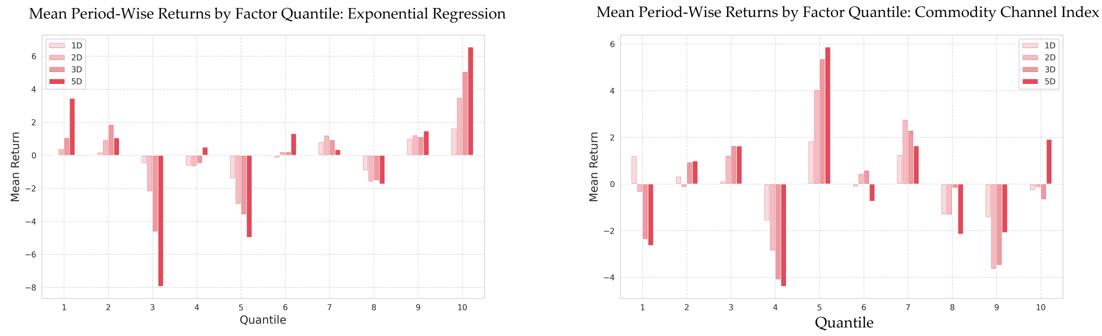

FIGURE 68: Alphalens example of good and relatively worse performance.

for split may be unstable in case of two or more variable are significantly correlated. This issue should be considered during the final feature evaluation. Particularly, different lags of return fall in this issue.

iv. Shapley Additive Explanation
Contrary to the features importances, Shapley Additive Explanation is based on the predictions made by the model rather than evaluation of trees themselves (Lundberg & Lee, [8]).
The concept is based on the framework of cooperative game theory. It evaluates how different input features influence the predictions made by the model. Each feature will receive a score representing its importance. Higher scores indicate that the feature has a greater influence on the model's predictions. Features can be ranked based on their importance scores that can be defined as follows:

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} \left(f_{S \cup \{i\}}(x_{S \cup \{i\}}) - f_S(x_S) \right). \tag{62}$$

This equation represents weighted sum of marginal contributions for each feature. Fraction stands for the combinatorial weight and second part of the equation stands for the marginal contribution. S is a subset of all features without feature i, F - set of all features.
$f_{S \cup \{i\}}(x_{S \cup \{i\}})$ - model prediction with feature
$f_S(x_S)$ - model prediction without feature.
So, marginal contribution of the features is defined as the level of performance improvement in case this particular feature is used for prediction. This scheme is multivariate since marginal contribution for one feature is estimated using large number of subsets with other features.

v. Alphalens Library
The last tool is based on the evaluation of how much returns the feature may generate if it is used as a factor for long-short trading. It is well known Python library developed by Quantopian, specifically designed for performance analysis of features and models. Alphalens provides extended backtest for different quantiles of factors. In particular we use the "Mean Period Wise Returns by Factors Quantile Analysis" provided by the library, which examine the average returns across different periods for assets grouped by their feature/model scores. This helps in assessing the effectiveness of the feature/model in predicting returns. The factor demonstrating consistency of quantile returns and have significant return in the last 10-th quantile an be considered as a good factor for long trading. The following is an example of 2 features, one that was accepted (Exponential Regression betas) and one that was rejected (Commodity Channel Index) based on Alphalens method: figure 68 reports the Mean Period Wise Returns by Factors Quantile Analysis of "Exponential Regression betas" and "Commodity Channel Index" features".

Light Gradient Boosting Model

Introduction
Light Gradient Boosting Machine (LightGBM) is one of the popular algorithms within the Gradient Boosting family (GBMs) of models. This is an Ensemble Method: gradient boosting models combine the predictions of multiple weak learners (usually decision trees) to create a strong predictive model.
Particularly, GBMs are well-suited for capturing complex non-linear relationships between features and the target variable. This flexibility means that they can model intricate relationships without the need for extensive feature engineering. Also, sequential nature of GBMs allows them to naturally capture interactions between features, even if they are not explicitly defined in the dataset.
Additionally, GBMs implement missing value support: they can inherently handle missing values without the need for extensive preprocessing, as tree-based methods can learn how to deal with missing data effectively. Also, they handle multicollinearity well as tree-based models are less affected by correlated input features.
First, we briefly describe the core gradient boosting algorithm and introduce LightGBM, which is used in our case study. Next, we provide details on the practical application of the model.

Algorithm
The idea of the algorithm is that new tree in the gradient boosting algorithm focuses on correcting the errors made by the previous trees, allowing for fine-tuning of predictions and improving the model's performance on the training set and unseen

data.
The process begins with the data in its original form. The gradient boosting algorithm iteratively builds a model by minimizing a specified loss function. For regression task the most common Loss Function specification is Mean Squared Error (MSE) that is used in our case:

$$L = (y_i - \gamma)^2, \text{ where } y_i \text{ - real value and } \gamma \text{ - predicted value.} \tag{63}$$

The process is initialized with the simplest prediction: assigning a single value for all data points.

$$F_0(x) = \text{argmin}_\gamma \sum_{i=1}^{n} L(y_i, \gamma). \tag{64}$$

Here, $F_0(x)$ stands for predictions of the output at first step. Right side specifies the variable γ, that minimizes Loss Function. Here, at the first step, it is simply the average of all the target data points.
Next, the algorithm calculates the residuals between the actual values and the predictions. This marks the beginning of the iterative procedure, where each new model is trained to correct the errors of the previous one.

$$r_{i,m} = -\left[\frac{\delta L(y_i, F(x_i))}{\delta F(x_i)}\right]_{F(x)=F_{m-1}(x)} = 2\left(y - F_{m-1}\right). \tag{65}$$

Here, the calculated residuals correspond to the negative gradient of the loss function. This means that residual values indicate direction and magnitude by which predictions should be adjusted to minimize loss function.
The adjustment of the residuals is determined by a decision tree. At each step, the decision tree aims to assign the residuals to the terminal nodes. The splits in the decision tree are made by selecting the features that best partition the residuals, aiming to minimize the loss at each node. Splits evaluation is described in the Feature Importance annex.
After the separation, algorithm searches for γ (equation 66) the that minimizes aggregating loss in the terminal nodes. With MSE, it is simply average of residuals in the terminal node.
The construction of decision trees is a complex process that includes numerous parameters and settings. For example, these setting may include criteria for the termination of trees building, logic deriving sequence of the splits and many other. Generally, this tools are applied to reduce overfitting of the model: Prediction of the trees should be generalizable to unseen data rather than precisely fit the sample data. Additionally, various advanced techniques have been implemented to reduce computational complexity while preserving predictive power. We will highlight advantages of LightGBM later in this annex.
The output of decision tree is then allocated to predictions for each data sample. The obtained corrections across all terminal nodes $\gamma_{j,m}$ are added to predictions from the previous step:

$$F_m(x) = F_{m-1}(x) + v \sum_{j=1}^{J_m} \gamma_{j,m} * I(x \in R_{jm}). \tag{66}$$

The corrections are scaled by a parameter v - learning rate, which typically takes a relatively small value around 0.01. This technique, called "shrinkage," prevents the model from overfitting the training data too quickly. Instead, it allows the gradient boosting model to generalize observed patterns from the features to unseen data.
This process of calculating residuals and applying decision trees is repeated iteratively until a termination criterion is met. Once training is complete, the data can be run through the model's trees to generate predictions.

Light Gradient Boosting Unique Tools
As mentioned before, many settings were applied to avoid overfitting and decrease computation time. Particularly, these settings affect evaluation of possible splits in a decision tree. It is one of the most hard and time consuming part of the algorithm. We will discuss the unique features of LightGBM, which distinguish it from other popular gradient boosting algorithms like XGBoost and CatBoost.
i. Gradient-Based One Side Sampling
The core idea of this implementation is to focus on the data points that were poorly predicted at the previous step. In this case, the decision tree evaluates splits based on a reduced subset of data points. Formally, algorithm sort residuals and select top k points with highest absolute gradient. From the remaining points algorithm selects reduced random sample.
In summary, this method allows to avoid extensive computation and insert faster corrections to the poorly predicted points.
ii. Exclusive Feature Bandling
The next tool is designed to reduce number of features, that also decrease computational effort. It merges mutually exclusive features into one single. Mutually exclusive features stands for sparse features that rarely take on non-zero values simultaneously.
In this way, evaluation of splits considers less variables without the loss in accuracy. Particularly, it is useful to categorical variables but can include also continuous ones.
iii. Leaf-Wise Construction
Another specific of LightGBM is leaf-wise construction of decision trees. At each step of building the tree, leaf-wise construction focuses on selecting the best possible split based on the leaf nodes of the tree, regardless of the depth already constructed.
In contrast, the depth-wise setting constructs a balanced tree, where the algorithm builds the tree level by level, ensuring that all leaf nodes are expanded uniformly. This means that all branches of the tree are developed to the same depth before the algorithm moves on to the next level.
This implementation makes trees deeper and usually also decrease computations.

Application of the model to the Case Study
Model Calibration in our case study is applied to each 5 years - 3 months fold from time aware cross validation.
In the first step, we specified parameters for constructing the decision trees, including pruning settings and feature fraction.

Metric	Description of the Metric
Annual return	The mean annualized rate of returns over the period.
Cumulative returns	The cumulative return over the entire period
Annual volatility	The annualized standard deviation of daily returns
Tracking Error Volatility	The annualized daily standard deviation of the difference in returns between the strategy and the benchmark (measuring how closely the portfolio tracks the benchmark)
Average Beta	The sensitivity of the portfolio to market risk. It reflects change in the portfolio returns relative to the change in market returns
Sharpe ratio	Risk-Adjusted metric: expected annual return (adjusted for risk-free rate) per unit of total volatility
Treynor ratio	Risk-Adjusted metric: expected annual return (adjusted for risk-free rate) relative to systematic risk exposure (beta). For the benchmark equal to annual rate of return minus risk-free rate
Calmar ratio	Risk-Adjusted metric: expected annual return per unit of maximum drawdown
Information ratio	Portfolio's excess return over the benchmark, adjusted for the volatility of those excess returns.
Information ratio (Sortino method)	Extension of previous metric: Portfolio's excess return over the benchmark, adjusted for the only downside volatility of those excess returns.
Stability	The R-Squared of linear time trend fitted to observed cumulative returns, measuring the consistency of returns over time.
Max drawdown	Maximum decline of portfolio value from cumulative maximum of portfolio value
Omega ratio	Comparison of probability distributions related to gains and losses
Sortino ratio	Risk-Adjusted metric: modified Sharpe Ratio, penalizes only downside volatility
Skew	Asymmetry of returns distribution for portfolio
Kurtosis	Tail estimation of returns distribution for portfolio
Tail ratio	Ratio of high and low quantiles (95th percentile to 5th) of returns distribution
Daily value at risk	The estimated portfolio loss at specified confidence level 5% calculated historically

TABLE 17: Description of Indicators Used in Portfolio Backtest.

Pruning is a technique in decision trees used to reduce model complexity and enhance generalization by removing branches that have little significance to the predictive outcome. This is achieved by elimination of nodes that contribute minimal additional predictive value. Particularly, the most important settings for our context are maximum number of terminal nodes and minimal number of data points that can remain in nodes. Feature fraction parameter stochastically limit the features that can be used for split evaluation. This method is also intends to avoid overfitting and enhance generalization.
We explore different combinations of parameter values, creating an extensive list of possibilities. If generating the full list is computationally intensive, it is possible to reduce the number of combinations through random sampling. For each fold in the cross-validation, we run models using parameter combinations selected from the final list.
At the second step we utilize the results from the models trained on the 5 years data. Each model is evaluated by running the unlabeled test data through it and calculating the Information Coefficient between the real and predicted values. . The final predictions are obtained by averaging the outputs of the top-performing models from the most recent years of the training period.

Portfolio Indicators

Here, in table 17 we provide short description of each indicator used in portfolio backtest.

Asymmetric GARCH Implementation

In this annex, we outline the approach that forms the basis of the portfolio implementation known as "Active Asset Allocation." This method allocates a portion of the portfolio value to the benchmark SPY ETF when forecasted market volatility is high. We identify the VIX as a proxy for the market's expected future volatility, and to forecast changes in this expected volatility, we implement an Asymmetric GARCH model on the percentage change in the daily VIX index.
The classic form of the Generalized Auto-Regressive Conditional Heteroscedasticity (GARCH) model estimates the variance of a time series in an autoregressive manner. Specifically, the variance at time (t) depends on the shock from the previous period as well as the variance estimate from the prior period. The complete equation is expressed as follows:

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2. \tag{67}$$

The model indicates that variance patterns can persist over time rather than dissipate immediately. This characteristic is particularly significant for stock price time series, as volatility frequently demonstrates a clustered structure, leading to periods of relatively high or low volatility.
The Asymmetric GARCH extension introduces an additional term that accounts for whether the deviation in the previous period upside or downside was:

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \gamma\epsilon_{t-1}^2 I_{[\epsilon_{t-1}<0]} + \beta\sigma_{t-1}^2. \tag{68}$$

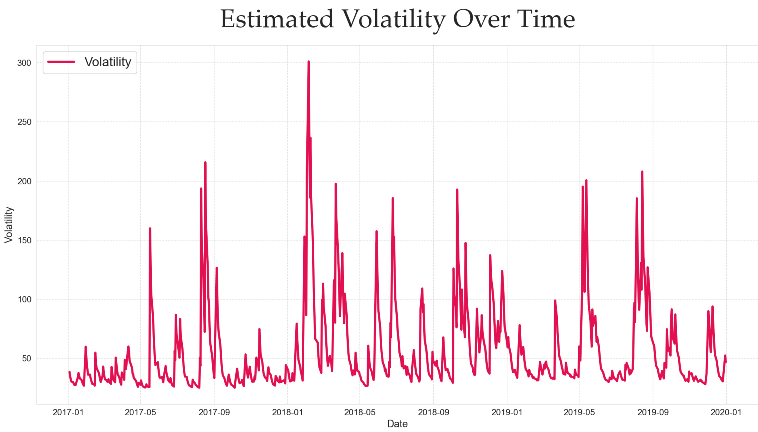

FIGURE 69: Appliation of Asymmetric Garch on the VIX.

This is the extended equation of Classical GARCH with the mentioned additional coefficient γ that corrects estimation depending on the ϵ_{t-1} value Indicator:

$$I_{\epsilon_{t-1}} = \begin{cases} 1, \text{ if } \epsilon_{t-1} < 0 \\ 0, \text{ if } \epsilon_{t-1} > 0 \end{cases} . \tag{69}$$

This extension is particularly well-suited for financial data because it differentiates between upside and downside shocks. Upside shocks are associated with stock growth, while the primary concerns are generally focused on negative shocks, which can lead to significant losses. The model is applied to the percentage changes of the VIX time series in a rolling manner. Forecasts for each day are generated by the model trained on a fixed number of previous days to avoid lookahead bias. The resulting dynamics are illustrated in Figure 69.

Here, we can observe the aforementioned clusters of volatility, characterized by low to moderate values that persist over relatively long periods. Extreme spikes in volatility are noted over 1-2 days and are not deemed significant for systematic settings; it is sufficient for volatility to exceed a certain threshold to be considered higher than during the prevailing market trend. However, these spikes are indicative of significant market drawdowns.

We have established reallocation thresholds at 40 and 60 to systematically differentiate between moderate and extreme volatility. In summary, if the value is between 0 and 40, no reallocation is performed. If the value falls between 40 and 60, a portion of the portfolio is reallocated to the SPY ETF. If volatility exceeds 60, a larger portion is reallocated to the SPY ETF.

Computation of RWAs for Securitisation Exposures: the iSEC Calculation Engine

About the Authors

Giuseppe Morisani
Senior Business Analyst
He holds a BSc in Economics&Management from the University of Pavia and a MSc in Finance&Investments from the University of Leeds Business School. He has been working as an analyst for the Morgan Stanley Market Risk Portfolio Team located in London, Budapest, Frankfurt, and New York; and supported other global consulting firms as well as wealth management funds in the following areas: risk management, research, asset management, IT solutions and project management. Currently, he is working on various strategic projects at a major Italian bank.

Mattia Bainotti
Quant Analyst
He holds a MSc in Finance from Università del Piemonte Orientale paired with a postgraduate specialization degree in Quantitative Finance and Risk management from Università di Torino, Collegio Carlo Alberto.

Caterina Papetti
Quant Analyst
She holds a M.Sc. in Mathematical Engineering, Quantitative Finance, at Politecnico di Milano. She joined Iason in 2024 and, as a Quantitative Analyst, she is currently working on various strategic projects at a major Italian bank.

Computation of RWAs for Securitisation Exposures: the iSEC Calculation Engine

Giuseppe Morisani Mattia Bainotti Caterina Papetti

This paper provides a detailed analysis of the various approaches that banks must use for computing the Risk Weighted Assets (RWAs) for securitisation exposures within the current regulatory framework and presents an integrated solution developed by iason to ensure the appropriate and efficient calculation of those measures. The work starts with a brief overview of the basic concepts related to Securitisations and the related risks. Subsequently, the various models and procedures for calculating risk weights are explained in detail. The article concludes with a short presentation of the user-friendly Python-based tool named iSEC elaborated by iason for the precise and automatic computation of capital requirements related to securitisation exposures with the SEC-SA and SEC-ERBA methodologies. iSEC is available upon request and a demo can be obtained by contacting the authors of the paper. The iSEC tool can be integrated upon request to allow the computation of capital requirements under the SEC-IRBA approach based on banks specific needs and internal models.

SECURITISATION is a financial mechanism that consists in repackaging illiquid assets or group of assets into tradable interest-bearing securities.
Practically, the owner (usually a bank) of a particular financial asset, such as a pool of corporate loans, conveys them to a Special Purpose Vehicle (SPV), typically set up us a trust, that issues a new financial product in the capital markets in the form of securitisation, allowing banks and other types of investors to purchase it.
Therefore, securitisation allows the original lender (also named originator or creditor) to eliminate risky assets from its balance sheet, enabling the allocation of the correspondent amount of risk into new businesses, such as conventional lending.
Securitisation types are defined on the base of the category of the underlying pool of assets. Table 18 provides a few examples.
Securitisation can be sold to the public in different forms:

- *Single-class offering:* all investors receive a pro rata interest in the revenus generated by the asset pool;
- *Multi-class offering:* two or more classes (also called tranches) with different payout and risk features are created and offered to the public.

Tranches are fractions of a securitisation, that differ in terms of risk or other characteristics in order to be marketable to different investors. Tranches carry different maturities, yields, degrees of risk and privileges in repayment in case of default.
In particular, there are three common categories of tranches: senior, mezzanine and junior. The former typically contains assets with higher credit ratings than the junior tranche, as well as a lower risk premium.
With regard to multi-class offering, it becomes crucial to understand the concept of "cash flows waterfall", that represents the mechanism that stands behind the attribution of cash flows and losses related to each tranche of the securitisation.
The cash flow waterfall can be described as follows: senior tranches are the first to receive the interest payments. Junior tranches receive payments only after the senior and mezzanine tranches have been paid in full. Therefore, junior tranches are the first to absorb the losses related to the underlying pool of assets (e.g. in case of default).
It should also be noted that in the specific case of securitisations with "pro rata" payments, the interests are paid (and the losses are absorbed) with the same priority (and at the same time) for all noteholders of the same category (e.g. mezzanine noteholders).
Depending on some transaction-specific features and the asset category, each securitisation may be supported by various credit risk mitigation elements that represent tools used by financial institutions to improve the credit risk profile of the securitisation. For example, the overcollateral, which refers to the difference between the face amount of the asset pool and the nominal amount of backed securities, can be used to absorb the first portion of losses in case of default on the underlying pool.
Another type of credit risk mitigation element might be any additional spread between the interest paid to the holder of the securitisation tranche and the net amount of interest payments from the underlying assets.
Bank securitisation (especially Mortgage Backed Securities) have been one of the major driver of the 2007-2008 financial crisis. In particular, most analysts argue that some banks have lowered their credit standards, increasing the risk profile of their credit book by exploiting the possibility to transfer those risks to other investors in the form of securitisation [2].
There are four major categories of potential risks related to securitisation exposures: credit risk, counterparty risk, market risk and legal risk.

- *Credit Risk* emerges when borrowers in a loan pool fail to make payments due to either their inability or unwillingness to do so. By examining the nature of the underlying asset class, the strength of the origination processes, the past performance of the originator's overall portfolio, and the characteristics

Underlying Pool of Assets	Securitisation Type
Mortgage Loans	Mortgage Backed Securities (MBS)
Consumer Loans	Asset Backed Securities (ABS)
Corporate Loans	Credit Loan Obligation (CLO)
Corporate Bonds	Collateralised Bond Obligation (CBO)

TABLE 18: Securitisation Types.

of the loan pool, one can gain valuable insights into the credit risk associated with these borrowers.

- *Counterparty Risk* arises from the non-performance of parties involved in a transaction. The primary counterparties to be analyzed include the servicer, the designated bank, and the swap counterparties.
- *Market Risk* is related to any exogenous factor that may have an impact on the securitisation transaction, such as changes in the macroeconomic landscape, interest rates volatility or loans prepayments.
- *Legal Risk* pertains to the sale and transfer of receivables from the originator to the SPV as well as the possibility that the originator goes bankrupt and the cash flows related to the underlying pool of assets are not assigned to the securitisation's underwriters.

All of those risks can be partially mitigated by taking appropriate actions, after a meaningful risk mapping and evaluation procedure.

Risk Weighted Assets Calculation

Quantifying regulatory capital requirements in the context of securitisation is far from trivial. This section provides a comprehensive guide to calculate Risk-Weighted Assets (RWAs) for securitisation exposures in compliance with the industry regulations set out in the Capital Requirements Regulation (CRR) [3] and its amendments, that supplies the regulatory framework within which financial institutions operate, ensuring both proper risk management and sectoral robustness.
This section delves into three key methodologies: Standardized Approach for Credit Risk - Internal Ratings-Based Approach (SEC-IRBA), Standardized Approach for Credit Risk - Standardized Approach (SEC-SA), and Standardized Approach for Credit Risk - External Ratings-Based Approach (SEC-ERBA), aiming to understand how these methods can be applied purposefully to securitisation, ensuring accurate assessment and prudent capital management.
This work presents key formulas and processes associated with RWAs calculation, highlighting how these methodologies can be adapted and effectively applied for estimating transactions profitability.
Firstly, we explore the hierarchy of Risk-Weighted Assets (RWAs) calculation methods suggested by the regulator. This hierarchical approach allows us to contextualize the three main methods under consideration.
Subsequently, a detailed summary of the inputs necessary for the implementation of each method is provided. This summary serves as a practical guide for financial institutions, offering a clear overview of essential data for accurate RWAs calculation.

The last subsection relates to the explanation of the necessary formulas for RWAs calculation under the SEC-IRBA, SEC-SA and SEC-ERBA approaches: it is important to specify that, since the SEC-IRBA represents a method based on banks internal models, this paper do not specify all the steps needed to determine the risk weights under this approach, hence concentrates on illustrating all the fundamental information and parameters that banks must consider for developing those types of models.

Hierarchy of Methods

Banks must follow a specific hierarchy when calculating the risk weights related to a securitisation exposure: the Internal Ratings-Based Approach (SEC-IRBA) stands at the top of the hierarchy. In order to use SEC-IRBA, banks must have a supervisory approved IRB model for the type of underlying exposure in the securitisation pool and sufficient information to estimate the capital charge for these underlying exposures. Specifically, institutions shall use the SEC-IRBA approach to calculate risk-weighted exposure amounts in relation to a securitisation position when the conditions set out in Article 258 [4] are met. In particular:

- The position is backed by an IRB pool or a mixed pool, provided that, in the latter case, the institution is able to calculate K_{IRB} for a minimum of 95% of the underlying exposure amount. Details on the K_{IRB} calculation are available for consultation in Article 255 [4] and Subsection II [3].
- There is sufficient information available in relation to the underlying exposures of the securitisation for the institution to be able to calculate K_{IRB}.
- The institution has not been precluded by the competent authorities from using the SEC-IRBA in relation to specified securitisation positions on a case-by-case basis, because of highly complex or risky features related to those securitisations, that are described in Article 258 [4].

Where the SEC-IRBA cannot be used, banks shall use the Securitisation - Standardised Approach (SEC-SA).
Instead, banks shall use the Securitisation - Exteral Rating Based Approach (SEC-ERBA) for rated positions or positions in respect of which an inferred rating may be used.
Moreover, an institution shall use the SEC-ERBA instead of the SEC-SA in each of the following cases[23].

- Where the application of the SEC-SA would result in a risk weight higher than 25% for positions qualifying as positions in an STS securitisation;
- Where the application of the SEC-SA would result in a risk weight higher than 25% or the application of the SEC-ERBA would result in a risk weight higher than 75% for positions not qualifying as positions in an STS securitisation;
- For securitisation transactions backed by pools of auto loans, auto leases and equipment leases.

[23] Article 254, [4].

For a position in re-securitisation[24], institutions shall apply the SEC-SA approach with the modifications set out in Article 269 of the Regulation (EU) 2017/2401 of the European Parliament and of the Council[25].
In all other cases, a risk weight of 1250% shall be assigned to securitisation positions[26].

Models Inputs

In this paragraph we propose a short description of the input variables used by the models presented in Section "Models Pipeline". This part is divided into four sections, depending on which model applies the input parameters described in each of them. The first passage contains a description of the inputs that apply to all the three models presented in Section"Models Pipeline".

Common Inputs

Attachment Point (*A*)[27] shall be expressed as a decimal value between zero and one and shall be equal to the greater of zero and the ratio of the outstanding balance of the pool of underlying exposures in the securitisation minus the outstanding balance of all tranches that rank senior or pari passu to the tranche containing the relevant securitisation position including the exposure itself to the outstanding balance of all the underlying exposures in the securitisation. Therefore, A represents the threshold at which losses within the underlying pool would be allocated to the *i-th* tranche of the securitisation at a given time *t*.
The concept of the Attachment Point (A) can be expressed mathematically. While Regulation (EU) 2017/2401 establishes the Attachment Point as a key metric, it does not recommend any specific formula for its calculation.
A possible representation of the Attachment Point equation could be:

$$A_{i,t} = Max\left(0; \frac{C_t - \sum_{j=i}^{n} L_{j,t}}{C_t}\right) \; for\ i = 1,..., n, \tag{70}$$

- n represents the total number of tranches of the securitazion (*j = 1* indicates the most junior tranche, *j = n* the most senior);
- C_t represents the collateral value at time t;
- $L_{j,t}$ represents the *j-th* tranche amount.

Detachment Point (*D*)[27] shall be expressed as a decimal value between zero and one and shall be equal to the greater of zero and the ratio of the outstanding balance of the pool of underlying exposures in the securitisation minus the outstanding balance of all tranches that rank senior to the tranche containing the relevant securitisation position to the outstanding balance of all the underlying exposures in the securitisation. Therefore, D represents the threshold at which losses within the underlying pool result in a total loss of principal for the *i-th* tranche of the securitisation at a given time *t*.
The CRR does not recommend any specific formula for the calculation of the Detachment Point. Below a possible representation of the Detachment Point equation:

$$D_{i,t} = \begin{cases} Min\left(1, \frac{C_t - \sum_{j=i+1}^{n} L_{j,t}}{C_t}\right) & \text{if i} \in [1, n-1] \\ 1 & \text{if i = n} \end{cases}, \tag{71}$$

- n represents the total number of tranches of the securitazion (*j = 1* indicates the most junior tranche, *j = n* the most senior);
- C_t represents the collateral value at time t, including overcollateral;
- $L_{j,t}$ represents the *j-th* tranche amount.

Figure 71 represents the Attachment and Detachment points for a securitisation with four tranches and €100,000 collateral, given an equity tranche of €10,000, two mezzanines of €20,000 and a senior of €50,000.

Delinquency Ratio[28] **(*W*)** represents the percentage of delinquent underlying exposures with respect to the total amount of underlying exposures. Delinquent underlying exposures are assets that are 90 days or more past due, subject to bankruptcy or insolvency proceedings, in the process of foreclosure, held as real estate owned, or in default, where default is defined within the securitisation deal documents. With regard to resecuritisation exposures, the parameter W is set to zero. **Exposure at Default** (EAD) is the predicted amount of loss a bank may face in the event of, and at the time of, the borrower's default. **Loss Given Default** (LGD) is a lender's (creditor) as projected loss in the event that a borrower triggers an event of default.

SEC-IRBA Inputs

Capital Charge (K_{IRB}) should be calculated in accordance with *Article 255* [4] by multiplying the risk-weighted exposure amounts in respect of the underlying exposures as if they had not been securitised by 8% divided by the exposure value of the underlying exposures.

SEC-SA Inputs

Risk Weight Standard (*RW_{STD}*) The allocation of the RW_{STD} is carried out pursuant to Regulation (EU) No 575/2013 from Article 114 to 134 [3], in accordance with the amendments that may come into force in 2025 following the approval of the Proposal for a Regulation of the European Parliament and of the Council amending the Regulation (EU) No 575/2013 [4] and represents the risk weight assigned for the calculation of capital requirements under the Standardised Approach to credit risk in relation to the underlying exposures as if they had not been securitised. Table 20 represents a summary of the risk weights suggested in the CRR with regard to the major asset classes.
The $RW_{STD,i}$ for each subpool i is calculated via a weighted average between its performing portion's risk weighting factor and its default portion's risk weighting factor. With regard to exposures in a state of default, the risk weighting factor assigned is equal to 100% if the underlying portfolio refers to exposures fully and totally guaranteed by mortgages on residential or non-residential properties. In the event of default of underlyings belonging to all other categories, the risk weighting factor assigned is equal to

[24] Article 258, [4].
[25] Article 269, [4].
[26] Article 258, [4].
[27] Article 256, [4].
[28] Article 261, [4].

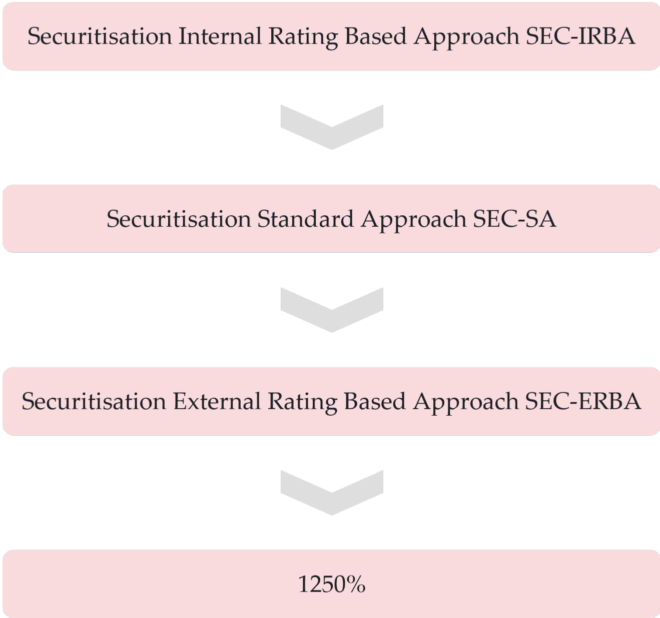

FIGURE 70: *Process for the approach definition.*

150%:

$$RW_{STD,i} = \left(\left(\frac{1 - EAD_i^d}{EAD_i}\right) \cdot RW_{B,i} + \frac{EAD_i^d}{EAD_i} \cdot RW_{DEF,i}\right)$$
$$\text{for } i = 1, .., M, \tag{72}$$

where:

- M is the total number of subpools;
- EAD_i^d and EAD_i are the exposure at default for the fraction of the subpool in default and in bonis respectively;
- $RW_{B,i}$ is the Risk Weight of the performing portion;
- $RW_{DEF,i}$ is the Risk Weight of the default portion.

Furthermore, in the specific case of exposures to SMEs not in default, the risk weighted factor for the performing portion of the securitized portfolio to be used for the determination of K_{SA} is adjusted based on what is indicated in *Article 501* [5] is modified as follows:

$$RW_{STD,i} = \left(\left(\frac{1 - EAD_i^d}{EAD_i}\right) \cdot RW_{B,i} \cdot SME_{factor} + \frac{EAD_i^d}{EAD_i} \cdot RW_{DEF,i} \text{ for } i = 1, .., M, \right. \tag{73}$$

where:

$$SME_{factor} = \frac{\min\{E^*; 2,500,000\} \cdot 0.7619 + \max\{E^* - 2,500,000; 0\} \cdot 0.85}{E^*}, \tag{74}$$

where E^* is the total amount, including any exposures in default, owed to the entity, its subsidiaries, its parent companies, or other subsidiaries of such parent companies, excluding, however, loans or potential loans secured by residential real estate, by SMEs, or by the SME's connected group of clients.

Resecuritazion[29] **(p_{SA})** represents whether or not it is a resecuritisation or an STS securitisation:

$$p_{SA} = \begin{cases} 1 & \text{if Standard} \\ 0.5 & \text{if STS securitisation} \\ 1.5 & \text{if Resecuritisation} \end{cases}. \tag{75}$$

SEC-ERBA Inputs

Credit Quality Step is related to the standard used for mapping ECAIs' ratings to a uniform scale.
Legal Maturity (M_L) is the legal maturity of the tranche.

Models Pipeline

In order to calculate the RWA for securitizations, the Risk Weight calculated in the following sections must be multiplied by the Exposure at Default:

$$RWA = EAD \cdot RW. \tag{76}$$

SEC-IRBA

The model specifications are available on *Article 259* [4]. The methodology for the computation of the Risk Weigth within the SEC-IRBA framework for securitisation is defined as follows.
The first step consists in the computation of the capital requirement per unit of securitisation exposure ($K_{SSFA(IRB)}$[30]) which is defined as:

$$K_{SSFA(IRB)} = \frac{e^{a \cdot u} - e^{a \cdot l}}{a \cdot (u - l)}, \tag{77}$$

where:

- $a = -\frac{1}{(p_{IRB} \cdot K_{IRB})}$;
- $u = D - K_{IRB}$;
- $l = max(A - K_{IRB}, 0)$;
- K_{IRB} should be calculated in accordance with Section "SEC-IRBA Inputs",

[29] Article 261, 262 and 296, [4].
[30] Article 255, [4].

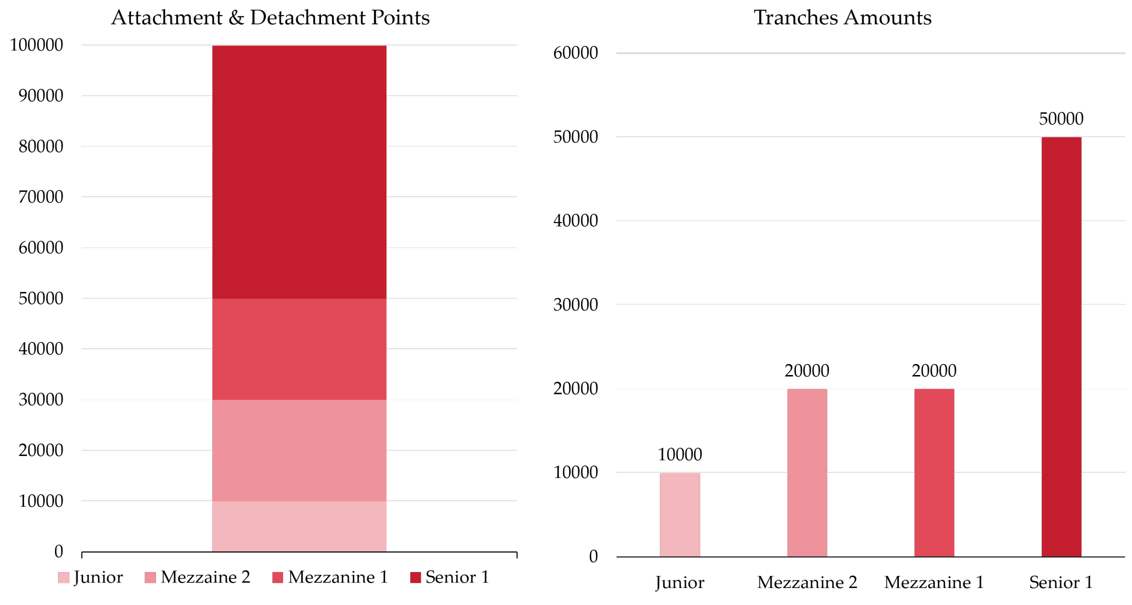

***FIGURE 71:** Attachment and Detachment Points.*

where:

$$p_{IRB} = max(0.3; \\ (A_{IRB} + B_{IRB} \cdot \frac{1}{N} + C_{IRB} \cdot K_{IRB} + D_{IRB} \cdot LGD + E_{IRB} \cdot M_T)), \tag{78}$$

and for STS securitisation

$$p_{IRB} = max(0.3; \\ 0.5 \cdot (A_{IRB} + B_{IRB} \cdot \frac{1}{N} + C_{IRB} \cdot K_{IRB} + D_{IRB} \cdot LGD + E_{IRB} \cdot M_T), \tag{79}$$

where:

- N is the effective number of exposures in the pool of underlying exposures:

$$N = \frac{(\sum_i EAD_i)^2}{\sum_i EAD_i^2}. \tag{80}$$

- LGD is the exposure-weighted average loss-given-default of the pool of underlying exposures:

$$LGD = \frac{\sum_i EAD_i \cdot EAD_i}{\sum_i EAD_i}. \tag{81}$$

- M_T is the maturity of the tranche[31]:

$$M_T = 1 + (M_L - 1) \cdot 80\%. \tag{82}$$

- The parameters A_{IRB}, B_{IRB}, C_{IRB}, D_{IRB}, and E_{IRB} shall be determined according to the look-up Table 19 in Appendix.

Finally, the Risk Weight[32] for the securitisation tranche is computed as:

$$RW = \begin{cases} K_{SSFA(IRB)} \cdot 12.5 & \text{if } K_{IRB} \leq A \\ \frac{(K_{IRB} - A)}{(D-A)} \cdot 12.5 + \frac{(D - K_{IRB})}{(D-A)} \cdot 12.5 \cdot K_{SSFA(IRB)} & \text{if } A < K_{IRB} < D \\ 1250\% & \text{if } K_{IRB} \geq D \end{cases}, \tag{83}$$

with regard to the mezzanine tranches, after the computation of the Risk Weights it might be necessary to perform a Seniority Reclassification as described in Section "Seniority Reclassification". Moreover the Risk Weights computed as described are subject to the following constraints:

- The Risk Weight cannot be lower than 15%[32];
- With respect to the STS securitisation[33], the Risk Weight cannot be lower then 10%.

SEC-SA

Example - Application of SEC-SA Approach

The model specifications are available on *Article 261* [4] amending Regulation (EU) No 575/2013 on prudential requirements for credit institutions and investment firms [3]. The methodology for the computation of the Risk Weight within the SEC-SA framework for securitisations is defined as follows.

The first step is the computation of the Weighted Average Capital Charge[34] (K_{SA}) of the entire portfolio of underlying exposures. This variable reflects the effects of any credit risk mitigant that is applied to the underlying exposures (either individually or to the entire pool), and hence benefits all of the securitisation exposures. The computation of K_{SA} is given by:

$$K_{SA} = 8\% \cdot \sum_{i=1}^{M} \frac{RW_{STD,i} \cdot EAD_i}{EAD}, \tag{84}$$

where EAD_i and EAD are the exposure of default of the single subpool and of the entire portfolio, respectively.

The next step consists on the computation of K_A[35]:

$$K_A = (1 - W) \cdot K_{SA} + 0.5 \cdot W. \tag{85}$$

In case the delinquency status is not known for no more than 5% of underlying exposures in the pool, the SEC-SA

[31] Article 257, [4].
[32] Article 259, [4].
[33] Article 260, [4].
[34] Article 255, [4].
[35] Article 261, [4].

approach can still be employed by adjusting the calculation of K_A as follows:

$$K_A = \left(\frac{EAD^{Wknown}}{EAD_{Total}} \cdot K_A^{Wknown} \right) + \frac{EAD^{Wunknown}}{EAD^{Total}}, \quad (86)$$

where:

- K_A^{Wknown} is computed according to Eq. (85);
- EAD^{Wknown} refers to the exposure at default for which the information on the delinquency status is available;
- $EAD^{Wunknown}$ refers to the exposure at default for which the information on the delinquency status is not available;
- EAD^{Total} refers to the total exposure at default in the pool.

The following step consists in the computation of the capital requirement per unit of securitisation exposure (K_{SSFA}[35]) which is defined as:

$$K_{SSFA} = \frac{e^{a \cdot u} - e^{a \cdot l}}{a \cdot (u - l)}; \quad (87)$$

where:

- $a = -\frac{1}{(p_{SA} \cdot K_A)}$;
- $u = D - K_A$;
- $l = max(A - K_A, 0)$.

Finally, the Risk Weight[35] for the securitisation tranche is computed as:

$$RW = \begin{cases} K_{SSFA} \cdot 12.5 & \text{if } K_A \leq A \\ \frac{(K_A - A)}{(D-A)} \cdot 12.5 + \frac{(D-K_A)}{(D-A)} \cdot 12.5 \cdot K_{SSFA} & \text{if } A < K_A < D \\ 1250\% & \text{if } K_A \geq D \end{cases}, \quad (88)$$

with regard to the mezzanine tranches, after the computation of the Risk Weights it might be necessary to perform a Seniority Reclassification as described in Section "Seniority Reclassification".
Moreover the Risk Weights computed as described are subject to the following constraints:

- The Risk Weight cannot be lower than 15%[35];
- The Risk Weight for the senior tranches cannot be higher then the weighting factor of the underlying portfolio (RW_{STD})[36];
- With respect to the STS securitisation, the Risk Weight for the senior tranches cannot be lower then 10%[37] ;
- Resecuritisation exposures are subject to a floor risk weight of 100%[38] ;
- The Risk Weight for a position in an NPE securitisation position shall be subject to a floor of 100%.

The following is an example of how to calculate risk weights through the SEC-SA approach for a securitisation with the following characteristics. The iSEC calculation engine has been used to produce the outputs.

- Collateral Amount: 100,000 €;
- Type of Securitisation: Standard;
- Payment Type: Sequential;
- Number of Subpool: 2;
- Tranches Amounts:
  - Junior: 10,000 €;
  - Mezzanine 2: 20,000 €;
  - Mezzanine 1: 20,000 €;
  - Senior: 50,000 €.
- Percentage of underlying assets for which delinquency status is known: 100%;
- Subpool characteristics:
  - Subpool 1:
    * EAD: 60%;
    * Delinquency Ratio: 0%;
    * Underlying Typology: Corporate.
  - Subpool 2:
    * EAD: 40%;
    * Delinquency Ratio: 0%;
    * Underlying Typology: Mortage non performing.

Results:

Seniority	Risk Weight
Junior	1,221.20%
Mezzanine 2	357.44%
Mezzanine 1	100%
Senior	100%

[36] Article 267, [4].
[37] Article 262, [4].
[38] Article 269, [4].

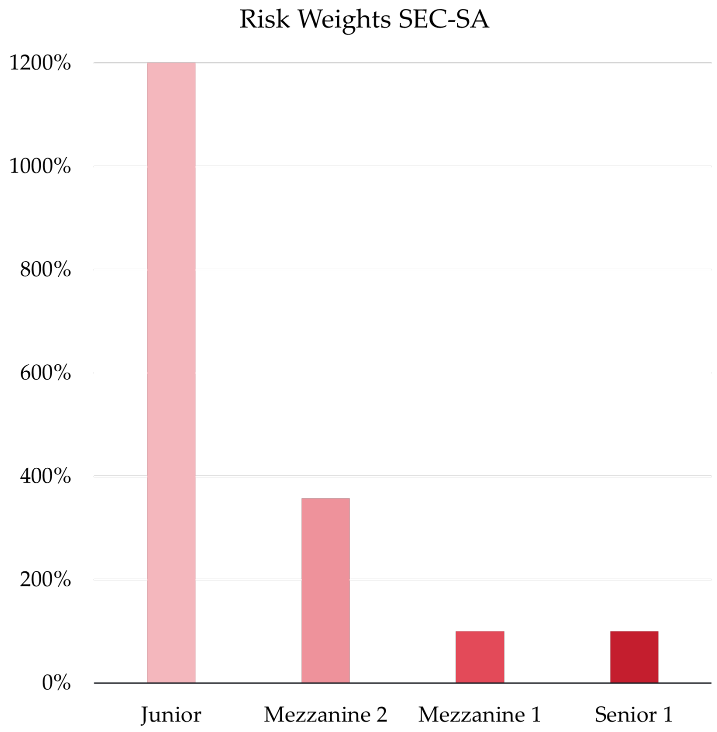

FIGURE 72: *Risk Weights from iSEC Engine.*

SEC-ERBA

Under the SEC-ERBA, the risk-weighted exposure amount for a securitisation position shall be calculated by multiplying the exposure value of the position[39] by the applicable risk weight.
When applying the SEC-ERBA, the RW is defined by the external rating of the tranche, its seniority, thickness and maturity[40].
The bank should apply the risk weights in the Tables reported in Appendix. In particular:

- For exposures with short-term credit assessment or when a rating based on a short-term credit assessment may be inferred, the bank should apply[41] the risk weights in Table 21 (or Table 23 in the case of STS).
- For exposures with long-term credit assessments or when a rating based on a long-term credit assessment may be inferred[41], the bank should apply the risk weights set out in Table 22 (or Table 24 in the case of STS), adjusting for maturity in accordance with the below model.

Moreover, in all the above cases, banks should adjust for tranche thickness the risk weights calculated with regard to non-senior tranches, as specified in the model.
Firstly, the tranche legal maturity M_L, is adjusted in accordance with the following formula[42]:

$$M_T = 1 + (M_L - 1) \cdot 80\%. \tag{89}$$

For the purposes of the Risk Weight calculation within the SEC-ERBA framework the determination of a tranche maturity (M_T) is subject in all cases to a floor of 1 year and a cap of 5 years.
Moreover, in order to determine the risk weight for tranches with a maturity between 1 and 5 years, institutions shall use linear interpolation between the risk weights applicable for 1 and 5 years maturities. Thus, the Risk Weight related to maturity M, RW_M, is given by:

$$RW_M = RW_1 + (M_T - 1) \cdot \frac{RW_5 - RW_1}{5 - 1}, \tag{90}$$

where:

- M_T is equal to tranche maturity computed according to Eq. 89;
- RW_1, RW_5 are respectively the one and five years Risk Weights according to Table 22 (or Table 24 in case of STS securitisation).

Finally, in order to calculate the Risk Weight for non-senior tranches it is necessary to adjust for tranche thickness the Risk Weight calculated in accordance to Eq. 90:

$$RW = RW_M \cdot (1 - \min(T; 50\%)), \tag{91}$$

where T is the tranche thickness defined as:

$$T = D - A, \tag{92}$$

with A and D attachment and detachment point defined in Section "Seniority Reclassification".
With regard to the mezzanine tranches, after the computation of the Risk Weights it might be necessary to perform a Seniority Reclassification as described in Section "Seniority Reclassification". Moreover the Risk Weights computed as described are subject to the following constraints:

- The Risk Weight cannot be lower than 15%;

[39] Article 248, [4].
[40] Article 68, [1].
[41] Article 263(7), [4].
[42] Article 257, [4].

Seniority	Risk Weight
Junior	72.59%
Mezzanine 2	64.53%
Mezzanine 1	64.53%
Senior	36.10%

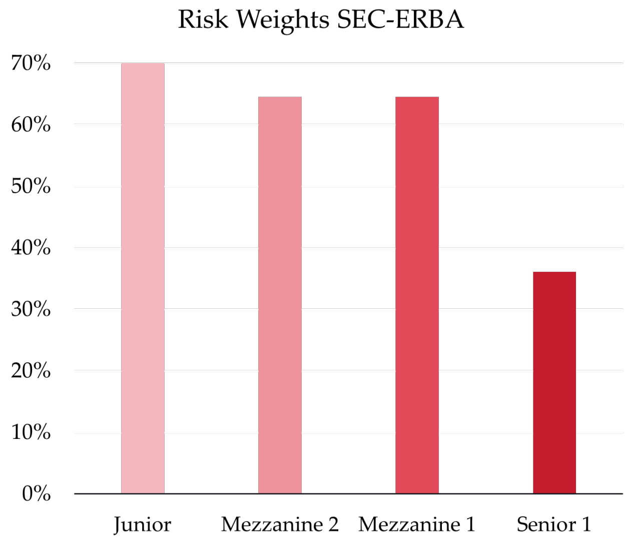

FIGURE 73: Risk Weights from iSEC Engine.

- The resulting risk weights shall be no lower than the Risk Weight corresponding to a hypothetical senior tranche of the same securitisation with the same credit assessment and maturity.

Example - Application of SEC-ERBA Approach

The following is an example of how to calculate risk weights through the SEC-ERBA approach for a securitisation with the following characteristics. The iSEC calculation engine has been used to produce the outputs.

- Collateral Amount: 100,000 €;
- Type of Securitisation: Standard;
- Payment Type: Sequential;
- Number of Subpool: 2;
- Tranches Amounts:
  - Junior: 10,000 €;
  - Mezzanine 2: 20,000 €;
  - Mezzanine 1: 20,000 €;
  - Senior: 50,000 €.
- Maturity: 30 June 2027 (long);
- Credit Quality Step: 4.

Results: See Figure 73.

Seniority Reclassification

Traditionally, securitisation tranches are classified into senior, mezzanine, and junior tranches based on their priority in the cash flow waterfall and associated risk levels. However, the Regulation (EU) 2017/2401 [4] establishes specific risk weight thresholds that allow mezzanine tranches to be reclassified into senior or junior ones, reflecting their "real" risk level and potential impact on capital requirements. The thresholds are as follows:

- **Mezzanine to Senior Reclassification**: a mezzanine tranche is reclassified as a senior tranche if its risk weight (RW) is less than 25%, which is considered a relatively low risk level that allows the treatment of the securitisation tranche to be amended in order to make it equivalent to that of a senior tranche.
- **Mezzanine to Junior Reclassification**: a mezzanine tranche is reclassified as a junior tranche if its risk weight (RW) exceeds 1250%. This reflects a high-risk level, similar to that of a traditional junior tranche.

Where mezzanine tranches are reclassified into senior ones, it is necessary to verify that all the constraints that belong to senior tranches and described in paragraph "Models Pipeline"are satisfied.

iSEC Calculation Engine

The Python-based calculation engine complies with the methodologies described in the previous paragraphs and it is structured as follows:

- *Home Page*: in this page the user must insert all the general information related to the securitisation that he would like to evaluate, such as collateral amount, typology and number of subpools;
- *Attachment and Detechment Points*: in this section the user can visualize the Attachment and Detachment

Points related to the securitisation he/she has created;

- *SEC-ERBA*: in this module the user should enter the specific inputs related to the SEC-ERBA model that basically depend on the rating assigned to the pool by the selected ECAI that should be converted to a Credit Quality Step and its maturity. By combining these information the tool executes the calculation of the Risk Weighted Assets under the SEC-ERBA approach;
- *SEC-SA*: in this page the user should insert the specific inputs related to the SEC-SA methodology with regard to each subpool and the delinquency status of the entire pool. Subsequently, the user can compute the Risk Weighted Assets. The tool uses the RW_{STD} reported in Table 20, which represents a summary of the risk weights suggested in the CRR with regard to the major asset classes. The flexibility of the tool and the underlying Python code would permit to easily make changes to the structure and the values reported in the table if needed.

The tool is available upon request.

References

[1] **Basel Committee on Banking Supervision.** *Revisions to the securitisation framework - Amended to include the alternative capital treatment for "simple, transparent and comparable" securitisations*. BIS, July 2016.

[2] **David Marques-Ibanez.** *Securitisation, credit risk and lending standards revisited*. ECB, Research Bulletin N. 32, March 2017.

[3] **European Parliament and Council Of The European Union.** *Regulation (EU) N. 575/2013 on prudential requirements for credit institutions and investment firms and amending Regulation (EU) N. 648/2012*. Official Journal of the European Union, June 2013.

[4] **European Parliament and Council Of The European Union.** *Regulation (EU) N. 2017/2401 amending Regulation (EU) N. 575/2013 on prudential requirements for credit institutions and investment firms*. Official Journal of the European Union, December 2017.

[5] **European Parliament and Council Of The European Union.** *Regulation (EU) N. 2019/876 amending Regulation (EU) N. 575/2013 as regards the leverage ratio, the net stable funding ratio, requirements for own funds and eligible liabilities, counterparty credit risk, market risk, exposures to central counterparties, exposures to collective investment undertakings, large exposures, reporting and disclosure requirements, and Regulation (EU) N. 648/2012*. Official Journal of the European Union, May 2019.

Annex

		A_{IRB}	B_{IRB}	C_{IRB}	D_{IRB}	D_{IRB}
Non Retail	Senior. granular (N>=25)	0	3.56	-1.85	0.55	0.07
	Senior. non-granular (N<25)	0.11	2.61	-2.91	0.68	0.07
	Non-senior. granular (N>=25)	0.16	2.87	-1.03	0.21	0.07
	Non-senior. non-granular (N<25)	0.22	2.35	-2.46	0.48	0.07
Retail	Senior	0	0	-7.48	0.71	0.24
	Non-senior	0	0	-5.78	0.55	0.27

TABLE 19: *SEC-IRBA parameters.*

Deal Type	RW_{STD}	RW_{DEF}
Corporate	100%	150%
SME Corporate	100%	150%
Retail	75%	150%
SME Retail	75%	150%
Mortgage	35%	100%
Mortgage Non Performing	100%	100%
Commercial Mortgage	50%	100%
Non Performing	150%	150%
Unrated Institutions	100%	150%
Banks	0%	150%

TABLE 20: *Risk Weights SEC-SA.*

Credit Quality Step	1	2	3	All other
Risk Weight	15%	50%	100%	1,250%

TABLE 21: *Short-term credit assessment.*

Credit Quality Step	1 year senior	5 years senior	1 year non-senior	5 years non-senior
1	15%	20%	15%	70%
2	15%	30%	15%	90%
3	25%	40%	30%	120%
4	30%	45%	40%	140%
5	40%	50%	60%	160%
6	50%	65%	80%	180%
7	60%	70%	120%	210%
8	75%	90%	170%	260%
9	90%	105%	220%	310%
10	120%	140%	330%	420%
11	140%	160%	470%	580%
12	160%	180%	620%	760%
13	200%	225%	750%	860%
14	250%	280%	900%	950%
15	310%	340%	1,050%	1,250%
16	380%	420%	1,130%	1,130%
17	460%	505%	1,250%	1,250%
All other	1,250%	1,250%	1,250%	1,250%

TABLE 22: Long-term Credit Assessment.

Credit Quality Step	**1**	**2**	**3**	**All other**
Risk Weight	10%	30%	60%	1,250%

TABLE 23: Short-term credit assessment for STS.

Credit Quality Step	1 year senior	5 years senior	1 year non-senior	5 years non-senior
1	10%	10%	15%	40%
2	10%	15%	15%	55%
3	15%	20%	15%	70%
4	15%	25%	25%	80%
5	20%	30%	35%	95%
6	30%	40%	60%	135%
7	35%	40%	95%	170%
8	45%	55%	150%	225%
9	55%	65%	180%	255%
10	70%	85%	270%	345%
11	120%	135%	405%	500%
12	135%	155%	536%	655%
13	170%	195%	645%	740%
14	225%	250%	810%	855%
15	280%	305%	945%	945%
16	340%	380%	1,015%	1,015%
17	415%	455%	1,250%	1,250%
All other	1,250%	1,250%	1,250%	1,250%

TABLE 24: *Long-term credit assessment for STS.*

Ensuring the Future: the Potential Evolution of the Insurance Market

About the Authors

Valerio Ciminelli:
Business Analyst
He holds a MSc in Finance, Intermediaries and Markets from Alma Mater Studiorum University of Bologna. He has experience in the management of complex IT and strategic projects for major clients in the Financial sector gained working closely with Securities Services and Credit Management.

Jacopo Figuriello:
Quant Analyst
He holds a master degree in Mathematical Engineering with a specialization in Quantitative Finance at Politecnico di Milano. After the graduation, he began his career at Iason where he has been working for about 2 years. He is currently working on a high-impact IT implementation Project for one of Italy's largest bank.

This document was prepared in collaboration with Nicola Mazzoni, who at the time was working for Iason Consulting.

Ensuring the Future: the Potential Evolution of the Insurance Market

Valerio Ciminelli Jacopo Figuriello Nicola Mazzoni

Financial markets are facing deep changes driven by the rapid advancement of digital innovation. In this ever-changing environment, new potential benefits and challenges arise and the organizations' need to develop more adaptive and agile solutions and business models has become even more crucial. Within this context, the rise of new risks, such as climate change, data security, technological-related risks, and economic volatility, and the issues related to traditional industry approaches, such as high operational costs and slow response times, will lead the insurance industry to rethink their business models and traditional risk management procedures. At the same time, the possibility given by the growth in technology and financial innovation could help tackle these risks, particularly through the integration of DeFi and Open Insurance. Through the implementation of environments that rely on blockchain, smart contracts, Oracles, and other data-sharing protocols, DeFi and Open Insurance will foster automated claims processing, better data gathering, and reduce information asymmetry. The research aims to explore how DeFi, Open Insurance frameworks, as well as their integration, could provide decentralized, automated solutions that address the limitations of conventional insurance approaches exploring future solutions where digital innovation empowers insurers to effectively navigate and mitigate modern risks.

Advancements in digital innovation are significantly transforming financial markets, introducing new potential benefits and challenges while reshaping traditional business models. In meanwhile the rise of emerging risks and the limited versatility of traditional industry practices are driving the need for new solutions and products within the insurance sector. This analysis explores the potential positive impacts of integrating technologies and paradigms such as asset tokenization, decentralized finance, and Open Insurance within the insurance industry. The research begins by examining the foundations of insurance, explaining the main drivers that lead the demand and the supply of insurance policies and providing an overview of industry trends and developments (a reader who is already comfortable with the concepts related to insurance markets could consider to go after to the next parts of the analysis or reading the chapter as a "refresher"). Next, a brief overview of Digital Ledger Technologies (DLTs) and asset tokenization is provided to set the stage for an analysis of insurance DeFi environments. The section dedicated to insurance DeFis will proceed in explaining the composition and the functioning of these environments, covering also the potential for the insurance market. After that, the focus will be put on Open Insurance, illustrating the main concepts of this new paradigm. Before concluding, the research will bring an illustrative view of some of the most relevant projects that rely on insurance DeFis and Open Insurance. Conclusions will focus on analyzing the potential benefit that the integration of insurance DeFis and Open Insurance could bring in the industry.

Insurance Foundations

Insurance Market

Looking back at history, we can see that risk coverage has always been part of human economic activities; indeed insurance policies were perhaps one of the first financial products ever developed, grounding its roots more than 3.000 years ago in the Babylonian Age with the first written evidence of pooling risks documented in the Hammurabi Code. Over centuries insurance market have gained in dimension and complexity evolving with the changing needs of societies becoming nowadays an industry managing trillions of dollars in assets. Despite the continuous technological innovation and the creation of countless insurance tools, the main element of the insurance industry remains the essential need it satisfies, the same need that led to its inception in the early ancient merchant communities: the need to protect one's assets (tangibles and intangibles) from the risks arising from the occurrence of uncertain events. This first chapter aims to analyze the key elements of insurance products and their main classifications, as well as to investigate the relationship between the insurance industry and the macroeconomic context to identify potential weaknesses in the sector that could become the targets for innovation introduced by blockchain technology.

According to academic literature, the main functions of insurance are:

- **Risk Allocation;**
- **Wealth Protection;**
- **Investment;**
- **Mobilization of Financial Resources;**
- **Improvement of Governance;**
- **Reduction of Public Expenditure.**

To better understand the achievement of the aforementioned functions, it is useful to focus on the components of demand and supply in the global insurance market.

Demand Analysis

Predominant academic literature identifies two main types of components in the demand within the insurance market:

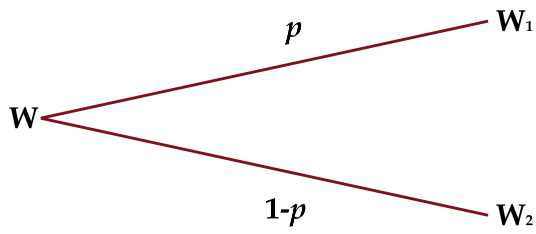

FIGURE 74: Monetary Lottery W.

- **Objective Components.** The risk of loss associated with owning an asset, external macroeconomic factors, and the level of premiums.
- **Subjective Components.** The risk aversion of individual economic agents.

The basic model of the demand for insurance coverage hypothesizes the presence of a binary risk scenario. Essentially, a loss of amount L on a particular asset will occur with probability p or, conversely, will not occur with probability $1-p$. This model is very similar to the classical microeconomic model of consumer choice (also the two goods model), where the goods represent, in our case, two alternative levels of asset values in the absence of insurance coverage:

- W_1**:** where the loss does not occur (with probability p).
- $W_2 = W_0 - L$**:** where the loss occurs with probability (1-p).

The model could be represented by the following monetary lottery W:

$$W = (W_1, W_2; p, 1-p).$$

The demand for insurance is influenced by a myriad of factors that operate both on individual and aggregate levels:

- **Individual-Level Drivers:**
  1. **Income and Wealth.** Income levels are a primary determinant of insurance demand. As individuals' income increases, they are more likely to purchase insurance because they have more assets that need to be protected and greater financial means to afford higher premiums. Wealthier individuals often pursue broader coverage policies, including health, life, and property insurance, to safeguard their accumulated assets. Higher-income earners are also more likely to opt for policies with higher coverage and additional benefits. More disposable income allows families to invest in financial planning and risk management tools, including insurance. This is because, after covering the cost of living, they have surplus funds that can be used to purchase insurance coverage. Higher-income earners, due to a higher level of education, often have a greater awareness of financial risks and a stronger desire to protect their wealth and assets. This drives them to purchase insurance as a mean of securing their financial stability and protecting their tangible and intangible assets
  2. **Risk Aversion.** Another factor that has been widely debated in academic literature, as a determinant of insurance demand, is risk aversion. An individual's risk aversion directly influences the perceived need to purchase insurance coverage. Those who perceive higher risks, whether due to health concerns, the nature of their occupation, or the safety of their assets, are more likely to seek insurance coverage. To maximize utility, risk-averse individuals are willing to pay a premium to avoid the potential loss of a large amount of wealth. This is the essence of insurance demand: by transferring risk to the insurer, individuals effectively "buy" safeness and preserve their utility against uncertainty. Risk aversion, measured in academic literature through the RRA (Relative Risk Aversion) variable, and its relationship with the demand for insurance products has been one of the most debated topics for decades in the economic field known as "behavioral finance." Risk perception is a crucial determinant of insurance demand. It refers to the subjective process by which an individual economic agent assesses the probability and magnitude of a potential loss due to adverse events. This perception influences his decisions regarding whether to purchase insurance and, if so, what types and amounts of coverage to seek[43]. In the 2021 paper by Jaspersen et al[13], the authors found moderately sized and statistically significant predictive power for the three most discussed primitive preference motives of insurance demand: utility curvature, probability weighting, and loss aversion.
  3. **Education Level.** Education level plays a pivotal role in determining individuals' financial decisions, including the demand for insurance coverage. Several academic papers have analyzed the relationship between education level and demand for insurance coverage. The evidence from these papers suggests a positive correlation between an individual's level of education and their participation in the insurance market. Education affects an individual's risk perception, understanding of financial products, and informed decision-making about protection against uncertain events. Guiso and Paiella (2008) [10] highlight that individuals with higher education tend to have a better assessment of the probability of adverse events and are better able to evaluate the benefits of insurance as a risk management tool. This improved risk assessment leads to a more rational decision-making process regarding insurance, where individuals weigh the cost of

[43] For a review of the empirical literature on risk aversion and with a particular focus on insurance demand or consumption see "J. François Outreville Risk Aversion, Risk Behaviour and demand for Insurance: a survey."

premiums against the potential financial consequences of unexpected events.

- **Aggregate-Level Drivers.** The demand for insurance is significantly affected by a range of macroeconomic factors, which include broader economic conditions and trends that influence both individuals' and businesses' financial behaviors and approaches to risk management. Financial literature understanding these macroeconomic influences is essential for comprehending how varying economic environments impact insurance markets. Below we will see how the main macroeconomic variables influence the demand for life and non-life insurance. The relationship between the broader economic environment and insurance demand is well-documented in the literature, with key factors including economic growth, inflation, interest rates, financial market development, and government policy playing prominent roles.

  1. **Economic Growth.** One of the most widely studied macroeconomic drivers of insurance demand is economic growth, obviously measured by GDP. Beck and Webb (2003) [2] conducted an empirical study across 68 countries and found that economic growth is a major determinant of life insurance demand. The authors used real GDP per capita as an indicator of permanent income, calculated as a predicted value from a regression of the log of each Country's real GDP per capita on a time trend (40 years observation period). The results of the regression indicate that a 10% increase in real income per capita increases life insurance penetration by 5.7%, confirming the theory that life insurance is a luxury good. As a Country's economy grows, individuals and businesses accumulate wealth, which increases the need to protect their assets and income streams through insurance products. To better understand the correlation between GDP per capita and insurance demand, as well as the correlation with other factors, it is necessary to introduce the variable typically used by academia and industry to quantify demand: the gross premium. The gross premium is the amount the insured pays for an insurance policy that is not the amount the insurance company earns for writing the policy. Gross premiums are typically adjusted upwards to account for commissions, selling expenses like discounts, and other insurer's expenses. In a 2014 paper (Cristea, Marcua, and Cârstina) [3], the authors found, using a multiple linear regression model, a high correlation between GDP per capita and gross premium in both the life and non-life sectors, with a stronger correlation observed in the life sector.
  2. **Inflation.** Inflation is another key macroeconomic variable that potentially impacts insurance demand, although its effects can be complex depending on the type of insurance product considered. In a 2002 research paper, Ward and Zurbruegg (2002) [15] state that inflation creates uncertainty in the insurance market, leading to higher premium costs. Insurers raise premiums to compensate for the expected increase in claims costs due to inflation, which can reduce the affordability of insurance and subsequently lower demand. As inflation erodes the cash value of any sums received in the future, the benefits of purchasing life insurance diminish. Further, higher levels of inflation are associated with macroeconomic uncertainty and as a result the discounted value of financial assets, including life insurance, will be less.
  3. **Interest Rates.** It comes as no surprise that interest rates play a critical role in the functioning of the financial markets and have also a profound influence on the demand for insurance, with particular reference to life insurance products which often include significant savings or investment components. Feyen et al. (2011) [9] argue that life insurance demand is inversely related to interest rates, particularly in developed economies. In a low-interest-rate environment, returns from life insurance policies are relatively more attractive compared to other financial products. As a result, individuals may prefer to purchase life insurance for both its risk coverage and savings benefits. Inversely, when interest rates rise, alternative financial products like bonds or savings accounts offer better returns, making life insurance policies less attractive.

Supply Analysis

Once we have completed the quick overview of the main drivers of demand for insurance products, the focus will now be provided on the supply side, thus turning our attention to insurance companies and their premium determination policies based on what is now a well-established academic theory. Insurance pricing is a critical aspect of the insurance industry, determining how premiums are set for various insurance products. It involves complex calculations that account for risk assessment, administrative costs, profit margins, and market conditions; the pricing of insurance products is so impactful and complex that it is the focus of an entire discipline: actuarial science. The primary objective of the pricing processes for insurance products is to ensure the economic sustainability of insurance companies' supply. Let's now take a closer look at the factors that contribute to determining the premiums of insurance products:

- **Risk Assessment:** The classical theory of insurance risk, in defining models for calculating insurance premiums, emphasizes the underwriting activities of insurance companies. Considering an Insurance Portfolio, actuaries form a stochastic process in time, consisting of two components:

  - Uncertain Number of Claims;
  - Uncertain Amount of a Claims.

  Here we describe a typical stochastic claim process: At time t_0, the insurer has a surplus, S_0, consisting of equity capital and accumulated reserves. Until the first claim occurs, the only inflow in the process is from premium payments, causing the surplus to grow in line with the amount of the premium over a given period:

  $$(1+\lambda)\pi,$$

  where π represents the fair premium income and λ is the surcharge per monetary unit which is a safety loading designed to mitigate the risk of insolvency. At time T and in the absence of any claim, the surplus would be:

  $$S(T) = S_0 + (1+\lambda)\pi \cdot T.$$

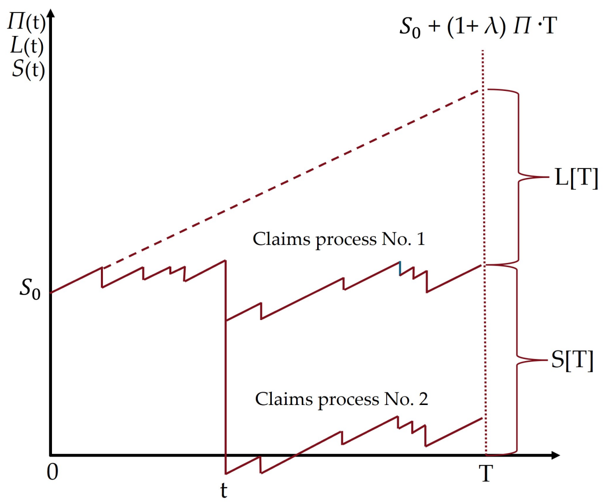

FIGURE 75: *Stochastic Claim Process.*

When the first claim occurs, the surplus decreases by the amount of loss payment. After that, the surplus process continues with the same slope until another claim is submitted. The vertical distance from the premium income process represents the cumulative value of losses paid L_T as clearly shown in Figure 75. Since this is a stochastic process, the claims process can take on different variables. For example, in Fig.2, we can see that claims process No. 2 is significantly less favorable from the insurer's perspective, with a high claim occurring at time t_c, causing the surplus to drop to negative values for a certain period and resulting in a concrete risk of insolvency for the insurer. It is precisely this concrete risk of insolvency that guides the insurer in the pricing process. In fact, by generalizing the claims process No. 2 described above, the latter can be represented by a loss distribution function f(x). Given an estimated loss distribution, the underwriters typically apply one of the several premium principles developed by academic literature to determine a fair premium $\pi(x)$. Considering the importance of the underwriting and risk assessment process in determining the insurance premium, it would be useful to explore the potential improvements that blockchain could bring to the process.

- **Regulatory Framework.** One of the most critical components of insurance regulation that affects supply is solvency regulation, which ensures that insurance companies maintain sufficient capital to meet their liabilities. The Solvency II Directive in the European Union[44], and similar frameworks like Risk-Based Capital (RBC)[45] requirements in the United States, are intended to maintain the financial stability of insurance companies by mandating a capital buffer proportional to the risk they underwrite. As well documented in an Eling & Schmeiser paper of 2010[7], solvency requirements have a direct effect on the capability of insurers to provide coverage, indeed, tighter capital requirements often lead to a reduction in supply, as insurers are forced to either raise new capital or limit the amount of risk they take on to maintain solvency levels.
- **Financial Markets.** Insurance companies are among the main players, alongside commercial and investment banks, in global financial markets. Their role stems from the need for capital reserves, as described in the previous two points, which drives insurers to constantly seek investment returns, particularly in debt instruments. It is therefore not surprising that there is a correlation between interest rate trends and the amount of insurance product supply. Considering the context of rising interest rates, insurance companies can rely on higher returns from their investments and expand their underwriting activities. Empirical studies have found that life insurance companies are particularly sensitive to changes in interest rates, as their long-term liabilities must be matched with assets generating consistent returns (Koijen & Yogo, 2016)[14].

Global Insurance Market: Outlook and Trends

Once the discussion of the main determinants of insurance product supply is completed, it is appropriate to conduct a deep dive into the trends of the global insurance market and the key industry developments. This will provide, by the end of the chapter, an initial "glance" at the role that tokenization could play in the insurance sector, also considering the relevant economic context.
The key figures and economic indicators for this section are mainly taken from the annual reports of the IAIS (International Association of Insurance Supervisors)[11] and EIOPA (European Insurance and Occupational Pensions Authority). The typical indicators considered by Market Authorities to monitor the health of the insurance industry are Solvency, Liquidity, and Profitability. The solvency ratio is the ratio of eligible own funds (Eligible Own Fund - EOF) to the solvency capital requirement (Solvency Capital Requirement - SCR) and expresses the level of capitalization of insurance companies. The SCR is the capital required for insurance companies to remain solvent with a probability of 99.5% over a one-year time horizon.

$$SolvencyRatio = \frac{EligibleOwnFund}{SolvencyCapitalRequirement}.$$

[44]For a complete overview see "Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II)".
[45]See NAIC's "Risk Based Capital (RBC) for Insurers Model Act", initially adopted in 1993 (latest revision, 2011).

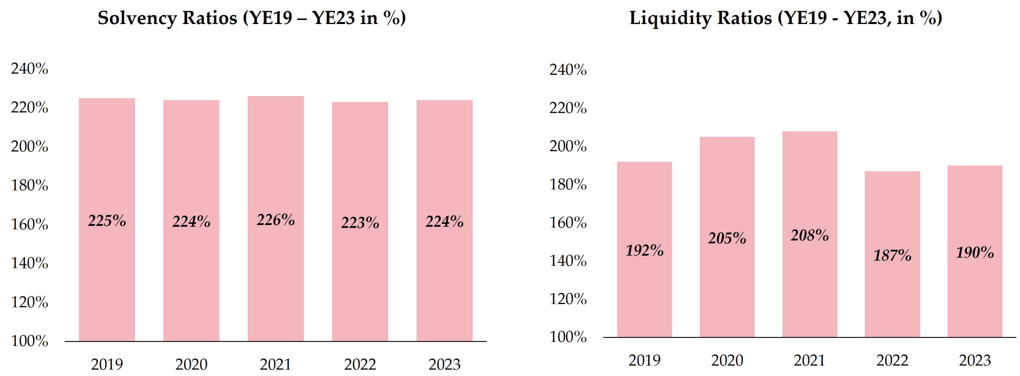

FIGURE 76: *Solvency Ratios and Liquidity Ratios.*

On the other hand, the ILR, as defined by IAIS, is the ratio of an insurer's liquidity sources and needs over a selected time horizon of assumed liquidity stress:

$$InsuranceLiquidityRatio = \frac{LiquiditySources}{LiquidityNeeds}.$$

The IAIS accurately identifies the various components that contribute to forming the Liquidity Sources, assigning each of them a specific weight to calculate the LCR (e.g. Cash 5%, Sovereign Debt 65%, Corporate Debt 30%).
Based on data collection from a pool of the 60 largest insurance groups worldwide, in its annual report, the IAIS notes that, during the observation period from 2019 to 2023, the solvency and liquidity of the global insurance market demonstrated considerable resilience to macroeconomic changes, showing a rather stable trend as illustrated in Figure 76.
The primary factors influencing solvency and profitability include increased interest rates in many regions (which can reduce the present value of liabilities, especially for life insurers offering long-term products), higher premium income, reduced dividend payouts by insurers, and an upswing in financial markets resulting in higher investment returns. Despite 2023 being a year of significant macroeconomic uncertainty and geopolitical instability, the insurance market and, in fact, the global economy, demonstrated relative stability, with a positive outlook for the next two years. This growth is expected to be supported by a decline in inflation resulting from monetary policy actions.
Now that we have analyzed the health of the global insurance market, it could be useful to provide a summary, non-exhaustive overview of the upcoming trends that insurance companies should embrace, as well as the risks that need closer monitoring.

- **High Interest Rates and Inflation.** According to IMF World Economic Outlook Update, the global economy is now stabilizing but we are not out of the woods yet, indeed, even if inflation is coming down it remains high in many countries, forcing their Central Banks to keep interest rates elevated. Furthermore, the complex geopolitical context must be considered, as it risks making any economic policy action much less effective; one only needs to think about the tensions that various conflict scenarios around the world can generate in international trade. These developments might mean that inflation is likely to remain a global concern for some time and any interest rate increase has potentially significant impacts on insurers. On one hand, Property and Casualty (P&C) insurers should face higher difficulties due to the potential rise of claims costs and they may need to minimize any shortfalls between premium revenue and claims payouts by raising premiums; on the other hand for Life and Retirement sector, higher interest rates could reduce the need to reinsure or transfer interest rate risks to other parties pushing insurers to potentially underwrite more new business. The "P&C" sector is therefore the most sensitive to rising interest rate dynamics. One additional tool available to insurance companies compared to the past could be leveraging the potential of new technologies in their information systems to obtain reliable data on "loss costs" and promptly guide business decisions aimed at minimizing this negative trend.
- **Credit Risk.** As revealed by a recent round of inspections by the ECB, there is a concrete credit risk associated with the deterioration of exposures secured by Commercial Real Estate, stemming from a collapse in CRE prices that began with the pandemic in 2020, leading to a drop in demand and further exacerbated by higher borrowing costs due to the inflationary context. This has implications for the insurance sector, which often holds substantial exposure (higher for life insurers) to these assets to generate income and manage risk. Furthermore, there is a business impact on the securitization market for real estate assets, which is also under downward pressure. In a 2024 scenario analysis conducted by the Federal Reserve Bank of Chicago[8], it has been estimated that the life insurance sector could face combined losses of about $36.3 billion from direct and indirect exposures to CRE.
- **Climate-Related Risks.** Climate change is contributing to heavy alterations in the Earth's climate system, leading to an increase in the frequency and intensity of catastrophic natural events. According to 2022 EIOPA's dashboard on the insurance protection gap for natural catastrophes, only around a quarter of the total economic losses caused by extreme weather and climate-related events are insured, leading to a substantial insurance protection gap. As the cost and frequency of claims increase, insurers are likely to increase premiums, potentially making insurance more expensive and even unaffordable. In extreme cases, insurers may withdraw from market segments as it becomes uneconomical to offer insurance, resulting in an even larger protection gap, as has happened recently in Florida in the aftermath of Hurricanes Helene and Milton. The increased frequency and severity of natural disasters make it harder for insurers to accurately predict future losses and appro-

priately price insurance products. Climate risk is a full-fledged source of financial risk, and as such, authorities are calling on insurance players to take an active role in the identification, measurement, mitigation, and coverage of climate risks, also with a view to contributing to the net-zero transition. However, there is often a gap between regulatory mandates and the actual achievement of targets due to market dynamics, which responds to the logic of utility maximization and risk aversion. A key to trying to influence market dynamics could be the introduction of innovative financial instruments by insurance companies that incentivize climate risk prevention, such as policies where the premium is tied to the implementation of risk exposure reduction measures by the insured (e.g., fire prevention systems, flood barriers, alarm systems...). A key challenge in offering these products is collecting high-quality, reliable data on the insured. To this end, insurers could leverage Open Insurance and tokenization to reduce data collection costs and significantly lower information asymmetry. An adequate and cost-effective data collection process could also reduce the risk for insurers related to incorrect policy pricing, possibly through the implementation of scenario analysis to properly calibrate their models and limit underwriting losses.

- **Impact of Digitalization and AI.** The rapid advancement of technologies such as big data, artificial intelligence (AI), machine learning (ML), blockchain, and the Internet of Things (IoT) is set to bring profound changes in how insurance companies operate, engage with customers, and manage risks. Each of these technologies can potentially bring benefits to both insurance companies and policyholders by introducing innovative services or simply making distribution channels more efficient. According to a 2024 EIOPA's Report on the digitalization of the European Insurance Sector[6], more than 52% of insurance companies that participated in the survey already have a dedicated digital transformation strategy in place, that covers areas such as customer service/user experience, digital infrastructure, data interfaces, digital platforms, and ecosystems, Application Program Interfaces (APIs), digital distribution channels, interaction with sales agents or IT security issues. Going into more detail, we can observe that while 50% of European insurance players in the non-life segment are already using artificial intelligence tools, only 15% of respondents reported actual usage of blockchain. Moreover, half of these use cases are still at the proof-of-concept stage. A noteworthy example is that of home insurance policies linked to mortgages. This is a collaborative project between the banking and insurance systems, which ensures access to a DLT database containing information related to mortgage loans. This allows for an automatic insurance offering to customers who take out a mortgage, using smart contracts. We believe, with a certain level of confidence, that the use of Open Insurance processes combined with the implementation of DLT technologies could bring numerous benefits to the insurance industry by reducing operational and agency costs through faster and more efficient processes; furthermore, digitalization could enable providers to automatize their interaction with clients, including at the claims stage, facilitating the accessibility for customers, and speeding the overall claims management processes.

Digital Ledger Technology

In recent years, digital assets and asset tokenization have captured the interest of many, especially in the financial sector, for their potential disruptive applications in a wide range of markets. The foundation that has made the spread of Digital Assets possible is the DLT (Digital Ledger Technology). In analyzing the potential benefits that insurance DeFis could bring to the insurance market, it seems necessary to begin by exploring the concepts related to DLTs. DLTs are nothing less than a digital ledger that differs from the classical centralized digital ledger by its core feature of being distributed in identical copies between the nodes that compose the network. One of the main characteristics of DLTs is that the information of the ledgers is constantly and simultaneously updated on each node of the network for every transaction. The following paragraphs will focus on explaining:

- **Consensus Mechanism;**
- **DLTs Architectures;**
- **Smart Contracts;**
- **Tokens;**
- **Asset Tokenization.**

Consensus Mechanism

One of the core features of DLTs is that every node in the network has stored the same information as the other nodes. When an event changes the information embedded in a node, each other node of the network must be updated in the same way. In a decentralized environment, one of the biggest challenges is to accord participants in the way a decision, such as the update of a node, should be taken and accepted. To prevent chaotic situations and address double-spending issues, DLTs rely on a particular cluster of algorithms that helps the network to agree on the information that must be updated in the nodes (the introduction of consensus mechanisms helps to overcome the so-called, Byzantine Generals Problem[46]. In the DLTs landscape, a lot of different consensus algorithms have been adopted to solve the need to coordinate the network's nodes update, in particular, we can cite as the most notables:

- **The Proof of Work (PoW).** The PoW mechanism relies on a "Mining Process that requires the network components to solve high-level computational problems to validate the information and update the ledgers"[16]. The first participant in the network who correctly solves the hash puzzle receives a reward (e.g. a unit of a cryptocurrency). The other participants can easily prove the solutions proposed and verify if it's correct or not. After the correctness is verified, the new information is updated across the entire network. The PoW ensures network trust and robustness against fraud due to its high computational and energy costs. Despite its high costs and its low scalability PoW represents one of the most famous consensus mechanisms and is the one on which Bitcoin relies on.

[46] An informatic problem that describes a scenario where multiple participants (the generals) have to agree on a coordinated decision, but some of the agents should be against the system. The challenge is to find a consensus even when some of the participants act against the system.

- **The Proof of Stake (PoS).** The PoS mechanism does not rely on the computational power of the validators, instead it requires the participants of the network that want to validate transactions to store as a collateral part of their tokens in the network. In case of malevolent behavior, the "stake" will be "slashed", and his tokens will be detained within the network.
- **The Proof of Authority (PoA).** In the PoA mechanism the validators are predefined trusted public entities. This configuration permits a faster and less costly environment in exchange for a reduction of the degree of decentralization of the network
- **The Proof of Capacity (PoC).** The PoC mechanism is based on the storage power of the network participants. The validators must store a pool of hash solutions in a memory space and every time a new block of the chain is created, they must search for the solutions inside their storage. The bigger the memory space they have at their disposal the bigger the probability is of having the solution of the hash in the downloaded pool of solutions.
- **The Proof of Burn (PoB).** The PoB bases its trusted validation mechanism on the requirement of "burning" tokens to become a network validator. This configuration requires that validators destroy a certain number of tokens (that are removed from the network circulation) to participate in the "mining" of new nodes. Typically, the reward for the miners is the mined fresh new coins.
- **Direct Acyclic Graph (DAG).** In DAG environments each transaction must be linked to previous validated transactions to be accepted by the network. Relying on its network for the connection between transactions DAGs do not require miners or validators resulting in an energy-efficient environment but the absence of validators hides an increase in security costs to maintain a trusted network
- **Byzantine Fault Tolerance Models (BFT-Models).** The BFT Models refer to a multitude of architectures that aim to solve the Byzantine Generals' Problem ensuring that the consensus within a network is reached even if a third of the participants act malevolently. The most notable are:
  - **Practical Byzantine Fault Tolerance (PBFT).** PBFT environments are based on a voting-round system where at least two-thirds of the network must agree on the validity of a transaction.
  - **Delegated Byzantine Fault Tolerance (DBFT).** In the DBFT environment network participants choose a delegation of participants that has to carry the burden of reaching the consensus for the network.

Architecture

Usually, when analyzing DLTs two main categories based on the network access control profile are taken into consideration:

- **Permissionless (or Public) DLTs.** In this environment, every network participant could work as a validator. This form of DLTs represents "the purest form of decentralized ledgers; without a centralized authority that manages the network, anyone can join it"[16] and relies on consensus mechanisms in order to validate the transactions within the network. The most famous Permissionless DLT is Blockchain, a DLT based on a data block structure where every time new information is validated it is added to the "chain" as a new block of it.
- **Permissioned DLTs.** In these environments, only a restricted pre-defined group of trusted participants could take part in the consensus mechanism. These environments are also divided into:
  - **Permissioned Private DLTs.** In these DLTS "the transaction validator role is in the hand of the Central Authority (typically the Network owner)"[17].
  - **Permissioned Consortium DLTs.** The role of validators in these configurations is in the hands of a small group of participants who act as trusted validators of the network.

Despite the taxonomy explained covering the majority of DLT networks is there the possibility of hybrid configurations that could present a mix of features from both permissionless and permissioned DLTs.

Smart Contracts

The core of the automatization of the functions embedded into a DLTs is represented by a family of computer protocols that under the satisfaction of an agreement execute some predetermined function, such as the payment of coupon or the exchange of a token, these protocols are the well-known smart contracts. Smart contracts play a crucial role within a DLT network, acting as its rule vault. In fact, the conditions embedded into smart contracts are the ones that enable the automation of the functions that rule exchanges and interactions within the network participants. To summarize, "a smart contract is a self-executing program that automates the tasks outlined in an agreement or contract; once executed, these transactions are traceable and cannot be reversed. Smart Contracts emerged as a solution to remove intermediaries in trustless third-party transactions"[18].
It is easy to figure out that smart contracts, and their nature of if/then condition protocols, could see their application in many markets, especially in insurance and financial ones where payments and claims are based on the trigger of pre-determined contractualized events.

Tokens

Tokens are the native digital assets exchanged among DLT network participants that can represent any form of value or claim such as a real-estate property or an IP, through the support of the conditions embedded in the smart contracts that foster their creation, issuing, and management. Tokens are composed of two main layers that determine the features and the rules of the token within the network and the token's characteristics:

- **Core Layer.** This specifies the logic and the rules that the token will be subject within the platform (e.g. the interoperability and Cross-Chain Functions or the token regulatory conditions).
- **Service Layer.** This contains the determining features of the token (e.g. ownership rights, issuance value, etc.).

Tokens could serve a vast variety of purposes and could represent a barely unlimited group of claims depending on both their design and reference platform. Considering so the token taxonomy distinguishing a vast plethora of possible tokens, following a brief view of the most famous ones:

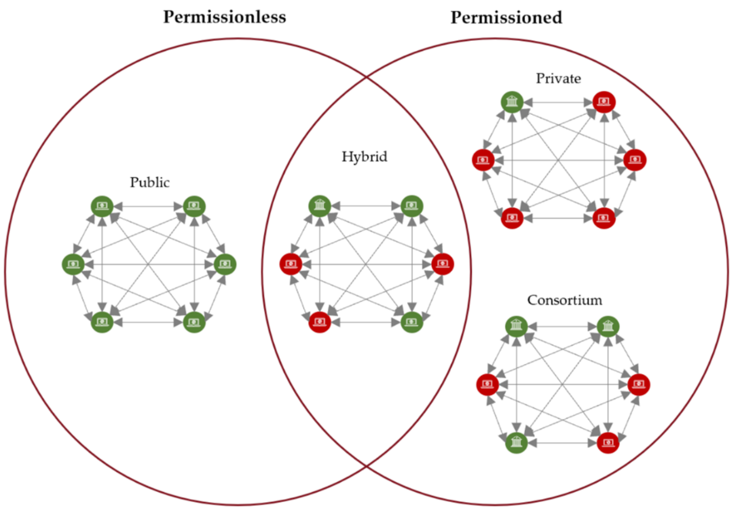

FIGURE 77: DLTs Architecures.

- **Cryptocurrencies.** Decentralized Digital currencies that rely on cryptographic algorithms to ensure the safeness of the transactions and the robustness of their network. Notable examples of cryptocurrencies are Bitcoin and Ethereum.
- **Security Tokens.** Tokens that represent a financial asset, such as equity shares, bonds, or real estate assets.
- **Stablecoins.** Tokens directly pegged to an underlying asset, usually a fiat currency, to which they ensure a par-exchange.
- **NFT (Non-Fungible Tokens).** This typology of tokens incorporates the ownership or the claim on a unique, real-world or digital, asset such as an art masterpiece.
- **Commodity Tokens.** Similarly to security tokens, these tokens are directly linked to the value of a particular commodity such as rare metals (e.g. gold, silver, etc.), oil, or other commodities (e.g. oat, pork belly,etc.).
- **Governance Tokens.** This typology embedded the voting rights that permit the network users to participate in the decision-making process of the decentralized environment.
- **Insurance Tokens.** These tokens "are designed to provide coverage against specific events and are typically used to create a kind of insurance contract into a DeFi Environment. They could also represent Tokenized ownership rights of insurance-related Assets, such as insurance policies or risk pools" [19].

Asset Tokenization

The process that allows a real-world asset, such as a real-estate asset or an art masterpiece, to be moved on a DLT is called asset tokenization. Technically speaking the process locks the real-world asset on its origin ledger as collateral for the tokens and, through a so-called "ramp", the asset is moved on the DLT. The process will guarantee the existence of the asset off the chain, while its related claims will be moved on the chain. The flow that will allow the tokenization of an asset is divided into six main phases:

- **Asset Selection.** The first step will involve the selection of the type of asset that should be tokenized. As seen before, asset tokenization could involve any real-world asset so choosing one or another typology of assets will affect every other step of the process influencing the valuation methodologies, the regulatory framework, the platform selection, and so on.
- **Asset Evaluation.** Starting from the asset that will be moved on the chain, this step regards first the choice of the evaluation methodologies of the assets then the potential demand analysis and other key evaluation steps such as the potential future revenues. To understand the importance of the right evaluation, it is important to state that discrepancies between the evaluation of the off-the-chain asset and its digitalized claims could lead to potential speculative and disruptive results.
- **Regulatory Analysis.** The scope of this process phase is to ensure that the right regulatory framework within the reference country for the chosen asset is identified. This also regards assessing the ALM (Anti Money Laundering), KYC (Know Your Customer), and data protection regulations that are in force in the reference country.
- **Platform Selection.** The core of the entire environment for the tokenized asset passes through the decisions that will be made regarding the platform and the technology on which it will rely, in particular, during this step, it will be defined if the native platform will be permissioned or permissionless, blockchain-based or not. The right platform and its configuration will directly affect several parts of the token distribution determining the transaction scalability, costs, and typologies (e.g. cross-boarders transactions).
- **Smart Contracts Development.** The rules and logic of the environment will be embedded, as already stated, in the smart contracts, which are going to be in charge of deploying the necessary actions for the token issuance, transaction management, determinate governance, and voting rights. Practically, this phase will define the rules on will govern the platform.
- **Token Creation and Issue.** The token creation and the following issue will end the process. At this stage, mostly all the decisions have been taken but this final part will be as crucial as the others defining

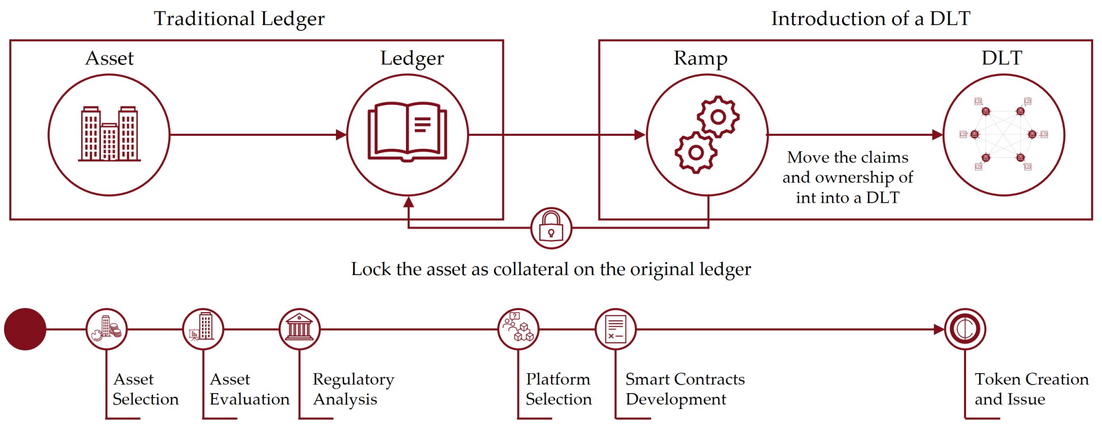

FIGURE 78: Asset Tokenization Process.

the property tokenization model opting for fractional ownership of the asset (and the number of tokens that will be issued) or for the individual ownership of the asset. The last part of the process will be the determination of the offering structure of the tokens, with the decision between an initial public offering (IPO) or a private placement.

Asset tokenization could foster the development of the economic system through several potential benefit effects. First, the fractionality of the ownership that could be reached through asset tokenization could lead the providing liquidity in historically un-liquid markets, such as real-estate markets or art markets. The possibility of dividing the ownership of an undividable asset could bring more retail investors into markets where they were excluded by the high barrier costs, boosting the liquidity of these markets.The market's efficiency could also be supported and boosted by the automatism guaranteed by smart contracts that could enforce faster and frictionless transactions. For instance, "a representative case of these possible improvements is the atomic settlement condition that could be coded into a smart contract that will permit an instantaneous exchange between two tokens once both parties submit their transaction"[19] Also, the lack of a central authority could indeed bring a reduction in transaction costs while the immutability of the ledger records could foster a more transparent environment, and the consensus mechanisms will ensure the safety of the network.

Decentralized Finance (DeFi)

The possibility unlocked by the scalability of programmability and the rise of DLTs has made possible the creation of decentralized environments that offer different financial services without the need to rely on financial intermediaries, these "ecosystems" are the DeFis. Decentralized finance environments leverage the power of smart contracts to establish protocols offering various financial services. The DeFi DLTs Network and the automatisms generated by smart contracts make it possible to have financial services offered that directly link the network participants without the need for "middlemen". The technical composition of a DeFi is correctly described by the DeFi Stack Reference Model (DSRM)[1]. The model identifies three main layers:

- **Settlement Layer.** The basement of the DeFi that is made by the native DLT on which the network is built and defines the transaction execution and validation rules ensuring the updating of the nodes and the safety of the environment.
- **Application Layer.** This layer defines the protocols and the services that are offered within the network, including also the native crypto asset that is exchanged within the environment.
- **Interface Layer.** The end-user interfaces.

Insurance DeFi

For the research purpose, it is important to enlighten the features of the so-called insurance DeFis. These environments rely, as seen before, on the DLTs and the automation of smart contracts to offer a wide plethora of decentralized insurance services. Two main families of insurance DeFi can be recognized in the market:

- **Traditional Insurance-Like DeFi.** Relies on the potential of DLTs to offer traditional insurance products substituting insurance agencies with a decentralized environment
- **DeFi's Risks Coverage Insurance DeFi.** Shields the network participants from risks directly embodied in a decentralized digital network such as hacking events or other technical issues.

Insurance DeFis exploit DLTs and smart contracts to create a liquidity pool, powered by the crypto asset flows of the liquidity provider, that serves as the monetary coverage for the risks that want to be covered by the insurance buyer. Usually, an insurance DeFi environment works as follows:

- **Liquidity Pool.** The liquidity of the environment is fed by users, the liquidity providers, that deposit funds (e.g. crypto assets) into a smart contract. Usually, the liquidity providers will receive a pool token that will represent the share of their participation in the pool. The liquidity pool represents the vault that will be used to cover the payment of the insurance claims to the buyers.
- **Insurance Offer.** The insurance DeFi could cover, as said, specific different risks, that range from smart contract failures to flight delays, to which the pools are linked. The liquidity providers choose which risk they want to cover with their funds and the smart contract associated with the pool will manage the eventual payment only on the designed specific risk
- **Price Evaluation.** The price that will be required to be insured is calculated through the support of smart contracts. The models applied are usually based on:

- The evaluation of the probability of occurrence of the risk on the duration of the contracts;
- The liquidity of the pool at the moment where the coverage is required.

Automated Market Maker Models (AMM) are the most famous price algorithms in the DeFi landscape, including the insurance DeFi. These models calibrate the price of the insurance, and the most famous one is Constant Product AMM. The Constant Product AMM requires that the value of the product of two variables remains constant during time, specifically, the following equation must be always satisfied:

$$x \cdot y = K,$$

where:

- **x** represents the total liquidity of the pool;
- **y** represents the coverage demand, as the total value of contracts in charge in a moment t;
- **k** represents the constant value.

If the demand for coverage rises the liquidity available to cover claims decreases, so to maintain the constant relation, the insurance premium increases. On the other hand, if the demand lowers the liquidity of the pool rises and the price for coverage falls as well. These models allow the definition of the prices through the simple mechanisms of offer-demand fostering the transparency of the pricing and avoiding intermediaries' arbitrage. Despite their simplicity, Constant Products AMM defines insurance prices by assessing the true liquidity risk of the environment. On the other hand, these models suffer from impermanent loss risk, which is defined as the loss generated by the significant movements in the liquidity pool deposited asset's prices. In these cases, the withdrawal of the assets could generate a loss in value for the liquidity provider's respect if they have detained their assets outside the pool.

- **Revenues for the Liquidity Providers.** Liquidity providers' revenues usually come through two channels:
  - **Insurance Premiums.** Based on their share of the liquidity pool the providers receive a proportional part of the revenues generated by the insurance's premiums.
  - **Interests.** DeFi environments usually pledge the liquidity pool funds (or a part of it) in other activities, to create passive revenues for the liquidity providers, for instance utilizing practices such as:
    * **Yield Farming.** The funds of the providers are pledged into other DeFi environments to support other financial services that generate passive incomes for the liquidity providers.
    * **Stacking.** The funds of the pool are stacked into a PoS (Proof of Stake) to receive the revenues for participating in the validation process of another DLT environment.
- **Insurance Payments.** If the event covered by the insurance occurs usually smart contracts provide automated payment to the insured. Smart contracts rely on oracles to determine if an event occurs. Oracles are an external party that permits the linkage between smart contracts and real-world data, their function is to permit the smart contracts' algorithms to determine if a claim should be paid or not. For instance, considering a DeFi that covers the risk of train delays, the oracles will link the smart contracts to the train schedule of the train company permitting them to define if a train is on late or not. Usually, we can distinguish between two typologies of oracles:
  - **Inbound Oracles.** Which take real-world data and transmit them to the chain.
  - **Outbound Oracles.** Which take information from the chain and transmit it outside of it.
- **Governance Tokens.** The liquidity pool providers are usually rewarded with governance tokens that enable them to participate in the management of the DeFi environment and to definition of its rules.

The application and development of Insurance DeFi could boost the existing insurance market thanks to some of their typical features, in particular:

- **Costs Reduction.** The development of a decentralized environment allows the reduction of agency and transaction costs typically associated with insurance contracts (and, more broadly, with all financial agreements). Additionally, the use of smart contracts and oracles significantly accelerates the refund process lowering the perceived recovery costs for policyholders. Traditionally, the time required to access these funds can be lengthy, often involving extensive bureaucratic procedures. However, the speed of transactions enabled by DeFi insurance will bring the immaterial costs of fund recovery close to zero.
- **Acces to the Market.** The reduction of costs, in general, will help to open up the market to underinsured or uninsured consumers, boosting both the financial inclusion of the economic system and the market growth.
- **Transparency.** In an insurance DeFi environment the insurance conditions are directly embedded into smart contracts meaning that there cannot be unilateral modification of the policy terms. Smart contracts also remove opacity in the policies pricing system to guarantee a fair price based on the true risk of the platform and/or of the event ensured (depending on the pricing model).
- **Reduction of Systemic Risk.** The total risk of the DeFi is distributed all along the liquidity providers, reducing the risk concentration, which could be sensibly higher in a traditional centralized environment, and making the system more resilient to adverse scenarios.
- **Flexibility and Innovation.** Relying on smart contracts it is possible to define coverage for any kind of scenario and events, boosting the possibility of widening the plethora of insurance events covering new market niches and deploying more tailored services that could attract more users to open an insurance contract. The flexibility guaranteed by smart contracts could also support the creation of new innovative solutions to cover more specific risks with the possibility of creating new markets.

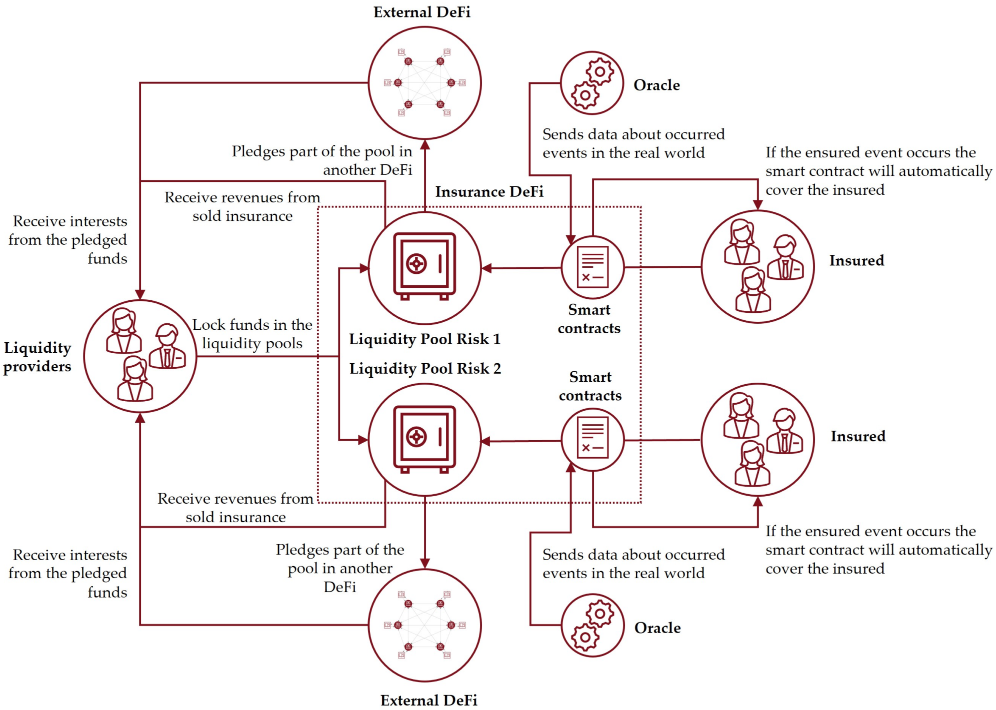

FIGURE 79: Insurance DeFi Enviroment.

Open Insurance

Open Insurance could be defined as accessing and sharing insurance-related data and services between all the actors involved in the process (consumers, insurance, third parties) to build better applications and services. This process could encourage consumers to compare the products in the market and to buy the best for their own interests; at the same time, this could push the growth of new insurance products and services tailored to the needs of the consumers.
Summarizing what we introduce in the above description, we can understand that Open Insurance refers to the exchange of data between consumers and third parties. This will potentially lead to positive effects on the growth of insurance markets, for instance through the development of new products and services or the reduction of information asymmetry'costs. So, the key aspects of Open Insurance are:

- **Data Sharing.** Insurers and other stakeholders share customer data, under their consensus, securely allowing for personalized and innovative insurance products and services.
- **Customer-Centric.** Customers gain more control over their data and can benefit from tailored insurance solutions that better meet their needs.
- **Innovation.** By opening up their data and services, insurance companies can foster innovation, leading to new products and services that can enhance customer experience and efficiency.

The services of Open Insurance are now offered mainly using APIs (Application Programming Interface), where insurance and third parties (such as big tech companies or other players) exchange data of their consumers. The first step in the Open Insurance implementation through API is the so called API Layer. This means making all company data and services available via APIs to the ecosystem of partners and company channels. Companies require two additional layers to address the challenges of Open Insurance through API:

- **The Data Management Layer.** That serves as a decoupling mechanism that collects data from the underlying systems and ensures its availability 24/7, regardless of the systems' availability.
- **The Business Logic Layer.** Where companies develop, reuse, and enhance microservices, as well as potential frontend components, which can be combined to generate unique services for partners.

An example of the actual structure of Open Insurance built through API can be represented by the example of car insurance reported in EIOPA's discussion paper on Open Insurance[4]. Assume having an insurance dashboard that could collect:

- Consumer existing insurance policies;
- Product information of the insurance companies/intermediaries who agree to use the dashboard.

The consumer can decide that his data could be visible only to the companies with which the client has a specific contract or can be visible to all the players belonging to the dashboard. In this second way, when the consumer needs to renew his car insurance, he can see and compare all the products offered by the insurance companies and make an informed choice. All the exchange of data between the client, the dashboard, and the insurance companies is performed through APIs. In Figure 80 is portrayed how the process works.
After this brief description of Open Insurance, we can highlight some advantages and disadvantages.
Starting with the benefits we have:

- Reduction of the information asymmetry between the insured and the insurer with a potential reduction in the information costs and a consequential decrease of policy prices which could lead to a widening of the potential individuals that could underwrite a policy;

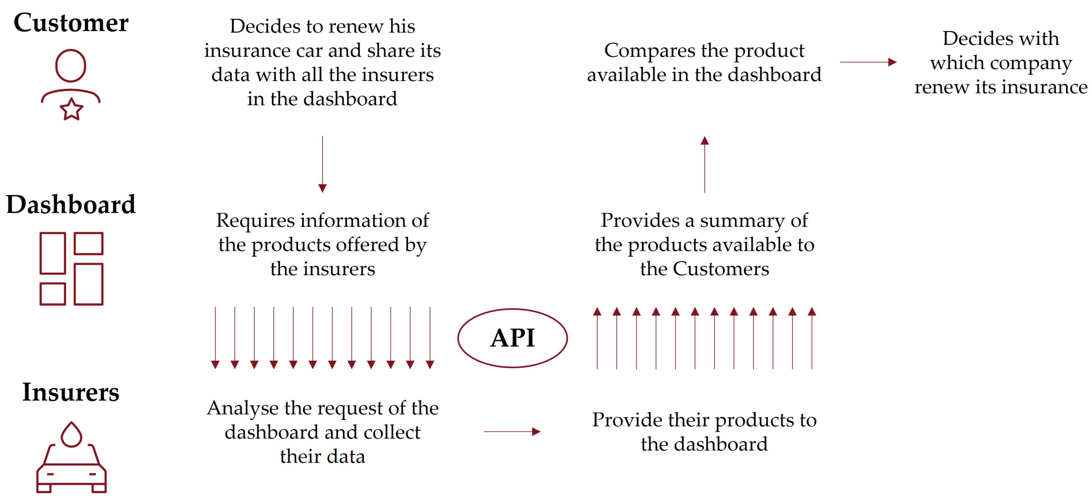

FIGURE 80: Open Insurance.

- More tailored products related to specific events (e.g. travel insurance when the consumer buys a train ticket; the proposal of insurance policy when the consumer reaches a certain age, etc.);
- The comparison between different companies and insurances is simplified. Moreover, also the change of insurance and product could be simplified by the comparison between more products and the choice of the product that better fits the interest of the consumer.
- The availability of new services is more specialized due to the needs and the situations in which they are offered to the customers. Having more data for their analysis, insurance companies can provide a better service due to the situation in which the customer needs of an insurance service.
- Enhanced competition between all the players to provide better products. Indeed, as stated before, the client has simplified opportunities to compare all the insurance's products, so the companies have to attract the interest of the client by proposing products that are better and more tailored to the needs of the customers.
- Development of new business models. Indeed, the insurance industry could implement new ways of studying the new data shared by the customers and new ways to offer their products; at the same time, the customer could receive and evaluate different proposals by simply comparing them in a unique environment.

There are also some risks related to Open Insurance:

- Data security and privacy, indeed all the parties involved in the process could get unauthorized access to data about the consumers.
- Common risks of the digitalization process, such as not being properly advised in all the steps of the subscription of a contract and receiving an incredible amount of offer based on own interests.
- The concentration of data of high quality could prevent the entrance into the market of new players narrowing the offer to the players who are already embedded and widely diffused.

All these risks are characterized and encouraged by the lack of a regulatory framework. This absence of a specific regulatory framework governing sharing data systematics in the insurance sector ensures that all Open Insurance projects are left to private initiative and are therefore implemented with special contracts stipulated by the Subjects involved. This has been highlighted by EIOPA[4] in his work and, despite this absence, the common sentiment is that Open Insurance will continue to expand in the following years. This absence of regulation underlined in Open Insurance is contraposed to the regulation born in the last years in Open Finance. Open Finance expands on the principles of Open Banking, extending the concept of data sharing and interoperability beyond just banking to encompass a broader range of financial services. This includes not only banking but also investments, pensions, insurance, and other financial products. The main goal of Open Finance is to create a more inclusive, competitive, and innovative financial ecosystem. Open Finance regulation in Europe is largely driven by the principles and frameworks established for Open Banking, particularly through directives such as the Revised Payment Services Directive (PSD2). However, as the concept of Open Finance expands beyond banking, additional regulatory considerations and frameworks are being developed.
The key regulatory framework is constituted by:

- **PSD2 (Revised Payment Services Directive).** Emphasizes strong customer authentication (SCA) and explicit consent, ensuring data privacy and security.
- **GDPR (General Data Protection Regulation).** GDPR's principles of data protection, transparency, and user consent is crucial for Open Finance, ensuring that consumer data is handled securely and ethically.
- **European Data Strategy.** Promotes the development of common data spaces, including finance, where data can be shared and utilized to foster innovation.
- **Digital Finance Strategy.** It supports the development of Open Finance, emphasizing the need for regulatory frameworks that facilitate data sharing across various financial services beyond banking.
- **MiFID II (Markets in Financial Instruments Directive II).** It enhances transparency and investor protection, which are important for the integration of investment services within the Open Finance framework.

The regulatory landscape for Open Finance in Europe is evolving, building on the foundations of PSD2 and GDPR while expanding to encompass a broader range of financial services. The focus is on creating a secure, transparent, and innovative financial ecosystem that benefits consumers and promotes competition. As Open Finance develops, ongoing regulatory adjustments and collaborations among EU institutions will be crucial to address emerging challenges and opportunities. Given that also the regulation on Open Insurance will start to grow according to Article 1(6) of the Regulation establishing the EIOPA (Regulation (EU) No 1094/2010)19 requires the EIOPA to contribute to promoting a sound, effective, and consistent level of regulation and supervision, ensuring the integrity, transparency, efficiency and orderly functioning of financial markets, preventing regulatory arbitrage and promoting equal competition. In addition, Article 9(2) requires the EIOPA to monitor new and existing financial activities. The above is the key motivation underpinning EIOPA´s work on digitalization. As stated before, Open Insurance is a field that is in continuous expansion and a process that can increase the development of this field is its integration in insurance DeFi environments. Indeed, the integration process of these two innovative paradigms could open up to several positive effects that could foster the growth of the insurance market. The analysis on the effects that integrating insurance DeFis and Open Insurance could have on the evolution of the insurance market will be analyzed in the final parts of the paper, while the next section will focus on illustrating some interesting cases of implementation of both insurance DeFi and Open Insurance environments.

Market Applications

This section of the analysis aims to illustrate some real-word examples of insurance solutions that rely on the technical frameworks that have been illustrated in the previous paragraph, trying to give a view on what are the potential unlocked by the technological innovation of insurance DeFis and Open Insurance.

Etherisc

Etherisc provides a complete suite of solutions to build, manage, and inspect decentralized insurance products. The essential utility token of the Etherisc ecosystem is the DIP token. DIP tokens give users access to the Decentralized Insurance Platform; by staking them, participants provide collateral for risk pools and guarantee future performance, availability, and service level. A risk pool, as previously discussed, is a smart contract that aggregates ("pools") various risks, represented by policy objects, and links them to risk capital. These pools gather collateral from investors contributing to capital rise. Risk investors stake DIP tokens and/or stablecoins into the pool, locking their assets in exchange for a reward. If a loss occurs, payouts are made from the risk pool, putting the stablecoin portion of the pool at risk. Investors have the flexibility to increase their investments in the pool or withdraw funds. However, before withdrawing, any associated risks must either expire or be settled. The Etherisc ecosystem is based on three pillars:

- **Risk Trasfer Market.** Investors will lock a certain amount of DIP tokens. The staked DIP token is a prerequisite to investing the actual risk capital in DIP or stablecoins. It is demanded that parties who profit from the ecosystem also own a share by owning and staking DIP tokens.
- **Regulatory Framework.** Insurance companies are highly regulated worldwide to protect customers as well as investors. The legal framework must be considered for each project, product, and jurisdiction, and the product owner is responsible for the proper implementation.
- **Techincal Framework.** The GIF developed and maintained by Etherisc allows for modeling, deploying, and operating insurance products based on blockchain in a decentralized and transparent way. Using GIFs, interested parties may quickly implement and securely operate their insurance products.

GIF stands for Generic Insurance Framework and consists of a collection of open-source smart contracts that implement essential functions of the lifecycle of insurance products and policies. Thus, GIF enables the modeling of a wide variety of insurance types. It is a basic implementation that can be used to create blockchain-based insurance applications. To be able to design insurance products quickly and easily, processing steps that run similarly in all products have been identified and made available as modules. Thus, only product-specific aspects, such as pricing models, need to be implemented for each product. To operate insurance products, including selling policies, collecting premiums, calculating trigger events, and handling payouts, a complete execution environment is needed in addition to the smart contract collections that define products and policies. GIF provides these generic functions for all sub-steps in the lifecycle of an insurance policy, thus enabling an automated workflow that controls the sequence of processing steps. The stakeholder roles in Etherisc are the following:

- **Insured/Customer.** Policyholder who wants to pass his risk to the risk pools. He is a customer of the insurance company.
- **Investor.** Investors have an interest in participating in risk pools to balance/diversify their risk portfolios. They provide collateral for risk pools in exchange for interest payments.
- **Oracel Owner.** Provides oracles that interface between the blockchain smart contracts and external data sources. For example, in the case of flight delay insurance, the oracle informs the smart contract whether the flight landed in time, how much it was delayed, or if it was canceled entirely.
- **Product Owner.** Designs and operates one or more products. This would be an insurance company or an MGA (Managing General Agent) in the traditional insurance industry. Due to the multi-client capability, a product owner can use all oracles located on the respective platform by the oracle owners.
- **Risk Pool Keeper.** Manages one or more risk pools.
- **Instance Operator.** Key role that operates a specific GIF instance.

Any instance of the GIF maintains:

- **Products.** A product is a specific smart contract that implements the functionality of its specific requirements, or it can use the generic functionality of GIF. After the product is technically developed and deployed to the blockchain, it must be registered in the GIF instance. This action is typically integrated into the deployment process.
- **Oracles.** Oracles form a vital part of the GIF, as they link the blockchain-based smart contracts and the index/parameter information necessary to operate real-world insurance products. Products can utilize product-specific oracles, but they can also make use of generic oracles, which can, in turn, be implemented by many different parties. For example, the FlightDelay rating oracle has one input parameter, the carrier/flight number combination, and one output parameter, an array of integers that represent the

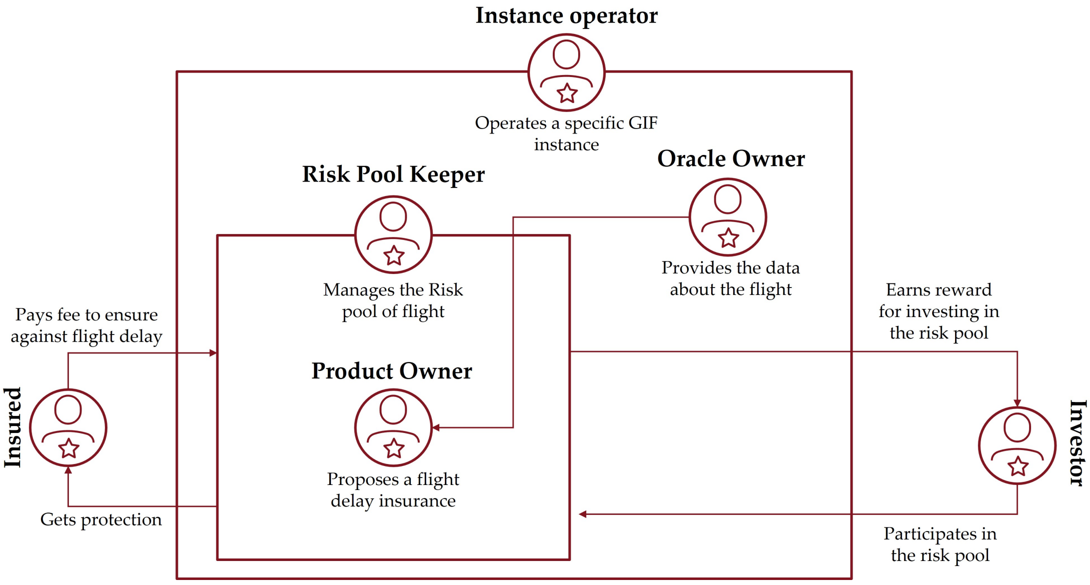

FIGURE 81: Etherisc.

historical number of delays for different amounts of delays.

- **Risk Pools.**

The GIF instance is agnostic about the way payments are made. Therefore, Etherisc does not offer specific functionality for this.

Lemonade

As described in the section dedicated to the main trends in global insurance markets, climate change-related risks represent a significant driver in the future business scenarios of insurance companies. The greatest challenge is to provide adequate coverage for climate risks without undermining the financial stability of insurers. In recent years, several initiatives have emerged in the insurance sector aiming to explore the sustainability of products covering climate risks. For the purposes of this paper, we consider the case study of the Lemonade Crypto Climate Coalition to be of particular interest for two reasons: the first relates to the blockchain technology used by Lemonade in this project, and the second is the high social impact that this initiative aims to achieve. Founded in 2022 by Lemonade Foundation, the non-profit arm of Lemonade Insurance Company, LCCC has the ambitious mission to use exponential technologies like blockchain to create social value in the real world. The pilot project was launched in Kenya to offer protection to smallholder farmers who, despite being the most exposed to extreme weather events such as droughts and floods, do not have access to insurance coverage products in the traditional market. In Sub-Saharan Africa, crop insurance is not very widespread; less than 3% of farmers have access to insurance products, according to the latest ISF report.[12] This happens because, for traditional insurance companies, offering products in the most rural areas is extremely costly and unsustainable: underwriting is very complex, and with low premiums due to limited demand, it is difficult to find agents willing to distribute their products. Additionally, the payout from the claim process is often lower than the costs incurred for the process itself. Last but not least, assessing and processing claims in the traditional insurance business model involve extensive documentation and verification processes which are not always feasible in rural contexts.
By the combination of Blockchain Technology, Data Intelligence, and Insurance expertise, Lemonade is trying to close this insurance protection gap. Using DAO (Decentralized Autonomous Organization) instead of a traditional corporate structure, smart contracts instead of policy documents, Oracles instead of claims adjusters, and providing services at cost instead of looking for a profit, LCCC managed to take a product that is unsustainable to something that is infinitely scalable; once the blockchain infrastructure is built, it can be easily replicated in any other market in the world.
As previously mentioned, key players in the blockchain market such as Avalanche and Etherisc are part of the coalition. The former provides a technically sustainable infrastructure from an environmental standpoint, while the latter offers blockchain-based insurance solutions and the user interface for the project. The enabling paradigm to provide adequate risk coverage to typically "uninsured" populations is, once again, the reduction of costs associated with the claims management process. The key to lowering costs lies in the implementation of smart contracts, programmed to automatically trigger payouts when predefined conditions are met, verified by external data sources provided by Chainlink. This ensures timely support for the farmers enrolled in the program. To overcome the challenges posed by the complexity of blockchain in a context with limited technological and financial literacy, the policy subscription process has been made as simple as possible: the entire onboarding takes no more than 30 seconds and is done directly from the farmer's phone via a text-based interface, as the vast majority of farmers in the region do not own smartphones but feature phones.
Once the farmer purchases insurance directly from his phone[47], at the end of the harvest season, the LCCC collects parametric data on weather and crop yield, which is then transferred to an Oracle on the Avalanche blockchain. This data triggers smart contracts and eventually sends money directly into the mobile wallet linked to the farmer's phone. At the moment these smart contracts are essentially re-insured by the Lemonade Foundation which basically

[47] Purchases are made through M-Pesa mobile/wallet, a mobile banking service, particularly diffused in Sub Saharan Africa, that allows users to store and transfer money through mobile phones.

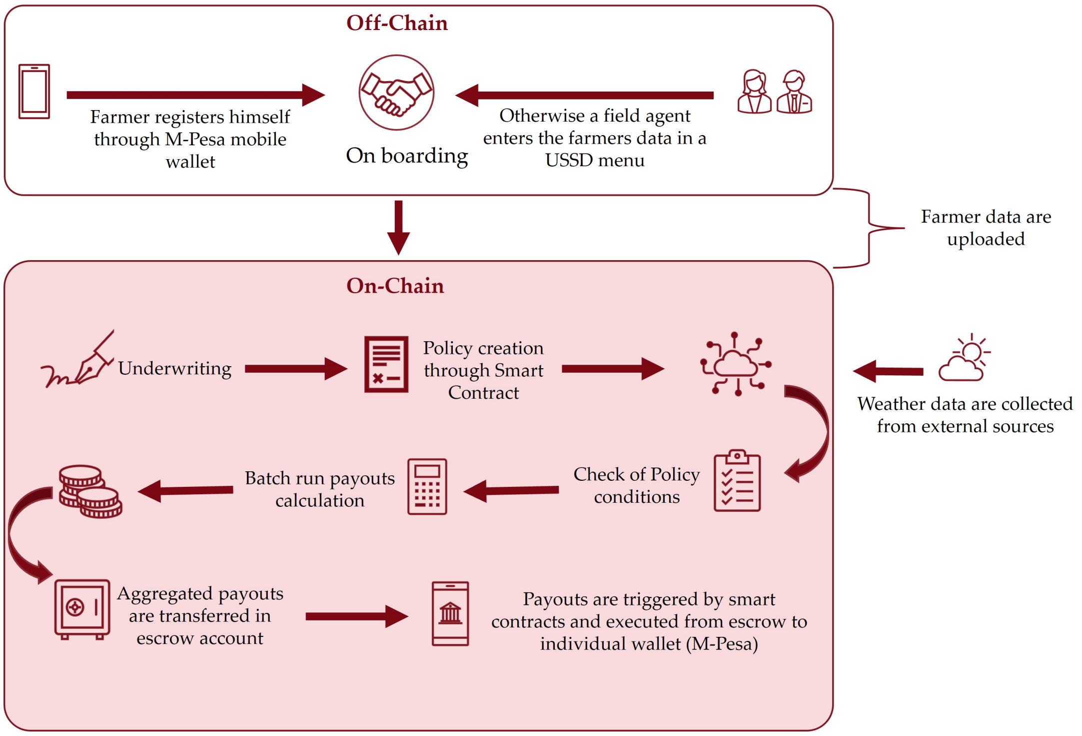

FIGURE 82: Lemonade.

provides funding for it, this is also because the first year saw a loss ratio of 173%, with payouts exceeding the premiums collected (the average premium is 5$ per farmer); thus, the greatest challenge for the project will be to make the model financially sustainable on its own. In its first year, the program collected premiums of 98,000$, reaching over 6,400 farmers, but with a total payout of 170,000$. These figures demonstrate the significant need for coverage in these communities, which are exposed to a real risk of food insecurity due to the unpredictability of weather events. A possible step toward market integration could be the issuance of a token, allowing retail investors to participate in the risk pool and help provide insurance coverage to small-scale farmers in emerging countries.

Solace

As noted by EIOPA[4] one of the most widespread applications of Open Insurance, although still in its early stages, in the European insurance market is claim management. The survey reveals that insurance players are preparing to leverage the efficiency improvements that APIs can bring to traditional claim management systems. It is not surprising, as improving claim management processes, in addition to reducing the overall costs of insurance products, would represent a significant enhancement of the user experience for policyholders. In fact, according to recent studies, more than 80% of customers surveyed considered changing their insurer after a negative experience during the claim submission process when the covered event occurred. In this context, APIs can be the enabler of new claim management models, particularly by leveraging applications based on event-driven architectures (EDA). An example is the PubSub+ solution created by Solace. Solace is a Canadian IT company founded in 2001, and a leader in providing event management platforms for clients across various sectors (e.g., capital markets, aerospace, retail, automotive).
Event-driven architecture is a software design pattern that sends messages based on specific events, in IT terms events represent a status change in data such as a field changing in a database, a bank deposit being completed, or a checkout button being clicked in an e-commerce app. There are three main elements in event-driven architecture:

- **The Publisher.** The entity that sends or publishes a message (also called a producer);
- **The Event Broker.** The message distributor from the publisher to the subscriber;
- **The Subscriber.** The message receiver.

In this architecture, the message represents the information that the publisher wants to send. Messages often contain event data, but can also carry queries, commands, and other information; in an event-driven architecture, a message typically has a destination that distinguishes the publisher from the subscriber.
The key advantage of EDAs compared to traditional APIs (also known as REST APIs) is that the subscriber only needs to "subscribe" to the event notification, specifying the endpoint that the publisher can call to send the message once the event has occurred and been confirmed. In contrast, with REST APIs each request is sent by the consumer to the provider, generating multiple potential negative responses until the event occurs. It is easy to understand how an event-driven architecture requires less processing and storage, while also reducing the time and costs associated with the event notification process.
Although the main documented applications of Solace's event management systems so far concern banks and capital markets, particularly in the development of data distribution layers for high-speed trading, we believe that insurance companies could leverage this architecture to significantly improve customer experience during the claim process. Let's take a closer look at how Solace's application works, to outline a potential claim management process for a health insurance company. We previously mentioned the endpoint that the publisher can call to send a message once the event has occurred. In the specific case of Solace PubSub+, these

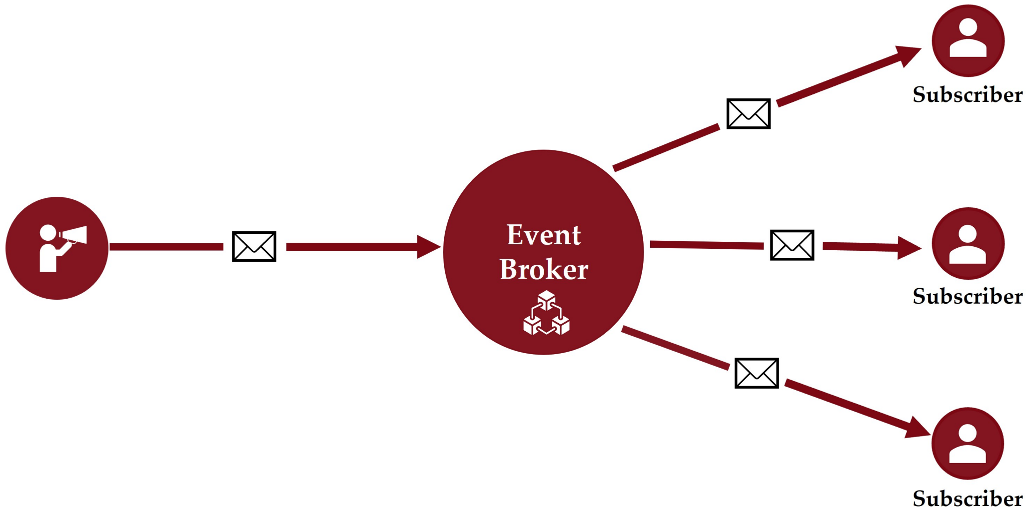

FIGURE 83: EDA's elements.

endpoints are, in turn, accessed via so-called topics. From a technical point of view, Solace topics are simply strings composed of one or more levels added as metadata in a message header that let publishers classify messages and let subscribers specify what they want to receive messages about. This publish-subscribe model enables event brokers to use topics as routing information to send event messages everywhere they need to go. On the consumer side, when an application connects to PubSub+ and specifies a topic subscription, the Event Broker maintains it in a list of all subscriptions for all clients. A Solace topic is a string with the format a/b/c/.../n where a/b/c and so on are levels in a hierarchy of information used to describe events; one or more topic levels can be defined using variables that are replaced with properties specific to the event when it occurs. Let's suppose we want to design an event broker that allows the insurer to speed up the claim management process in the specific case of health insurance. In this scenario, one could implement an event-driven messaging system between the hospital's TPA (Third Party Administrator) and the insurance company's underwriting service, leveraging the scalability of Solace PubSub+. In this setup, the hospital would take on the role of publisher, and the insurer would act as the subscriber. In the traditional claim process, the insured retrieves the claim form from the hospital's TPA or the insurer's website, fills it out with all the necessary information, collects the required documentation from the hospital (e.g., medical records, prescriptions...), and then submits the claim to the insurer. With the use of PubSub+, the hospital's TPA, acting as a publisher, can create specific topics to which, potentially, multiple insurers can subscribe, receiving timely information about the occurrence of the insured event almost in real-time. Let's consider a typical insurance coverage that offers reimbursement for illness or injury expenses in the case of hospitalization or surgery. In this case, the Publisher can create a topic with a hierarchical structure to standardize the messaging related to hospitalization events, such as the string shown below:

Claim/[Hospitalisation]/[Policy Number]/[Hospital]/
[Certificate Number]/[Hospital Bill]

The brackets "[ ]" indicate a variable in the topic that is then replaced in the message by data specific to the event.

- **Claim.** Indicates the type of event.
- **Hospitalisation.** Indicates the action that generates the event.
- **Policy Number, Hospital, Certificate Number, Hospital Bill.** Are properties of the specific event.

In this type of application, the insurance company will simply need to define its topic subscription to capture the published events that are relevant to it. The event broker will then route topics without deserializing, decoding, or interpreting the event. When the event occurs, the TPA publishes the message, including in the body all the information required by the insurer to verify the eligibility for claim acceptance, such as receipts and medical records. The event broker, upon receiving the message, reads the topic and routes it to all subscribers of the event almost in real-time. The message is then received by the insurer's claim service, which proceeds with the document assessment and promptly provides the outcome of the review to the policyholder (see process below). It is reasonable to assume that a further acceleration of the process described here could be achieved through the introduction of blockchain in event-driven architecture. Let's suppose that in Step 3 of the process, blockchain technology is used. Smart contracts can be deployed to minimize claims management expenses, mitigate claims fraud, and automate the verification of claims when the conditions specified in the insurance contract are met.

InsurAce

As mentioned in paragraph 3.1, in addition to the traditional family of insurance products, DeFi insurance finds another extensive application in covering risks arising from the use of blockchain technology itself. The case of InsurAce deserves further examination in this regard. As platforms offering services on blockchain grow, along with the increasing number of holders of so-called crypto assets, the demand for coverage of the risks to which crypto holders are exposed also rises. Consider that the TVL (Total Value Locked) in DeFi protocols as of October 2024 exceeds 86$ billion, of which only a small percentage is actually insured. InsurAce aims to close this protection gap by offering the possibility to purchase a wide range of mutual protection products for digital assets against various types of risks, below a non-exhaustive list:

- **Smart Contract Vulnerability Cover.** This type of product offers protection against any claimable loss

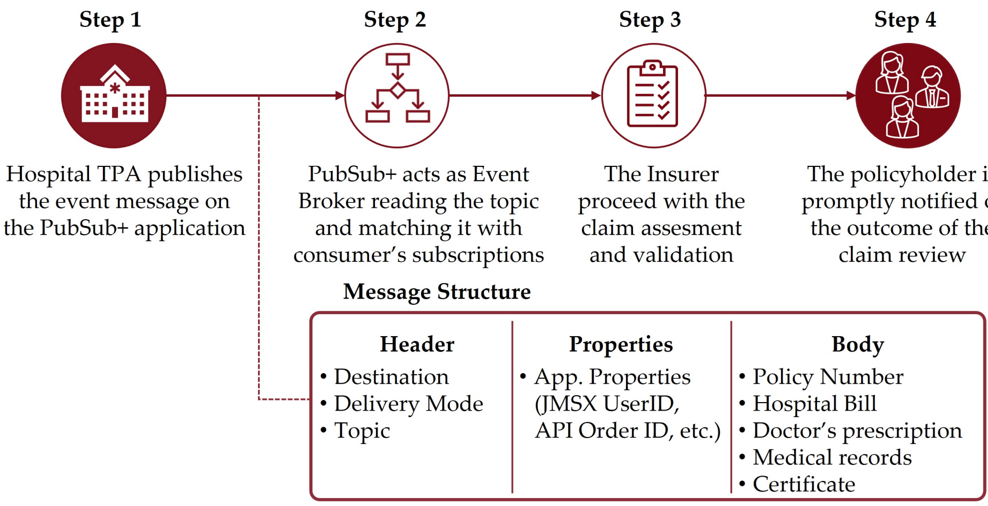

FIGURE 84: Claim Management Process.

due to malfunction, programming flaws, unauthorized, malicious, criminal attacks or hacks of the designated smart contract.

- **Stablecoin De-Peg Cover.** Refers to the risk that a stablecoin, which is designed to maintain a fixed value (typically pegged to a fiat currency like the US dollar), deviates or "depegs" from its intended value. For example, InsurAce offers a $USDT De-Peg Cover that compensates the policyholders for claimable loss realized in selling $USDT below the US$ 1.00 per $USDT peg between the claimable risk event and the claim deadline.
- **Ethereum Slashing.** coverage on a basket of different DeFi protocols, if they misbehave, a portion of their staked ETH can be slashed (confiscated).
- **Bridge Cover.** This kind of product aims to offer protection against the risk associated with using cross-chain bridges, which allow the transfer of assets and data between different blockchain networks. Claimable risk events protected are loss of tokens in transit to bridge malfunctions, hack or vulnerability exploits and loss of tokens in transit to error in slippage reported by bridge and/or DEX for tokens received at bridge or DEX on destination chain.

InsurAce's approach, based on covering the various risks associated with holding assets on the blockchain, allows it to offer policyholders comprehensive portfolio coverage on a basket of different DeFi protocols, creating a well-diversified risk management tool for investors.
Risk is currently shared in two mutual pools under the InsurAce protocol: cover payment pool and underwriting mining pool (with the latter offering higher returns to investors) which are governed by its members where membership rights are represented by the $INSUR token. More specifically, InsurAce's business model consists of two functional branches: the Cover arm that manages a low risk capital pool which helps maintain the InsurAce protocol's solvency and therefore ability to meet its cover obligations and the Investment arm which manages high risk investment pools which generates returns in order to finance possible claim payouts and attract investment capital; unlocked capital in the Cover capital pool may be transferred to the Investment pool to earn higher yields and subsidize users' costs from cover payments. According to InsurAce's official documentation, the project's key strength from a management perspective is its pricing model for the insurance products offered. InsurAce claims to provide more competitive prices compared to major competitors, who use models based on the staked value in their respective protocols (an inversely proportional relationship between staked value and insurance price), thanks to its use of actuarial pricing models. The price of the different coverages offered by InsurAce consists of two components: the base price, determined using an Aggregate Loss Distribution model, plus a dynamic component calculated based on the supply and demand volume for the insurance coverage. As for the base component of the price, it is the output of a model whose workflow is described in the figure below, based on the estimation of expected loss. The model's input factors include the number and amount of claims, the number and amount of exposures during a specific observation period, as well as assumptions about the level of inflation. These inputs are used to develop and train two distinct models: the frequency model and the severity model. The frequency model estimates the probability of a certain number of losses occurring within a given period, while the severity model generates the distribution of loss amounts and determines deductibles and coverage limits. Once both models are accurately calibrated, they are combined to calculate aggregate loss; the aggregate loss is then integrated with protocol-specific risk factors (such as bridge risk or de-peg risk) in order to calculate the actual base price.

Nexus Mutual

Nexus Mutual is a decentralized insurance platform that allows members to join and share risks. Members can purchase cover products that protect against different kinds of risk. The Nexus Mutual protocol is built on Ethereum and provides the infrastructure for members to buy cover, underwrite risk, assess claims, and build risk management businesses. The token on Nexus Mutual protocol is the NXM Token which is a governance and utility token backed by crypto assets held in the capital pool contract. Members can purchase cover products that protect against various types of risk:

- **Protection Against a Range of Loss Events Caused by Smart Contract Hacks/Exploits, Oracle Manipulation or Failure, Severe Liquidation Failures, and Governance Takeovers (Protocol Cover, Bundled Protocol Cover, DeFi Pass Cover, Native Protocol Cover).** These protections guard against

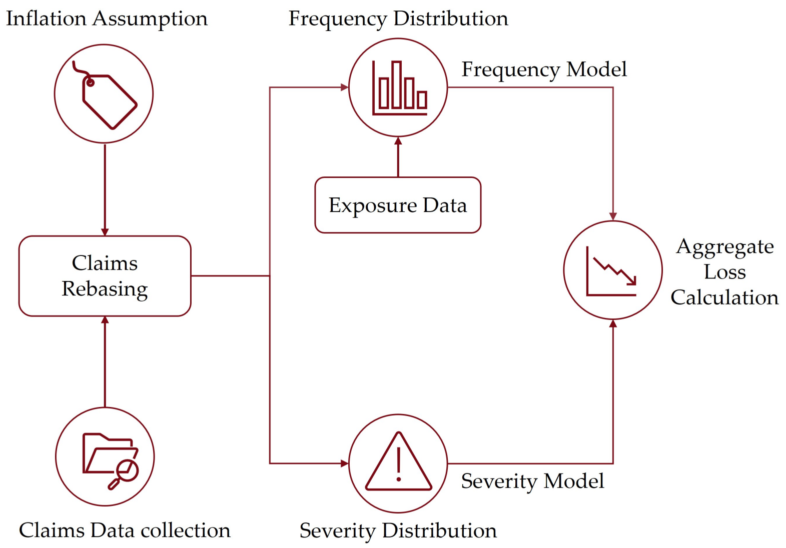

FIGURE 85: InsurAce's Aggregate Loss Model.

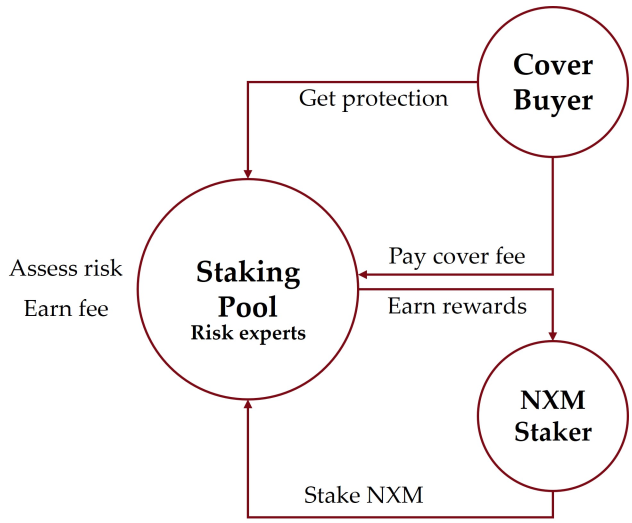

FIGURE 86: Nexus Mutual.

smart contracts exploits/hacks, severe oracle failure/manipulation, severe liquidation failure, governance attacks. For example, when you hold Protocol Cover and suffer a loss of funds, you can file a claim and claim assessors will review your claim submission to determine whether your claim is valid. The process works in the following way:

- If you hold Protocol Cover at the time the loss event occurs, you can submit a claim with supporting evidence, otherwise referred to as proof of loss. You will be able to include written details, links to supporting documentation, and/or upload screenshots or other files in the Incident Details section of the claim submission process. You will choose to either sign a message from the affected address or send a 0-value transaction from the affected address to prove your own and control the affected address.
- Claim assessors will review, discuss and vote to approve claims where proof of loss shows that you have indeed suffered a loss of funds. If your claim is approved, you will be able to redeem your claim payout after the 24-hour cooldown period passes in the Your Covers menu. You can also check your Dashboard to see the status of any active claims; if your claim is denied, you will be able to file another claim with more supporting evidence.

- **Protection for Underlying Risk Underwritten by a Cover Provider (Quota Share Cover).** On-chain cover providers that can protect their underwriting capital against the underlying risks they offer coverage for. When a covered organization pays out claims for an underlying risk that is included in the Quota Share Cover terms, they can file a claim in the Nexus Mutual user interface. Your organization will need to wait 14 days for the cooldown period to pass. The process works in the following way:
  - Your organization will file a claim using evidence of their payments for the underlying risks covered;
  - Claim assessors will review, discuss, and vote

to approve claims where proof of loss shows that your organization has indeed paid out claims for covered underlying risks above the deductible. If the claim is approved, your organization will be able to redeem the payout after the 24-hour cooldown period ends; if the claim is denied, your organization will be able to file another claim with more supporting evidence.

- **Protection Against Underwriting Risk in Traditional Markets Using On-Chain Capital.** This type of protection is expected to go live soon but currently is unavailable.

The Nexus Mutual DAO operates as a discretionary mutual where people are required to join as members before they can interact with the protocol. The process works in the following way:

- **Staking Pool Managers.** Members with risk and pricing expertise that create and manage risk-staking pools;
- **NXM Staker.** Delegates their staked NXM to risk experts that manage staking pools;
- **Cover Buyer.** Members which buy the protection paying a fee. The fee is distributed both to the staking pool managers and to the NXM Staker.

Integrating Insurance DeFis and Open Insurance

In the previous sections of the research, we have delved into the potential of both asset tokenization and Open Insurance in developing much more innovative and resilient insurance markets. This final section of the analysis will leverage on a simple model based on choice behaviors under uncertainty in order to assess the potential benefits for the insurance market that could be generated by developing solutions that leverage both asset tokenization and Open Insurance. First, let's assume an economic system where individuals are characterized by risk aversion and hold a capital endowment of K. Within the economic system, there is a probability that an event may occur with a negative effect on the capital endowment K of the individuals, which could be represented by the non-degenerate monetary lottery A (Figure 87) with:

- **K**= Capital Endowment;
- **L**= Loss.

The expected value of this lottery, EV(A) is defined as follows:

$$EV(A) = pK + (1-p)(K-L).$$

In this environment, individuals can buy an insurance policy offered by an agent that guarantees the full coverage of the loss L at the cost C. The payoff generated by the purchase of the insurance could be represented by the monetary lottery B_1 (Figure 88a).
The expected value of B_1 could be represented as:

$$EV(B_1) = p(K-C) + (1-p)(K-L-C+L);$$

$$EV(B_1) = p(K-C) + (1-p)(K-C);$$

$$EV(B_1) = K - C.$$

With the policy purchase, the individuals transform A into the degenerate lottery B_1. Now let's assume that in fixing the price C the insurance agent overturns the costs associated with the agency costs and the cost related to the information asymmetry of the contract to insured individuals. That considered, we can determine:

$$C = I + U + A,$$

where:

- **I** represents the risk premium required to be assured of a specific risk;
- **U** represents the price markup for the information asymmetry and the moral hazard risk;
- **A** represents the agency costs (e.g. bureaucracy costs).

Other than these, let's assume that the individuals bear a non-material cost R, which represents the claims processing cost for retrieving the insurance payout once the negative event occurs. Considering this, we can redefine the lottery B including the cost decomposition and R (Figure 88b) with:

$$EV(B_2) = K - (I + U + A) - R.$$

As assumed before, the economic system is characterized by risk-averse individuals which implies that they prefer a lottery only if its expected value is strictly greater than their certain equivalent. In order to have individuals underwrite the insurance policy the following equation has to be satisfied:

$$u(\overline{K}) < u(K - C - R).$$

With $\overline{K}$ = Certain equivalent of the lottery A
And so, considering that a VNM utility function is monotonously increasing:

$$\overline{K} < K - C - R.$$

At this point we have to introduce a new assumption: in this economic system individuals are characterized by different levels of risk aversion. This implies that at the price at which the insurance is offered C, only the portion of individuals that have $u(\overline{K}) < u(K - C - R)$ will purchase the insurance. Now, as shown earlier in the analysis, insurance DeFis could reduce the recovery cost R to zero, thanks to oracles and smart contracts automatization, while relying on a decentralized environment will prevent the agency costs from increasing the cost of policies. Considering this, the development of an insurance DeFi could reduce the costs for individuals, setting, for simplicity:

- **R**= 0,
- **A**=0.

With a new lottery T (Figure 89a).
With: $C_1 = I + A + U + R > CT = I + U$
and so:

$$u(\overline{K}) < u(K - C_1) < u(K - C_T).$$

The integration of the Open Insurance practices into a tokenized environment thanks to the sharing through API of the information of the individual's behaviors could also set to zero the costs related to the information asymmetry. Regarding this, we can define the risk premium modified by the integration of both asset tokenization and Open Insurance as:

$$C_{OT} = I.$$

With the new lottery OT (Figure 89b) and:

$$C_1 = I + A + U + R > C_T = I + U > C_{OT} = I.$$

These considerations will lead to the following relation:

$$u(\overline{K}) < u(K - C_1) < u(K - C_T) < u(K - C_{OT}).$$

The reduction of the costs required to ensure capital K against the loss L will bring even the more risk-averse individuals, that were cut off by the market, to underwrite

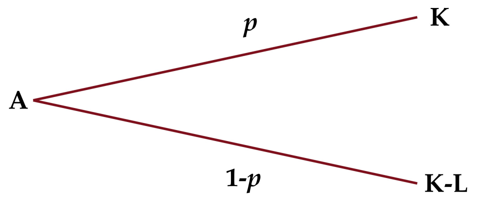

FIGURE 87: Lottery A.

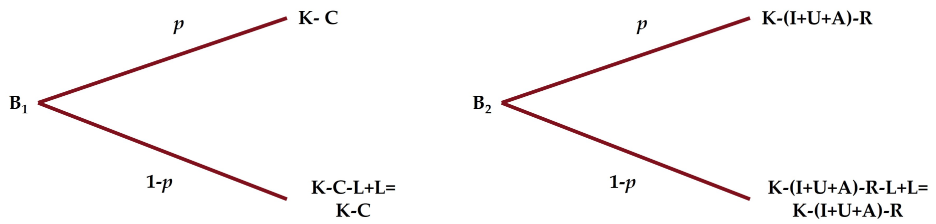

(a) Lottery B_1 **(b)** Lottery B_2

FIGURE 88: Lotteries B_1 and B_2.

the insurance policy. With a new lottery T (Figure 89a) with $C_1 = I + A + U + R > CT = I + U$ and so:

$$u(\overline{K}) < u(K - C_1) < u(K - C_T).$$

The integration of the Open Insurance practices into a tokenized environment thanks to the sharing through API of the information of the individual's behaviors could also set to zero the costs related to the information asymmetry. Regarding this, we can define the risk premium modified by the integration of both asset tokenization and Open Insurance as:

$$C_{OT} = I.$$

With the new lottery OT (Figure 89b) and:

$$C_1 = I + A + U + R > C_T = I + U > C_{OT} = I.$$

These considerations will lead to the following relation:

$$u(\overline{K}) < u(K - C_1) < u(K - C_T) < u(K - C_{OT}).$$

The reduction of the costs required to ensure capital K against the loss L will bring even the more risk-averse individuals, that were cut off by the market, to underwrite the insurance policy.
On the other hand, the insurers face the lottery E_1 (Figure 89c).
As already explained, insurers entirely shift the agency costs, and the costs related to information asymmetry onto the consumers. Considering this, the lottery E_1 could be transformed in the lottery E_2 (Figure 89d).
Thus, the profit function of insures could be described by the following equation:

$$EV(E_2) = W = N[p(I) + (1 - p)(I - L)].$$

With:

- **N**= number of individuals that underwrite the policy.

Considering that the insurers maximize their profits, and their profit function W is a growing function in N, it can be stated that the growth of individuals willing to underwrite insurance policies given by the introduction of insurance products that rely on both insurance DeFis and Open Insurance practices could also boost the total profit of the insurance market.
The reduction of prices, guaranteed by the integration of both insurance DeFi and Open Insurance and the consequent increment of individuals who are willing to be insured, will lead to the growth of the total demand for insurance policies. Consequentially, the effect of technological improvements in these new environments will bring a reduction of the costs that insurers will bear[48] ($C_1 > C_T > C_{OT}$); also, the new technology will potentially lower entry barriers with new players from other sectors that could be encouraged to break into the market. These factors could lead to an increment in the total supply of the market. This will permit it to reach a new higher economic equilibrium with benefits for the entire market. This model is an oversimplification of the real world but could help illustrate the possible effects that the introduction of innovative solutions - such as Open Insurance and insurance DeFis - could bring to the markets.

[48] In the model hypothesis we have considered that the insurer will completely shift its cost on the insurer, but in a more realistic environment there will be some costs that cannot be entirely shifted on the price of a good.

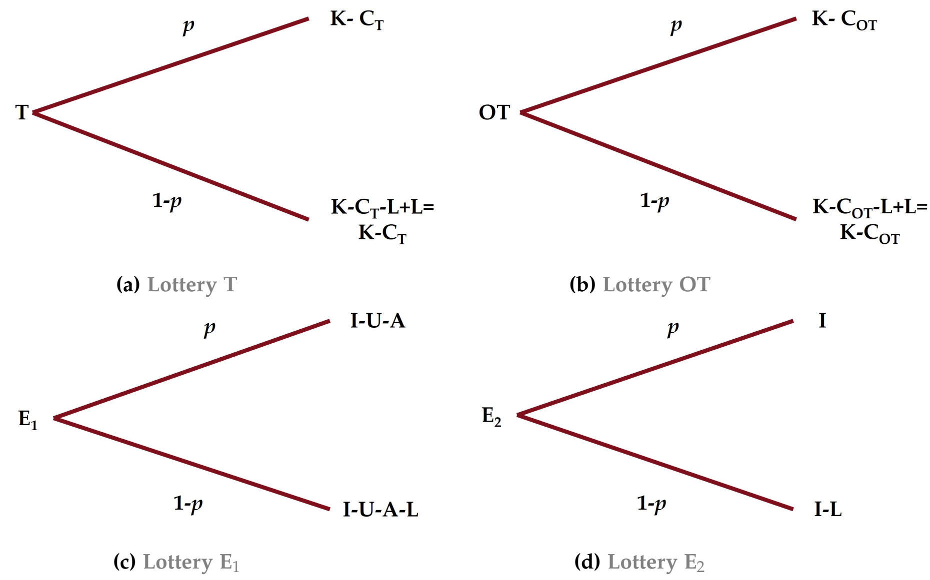

(a) Lottery T (b) Lottery OT

(c) Lottery E_1 (d) Lottery E_2

FIGURE 89: Lotteries T, OT, E_1 and E_2.

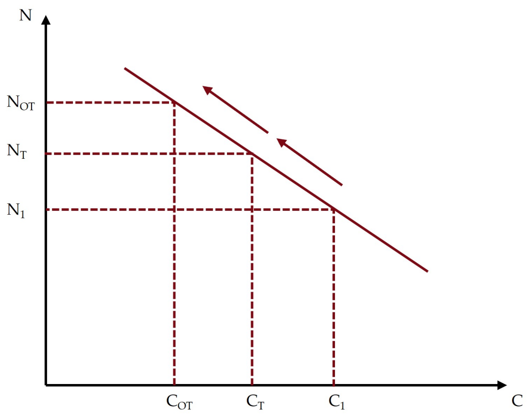

FIGURE 90: Increase of the Number of Insurend.

Conclusions

In conclusion, the analysis delved into the main features of insurance DeFis and Open Insurance, trying also to figure out the advantages that the embracement of these two paradigms could bring into the insurance markets. To summarize, we can highlight the following potential benefits of integrating Open Insurance and insurance DeFi:

- **Cost Reduction:** The integration of Open Insurance and DeFi insurance could lead to a cost reduction through:
  - **The Reduction of Manual Processes and Agency Costs.** Asset tokenization, thanks to smart contracts automatization, could eliminate manual processes and reduce the needs of intermediaries. Smart contracts, with Oracles' support, directly execute insurance policies when predefined conditions are met, cutting down agency costs significantly.
  - **Reduction of Information Asymmetry'Costs:** Open Insurance, relying on APIs, enables the access and sharing of customers' data. This access to real-time and historical data could lead to a bigger transparency, reducing the infamous costs of adverse selection and moral hazard. This reduction of the information asymmetry will also lead to more accurate pricing decisions.
  - **Lowering Information-Gathering Costs through Smart Contracts:** The integration of Smart contracts data collection, verification, and claims execution automatisms with Open Insurance procedures will facilitate the data sharing between multiple entities minimizing the need for manual data collection and thus reducing the associated costs.
- **Increase in Transparency.** The integration of Open Insurance and DeFi insurance could lead to an increase in market transparency through:
  - **Reduce Information Asymmetry.** As already seen, Oracles connect real-world events to smart contracts, helping the development of automated processes in event verification and policy settlement lowering the information gaps between insurers and policyholders.
  - **A Better Risk Assessment.** The pillar of an Open Insurance environment is the insured's

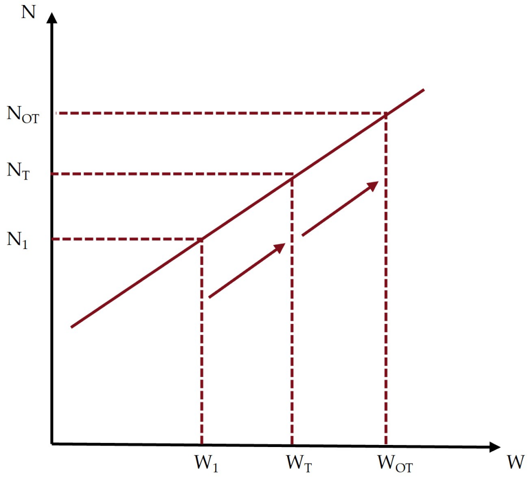

FIGURE 91: *Increase in Insurers Profits.*

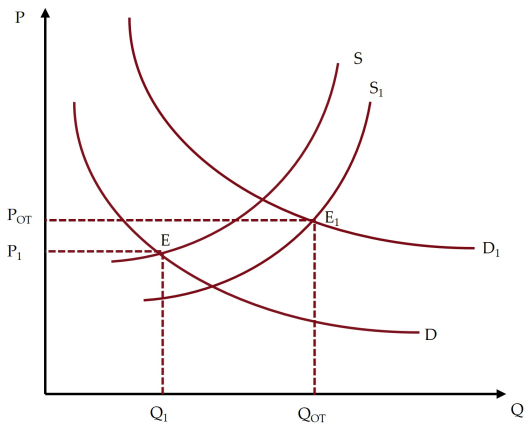

FIGURE 92: *Market Equilibrium.*

data sharing between organizations fostered by APIs. Thanks to real-time data access and the historicization of data series the market will potentially benefit by a lowering in moral hazard and adverse selection, bringing insurers to offer more tailored policies thanks to a more accurate risk assessment.

– **Tailor-Made and Transparent Prices.** The deeper data granularity unlocked by oracles and Open Insurance will lead to a better risk assessment allowing to design of customized price brackets much more specific to true individual's or entity's risk exposure.

- **Increase in Financial Inclusion.** The possibility of developing more tailored policies and the reduction of costs, as simplified in the behavioral model, will potentially lead to an increase in the number of insured. The individuals that were excluded by the market will benefit from policies that reflect their real risk and offer a consequent fair price.
- **Fraud Reduction.** The integration of DeFi insurance and data sharing of Open Insurance could improve fraud prevention, thanks to the traceability guaranteed by DLT. Moreover, the possibility of verifying data from multiple sources will make it harder to falsify information or manipulate insurance processes.
- **Market Innovation.** The development of highly innovative environments that rely both on DeFi insurance and open data sharing will lead to the deployment of new insurance solutions and automated policies for extreme market events or for specific risks related to emerging technologies.

References

[1] **Auer, R., Haslhofer, B., Kitzler, S., Saggese, P. and Victor, F.** *The Technology of Decentralized Finance (DeFi).* BIS Working Papers N. 1066, January 2023.

[2] **Beck, T. and Webb, I.** *Economic, Demographic and Institutional Determinants of Life Insurance Consumption across Countries*. The World Bank Economic Review, Vol. 17, N. 1, 2003.

[3] **Cristea, M., Marcu, N. and Carstina, S.** *The Relationship between Insurance and Economic Growth in Romania Compared to the Main Results in Europe - A Theoretical and Empirical Analysis*. Procedia Economics and Finance, December 2014.

[4] **EIOPA.** *Discussion Paper on Open Insurance: an exploratory use case in the insurance sector*. July 2023.

[5] **EIOPA.** *Open Insurance: accessing and sharing insurance-related data*. January 2021.

[6] **EIOPA.** *Report on the Digitalisation of the European Insurance Sector*. April 2024.

[7] **Eling, M. and Schmeiser, H.** *Insurance and the Credit Crisis: Impact and Ten Consequences for Risk Management and Supervision*. The Geneva Papers on Risk and Insurance - Issues and Practice, January 2010.

[8] **Federal Reserve Bank of Chicago.** *Economic Perspectives, No.5*. August 2024.

[9] **Feyen, E., Rodney, L., Rocha, R.** *What Drives the Development of the Insurance Sector? An Empirical Analysis based on a panel of developed and developing Countries*. The World Bank, Policy Research Working Paper, February 2011.

[10] **Guiso, L. and Paiella, M.** *Risk Aversion, Wealth and Background Risk*. Journal of the European Economic Association, December 2008.

[11] **IAIS.** *Global Insurance Market Report (GIMAR), Mid-Year Update*. July 2024.

[12] **ISF.** *State of the sector: Agri-Insurance for smallholder farmers*. March 2022.

[13] **Jaspersen, J., Ragin, M. and Sydnor, J.** *Predicting insurance demand from risk attitudes*. Journal of Risk and Insurance, April 2021.

[14] **Koijen, R. and Yogo, M.** *Shadow Insurance*. Econometrica, Vol. 84, N. 3, May 2016.

[15] **Ward, D. and Zurbruegg, R.** *Law, Politics and Life Insurance Consumption in Asia*. The Geneva Papers on Risk and Insurance, Vol. 27, N. 3, July 2002.

[16] **Mazzoni, N., Ciminelli, V., Morisani, G.** *Asset Tokenization: Potential Application*. Argo Magazine n. 25, April 2024.

[17] **Mazzoni, N., Mori, G., Bainotti, M., et al.** *Exploring the Digital Renminbi: Insights into China's CBDC*. Argo Magazine n 26, July 2024.

[18] **Mori, G., Bainotti, M.** *The Crypto Ecosystem: Key Elements and Risks*. Just in Time, November 2023.

[19] **Mazzoni, N., Ciminelli, V., Morisani, G.** *Asset Tokenization: Potential and Challenges of a Disruptive Technology*. Just In Time, July 2023.

In the previous issue

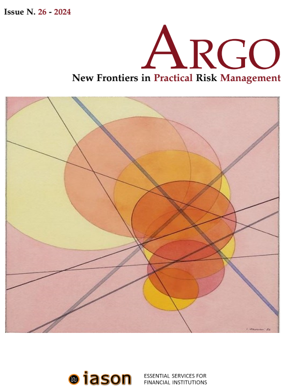

Issue N. 26 - 2024

TECHNOLOGY

Artificial Intelligence Act (AI Act)

Exploring the Digital Renminbi: Insights into Chinas CBDC

CREDIT RISK

Introducing Sectoral PD Satellite Models through Constrained BACE

Last issues are available at www.iasonltd.com/research/argo/