Advancing credit portfolio strategies : insights on corporate actions, asset ranking and interest rate hedging

Mathis Linger

Université d'Orléans

École doctorale de Sciences de la Société : Territoires, Économie, Droit
Laboratoire d'Économie d'Orléans

Thèse présentée par Monsieur Mathis Linger
soutenue le 27 juin 2025
pour obtenir le grade de Docteur de l'Université d'Orléans
Discipline/Spécialité Sciences Économiques

Advancing Credit Portfolio Strategies: Insights on Corporate Actions, Asset Ranking and Interest Rate Hedging

Thèse dirigée par :

Monsieur Marcel Voia	Professeur, Université d'Orléans

Co-encadrée par :

Madame Denisa Banulescu-Radu	Maître de Conférences, Université d'Orléans

Rapporteurs :

Monsieur Thanasis Stengos	Professeur, Guelph University, Canada
Monsieur Sylvain Benoit	Maître de Conférences, Université Paris Dauphine

Jury :

Monsieur Christophe Rault	Professeur, Université d'Orléans, Président du jury
Madame Miruna Pochea	Professor, Babeș-Bolyai University, Roumanie

Acknowledgments

First and foremost, I would like to express my deepest gratitude to Samer Comair and Guillaume Boulanger, whose guidance and expertise were instrumental in introducing me to the world of credit trading. Their support, patience, and insights have been invaluable in both my professional and academic journey.

I am also profoundly thankful to my academic supervisors, Marcel Voia and Denisa Banulescu-Radu, for their unwavering support, flexibility, and encouragement. Their mentorship and dedication have played a crucial role in shaping this thesis, and I am truly appreciative for the generous time and effort they invested in guiding me toward this achievement.

I would like to extend my heartfelt thanks to my co-authors and collaborators on this thesis: Omar Ben-Said, Pierre-Alexis Dufourg, Khalil Sbai, Thibaut Metz, Ilias Mellouki, Jamyang Bhutia, and Axel Pinçon. Their contributions, ideas, and hard work have been invaluable to this project, and it has been a privilege to work alongside them.

Finally, I extend my gratitude to my colleagues and peers, as well as to my family and friends, whose encouragement and understanding have been a constant source of motivation.

This thesis, conducted within the framework of the Convention Industrielle de Formation par la Recherche (CIFRE) program, would not have been possible without the collaboration between Drakai Capital and the University of Orléans. I am sincerely grateful for the opportunity to benefit from both academic and industry perspectives, which greatly enriched this work.

Contents

List of Figures

List of Tables

Chapter 1

General Introduction

The credit market has long been characterized by its opacity and reliance on over-the-counter (OTC) trading, where transactions were predominantly conducted via bilateral agreements and telephone negotiations. This structure fostered a "Banks-versus-All" paradigm, in which banks held a dominant position, controlling both the flow of information and market liquidity. The financial crisis of 2008 marked a turning point, as regulatory bodies across the world implemented reforms aimed at increasing market transparency, promoting standardization, and accelerating the digitalization of trading processes. These efforts culminated in the gradual "electronification" of the credit market, facilitating a transition towards an "All-to-All" trading model that mirrors the more transparent structure of equity markets.

This structural evolution has created new opportunities for innovation, particularly in the application of data science and quantitative research methods to credit investing. With the proliferation of electronic platforms and the surge in data availability, market participants can now leverage advanced technologies to enhance decision-making and strategy formulation. The integration of artificial intelligence (AI) and machine learning (ML) within this context has enabled the development of systematic approaches to credit investment that were previously unattainable due to data limitations and market opacity.

1.1 Historical context

1.1.1 Regulatory developments in the credit default swap market

The global financial crisis of 2008 had a profound impact on the financial markets, particularly the fixed income and credit default swaps (CDS) market. The CDS market had largely unchanged contract structure and trading mechanisms since the early 1990s, which, coupled with inadequate regulation and lack of transparency, contributed to the severity of the crisis. In response, a number of industry- and regulatory-initiated changes were established to address the issues in the CDS market.

The first reforms that were implemented in response to the financial crisis of 2008 were the Big-Bang and Small-Bang Protocols, in April and July 2009. They aim to create greater standardization

Date	Event Description
Apr. 2009	Big Bang Protocol
Jul. 2009	Small Bang Protocol
Jul. 2012	Dodd-Frank Wall Street Reform and Consumer Protection Act
Jan. 2013	Reporting of swap transactions and pricing for Index CDS
Mar. 2013	Mandatory clearing CDX.NA
Apr. 2013	Mandatory clearing iTraxx
Aug. 2013	Swap Execution Facilities (SEF)
Feb. 2014	Mandatory SEFs execution
Sep. 2014	ISDA 2014
Sep. 2016	Mandatory initial margin for non-cleared positions
Nov. 2021	Security-based swap dealers and major security-based swap participants are required to register with the SEC
Feb. 2022	Public dissemination of security-based swap transactions

Note: This table outlines significant milestones in the evolution of the CDS market following the 2008 financial crisis.

Table 1.1: CDS Market Evolution after 2008 Crisis

and substitutability by revising the ISDA Master Agreement, the document specifying CDS contract terms. They improved operational efficiency by eliminating redundant trades and facilitating centralized clearing. The Big Bang Protocol introduced four significant changes to the North American CDS market, including an auction mechanism to determine recovery rates, a Determinations Committee to determine credit events, a retrospective period for protection assessment, and fixed coupons (either 100 or 500 basis points) for single-name North American contracts. The Small Bang Protocol introduced similar changes to European corporate and Western European sovereign CDS contracts. In addition, the Big Bang Protocol removed restructuring as a credit event for North American corporate single-name contracts. These reforms aimed to align single-name CDS contracts with the standard corporate CDS indexes.

Later, in 2012, the US CDS market underwent significant changes through Dodd-Frank regulatory reforms, with the aim of reducing the risk posed by complex inter-institutional connections through OTC credit derivatives. Title VII of the Dodd-Frank Act established a complete regulatory framework for the OTC swap markets, which included CDS, with the goal of decreasing counterparty risk and enhancing pricing transparency. As a part of the Dodd-Frank Act, the bill introduced (i) requirement that standardized CDS contracts must be cleared through central counterparties (CCPs). The CCPs act as intermediaries between buyers and sellers of CDS contracts, reducing counterparty risk and increasing market transparency. CCPs also enable the clearing of offsetting trades, making cleared transactions more appealing to market participants than uncleared transactions ; (ii) rules for the registration and regulation of swap dealers and major swap participants: market participants trading in CDS contracts now have to report their trades to a trade repository (swap data repository SDR), enhancing market transparency and facilitating regulatory oversight. On January 3, 2013 the Depository Trust & Clearing Corporation (DTCC) required that all US registered swap dealers active in credit and interest rate trading have to report and publicly disseminate trade information to DTCC's SDR as soon as technically possible; generally within one to two minutes of execution time. This has contributed to increase post-trade transparency of centrally cleared obligors by providing public access to daily trading volume, open interest, and settlement prices.

On May 16, 2013, the Commodity Futures Trading Commission (CFTC) issued rules governing Swap Execution Facilities (SEFs), electronic platforms designed to promote transparency and competition in swaps trading, with implementation beginning in August 2013. Following these reforms, the "made-available-to-trade" determination in January 2014 required certain CDS contracts to be executed on SEFs. By February 2014, the current on-the-run and first off-the-run series of CDX North America Investment Grade, CDX North America High Yield, iTraxx Europe, and iTraxx Europe Crossover became mandatory for SEF trading, marking a significant step toward greater market transparency and efficiency. According to the CFTC, the adoption of SEFs has increased trading volumes and resulted in lower transaction costs for market participants.

Besides, the introduction of the 2014 ISDA Credit Derivatives Definitions brought additional significant changes to the credit derivatives market. One of the most significant changes was the introduction of a credit event for government bail-in, which allows protection buyers to be compensated in the event that the reference entity experiences a debt write-down or conversion as a result of government intervention. Additionally, the new definitions introduced a standard restructuring credit event, replacing the previous multiple and complex restructuring credit events, and the settlement process was simplified by allowing parties to choose cash settlement in place of physical settlement. The introduction of these new definitions brought greater clarity and consistency to the market, improving transparency and reducing operational risks.

In the most recent past, the introduction of mandatory initial margins for non-cleared trades added further safeguards. It was implemented in the US swap market in September 2016 and globally in March 2017. These regulatory reforms have led to an increase in bilateral trading costs as both dealers and customers are required to post margin, and the initial margin levels are higher than in comparable cleared contracts. The result is that market participants are now incentivized to shift towards cleared trades for instruments that are eligible for clearing. This trend is in line with the objective of the regulatory reforms to promote central clearing and reduce systemic risks in the derivatives market.

Despite the CFTC's successful implementation of SDRs as a cornerstone of the derivatives market infrastructure for other asset classes, it is noteworthy that SDRs specifically for CDS remained unestablished until 2022. This gap was finally addressed on February 14, 2022, when the SEC's regulations mandated public dissemination of security-based swap transactions. This milestone provided market participants and the public with granular, real-time access to transaction-level information, marking a significant step toward greater transparency in the CDS market[1].

1.1.2 Evolution of bond market trading

The bond market, like the CDS market, has undergone significant changes in execution, clearing, and settlement over the past decades, with the most important shift being the electronification of trading. Historically, corporate bonds were predominantly traded via voice-based interactions between institutional investors and a small number of dominant dealers, with inter-dealer brokers playing a key intermediary role. However, advancements in data availability and new technologies have driven a significant tran-

[1] See the statement from Chairman Gary Gensler.

sition toward electronic platforms, such as those using the RFQ system, disrupting these long-standing practices. From 2010 to 2017, the share of electronically executed trades increased from 6% to over 13% by volume and from 9% to 25% by trade count, signaling meaningful adoption of digital channels ([77]). This trend is particularly evident in US treasuries, where electronic trading now accounts for 65% of transactions, with the majority executed through RFQ protocols, according to Flow Traders. Similarly, the adoption of electronic trading is growing in other bond markets, including investment-grade and high-yield corporate bonds, highlighting a broader industry shift towards more efficient, data-driven trading mechanisms.

These changes have significantly influenced liquidity and pricing. The move to "All-to-All" platforms has facilitated a more competitive trading environment, enabling various market participants—including asset managers, banks, and brokers—to trade directly with each other. This has lowered transaction costs, improved liquidity, and provided better price discovery, particularly for smaller trade sizes. As evidenced by [77], over 50% of odd-lot trades in investment-grade bonds now occur electronically, reducing search costs and enhancing price transparency. Transaction costs have fallen markedly, with high-yield bond trades dropping from 35 basis points in 2010 to below 20 basis points by 2017, while voice trading costs have also declined, albeit more gradually, influenced by competition from electronic venues.

The expansion of fixed-income ETFs has also played a crucial role in improving liquidity, as these products offer both trading opportunities and liquidity to the underlying bond market. About 20% of ETF volumes are traded directly in the underlying bonds, according to State Street Global Advisors. Furthermore, portfolio trading—a protocol allowing investors to trade blocks of credit risk across multiple bonds—has surged in popularity, accounting for about 5–6% of corporate bond volumes as of 2023, compared to nearly zero in 2018, according to FINRA. This growth highlights the potential for larger trades to be executed more efficiently.

These advancements in electronic trading and data availability are enabling a more systematic approach to fixed-income investing. Traders now have access to real-time pricing, predictive analytics, and ML tools that help inform decisions and optimize trade execution. As more participants adopt algorithmic strategies, there is a noticeable shift away from traditional voice-based negotiations, especially for smaller trades, leading to greater market transparency. The increasing use of these automated tools supports more reliable execution, allowing for better pricing and reduced slippage. According to Tradeweb, the average time to execute a basket trade has decreased significantly from 90 minutes in early 2021 to less than an hour by mid-2023, underscoring the growing efficiency of electronic trading in the bond market.

As shown in Figure 1.1, the corporate bond market, particularly US and European investment-grade (IG) and high-yield (HY) segments, appears to be nearing a critical inflection point. Drawing parallels to the rapid increase in electronification and trading volumes observed in FX spot, US treasuries, cash equities, and equity options after their own inflection points, the corporate bond market seems poised to follow a similar trajectory. As it transitions, it could emulate the path of US and EU treasuries, where widespread adoption of electronic trading significantly boosted efficiency and liquidity.

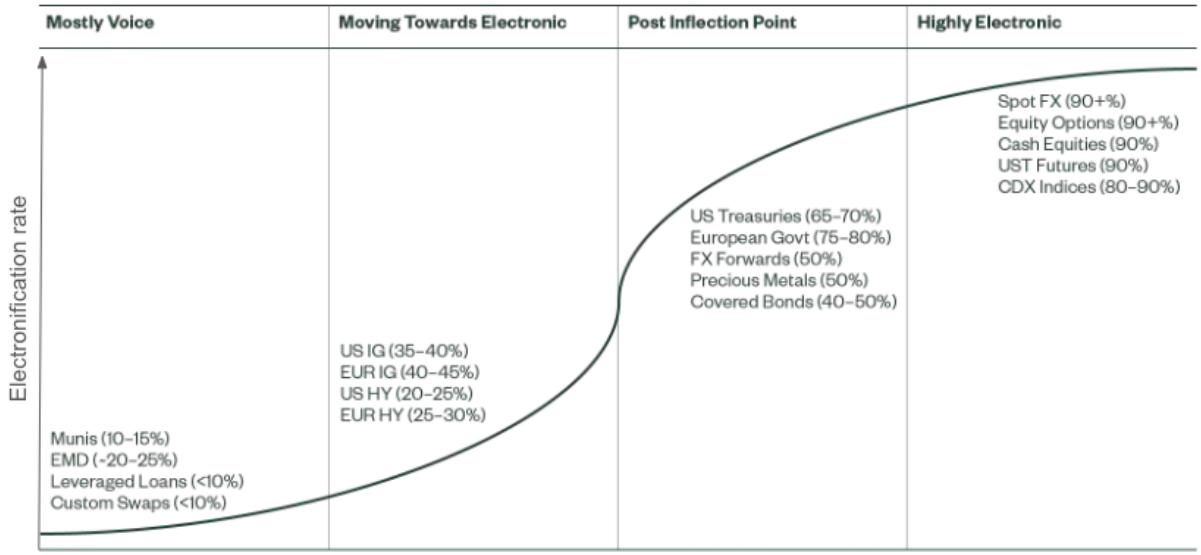

Source: Flow Traders "New Frontiers in Fixed Income Liquidity: The Evolution of the Ecosystem and the Opportunities for the Buy-side".

Figure 1.1: Electronification rate across markets

1.2 Credit market structure: from "Banks-versus-All" to "All-to-All" paradigm

1.2.1 CDS market

The structure of the CDS market has historically been highly concentrated and dominated by a small number of large dealers, reflecting a "Bank-versus-All" framework. In this structure, dealers acted as central intermediaries, controlling both liquidity and counterparty risk. In [37], evidence from the CDS market in 2012 reveals that the top 10 dealers accounted for 72.2% of all buying activity and 77.4% of all selling activity, demonstrating a strong reliance on a limited set of institutions for market-making. This concentration raised concerns about systemic risk, as defaults by key dealers could trigger cascading failures across the market. For instance, the five largest buyers (sellers) accounted for 45.5% (49.8%) of the gross notional amount traded, underscoring the centrality of these entities in market transactions. Such interconnectedness made the CDS market vulnerable to contagion effects, where distress at one institution could rapidly spread to others.

To better understand the structural evolution of the CDS market, Figure 1.2 from the referenced research [37] illustrates the network topology of dealer and non-dealer interactions, both across all CDS reference entities and separately for corporate and sovereign CDS. The first figure highlights the overall network structure, demonstrating the dominance of major dealers in the CDS market. The size of the nodes represents the trading volume, while the thickness of connections indicates the notional value traded between counterparties. The visualization confirms that a small group of dealers acts as central hubs, with non-dealer entities predominantly trading through them. This pattern exemplifies the "Bank-

versus-All" model, where liquidity and risk management are concentrated in a few key institutions. The second figure, which focuses on corporate CDS, reveals a slightly more dispersed network compared to the aggregate market view. Corporate CDS trading exhibits relatively higher connectivity among smaller participants, suggesting early signs of decentralization. However, the larger dealers continue to dominate, reflecting the persistence of the traditional intermediary-driven structure in corporate markets. The third figure shifts attention to sovereign CDS, which displays a more concentrated network topology. Here, the largest dealers maintain even stronger connections, emphasizing their centrality in managing sovereign credit risk. Notably, the network structure for sovereign CDS shows fewer links among non-dealers, highlighting the importance of dealers in providing liquidity for this segment.

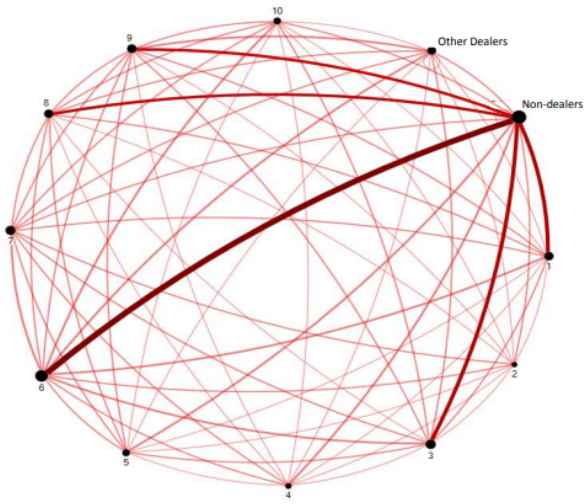

All reference entities

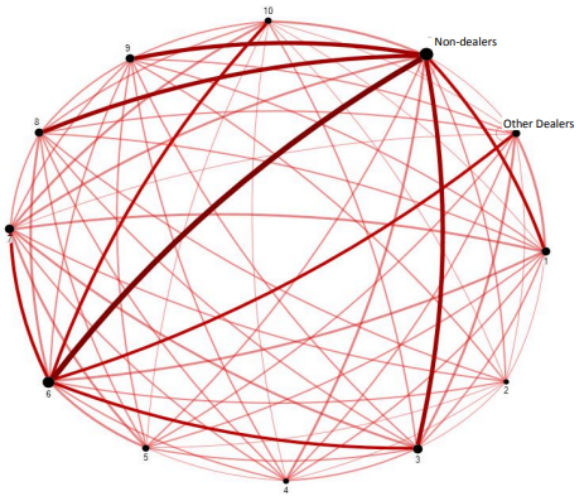

Corporate reference entities

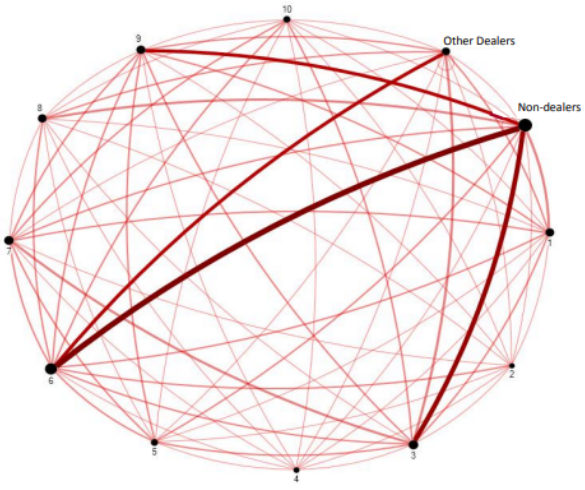

Sovereign reference entities

Source: [37]

Figure 1.2: Dealer Topology

The regulatory reforms implemented in the wake of the 2008 financial crisis have played a pivotal role in addressing systemic risks stemming from highly interconnected market networks. These reforms, as outlined in Section 1.1.1, aimed to increase market transparency, enhance operational efficiency, and reduce vulnerabilities, as shown in [24]. Early initiatives such as the Big Bang and Small Bang Protocols focused on standardizing contract terms, which laid the groundwork for more resilient market practices. Building on these early steps, the Dodd-Frank Act introduced crucial measures, including mandatory central clearing, trade reporting, and the creation of SEFs, all designed to enforce transparency both pre- and post-trade. In 2022, additional requirements were introduced for the public dissemination of security-based swap transactions, providing market participants with real-time access to trading data and enabling more informed decision-making. These cumulative regulatory changes, by increasing transparency and lowering barriers to entry, are dismantling the traditional "Bank-versus-All" model, fostering greater competition, improving liquidity, and enhancing overall market resilience. Moreover, this shift has led to a fundamental change in the behaviors of market participants.

1.2.2 Corporate bond market

As detailed in Section 1.1.2, the corporate bond market, traditionally dominated by a dealer-driven framework, is steadily transitioning toward an "All-to-All" trading paradigm. This shift reflects broader trends of technological innovation and enhanced market connectivity. Historically, corporate bonds were traded primarily through voice-based interactions between institutional investors and a few dominant dealers. However, as highlighted in [77], the rise of electronic trading platforms, including those utilizing the RFQ system, has begun to disrupt these longstanding practices.

Despite these advancements, the shift to an "All-to-All" trading model—where buyers and sellers interact directly without dealer intermediation—remains in its infancy. As of 2017, "All-to-All" trading accounted for less than 2% of total volume. The persistence of the dealer-centric structure can be attributed to several barriers, including the complexity of executing large block trades, the information-sensitive nature of high-yield bonds, and the reliance on dealer relationships for liquidity during periods of market stress. These factors highlight the limitations of current electronic platforms in addressing the needs of institutional investors managing large or illiquid positions. Another key aspect of this transition is its impact on execution quality. Recent research [79] reveals significant disparities in execution prices for institutional investors. For instance, less active traders—defined as those with smaller aggregate bond portfolios—paid 0.17% more on average for purchases and received 0.36% less for sales compared to their more active counterparts, even when controlling for trade size and dealer identity. This underscores the persistence of inefficiencies in the corporate bond market, particularly in opaque OTC environments where dealer market power plays a central role.

The inter-dealer market has also experienced notable changes. Dealers historically used inter-dealer trades to manage inventory risk, particularly for large transactions. However, electronic trading has reduced reliance on this mechanism, with inter-dealer trading's share of total investment-grade transactions dropping from 42% in 2010 to 28% in 2017 [77]. This shift underscores the role of electronic platforms in streamlining inventory management and facilitating direct matches between end-users. Yet, challenges remain. The adoption of electronic trading is still concentrated among smaller trade sizes

and investment-grade bonds, while larger trades and high-yield securities continue to rely heavily on voice-based negotiations. Furthermore, the RFQ system's dependence on pre-established relationships between participants limits the competitive entry of smaller dealers and the broader adoption of "All-to-All" trading. Indeed, the top 10 dealers maintained their dominance over both voice and electronic trading during the study period, capturing 70% of the market's volume. The high concentration of dealer market power exacerbates these inefficiencies. [79] indicates that the largest dealer executes 69% of the volume for the average bond issue, with the top three dealers capturing 92% of the market share. This level of concentration, measured by a Herfindahl-Hirschman Index of 61%, is far above the threshold for what is considered a highly concentrated market.

Overall, while the corporate bond market is moving toward greater connectivity and reduced intermediation, the transition is evolutionary rather than revolutionary. The nascent "All-to-All" paradigm holds promise for democratizing access and fostering competition, but its growth will likely depend on overcoming the entrenched structures and operational complexities unique to the bond market.

1.3 Industrial perspectives

Drakai Capital operates at the forefront of this transformation, focusing on leveraging these technological advancements to address core challenges in the credit market. Specifically, its research agenda targets three critical areas of innovation:

- **Market-Neutral Arbitrage Opportunities**: Drakai Capital seeks to identify and exploit arbitrage opportunities within the credit market using AI-driven models. These models incorporate a diverse range of data inputs, including fundamental financial indicators, technical analysis metrics, macroeconomic variables, and alternative data sources. The goal is to achieve a market-neutral stance, ensuring returns are decoupled from broader market movements, thereby minimizing risk exposure.
- **Portfolio Diversification and Risk Management**: Drakai Capital emphasizes the construction of diversified credit portfolios by combining domain expertise with ML algorithms. These portfolios aim to optimize risk-adjusted returns while ensuring stability and low volatility. The integration of ML tools allows for dynamic rebalancing and enhanced responsiveness to market changes, ultimately providing superior diversification benefits to investors.
- **Liquidity Analysis and Optimal Execution**: Liquidity in the credit market remains fragmented, with significant heterogeneity in pricing across instruments. Drakai Capital conducts in-depth studies of liquidity dynamics to ensure effective execution strategies. By understanding the evolving nature of liquidity and the impact of new market participants, the firm can minimize transaction costs and price slippage in an increasingly electronic trading environment.

Achieving success in these areas requires overcoming several scientific and technical hurdles. Drakai Capital's mandate is to leverage cutting-edge quantitative techniques and technological innovations to address core challenges in credit trading. Central to this mission is the development of advanced models

that can navigate the complex dependencies between financial assets, particularly under conditions of market stress. Periods such as the subprime mortgage crisis of 2008 and the market disruptions of 2020 have underscored the importance of capturing dynamic correlations, as heightened interconnectedness during these times amplified systemic risks. Drakai Capital also prioritizes the integration of AI and ML to enhance decision-making in credit investing, while addressing critical issues of model interpretability, robustness, and real-world applicability.

This thesis, conducted as part of the convention industrielle de formation par la recherche (CIFRE) collaboration between Drakai Capital and the Laboratoire d'Économie d'Orléans (LEO), addresses key challenges in constructing ML strategies for credit trading. By bridging academic research with real-world financial applications, this work not only advances the practical use of quantitative methods in credit investing but also enriches the broader academic discourse on financial innovation. The insights gained from this thesis have the potential to inform both industry practices and regulatory frameworks, fostering a more efficient and resilient credit market.

1.4 Key stages in developing quantitative trading models

The development of quantitative trading strategies typically follows a standardized and iterative workflow, yet each asset manager or hedge fund tends to adopt a unique setup and approach to constructing systematic strategies. In fact, it can be stated that there are as many configurations as there are asset managers, each reflecting distinct philosophies, tools, and objectives. This section presents what is, from my perspective, the most common setup observed across asset management firms, with the goal of outlining a typical workflow in quantitative strategy development (Figure 1.3). This thesis, in turn, seeks to enhance several critical stages of this process, particularly the definition of the trading universe, the design of quantitative models and portfolio construction, aiming to improve both their efficiency and predictive power within the context of systematic investing.

The first step in developing a quantitative investment strategy is defining the trading universe, which involves selecting the set of assets to be analyzed and included in the portfolio, as well as determining the relevant time periods during which these assets will be traded. These considerations are crucial because the defined universe directly shapes the dataset used to train the quantitative model, influencing its performance and ability to generalize across unseen data. When selecting instruments, it is common to focus on specific tenors of CDS contracts known for their higher liquidity and stable behavior over time. This simplifies the learning process for the model, as well-behaved, homogeneous data sets allow for better generalization. However, it is important to avoid over-filtering the data, as excluding too many time periods or specific assets can lead to distribution mismatches between the training and test sets, undermining the model's ability to perform in real-world trading. Overfitting, a major challenge in quantitative finance, occurs when a model is overly tailored to historical data and fails to perform well in live trading. This is often due to the model being too closely tailored to the idiosyncrasies of the training data, failing to adapt to new, unseen market conditions. Therefore, careful selection and balance in defining the trading universe are essential to developing a robust model that can effectively handle diverse market scenarios and avoid overfitting.

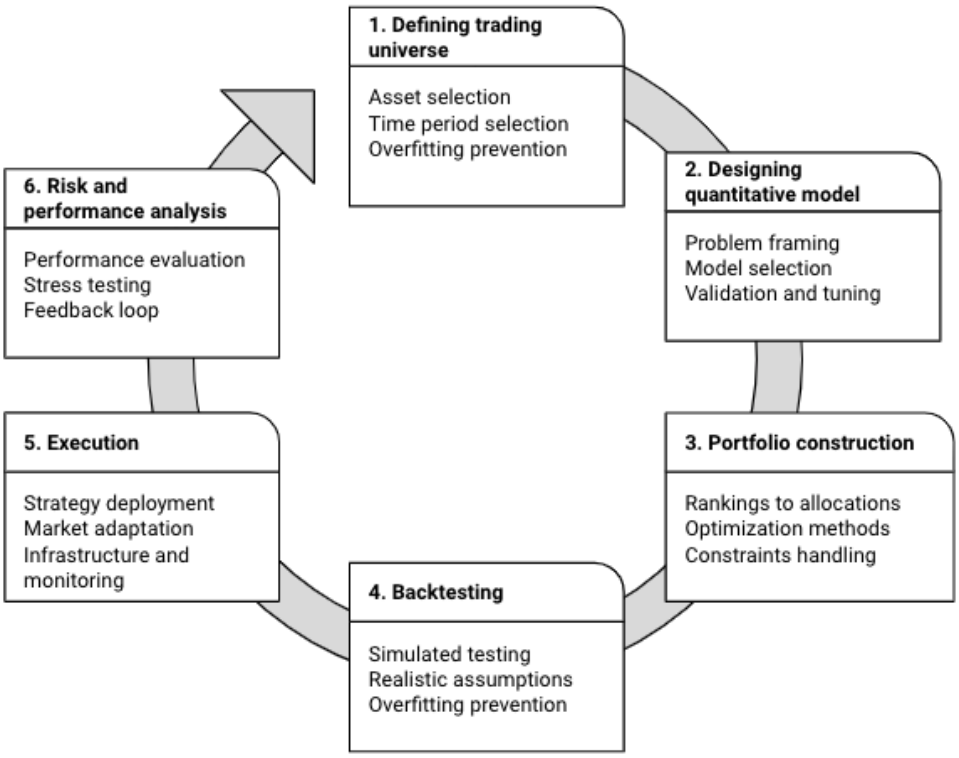

Figure 1.3: Quantitative trading strategy workflow

The second step in developing a quantitative trading strategy involves designing the quantitative model, which begins with a clear definition of the problem to be solved. This entails specifying the prediction target, such as forecasting asset returns over a defined time horizon, and determining the model's objective—whether it will focus on classification, regression, or ranking. Once the problem is defined, selecting an appropriate algorithm becomes crucial. Depending on the task, one might choose simpler methods like linear regression for interpretability, tree-based models to capture non-linear relationships, or more complex approaches, such as neural networks, to uncover intricate patterns in the data. The next phase centers on optimizing the loss function to train the model effectively, often using techniques like cross-validation to evaluate its performance across different data subsets. Fine-tuning hyperparameters is essential during this phase, ensuring a balance between bias and variance, and enhancing the model's robustness while mitigating the risk of overfitting. This entire process forms the basis for creating predictive models that can identify patterns and generate practical trading signals from historical and real-time data.

The third step focuses on portfolio construction, which translates the model's predictions into actionable asset allocations. In asset management firms, this stage is often handled separately from the modeling process. Initially, the trained model generates predictive outputs, such as rankings or expected returns for individual assets, which serve as the foundation for applying portfolio construction methods. These methods are then employed to determine an optimal asset allocation that aligns with the model's forecasts. This process typically utilizes heuristics or optimization techniques designed to balance expected returns with the portfolio's risk profile. The allocation strategy takes into account various constraints, including investor preferences, risk limits, turnover restrictions, and regulatory requirements,

ensuring the constructed portfolio not only reflects the firm's strategic goals but also complies with necessary guidelines. By using these techniques, portfolios are built to be diversified and to optimize the trade-off between risk and return in a structured and efficient manner.

The fourth step is backtesting, a crucial phase that allows for the evaluation of both the model and the portfolio construction before live deployment. Backtesting involves simulating the strategy on historical data to assess how it would have performed under past market conditions. This process provides valuable insights into the strategy's potential strengths and weaknesses, focusing on key performance metrics such as returns, volatility, drawdowns, and risk-adjusted measures. Additionally, backtesting tests the strategy's robustness by examining its performance across various time periods, market regimes, and stress scenarios, thus ensuring it remains resilient under different conditions. To enhance reliability, it is essential to incorporate realistic trading assumptions—such as transaction costs, slippage, and liquidity constraints—into the backtesting process. These factors can significantly impact results and provide a more accurate reflection of how the strategy would behave in live markets. However, a common pitfall of backtesting is the risk of overfitting, where the model performs well on historical data but fails to generalize to unseen, future data. To address this challenge, it is critical to use validation techniques like out-of-sample testing and walk-forward analysis. These methods help ensure that the observed performance is not merely the result of data-mining or curve-fitting, and they promote a more reliable and realistic assessment of the strategy's true potential in live trading environments.

The fifth step in the development of a quantitative trading strategy is execution, an essential phase where the strategy transitions from theoretical design to real-world implementation. At this stage, the focus shifts to translating the portfolio allocations determined in earlier steps into actual trades, ensuring that orders are placed efficiently and at minimal cost. Execution is particularly pivotal in credit markets, where trading practices have undergone significant transformations, as discussed in Section 1.1.2. The rise of electronic trading platforms has revolutionized liquidity, transparency, and accessibility, enabling the execution of systematic strategies in fixed-income markets with greater ease and efficiency. Effective execution in this context requires more than simply placing trades; it demands the integration of real-time market data, as well as the maintenance of a robust trading infrastructure. Low-latency systems are essential for ensuring timely execution, while comprehensive risk controls are necessary to navigate dynamic market conditions and mitigate potential disruptions. The ability to adapt to these changing conditions is vital, particularly in the face of unforeseen market events. Moreover, continuous performance monitoring is indispensable to evaluate execution quality, identify deviations from expected costs, and assess market impact. This monitoring process serves as a feedback loop, enabling practitioners to refine earlier stages of the strategy development process and optimize the execution approach over time, ensuring that the strategy aligns with its intended performance goals.

The final step in the development of a quantitative trading strategy is risk and performance analysis, which primarily serves to ensure that the performance of the model aligns with the theoretical performance observed during backtesting. This phase helps identify any shifts in performance once live trading begins, providing insight into whether the model is deviating from the expected results. It acts as a diagnostic tool, helping to detect if the model is drifting abnormally, if the assumptions made during backtesting were incorrect, or if the model is not executing as planned. Key performance metrics, such as the Sharpe ratio or maximum drawdown, are used to monitor the strategy's live performance and compare it against

the backtest. The analysis also serves to pinpoint potential weaknesses, such as underperformance in specific market regimes or instability in response to changing market conditions. Continuous monitoring is crucial, as it allows for the detection of model drift, recalibration of parameters, and adjustment of constraints as market conditions evolve. This diagnostic approach ensures the model remains on track, allowing for timely adjustments to maintain alignment with the expected performance and long-term objectives.

In conclusion, the development of a robust and effective quantitative trading strategy is an iterative process, where each stage plays a critical role in the overall success of the system. From defining the trading universe, designing predictive models, portfolio construction, backtesting, execution, and performance evaluation, each step is interdependent. A flaw or oversight in any one of these stages can compromise the entire strategy, highlighting the difficulty of creating a consistent and high-performing trading workflow. In quantitative finance, the greatest challenge is not the execution of a well-designed strategy, but rather the creation of a truly significant one—a strategy that performs well not only in backtests but also in live markets. Achieving this requires a delicate balance between predictive power and robustness, ensuring that the strategy can adapt to changing market conditions while consistently generating value. This is especially challenging due to the dynamic nature of markets, influenced by numerous unpredictable factors, which makes it difficult to build a model that can reliably deliver positive results over time. Moreover, while advanced techniques like ML and algorithmic modeling are powerful, they do not guarantee success. The true challenge is identifying signals that are genuinely predictive and resilient enough to withstand the noise and volatility inherent in financial markets. Many strategies that appear promising in backtesting fail to perform in real-world conditions due to overfitting or an inability to adjust to market shifts. Therefore, developing a sustainable trading strategy requires not only a deep understanding of market mechanics but also the ability to manage inherent risks and uncertainties.

This thesis aims to address some of these technical challenges by providing novel solutions that enhance the robustness and adaptability of fixed income quantitative strategies. Through providing insights into credit universe selection, model design, and portfolio construction techniques, the goal is to offer asset managers practical insights and methodologies that can help improve the reliability and long-term success of quantitative trading systems in navigating the complexities of real-world credit markets.

1.5 Key contributions to advancing quantitative trading strategies

Structured into four chapters, the thesis examines key real-world challenges faced by Drakai Capital, offering insights into the complexities of modern credit markets—insights that may also prove valuable to asset managers and regulatory bodies more broadly.

Chapter 2 *Corporate Actions and Credit Risk Dynamics: Evidence from Stock Buybacks, Cash Dividends and Mergers & Acquisitions* aims to improve stage *"1. Defining trading universe"* in Figure 1.3. More precisely, it examines how corporate action announcements influence CDS spreads, which often exhibit abnormal fluctuations around these events. Such movements—frequently characterized as

outliers—can distort data distributions, introduce noise, and degrade the performance of ML models when included in the training process. By analyzing the timing, magnitude, and persistence of these abnormalities, this chapter enhances our understanding of the dynamics surrounding CDS spread reactions to corporate actions announcements. The focus of this research is twofold. First, it explores the temporal patterns of CDS spread adjustments following key corporate events, including stock buybacks, mergers and acquisitions (M&A), and cash dividend announcements. Second, it disentangles the underlying drivers of these reactions, distinguishing between wealth effects—linked to changes in financial flexibility—and signaling effects, which reflect shifts in perceived creditworthiness. By investigating these mechanisms, the chapter identifies the specific characteristics of companies and events that amplify or mitigate credit risk responses. From a practical standpoint, this analysis has direct implications for Drakai Capital's trading framework. Understanding abnormal credit behavior during corporate event windows allows for the refinement of CDS trading strategies by avoiding volatile periods that could distort model performance. More specifically, it enables asset managers to define a trading universe that excludes unstable time frames, ensuring ML models are trained on more representative and robust data. Furthermore, these insights provide asset managers with tools to anticipate and interpret credit market reactions, improving decision-making and execution around corporate action announcements. Finally, this chapter's contributions extend beyond isolating the effects of individual corporate actions. It offers a comparative framework that evaluates multiple event types within a unified methodology, addressing gaps in the existing literature. Moreover, it highlights the importance of accounting for event timing and firm-specific attributes, shedding light on the heterogeneous nature of credit risk responses.

Chapter 3 *Enhancing Long-Short Portfolios: A Refined Approach using Learn-to-Rank Algorithms* is related to *"2. Designing quantitative model"* in Figure 1.3, and investigates inherent asymmetries in classical learning-to-rank (LTR) loss functions, which, while effective in information retrieval tasks, face notable limitations when applied to long-short portfolio construction. Traditional LTR approaches tend to emphasize accuracy at the top of the ranking spectrum, often overlooking the importance of correctly identifying underperforming assets at the bottom. This imbalance poses a critical challenge in the context of long-short strategies, where both long and short positions play equally vital roles in generating returns and maintaining market neutrality. Through a combination of theoretical insights and empirical evidence, this chapter identifies and quantifies these biases, shedding light on their adverse impact on portfolio performance. Building upon this analysis, we propose tailored modifications to three widely used LTR loss functions—ListNet, ListMLE, and ListFold—adapting them specifically to address the dual ranking requirements of long-short strategies. These adaptations are designed to enhance ranking precision across the entire asset distribution, ensuring a more symmetric and consistent evaluation of both outperformers and underperformers. Empirical validation using weekly China A-share equity data demonstrates the effectiveness of these modifications, showcasing significant improvements in both ranking performance and portfolio returns. While the empirical tests in this chapter focus on equity data for the purpose of academic publication, the insights and methodologies developed are directly applicable to Drakai Capital's fixed-income trading strategies. Specifically, the refined loss functions provide a robust framework for improving ranking precision and optimizing long-short portfolio construction in credit markets.

Chapter 4 *Unifying Asset Ranking and Portfolio Weighting through a Multi-Task Neural Network* introduces a novel framework that unifies asset ranking and portfolio allocation within a single, cohesive model using a multi-task neural network. Looking to Figure 1.3 , this work aims to unify stages *"2. De-*

signing quantitative model" and *"3. Portfolio construction"*. Traditionally, quantitative asset managers, including Drakai Capital, rely on a two-step process for portfolio construction, where asset ranking models are first used to predict relative performance, and portfolio weights are subsequently determined through heuristic rules or separate optimization routines. While effective, this sequential approach often fails to account for the interdependencies between ranking predictions and allocation decisions, potentially limiting overall performance. In response, this chapter proposes an integrated methodology that jointly optimizes both tasks within a unified learning framework. By shifting the optimization process directly into the learning framework, the model captures interactions between ranking and allocation, resulting in more coherent and robust investment decisions. Although the empirical validation presented in this chapter is based on equity data for the purpose of publication, the proposed framework is naturally extensible to fixed-income markets. This integration provides Drakai Capital with a scalable and adaptable tool for optimizing credit portfolios, aligning predictive modeling with real-world trading objectives and reinforcing its competitive edge in systematic strategies.

Chapter 5 *Understanding the Equity-Corporate Bond Nexus: A Framework for Risk Decomposition and Interest Rate Hedging* relates to stage *"3. Portfolio construction"* in Figure 1.3. It addresses the interplay between equity and corporate bond markets by decomposing their correlation into credit risk and interest rate sensitivity components. Recognizing that CDS spreads and sovereign bond yields reflect distinct dimensions of corporate bond risk, this chapter introduces a novel framework to isolate these effects. More specifically, it leverages the equity-CDS correlation to capture credit risk sensitivity and the equity-sovereign bond correlation to represent interest rate exposure, providing a clearer view of the fundamental drivers influencing the equity-corporate bond correlation under different market conditions. The empirical analysis highlights significant variations in the relative impact of these components across credit qualities and economic regimes. During stable periods, investment-grade bonds exhibit higher sensitivity to interest rate movements, whereas high-yield bonds predominantly respond to credit risk fluctuations. However, in times of heightened market stress—such as the COVID crisis—credit risk dominates for all corporate bonds, temporarily overriding the influence of interest rates. Post-COVID, interest rate sensitivity regains prominence, particularly for investment-grade securities, as central banks shift towards monetary tightening. From a trading perspective, these findings hold direct implications for Drakai Capital's risk management and hedging strategies. Understanding how the relative importance of credit and interest rate risk evolves over time allows for more precise adjustments in portfolio hedging techniques. To this end, the chapter also introduces and validates two adaptive hedging strategies—the Yield Ratio and Regression Model approaches—which dynamically adjust hedge ratios in response to shifting market conditions. Both strategies demonstrate superior performance over static, duration-convexity based hedging techniques, particularly in environments characterized by rapid shifts in interest rates or credit spreads. Beyond its immediate application to Drakai Capital, this research contributes to the broader asset management and risk management literature by offering a systematic approach to decomposing corporate bond risk and refining interest rate hedging methodologies.

Taken together, these contributions underscore the importance of integrating domain-specific insights with rigorous quantitative techniques to enhance trading strategies in modern credit markets. By addressing key challenges across asset selection, model design, and portfolio construction, this thesis not only refines existing methodologies but also paves the way for more adaptive and resilient systematic strategies. Ultimately, these advancements provide a robust foundation for both academic research and

real-world implementation, offering Drakai Capital and the broader investment community actionable tools to navigate an increasingly complex financial landscape.

Chapter 2

Corporate Actions and Credit Risk Dynamics: Evidence from Stock Buybacks, Cash Dividends and Mergers & Acquisitions[1]

Abstract

This chapter investigates the temporal dynamics and impact of corporate action announcements, including stock buybacks, cash dividends, and mergers and acquisitions, on credit risk. Relying on a dataset of single name credit default swap from 2012 to 2023, we employ a comprehensive event-study approach to examine different event windows and their comparative effects on credit risk. Key findings highlight the significant influence of stock buyback announcements, particularly within two days post-announcement. Dividend raises prompt an immediate and substantial response, while dividend cuts exhibit extended anticipation and impact. The temporal effects of mergers and acquisitions differ for seller and buyer companies, with sellers showing notable anticipation and buyers experiencing a more extended but subdued impact. We differentiate wealth transfer from signaling effects, revealing wealth transfer's dominance around these corporate event announcements. Empirical evidence supports the positive impact of stock buybacks and mergers and acquisitions magnitude on abnormal spread changes and the market's asymmetric response to cash dividend announcements. Finally, regional disparities in mergers and acquisitions announcements underscore the influence of governance structures on credit market responses.

[1]This chapter is based on the research article *Corporate Actions and Credit Risk Dynamics: Evidence from Stock Buybacks, Cash Dividends and Mergers & Acquisitions* co-authored with Guillaume Boulanger and presented at the first International Conference in Finance and Behavioral Decision Making, 10-12 October 2024, Cluj-Napoca, Romania.

2.1 Introduction

Credit derivatives, particularly credit default swaps (CDS), are fundamental instruments for managing credit risks in financial markets ([49]). CDS contracts offer insurance against default by a specific company or sovereign entity. In return for this insurance, the CDS buyer pays periodic premiums to the seller, gaining the right to sell bonds issued by the reference entity for their face value in the event of a predefined credit event. These instruments are considered a perfect risk indicator because they represent the cost to insure against the default of a particular debt issuer, providing a direct measure of market's perception of credit risk ([5], [69], [76]). On the other hand, corporate actions are pivotal events that can significantly influence a company's financial stability, capacity to meet debt obligations and therefore perceived creditworthiness. As such, each corporate action has the potential to alter a company's risk profile, shaping the trajectory of its single-name CDS spreads. Our study investigates three specific corporate actions: stock buybacks, mergers and acquisitions (M&A) and cash dividends. Buybacks involve a company repurchasing its own shares, cash dividends entail direct shareholder payments, and M&A activities encompass strategic and financial consolidation efforts between companies.

The interest in corporate events impacting credit risk pricing has seen a significant rise among scholars, financial regulators, and industry practitioners. This relationship is influenced by numerous factors such as company's financial resilience, market sentiment, economic conditions, and the unique characteristics of the specific corporate action at hand. For further examination, researchers frequently employ the event study methodology, a statistical method used to evaluate the impact of an event on the value of a company, providing insights into market efficiency and reactions. Recently, this approach has gained prominence in assessing the influence of events on credit risk, thanks to the expanding wealth of data and its increasing reliability ([4]). In the current literature, while there is often a focus on exploring individual corporate actions ([80], [87], [48]), there is a noticeable absence of comparative analyses that assess how various corporate actions influence credit risk within a unified evaluation framework. Moreover, the literature often lacks adequate justification for the chosen time frame of analysis, contributing to the incomplete exploration of the dynamics that influence the response of credit risk to corporate action announcements. Finally, the current body of research faces challenges in establishing a strong and consistent grasp of the specific attributes influencing the relationship between corporate actions and credit risk. This is highlighted by the common reliance on limited sample sizes.

This chapter contributes to the existing literature in several ways:

1. **Temporal dynamics of CDS spread reaction**: Our methodological framework delves into the temporal aspect of CDS spread reactions to corporate event announcements, offering a more nuanced understanding of the temporal dynamics involved.
2. **Disentangling wealth versus signaling effect**: Following the investigation of the temporal dynamics, an examination is conducted to test for the dominant effect shaping credit risk reactions, namely the wealth and signaling effects.
3. **Company characteristics impact**: Finally, we analyze the mechanisms through which corporate actions affect credit risk, examining company and corporate action characteristics and their connection to credit risk responses.

In contrast to most of the existing literature, this holistic approach enables a comparative analysis, shedding light on the relative effects of different corporate actions on credit risk. The remainder of the chapter proceeds as follows. We describe our hypotheses and theoretical framework in Section 2.2. Then, in Section 2.3, we describe our methodological approach. Finally, we present the data used in the study and empirical results in Section 2.4.

2.2 Theory and hypotheses

Our research endeavors to address three primary questions: (i) Around the timing of corporate action announcements, when do we observe abnormal and significant movements in credit risk? ; (ii) For each type of corporate action, which of the signaling effect or the wealth effect, prevails in influencing credit risk? ; (iii) What are the key company and corporate action characteristics that have the most significant impact on abnormal changes in credit risk? In the ensuing subsections, we develop hypotheses essential for addressing these inquiries.

2.2.1 Impact and temporal dynamics

First, we examine the dynamics of credit risk reaction to stock buybacks, cash dividends, and M&A announcements. More precisely, we ask if these corporate actions influence credit risk, and investigate the chronology through which these effects unfold, from the anticipation preceding a corporate action to the immediate aftermath, and the gradual ripple effect over time. It is important to note that the existing body of literature often falls short in providing a quantitative rationale for selecting specific event window durations in traditional event studies. These studies commonly rely on predetermined date intervals, with the standard practice involving a symmetric window centered around the announcement date, making the assumption that credit risk reacts both before and after the announcement. For instance, [80] examine the reactions of the CDS spread to stock buyback announcements across the short-term and mid-term windows using symmetric event windows of (-7;7), (-15;15), and (-30;30). To analyze the impact of cash dividends, [87] uses and event windows of (-7;7). Furthermore, [48] analyze the impact of M&A announcements using a symmetric event window of (-5;5) days and verify the robustness of their findings by considering 1, 2, 3, and 4 days around the announcement. We aim to build on these works by providing a more nuanced analysis of the temporal dynamics in credit risk reactions to corporate actions. We commence our investigation by assessing if each of the three corporate actions significantly impacts credit risk.

H1: *The announcement of the corporate action has a significant impact on credit risk.*

If a significant impact is found, we then consider its timing. The timing patterns help us understand how the market reacts to corporate events and how credit risk perceptions evolve. We consider three different time periods: before the corporate action announcement (anticipation effect), at the precise moment of the announcement (immediate effect), and in the aftermath of the announcement (delayed effect). Anticipation effects occur when there are significant market speculations or rumors about an upcoming corporate action before the official announcement date. Investors may receive or act upon

private information, leading to an abrupt movement in the CDS spread before the event date. In such cases, the market's reaction precedes the formal event announcement, reflecting early anticipation of the potential impact on the company's financial health and creditworthiness. In the case of dividend changes, the observation that CDS changes precede public announcement aligns with the concept of non-public information utilization in the CDS market, as previously posited by [1] and [84]. This earlier research highlights the substantial incremental information revelation within the CDS market, particularly in response to negative credit news. Furthermore, this pattern is also consistent with prior studies that have shown CDS markets typically exhibit reactions in advance of event announcements. For instance, [66] demonstrated that the CDS market responds to the announcement of central clearing ten days before the official announcement date.

H2a: *Corporate action announcement involves anticipation behaviors, causing a significant movement in the CDS spread prior to the announcement date.*

Immediate impacts are evident when the CDS spread reacts immediately following the corporate action announcement, and any abnormal changes tend to dissipate within a few days of the announcement. This swift dissipation signals that the market rapidly incorporates the new information, resulting in a quick adjustment of credit risk perceptions. During this short period, the market processes the information conveyed by the corporate action announcement, and any significant abnormal movements in the CDS spread within this window reflect the market's rapid response to the new information.

H2b: *Corporate action announcement involves a direct and significant movement in the CDS spread, reflecting the immediate market reaction.*

Delayed effects manifest when the impact of the corporate action announcement unfolds gradually over time. Instead of a sudden movement right after the announcement, the CDS spread experiences a cumulative change that becomes more pronounced over several days after the announcement date. This pattern may arise due to delayed market reactions, or an extended assessment of the event's implications on the company's creditworthiness. For instance, in their study, [80] consider delayed responses of CDS spread to stock buybacks, when the CDS spread may not respond efficiently in a short time, by exploring a one-month and two-month event windows.

H2c: *Corporate action announcement involves a delayed and cumulative movement in the CDS spread after the corporate action announcement date.*

2.2.2 Wealth versus Signaling effects

In the realm of corporate actions and their influence on credit risk, two essential concepts are the signaling effect and the wealth effect. The signaling effect pertains to how certain decisions undertaken by a company can convey valuable information to investors and the market at large. Positive signals indicate that a company is performing well, has strong future prospects, or possesses favorable financial health. In contrast, negative signals may suggest otherwise. The signaling effect can influence investor perceptions and potentially impact financial indicators, including credit risk assessments and CDS spreads. The wealth effect refers to the alteration in the wealth distribution between different stakeholders of a company,

particularly between debtholders and shareholders. Corporate actions can lead to changes in a company's capital structure or financial position, consequently affecting the distribution of wealth. As an illustration, when a company takes actions that raise the allocation of profits or assets to its shareholders, it can enhance the wealth of shareholders, potentially resulting in a corresponding reduction in the wealth of debtholders. Then, the wealth effect can influence credit risk perceptions and, consequently, CDS spreads.

The initial research exploring these effects following stock buybacks encompasses [73] and [42]. [73] reveal positive impacts on stock prices and negative effects on bond prices due to buyback announcements, substantiating the wealth transfer hypothesis. However, they cannot dismiss the signaling effect since they observe that firm value also experiences an increase during open-market stock buyback announcements. [42] employ a refined methodology to validate the positive signaling effect and partially support the wealth transfer effect. In these studies, shifts in bondholders' wealth are gauged by changes in abnormal bond returns or yield spread fluctuations. Recently, [80] argue that CDS spreads offer a superior measure than bond returns or yield spreads for bondholders' wealth. They investigate the impact of stock buyback announcements on CDS spread changes under different macroeconomic conditions using US data. It reveals that smaller firms that announce a larger share repurchase ratio see an increase in abnormal CDS spreads during normal economic periods. In contrast, larger firms see a decrease during crises, particularly over short and medium-term horizons. The findings suggest that wealth transfer effects dominate for smaller firms during normal times, while positive signaling effects are stronger for larger firms' bondholders during crises.

Earnings disclosures, whether positive or negative, have a notable impact on credit risk during their announcement ([16]). The increase in dividends presents a dual impact on stakeholders, providing benefits to stockholders while potentially introducing drawbacks for debtholders, especially in the context of bankruptcy. Consequently, dividend alterations result in a wealth redistribution, impacting both stockholders and debtholders. This perspective, designating such redistribution as the primary consequence of dividend changes, is commonly acknowledged as the wealth transfer effect. An alternative viewpoint, the signaling effect, is advanced by scholars like [8], [9], [74] and [54]. According to this notion, firms adjust dividend levels as a means to communicate shifts in their future prospects. In instances where the market correctly comprehends this signal, dividend raises (cuts) lead to concurrent increases (decreases) in both equity and debt values. Consequently, these two hypotheses converge in their implications for equity valuation, while diverging in their effects on debt valuation. The existing literature has not conclusively determined which effect holds dominance. For instance, [92] and [44] find that unexpected dividend increases (decreases) are associated with positive (negative) bond returns, hence suggesting that the signaling effect dominates the wealth transfer effect for dividend announcements. In contrast, [26] and [59], find that large dividend decreases (increases) are associated with positive (negative) bond excess returns, which suggests that the wealth transfer effect dominates the information content effect in the bond market. Lately, [87] analyze the impact of payout policy on credit risk, drawing insights from the CDS market. They focus on how CDS spreads are influenced by various dividend-related announcements. Notably, their findings reveal a significant elevation in CDS spreads following dividend reduction disclosures, with this effect particularly pronounced during economic downturns and within companies encountering financial difficulties. The magnitude of this reaction is more pronounced for financial institutions, which tend to possess greater opacity. The results show that the information effect of dividend changes dominates the wealth transfer effect.

In contrast to stock buybacks or earning announcements, the influence on credit risk of M&As varies depending on the company's role, whether it is the acquirer or the seller. [35] find that M&As increase the acquirer's risk, especially through risk transfer from the target and leverage changes after M&As. It is also mentioned that M&As generally occur when risk is increasing for other circumstances, as a further reason for the rising default risk of the acquirer. Moreover, [10] and [57] show that if the target has a lower rating than the acquirer, a risk transfer occurs, which might result in a higher risk for the acquirer after the transaction. In contrast, the bonds of target firms tend to increase in value if the acquiring firm's bonds have a stronger rating. Lately, [48] demonstrate that M&A announcements result in an increase default risk for acquirers and a decrease default risk for targets, supporting the wealth transfer effect. Diverging from the diversification literature, mergers within the same industry appear to spread risk more effectively than those across industries. Furthermore, transactions involving stock payment negatively influence the perceived creditworthiness of both acquirer and target companies, possibly due to concerns about overvalued stock.

H3: *The wealth effect holds greater significance than the signaling effect in driving changes of the CDS spread following corporate action announcements.*

2.2.3 Influence of company and deal characteristics

The last set of hypotheses delves into the role of company and deal characteristics in shaping credit risk reactions to corporate actions. Our focus is set on empirically testing variables that have emerged as pivotal factors in existing literature, namely the size of the company, its leverage, liquidity, credit rating, location, and the magnitude of the corporate action announced. Employing these variables, we aim to explain the credit risk response to each of the three corporate actions under examination. The subsequent sections describe the principal hypotheses extracted from the current financial literature.

Buybacks

First, examining the relationship between stock buyback announcements and CDS spreads, the literature underscores the importance of the buyback's magnitudes in credit risk evaluations. We propose that stock buybacks characterized by a high repurchase ratio, indicating a substantial magnitude of repurchased shares, lead to a greater abnormal change in CDS spreads. The underlying assumption is that the magnitude of the buyback itself, regardless of its potential signaling or wealth effects, can independently affect market perceptions of credit risk. Companies engaging in substantial stock buybacks may alter market expectations about their financial stability, operational performance, and future prospects, thereby influencing credit risk evaluations and resulting in larger shifts in CDS spreads. Previous studies ([50] and [73]) highlight that companies with higher repurchase ratios tend to experience improved stock performance following buyback announcements. If this finding extends to CDS, in case of wealth effect holding, we anticipate a positive relationship where a high repurchase ratio results in higher abnormal spread change returns. In addition, [25] investigates CDS spread changes for 53 firms within the S&P 100 index over the sample period from 2011 to 2018. [25] finds significant average abnormal CDS spread changes around buyback announcements. However, [25] cannot find any significant relationship with the

stock buyback magnitude, debt to assets ratio and market capitalization. Lately, [80] nuance this finding by stating that abnormal CDS spreads increase for small-sized firms announcing to repurchase a higher share ratio during the normal period. In contrast, abnormal CDS spreads decrease for big-sized firms regardless of the magnitude of the repurchase ratio during the crisis period.

H4: *Higher magnitude stock buyback announcements cause higher abnormal CDS spread change.*

Besides, we suggest that larger-sized companies are expected to exhibit lower abnormal changes in CDS spreads. This hypothesis is centered on the notion that larger companies often possess greater financial stability, access to resources, and diversified revenue streams, which can act as mitigating factors against fluctuations in credit risk perceptions. The size of a company, irrespective of its signaling or wealth implications, could act as a safeguard against sudden adverse events or market sentiment shifts, resulting in relatively smaller deviations in CDS spreads. Besides, [60] observe that for smaller firms, stock prices exhibit more pronounced reactions to buyback announcements, with strong signaling effects for future profitability. Following [80], we propose that stock buyback announcements from larger companies result in less significant abnormal CDS spread changes.

H5: *Larger companies demonstrate less pronounced abnormal changes in CDS spreads, following stock buyback announcements.*

Cash Dividends

Regarding cash dividend announcements, we suggest that the credit market's response is asymmetric, with negative news having a more substantial impact on CDS spread changes compared to positive news. This hypothesis draws upon behavioral finance theories, suggesting that investors tend to react more strongly to negative information due to loss aversion and heightened risk perception. Then, unfavorable cash dividend announcements might trigger concerns over a firm's ability to meet its financial obligations, causing a more pronounced elevation in CDS spreads. On the other hand, positive dividend news, while potentially indicating healthy financial performance, might not yield a proportional decrease in CDS spreads. This hypothesis has been confirmed in [87], where it is found that CDS spreads exhibit a notably higher reaction to dividend cuts than to dividend raises. [87] also show that the increase in CDS spreads in response to dividend cuts is stronger among firms with high credit risk and negative past stock performance. [87] support their findings using two widely accepted financial economics concepts. First, the debt value is a concave function of a firm's asset value, and thus debt reacts more strongly when the asset value is low. Second, dividend decisions are not exogenous. Dividend cuts are more likely to occur when firms are experiencing distress and closer to the default boundary, while dividend raises or share repurchases are more likely to occur when firms are performing well.

H6: *Negative news surrounding cash dividend announcements have a greater impact on CDS spreads compared to positive news.*

Then, within the context of cash dividend announcements, we state that reference entities possessing lower credit ratings are prone to experience more pronounced market reactions to earnings news, in comparison to their counterparts with higher credit ratings. This hypothesis aligns with the idea that

the creditworthiness of a firm significantly influences market perceptions of its financial health. A lower credit rating implies a higher likelihood of default, and as such, investors are likely to be more sensitive to any signals that might impact the firm's ability to meet its debt obligations. This heightened sensitivity might translate into larger price swings and, consequently, more significant fluctuations in CDS spreads for firms with lower credit ratings. Firms with higher credit ratings, on the other hand, might have a more established track record of financial stability, leading to relatively muted market reactions. This hypothesis has been verified by [87], where they show that CDS reactions to dividend cuts are stronger when firms are in financial distress, both in time series when the economy is in and out of recessions and in cross sections for firms with different levels of default risk.

H7: *Reference entities with lower credit ratings experience more pronounced market reactions to earnings announcements.*

Mergers & Acquisitions

Regarding M&A announcements, we propose that larger acquiring firms, endowed with robust financial resources and established market positions, are better equipped to navigate M&A complexities, mitigating concerns and diminishing the impact on CDS spreads. Conversely, we propose that an increase in the size of the selling firm amplifies abnormal CDS spread changes due to the complexity of financial structures and business segments, leading to pronounced market reactions and subsequent changes in CDS spreads. Lately, [48] confirm that bigger targets are more likely to transfer risk to the acquirer. In their study, [85] emphasize how M&A announcements can influence anticipated efficiency gains, subsequently affecting the combined firm's ability to fulfill its fixed debt obligations. They highlight the fact that the proportional size of the target and the bidder is a crucial factor influencing CDS spread reaction. On one hand, larger targets may offer increased coinsurance of cash flows and contribute more assets to the combined entity, thereby enhancing debt capacity ([47]). However, there exists a limit to the absorption capacity of bidding firms, making the successful implementation of large deals challenging. Consequently, the efficiency gains associated with acquiring smaller targets are expected to be relatively more significant ([7]). Notably, target bondholders may stand to benefit more when the bidding firm is relatively large, as larger bidders are generally more diversified, resulting in lower credit risk at a given leverage ratio ([32]). Consequently, larger deals size, measured as the relative ratio between the size of the deal and the market capitalization of the acquirer, are generally associated with higher CDS spreads in the market ([6] [75]).

H8: *Larger deal magnitude announcements lead to elevated abnormal CDS spreads.*

In common law countries, for instance within the Anglo-American sphere, robust shareholder rights and rigorous disclosure requirements have fostered market-oriented corporate governance models. These models fundamentally consider creditors and various stakeholders as independent entities engaging in arms-length contractual relationships with the company ([51]). Conversely, within the civil law-based governance systems found in Continental Europe, there exists a distinctly different dynamic in the firm-creditor relationship. Here, banks assume the role of concentrated lenders and delegated overseers, playing a critical role in mitigating informational imbalances and addressing agency issues ([27]). This approach diminishes the effectiveness of external market mechanisms. Additionally, other stakeholders cultivate long-term affiliations with the firm, often characterized by closely-held equity structures and

intricate group memberships. Consequently, the augmented influence of banks and other risk-averse stakeholders in shaping corporate decisions should make M&As inherently more bondholder-friendly within the stakeholder-oriented governance frameworks prevailing in Continental Europe. Supporting this perspective, [85] show that bondholder gains in both bidders and targets are systematically higher ceteris paribus in M&As that involve Continental European firms.

H9: *M&A announcements in continental Europe, characterized by stakeholder-oriented governance frameworks, lead to comparatively smaller changes in CDS spreads than in the Anglo-American context.*

2.3 Event study methodology

2.3.1 Measuring abnormal spread change

Drawing on [4], we estimate abnormal changes in CDS spreads (ASC) through a structured approach. Firstly, we determine spread changes that reflect fluctuations in the premium between two successive trading days by calculating the difference in the logarithm of the CDS spread. This process can be expressed as follows:

$$\Delta S_{i,t} = \ln(S_{i,t}) - \ln(S_{i,t-1})$$

where $S_{i,t}$ and $S_{i,t-1}$ represent the spread level (in basis points) for contract i, at time t and $t-1$, respectively.

We use the difference in logarithm instead of the absolute difference, as we expect the market reaction to be proportional to the CDS initial credit risk level. Subsequently, abnormal changes in CDS spreads are defined as the difference between the realized spread change and the normal spread change over the event window.

$$ASC_{i,t} = \Delta S_{i,t} - E[\Delta S_{i,t}|\Omega_t]$$

where $ASC_{i,t}$ represents the abnormal spread change for contract i at time t, $\Delta S_{i,t}$ the realized spread changes and $E[\Delta S_{i,t}|\Omega_t]$ the normal spread changes that would have been observed in the absence of the corporate event, conditioned on past CDS changes Ω available at time t.

In order to modelize the expected change in CDS spread, $E[\Delta S_{i,t}|\Omega_t]$, we explore three primary model specifications centered on the concept of CDS index. The Market model, rooted in [68], seeks to forecast shifts in CDS spreads through the utilization of established CDS indices. Additionally, the sector index model, adapted from [80], establishes a sector index through an equally weighted portfolio of firms within a same sector. Moreover, leveraging random matrix theory on correlation matrices, we introduce a novel index to capture the correlation pattern in individual CDS returns. These models are further enriched by the inclusion of supplementary macroeconomic variables, aligning with the framework proposed by [4].

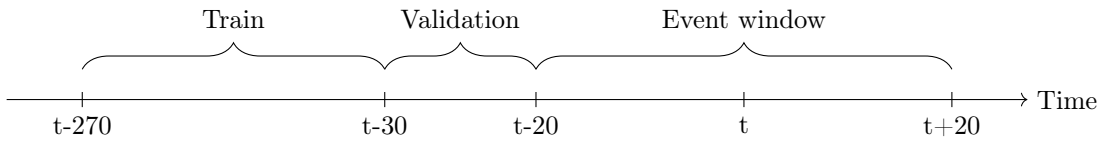

Figure 2.1: Event study periods

In order to assess the validity of each of these different models, we divide our data into distinct periods: a training period, a validation period, and an event window, as depicted in Figure 2.1. The training period spanned from -273 business days (equivalent to approximately one year, [53]) prior to the occurrence of corporate events. This provides a substantial historical context for the models to learn the relationship between the specified variables and the changes in companies CDS spreads. The validation period encompasses the time interval from -30 to -20 business days. This allows for an evaluation of the models' predictive performance on unseen data, specifically focusing on their ability to predict the changes in CDS spreads during a period that remains unaffected by the event. Finally, the event window spans from -20 to 20 business days around each corporate action announcement. This window incorporates relevant data points prior to the event, during the event, and after the event, enabling a comprehensive understanding of the dynamics surrounding the corporate event announcement and its impact on CDS spreads. By employing this structured methodology with a training window of -273 to -30 business days, a validation window of -30 to -20 business days, and an event window of -20 to 20 business days, we aim to build and select the best model at forecasting expected change in CDS spreads, allowing for a robust estimation of abnormal CDS move around corporate actions announcements.

Models specifications

Market Model

The market model is defined as follows:

$$\Delta S_{i,t} = \alpha + \beta \Delta S_{\text{index},t} + \epsilon_{i,t} \qquad E[\epsilon_{i,t}] = 0 \quad \text{Var}[\epsilon_{i,t}] = \sigma^2_{\epsilon_i}$$

where $\Delta S_{i,t}$, $\Delta S_{\text{index},t}$ represent the spread change of the single-name CDS i and of the CDS market index at time t.

In their study, [4] use separate indices for North America and Europe due to distinct restructuring clauses. They employ the Europe iTraxx Europe High Yield index for European single-name CDS and the American CDX High Yield index for American single-name CDS. To validate these choices, we conducted a Pearson's correlation analysis between widely recognized CDS indices spread and our CDS universe using training data. Indices exhibiting the most robust correlation aligned with [4]'s recommendations of the Europe iTraxx Europe index for European single-name CDS and the American CDX IG index for American single-name CDS.

Sector Model

The foundation of this second model lies in portfolio matching, aiming to capture sector-specific dynamics. Mathematically, the sector-level spread change, denoted as $\Delta S_{\text{sector},t}$, is computed as the average of the spread changes of all the CDS within the sector, excluding the spread change of the specific company being analyzed. This can be expressed as:

$$\Delta S_{\text{sector},t} = \frac{1}{|\text{sector}| - 1} \sum_{j \in \text{sector}_i - \{i\}} \Delta S_{j,t}$$

where sector_i and $\Delta S_{j,t}$ represent respectively the sector to which company i belongs to and the spread change of CDS j at time t.

The equation representing the sector model is similar to the market model:

$$\Delta S_{i,t} = \alpha + \beta \Delta S_{\text{sector},t} + \epsilon_{i,t} \qquad E[\epsilon_{i,t}] = 0 \quad \text{Var}[\epsilon_{i,t}] = \sigma^2_{\epsilon_i}$$

where $\Delta S_{i,t}$ and $\Delta S_{\text{sector},t}$ represents the spread change of single-name CDS i at time t and the sector-level spread change at time t.

Correlation Matrix Index Model

This model utilizes the correlation matrix of returns from single-name CDS instruments. To begin, we analyze the eigenvectors of the correlation matrix and compare their corresponding eigenvalues to the Marchenko-Pastur distribution. The Marchenko-Pastur distribution arises as the distribution of eigenvalues for the correlation matrix of independent and identically distributed random variables ([70]). Following [3], the Marchenko-Pastur distribution serves as a reference distribution to assess the significance of eigenvalues in the correlation matrix of CDS returns. By comparing the observed eigenvalues to the upper limit of the support of the Marchenko-Pastur distribution, we identify eigenvalues that exceed the expected range for independent variables. These eigenvalues, which deviate from the Marchenko-Pastur distribution, indicate the presence of genuine correlations in the data and are considered meaningful signals for the construction of the index.

Let C be the correlation matrix of our CDS spread returns. We perform a matrix decomposition to obtain eigenvectors V and eigenvalues Λ satisfying:

$$C = V \Lambda V^{-1}$$

We retain eigenvectors having a corresponding eigenvalues greater than the upper limit of the support of the Marchenko-Pastur distribution $\lambda_+ = \left(1 + \sqrt{\frac{N}{T}}\right)^2$, where N and T represent respectively the number of CDS and the number of observations[2]. We define V_{kept} the matrix of the kept eigenvectors,

[2]To ensure the robustness of the correlation matrix, we need the condition $T > N$. Consequently, we expand the training period to two years for estimating the correlation matrix. However, to train our different correlation matrix index based model, the training period remains at one year.

$$V_{kept} = [v_1, v_2, \ldots, v_k]$$

Each retained eigenvector represents a weighted portfolio, which we use as explanatory variables in our linear model. We finally get the vector $\Delta S_{\text{corr},t}$ of dimension k at time t:

$$\Delta S_{corr,t} = V_{kept}^T \cdot \Delta S_t$$

Finally, we define our correlation matrix index model as:

$$\Delta S_{i,t} = \alpha + \beta \Delta S_{\text{corr},t} + \epsilon_{i,t} \qquad E[\epsilon_{i,t}] = 0 \quad \text{Var}[\epsilon_{i,t}] = \sigma_{\epsilon_i}^2$$

where $\Delta S_{i,t}$ represents the spread change of single-name CDS i at time t.

Controlling for past events

For each model presented, we account for the potential influence of past events on CDS spread changes, introducing a dummy variable taking the value of 1 if any corporate event occurred within a 20-days window prior to the observation date, and 0 otherwise. This dummy variable allows us to assess the specific impact of the event under investigation, while controlling for any residual effects of previous events. During the validation and testing periods, the dummy variable is set to 0 to ensure that the analysis focuses solely on the effects of the events being examined.

Factor Model

The three aforementioned models assume that spread changes are sufficiently explained by common pricing factors. Nevertheless, recent research has unveiled other crucial drivers of fluctuations in credit spreads. Following [4], we chose to include three key variables: the level of a risk-free yield curve, the slope of a risk-free yield curve, and equity volatility. To represent the level of the risk-free yield curve for US companies, we use the 5-year risk-free index USGG. For European companies, we use the GTDEM index. These indices provide a measure of the prevailing risk-free interest rates in their respective regions. The slope of the risk-free yield curve is calculated as the difference between the 10-year and 2-year yields. The same indices, USGG and GTDEM, are used to calculate this slope for US and European companies, respectively. Additionally, the VIX index is used to represent equity volatility for US companies, while the VSTOXX index is used for European companies. These indices measure the expected volatility of the underlying equity markets. By integrating these macroeconomic variables, the equation representing the enhanced models becomes:

$$\Delta S_{i,t} = \alpha + \beta_1 \Delta S_{\text{index},t} + \beta_2 \Delta \text{Level}_t + \beta_3 \Delta \text{Slope}_t + \beta_4 \Delta \text{Vol}_t + \epsilon_{i,t}$$

where $\Delta S_{i,t}$ and $\Delta S_{\text{index},t}$ represent the spread change of single-name CDS i and index CDS at time t. ΔLevel_t, ΔSlope_t, and ΔVol_t denote the changes of the respective macroeconomic variables at time t.

Model selection

In order to select the most appropriate model, we analyze three key metrics: adjusted R-squared (Adj. R^2) and root mean squared error (RMSE) on the training periods, and RMSE on the validation periods. The adjusted R-squared provides a measure of the model's explanatory power, considering the number of predictors included. A higher adjusted R-squared indicates a better ability to explain the variability in the dependent variable. Besides, the RMSE helps to evaluate the model's fit to the observed data. A lower RMSE value indicates a closer alignment between predicted and actual values within the training periods. However, to ensure the model's ability to generalize to new, unseen data, we also assess its performance on the validation periods using RMSE. A lower RMSE on the validation periods suggests the model's capability to accurately predict outcomes beyond the training data, enhancing its robustness and applicability. By considering these three metrics, we aim to select the model displaying the best explanatory power, fit to the training periods, and ability to generalize to new observations.

2.3.2 Measuring influence of corporate actions announcements

Our primary objective is to assess the statistical significance of the influence of corporate actions announcements on CDS spread. For this purpose, we calculate the cumulative abnormal spread change over a specified time interval and evaluate it against the null hypothesis, which posits that its value significantly deviates from zero. In the literature on event-study hypothesis testing, a wide range of tests is employed to investigate the impact of such events. These tests can be broadly classified into two categories: parametric and nonparametric tests. Parametric tests, particularly in the context of event studies, assume that the individual firm's abnormal spread changes are normally distributed. In contrast, nonparametric tests do not rely on any assumption about the distribution of abnormal spread change. This distinction is crucial because the assumption of normality may not always hold, and the presence of outliers can significantly affect the results of parametric tests. In line with the methodology recommended by [86], we employ both parametric and nonparametric tests to ensure the robustness of our analysis. Specifically, we use the t-test, the Wilcoxon sign rank ([91]) test and the Standardized Cross-Sectional (BMP) test ([11]). The t-test is a parametric method assuming normality and equal variance of abnormal spread change, allowing for precise estimation of mean differences, whereas the Wilcoxon Sign Rank test is nonparametric, making no distribution assumptions and is robust against deviations from normality. In addition, we employ the BMP test, a less common but powerful nonparametric test designed specifically for event studies. The BMP test assesses the distribution of abnormal spread changes, determining if it significantly differs from a symmetric distribution, which is particularly valuable when non-symmetrical abnormal returns are a concern. For this final test, we present its formulation to shed light on its distinct characteristics.

The BMP test is designed to address the bias arising from an event-induced increase in variance. Indeed, [11] show that when an event causes even minor increases in variance, the most commonly used methods tend to reject the null hypothesis of zero average abnormal return too frequently when it is true. For the BMP test, the test statistic t_{BMP} is computed as follow:

$$t_{BMP} = \sqrt{N}\frac{\overline{\text{SCAR}}}{\sigma_{\text{SCAR}}}$$

Here, $\overline{\text{SCAR}}$ and σ_{SCAR} represents respectively the mean and the standard deviation of the standardized cumulative abnormal return, defined as below.

$$\overline{\text{SCAR}} = \frac{1}{N}\sum_{i=1}^{N}\text{SCAR}_i$$

$$\sigma^2_{\text{SCAR}} = \frac{1}{N-1}\sum_{i=1}^{N}(\text{SCAR}_i - \overline{\text{SCAR}})^2$$

These statistics rely on the computation of SCAR_i, which represents the forecast-error-corrected cumulative abnormal return. SCAR_i is calculated as:

$$\text{SCAR}_i = \frac{\text{CAR}_i}{\sigma_{\text{CAR}_i}}$$

Where CAR_i is the cumulative abnormal return and σ_{CAR_i} is the forecast-error-corrected standard deviation for CDS i. For instance, for the market model:

$$\sigma^2_{\text{CAR}_i} = \sigma^2_{\text{AR}_i}\left(L + \frac{L}{M} + \frac{\sum_{t\in\text{EW}}(\Delta S_{\text{index},t} - \overline{\Delta S_{\text{index},t}})^2}{\sum_{t\in\text{TP}}(\Delta S_{\text{index},t} - \overline{\Delta S_{\text{index},t}})^2}\right)$$

Here, $\sigma^2_{AR_i}$ represents the variance of abnormal returns, while M and L represent the respective numbers of days seen in the training period and event window. $\Delta S_{\text{index},t}$ and $\overline{\Delta S_{\text{index},t}}$ denote the index returns and their average, respectively. TP and EW represent the training period and the event window respectively. Finally, the approximate null distribution of the test statistic t_{BMP} follows a t-distribution with $M-1$ degrees of freedom.

2.4 Results and discussions

2.4.1 Data

We obtain CDS data for the period from 2012 to 2023 from Markit. Our analysis focuses specifically on clearable[3] single-name CDS as defined by ICE (Intercontinental Exchange), with maturities of five years as they are the most liquid ones. In addition, to enhance the homogeneity of the sample, we limit our analysis to CDS contracts that are denominated in US dollar and EURO currency and that include the no-restructuring (XR) or modified-restructuring (MR) clause ([34]). To generate its single-name CDS composite spread and price for a given contract, Markit collects daily pricing data from a group of 22

[3]Clearable single-name CDS refer to contracts that are standardized and eligible for clearing through ICE Clear Credit, a leading central counterparty for credit derivatives.

global banks. This data comprises daily closing prices from contributing banks' books of record or the most recent prices from their trading desks. By averaging the CDS spreads from at least two contributing banks that pass its data cleaning tests, Markit produces the daily CDS composite spread and price.

Data related to equity, macroeconomic variables and corporate actions are collected from Bloomberg. The corporate actions data from Bloomberg is particularly valuable due to its granularity, providing essential details such as the event date, corporate action type, and associated financial indicators. Our final sample consists of 3,461 stock buybacks, 3,542 earning announcements and 3,180 M&As by 339 companies over 10 sectors.

In order to reconcile the data from Markit and Bloomberg, we manually map the Markit tickers to the corresponding Bloomberg tickers. This involves extensive manual cross-referencing and verification to ensure accurate alignment between the two datasets. This process allows us to establish a reliable connection between the corporate entities in the Markit dataset and their counterparts in the Bloomberg dataset.

2.4.2 Measuring abnormal spread change

Table 2.1 presents performances of the various models we evaluate to forecast CDS spread changes. Notably, the sector index model, bolstered by macroeconomic variables, demonstrates superior performances. This model surpasses others, displaying elevated adjusted R^2 values and demonstrating lower RMSE across both the training and validation periods, for all corporate actions. This outcome emphasizes the pivotal role of considering sector-specific dynamics when developing models for predicting CDS spreads. It is important to highlight that the straightforward sector index model outperforms the more complex correlation matrix index model. Comparable findings are evident in the research conducted by [14], where they demonstrate that a simple methodology based on the market model performs well under a wide variety of conditions for predicting price performances. In some situations, even simpler methods such as the constant-mean model which do not explicitly adjust for marketwide factors or for risk perform no worse than the market model. This suggests that the variance of abnormal returns remains largely unchanged, indicating the limited effectiveness of employing more sophisticated models. All subsequent analyses and findings are grounded in the sector index model augmented with macroeconomic variables.

Corporate Action	Model	adj. R^2 train	RMSE train	RMSE val
Buyback	Sector	0.53	0.02	0.017
	Simple index	0.31	0.024	0.021
	Custom index	0.35	0.023	0.02
	Sector + macro	**0.54**	**0.019**	**0.017**
	Simple index + macro	0.33	0.023	0.021
	Custom index + macro	0.4	0.022	0.021
Cash dividends	Sector	0.51	0.02	0.017
	Simple index	0.32	0.02	0.02
	Custom index	0.34	0.022	0.019
	Sector + macro	**0.54**	**0.019**	**0.019**
	Simple index + macro	0.34	0.022	0.023
	Custom index + macro	0.4	0.021	0.022
M&A	Sector	0.56	0.021	0.018
	Simple index	0.36	0.025	0.019
	Custom index	0.37	0.023	0.019
	Sector + macro	**0.59**	**0.02**	**0.02**
	Simple index + macro	0.39	0.024	0.021
	Custom index + macro	0.43	0.022	0.022

Note: For stock buybacks, N_{train}=944,853, N_{val}=34,610.
For cash dividends, N_{train}=966,966, N_{val}=35,420.
For cash dividends, N_{train}=868,140, N_{val}=31,800.
Here, N corresponds to the number of observations involved in training and validating the different models.

Table 2.1: Models performances

2.4.3 Measuring influence of corporate actions announcements

Acknowledging the potential diversity in effects associated with various attributes of the corporate actions, we partition our samples to assess the impact of corporate action announcements on credit risk. For M&A, we follow [48] and divide samples into sellers and buyers companies, recognizing the potential divergence in credit risk impact between these two distinct roles. Seller companies are companies that have been bought during the corporate action. In contrast, the buyer company is the company buying. All the retained companies have CDS instruments traded on their debt. Besides, when scrutinizing the effects of cash dividends, we follow [87] and stratify our sample into two distinct categories: dividends raise, where ratio of dividends over market cap is greater than the ratio of the previous cash dividend exercise, and dividends cut, where this ratio falls below the one of the last exercise. For measuring the impact of stock buybacks, we refrain from sample division, as the effect's sign of this corporate event are assumed to be uniform and independent of specific buybacks attributes ([80]).

Impact and temporal dynamics

We recognize that each corporate action possesses unique characteristics, resulting in distinct dynamics surrounding its announcement. To comprehensively capture and evaluate these dynamics, we adopt a methodology which involve considering all possible event windows, spanning from 20 days preceding the announcement to 20 days following it. For each of the possible event windows, we calculate the cumulative abnormal spread changes (CASC), measuring deviations from the expected or normal behavior, as determined by our predictive model. Subsequently, we use the t-test, the Wilcoxon sign rank test and the BMP test in order to assess the significance of these cumulative changes.

To visualize the temporal dynamics, for each event window, we calculate the geometric mean of the p-values for the 3 statistical tests, and leverage heatmap representations to visualize the significance levels across the different event windows. These heatmaps highlight event windows during which CDS spreads exhibits strong and significant abnormal behavior in response to corporate action announcements. Our primary objective in studying these dynamics is to identify the most suitable start and end dates, forming the event window displaying the most substantial abnormal spread changes. To facilitate this selection, we also represent the mean of p-values across start (by columns) and end (by rows) dates at the top and right hand-side of the heatmaps. Specifically, we select dates where the mean p-value is initially low and then rapidly increases, signifying the point at which abnormal spread behavior returns to normal.

To illustrate our approach, we provide three theoretical heatmap examples[4]. Figure 2.2 depicts the expected heatmap in scenario where a significant abnormal spread change occurs on the day of the announcement, without preceding rumors, and with the spread returning to its usual pattern immediately afterward. Figure 2.3 illustrates situation where noteworthy movements occur on the announcement day, possibly influenced by rumors or shifts preceding the official announcement. Finally, Figure 2.4 illustrates the case of substantial abnormal changes on the announcement day, gradually dissipating in the days following the announcement, and with anticipation behaviors. We now delve into the dynamics associated with the announcement of the three corporate actions under examination.

[4]These heatmaps are illustrative and based on simulated data, intended solely to visualize stylized scenarios for explanatory purposes. They do not reflect actual empirical results.

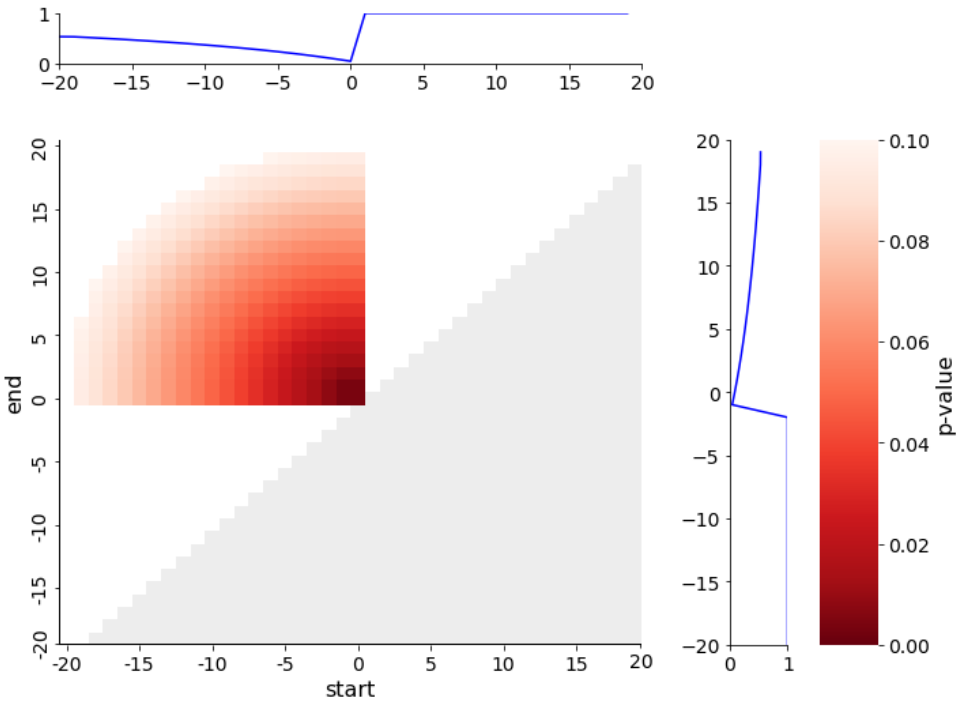

Figure 2.2: Simulated credit risk reaction at announcement date

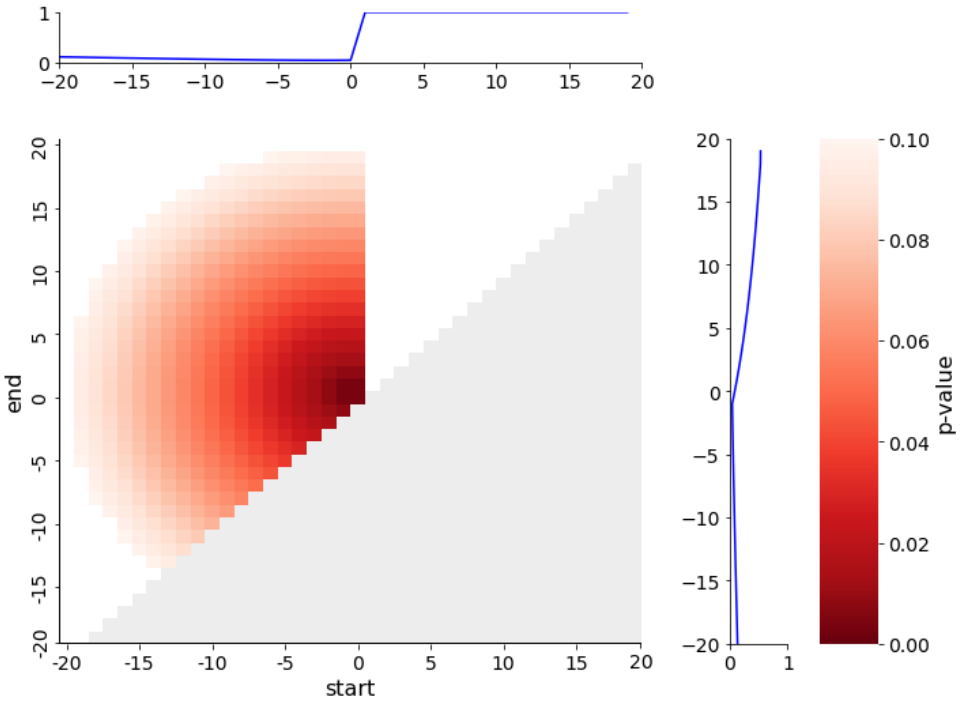

Figure 2.3: Simulated credit risk reaction before announcement date

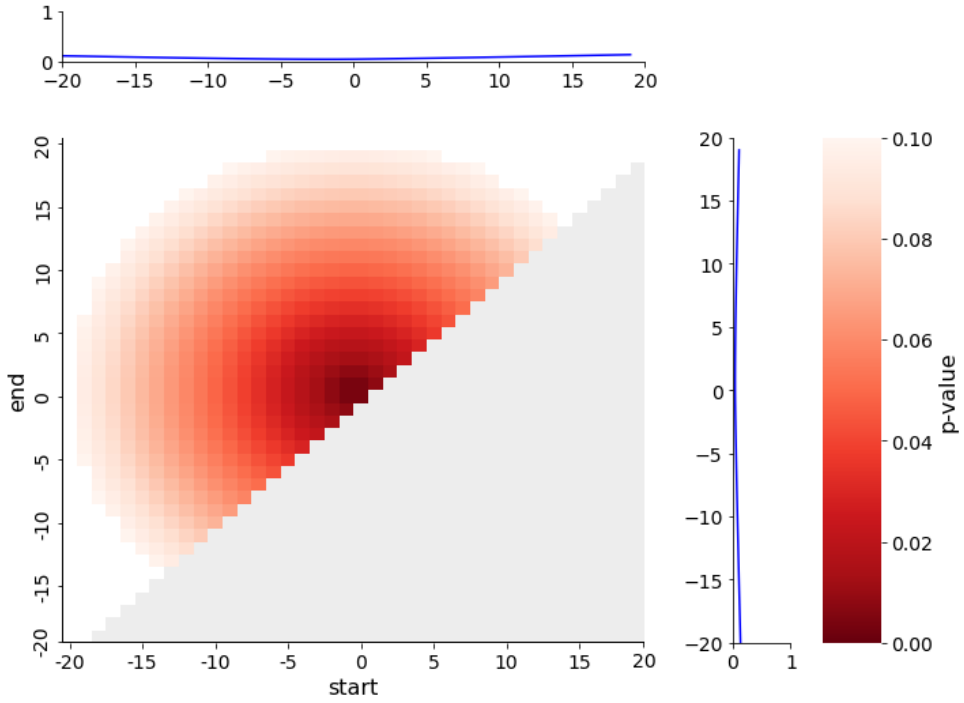

Figure 2.4: Simulated credit risk reaction before and after announcement date

Buyback

Figure 2.5 indicates a significant impact of stock buyback announcements on credit risk. We observe a notable abnormal spread increase for the event window spanning from the day of the announcement to the 2 days right after. We can draw a parallel between the behavior of CDS spreads in this scenario and the theoretical heatmap in Figure 2.2 where no significant abnormal spread change is observed prior event announcement, and where spread change returns to its usual pattern immediately afterward. Besides, it is important to be aware of the potential sensitivity of the statistical tests used in the event study analysis. As shown in Figure 2.5, event window starting 14 to 9 days before the announcement date exhibits high abnormal spread changes. This variation emphasizes the need of examining the broader dynamics of spread changes when selecting event windows, and not relying on arbitrary selected event windows, which might be influenced by random fluctuations. Given the notable change observed on the day of the announcement, we consider the event window spanning from 0 to 2 days after the announcement date when analyzing stock buybacks.

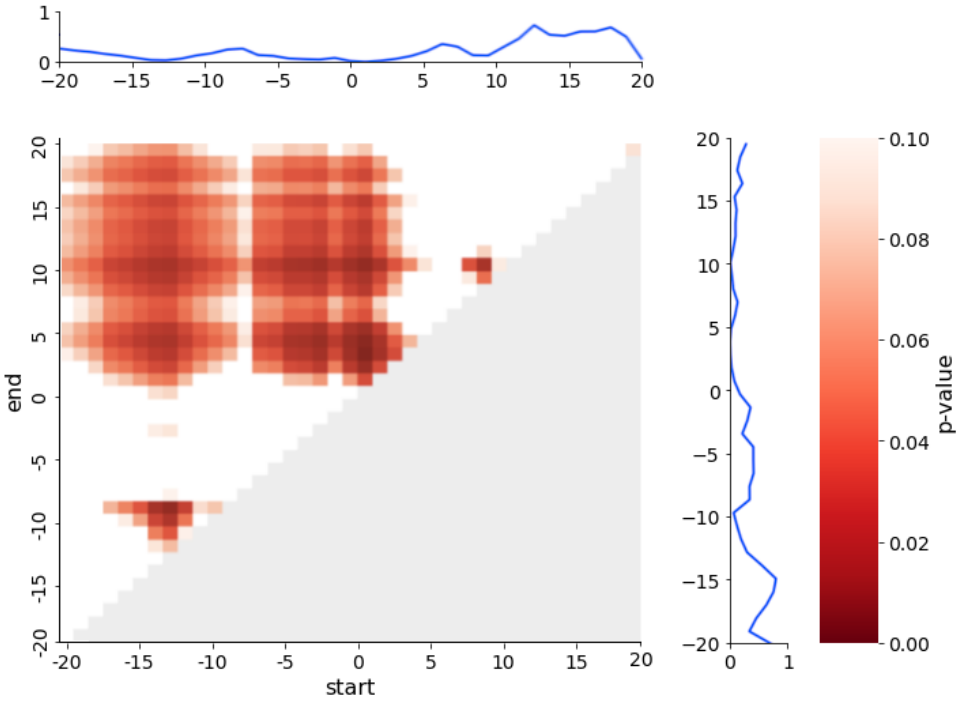

Figure 2.5: Dynamics around buyback announcement

Dividend raises

The analysis of Figure 2.6 reveals a pattern characterized by its immediacy and absence of anticipation behaviors preceding the event announcement. As depicted in Figure 2.2, the heatmap closely resembles scenarios where significant abnormal spread changes occur on the day of the announcement, without the influence of preceding anticipations, and with the CDS spreads returns returning to their normal pattern immediately afterward. This observation suggests that cash dividends raise events prompt an instantaneous and substantial response in the CDS market, with no delayed or cumulative effects. The data indicates that the impact is confined to a brief window, typically not extending beyond two days, highlighting the event's direct and immediate influence on credit risk perception.

Dividend cuts

In contrast to cash dividend raise events that appear to trigger a straightforward and rapid adjustment in CDS spreads, cash dividends cuts exhibit more complex dynamics, including extended anticipation and a more pronounced and lasting impact. As evident in Figure 2.7, the dynamics surrounding cash dividends cuts display a notable anticipation behavior, commencing approximately 9 days before the official announcement date. This anticipatory phase aligns with a phenomenon often observed in the literature, where market participants, in response to signals or expectations, adjust their positions ahead of corporate actions ([39]). In the case of cash dividends cuts, negative news often carry a higher degree of uncertainty and potential financial instability, which may contribute to heightened market volatility during the anticipation phase. Furthermore, we observe that the CDS spread returns to normal roughly 10 days after the announcement date. This represents a more extended adjustment period, compared to dividends raises where credit risk absorbs the news within 2 days. This extended impact duration

following cash dividends cuts aligns with existing research findings ([21]) that suggest that negative corporate events, such as dividend cuts, tend to elicit more prolonged reactions in credit markets due to the greater uncertainty. The literature has often highlighted the lingering effects of negative news on credit risk perception, underscoring the need for a longer adjustment period before the credit markets stabilize ([62]).

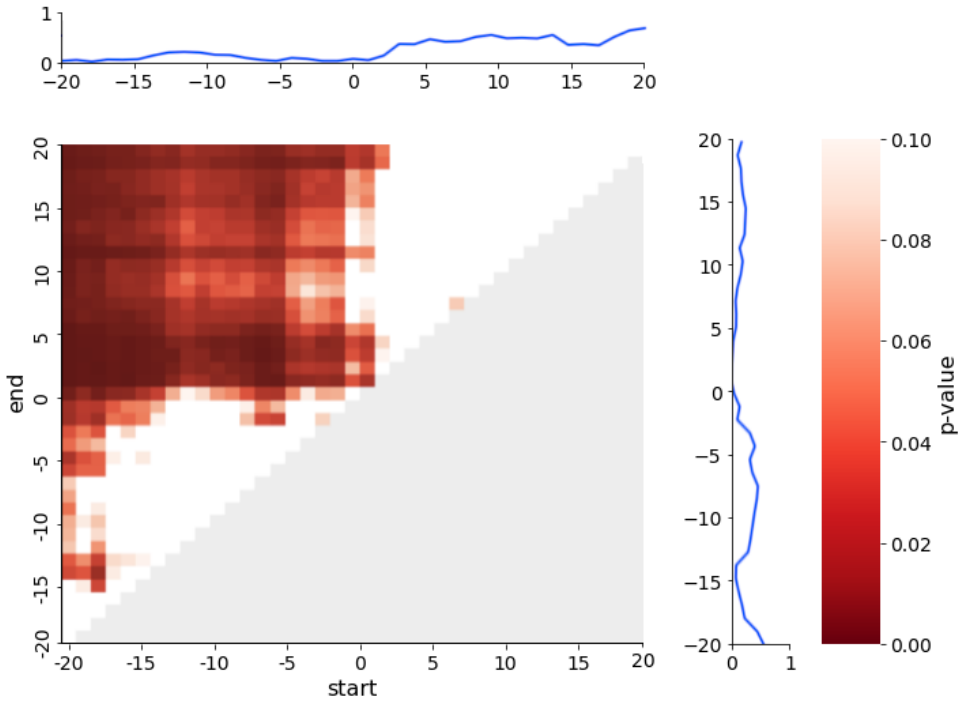

Figure 2.6: Dynamics around cash dividend raises

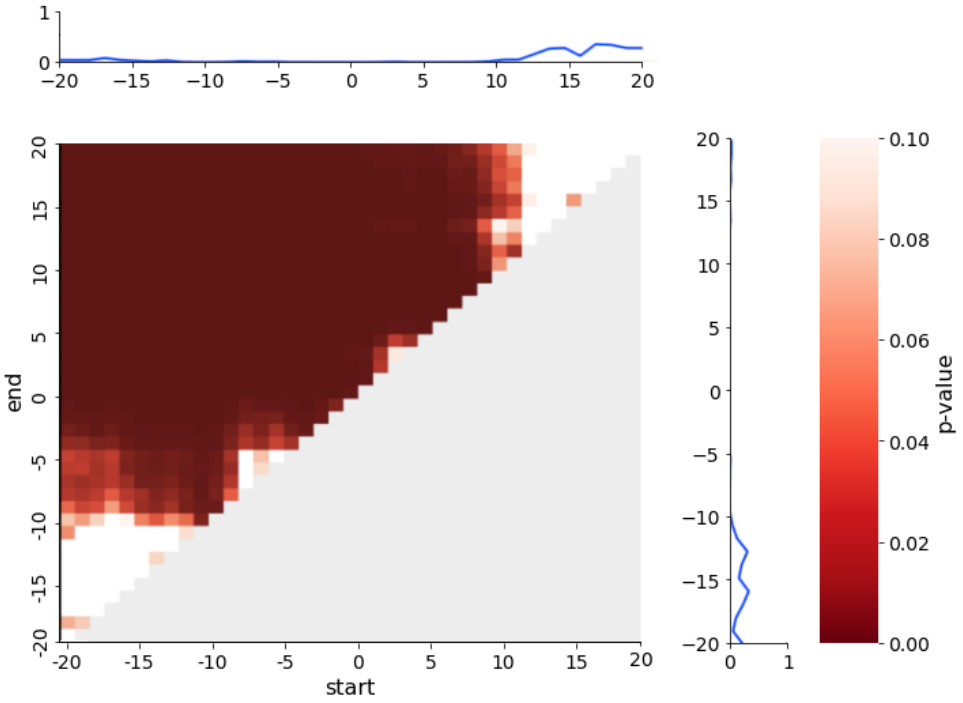

Figure 2.7: Dynamics around cash dividend cuts

M&A seller

The impact of M&A announcements for seller companies displays a significant temporal pattern, extending across a broad and asymmetric event window. More precisely, this effect spans from 15 days prior to 2 days after the announcement date, suggesting a substantial period of anticipation leading up to these announcements. This heightened sensitivity of seller companies' credit risk during M&A events might be influenced by factors like privileged information or speculative trading. The observed behavior corresponds to the trend seen in Figure 2.3, potentially linked to pre-announcement market adjustments or rumors, while no discernible impact is evident in data collected from 2 days after the announcement. This temporal asymmetry underscores the distinct nature of M&A activities for sellers. While the market demonstrates strong anticipation of such events, the impact is short-lived. To sum up, the dynamics of M&A announcements for seller companies feature a notable anticipatory phase and a rapid return to market stability.

M&A buyer

For buyer companies involved in M&A activities, the observed impact appears to be less pronounced in magnitude compared to seller companies. Similar to seller companies, buyer firms experience a notable abnormal spread change starting 15 days before the formal announcement, indicating a significant anticipatory phase. This early abnormal spread change suggests that market participants may anticipate M&A actions also for buyer companies, potentially driven by information leaks or strategic positioning. Moreover, in contrast to seller companies, the impact of M&A announcements on buyer companies seems less readily absorbed by the market, extending over four days post-announcement. This temporal asymmetry highlights the unique nature of M&A events for buyers, characterized by both anticipation and a

more prolonged impact.

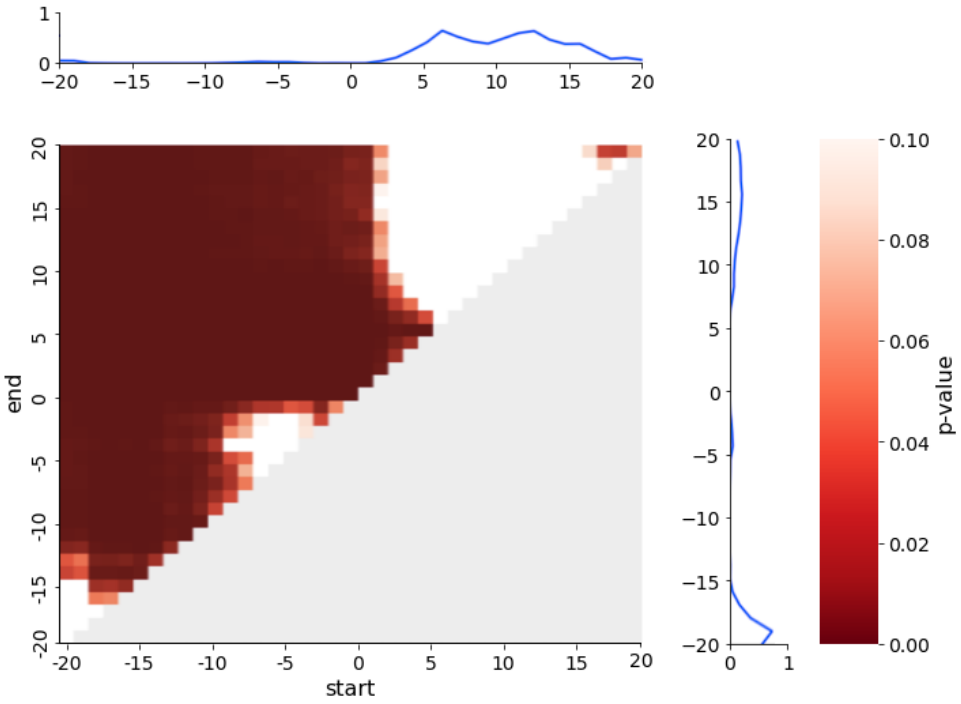

Figure 2.8: Dynamics around M&A seller companies

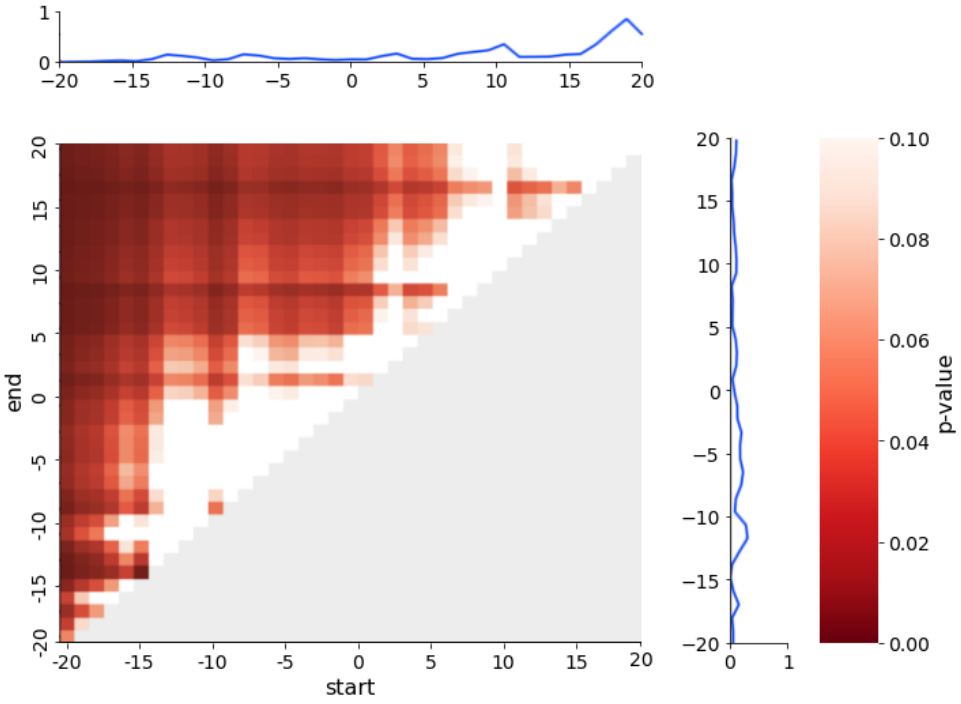

Figure 2.9: Dynamics around M&A buyer companies

In summary, our investigation underscores the importance of statistically isolating windows of interest for event studies. The method we propose to select event window avoid reliance on arbitrarily chosen event

	Buyback	Cash divs		Mergers & Acquisitions	
		Raise	Cut	Buyer	Seller
Hypothesis 1	Y	Y	Y	Y	Y
Hypothesis 2a			Y	Y	Y
Hypothesis 2b	Y	Y			Y
Hypothesis 2c			Y	Y	

Note: Y corresponds to verified hypothesis.

Table 2.2: Impact and temporal dynamics hypotheses

windows, which may be prone to outliers influence. Such precision is pivotal for event study analysis, contributing to robustness and validity of the conclusions.

Wealth versus Signaling effects

Regarding stock buybacks, Table 2.3 highlights a significant wealth transfer effect between shareholders and debtholders, consistent with earlier studies ([23], [90], [50], [73], [2]). The wealth transfer effect is represented by the significant and positive mean abnormal spread change of 0.0033 within the event window from 0 to 2 days after stock buybacks announcements. Both the t-test and the BMP test indicate a p-value of 0.01. Additionally, the geometric mean of p-values (0.03) further supports the statistical significance of the mean abnormal spread change at a 5% confidence level. While stock buybacks potentially enhance shareholder value by reducing outstanding shares, they concurrently reduce the value of the company's debt holdings. This reduction in cash reserves may raise concerns among debtholders and prompt the wealth transfer effect. However, it is essential to approach these findings cautiously, as the Wilcoxon statistical test fails to identify any significant abnormal spread change. This aligns with existing research on stock buybacks, suggesting that the distinction between signaling and wealth effects might not always be straightforward ([73], [42], [80]).

In the context of earnings announcements, our research once more underscores the preeminence of the wealth transfer effect. More precisely, we find a discernible increase of 0.0061 in abnormal spread changes within the 0 to 2 days following earnings call announcements when dividends are raised. This observation highlights the significance of the wealth transfer effect during favorable earnings announcements. The mechanisms behind this phenomenon are rooted in the impact of dividend raises on a company's financial structure. When dividends increase, the firm allocates a larger portion of its earnings to equity shareholders, leaving less cash available for debt repayment and potentially increasing leverage ([35]). This shift in wealth distribution triggers apprehensions among the debtholders, as it may raise concerns about the company's ability to meet its debt obligations, especially in situations involving financial distress or bankruptcy. Conversely, when dividends are cut, we observe a mean abnormal spread change of -0.0153 within the window of -9 to 10 days after earnings call announcements. This negative abnormal spread change underlines the wealth transfer effect associated with dividend cuts. In such cases, the company's decision to reduce dividends is perceived as a disadvantage to equity shareholders and is viewed positively by debtholders. A dividend cut may indicate a more conservative financial approach, potentially bolstering the company's ability to meet debt obligations. Our findings align with previous studies in the field such as [26] and [59]. Additionally, our study underscores the remarkable market asymmetry

in reactions to dividend cuts compared to raises. Indeed, the magnitude of the negative mean abnormal spread change for dividend cuts, when contrasted with the positive value for dividend raises, underscores the market's heightened sensitivity to negative news. This result has been confirmed in [87], where they find that CDS spreads exhibit a notably higher reaction to dividend cuts than to dividend raises.

Regarding M&A announcements, our research emphasizes the predominant role of the wealth transfer effect in shaping the dynamics of CDS spreads. Specifically, we find a noteworthy mean abnormal spread change of 0.0086 for buyer companies within the event window of -15 to 4 days surrounding the M&A announcement. This observation underscores the significance of the wealth transfer effect during such announcements when a company is positioned as the buyer. The mechanisms behind this phenomenon can be elucidated by examining the impact of M&A activities on a firm's financial structure. When a company announces its role as a buyer in an M&A transaction, it often implies that the firm will be allocating resources towards the acquisition, potentially leading to increased leverage and concerns among debtholders ([35]). This shift in wealth distribution between equity shareholders and debtholders, with an emphasis on the potential disadvantages for the latter, amplifies the importance of wealth redistribution during M&A announcements. In contrast, when a company is designated as the seller in M&A transactions, we observe a substantial mean abnormal spread change of -0.0263 within the event window of -15 to 2 days surrounding the announcement date. This negative abnormal spread change supports a wealth transfer effect associated with seller firms. In this context, the decision to divest assets or business segments is often interpreted as a wealth transfer from the seller to the buyer, and this may be perceived as disadvantageous to the equity shareholders of the selling company while potentially raising confidence among debtholders. Moreover, target firms typically experience positive effects by becoming integral parts of larger organizations, enabling them to overcome market frictions. This aspect highlights the potential of M&A to mitigate financial constraints ([30]). Consequently, M&A activities have the potential to facilitate enhanced financing for target firms, securing favorable options through improved access to capital markets ([45]). This expanded access to capital, in turn, augments investment opportunities for target firms, contributing to the maximization of firm value and potentially lowering the default probability ([64]). Our findings are consistent with prior research, such as [19] and [48]. Importantly, the magnitude of the coefficient for sellers, which stands at -0.0263, is notably greater than that for buyers (0.0086), signifying that the intensity of the impact on seller companies is more pronounced than on buyer firms.

Influence of company and deal characteristics

The preceding Section emphasizes the significant influence of the wealth transfer effect in shaping credit risk dynamics in response to corporate action announcements. We now delve into a detailed study of the factors underpinning credit risk fluctuations during these corporate event announcements. As explained in Section2.3.1, we leveraged market-related variables to measure abnormal spread changes. To understand the determinants of these abnormal spread changes, our focus shifts to company and deal characteristics. We begin by introducing a comprehensive set of company-specific and corporate action attributes, which include financial metrics and structural characteristics, all of which have the potential to act as determinants of credit risk fluctuations in the context of corporate actions. Following this introduction, we consider these variables as determinants of the changes in credit risk by running a

	Buyback	Cash divs		M&A	
		Raise	Cut	Buyer	Seller
Event window	[0 ; 2]	[0 ; 2]	[-9 ; 10]	[-15 ; 4]	[-15 ; 2]
N	3,461	1,763	1,779	2,050	1,130
Mean abnormal spread change	0.0033	0.0061	-0.0153	0.0086	-0.0263
t-test	0.01**	0***	0***	0.02**	0***
Wilcoxon	0.57	0.02**	0***	0.06*	0***
BMP	0.01***	0.32	0***	0.01***	0.09*
Geometric mean	0.03**	0.04**	0***	0.05**	0***

Note: * p < 0.1; ** p < 0.05; *** p < 0.01.
Here, N corresponds to the number of corporate actions involved in computing the different test statistics.

Table 2.3: Tests significance

	Buyback	Cash divs		Mergers & Acquisitions	
		Raise	Cut	Buyer	Seller
Hypothesis 3	Y	Y	Y	Y	Y

Note: Y corresponds to verified hypothesis.

Table 2.4: Wealth versus Signaling effects hypotheses

standard regression model, in which we use the cumulative abnormal change in CDS spreads within the defined event windows (as defined in Table 2.3) as the dependent variable.

Firstly, a key metric in our analysis is the cumulative abnormal stock returns (CAR), a measure that unveils the relative performance of individual stocks against the broader market. CAR is calculated through a regression model where stock returns R_i for company i is the dependent variable and index returns R_{index} act as the independent variable[5]:

$$R_i = \alpha + \beta R_{index} + \epsilon_i \qquad E[\epsilon_i] = 0 \quad \text{Var}[\epsilon_i] = \sigma^2_{\epsilon_i}$$

The cumulative residual of this regression CAR_i, encapsulates stock-specific variations unexplained by overall market movements.

$$CAR_i = \sum_{t=0}^{T} R_{i,t} - (\hat{\alpha} + \hat{\beta} R_{index,t})$$

where CAR_i represents the cumulative abnormal stock return of stock i during the event window spanning T days[6].

Although we refrain from interpreting the value of the coefficient of CAR, it allows us to account for

[5]We used the EURO STOXX 50 index for European companies and the S&P 500 index for US companies.
[6]The event windows used to compute cumulative abnormal stock returns are the ones selected in table 2.3

abnormal equity returns during the announcement window. This provides a robust means to measure the relationship between company characteristics and additional abnormal spread changes not explained by equity market movements, thereby offering insights into the incremental information conveyed to the credit market during corporate action announcements.

In addition to CAR parameter, our model captures company specific attributes. We incorporate well-established company-specific variables, widely recognized as significant determinants of CDS spreads in the literature. Following [80], to reflect company size, we use the market capitalization measured as the total value of the outstanding shares in the market. Companies of greater size, often linked to higher market capitalization and bolstered financial stability, resource access, and diversified revenue streams, have the potential to alleviate credit risk perceptions. Then, we use the leverage ratio to measure company's capital structure and the proportion of debt financing relative to equity financing. This metric measures the extent to which a company utilizes debt in relation to its equity. Following [58], the leverage ratio is computed as the ratio of the book value of debt[7] to the sum of the book value of debt and the market value of equity. Numerous studies ([31], [28], [36] and [58]) affirm that the leverage ratio is a key determinant of changes in CDS spreads. Additionally, previous literature ([88], [12] and [89]) reports that CDS spread changes are not fully explained by structural variables such as leverage ratio and that the unexplained parts can be associated with CDS liquidity. The liquidity reflects the ease with which CDS contracts can be bought or sold without significantly impacting their market prices. In our analysis, we measure the liquidity of the CDS contract for each firm using the natural logarithm of the CDS market depth, i.e. the number of quote contributors for the corresponding five-year CDS spread at the announcement date. Following [73] and [80], we measure companies credit risk using the credit ratings given by the major agencies such as S&P, Moody's, and Fitch. Our approach involves converting these ratings into binary variables, 1 denoting investment-grade status and 0 high-yield status. Then, we aggregate these binary values by calculating the mode, yielding a synthesized credit rating. To capture regional distinctions, we introduce a binary variable, taking a value of 1 when the company is identified as a US company, and 0 otherwise.

Finally, magnitude of the corporate action is incorporated into the regression as deal characteristic. The event magnitude is key as it quantifies the impact or significance of each corporate event. The calculation of event magnitude varies depending on the event type. For cash dividends, the magnitude is determined by dividing the dividend amount by the market capitalization of the company. For stock buybacks, the magnitude is defined as the repurchase ratio, i.e. the percentage of shares targeted for buyback at the announcement date. Finally, following [48] the magnitude for M&A events is the ratio between the value paid for the target and the market value of the acquirer 41 days prior to the announcement date.

We now delve into the empirical results, offering comprehensive insights into the relationships between abnormal stock returns, size, leverage, liquidity, credit ratings, region, magnitude and abnormal CDS spread changes following corporate action announcements. The regression results offer nuanced insights into the hypotheses related to stock buybacks. Hypothesis 4, positing that a higher magnitude of stock buyback announcements leads to a greater abnormal spread change, finds empirical support. The positive and significant coefficient (0.0612) for the variable *magnitude* in the regression table aligns with our expectations, implying that substantial buybacks, characterized by a high repurchase ratio, indeed

[7]The book value of debt includes "long-term debt" plus "debt in current liabilities" plus "preferred stock."

correspond to larger abnormal changes in CDS spreads. Regarding Hypothesis 5, which suggests that larger-sized companies are expected to exhibit lower abnormal changes in CDS spreads, the coefficient for *size* is positive but statistically insignificant. This suggests that, on average, company size does not significantly correlate with the impact of stock buybacks on CDS spreads. The negative sign for the *leverage* variable (-0.011) suggests that, holding other factors constant, companies with higher leverage tend to exhibit a relatively lower abnormal spread change after a buyback announcement compared to companies with a low leverage ratio at the beginning. While the sign of the *leverage* coefficient in our study is consistent with that in [80], our analysis uncovers a statistically significant coefficient, presenting a notable departure from the non-significant result reported in [80]. This is likely explained by the fact that highly leveraged firms may already be perceived as having higher credit risk before the buyback announcement. When these firms engage in buybacks, the market interprets it as a strategic move to enhance financial stability and alleviate credit risk concerns, leading to a relatively smaller abnormal spread change. It is worth noting that the adjusted R-squared values for the regression model is low but relatively high compared to other corporate actions regressions, indicating that the selected independent variables explain a relative high portion of the variation in the cumulative abnormal spread change, compared to other corporate actions.

Turning to cash dividends, the statistically significant negative coefficient for *const* (-0.0356) in the context of dividend cuts substantiates our expectation of a reduction in abnormal CDS spreads for companies making such announcements. This finding is consistent with the wealth transfer effect hypothesis, affirming that conservative dividend actions contribute to a decrease in perceived credit risk. The significance of the negative coefficient for dividend cuts, reaching the 5% level, stands in contrast to the non-significant coefficient for dividend raises. This asymmetry in significance levels, also observed in Figure 2.8 and 2.9, further underscores hypothesis 6, stating that negative news has a more substantial impact on CDS spreads compared to positive news. This empirical validation, consistent with [87] emphasizes the asymmetric market response to cash dividend announcements, highlighting the heightened influence of unfavorable news on credit risk perceptions. Consistent with our expectations outlined in Hypothesis 7, the coefficient for the *rating* variable is negative (positive) for dividend raises (cuts). However, the result lacks statistical significance, introducing some inconsistency with the findings in [87]. Despite the absence of statistical significance, it is noteworthy that the signs of the coefficients align with the anticipated direction. Indeed, the underlying premise suggested that lower-rated entities would experience more pronounced market reactions, aligning with the notion that creditworthiness significantly influences market perceptions. Moving to the *leverage* variable in the context of dividend raises, the negative and statistically significant coefficient (-0.0118) mirrors the findings previously established for stock buybacks. This result is consistent with the expectation that companies with higher leverage tend to exhibit a relatively lower abnormal spread change after a cash dividend raise announcement. Finally, contrary to expectations, the coefficient for company *size* is positive and statistically significant, indicating that larger companies tend to experience a higher increase in abnormal spread change following dividend raises compared to their smaller counterparts. This unexpected result suggests that factors associated with larger firms, such as their financial standing, market presence, or investor perceptions, may contribute to a more pronounced reaction in credit risk following positive dividend news. This finding challenges conventional assumptions and underscores the complexity of the relationship between company characteristics and credit risk dynamics during specific corporate events.

	Buyback	**Cash Dividends**		**Mergers & Acquisitions**	
		Raise	**Cut**	**Buyer**	**Seller**
const	-0.0013	-0.0053	-0.0356**	-0.0137	-0.0072
	(0.819)	(0.342)	(0.035)	(0.395)	(0.747)
CAR	-0.0035***	-0.0019***	-0.0073***	-0.0039***	-0.0051***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
size	0.0120	0.0487*	0.1149	0.0613	0.0088
	(0.657)	(0.093)	(0.236)	(0.318)	(0.914)
leverage	-0.0110*	-0.0118*	0.0183	0.0112	-0.0141
	(0.086)	(0.079)	(0.374)	(0.685)	(0.688)
ln(liquidity)	-0.0002	0.0005	0.0024	0.0008	-0.0105**
	(0.860)	(0.656)	(0.453)	(0.788)	(0.012)
rating	0.0027	-0.0015	0.0061	0.0026	0.0126
	(0.288)	(0.540)	(0.438)	(0.774)	(0.286)
region	0.0057	0.0058	0.0242	0.0169*	0.0162
	(0.136)	(0.084)	(0.012)	(0.081)	(0.154)
magnitude	0.0612***	-0.0508	0.0519	0.0173*	0.0110**
	(0.002)	(0.287)	(0.439)	(0.095)	(0.046)
N	3,461	1,763	1,779	2,050	1,130
Adj. R^2	0.134	0.042	0.088	0.013	0.068

Note: p-values are between parenthesis.
* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 2.5: Linear regression of cumulative abnormal spread changes CASC on company characteristics

Finally, moving on to M&A, both the hypothesis 8, the significance of the *magnitude* coefficients emphasizes that the relative impact of the deal, measured by the ratio of the value paid for the target to the acquirer's market value, plays a pivotal role in shaping abnormal spread changes. For buyer companies, the positive coefficient of 0.0173 implies that larger deals, relative to the acquirer's market value, correspond to more pronounced abnormal spread changes. This suggests that market participants perceive greater relative financial commitment and strategic significance in larger M&A transactions, leading to heightened credit risk concerns. This aligns with the notion that substantial deals, even for large firms, can lead to increased complexity, financial strain, or uncertainties, thereby influencing CDS spreads. Contrastingly, for seller companies, the positive and statistically significant coefficient of 0.011 on the magnitude variable indicates that a decrease in the relative value of the target company concerning the buyer's size corresponds to a significant decrease in credit risk, leading to less pronounced abnormal spread changes. In simpler terms, when smaller firms are being acquired by larger firms, the merger and acquisition announcement appears more robust, causing a higher decrease in abnormal spread change for the target company. This dynamic reflects market participants' perceptions that larger acquiring firms, with potentially more extensive financial resources and strategic capabilities, can effectively manage and absorb the challenges associated with substantial M&A transactions. Consequently, this leads to a muted impact on credit risk perceptions for the selling firm, contributing to a decrease in abnormal spread changes. In essence, these mechanisms underscore the importance of considering not only the inherent characteristics of the firms involved but also the relative significance of the M&A transactions in relation to their respective scales. Besides, the lack of significance in the coefficients for firm *size* suggests that the sheer magnitude of a company, in terms of market capitalization, might not be the primary driver of credit

risk reactions to M&A announcements. Large companies may not inherently command a more favorable or unfavorable response from the market in the context of M&A. Furthermore, hypothesis 9, stating that M&A announcements in Continental Europe lead to comparatively smaller changes in CDS spreads than in the Anglo-American context due to stakeholder-oriented governance frameworks, finds empirical support. The negative and statistically significant coefficient for the *region* variable in Table 2.5 indicates that M&A announcements involving Anglo-American companies are associated with higher changes in CDS spreads compared to those in Continental Europe. This outcome aligns with [85], highlighting the impact of governance structures on market reactions to corporate events. The results underscore the influence of stakeholder-oriented governance in Continental Europe, where long-term affiliations and concentrated lenders contribute to a more measured response to M&A announcements.

2.5 Conclusion

In conclusion, this chapter provides valuable insights into the impact and temporal dynamics of corporate action announcement on credit risk, addressing some limitations in the existing literature. Previous studies have often relied on fixed event windows, which can fail to capture the nuanced dynamics of how corporate action announcements influence credit risk. Our study adopted a comprehensive approach, considering a broad range of event windows from 20 days preceding to 20 days following the announcement. This allows us to discern distinct patterns in the response of CDS spreads to various corporate actions, offering a nuanced understanding of the temporal dynamics involved. Moreover, our research stands out by evaluating three specific corporate actions—stock buybacks, cash dividends and M&A—within the same evaluation framework. This holistic approach enables a comparative analysis, shedding light on the relative effects of different corporate actions on credit risk.

For stock buyback announcements, we identify a significant impact on credit risk, particularly in the event window from the announcement day to two days afterward. In the case of dividend raises, our analysis reveals an immediate and substantial response in the CDS market, confined to a narrow window, within two days of the announcement date. On the contrary, dividend cuts exhibit different dynamics with an extended anticipation period and a more prolonged impact, aligning with existing research on the lingering effects of negative news on credit risk perception. For M&A, the temporal influence varies between seller and buyer companies. Seller companies display a notable anticipatory phase, with the impact concentrated around the announcement date. Buyer companies, on the other hand, experienced a less pronounced but more prolonged impact, suggesting a unique nature of M&A events for buyers characterized by both anticipation and lasting consequences.

In addition to our examination of the temporal dynamics surrounding corporate actions, we further distinguish the wealth transfer effect from signaling effects. Specifically, we find that the wealth transfer effect dominates in stock buybacks, cash dividends and M&A announcements. In the context of stock buybacks, our findings suggest that the wealth transfer effect dominates, as indicated by the significant positive mean abnormal spread change within the 0 to 2 days event window. Then, whether through dividend raises or cuts, the observed abnormal spread changes underscore the profound influence of wealth redistribution during these announcements, surpassing the impact of signaling effects. Notably,

the market's heightened sensitivity to negative news, exemplified by the magnitude of the negative abnormal spread change for dividend cuts compared to raises, further accentuates the dominance of the wealth transfer effect. Furthermore, in the realm of M&A, both for buyer and seller companies, the observed mean abnormal spread changes underscore the significance of wealth redistribution during M&A announcements, with the magnitude notably greater for seller firms.

Finally, through an analysis of company and deal characteristics, we uncover determinants of abnormal spread changes repsonses to corporate event announcements. The empirical results validate several hypotheses, such as the positive influence of the magnitude of stock buybacks on abnormal spread changes and the asymmetric market response to cash dividend announcements. However, some unexpected findings, such as the positive correlation between company size and abnormal spread changes following dividend raises, highlight the complexity of the interplay between company characteristics and credit risk dynamics. Additionally, our study supports the notion that the relative magnitude of M&A, rather than just the size of the involved companies, plays a crucial role in shaping abnormal spread changes. The empirical validation of hypotheses related to regional differences in M&A announcements adds further depth to our understanding, emphasizing the impact of governance structures on market reactions.

Chapter 3

Enhancing Long-Short Portfolios: A Refined Approach using Learn-to-Rank Algorithms[1]

Abstract

This chapter delves into the challenges posed by ranking biases inherent to traditional learn-to-rank loss functions, particularly focusing on their impact on the construction of long-short portfolios. Through the analysis of synthetic data, we uncover inherent biases in these methods, particularly detrimental for long-short portfolios where equal importance lies with top and bottom-ranked assets. To address these challenges, we propose enhanced versions of learn-to-rank loss functions—ListNet-Fold, ListMLE-weighted, and ListFold-weighted. These adaptations, tailored for long-short strategies, draw inspiration from pairwise approaches and adjust weighting mechanisms. Empirical results using real-world data sourced from the China A-share market consistently reveal enhancements in ranking metrics, notably improving accuracy in ranking extreme assets, which are more traded in long-short portfolios. Furthermore, financial performance metrics validate the efficacy of these methods, demonstrating enhanced risk-adjusted returns, profitability, and robustness across varying numbers of assets included in the long-short strategy. This research offers valuable insights and practical remedies for mitigating biases in learn-to-rank algorithms, presenting promising tools for constructing long-short portfolios.

3.1 Introduction

In recent years, the integration of machine learning into cross-sectional systematic strategies has led to significant advancements in asset management. Cross-sectional strategies entail simultaneous analysis

[1]This chapter is based on the research article *Enhancing Long-Short Portfolios: A Refined Approach using Learn-to-Rank Algorithms* co-authored with Thibaut Metz, Khalil Sbai and Guillaume Boulanger and published in the *Journal of Financial Data Science*, Winter 2025, 7 (1) 76-97, DOI: 10.3905/jfds.2024.1.173.

of multiple assets at a given point in time, with the goal of identifying pricing inefficiencies or uncovering patterns that guide portfolio construction decisions. This strategy involves analyzing financial instruments to identify relationships, trends, or factors that may impact their relative performance. The introduction of machine learning techniques into cross-sectional strategies introduces a data-driven approach to decision-making, fostering a more nuanced understanding of market dynamics. This ultimately contributes to enhanced risk management and optimized portfolio construction.

Classical methodologies employed for asset ranking encompass cross-sectional momentum, regress-then-rank, and learning-to-rank (LTR). Cross-sectional momentum seeks to compare the relative performance of assets using historical performance indicators, typically represented by returns. The regress-then-rank method employs machine learning models to score assets, addressing tasks such as predicting future volatility-normalized returns. Recent empirical studies ([81]) have substantiated the outperformance of LTR algorithms over traditional regression-based approaches for asset ranking. LTR algorithms aim to directly learn the overall future performance ranking for a list of assets and can be categorized into pointwise, pairwise, or listwise methods.

Besides, from a ranked assets list, a classical approach for constructing portfolios involves implementing long-short strategy with an emphasis on market neutrality. Simultaneously taking long positions in the highest-ranked assets, expected to outperform, and short positions in the lowest-ranked assets, anticipated to underperform, ensures a balanced market exposure, mitigating the impact of broader market movements. Long-short portfolios provide a robust mechanism for capitalizing on relative asset performance while minimizing systematic risk. In contrast to traditional information retrieval tasks that predominantly focus on optimizing accuracy at the top of the ranked list, constructing long-short portfolios introduces distinct requirements for accuracy at both ends of the ranking spectrum. Long-short strategies necessitate not only the identification of assets with the highest expected returns but also the precise pinpointing of those with lower expected returns, reflecting the dual nature of this investment approach. This nuanced objective underscores the need to tailor LTR algorithms to the specific demands of long-short portfolio construction.

To date, a noticeable gap exists in the literature regarding a comprehensive analysis of the behavior of LTR loss functions, specifically tailored for constructing long-short portfolios. While recent studies ([94]) have made progress in adapting LTR algorithms for asset ranking purposes, a precise examination of how different loss functions perform in the context of long-short portfolio construction remains largely unexplored. Recognizing this knowledge gap, this chapter contributes to the existing literature in several ways:

1. **Identification of biases**: We present theoretical and empirical evidences highlighting biases inherent in existing LTR algorithms. Specifically, our focus is on their treatment of lower-ranked elements compared to higher-ranked ones.

2. **Adaptations for long-short strategies**: We introduce variations of ListNet, ListMLE, and ListFold loss functions tailored specifically for long-short portfolio construction. These adaptations address the identified biases and aim to optimize ranking accuracy within the context of long-short strategies.

3. **Empirical validation**: To validate the effectiveness of our proposed adaptations, we conduct an empirical study using weekly China A-share equity data. Our results demonstrate significant improvements in both ranking and portfolio performance metrics, providing robust empirical support for the utility of our enhanced LTR loss functions.

The remainder of the chapter proceeds as follows. In Section 3.2, we delve into the fundamentals of classical LTR algorithms. Section 3.3 examines biases within classical LTR algorithms and offers theoretical rationale. Our adaptations of LTR loss functions for long-short strategies are detailed in Section 3.4. Furthermore, Section 3.5 delineates our empirical methodology and presents our main results. Lastly, Section 3.6 summarizes our key findings and outlines potential avenues for future research.

3.2 Backgrounds: learn-to-rank algorithms

Originating from the information retrieval field, LTR algorithms take as input financial assets, represented by their feature vectors, in order to predict a rank list. This process employs a surrogate loss function to facilitate the learning of a scoring function, assigning scores to each asset based on their feature vectors. This ensures that sorting assets by their scores yields an accurate prediction of their rank. LTR algorithms can be categorized into pointwise, pairwise, and listwise approaches based on their respective loss functions. The pointwise (pairwise) strategy frames the ranking challenge as a classification, regression, or ordinal classification task for individual (pairs of) samples. In contrast, the listwise approach constructs the appropriate ranking model by utilizing ranking lists as inputs. Regarding ranking performance, the pointwise method is usually noted to be less effective compared to the latter two techniques ([63]), aligning with findings in the field of information retrieval.

Let $\mathcal{X} = \{x_1, \cdots, x_n\}$ be the objects to be ranked, in our case financial assets. Assume each object x_i is associated with a true score y_i, and let $\mathcal{Y} = \{y_1, \ldots, y_n\}$ denotes the set of these scores representing the returns of the financial assets. If $y_i > y_j$, it implies that the return of x_i surpasses that of x_j, and consequently, x_i is ranked lower than x_j in the overall ranking. Furthemore, we denote $y^{-1}(i)$ the index of the item that is truly ranked at the i-th position. Let $\mathcal{F}$ be the function class and $f \in \mathcal{F}$ be a ranking function. The optimal ranking function is learned from the training data by minimizing a certain loss function defined on the objects $\mathcal{X}$, their scores $\mathcal{Y}$, and the ranking function f. In the pointwise approach, the loss function is defined on the basis of single objects. For example, the mean squared error may serve as loss function, formulated as follows:

$$L^r(f, \mathcal{X}, \mathcal{Y}) = \sum_{i=1}^{n} \left(f\left(x_i\right) - y_i\right)^2$$

In the pairwise approach, the loss function is formulated by considering pairs of objects with disparate labels. RankNet ([15]) was the first attempt to train a model based on pairs of samples. RankNet involves minimizing the cross-entropy error specifically for pairs of samples, emphasizing the correct classification of relative rankings between two elements. The RankNet loss function has the following form:

$$L^p(f, \mathcal{X}, \mathcal{Y}) = \sum_{i=1}^{n-1} \sum_{\substack{j=1 \\ y_i > y_j}}^{n} \phi\left(f\left(x_i\right) - f\left(x_j\right)\right)$$

In the case of RankNet, ϕ is defined as $\phi(z) = \log\left(1 + e^{-z}\right)$, leading to a logistic distribution of errors, often used for binary classification tasks. Other methods encompass Ranking SVM ([46]), where ϕ is defined as $\phi(z) = (1 - z)_+$, typically associated with margin-based methods, and RankBoost ([33]), where ϕ is exponential $\phi(z) = e^{-z}$, implying an exponential distribution of errors.

In the listwise approach, the loss function is defined on the entire set of n objects, offering a holistic perspective on the ranking task. ListNet ([17]) addresses the challenges of computational costs and misaligned objectives associated with pairwise techniques, utilizing a probabilistic approach on permutations. The method involves calculating top one probability distributions over a list of predicted and ground truth scores, defining the loss as the cross-entropy between these distributions.

$$L^{lNet}(f, \mathcal{X}, \mathcal{Y}) = -\sum_{i=1}^{n} \frac{\phi(y_i)}{\sum_{j=1}^{n} \phi(y_j)} . \log\left(\frac{\phi\left(f\left(x_i\right)\right)}{\sum_{j=1}^{n} \phi\left(f\left(x_j\right)\right)}\right)$$

Moreover, [93] propose ListMLE, an advanced listwise LTR algorithm, establishing the probability distribution through the top-down application of the Plackett-Luce Model. Its objective is to employ a likelihood loss as the surrogate loss, articulated as follows:

$$\begin{aligned} L^{lMLE}(f, \mathcal{X}, \mathcal{Y}) &= -\log \prod_{i=1}^{n} \frac{\phi\left(f\left(x_{y^{-1}(i)}\right)\right)}{\sum_{k=i}^{n} \phi\left(f\left(x_{y^{-1}(k)}\right)\right)} \\ &= -\sum_{i=1}^{n} \left(\log \phi\left(f\left(x_{y^{-1}(i)}\right)\right) - \log \sum_{k=i}^{n} \phi\left(f\left(x_{y^{-1}(k)}\right)\right)\right) \end{aligned}$$

ϕ is taken to be exponential $\phi(z) = e^{-z}$ for both ListMLE and ListNet, leading to an exponential distribution of errors which is beneficial in modeling the decay of rank probabilities and prioritizing top-ranked items.

Lately, a novel model called ListFold has been introduced by [94]. Inspired by the long-short strategy, they propose a distinct surrogate loss function that assigns equal importance to both the top and bottom elements, recognizing their shared contribution to the portfolio's profit and loss. The aim of ListFold is to decompose a permutation into an ordered stepwise pair selection procedure: the first long-short pair, the second long-short pair until the n-th long-short pair. This surrogate loss function is defined as follows:

$$
\begin{aligned}
L^{lFold}(f, \mathcal{X}, \mathcal{Y}) &= -\log \prod_{i=1}^{n} \frac{\phi\left(f\left(x_{y^{-1}(i)}\right) - f\left(x_{y^{-1}(2n+1-i)}\right)\right)}{\sum_{i \leq u \neq v \leq 2n+1-i} \phi\left(f\left(x_{y^{-1}(u)}\right) - f\left(x_{y^{-1}(v)}\right)\right)} \\
&= -\sum_{i=1}^{n} \Bigg[\log \phi\left(f\left(x_{y^{-1}(i)}\right) - f\left(x_{y^{-1}(2n+1-i)}\right)\right) \\
&\quad - \log \sum_{i \leq u \neq v \leq 2n+1-i} \phi\left(f\left(x_{y^{-1}(u)}\right) - f\left(x_{y^{-1}(v)}\right)\right) \Bigg]
\end{aligned}
$$

The authors propose two versions of their ListFold loss function: ListFold-exp, where ϕ is exponential $\phi(z) = e^{-z}$, and ListFold-sgm, where ϕ is defined as the sigmoid function $\phi(z) = \frac{1}{1+e^{-z}}$. [94] show that ListFold-sgm is consistent with the binary classification loss. In all subsequent analyses, we consistently refer to ListFold-exp as ListFold, and only this version is considered. This choice is based on [94] research findings, which highlight that the exponential version of ListFold consistently outperforms its sigmoid counterpart in both ranking and financial performance.

3.3 Ranking bias analysis

In long-short portfolio construction, the simultaneous holding of both long and short positions imposes the need for an algorithm that impartially ranks assets. It is crucial to avoid bias in favoring certain parts of the list over others during the training process. To assess the effectiveness of various LTR loss functions, including RankNet, ListNet, ListMLE and ListFold, we conduct an empirical evaluation using synthetic data. Our objective is to evaluate the training performance of these loss functions in ranking different groups of elements within a list while avoiding the introduction of biases.

3.3.1 Experimental setup

For the experiment, we simulate a training dataset denoted as X, with dimensions $300 \times 74 \times 62$ from a uniform distribution $\mathcal{U}(-1, 1)$. In this configuration, 300 corresponds to the temporal dimension, 74 signifies the number of assets subject to ranking, and 62 denotes the number of features characterizing each asset[2]. Concurrently, we generate a matrix Y with dimensions 300×74, representing the returns for the 74 distinct assets over the 300 time periods. Specifically, the return of asset j at time i is computed using a quadratic form of the feature vector:

$$
Y_{ij} = \sum_{k=1}^{62} (X_{ijk})^2
$$

where X_{ijk} represents the k-th feature of asset j at time i.

[2]The choice of dimensions, namely 300, 74, and 62, is guided by the resemblance to the structure of the empirical dataset outlined in Section 3.5.1.

To learn the scoring function f, we follow the approach of [94] and employ a multilayer perceptron model with a network structure of $[136 \times 272 \times 34 \times 1]$[3]. Each hidden layer employs a Rectified Linear Unit (ReLU) activation function. In the final layer, we opt for a sigmoid activation function to promote symmetry in the final activation, aligning with our objective of avoiding bias in the ranking process. It is important to note that the identical network structure for the scoring function is consistently applied in all subsequent empirical analyses. Our training process extends over 4,000 epochs, with a batch size of 32. To ensure the robustness and consistency of our findings, the experiment is conducted 100 times for each loss function.

3.3.2 Results

Figures 3.1, 3.2 and 3.3 visually depict the evolution of the loss functions RankNet, ListNet, and ListMLE during training using our simulated data. Additionally, it illustrates the changes in accuracy for the top and bottom 10 list elements throughout the training epochs. The accuracy for the top (bottom) k elements, also known as the top-k accuracy metric, is defined as the number of elements at the intersection between the predicted ranks and the true ranks for the true top (bottom) k elements, divided by the total number of elements considered. This indicator offers an assessment of how accurately the predictions align with the actual rankings for the considered list subset. In other words, it measures the proportion of correctly identified top (bottom) k elements out of the k elements that are truly at the top (bottom) of the ranking.

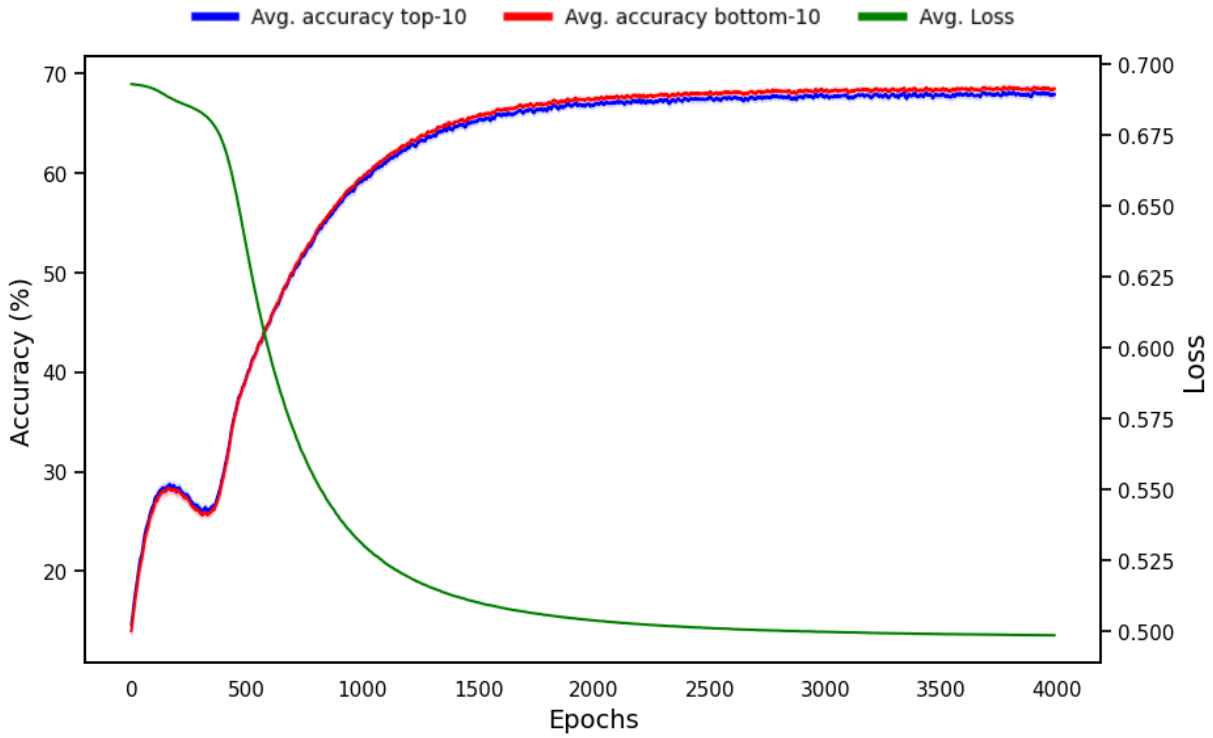

Figure 3.1: Training progress of RankNet: ranking accuracy and loss function evaluation

[3]The rationale for expanding feature dimensions in hidden layers is to promote interaction among factors, resulting in the generation of additional features.

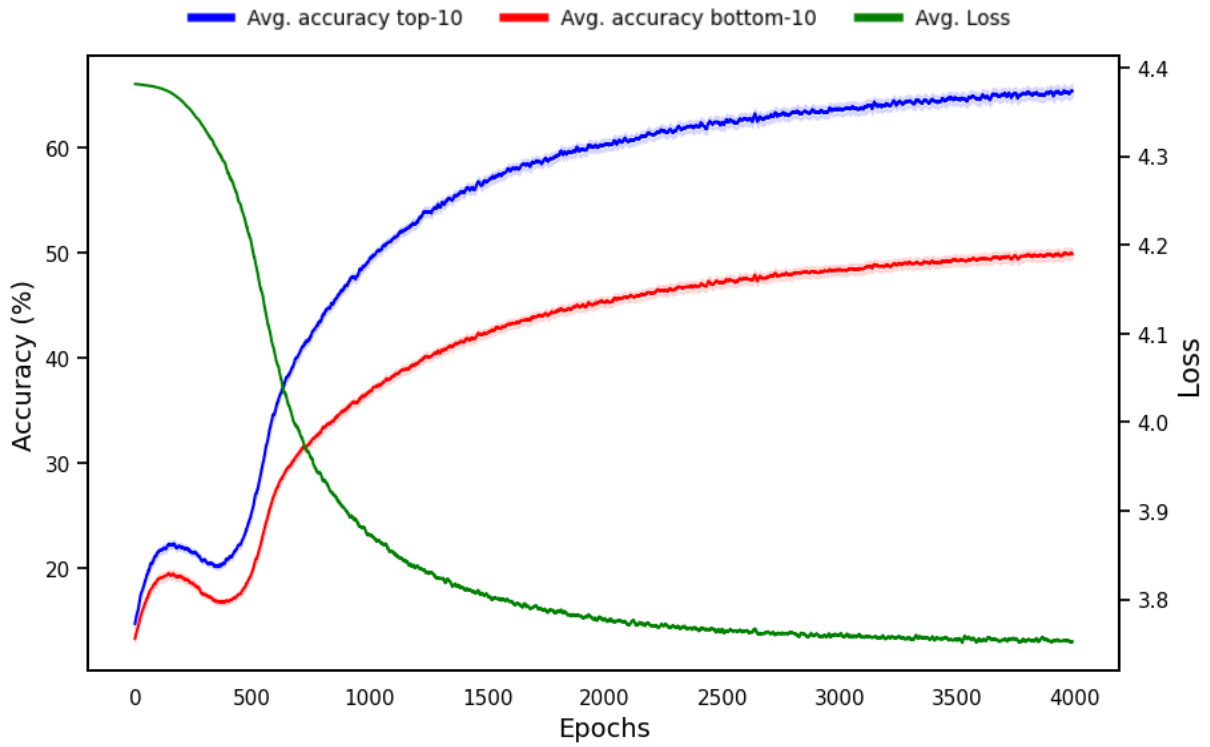

Figure 3.2: Training progress of ListNet: ranking accuracy and loss function evaluation

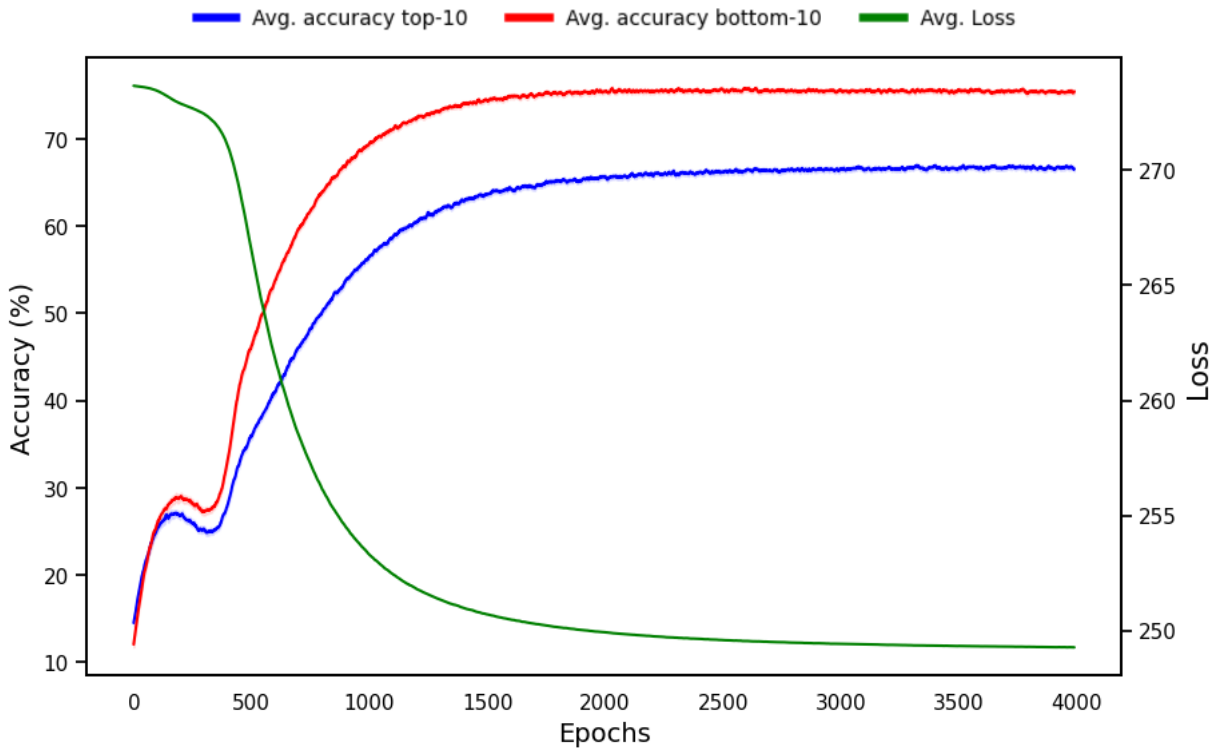

Figure 3.3: Training progress of ListMLE: ranking accuracy and loss function evaluation

Unlike listwise approaches, the pairwise RankNet loss function does not display any bias in assigning

more weight to the same prediction error, regardless of the top or bottom group of the list considered. Indeed, RankNet inherently offers a symmetrical approach to penalization due to its pairwise nature. This symmetry stems from the consideration of pairwise preferences, which ensures that errors in ranking are treated uniformly, regardless of their position within the list. The symmetry associated with the RankNet pairwise approach is discussed in Section 3.7.1.

The ListNet algorithm demonstrates a significant disparity in the ranking accuracy of top versus bottom-ranked elements. Indeed, we observe higher accuracy in ranking the top 10 items compared to the bottom 10. ListNet loss is a cross-entropy measure comparing observed and calculated probabilities, aiming to minimize the gap between predicted and true scores. However, if there is a slight deviation in scores, ListNet demonstrates a tendency to favor predictions that resemble top elements over those resembling bottom elements, thus introducing a bias in its weighting mechanism. This bias primarily stems from the exponential transformation applied in the softmax function, which accentuates differences, particularly in higher scores, leading to a greater impact on the loss when top elements are misranked compared to when bottom elements are misranked. Section 3.7.2 presents mathematical explanations for the ranking bias in ListNet.

In ListMLE, a bias towards being more accurate on ranking bottom-ranked items over top-ranked ones contrasts with ListNet's inclination towards lower-ranked items. This reflects a systematic bias inherent within the algorithm's sequential ranking process. Specifically, ListMLE commences by discerning the top-performing elements from a larger pool and subsequently progresses towards ranking lower-performing ones. This sequential approach inherently imposes greater complexity in distinguishing top-ranked elements due to the necessity of evaluating them against all the other elements of the list. Conversely, as the algorithm traverses towards the lower ranks, the task becomes increasingly simplified owing to the reduced number of elements requiring consideration. As a result, ListMLE tends to better rank lower-ranked elements than their top-ranked counterparts, revealing an asymmetry in how the algorithm handles different parts of the ranking spectrum. Section 3.7.3 offers mathematical explanation for the ranking bias in ListMLE.

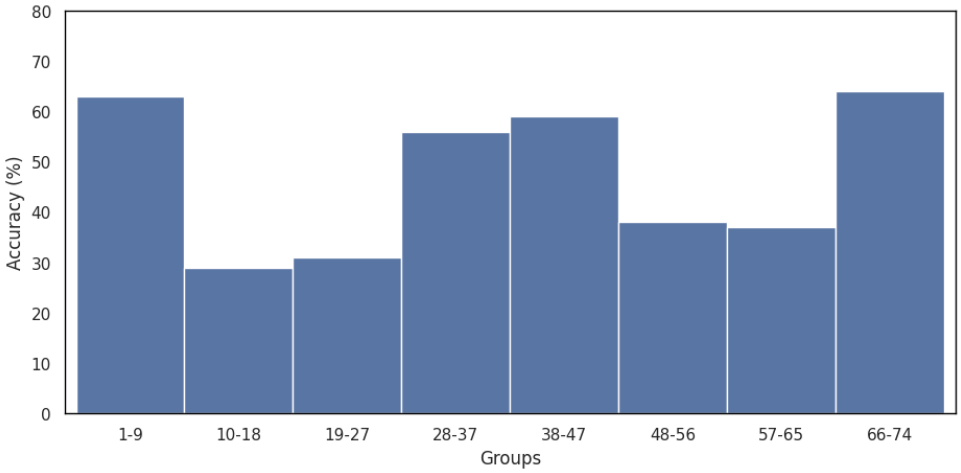

Figure 3.4: ListFold: group ranking accuracy of last training epoch

Finally, Figure 3.4 presents the distribution of ranking accuracy obtained during the final epoch of

ListFold training across various ranking groups. ListFold, inspired by ListMLE but operating with pairs rather than individual items, introduces bias and symmetry in the ranking. Upon closer examination of accuracy for different group items, a clear symmetrical pattern emerges between the two halves of the list (1-37 and 38-74). This symmetry arises from ListFold's pairwise nature, wherein it processes pairs starting from the extremes. Consequently, there exists a symmetrical accuracy pattern for list elements paired during the loss computation. Notably, the extreme elements of both halves (1-9 and 66-74) display the highest accuracy, reinforcing the impact of the pairwise approach. It is important to note that the middle groups (28-47), while exhibiting higher accuracy, are less relevant for long-short portfolio construction, as they are not actively traded.

3.4 Enhanced learn-to-rank loss functions for long-short ranking

In this section, we introduce various enhancements to classical LTR loss functions. Firstly, as discussed in Section 3.3, we identify systematic biases inherent in traditional listwise loss functions, especially impacting the treatment of top and bottom rankings. To address this issue, we propose corrections for the ListNet and ListMLE loss functions, aiming to rectify this asymmetry and optimize the algorithms for long-short strategies. Additionally, acknowledging the heightened significance of extreme elements compared to the middle section in a list, we present a weighting scheme for ListFold. This scheme amplifies attention towards the extremes while diminishing emphasis on the less critical middle section. Through these adjustments, our goal is to improve the adaptability of LTR algorithms, effectively mitigating biases and facilitating rankings tailored to the demands of long-short strategies.

3.4.1 ListNet-Fold

As outlined in Section 3.7.2, the softmax function used in the cross-entropy calculation of ListNet tends to accentuate errors occurring at the top positions. In order to adapt ListNet for long-short strategy, we leverage this inherent imbalance between top and bottom rankings. Taking inspiration from the ListFold framework, which integrates pairwise and listwise surrogate loss functions, we introduce ListNet-Fold. This variant is constructed as a cross-entropy for the differences in scores between every two pairs. By adopting this approach, we utilize the ListNet imbalance, focusing on enhancing rankings at the extremes while de-prioritizing pairs in the middle. This loss function enables us to strategically direct optimization efforts towards areas that have the most significant impact on the construction of long-short portfolios. The formulated loss function is expressed as:

$$L^{lNet-Fold}(f, \mathcal{X}, \mathcal{Y}) = -\sum_{i=1}^{n}\sum_{j=1}^{n} \frac{\phi(y_i - y_j)}{\sum_{k=1}^{n}\sum_{l=1}^{n}\phi(y_k - y_l)} \log\left(\frac{\phi(f(x_i) - f(x_j))}{\sum_{k=1}^{n}\sum_{l=1}^{n}\phi(f(x_k) - f(x_l))}\right)$$

where ϕ is taken to be exponential.

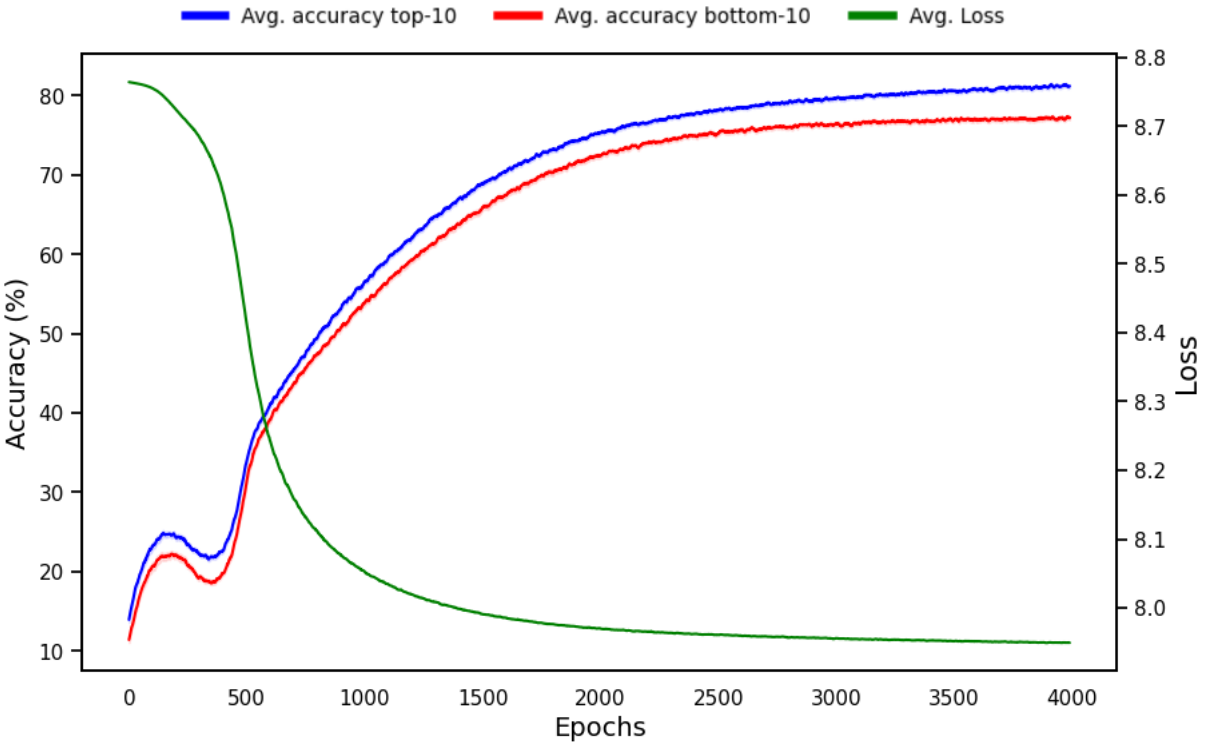

Figure 3.5: ListNet-Fold: Accuracy rankings and loss function evaluation

To demonstrate the effectiveness of our corrected version of ListNet, we examine the evolution of accuracies for the top and bottom 10 elements during training, as discussed in Section 3.3. From Figure 3.5, we can observe that our modified version of ListNet effectively addresses the asymmetry concerns.

3.4.2 ListMLE-weighted

For ListMLE, we suggest incorporating weights to emphasize the importance of higher-ranked items, aiming to counteract the bias introduced during sequential learning. Since sequential learning may diminish the significance of top-ranked items, the introduction of weighting aims to enhance their impact on the overall likelihood, thereby offsetting the inherent bias. The ListMLE-weighted loss function is expressed as follows:

$$L^{lMLE-w}(f, \mathcal{X}, \mathcal{Y}) = -\sum_{i=1}^{n} \lambda_i \left(\log \phi \left(f \left(x_{y^{-1}(i)} \right) \right) - \log \sum_{k=i}^{n} \phi \left(f \left(x_{y^{-1}(k)} \right) \right) \right)$$

where the weights $\lambda_i = n + 1 - i$ signify the importance of accurately predicting the rank of element i in the ranking process[4].

This new formula can also be interpreted probabilistically. Specifically, one may view the likelihood of selecting the element ranked i as repeated λ_i times. This is represented by the following probability

[4]It is important to note that this new formula maintains the consistency of the loss function. A similar modification is implemented by [61] to enhance *position awareness* within ListMLE.

distribution:

$$P(y|x;f) = P(y^{-1}(1)|x;f)^{\lambda_1} \prod_{i=2}^{n} P(y^{-1}(i)|x, y^{-1}(1), \ldots, y^{-1}(i-1);f)^{\lambda_i}$$

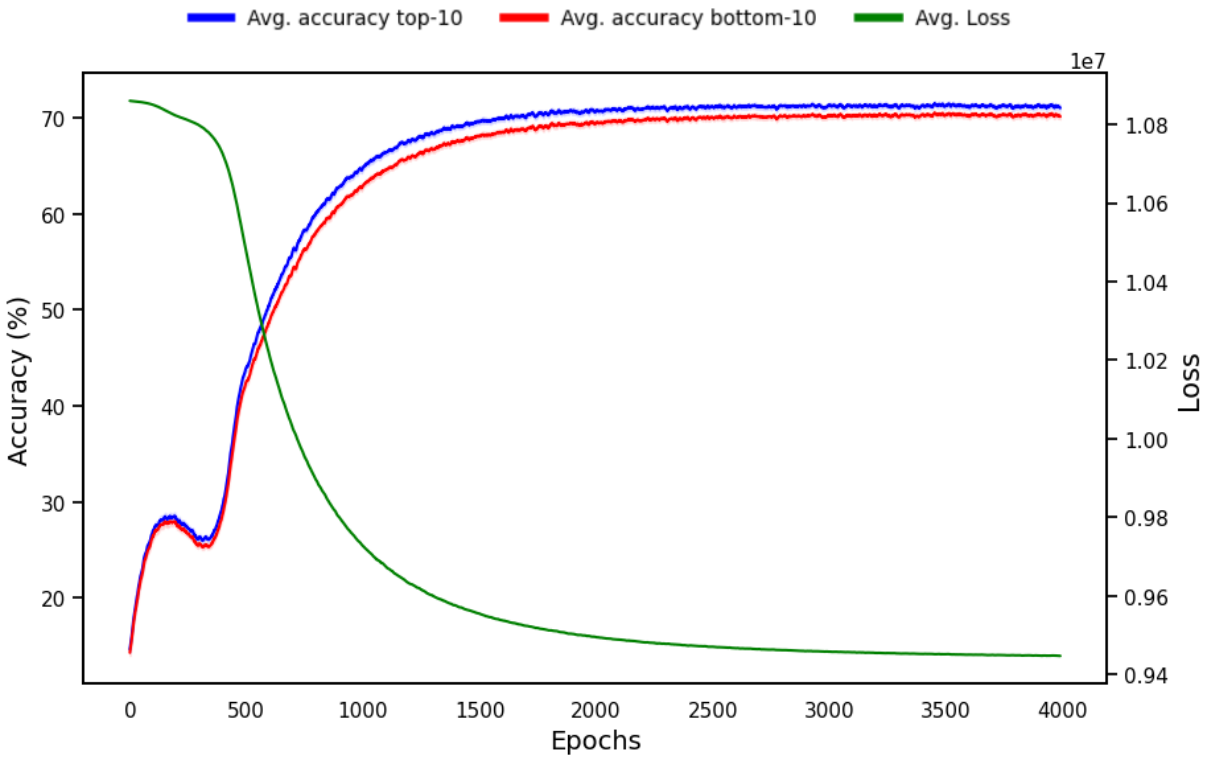

Figure 3.6: ListMLE-weighted: Accuracy rankings and loss function evaluation

Upon examination of Figure 3.6, we observe a significant reduction in the discrepancy between top and bottom accuracies. This indicates that our proposed method can effectively mitigates the bias introduced by the sequential learning of ListMLE.

3.4.3 ListFold-weighted

In ListFold, the sequential selection process initially tends to sort out pairs with the largest differences in scores, and then progresses to pairs with smaller differences. We propose to augment the ListFold loss function by incorporating weights, which involves repeating the i^{th} step in the sequential pairs selection process a number of times proportional to the differences in scores. This amplifies the significance of items crucial to the long-short objectives, countering their dilution in the likelihood calculation process.

We propose the following formulation for ListFold-weighted:

$$
\begin{aligned}
L^{lFold-w}(f, \mathcal{X}, \mathcal{Y}) &= -\log \prod_{i=1}^{n} \left(\frac{\phi\left(f\left(x_{y^{-1}(i)}\right) - f\left(x_{y^{-1}(2n+1-i)}\right)\right)}{\sum_{i \le u \ne v \le 2n+1-i} \phi\left(f\left(x_{y^{-1}(u)}\right) - f\left(x_{y^{-1}(v)}\right)\right)} \right)^{\lambda_i} \\
&= -\sum_{i=1}^{n} \lambda_i \left[\log \phi\Big(f\big(x_{y^{-1}(i)}\big) - f\big(x_{y^{-1}(2n+1-i)}\big) \Big) \right. \\
&\quad \left. - \log \sum_{i \le u \ne v \le 2n+1-i} \phi\Big(f\big(x_{y^{-1}(u)}\big) - f\big(x_{y^{-1}(v)}\big) \Big) \right]
\end{aligned}
$$

where weights $\lambda_i = 1 - \frac{i}{n}$ represent the significance of accurately predicting the rank of pairs of elements $(i, 2n - i)$ during the ranking process. The weight λ_i decreases as i increases, indicating diminishing importance as we progress through pairs of middle-list elements.

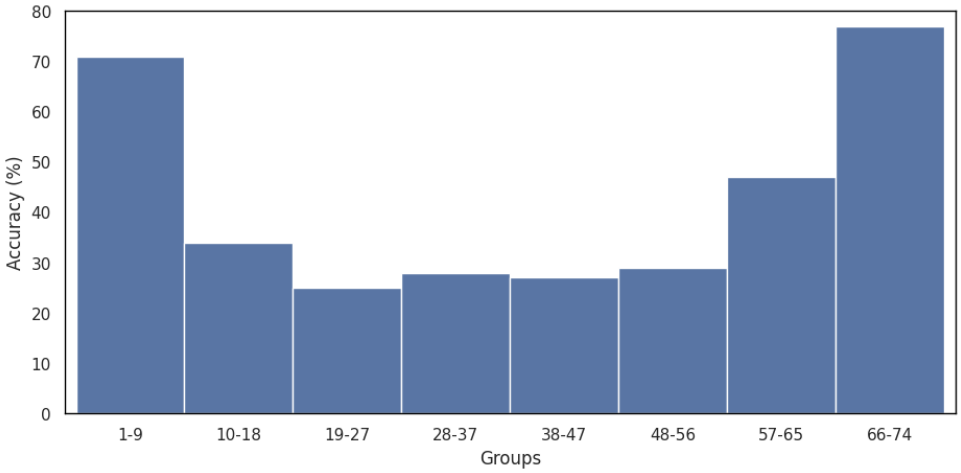

Figure 3.7: ListFold-weighted: group ranking accuracy of last training epoch

Figure 3.7 depicts the distribution of ranking accuracy during the final epoch of ListFold-weighted training. Notably, there is a decline in accuracy for items within the 28-47 groups, while accuracy increases for the top (1-9) and bottom (66-74) extreme groups. In the context of long-short portfolio construction, where the focus is primarily on the top and bottom extremes, this observed trade-off is deemed acceptable. The slight reduction in accuracy for middle groups is offset by the enhanced accuracy in extreme groups, containing elements of greater strategic importance for long-short portfolio construction.

3.5 Empirical studies

3.5.1 Dataset overview

To empirically assess the effectiveness of our proposed approaches, namely ListNet-Fold, ListMLE-weighted, and ListFold-weighted, we employ a dataset containing information about equities characterized

by various factors. This dataset is derived from the study by [94] and is available on GitHub[5]. Specifically, the dataset comprises 620 weekly observations related to 74 stocks within the China A-share market, each described by 62 factors. The data spans from December 29, 2006, to February 1, 2019. Notably, these 74 stocks are predominantly listed in the HS300 index and are characterized by high liquidity.

3.5.2 Training and backtesting procedure

We implement a rolling training strategy for the model introduced in Section 3.3.1. Specifically, we use a training set of 300 weeks, excluding the subsequent week to prevent potential correlations between the training and testing data. The following 16 weeks are designated as the test set. To ensure comparability among features with varying orders of magnitude, we apply standard scaler normalization. This normalization involves computing the parameters solely on the training set and then utilizing these parameters to normalize the test sets.

The models are trained using a mini-batch size of 32, spanning 2,000 epochs, with the Adam optimizer and a learning rate of 10^{-3}. The test set comprises 320 weekly data points, spanning from November 1, 2012, to January 31, 2019. For fair comparisons across different loss functions, we initialize the multi-layer perceptron model with identical weights. Additionally, each mini-batch presented to the model remains strictly identical across trainings.

Based on the obtained out-of-sample rank predictions, we construct our long-short portfolio, involving taking long positions on the top k stocks and short positions on the bottom k. Our allocation strategy places greater emphasis on stocks at the top and bottom of the predicted ranking. This is achieved through the following weighting formulas:

$$w_{j+} = \frac{k-j}{k(k+1)} \quad \text{and} \quad w_{j-} = -\frac{k-j}{k(k+1)}$$

Here, w_{j+} and w_{j-} represent the weights assigned to the j-th stock for a long and short position, respectively. We consistently allocate a nominal capital of \$1 to our investment activities on a weekly basis.

3.5.3 Performances

The evaluation of the various algorithms involves the use of two sets of metrics, both calculated during out-of-sample test periods. The first set is derived from literature on information retrieval and ranking. The second set encompasses standard financial metrics, as suggested by [81], evaluating aspects such as profitability, risk, and financial performance.

From a ranking perspective, we assess metrics like Information Coefficient (IC), Kendall Tau ($K\tau$), and weighted Kendall Tau ($K\tau_w$). In the case of weighted Kendall Tau, we aim for a ranking metric

[5]https://github.com/TCtobychen/ListFold.

that gives more weight to stocks at the extremes aligning with the objective of constructing long-short portfolio. We achieve this by using a quadratic[6] function of the true rank vector r as weights, represented by the formula $w_i = (r_i - \bar{r})^2$, r_i being the true rank of element list i.

In terms of profitability, we examine expected yearly returns adjusted to risk free ($\mathbb{E}[\text{Returns}] - r_f$) and the percentage of positive returns at the portfolio level. For risk metrics, we compute yearly volatility (V), maximum drawdown (MDD), and downside deviation (DD). Financial performance metrics for risk-adjusted evaluation include Sharpe ($\frac{\mathbb{E}[\text{Returns}]-r_f}{V}$), Sortino ($\frac{\mathbb{E}[\text{Returns}]-r_f}{DD}$), and Calmar ($\frac{\mathbb{E}[\text{Returns}]-r_f}{MDD}$) ratios. Additionally, we include the average profit divided by the average loss ($\frac{Avg.P}{Avg.L}$). The risk-free rate r_f that we use is equal to 3%.

Ranking metrics

	ListNet		ListMLE		ListFold	
	Classic	Fold	Classic	Weighted	Classic	Weighted
IC	0.058	**0.064**	**0.073**	0.068	0.057	**0.062**
$K\tau$	0.040	**0.044**	**0.051**	0.047	0.039	**0.047**
$K\tau_w$	0.053	**0.058**	0.062	**0.064**	0.052	**0.056**

Note: Bold font indicates the best metric.

Table 3.1: Ranking metrics

The results in Figure 3.1 provide a comprehensive comparison of classic ListNet, ListMLE and ListFold loss functions alongside their enhanced versions. The evaluation is based on three critical ranking metrics — Information Coefficient, Kendall Tau, and weighted Kendall Tau — for which higher values are desirable.

Starting with the Information Coefficient, ListMLE outperforms both ListNet and ListFold in their classic and enhanced forms. However, our proposed improvement, ListNet-Fold, exhibits a notable increase in Information Coefficient (+0.006) compared to its classic counterpart. Similarly, ListFold-weighted demonstrates improvement over the classic ListFold, emphasizing the efficacy of our modifications in enhancing the information coefficient. Turning to Kendall Tau, the same pattern emerges, with ListMLE outperforming ListNet and ListFold in their classic versions. Additionally, ListNet-Fold and ListFold-weighted demonstrate improvements over their classic counterparts, with ListFold-weighted showing a particularly noteworthy enhancement (+0.008). Finally, the analysis of weighted Kendall Tau reveals the outperformance of our enhanced LTR loss functions over the classical ones. Specifically, all our proposed improvements — ListNet-Fold, ListMLE-weighted, and ListFold-weighted — demonstrate advancements over their classical versions, underscoring the effectiveness of our modifications. It is important to highlight a nuanced observation regarding ListMLE. While its classical form surpasses ListMLE-weighted in both Information Coefficient and Kendall Tau, our modification of ListMLE shows

[6]We selected the quadratic form to emphasize the ranking accuracy of extreme assets due to its natural fit with our objectives, although other weighting functions, such as linear weights, were not tested.

improvement only in weighted Kendall Tau. The enhanced performance in weighted Kendall Tau suggests the efficacy of our weighting scheme in emphasizing the importance of extreme elements within the ranking process.

In summary, Figure 3.1 illustrates that our proposed modifications, namely ListNet-Fold, ListMLE-weighted, and ListFold-weighted, enhance the ranking performance of classic LTR algorithms. However, it is important to note that our modifications do not uniformly improve all three ranking metrics for ListMLE-weighted. Nevertheless, all three enhanced approaches demonstrate superiority, particularly in the case of the weighted Kendall Tau. This metric holds particular significance as it better represents the desired target ranking for long-short portfolio construction.

Portfolio metrics

	ListNet		ListMLE		ListFold	
	Classic	Fold	Classic	Weighted	Classic	Weighted
$\mathbb{E}[\text{Returns}] - r_f$	0.151	**0.165**	0.197	**0.223**	0.119	**0.175**
V	0.174	**0.171**	0.160	**0.157**	0.163	**0.157**
Sharpe	0.868	**0.965**	1.231	**1.42**	0.73	**1.115**
DD	**0.127**	0.130	0.110	**0.102**	0.116	**0.106**
MDD	0.352	**0.343**	**0.178**	0.183	0.210	**0.185**
Sortino	1.189	**1.269**	1.79	**2.187**	1.026	**1.651**
Calmar	0.429	**0.481**	1.107	**1.219**	0.567	**0.946**
% +ve Returns	0.581	**0.589**	**0.611**	0.610	0.579	**0.597**
Avg. P / Avg. L	1.098	**1.099**	1.132	**1.180**	1.045	**1.119**

Note: Bold font indicates the best metric.

Table 3.2: Average financial metrics for different long-short k ($1 \leq k \leq 10$)

Figure 3.2 provides a comprehensive view of the average financial performance metrics across various long-short portfolio constructions ($1 \leq k \leq 10$) using ListNet, ListMLE, and ListFold in both classic and enhanced forms.

Considering expected returns, ListMLE-weighted consistently outperforms its classic counterpart, showcasing the highest returns among all algorithms (0.223). ListNet-Fold and ListFold-weighted also demonstrate improvements over their classic versions. In terms of volatility, all enhanced versions of the loss function exhibit a reduction compared to their classic counterparts, signaling decreased risk exposure. ListMLE-weighted and ListFold-weighted display the lowest volatility at 0.157, leading to higher Sharpe ratios. The ListMLE-weighted model achieves the best Sharpe value (1.42), highlighting its effectiveness in enhancing risk-adjusted profitability. Besides, ListFold-weighted shows improvement over both downside deviation and maximum drawdown. However, ListNet-Fold and ListMLE-weighted do not improve these metrics uniformly. Specifically, ListNet-Fold does not improve downside deviation, while ListMLE-weighted does not enhance maximum drawdown. Nevertheless, the improvement in expected

returns compensates, resulting in enhanced Sortino and Calmar indices across all three improved versions of loss functions. Analyzing the percentage of positive returns (% +ve Returns), ListNet-Fold and ListFold-weighted outperform their classic counterparts, indicating a higher proportion of positive returns in the long-short portfolios. ListMLE-weighted, however, experiences a slight decrease in this metric compared to classic ListMLE.

In summary, the financial performance metrics suggest that ListNet-Fold, ListMLE-weighted, and ListFold-weighted consistently enhance the risk-adjusted returns and profitability of long-short portfolios. It is worth noting that, even if all three ranking performances of ListMLE-weighted are not uniformly improved, the resulting long-short portfolio demonstrates improved financial performance.

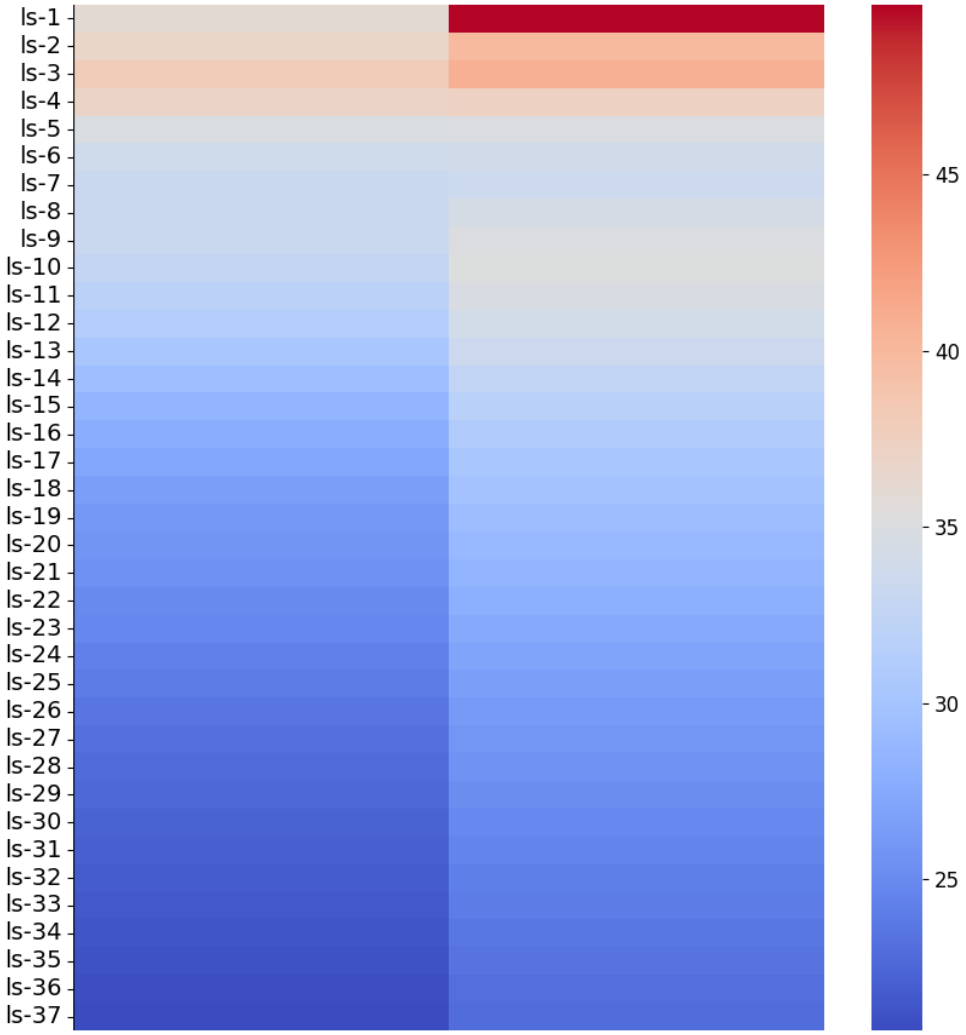

Figure 3.8: Heatmap of average weekly returns for long-short k pairs: ListNet (left) and ListNet-Fold (right)

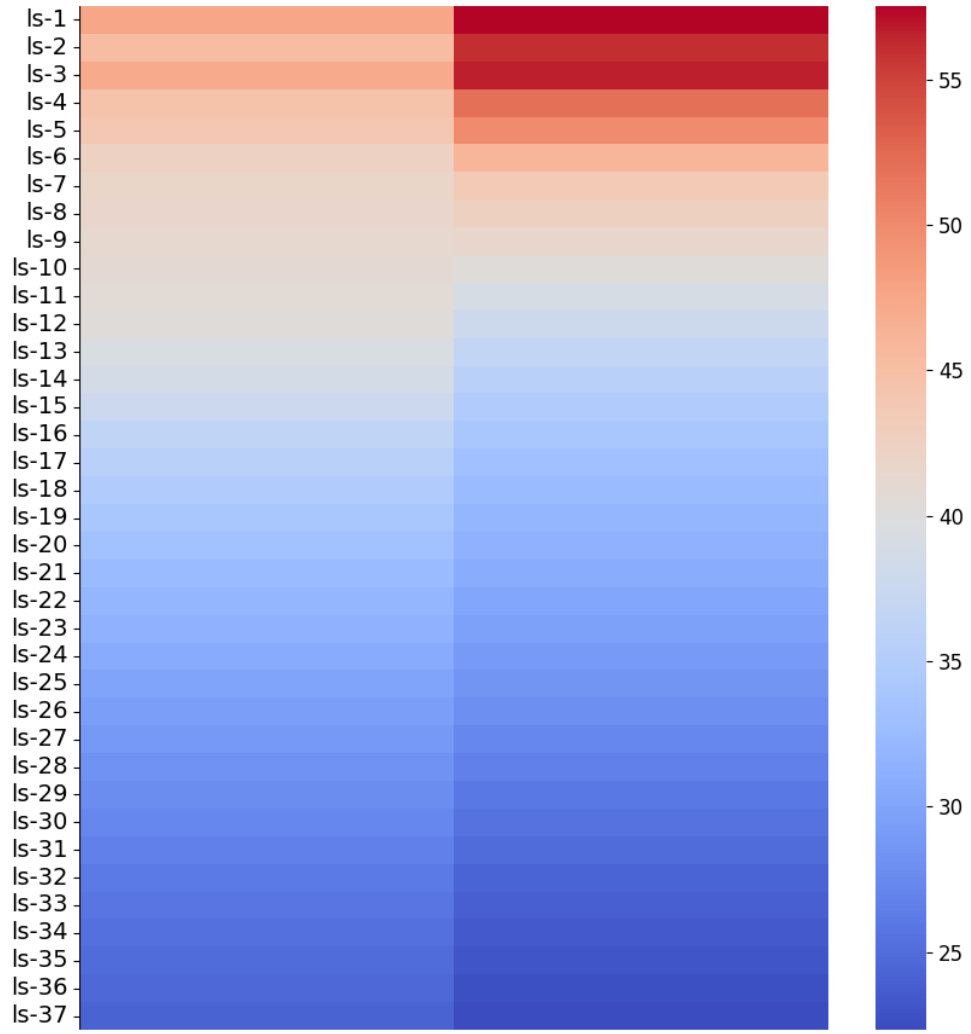

Figure 3.9: Heatmap of average weekly returns for long-short k pairs: ListMLE (left) and ListMLE-weighted (right)

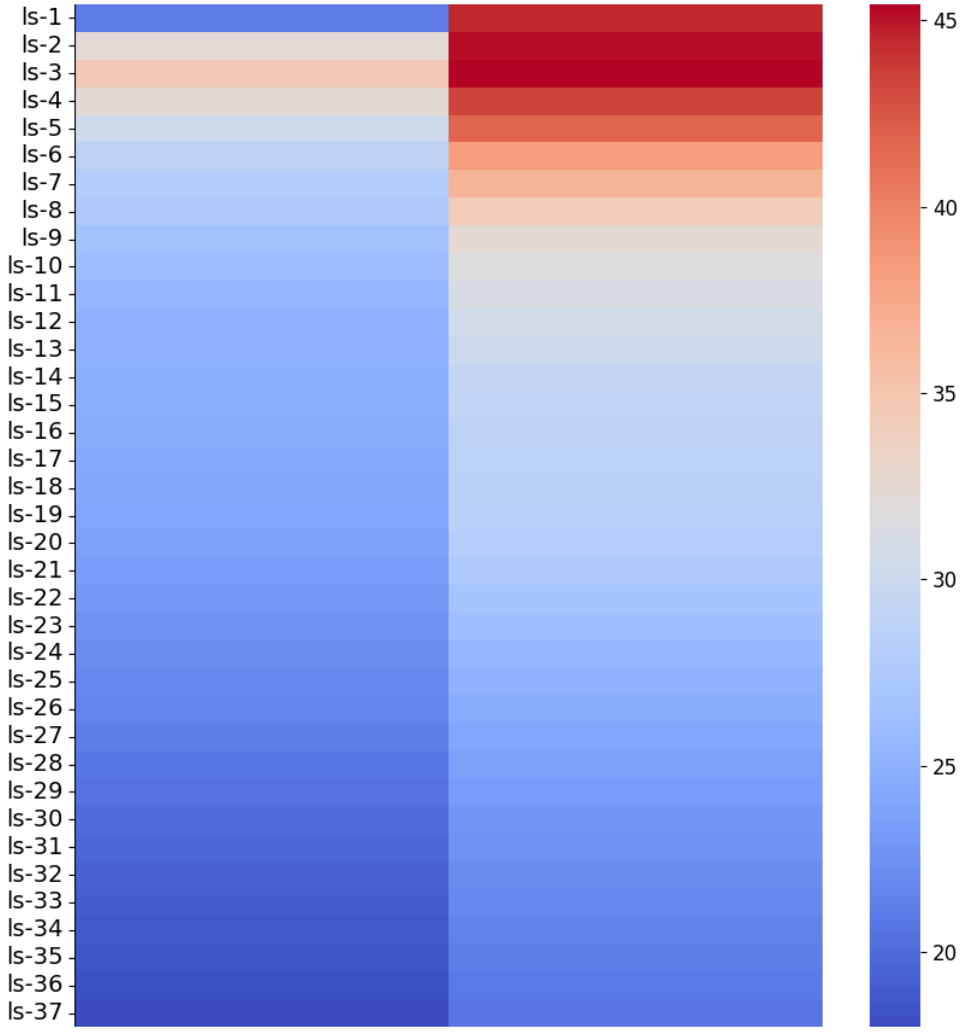

Figure 3.10: Heatmap of average weekly returns for long-short k pairs: ListFold (left) and ListFold-weighted (right)

Figures 3.8, 3.9 and 3.10 delves into the robustness analysis of the position cutoff parameter k, representing the number of stocks considered for both long and short positions in our portfolio. These figures showcase heatmaps generated for various models, each associated with different values of k. Columns represent different models, while each row corresponds to a specific position cutoff parameter k. For instance, in the 8^{th} row, we consider the top 8 stocks for long positions and the bottom 8 stocks for short positions. The color in each cell indicates the average weekly return (in basis points) achieved out of sample without transaction fees, with warmer colors denoting larger average returns.

In an ideal scenario where our predictions perfectly align with the ground truth, we would anticipate the portfolio's profit and loss (P&L) to gradually decrease as k increases from 1 to 37. Figures 3.8, 3.9 and 3.10 illustrate that as k increases, the average return decreases, indicating the robustness of all six rank methods with respect to k. Additionally, there is a discernible advantage for ListNet-Fold, ListMLE-weighted, and ListFold-weighted over their classic counterparts. Specifically, these methods exhibit greater average returns, especially when the long-short portfolio includes extreme top and bottom pairs. Notably, while ListMLE-weighted shows a higher average return at the top of the heatmap, classic

ListMLE yields higher average returns for a high k. This suggests that ListMLE-weighted, despite emphasizing extreme pairs, may not rank the overall list better than ListMLE, as detailed in Section 3.5.3.

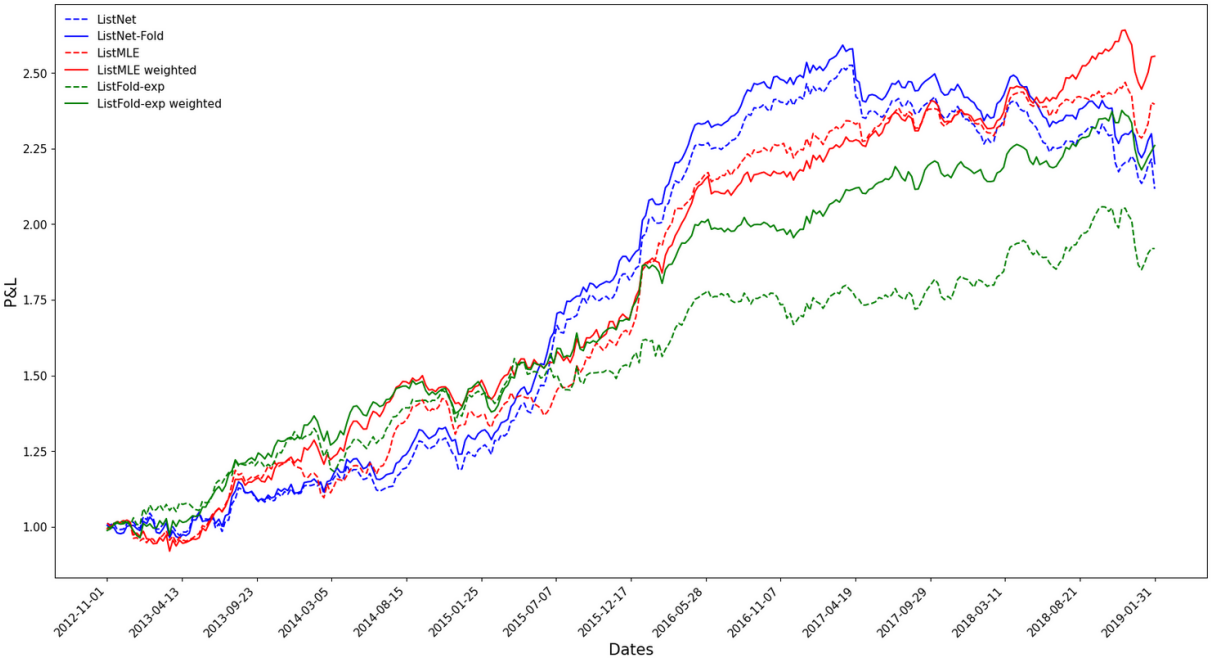

Figure 3.11: Average P&L for different long-short k ($1 \leq k \leq 10$)

Figure 3.11, depicting the average P&L for a long-short strategy across various values of k (from 1 to 10) provides valuable insights into the performance of our proposed enhancements to classic LTR loss functions. Notably, our modifications yield a discernible improvement in P&L during the testing period from November 1, 2012, to January 31, 2019, using China A-share market equity data. The enhanced versions of ListNet, ListMLE, and ListFold consistently demonstrate higher P&L values compared to their classic counterparts, emphasizing the efficacy of our adaptations in enhancing the profitability of the long-short portfolio strategy.

3.6 Conclusion

In this chapter, we conduct a comprehensive analysis of ranking bias in the context of long-short portfolio construction, focusing on classic LTR loss functions, including RankNet, ListNet, ListMLE, and ListFold. Our empirical evaluation, performed using synthetic data, reveals inherent biases in these methods, particularly in the treatment of top and bottom-ranked elements. To address these biases, we introduce enhanced versions of the classic LTR loss functions, namely ListNet-Fold, ListMLE-weighted, and ListFold-weighted. These enhancements aim to rectify asymmetries and optimize the algorithms for long-short strategies by adjusting the weighting mechanisms and introducing tailored modifications.

Our empirical experiments on a real-world dataset from the China A-share market demonstrate the effectiveness of our proposed enhancements. The ranking metrics, including Information Coefficient,

Kendall Tau, and weighted Kendall Tau, consistently show improvements in the enhanced versions, emphasizing their ability to better capture the desired ranking for long-short portfolio construction. Moreover, the financial performance metrics highlight the enhanced risk-adjusted returns and profitability of long-short portfolios constructed using the proposed modifications. ListNet-Fold, ListMLE-weighted, and ListFold-weighted consistently outperform their classic counterparts in terms of expected returns, volatility, Sharpe ratio, Sortino ratio and Calmar ratio.

In conclusion, our research provides valuable insights into the biases present in classic LTR loss functions in the context of long-short portfolio construction. The proposed enhancements successfully mitigate these biases, resulting in improved ranking accuracy and financial performance. The adaptability of ListNet-Fold, ListMLE-weighted, and ListFold-weighted makes them promising candidates for practitioners and researchers seeking robust solutions for long-short strategies in financial markets. Moving forward, the potential for further advancements in this field remains significant. In future investigations, researchers could delve into employing supplementary machine learning techniques, including multitasking, to enhance the optimization of pair selection for long-short positions. Additionally, extending our findings to analyze diverse financial data would provide valuable insights in this field.

3.7 Appendix

3.7.1 Unbiased ranking penalization in RankNet

To analyze RankNet's symmetric penalization, we start from a correctly ordered ranking represented by scores a where $a_1 > a_2 > \ldots > a_n$. We aim to compare the penalization for misranking two consecutive elements in the upper half of the ranking to that in the lower half [7]. For this purpose, we introduce scores h and g, representing misrankings in the upper and lower halves respectively. Specifically, h represents a misranking in the upper half where two consecutive elements are incorrectly ordered, while g represents a similar misranking in the lower half.

For the top half of the ranking, h is defined such that:

$$\begin{cases} h_j = a_{j+1} \ , \ h_{j+1} = a_j & \text{where } j < n/2 \\ h_k = a_k \text{ for } 1 \leq k \leq n & \text{where } k \neq j, j+1 \end{cases}$$

This configuration ensures a swap of consecutive items within the top half of the ranking.

For the bottom half of the ranking, g is defined as follows:

$$\begin{cases} g_i = a_{i+1} \ , \ g_{i+1} = a_i & \text{where } i \geq n/2 \\ g_k = a_k \text{ for } 1 \leq k \leq n & \text{where } k \neq i, i+1 \end{cases}$$

[7] Our analysis is predicated on the notion that any permutation can be broken down into adjacent transpositions. This rationale supports our emphasis on misrankings of consecutive indices, as adjacent transpositions are the fundamental operations needed to revert a permutation to its initial state.

Similar to h, this arrangement results in a swap of consecutive items within the bottom half of the ranking. The penalization for misranking two consecutive elements in the top half given a set of scores $L^{rNet}(h) - L^{rNet}(a)$ evaluates to:

$$\begin{aligned} L^{rNet}(h) - L^{rNet}(a) &= \sum_{p=1}^{n-1} \sum_{\substack{q=1 \\ y_p > y_q}}^{n} [\log(1 + \exp(h_q - h_p)) - \log(1 + \exp(a_q - a_p))] \\ &= \log(\frac{1 + \exp(a_j - a_{j+1})}{1 + \exp(a_{j+1} - a_j)}) \end{aligned}$$

As for a ranking error in the lower half, the penalization $L^{rNet}(g) - L^{rNet}(a)$ evaluates to:

$$\begin{aligned} L^{rNet}(g) - L^{rNet}(a) &= \sum_{p=1}^{n-1} \sum_{\substack{q=1 \\ y_p > y_q}}^{n} [\log(1 + \exp(g_q - g_p)) - \log(1 + \exp(a_q - a_p))] \\ &= \log(\frac{1 + \exp(a_i - a_{i+1})}{1 + \exp(a_{i+1} - a_i)}) \end{aligned}$$

When errors exhibit the same magnitude, denoted as $a_j - a_{j+1} = a_i - a_{i+1}$, our analysis reveals that RankNet's pairwise approach establishes a balanced penalization scheme. This scheme equally addresses errors in both the upper and lower halves of the ranking, highlighting the fairness of RankNet in assessing ranking errors across the entire ranking spectrum.

3.7.2 Ranking penalization bias in ListNet

Suppose $a = a_1, a_2, \ldots, a_n$ represents scores that provide a correct ranking, with $a_1 > a_2 > \ldots > a_n$. Additionally, let h denotes scores that misrank the top two elements, and g denotes scores that misrank the bottom two elements. For h, which misranks the top two elements:

$$\begin{cases} h_1 = a_2 & \text{indicating that the top-ranked score is now } a_2 \\ h_2 = a_1 & \text{indicating that the second-ranked score is now } a_1 \\ h_k = a_k & \text{for } 3 \leq k \leq n \text{ meaning that the remaining scores are unchanged} \end{cases}$$

For g, which misranks the bottom two elements:

$$\begin{cases} g_n = a_{n-1} & \text{indicating that the last-ranked score is now } a_{n-1} \\ g_{n-1} = a_n & \text{indicating that the second-to-last-ranked score is now } a_n \\ g_k = a_k & \text{for } 1 \leq k \leq n-2 \text{ meaning that the remaining scores are unchanged} \end{cases}$$

Besides, let $y = y_1, y_2, \ldots, y_n$ be the vector of true scores where $y_1 > y_2 > \ldots > y_n$. These returns serve as the basis for calculating the true probabilities in cross-entropy.

The penalization for misranking the top two elements given a set of scores $L^{lNet}(h) - L^{lNet}(a)$ evaluates

to:

$$L^{lNet}(h) - L^{lNet}(a) = \sum_{i=1}^{n} \frac{\exp(y_i)}{\sum_{k=1}^{n} \exp(y_k)} \left[\log\left(\frac{\exp(a_i)}{\sum_{k=1}^{n} \exp(a_k)} \right) - \log\left(\frac{\exp(h_i)}{\sum_{k=1}^{n} \exp(h_k)} \right) \right]$$
$$= \frac{1}{\sum_{k=1}^{n} \exp(y_k)} \left[\sum_{i=1}^{n} \exp(y_i) \log\left(\frac{\exp(a_i)}{\exp(h_i)} \right) \right]$$

since $h_1 = a_2$, $h_2 = a_1$, and $h_k = a_k$ for $3 \leq k \leq n$, then:

$$L^{lNet}(h) - L^{lNet}(a) = \frac{1}{\sum_{k=1}^{n} \exp(y_k)} \left[\exp(y_1) \log\left(\frac{\exp(a_1)}{\exp(a_2)} \right) + \exp(y_2) \log\left(\frac{\exp(a_2)}{\exp(a_1)} \right) \right]$$
$$= (a_1 - a_2) \left(\frac{\exp(y_1)}{\sum_{k=1}^{n} \exp(y_k)} - \frac{\exp(y_2)}{\sum_{k=1}^{n} \exp(y_k)} \right)$$

The penalization for misranking the bottom two elements given a set of scores $L^{lNet}(g) - L^{lNet}(a)$ evaluates to:

$$L^{lNet}(g) - L^{lNet}(a) = \sum_{i=1}^{n} \frac{\exp(y_i)}{\sum_{k=1}^{n} \exp(y_k)} \left[\log\left(\frac{\exp(a_i)}{\sum_{k=1}^{n} \exp(a_k)} \right) - \log\left(\frac{\exp(g_i)}{\sum_{k=1}^{n} \exp(g_k)} \right) \right]$$
$$= \frac{1}{\sum_{k=1}^{n} \exp(y_k)} \left[\sum_{i=1}^{n} \exp(y_i) \log\left(\frac{\exp(a_i)}{\exp(g_i)} \right) \right]$$

since $g_{n-1} = a_n$, $g_n = a_{n-1}$, and $g_k = a_k$ for $k < n-1$, then:

$$L^{lNet}(g) - L^{lNet}(a) = \frac{1}{\sum_{k=1}^{n} \exp(y_k)} \left[\exp(y_{n-1}) \log\left(\frac{\exp(a_{n-1})}{\exp(a_n)} \right) + \exp(y_n) \log\left(\frac{\exp(a_n)}{\exp(a_{n-1})} \right) \right]$$
$$= (a_{n-1} - a_n) \left(\frac{\exp(y_{n-1})}{\sum_{k=1}^{n} \exp(y_k)} - \frac{\exp(y_n)}{\sum_{k=1}^{n} \exp(y_k)} \right)$$

The differences in penalizations are :

$$\begin{cases} L^{\text{lNet}}(h) - L^{\text{lNet}}(a) = (a_1 - a_2) \left(\frac{\exp(y_1)}{\sum_{k=1}^{n} \exp(y_k)} - \frac{\exp(y_2)}{\sum_{k=1}^{n} \exp(y_k)} \right) \\ \\ L^{\text{lNet}}(g) - L^{\text{lNet}}(a) = (a_{n-1} - a_n) \left(\frac{\exp(y_{n-1})}{\sum_{k=1}^{n} \exp(y_k)} - \frac{\exp(y_n)}{\sum_{k=1}^{n} \exp(y_k)} \right) \end{cases}$$

Now, suppose we have a scoring error of the same magnitude $a_1 - a_2 = a_{n-1} - a_n = \delta_s$, and the difference of the returns between adjacent ranks is consistent $y_{n-1} - y_n = y_1 - y_2 = \delta_r$. In such a scenario, the penalization would be larger for misrankings at the top elements.

$$\begin{cases} L^{\text{lNet}}(h) - L^{\text{lNet}}(a) \propto \exp(y_1) - \exp(y_2) \\ L^{\text{lNet}}(g) - L^{\text{lNet}}(a) \propto \exp(y_{n-1}) - \exp(y_n) \end{cases}$$

This bias arises from the exponential nature of the softmax function. Indeed:

$$\frac{L^{\text{lNet}}(h) - L^{\text{lNet}}(a)}{L^{\text{lNet}}(g) - L^{\text{lNet}}(a)} = \frac{\exp(y_1) - \exp(y_2)}{\exp(y_{n-1}) - \exp(y_n)} = \exp(y_2 - y_n)$$

with,

$$\begin{aligned}\exp(y_1) - \exp(y_2) &= \exp(y_2 - y_n)(\exp(y_1 - y_2 + y_n) - \exp(y_n)) \\ &= \exp(y_2 - y_n)(\exp(\delta_r + y_n) - \exp(y_n)) \\ &= \exp(y_2 - y_n)(\exp(y_{n-1}) - \exp(y_n))\end{aligned}$$

Since $y_2 \geq y_n$, $\frac{L^{\text{lNet}}(h) - L^{\text{lNet}}(a)}{L^{\text{lNet}}(g) - L^{\text{lNet}}(a)} > 1$ and $L^{\text{lNet}}(h) \geq L^{\text{lNet}}(g)$. The difference in penalization is exponentially as significant as the difference of scores between the second-best item and the worst one[8].

3.7.3 Ranking penalization bias in ListMLE

ListMLE is a listwise ranking algorithm that operates within a probabilistic framework rooted in the Plackett-Luce model. At its core, ListMLE formulates the likelihood of observing a given ranking $(y^{-1}(1), ..., y^{-1}(n))$ given input data $(x_1, ..., x_n)$ and scoring function f as follows:

$$P(y|x; f) = P(y^{-1}(1)|x; f) \prod_{i=2}^{n} P(y^{-1}(i)|x, y^{-1}(1), \ldots, y^{-1}(i-1); f)$$

This formulation decomposes the probability of a ranking into stepwise conditional probabilities. Each conditional probability represents the likelihood of placing an object in its accurate position within the sequence, considering the prior positioning of preceding elements. Specifically, $P(y^{-1}(i)|x, y^{-1}(1), \ldots, y^{-1}(i-1); f)$ denotes the probability of correctly positioning the i-th object given the rankings of previously placed objects. Incorporating this probabilistic framework, the ListMLE loss function $L^{lMLE}(f, \mathcal{X}, \mathcal{Y})$ is defined as the negative logarithm of the likelihood of $P(y|x; f)$:

$$L^{lMLE}(f, \mathcal{X}, \mathcal{Y}) = -\log P(y|x; f) = -\log \left(\prod_{i=1}^{n} \frac{\exp(f(x_{y^{-1}(i)}))}{\sum_{k=i}^{n} \exp(f(x_{y^{-1}(k)}))} \right)$$

To understand why ListMLE demonstrates superior ranking performance for bottom elements as opposed to top elements, we investigate the variation in the loss value by interchanging two elements from the perfectly ranked list—one from the top and one from the bottom. This analytical comparison illuminates the distinct impact of ranking errors on the loss function for both upper and lower segments of the list. Suppose $a = a_1, a_2, \ldots, a_n$ represents scores that provide a correct ranking, with $a_1 > a_2 > \ldots > a_n$. Additionally, let h denotes scores that misrank the top two elements, and g denote scores that

[8]The distribution's skewness adds complexity this phenomenon. In negatively skewed distributions, larger errors at the bottom counteract the exponential bias, lessening its impact. Conversely, in positively skewed distributions — commonly observed in financial returns — larger errors at the top intensify the exponential bias, leading to an amplified penalization effect.

misrank the bottom two elements. For h, which misranks the top two elements:

$$\begin{cases} h_1 = a_2 & \text{indicating that the top-ranked score is now } a_2 \\ h_2 = a_1 & \text{indicating that the second-ranked score is now } a_1 \\ h_k = a_k & \text{for } 3 \leq k \leq n \text{ meaning that the remaining scores are unchanged} \end{cases}$$

For g, which misranks the bottom two elements:

$$\begin{cases} g_n = a_{n-1} & \text{indicating that the last-ranked score is now } a_{n-1} \\ g_{n-1} = a_n & \text{indicating that the second-to-last-ranked score is now } a_n \\ g_k = a_k & \text{for } 1 \leq k \leq n-2 \text{ meaning that the remaining scores are unchanged} \end{cases}$$

The penalization for misranking the bottom two elements given a set of scores $L^{lMLE}(g) - L^{lMLE}(a)$ evaluates to:

$$\begin{aligned} L^{lMLE}(g) - L^{lMLE}(a) &= \log\left(\frac{\prod_{i=1}^{n}\exp(a_i)}{\sum_{k=i}^{n}\exp(a_k)}\right) - \log\left(\frac{\prod_{i=1}^{n}\exp(g_i)}{\sum_{k=i}^{n}\exp(g_k)}\right) \\ &= \sum_{i=1}^{n}\log\left(\frac{\sum_{k=i}^{n}\exp(g_k)}{\sum_{k=i}^{n}\exp(a_k)}\right) \\ &= \log\left(\frac{\exp(a_n)+\exp(a_{n-1})}{\exp(a_{n-1})+\exp(a_n)}\right) + \log\left(\frac{\exp(a_{n-1})}{\exp(a_n)}\right) \\ &= \log\left(\frac{\exp(a_{n-1})}{\exp(a_n)}\right) \\ &= a_{n-1} - a_n \end{aligned}$$

The penalization for misranking the top two elements given a set of scores $L^{lMLE}(h) - L^{lMLE}(a)$ evaluates to:

$$\begin{aligned} L^{lMLE}(h) - L^{lMLE}(a) &= \log\left(\frac{\prod_{i=1}^{n}\exp(a_i)}{\sum_{k=i}^{n}\exp(a_k)}\right) - \log\left(\frac{\prod_{i=1}^{n}\exp(h_i)}{\sum_{k=i}^{n}\exp(h_k)}\right) \\ &= \sum_{i=1}^{n}\log\left(\frac{\sum_{k=i}^{n}\exp(h_k)}{\sum_{k=i}^{n}\exp(a_k)}\right) \\ &= \log\left(\frac{\exp(a_2)+\exp(a_1)+\sum_{k=3}^{n}\exp(a_k)}{\exp(a_1)+\exp(a_2)+\sum_{k=3}^{n}\exp(a_k)}\right) + \log\left(\frac{\exp(a_1)+\sum_{k=3}^{n}\exp(a_k)}{\exp(a_2)+\sum_{k=3}^{n}\exp(a_k)}\right) \\ &= \log\left(\frac{\exp(a_1)+\sum_{k=3}^{n}\exp(a_k)}{\exp(a_2)+\sum_{k=3}^{n}\exp(a_k)}\right) \\ &= a_1 - a_2 + \underbrace{\left[\log\left(1+\frac{\sum_{k=3}^{n}\exp(a_k)}{\exp(a_1)}\right) - \log\left(1+\frac{\sum_{k=3}^{n}\exp(a_k)}{\exp(a_2)}\right)\right]}_{A} \\ &\leq a_1 - a_2 \end{aligned}$$

Given the same magnitude of error, i.e. $a_1 - a_2 = a_{n-1} - a_n = \delta \geq 0$, the penalization for misranking the bottom two elements is directly δ . In contrast, the penalization for misranking the top two elements is less than δ, influenced by A. This indicates that ListMLE inherently places more weight on errors at the

bottom of the list. Specifically, for a given error magnitude δ, the penalization is unequivocally larger or equal for a misranking at the bottom than for a misranking at the top of the list, i.e. $L^{lMLE}(g) = L^{lMLE}(a) + \delta > L^{lMLE}(h)$.

As a consequence of ListMLE's heightened sensitivity to errors at the bottom of the list, the trained algorithm exhibits a bias, resulting in fewer errors at the bottom where accurate ordering of lower-ranked items is particularly crucial.

Chapter 4

Unifying Asset Ranking and Portfolio Weighting through a Multi-Task Neural Network[1]

Abstract

In this chapter, we present a novel approach that integrates asset ranking and portfolio weighting within a single framework. Unlike traditional methods, which separate asset ranking from portfolio weighting, this research employs a multi-task neural network to concurrently learn asset rankings and optimize the number of assets for long and short positions. This innovation aims to better align the investment strategy with investor preferences right from the model prediction phase. To assess its effectiveness, we conduct experiments using historical weekly market data from China A-shares. The findings indicate that incorporating portfolio weighting into a multi-task learning framework significantly improves out-of-sample financial performance compared to benchmark methods that rely on heuristics or historical estimations.

4.1 Introduction

Although the integration of machine learning into asset management continues to grow steadily, most applications remain limited to predicting future asset performance. Among machine learning techniques, learning-to-rank (LTR) approaches demonstrate superior performance compared to traditional heuristic-based or regress-then-rank methods ([81]). These LTR approaches are particularly effective for cross-sectional strategies, allowing for more accurate asset ranking based on expected returns. However, many studies on LTR do not directly address portfolio weighting. Instead, after determining the rankings,

[1]This chapter is based on the research article *Unifying Asset Ranking and Portfolio Weighting through a Multi-Task Neural Network* co-authored with Ilias Mellouki and Guillaume Boulanger and published in the *Journal of Financial Data Science*, Spring 2025, 7 (2) 84-104, DOI: 10.3905/jfds.2025.1.190.

asset allocation is typically executed using heuristics or straightforward optimization techniques. These optimization processes often incorporate various constraints such as investor preferences, risk aversion parameters, and other specific requirements ([20]). This two-step process—predicting asset ranks and then optimizing portfolio weights—has been the standard practice, but leaves room for improvement.

In this chapter, we propose a novel approach that leverages a multi-task neural network model to enhance the traditional LTR framework. Our method combines the tasks of ranking assets and determining the optimal allocation to each asset, whether long or short. By integrating these tasks into a single, cohesive model, and shifting more of the optimization burden to the machine learning model, we seek to enhance the robustness and accuracy of investment decisions (see Figure 4.1).

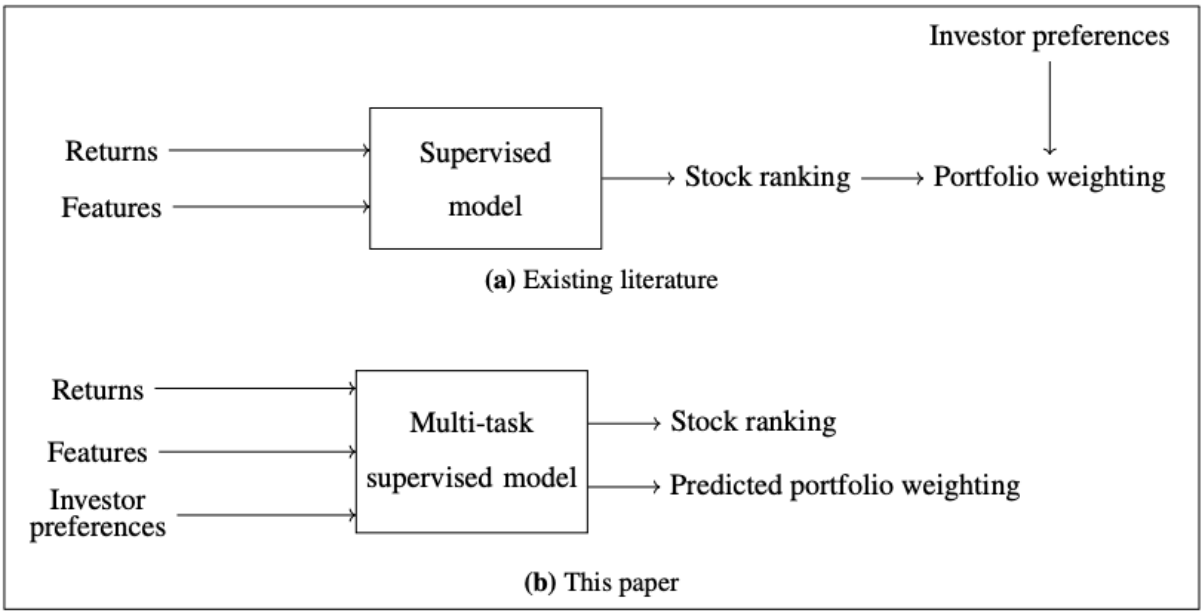

Figure 4.1: **Comparative diagram**. The upper panel illustrates the conventional approach where portfolio construction follows the asset ranking prediction. The lower panel depicts the core concept introduced in this chapter.

This chapter contributes to the existing literature in several ways:

1. **Enhancing financial performance**: We demonstrate that incorporating portfolio weighting into the learning process can significantly enhance overall portfolio financial performance.
2. **Meeting target metrics**: Our system learns portfolio constraints and generates more efficient portfolios that align with investors preferences.
3. **Asset ranking versus portfolio weighting**: We investigate the potential trade-off between learning assets ranking and optimizing portfolio weights. The existence of an optimal balance supports optimizing both tasks.

The remainder of the chapter proceeds as follows. In Section 4.3, we carry out a dataset overview, then present the model architecture and supervision procedure. In Section 4.4, we discuss the training

and backtesting methodologies. Finally, Section 4.5 presents benchmark approaches and evaluates the performance of the multi-task model.

4.2 Related literature

The use of LTR algorithms has been increasingly explored in recent studies. [81] emphasize the critical importance of accurately ranking assets before portfolio construction in cross-sectional systematic strategies. Traditional techniques usually use simple heuristics, outputs from standard regression, or classification models for ranking, which have been shown to be suboptimal in other domains such as information retrieval. To overcome this shortcoming, [81] propose a framework that incorporates LTR algorithms to enhance cross-sectional portfolios, demonstrating significant improvements in ranking accuracy and trading performance through pairwise and listwise learning structures. Similarly, [94] discuss the application of LTR in factor strategies. They propose a new listwise LTR loss function that emphasizes both the top and bottom of a ranked list, suitable for long-short strategies. Their method, tested on a dataset of 68 factors in the China A-share market, shows substantial outperformance with a high Sharpe ratio, illustrating the strength of using LTR algorithms in predicting ranked asset lists instead of direct return predictions.

In all these approaches, the portfolio construction task is not directly tackled, and often involves longing and shorting a fixed number of assets at each rebalancing timestamp. Recent research explore the direct optimization of portfolio weights using machine learning algorithms. [95] propose optimizing Sharpe ratios via gradient ascent in a neural network architecture, deriving portfolio weights from a given set of predictors. Their framework avoids the need for forecasting expected returns, allowing for direct optimization of portfolio weights. This approach, tested on ETFs of market indices, demonstrates superior performance over traditional methods, especially during periods of financial instability. Later, [20] present an asset allocation strategy that engineers optimal weights before feeding them to a supervised learning algorithm. This approach contrasts with traditional methods by allowing the model to learn risk measures, preferences, and constraints within a flexible, forward-looking, and non-linear framework. Their empirical analysis shows that predicting optimal weights directly results in more stable portfolios with better risk-adjusted performance.

Besides, some fields of research are investigating the use of multi-task models for portfolio construction. [78] demonstrate that a diversified risk-adjusted time-series momentum portfolio, utilizing multi-task learning in a deep neural network architecture, can deliver substantial abnormal returns and provide tail risk protection. Their approach jointly learns portfolio construction and various auxiliary tasks related to volatility, showing that adding auxiliary tasks enhances portfolio performance. Linked to LTR approaches, [67] propose using auxiliary tasks to aid in ranking assets through LTR algorithms. They introduce a deep multi-task learning solution called Multi-Task Stock Ranking (MTSR), which leverages easily-trained auxiliary tasks to assist the learning of the main listwise ranking task. Their experiments on real-world stock data demonstrate the superiority of MTSR over several state-of-the-art methods.

The present chapter aims to propose a general framework to unify asset ranking and portfolio weighting through the use of a multi-task neural network, combining the strengths of LTR algorithms and direct

portfolio weight prediction to achieve improved investment performance.

4.3 Model overview

4.3.1 Dataset overview

To empirically assess the effectiveness of our proposed approach, we employ a dataset containing information about equities characterized by various factors. This dataset is derived from the study by [93] and is available on GitHub[2]. Specifically, the dataset $\mathcal{D}$ comprises $T = 630$ weekly observations related to $N = 74$ assets within the China A-share market, each described by $M = 62$ factors. The dataset spans from December 29, 2006, to February 1, 2019. Notably, these 74 assets are predominantly listed in the HS300 index and are characterized by high liquidity.

4.3.2 Model architecture

We propose a multi-task neural network, extending previous LTR frameworks ([94]). The network is given $\mathcal{D}$ as input, which possesses a three-dimensional structure with dimensions $T \times N \times M$. It employs a structure of $[124 \times 248 \times 31]$ for the initial three fully connected hidden layers, which are shared across three specific tasks. After these shared layers, the network branches into three separate heads, each dedicated to a distinct task.

The first head is designed to predict the ranking of assets, and includes a single fully connected output layer with a ReLU activation function, producing an output vector of size N, where each element corresponds to the predicted score of an asset. The second and third heads are designed to forecast the optimal number of assets to long and short. Both heads consist of one additional fully connected hidden layer. The output of this hidden layer is then flattened before passing through a final fully connected output layer, which employs a softmax activation function, ensuring that the output is a probability distribution over the possible number of assets to long and short. Consequently, the final outputs for each head is a vector of size $(k_{\max} + 1)$, where $k_{\max}$ represents the maximum number of assets allowed to be longed or shorted in the strategy[3]. We set $k_{\max} = \frac{N}{2}$ to allow the trading of all the N assets. This choice empowers the model to determine the exact number of assets to trade from the full set of possibilities.

The network architecture is illustrated in Figure 4.2. This architectural design enables simultaneous learning across multiple tasks. By combining the loss functions of the three tasks, the model ensures that the scoring function benefits from a comprehensive set of investment-related objectives. The rationale behind this multi-task approach is to unify assets ranking and portfolio weighting, which is expected to enhance the overall performance of the trading system.

[2]https://github.com/TCtobychen/ListFold.

[3]We include the option of not taking any long or short positions, which explains why the dimensionality is $(k_{\max} + 1)$ rather than $k_{\max}$.

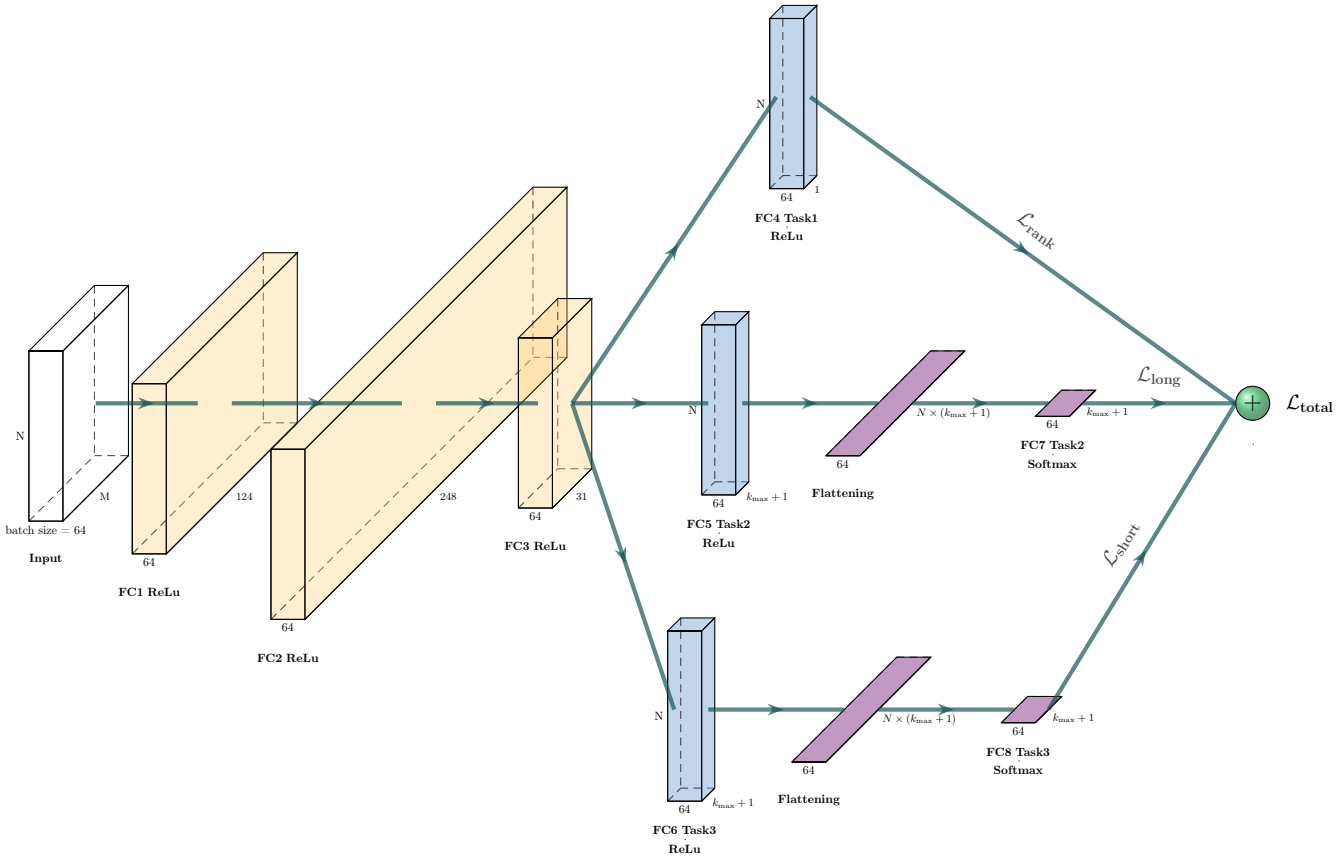

Figure 4.2: Architecture of the multi-task neural network

4.3.3 Loss functions

Let $\mathcal{X} = \{x_1, \cdots, x_N\}$ represent the objects to be ranked, which in our context are financial assets. At timestamp t, each object x_i is associated with a true score $y_{i,t}$, for instance representing the returns of the financial asset. The set of these scores at time t is denoted as $\mathcal{Y}_t^{\text{rank}} = \{y_{1,t}, \ldots, y_{N,t}\}$. If $y_{i,t} > y_{j,t}$, it implies that the score of x_i surpasses that of x_j, and consequently, x_i achieves a superior ranking than x_j, with 1 indicating the best ranking and N representing the lowest rank. Furthermore, at time t, we denote $y_t^{-1}(i)$ the index of the item that is truly ranked in the i-th position. Hence, $x_{y_t^{-1}(1)}$ is the element that is truly ranked at the best position at time t and $x_{y_t^{-1}(N)}$ is the element that is truly ranked at the worst position at time t. Let f be the output function of the multi-task neural network. We aim to learn the optimal function f from the training dataset by minimizing a combined loss function from three different tasks. To achieve this, we define three separate loss functions, one for each task, and sum them to form a single comprehensive loss function.

For the ranking task, we use the ListMLE loss function introduced in [93].

$$\begin{aligned}\mathcal{L}_t^{\text{rank}}(f, \mathcal{X}, \mathcal{Y}_t^{\text{rank}}) &= -\log \prod_{i=1}^{N} \frac{\phi\left(f\left(x_{y_t^{-1}(i)}\right)\right)}{\sum_{k=i}^{N} \phi\left(f\left(x_{y_t^{-1}(k)}\right)\right)} \\ &= -\sum_{i=1}^{N} \left(\log \phi\left(f\left(x_{y_t^{-1}(i)}\right)\right) - \log \sum_{k=i}^{N} \phi\left(f\left(x_{y_t^{-1}(k)}\right)\right)\right)\end{aligned}$$

where $\phi(z) = e^z$.

For the second and third tasks, we aim at predicting the optimal number of assets to long and short. We frame this problem as a multiclass classification problem, where each of the $(k_{\max} + 1)$ possible outcomes represents a distinct class. The class 0 indicates that no asset will be longed (or shorted). Hence, tasks 2 and 3 focus on predicting the optimal number of assets to long and short, respectively, and are trained using the classical categorical cross-entropy loss function:

$$\mathcal{L}_t^{\text{long}}(f, \mathcal{X}, \mathcal{Y}_t^{\text{long}}) = -\sum_{i=0}^{k_{\max}} y_{i,t}^{\text{long}} \log(p_{i,t}^{\text{long}})$$

$$\mathcal{L}_t^{\text{short}}(f, \mathcal{X}, \mathcal{Y}_t^{\text{short}}) = -\sum_{i=0}^{k_{\max}} y_{i,t}^{\text{short}} \log(p_{i,t}^{\text{short}})$$

where:

- $\mathcal{Y}_t^{\text{long}} = \left\{ y_{0,t}^{\text{long}}, \ldots, y_{k_{\max},t}^{\text{long}} \right\}$ (resp. $\mathcal{Y}_t^{\text{short}}$)
  with $y_{i,t}^{\text{long}} = \begin{cases} 1 & \text{if } i \text{ is the optimal number of assets to long at time } t \\ 0 & \text{otherwise.} \end{cases}$
- $p_{i,t}^{\text{long}}$ (resp. $p_{i,t}^{\text{short}}$) represents the predicted probability of longing (resp. shorting) asset i at time t.

Finally the total loss is expressed as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rank}} + \mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$$

To summarize, learning f over a training/validation time horizon $\mathcal{T} = \mathcal{T}_{\text{train}} \cup \mathcal{T}_{\text{val}}$ depends on $\mathcal{X}$, f, and three sets of labels $\mathcal{Y}_{\text{rank}}$, $\mathcal{Y}_{\text{long}}$, $\mathcal{Y}_{\text{short}}$ where:

$$\mathcal{Y}^{\text{rank}} = \bigcup_{t \in \mathcal{T}} \mathcal{Y}_t^{\text{rank}} \text{ and } \mathcal{Y}^{\text{long}} = \bigcup_{t \in \mathcal{T}} \mathcal{Y}_t^{\text{long}} \quad (\text{resp. } \mathcal{Y}^{\text{short}})$$

In the next section, we present the methodology used to supervise each of the three tasks, i.e. to generate $\mathcal{Y}^{\text{rank}}$, $\mathcal{Y}^{\text{long}}$ and $\mathcal{Y}^{\text{short}}$.

4.3.4 Ranking labels $\mathcal{Y}^{\text{rank}}$

The labels $\mathcal{Y}^{\text{rank}}$ designate the scores given to each asset, which are used to establish their rankings. We determine the score of each asset at time t based on its cumulative return over a future time window τ. More precisely, the score of asset i at time t is expressed as:

$$y_{i,t}^{\text{rank}} = \prod_{j=t+1}^{t+1+\tau} (1 + r_{i,j})$$

where $r_{i,j}$ represents the return of asset i at time j.

These cumulative returns provide a measure of the asset's future performance over the specified future window, which is then used as the score for learning the future ranks of the assets. Introducing the window τ instead of relying on one-step returns aims to provide more robustness into the supervision of the ranking task.

4.3.5 Long/short labels $\mathcal{Y}^{\text{long}}/\mathcal{Y}^{\text{short}}$

In this section, we explore two distinct methodologies for calculating $\mathcal{Y}^{\text{long}}$ and $\mathcal{Y}^{\text{short}}$, focusing on portfolio optimization strategies. These approaches aim to provide flexible portfolio optimization techniques, allowing investors to tailor strategies to their specific goals and risk tolerance. The first approach involves either maximizing expected returns or minimizing portfolio volatility, offering a straightforward strategy based on a single objective. The second approach, however, introduces optimization under specific constraints, such as maximizing expected returns while maintaining a target volatility level. This constrained optimization allows for a more controlled, customized strategy, addressing the need for balancing risk and return in line with the investor's preferences.

Let $P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})$ represent the portfolio at time t, that includes long positions in the top k^{long} assets and short positions in the bottom k^{short} assets, based on their $\mathcal{Y}^{\text{rank}}$ scores. Furthermore, the performance of this portfolio over a defined window τ is denoted by $g_\tau[P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})]$ where g is a specific financial metric of interest, such as the expected return (μ) or the volatility (σ). For consistency, the window τ for ranking labels $\mathcal{Y}^{\text{rank}}$ must be the same as the one used to compute long-short labels. We first present the optimization problems intended to generate the labels $\mathcal{Y}^{\text{long}}$ and $\mathcal{Y}^{\text{short}}$. These optimization problems are designed to enhance one or multiple portfolio performance metrics, considering both constrained and unconstrained scenarios.

Unconstrained portfolio optimization

For a single criterion, denoted by the financial metric of interest g and a fixed window τ, determining the optimal number of assets to long-short at time t involves solving the following optimization problem:

$$(k_t^{\text{long}*}, k_t^{\text{short}*}) = \underset{(k^{\text{long}}, k^{\text{short}})}{\arg\max}\ g_\tau[P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})] \tag{4.1}$$

Considering multiple criteria $\mathcal{C} = [g_\tau^1, \dots, g_\tau^n]$, we rely on a more holistic approach capable of optimizing multiple financial metrics concurrently, called the Technique for Order Preference by Similarity to Ideal Solution (TOPSIS). This framework ([52]) is an effective multi-criteria decision-making method used for ranking and selection among portfolio alternatives. In our context, each alternative represents a unique combination of k^{long} and k^{short}. Let us now illustrate TOPSIS practical application for our labeling process.

At time t, we construct a decision matrix X_t with $K = (k_{\max} + 1)^2 - 1$ rows [4], and n columns. Each row represents a different portfolio configuration (i.e. different combinations of k^{long} and k^{short}), and each column represents a performance criterion.

$$X_t = \begin{bmatrix} \varepsilon_1 g_\tau^1[P(\mathcal{Y}_t^{\text{rank}}, 0, 1)] & \dots & \varepsilon_N g_\tau^n[P(\mathcal{Y}_t^{\text{rank}}, 0, 1)] \\ \vdots & \dots & \vdots \\ \varepsilon_1 g_\tau^1[P(\mathcal{Y}_t^{\text{rank}}, k_{\text{long}}, k_{\text{short}})] & \dots & \varepsilon_n g_\tau^n[P(\mathcal{Y}_t^{\text{rank}}, k_{\text{long}}, k_{\text{short}})] \\ \vdots & \dots & \vdots \\ \varepsilon_1 g_\tau^1[P(\mathcal{Y}_t^{\text{rank}}, k_{\max}, k_{\max})] & \dots & \varepsilon_n g_\tau^n[P(\mathcal{Y}_t^{\text{rank}}, k_{\max}, k_{\max})] \end{bmatrix}$$

It is important to define the metrics such that a higher value is indicative of a better outcome. Therefore, for a criterion i that needs to be maximized, $\varepsilon_i = 1$, and for a criterion that needs to be minimized $\varepsilon_i = -1$. Additionally, to ensure that all criteria are on the same scale, we normalize the decision matrix. The normalized value $x_{i,j}^{\text{normalized}}$ for the element $x_{i,j}$ is calculated as:

$$x_{i,j}^{\text{norm}} = \frac{x_{i,j}}{\sqrt{\sum_{i=1}^{K} x_{ij}^2}}$$

Next, we construct the weighted normalized decision matrix Z_t by multiplying each column of the normalized matrix X_t^{norm} by a criterion weight. This allows us to assign different weights w_i to each criterion, reflecting their relative importance. We then define:

$$Z_t = \begin{bmatrix} w_1 x_{0,1}^{\text{norm}} & \dots & w_n x_{0,1}^{norm} \\ \vdots & \dots & \vdots \\ w_1 x_{i,j}^{\text{norm}} & \dots & w_n x_{i,j}^{\text{norm}} \\ \vdots & \dots & \vdots \\ w_1 x_{k_{\max},k_{\max}}^{\text{norm}} & \dots & w_n x_{k_{\max},k_{\max}}^{\text{norm}} \end{bmatrix}$$

where $w \in [0,1]^n$, with $\sum_{i=1}^{n} w_i = 1$.

By emphasizing certain metrics over others, we can tailor the labeling to prioritize specific objectives. For instance, assigning a higher weight to expected returns focuses on potential gains, while prioritizing volatility emphasizes stability. This weighting aligns the decision matrix with investors constraints.

[4] since we do not allow the case $(k^{\text{long}}, k^{\text{short}}) = (0, 0)$.

Besides, we define the ideal solution z^+ and the anti-ideal solution z^- as follows:

$$z^+ = \left(\max_{1 \leq i \leq K} z_{i,1}, \ldots, \max_{1 \leq i \leq K} z_{i,n} \right)$$
$$z^- = \left(\min_{1 \leq i \leq K} z_{i,1}, \ldots, \min_{1 \leq i \leq K} z_{i,n} \right)$$

If we denote by $z_{k_{\text{long}},k_{\text{short}}}$ the row of Z_t corresponding to the portfolio $P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})$, then the best pair according to this framework is the one maximizing the relative closeness to the ideal solution:

$$(k_t^{\text{long}*}, k_t^{\text{short}*}) = \underset{(k_{\text{long}},k_{\text{short}})}{\arg\max} \frac{\|z_{k_{\text{long}},k_{\text{short}}} - z^-\|}{\|z_{k_{\text{long}},k_{\text{short}}} - z^-\| + \|z_{k_{\text{long}},k_{\text{short}}} - z^+\|} \tag{4.2}$$

where $\|.\|$ is the euclidian distance.

It can be argued that this alternative has the farthest distance from z^- and is the closest alternative to z^+. In addition, one can notice that solving problem (4.1) is equivalent to choosing $\mathcal{C} = [g_\tau]$ following TOPSIS methodology.

Constrained portfolio optimization

In addition to exploring unconstrained optimization, we investigate our framework's ability to optimize a financial metric g^i while targeting a specific level for another metric g^j:

$$\begin{aligned}(k_t^{\text{long}*}, k_t^{\text{short}*}) = \underset{(k^{\text{long}},k^{\text{short}})}{\arg\max}\ & \varepsilon_i g_\tau^i[P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})] \\ \text{s.t.}\quad & \varepsilon_j g_\tau^j[P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})] \geq \varepsilon_j g_{\text{tgt}}^j\end{aligned} \tag{4.3}$$

For instance, a common portfolio optimization problem aims at maximizing returns while maintaining portfolio volatility at or below a specified target level σ_{tgt}.

$$\begin{aligned}(k_t^{\text{long}*}, k_t^{\text{short}*}) = \underset{(k^{\text{long}},k^{\text{short}})}{\arg\max}\ & \mu_\tau[P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})] \\ \text{s.t.}\quad & \sigma_\tau[P(\mathcal{Y}_t^{\text{rank}}, k^{\text{long}}, k^{\text{short}})] \leq \sigma_{tgt}\end{aligned} \tag{4.4}$$

The computation of the optimal pair for problems (4.1), (4.2) and (4.3) is done using a grid search, since there is a finite number of candidates K. Finally, $\mathcal{Y}_t^{\text{long}}$ and $\mathcal{Y}_t^{\text{short}}$ represent the corresponding one-hot encoded labels at time t. When no solution exists among the candidates, we take the closest solution to satisfy the constraint.

When determining the optimal number of assets to long and short ($k^{\text{long}*}$, $k^{\text{short}*}$), the amount invested in each asset is crucial, as it affects the financial metric values calculated over the window τ. One approach is to invest the same amount in all selected assets, corresponding to the use of a uniform distribution over the ranks. Alternatively, we can assign greater weight to assets at the extremes of the ranking spectrum using an exponential distribution. These two weighting schemes are detailed further in

Section 4.4.2.

4.4 Training and backtesting methodologies

4.4.1 Training

We adopt a temporal rolling strategy to train, validate and test the model discussed in Section 4.3.2. Specifically, our training period spans $|\mathcal{T}_{\text{train}}| = 300$ weeks, while both the validation ($|\mathcal{T}_{\text{val}}|$) and test ($|\mathcal{T}_{\text{test}}|$) periods consist of 16 weeks each. As shown in Figure 4.3, successive folds are separated by $|\mathcal{T}_{\text{test}}|$ weeks. This process is repeated until the end of the test set is reached. Furthermore, we forecast $\tau = 20$ weeks ahead to generate labels[5]. To prevent potential data correlations, we exclude the 20 weeks immediately following both the training and validation periods when defining the validation and test sets.

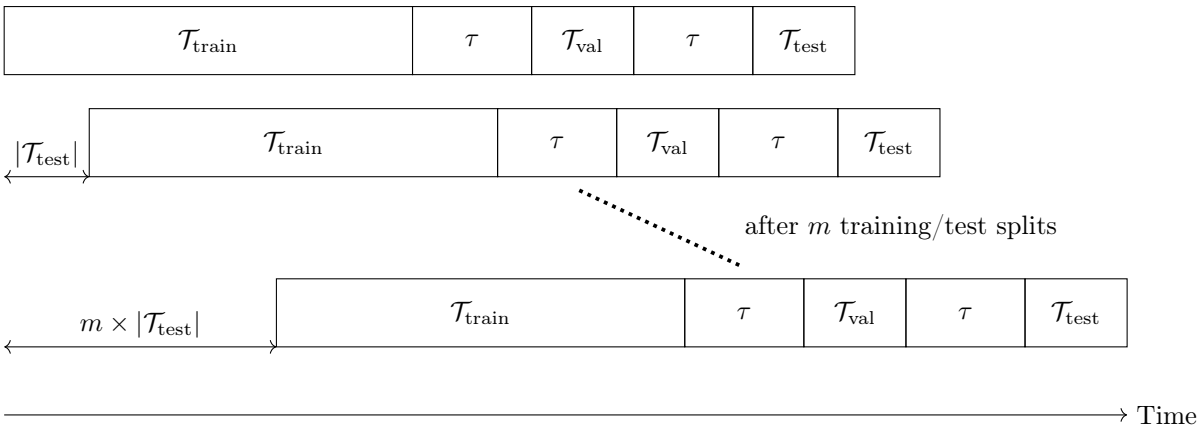

Figure 4.3: Rolling training/validation/testing strategy of the model

To ensure comparability among features with varying orders of magnitude, we apply standard scaler normalization. The model training is conducted with a mini-batch size of 64. We employ early stopping as a regularization technique to prevent overfitting, with a patience parameter set to 5 epochs and a minimum delta of 0.001. The training process stops if the validation loss does not improve by at least 0.001 over 5 consecutive epochs. In addition, the Adam optimizer, with a learning rate of 10^{-4}, is employed to ensure efficient convergence. We initialize the network biases to be 0.05 and the weights $W_{i,j}$ at each layer with the following commonly used heuristic ([38]):

$$W_{i,j} \sim \mathcal{U}(-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}})$$

where d is the input dimension of the corresponding layer, and $\mathcal{U}$ designates the uniform distribution.

[5]The choice of $\tau = 20$ allows for a robust computation of expected future returns and volatility, which are essential for calculating $\mathcal{Y}^{\text{rank}}$, $\mathcal{Y}^{\text{long}}$, and $\mathcal{Y}^{\text{short}}$. While this selection provides a reasonable foundation, further analysis on the robustness of τ across different values has not yet been conducted. This could be a valuable direction for future research.

4.4.2 Backtesting

For backtesting, we perform weekly portfolio rebalancing, applying the model to generate predictions for each timestamp in the test set. The optimal number of assets to long or short at each time is determined based on the model predictions. Assets are selected for longing or shorting based on their rankings, with a portion of \$1 invested in each selected asset. The excess return at the end of each day is calculated based on the returns of the longed and shorted assets, adjusted for transaction fees. Let us define the model predictions on the test set at a certain time t:

- $(x_1^*, \dots, x_N^*)$ the ranking of the assets in decreasing order.
- $(k_t^{\text{long}*}, k_t^{\text{short}*})$ the optimal number of assets to long-short.

The assets $(x_1^*, \dots, x_{k_t^{\text{long}*}}^*)$ are longed, whereas the assets $(x_{N-k_t^{\text{short}*}+1}^*, \dots, x_N^*)$ are shorted. Let us denote $\mathcal{I}_t^{\text{long}} = \{1, \dots, k_t^{\text{long}*}\}$ and $\mathcal{I}_t^{\text{short}} = \{N - k_t^{\text{short}*} + 1, \dots, N\}$ the indices of the assets to be longed and shorted. At each time t, \$1 is invested to long-short the selected assets. Between each rebalancing period, assets are bought and sold to align with the updated portfolio allocation. For $i \in \{1, \dots, N\}$, let us call $s_{i,t}$ the amount of money used to long or short the asset i at time t. Positive values of $s_{i,t}$ indicate long positions, while negative values indicate short positions. If $s_{i,t} = 0$, no investment is made in the asset. We have $\sum_{i \in \mathcal{I}_t^{\text{long}} \bigcup \mathcal{I}_t^{\text{short}}} |s_{i,t}| = \$1$. The excess return made out of this investment strategy at time $t+1$ is:

$$\delta_t = \sum_{i \in \mathcal{I}_t^{\text{long}}} s_{i,t} r_i^{t+1} + \sum_{i \in \mathcal{I}_t^{\text{short}}} s_{i,t} r_i^{t+1} - F_t$$

Where $F_t = \frac{1}{2} \times \mathcal{TO}_t \times 10$ bps are the transaction fees at time t and $\mathcal{TO}_t = \sum_{i=1}^{N} |s_{i,t} - s_{i,t-1}|$ is the weekly turnover.

We consider two approaches for determining the amount of money invested in each asset: uniform density and exponential density. The uniform density assigns equal weights to all selected assets:

$$\forall i \in \mathcal{I}_t^{\text{long}} \cup \mathcal{I}_t^{\text{short}} : |s_{i,t}| = \frac{1}{|\mathcal{I}_t^{\text{long}} \bigcup \mathcal{I}_t^{\text{short}}|} = \frac{1}{k_t^{\text{long}*} + k_t^{\text{short}*}}$$

The exponential density assigns more weight to assets ranked at the edge from both long and short positions. Specifically, the exponential density function is computed using a softmax to ensure that the invested amounts sum up to \$1.

$$\forall i \in \mathcal{I}_{\text{long}} : \quad s_{i,t} = \frac{f(i, \theta)}{\sum_{j \in \mathcal{I}_t^{\text{long}}} f(j, \theta) + \sum_{j \in \mathcal{I}_t^{\text{short}}} f(n - j + 1, \theta)}$$

$$\forall i \in \mathcal{I}_{\text{short}} : \quad s_{i,t} = \frac{-f(n - i + 1, \theta)}{\sum_{j \in \mathcal{I}_t^{\text{long}}} f(j, \theta) + \sum_{j \in \mathcal{I}_t^{\text{short}}} f(n - j + 1, \theta)}$$

Where $f(x, \theta) = \theta e^{-\theta x}$ and $\theta = 0.1$.

4.5 Performances

The evaluation of our results involves the use of two sets of metrics, both calculated for out-of-sample test periods. The first set encompasses standard financial metrics, as suggested by [81], evaluating aspects such as profitability, risk, and financial performance. The second set is derived from literature on information retrieval and ranking.

In terms of profitability, we examine expected yearly returns adjusted to risk free ($\mu_{\text{ann}} - r_f$) and the percentage of positive returns at the portfolio level. For risk metrics, we compute the annualized volatility (σ_{ann}), maximum drawdown ($\mathcal{MDD}$). Financial performance metrics for risk-adjusted evaluation include the Sharpe ratio ($\frac{\mu_{\text{ann}} - r_f}{\sigma_{\text{ann}}}$) and the Sortino ratio ($\frac{\mu_{\text{ann}} - r_f}{Downside\ Deviation}$). Additionally, we include the average profit divided by the average loss. The risk-free rate r_f that we use is equal to 3%. We also investigate the averaged weekly turnover $\mathcal{ATO}$ ($\frac{1}{T_{\text{test}}} \sum_{t=1}^{T_{\text{test}}} \mathcal{TO}_t$)[6] which evaluates the magnitude of asset rotation in the portfolios. High rotation is naturally associated with higher trading costs. Moreover, from a ranking perspective, we assess metrics like Information Coefficient (IC) and the Kendall Tau ($K\tau$).

4.5.1 Benchmark strategies

To validate the efficiency of the multi-task model, we compare its performance against a baseline single-task model focused solely on asset ranking ($\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rank}}$). We employ two distinct benchmark approaches, both relying on the baseline model for asset ranking but differing in their methods for selecting the number of assets to long-short. This comparison is crucial to illustrate how integrating portfolio weighting tasks can improve financial metrics under different constraints.

The first benchmark approach employs a *naive strategy*, where a fixed number k of assets are longed and shorted each day. Since the optimal value of k is unknown, we evaluate performance across all $k \in \left\{1, \cdots, \frac{N}{2}\right\}$ and compare our model's results to the best-performing k. The second benchmark employs a *historical strategy*. It uses the same labeling methodology as the multi-task model for creating $\mathcal{Y}^{\text{long}}/\mathcal{Y}^{\text{short}}$, but estimates them based on historical data. This involves solving the problems (4.1), (4.2), and (4.3), with $-\tau$ instead of τ, indicating that the estimation is based on historical data up to and including the current time t.

4.5.2 Unconstrained portfolio optimization

In this section, we aim to show that adding two additional tasks to the model can improve the optimization of one or more financial metrics of interest compared to benchmark methods.

Maximizing expected returns μ_τ

Table 4.1 is obtained by creating labels that focus on maximizing a single criterion, specifically the expected return: $\mathcal{C} = [\mu_\tau]$. For the naive investment strategy, the optimal value of k maximizing the

[6] T_{test} designates the length of the whole test set.

expected return is $k = 1$. This is a natural outcome, as with an effective ranking system, the optimal strategy would be to go long on the top-ranked asset and short on the bottom-ranked asset. A key observation is our ability to achieve the highest value of expected return: 40.2% for the uniform density and 41.4% for the exponential density. Our strategy significantly outperforms the benchmarks for both distributions. Notably, the exponential distribution leads to a higher expected return, suggesting that the two additional tasks add value to the trading system by choosing more appropriate porfolio weights. Furthermore, for both distributions, we observe that our approach achieves a significantly lower average weekly turnover ($\mathcal{ATO} = 0.571$) compared to the benchmarks. This reduction can be attributed to the dynamic nature of the system, which judiciously opts to refrain from taking long or short positions in certain instances. Such strategic inaction contributes to a lower turnover, which leads to less transaction costs and amplifies the expected return.

	Uniform			Exponential		
	$k = 1$	Historical	Predictive	$k = 1$	Historical	Predictive
$\mu_{\text{ann}} - r_f$	0.202	0.246	**0.402**	0.202	0.269	**0.414**
Volatility (ann.)	**0.284**	0.342	0.398	**0.284**	0.345	0.399
Sharpe	0.711	0.721	**1.010**	0.711	0.782	**1.038**
Sortino/$\sqrt{2}$	0.741	0.741	**1.092**	0.741	0.817	**1.125**
MDD	**0.360**	0.387	0.372	**0.360**	0.382	0.372
% +ve Returns	0.583	0.599	**0.620**	0.583	0.599	**0.620**
Avg. P / Avg. L	**0.981**	0.913	0.932	**0.981**	0.937	0.941
Avg. Weekly TO	0.965	0.825	**0.571**	0.965	0.804	**0.551**

Note: Bold font indicates the best metric.

Table 4.1: Financial metrics for $\mathcal{C} = [\mu_\tau]$

	Uniform		Exponential	
	Benchmarks	Predictive	Benchmarks	Predictive
IC	0.1239	**0.1242**	**0.1239**	0.1231
$K\tau$	0.0862	**0.0864**	**0.0862**	0.0857

Note: Bold font indicates the best metric.

Table 4.2: Ranking metrics for $\mathcal{C} = [\mu_\tau]$

Minimizing volatility σ_τ

Table 4.3 presents the financial performances when focusing on minimizing volatility $\mathcal{C} = [\sigma_\tau]$. Within the naive investment strategy, the optimal k minimizing the volatility is $k = 37 = \frac{N}{2}$. This result aligns with the principle that diversification typically reduces risk, as spreading investments across a larger number of assets tends to lower overall volatility. When employing the exponential distribution, our strategy achieves the lowest volatility compared to both benchmarks. With the uniform density, the predictive strategy significantly outperforms the historical benchmark, while the naive strategy slightly edges out

our approach by 0.1%. However, the rest of the metrics are significantly different, indicating that the model learns to dynamically select the number of assets to long-short rather than adhering to a constant $k = 37$. Synthesizing the results from both tables, we conclude that the system is adept at maximizing the expected return μ and minimizing the volatility σ.

	Uniform			Exponential		
	$k = 37$	Historical	Predictive	$k = 37$	Historical	Predictive
$\mu_{\text{ann}} - r_f$	0.017	**0.029**	0.028	0.056	**0.076**	0.053
Volatility (ann.)	**0.048**	0.062	0.049	0.076	0.088	**0.075**
Sharpe	0.354	0.467	**0.571**	0.737	**0.864**	0.705
Sortino/$\sqrt{2}$	0.431	0.443	**0.672**	0.927	**1.054**	0.863
MDD	0.112	0.142	**0.107**	**0.145**	0.152	0.149
% +ve Returns	0.508	0.515	**0.516**	0.508	0.500	**0.524**
Avg. P / Avg. L	1.414	1.382	**1.475**	1.481	**1.588**	1.372
Avg. Weekly TO	**0.250**	0.357	0.263	**0.419**	0.474	0.435

Note: Bold font indicates the best metric.

Table 4.3: Financial metrics for $\mathcal{C} = [\sigma_\tau]$

	Uniform		Exponential	
	Benchmarks	Predictive	Benchmarks	Predictive
IC	**0.1239**	0.1219	0.1239	**0.1246**
$K\tau$	0.0862	**0.0849**	0.0862	**0.0868**

Note: Bold font indicates the best metric.

Table 4.4: Ranking metrics for $\mathcal{C} = [\sigma_\tau]$

Maximizing expected returns μ_τ while minimizing volatility σ_τ

Now we assume that the chosen portfolio criteria to optimize are both the expected return and volatility: $\mathcal{C} = [\mu_\tau, \sigma_\tau]$. Let w_1 be the weight assigned to μ_τ. The vector of weights assigned to each criterion can be written $w = (w_1, 1 - w_1)$. We let w_1 vary from 0 to 1 in increments of 0.1. For each value, we generate the corresponding labels $\mathcal{Y}^{\text{long}}(w_1)$ and $\mathcal{Y}^{\text{short}}(w_1)$. Subsequently, we compute the out-of-sample cumulative return and volatility yielded by the model. The results are depicted in Figures 4.4 and 4.5:

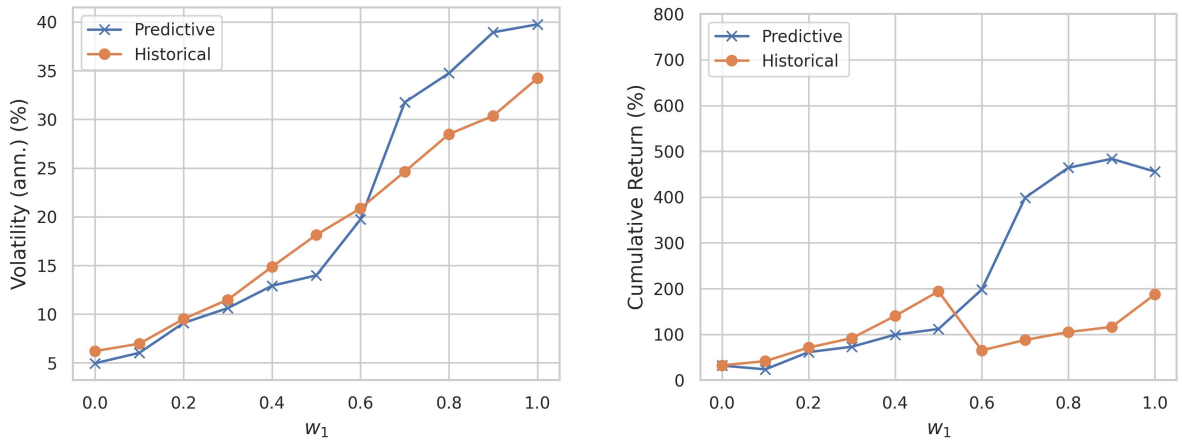

Figure 4.4: Realized volatility and cumulative return for different values of w_1, using a uniform density

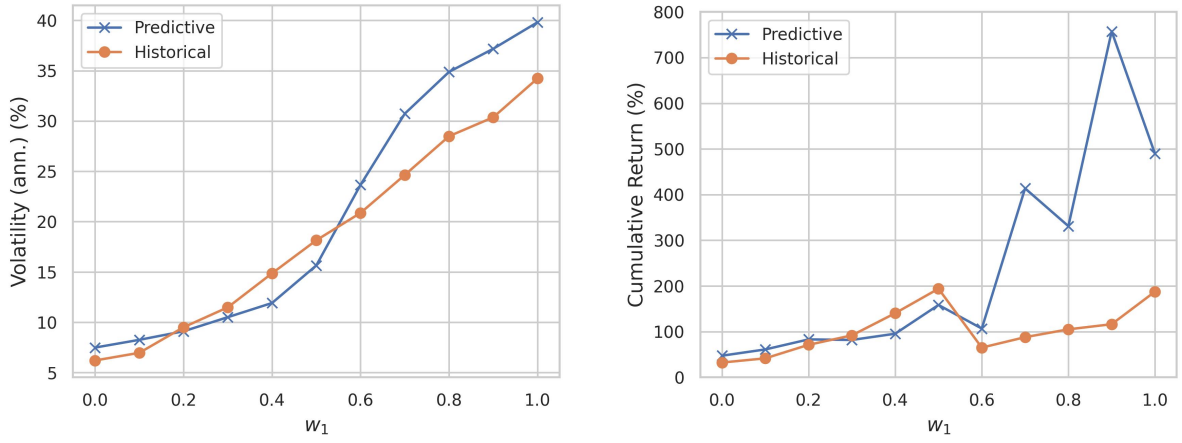

Figure 4.5: Realized volatility and cumulative return for different values of w_1, using an exponential density

The analysis of both figures reveals a clear trend: as we increase the emphasis on minimizing volatility by assigning more weight to σ_τ (through smaller values of w_1), we successfully achieve lower volatility levels. However, this focus on volatility comes at the expense of cumulative returns, which tend to be lower as they are not prioritized in the label construction. Conversely, when we shift our priority towards maximizing the expected return by increasing w_1, we observe a corresponding rise in cumulative returns. This shift leads to an increase in volatility, which is a consequence of reducing the focus on σ_τ within the labeling process. This trade-off between return and volatility is a fundamental aspect of portfolio optimization, reflecting the balance between risk and reward.

When w_1 is set to either high or low extremes, the predictive approach consistently surpasses those of the historical approach. This indicates that the inclusion of the two additional tasks in the learning process effectively aids in reaching the desired performance metrics. However, for intermediate values of w_1, we observe that the historical approach marginally outperforms the predictive model. This suggests that the model faces uncertainty in balancing its focus between returns and volatility, leading to less

optimal outcomes when neither metric is distinctly prioritized.

4.5.3 Constrained portfolio optimization

We now focus on our ability to target specific portfolios that maximize expected return subject to a volatility constraint.

	Uniform			Exponential		
	$k = 5$	Historical	Predictive	$k = 5$	Historical	Predictive
$\mu_{\text{ann}} - r_f$	0.129	0.110	**0.151**	0.124	0.096	**0.135**
Volatility (ann.)	0.127	0.145	**0.120**	0.130	0.145	**0.123**
Sharpe	1.016	0.758	**1.258**	0.954	0.662	**1.098**
Sortino/$\sqrt{2}$	1.241	0.791	**1.482**	1.151	0.685	**1.293**
MDD	0.339	0.195	**0.228**	0.208	0.341	**0.288**
% +ve Returns	0.555	0.591	**0.591**	0.567	0.575	0.563
Avg. P / Avg. L	1.272	0.992	**1.321**	1.184	1.023	**1.251**
Avg. Weekly TO	**0.587**	0.872	0.689	**0.614**	0.821	0.692

Note: Bold font indicates the best metric.

Table 4.5: Financial metrics for constrained problem (4.4) with $\sigma_{\text{tgt}} = 0.15$

	Uniform		Exponential	
	Benchmarks	Predictive	Benchmarks	Predictive
IC	**0.1239**	0.1226	**0.1239**	0.1215
$K\tau$	**0.0862**	0.0849	**0.0862**	0.084

Note: Bold font indicates the best metric.

Table 4.6: Ranking metrics for constrained problem (4.4) with $\sigma_{\text{tgt}} = 0.15$

Table 4.5 shows the summary statistics of both our approach and the benchmark for the constrained problem (4.4), when setting $\sigma_{\text{tgt}} = 0.15$. Under the naive investment strategy, we identify that selecting $k = 5$ optimizes the expected return while adhering to the volatility constraint. For both the uniform and exponential distributions, the predictive approach clearly outperforms the benchmarks. This superior performance is evidenced by not only a higher return but also a lower volatility, all while adhering to the constraint $\sigma_\tau \leq 0.15$.

To confirm the robustness of our approach, for each σ_{tgt} value varying between 8% and 24% with increments of 1%, we generate the corresponding labels $\mathcal{Y}^{\text{long}}(\sigma_{\text{tgt}})$ and $\mathcal{Y}^{\text{short}}(\sigma_{\text{tgt}})$. Then, we compute the cumulative return, and volatility obtained by the multi-task model. The results are displayed in Figures 4.6 and 4.7:

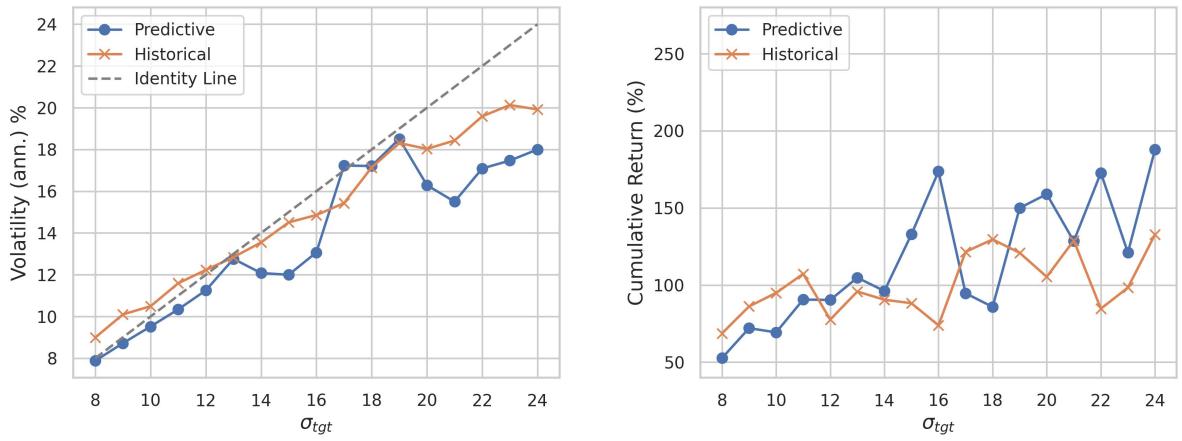

Figure 4.6: Realized volatility and cumulative return for different values of σ_{tgt}, using a uniform density

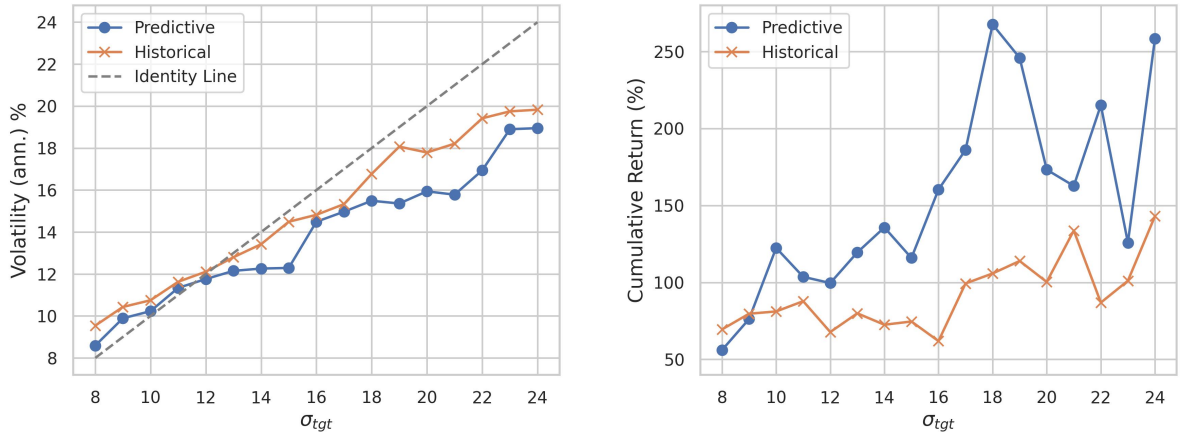

Figure 4.7: Realized volatility and cumulative return for different values of σ_{tgt}, using an exponential density

Figure 4.6 illustrates that under the uniform density, the predictive approach manages to perfectly meet the volatility constraint. Besides, when examining the cumulative returns, our strategy demonstrates superior performances on average. However, for certain values of σ_{tgt}, particularly the lower ones, the system performs less effectively.

In contrast, under the exponential density (Figure 4.7), the multi-task model is notably more robust and reliable. It adheres closely to the volatility constraint, consistently achieving smaller values than the target. Notably, not only does it outperform the historical strategy in terms of return, but it also achieves that with a reduced volatility, demonstrating superior stability and efficiency in its results. For instance, for $\sigma_{\text{tgt}} = 0.18$, the predictive approach achieves a cumulative return that is 160% higher than the historical approach, while simultaneously reducing volatility by 1.3%.

According to Tables 4.2, 4.4, 4.6, the ranking metrics realized by the predictive model compared to the one-task based model have the same order of magnitude. This supports the idea that with the

current configuration, the two additional tasks are not degrading the ranking system, while also bringing additional value through a better portfolio weighting.

4.5.4 Balancing loss components

Finally, in an effort to outline the tradeoff between optimizing asset ranking and portfolio weighting, we alter the total loss expression by incorporating a new hyperparameter c. The revised loss function is expressed as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rank}} + c(\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}})$$

c represents the relative weight assigned to the losses $\mathcal{L}_{\text{long}}$ and $\mathcal{L}_{\text{short}}$ with respect to $\mathcal{L}_{\text{rank}}$. Specifically, c controls how much emphasis is placed on minimizing the losses associated with predicting the optimal number of assets to long and short relative to the ranking task. It is important to note that the impact of c on the overall loss depends on the specific ranking loss function used, as different ranking losses (e.g., ListMLE) have varying scales and magnitudes. This means that the absolute value of c should not be interpreted in isolation but rather in the context of the magnitude of $\mathcal{L}_{\text{rank}}$. To provide a more meaningful interpretation of c, we prefer to frame it in terms of the percentage contribution of the term $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$ to the total loss $\mathcal{L}_{\text{total}}$. For a given value of c, we posit that the combined loss $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$ contributes approximately $x\%$ of the total loss. The choice of c can be tuned such that this contribution aligns with the desired balance between achieving better ranking accuracy and optimizing the predictions for the number of assets to long and short. The methodology for computing c to achieve a specific contribution percentage is detailed in Section 4.7.1.

Using the labels $\mathcal{Y}^{\text{long}}(\sigma_{\text{tgt}} = 0.15)$ and $\mathcal{Y}^{\text{short}}(\sigma_{\text{tgt}} = 0.15)$, we train 200 models, each with a different value of c, resulting in a distinct emphasis of $x\%$ on $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$. The relative contribution of $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$ over $\mathcal{L}_{\text{total}}$ is spaced uniformly every 0.5%, ranging from 0% (where $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rank}}$) to 100% (where $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$). We then calculate the average training losses $\mathcal{L}_{\text{rank}}$, $\mathcal{L}_{\text{long}}$, and $\mathcal{L}_{\text{short}}$ across the entire data sample. To visualize the impact of assigning more weight to the portfolio weighting tasks by increasing c, we created two scatter plots showing $\mathcal{L}_{\text{rank}}$ against $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$, with a color gradient representing the realized Sharpe ratio, and the percentage contribution associated to c (see Figures 4.8 and 4.9). The plot reveals a clear trade-off between the ranking task and the portfolio weighting tasks. Interestingly, intermediate values of c are associated with higher Sharpe ratios, indicating an optimal balance between both tasks. On the other hand, extreme values of c (either too low or too high) tend to correspond to lower Sharpe ratios. This suggests that a slight reduction in ranking accuracy can be offset by improvements in portfolio weighting efficiency.

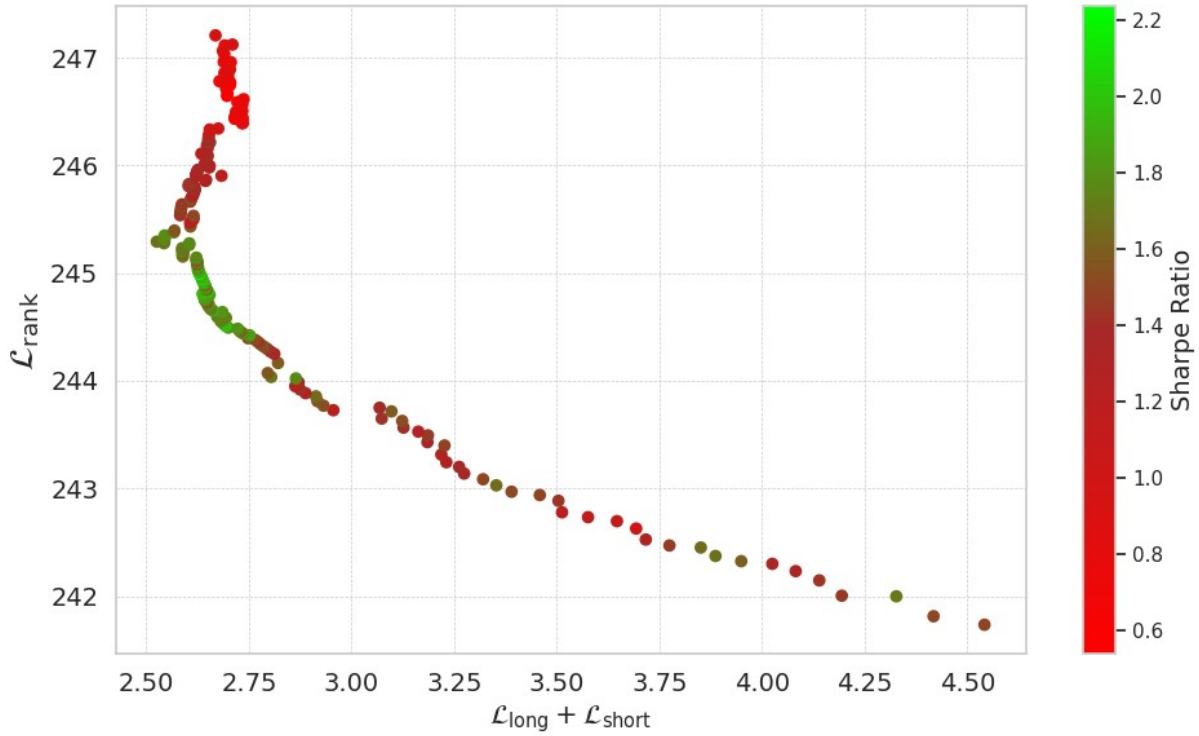

Figure 4.8: Scatter plot displaying $\mathcal{L}_{\text{rank}}$ vs $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$ and the realized Sharpe ratio for 200 distinct models trained using different values of c

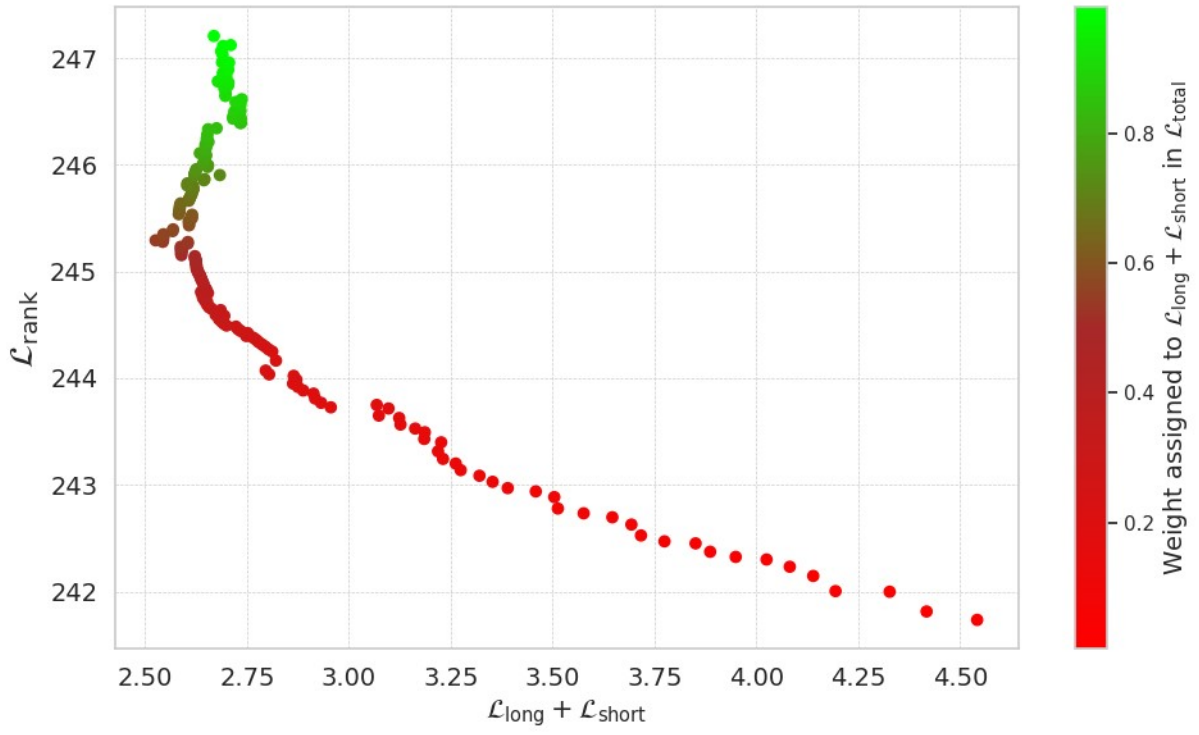

Figure 4.9: Scatter plot displaying $\mathcal{L}_{\text{rank}}$ vs $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$ and the contribution for 200 distinct models trained using different values of c

4.6 Conclusion

In conclusion, by employing LTR algorithms in conjunction with a multi-task neural network model, this chapter presents a novel approach to cross-sectional systematic strategies. We introduce a method that not only predicts the ranking of assets, but also dynamically determines the optimal number of assets to long and short. This innovative strategy represents a significant departure from traditional methods that first predicts asset ranks and then performs portfolio weighting, as well as from newer methods that exclusively predicts optimal portfolio weights.

Our findings indicate that this integrated approach leads to improved financial metrics considering unconstrained and constrained portfolio optimization scenarios. Indeed, our trading system outperforms benchmark approaches that either rely on a fixed number of assets to long-short, or on historical portfolio weights estimation. Furthermore, the system's ability to learn and adjust to portfolio constraints ensures the generation of efficient portfolios that are better aligned with investor preferences. Finally, we underline the existing trade-off between asset ranking and portfolio weighting, suggesting that particular attention must be paid to the relative importance of both tasks in the model training.

For future study, it is worthwhile to further investigate the consistency of the approach across different types of data, including daily records and other asset classes. Additionally, extending the framework to incorporate diverse rankers for computing the ranking labels could provide valuable insights. Lastly, exploring the use of three or more criteria in selecting $\mathcal{C}$ when generating the TOPSIS labels would also enhance the robustness of the framework.

4.7 Appendix

4.7.1 Methodology for computing c

Consider the revised total loss function:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rank}} + c(\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}})$$

where c controls the weight assigned to the portfolio weighting loss functions ($\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$) relative to the total loss ($\mathcal{L}_{\text{total}}$). In order to compute the value of c corresponding to a relative contribution of $x\%$ on $\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$ over $\mathcal{L}_{\text{total}}$, we first calculate the magnitude of $\mathcal{L}_{\text{rank}}$, using a baseline predictor f_{base} that assigns the same score to each asset.

The ListMLE loss function is defined as:

$$L_{\text{rank}}^{\text{ListMLE}}(f, \mathcal{X}, \mathcal{Y}) = -\log \prod_{i=1}^{n} \frac{\phi(f(x_{y^{-1}(i)}))}{\sum_{k=i}^{n} \phi(f(x_{y^{-1}(k)}))} = -\log P(y|x; f)$$

Recalling that:

$$P(y|x;f) = P(y^{-1}(1)|x;f)\prod_{i=2}^{n} P(y^{-1}(i)|x;y^{-1}(1),\ldots,y^{-1}(i-1);f)$$

and noting that the probability of randomly ranking the i^{th} asset knowing that the first $i-1$ assets have been ranked is $\frac{1}{n-i+1}$, we obtain:

$$L_{\text{rank}}^{\text{ListMLE}}(f_{\text{base}},\mathcal{X},\mathcal{Y}) = -\log P(y|x;f_{\text{base}}) = -\log(\frac{1}{n!}) = \log(n!) \approx n\log(n) - n$$

Next, we recall the expressions for the portfolio weighting loss functions $\mathcal{L}_{\text{long}}$ and $\mathcal{L}_{\text{short}}$:

$$\mathcal{L}_t^{\text{long}}(f,\mathcal{X},\mathcal{Y}_t^{\text{long}}) = -\sum_{i=0}^{k_{\max}} y_{i,t}^{\text{long}}\log(p_{i,t}^{\text{long}})$$
$$\mathcal{L}_t^{\text{short}}(f,\mathcal{X},\mathcal{Y}_t^{\text{short}}) = -\sum_{i=0}^{k_{\max}} y_{i,t}^{\text{short}}\log(p_{i,t}^{\text{short}})$$

For a random predictor, which assigns equal probabilities over the $k_{\max}+1$ categories, we have:

$$\mathcal{L}_t^{\text{long}}(f_{\text{base}},\mathcal{X},\mathcal{Y}_t^{\text{long}}) = \mathcal{L}_t^{\text{short}}(f_{\text{base}},\mathcal{X},\mathcal{Y}_t^{\text{short}}) = \log(k_{\max}+1)$$

Given that $n = 74$ and $k_{\max} = 37$, the magnitude of the losses for the baseline predictor can be summarized in the following table:

$L_{\text{rank}}^{\text{ListMLE}}$	$\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}$
244.2	7.28

This table provides the basis for determining the value of c such that the portfolio weighting losses account for $x\%$ of the total loss $\mathcal{L}_{\text{total}}$. Finally, we compute the value of c such that:

$$\frac{c(\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}})}{\mathcal{L}_{\text{total}}} = x\%$$

that is,

$$c = \frac{x\% \times \mathcal{L}_{\text{total}}}{\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}}} = \frac{x \times \mathcal{L}_{\text{rank}}}{(1-x)(\mathcal{L}_{\text{long}} + \mathcal{L}_{\text{short}})}$$

Chapter 5

Understanding the Equity-Corporate Bond Nexus: A Framework for Risk Decomposition and Interest Rate Hedging[1]

Abstract

This chapter explores the complex dynamics between equity and corporate bond markets, focusing on their correlation's decomposition into credit risk and interest rate sensitivity components. By utilizing equity-credit default swap and equity-sovereign bond correlations, we identify macroeconomic factors—such as growth and inflation uncertainties—that influence these interactions across different credit qualities and market regimes. Empirical findings reveal significant variations in the equity-corporate bond correlation during periods of market stress and rising interest rates. Additionally, we propose two novel dynamic hedging strategies, leveraging adaptive hedge ratios to isolate credit risk while mitigating interest rate exposure. The results demonstrate that these strategies outperform traditional approaches, offering robust tools for portfolio management amidst fluctuating economic conditions.

5.1 Introduction

The relationship between equity and fixed income instruments has long been a central topic of financial research, driven by the need to understand their diversification potential and shared exposure to economic factors. Among fixed income assets, the correlation between equities and sovereign bonds has attracted

[1]This chapter is based on the research article *Understanding the Equity-Corporate Bond Nexus: A Framework for Risk Decomposition and Interest Rate Hedging* co-authored with Jamyang-Dorje Bhutia, Axel Pinçon and Marcel Voia and published in the *Journal of Fixed Income*, DOI: 10.3905/jfi.2025.1.210.

significant attention due to its implications for portfolio risk management. In recent years, notable shifts in this correlation have sparked renewed interest and debate within the financial community. Historically, equities have served as the primary return driver in portfolios, while sovereign bonds have acted as a key diversifier. From 1900 to 2000, equities and sovereign bonds generally moved in the same direction, offering some diversification benefits but often responding similarly to market conditions. In the early 2000s, however, this relationship shifted, with sovereign bonds frequently moving in the opposite direction of equities and providing a more reliable counterbalance ([18]). In recent years, the relationship between equities and sovereign bonds has shifted significantly, particularly in light of the interest rate hikes that began in March 2022. As shown in Figure 5.1, this period has seen both asset classes increasingly moving in tandem, leading to a rise in positive correlation between them. As a result, traditional assumptions about the diversification benefits of equity-sovereign bond portfolios are being challenged, prompting a reevaluation of how these assets interact and the implications for portfolio management strategies. Recent findings indicate that this evolving correlation is influenced by various macroeconomic factors and market regimes. [13] underscores that equity-sovereign bond correlation is primarily driven by growth and inflation uncertainties; equities and sovereign bonds tend to act as stronger diversifiers when growth news is more prevalent, whereas their effectiveness diminishes when inflation concerns dominate. Furthermore, [82] reveals that sovereign credit risk often correlates with positive equity-sovereign bond movements, highlighting the complexity of correlation behavior under varying market conditions. Moreover, findings indicate that in typical market conditions, the correlation may approach zero, whereas extreme market regimes tend to heighten these correlations. This changing relationship emphasizes the importance of understanding the dynamics driving asset correlations, as they directly impact portfolio construction and risk management strategies.

Besides sovereign bonds, several studies have found a significant negative relationship between firm-level equity and credit default swap (CDS) prices ([71], [72], [43]). At the firm level, an expanding credit spread typically corresponds with a declining equity price, as lower equity prices suggest reduced potential future earnings, thereby increasing the probability of default. Building on these insights, additional research highlights a relatively stable yet nuanced correlation between equity returns and CDS spreads, which varies with market stress, firm-specific factors, and economic cycles. While equity-CDS correlation is generally more stable over time compared to equity-sovereign bond correlation (see Figure 5.1), it tends to intensify during market stress. For instance, in the 2008 financial crisis, CDS spreads widened significantly while equity returns dropped, amplifying the negative correlation between them. Besides, [65] demonstrate that factors such as firm leverage and cash flow news play a significant role in shaping equity-CDS correlation, with speculative-grade firms showing a stronger relationship than investment-grade firms due to their heightened sensitivity to credit risk. Additionally, [65] find that equity-CDS correlation follows cyclical patterns, often intensifying during periods of economic distress. Using regime-switching models, their study reveals that correlations strengthen in volatile market regimes, such as crises, when cash flow uncertainty is at its highest.

Although numerous studies have explored the correlations between equity-sovereign bond and equity-CDS, there seems to be a limited emphasis in the existing literature on the specific relationships between equity returns and corporate bond returns. Given the role corporate bonds play as intermediaries between sovereign debt and CDS instruments, this unexplored correlation could hold significant implications for portfolio allocation, hedging, and risk assessment. Figure 5.1 highlights this point, by displaying the

correlation between the S&P 500 and various corporate bond indices. Specifically, the correlation with the ICE BofA 5-Year AAA corporate bond index is more closely aligned with that of US Treasuries 5-Year rates, while the correlation with the ICE BofA 5-Year CCC shows greater similarity to that of CDX 5-Year High Yield Index. Besides, as shown in Figure 5.1, during market stress events such as the 2008 financial crisis and the COVID pandemic, these correlations tend to diverge significantly. These observations underscore the complexity of the relationship between equity and corporate bonds, suggesting that the risk profiles of different corporate bonds, along with the influence of market stress, significantly affect their interaction with equity returns.

Building on this foundation, this chapter addresses two interconnected questions: (1) What drives the correlation between equities and corporate bonds, and how can this be decomposed into credit and interest rate components? Furthermore, through which channels—credit or interest rate—do macroeconomic drivers influence this correlation? (2) How can interest rate risk be effectively neutralized in corporate bonds to isolate their credit risk, and how does this required level of hedging evolve across credit qualities and market conditions? In the first part of the study, we propose a framework that decomposes the equity-corporate bond correlation into two distinct components: the equity-CDS correlation, which captures credit risk, and the equity-sovereign bond correlation, which reflects interest rate sensitivity. This framework not only enables us to isolate the relative contributions of credit and interest rate factors but also allows us to investigate the underlying macroeconomic drivers that influence these components. By examining these dynamics, we aim to identify periods when either credit or interest rate factors dominate. Additionally, we seek to determine the channels through which macroeconomic conditions shape the equity-corporate bond relationship across time and different levels of credit quality. In the second part, we extend this analysis to explore practical strategies for optimizing interest rate hedging for corporate bonds. This focus on isolating credit risk aims to create a stable correlation between equities and hedged corporate bonds, similar to the equity-CDS correlation. We introduce the concept of a hedge ratio (HR) to measure the level of interest rate exposure that should be hedged, and we examine how it evolves over time and across different credit ratings. This approach is particularly useful for asset managers seeking to invest solely in credit risk, as it allows them to isolate and better manage this specific exposure without the confounding effects of interest rate risk. By integrating these two perspectives, this study offers a comprehensive framework for understanding and managing the nuanced interplay between equities and corporate bonds. More precisely, this chapter contributes to the existing literature in several ways:

1. **Decomposing equity-corporate bond correlation**: A detailed framework is introduced to decompose equity-corporate bond correlation into credit risk and interest rate risk.
2. **Impact of macroeconomic factors**: Our study reveals how macroeconomic factors such as inflation and growth uncertainties influence this correlation through credit and interest rate channels.
3. **Adaptive hedging strategies**: We propose two adaptive interest rate hedging strategies for corporate bond portfolios that outperform standard approaches.

The remainder of this research is organized as follows. Section 5.2 outlines the dataset. Section 5.3 examines the equity-corporate bond correlation by decomposing it into credit and interest rate components, exploring macroeconomic drivers, and presenting empirical findings. Section 5.4 focuses on improving

classical interest rate hedging methodologies for corporate bonds, detailing theoretical considerations and empirical evaluations. Finally, Section 5.5 concludes the study by summarizing the key findings.

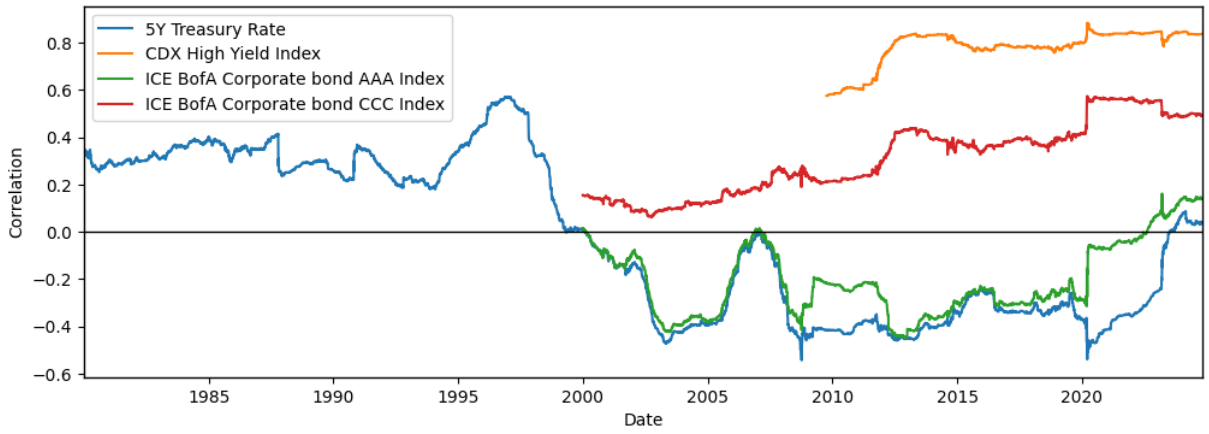

Notes: Rolling 3-year series using overlapping daily returns: US equities are the S&P 500; US Treasuries are 5-Year rates; US Corporate bonds are represented by ICE BofA 5-Year AAA and CCC & Lower High Yield indices. US CDS uses the CDX 5-Year High Yield Index. CDS returns assume a short position; equity, corporate bond, and sovereign bond returns assume a long position.

Figure 5.1: Rolling Correlation between US Equity and US Treasury rates, US Corporate bond and US CDS Returns, January 1, 1980–October 25, 2024

5.2 Dataset overview

This study focuses on 534 US companies, including a sample of 8,394 corporate bonds, selected from the Bloomberg US Corporate and Bloomberg US Corporate High Yield indices, representing a broad range of credit qualities. We leverage three primary data sources—Bloomberg, Markit, and Financial Modeling Prep—covering the sample period from 2013 to 2024. Equity data is sourced from Bloomberg, while US sovereign bond, corporate bond, and CDS data are obtained from Markit. Since bond data availability begins in 2013, our analysis is naturally constrained to this timeframe. Furthermore, to complement our analysis, we incorporate ICE BofA indices sourced from Financial Modeling Prep.

In addition, we dynamically classify firms as investment-grade (IG) or high-yield (HY) over time. This classification is based on comparing each firm's corporate bond spread to the ICE BofA BBB Index. Firms with spreads above the index are classified as HY, while those with spreads below the index are classified as IG.

Finally, to enable a meaningful comparison of profit-and-loss (P&L) across corporate bonds, CDS, and sovereign bonds, we build composite CDS and sovereign bond P&L that are aligned with the time to maturity of the corresponding corporate bonds. This approach ensures that all instruments in our analysis are matched by maturity, enabling a consistent evaluation of performance across asset classes while controlling for differences in their durations.

5.3 Understanding equity-corporate bond correlation

In this section, we introduce a methodology that decomposes the P&L of a corporate bond into two key components: sensitivity to interest rate movements, as captured by fluctuations in US sovereign bond prices, and exposure to credit risk, represented by changes in the issuing firm's CDS prices. This decomposition enables us to evaluate the relative contributions of interest rate and credit risk fluctuations to equity-corporate bond co-movements. In addition, we examine the channels through which macroeconomic factors impact the equity-corporate bond relationship, either via interest rate or credit risk.

5.3.1 Decomposing the equity-corporate bond correlation: interest rate versus credit components

The yield of a corporate bond can be decomposed into two primary components. The first is the interest rate, typically represented by the yield of a sovereign bond with a similar maturity, reflecting the time value of money. The second component is the credit spread, which is the additional risk premia that investors require for taking on the credit risk associated with the corporate bond. This spread compensates investors for the risk of default and is often measured by the spread on the issuer's CDS, which provides a market-based indication of the issuer's creditworthiness. This yield decomposition naturally leads to our P&L formula, where we represent corporate bond P&L as a combination of sovereign bond and CDS P&Ls, capturing the effects of both interest rates and credit risk:

$$\Delta P^B = \Delta P^R - \Delta P^C \tag{5.1}$$

where ΔP^i denotes the P&L of asset i. B, R and C represent corporate bond and comparable-maturity sovereign bond and CDS[2]. This expression provides a practical approximation of ΔP^B, capturing the two primary drivers of corporate bond performance:

- Interest rate sensitivity: the corporate bond's P&L in a long position exhibits similar sensitivity to interest rate movements as a sovereign bond of comparable maturity. When interest rates increase, both sovereign bond and corporate bond values typically decrease, resulting in negative P&L, while a decrease in interest rates generates positive P&L.
- Credit risk sensitivity: the corporate bond's P&L is inversely related to that of a long CDS position with the same credit exposure. As credit risk rises, the corporate bond's value falls, yielding a negative P&L, whereas the CDS, which protects against default, gains in value, generating positive P&L.

Following Equation 5.1, we can write the equity-corporate bond correlation $\rho_{E,B}$ as:

$$\rho_{E,B} = \frac{\sigma_{E,B}}{\sigma_E \cdot \sigma_B} = \frac{\sigma_{E,R-C}}{\sigma_E \cdot \sigma_{R-C}} = \frac{\sigma_{E,R} - \sigma_{E,C}}{\sigma_E \cdot \sqrt{\sigma_R^2 + \sigma_C^2 - 2 \cdot \sigma_{R,C}}}$$

[2]We assume all positions are long and cover the same notional amount.

$$= \underbrace{\left(\frac{\sigma_R}{\sqrt{\sigma_R^2 + \sigma_C^2 - 2 \cdot \sigma_{R,C}}} \cdot \rho_{E,R} \right)}_{A_1} - \underbrace{\left(\frac{\sigma_C}{\sqrt{\sigma_R^2 + \sigma_C^2 - 2 \cdot \sigma_{R,C}}} \cdot \rho_{E,C} \right)}_{A_2}, \tag{5.2}$$

where $\rho_{i,j}$ (resp. $\sigma_{i,j}$) represents the correlation (resp. the covariance) between P&L of asset i and P&L of asset j. σ_i represents the standard deviation of asset i's P&L.

With Equation 5.2, we arrive at a decomposition of the equity-corporate bond correlation $\rho_{E,B}$ into two distinct terms. The first term A_1 represents the correlation between equity and interest rates $\rho_{E,R}$, weighted by the relative influence of sovereign bond volatility to corporate bond volatility. The second term A_2 represents the correlation between equity and credit risk $\rho_{E,C}$, weighted by the relative influence of CDS volatility to corporate bond volatility. By isolating these components, we can assess which of these two factors—interest rate movements or credit risk fluctuations—dominates in determining the equity-corporate bond co-movements.

5.3.2 Macroeconomic drivers of the equity-corporate bond correlation

To extend the analysis that identifies whether $\rho_{E,R}$ or $\rho_{E,C}$ primarily drives $\rho_{E,B}$, we propose to examine the underlying macroeconomic determinants of each component. To model $\rho_{E,R}$, we employ a Growth-Inflation framework as proposed by [13]. This approach assumes that $\rho_{E,R}$ is driven primarily by three factors: uncertainty in inflation σ_π^2, uncertainty in growth σ_g^2, and the correlation between growth and inflation $\rho_{g,\pi}$.

According to [13] and [82], equity prices generally respond positively to growth shocks as they boost investor expectations, but negatively to inflation shocks, which raise credit risk. Sovereign bond prices, on the other hand, tend to decline in response to both growth and inflation shocks. Growth shocks lead to anticipated increases in short-term interest rates, while inflation shocks directly raise interest rates, both of which cause sovereign bond prices to fall. Finally, recasting their approach in a simplified form, we are able to pose the following regression model:

$$\rho_{E,R} = \beta_0 + \beta_1 \cdot \sigma_g^2 + \beta_2 \cdot \sigma_\pi^2 + \beta_3 \cdot \rho_{g,\pi} + \epsilon \tag{5.3}$$

To analyze $\rho_{E,C}$, we apply a similar framework, acknowledging the significant role that macroeconomic factors play in shaping this relationship. Economic conditions such as inflation and growth influence $\rho_{E,C}$ by impacting firms' financial health and the market's perception of credit risk. For example, during economic downturns, firms may face heightened credit risk, leading to wider CDS spreads and falling equity prices, thus strengthening the negative correlation between equity and CDS, as noted by ([65]). To model this relationship, we adopt the same Growth-Inflation framework:

$$\rho_{E,C} = \beta_0 + \beta_1 \cdot \sigma_g^2 + \beta_2 \cdot \sigma_\pi^2 + \beta_3 \cdot \rho_{g,\pi} + \epsilon \tag{5.4}$$

To summarize our approach, we first express $\rho_{E,B}$ in Equation 5.2 as a function of $\rho_{E,R}$ and $\rho_{E,C}$. Next, we model $\rho_{E,R}$ (Equation 5.3) and $\rho_{E,C}$ (Equation 5.4) as functions of σ_g^2, σ_π^2 and $\rho_{g,\pi}$. This framework enables us to establish a connection between these macroeconomic indicators and $\rho_{E,B}$, either via the

interest rate channel (through $\rho_{E,R}$) or the credit risk channel (through $\rho_{E,C}$). The methodology used to estimate the contributions of explanatory variables in Equations 5.2, 5.3, and 5.4 is provided in Appendix 4.7.1.

5.3.3 Empirical results

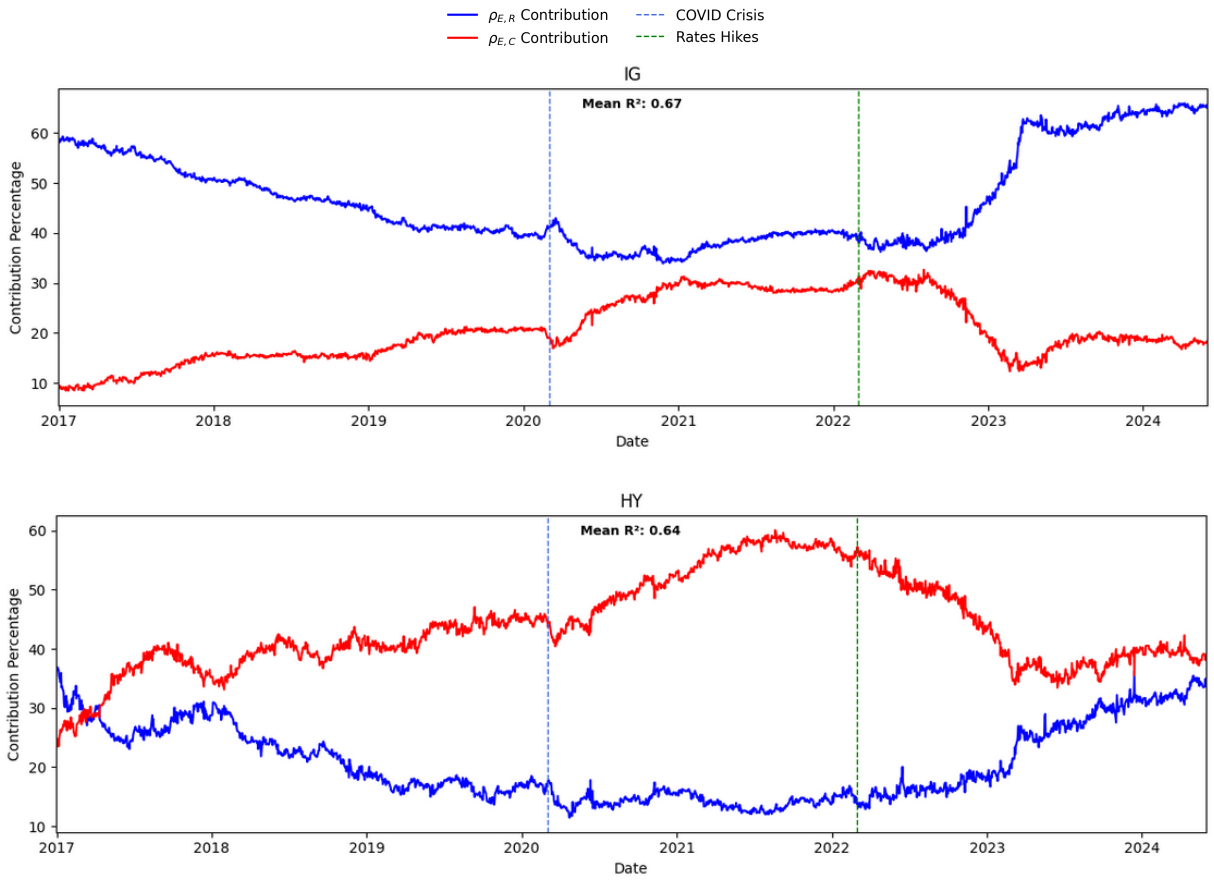

Figure 5.2: Contribution of $\rho_{E,R}$ and $\rho_{E,C}$ in explaining $\rho_{E,B}$ for IG and HY firms

Using the methodology detailed in Appendix 5.6.1, Figure 5.2 presents a comparative analysis of the relative contributions of $\rho_{E,R}$ and $\rho_{E,C}$ to explaining variations in $\rho_{E,B}$ for IG and HY firms. The first panel examines IG companies, revealing that $\rho_{E,R}$ consistently emerges as the dominant driver in explaining $\rho_{E,B}$. This result is intuitive, as IG bonds are inherently more sensitive to interest rate fluctuations due to their lower credit spreads and minimal default risk, making them behave similarly to sovereign bonds. However, it is essential to account for credit risk, as corporate bonds, even within the IG category, are not immune to default concerns. This vulnerability becomes particularly evident during periods of heightened market stress. A striking deviation occurs during the COVID crisis, where the contribution of $\rho_{E,C}$ increases sharply. This is consistent with the findings of [65], who show that the contribution of $\rho_{E,C}$ often intensifies during periods of economic distress. This shift reflects a temporary surge in credit risk concerns, as markets reassessed the resilience of IG firms amidst unprecedented economic disruptions. Post-crisis, the landscape shifts again. Following the interest rate hikes of 2022, the

influence of interest rate dynamics—captured by $\rho_{E,R}$—reasserts itself, not only recovering but surpassing pre-pandemic levels. This rebound underscores the renewed role of $\rho_{E,R}$ as a key driver of variations in $\rho_{E,B}$ for IG firms, reflecting the market's swift adaptation to shifting rate environments and reaffirming the sensitivity of IG bonds to interest rate dynamics.

The second panel reveals a distinct pattern for HY companies, where $\rho_{E,C}$ consistently plays a more significant role in explaining $\rho_{E,B}$ than for IG companies. Although the overall dynamics align with those observed for IG firms, the contribution of $\rho_{E,C}$ is notably stronger in HY markets. During the COVID crisis, this distinction becomes even more pronounced, as $\rho_{E,C}$ accounts for a larger share of the observed correlation, reflecting heightened default risk expectations and the widening CDS spreads for HY issuers. After the 2022 interest rate hikes, $\rho_{E,R}$ also gains importance in explaining $\rho_{E,B}$. However, its influence remains secondary to that of $\rho_{E,C}$, underscoring the continued dominance of credit risk in the HY market.

These findings emphasize that $\rho_{E,B}$ is governed by a dynamic interaction between interest rate and credit risk factors, with their relative contributions shifting over time in response to macroeconomic conditions. Specifically, we delineate three distinct regimes—pre-COVID, COVID, and post-COVID—each exhibiting unique characteristics in the behavior of $\rho_{E,B}$. While the underlying mechanisms driving $\rho_{E,B}$ are consistent across credit quality tiers, notable differences emerge in the relative significance of $\rho_{E,R}$ and $\rho_{E,C}$. In particular, IG bonds are predominantly influenced by interest rate dynamics, whereas HY bonds are more sensitive to credit risk. Despite these disparities, the structural determinants of $\rho_{E,B}$ exhibit a consistent framework across credit classes.

		pre-COVID			COVID			post-COVID		
		$\rho_{E,R}$	$\rho_{E,C}$	Total	$\rho_{E,R}$	$\rho_{E,C}$	Total	$\rho_{E,R}$	$\rho_{E,C}$	Total
IG	$\rho_{g,\pi}$	0.20	0.05	**0.25**	0.04	0.02	0.06	0.04	0.05	0.09
	σ_g^2	0.07	0.03	0.10	0.22	0.15	**0.37**	0.17	0.06	0.23
	σ_π^2	0.08	0.03	0.11	0.06	0.06	0.12	0.23	0.04	**0.27**
	ϵ	0.16	0.03	0.19	0.06	0.05	0.11	0.09	0.07	0.16
	Total	**0.51**	0.14	0.65	**0.38**	0.27	0.65	**0.53**	0.22	0.75
HY	$\rho_{g,\pi}$	0.10	0.12	**0.22**	0.02	0.06	0.08	0.02	0.10	0.12
	σ_g^2	0.03	0.06	0.09	0.07	0.26	**0.33**	0.08	0.12	**0.20**
	σ_π^2	0.05	0.07	0.12	0.02	0.10	0.12	0.10	0.10	**0.20**
	ϵ	0.07	0.12	0.19	0.03	0.11	0.14	0.04	0.10	0.14
	Total	0.25	**0.37**	0.62	0.14	**0.53**	0.67	0.24	**0.42**	0.66

Notes: Bold font indicates the best metric.
Periods: pre-COVID (2016-12-30 to 2019-03-01), COVID (2020-03-01 to 2022-03-17), and post-COVID (2022-03-17 to 2024-07-19).
Inflation and Growth refer to 3Y rolling volatility of CPI YoY and Industrial Production YoY, respectively. GIC represents the 3Y rolling correlation between CPI YoY and Industrial Production YoY.
ϵ denotes the residuals, capturing the unexplained contribution in the two-step rolling regressions.

Table 5.1: Mean of the absolute contributions of macroeconomic features to $\rho_{E,R}$ and $\rho_{E,C}$ for IG and HY firms

Table 5.1 quantifies the contributions of growth (σ_g^2) and inflation (σ_π^2) uncertainties, and growth-inflation correlation ($\rho_{g,\pi}$) to $\rho_{E,B}$ through the interest rate ($\rho_{E,R}$) and credit ($\rho_{E,C}$) channels for both IG and HY firms. The methodology employed to compute these contributions is detailed in Appendix 5.6.2. Notably, for IG bonds, the interest rate channel consistently dominates across all periods, while the credit risk channel primarily drives HY bonds. This is consistent with the findings of [56], which show that as credit ratings decline, credit spreads widen, amplifying the impact of credit risk and reducing sensitivity to interest rate movements. In the pre-COVID period, considered stable, the growth-inflation correlation explains most of $\rho_{E,B}$, both for IG (0.25) and HY (0.22) companies. This aligns with the findings of [13], who similarly concludes that the correlation between growth and inflation serves as a primary driver of $\rho_{E,R}$ during periods of stability.

During the COVID crisis, growth uncertainty became the dominant driver of $\rho_{E,B}$. For IG bonds, this uncertainty played a predominant role in shaping the dynamics of the $\rho_{E,B}$, with contributions transmitted through both the interest rate channel (0.22) and the credit risk channel (0.15). In contrast, for HY bonds, the transmission occurred solely through the credit risk channel (0.26). This shift reflects heightened fears of an economic recession, exacerbated by central banks' aggressive rate cuts to support the economy, while concerns about credit risk surged.

Post-COVID, inflation uncertainty began to take center stage, becoming the dominant driver of $\rho_{E,B}$. This shift was amplified by inflationary pressures arising from supply-chain disruptions and the economic

reopening, prompting central banks to initiate significant interest rate hikes starting in 2022, which continued into 2023. These interest rate hikes further fueled inflation uncertainty, which strongly influenced $\rho_{E,B}$ during this period (0.27 for IG and 0.20 for HY companies). However, right after the interest rate hikes, growth uncertainty still played a pivotal role in explaining these correlations. This was evident in the ongoing impact on both IG (0.23) and HY (0.20) bonds, suggesting that concerns about economic growth persisted even as inflationary factors took hold. It was not until early 2023 that inflation uncertainty overtook growth uncertainty as the primary driver of $\rho_{E,B}$, marking a shift in the dominant macroeconomic factor influencing market dynamics.

These observations underscore the pivotal role of macroeconomic variables in shaping $\rho_{E,B}$, primarily through two distinct channels—the interest rate channel and the credit risk channel. Their relative influence varies depending on credit quality and evolves in response to broader global economic conditions. Given these dynamics, a natural question arises: How can interest rate risk be effectively neutralized in corporate bonds to isolate their credit risk, and how does the required level of hedging evolve across credit qualities and market conditions? This question is particularly relevant for capital structure arbitrage strategies, where accurately isolating credit risk is essential for exploiting mispricings between a firm's equity and debt instruments. The following sections aim to explore this issue in greater detail, providing insights into the mechanisms and tools required to achieve effective interest rate risk hedging across varying market environments.

5.4 Enhancing classical interest rates hedging for corporate bonds

5.4.1 Why hedging interest rates?

As detailed in the preceding section, corporate bond returns can be theoretically decomposed into two components: movements driven by interest rate risk and movements driven by credit risk. The intuition is straightforward: a corporate bond, like a sovereign bond, is a fixed-income asset that pays a series of cash flows over time. However, it differs from a sovereign bond due to its issuer—a company subject to credit risk. Classical capital structure arbitrage strategies aim to exploit opportunities from credit mispricing by isolating the credit component and eliminating exposure to interest rate risk. This is achieved by constructing an interest rate neutral position through a combination of corporate bonds and sovereign bonds.

To formalize this idea, we first recall the classical bond pricing formula:

$$P = \sum_{t=1}^{T-1} \frac{CF_t}{(1+y)^t} + \frac{CF_T}{(1+y)^T} = P(y) = \sum_{t=1}^{T-1} \frac{CF_t}{(1+s+r)^t} + \frac{CF_T}{(1+s+r)^T} = P(r,s)$$

where P, T and $(CF_t)_{t\in[1,T]}$ represent the price, maturity and annual cash flows of the bond. The final cashflow CF_T contains the face value of the bond. In this expression, y is the bond yield, which can be decomposed into two components, as described in the following formula: $y = r + s$, where r is the risk-free interest rate (e.g. the yield on a sovereign bond) and s is the credit spread, which reflects the

additional risk associated with the bond issuer.

Interest rate risk arises because changes in the interest rate r directly impact the bond price $P(r, s)$, given the inverse relationship between bond prices and interest rates. To hedge against this risk, we aim to construct a portfolio such that:

$$\frac{\partial \text{Portfolio}}{\partial r} = 0 \quad \text{(first-order hedging, or duration hedging)}$$

ensuring the portfolio is insensitive to small changes in interest rates. In cases where large or nonlinear interest rate movements are a concern, we extend this to second-order hedging:

$$\frac{\partial^2 \text{Portfolio}}{\partial r^2} = 0 \quad \text{(second-order hedging, or convexity hedging)}$$

To implement duration or duration-convexity hedging ([22]), we adjust the corporate bond position by adding one sovereign bond leg for delta hedging or two sovereign bond legs with distinct rate sensitivities for delta-gamma hedging. This approach enables the resolution of a system of one or two equations, respectively, to neutralize the portfolio's exposure to interest rate movements at the first or second order. Consequently, the interest rate neutral portfolio effectively isolate the credit risk component, enabling the strategy to focus on exploiting credit mispricing while mitigating interest rate risk.

5.4.2 Theoretical considerations

The foundation of interest rate risk hedging lies in approximating the corporate bond price sensitivity to small changes in interest rates. This sensitivity can be captured using a second-order Taylor expansion. For a small change in the interest rate Δr, the corporate bond price P^B can be expressed as:

$$P^B(r + \Delta r, s) \approx P^B(r, s) + \frac{\partial P^B}{\partial r}\Delta r + \frac{1}{2}\frac{\partial^2 P^B}{\partial r^2}(\Delta r)^2$$

The percentage change in the corporate bond price is then given by:

$$\frac{P^B(r + \Delta r, s) - P^B(r, s)}{P^B(r, s)} \approx -(\Delta r) \cdot \frac{\frac{\partial P^B}{\partial r}}{P^B} + \frac{(\Delta r)^2}{2} \cdot \frac{\frac{\partial^2 P^B}{\partial^2 r}}{P^B}$$

To further simplify the analysis and directly connect interest rate changes to bond yield movements, we introduce the following constant credit risk assumption:

Hypothesis 1 $s = const$

Under Hypothesis 1, the corporate bond price sensitivity with respect to r becomes equivalent to its

sensitivity to the yield y. Mathematically:

$$\frac{\partial P^B}{\partial r} = \frac{\partial P^B}{\partial y} \cdot \frac{\partial y}{\partial r} = \frac{\partial P^B}{\partial y} \cdot \frac{\partial r}{\partial r} = \frac{\partial P^B}{\partial y}$$

This assumption simplifies the hedging analysis. It allows us to use standard yield-based metrics, such as duration and convexity, to approximate the corporate bond's sensitivity to interest rate movements without having to explicitly separate the impact of changes in r and s. Classical first and second order yield sensitivity metrics encompass[3]:

- Modified Duration:
$$MDur = \frac{1}{P} \cdot \sum_{t=1}^{T} \frac{CF_t}{(1+y)^{t+1}} = -\frac{\frac{\partial P}{\partial y}}{P}$$
which represents the percentage price change of a bond for a 100 basis points change in yield.

- Convexity:
$$Conv = \frac{1}{P} \cdot \sum_{t=1}^{T} \frac{CF_t}{(1+y)^{t+2}} = \frac{\frac{\partial^2 P}{\partial y^2}}{P}$$
which represents the basis point change in the bond's modified duration for a 100 basis points change in yield.

While Hypothesis 1 simplifies the analysis, it may not hold exactly in practice. Indeed, yield movements across corporate bonds can differ significantly due to fluctuations in credit spreads, which are bond-specific and subject to both issuer-specific and market-wide factors. In some market conditions, these spread movements may account for a substantial portion of the total yield change, making duration-convexity hedging less applicable in practice. By decomposing the change in corporate bond price with respect to yield, we can separate the effects of interest rate and credit risk, as shown below:

$$\frac{P(y+\Delta y) - P(y)}{\Delta y} = \frac{P(y+\Delta r+\Delta s) - P(y)}{\Delta y} \approx \frac{\Delta r + \Delta s}{\Delta y} \cdot P'(y) = \underbrace{\frac{\Delta r}{\Delta y} \cdot P'(y)}_{\text{Interest rate move}} + \underbrace{\frac{\Delta s}{\Delta y} \cdot P'(y)}_{\text{Credit risk move}}$$

5.4.3 Hedging system

For duration hedging, we combine the corporate bond with one sovereign bond to solve the following equation:

$$N^B \cdot \frac{\partial P^B}{\partial r} + N^R \cdot \frac{\partial P^R}{\partial r} = 0$$

where N^i stands for the notional quantities allocated to asset i. B and R represents corporate and sovereign bond respectively. This is equivalent to:

$$N^B \cdot P^B \cdot MDur^B + N^R \cdot P^R \cdot MDur^R = 0$$

since $\frac{\partial P^B}{\partial r} \approx \frac{\partial P^B}{\partial y}$ (because of Hypothesis 1) and $\frac{\partial P^R}{\partial r} \approx \frac{\partial P^R}{\partial y}$ (because $y \approx r$ for sovereign bonds).

[3]These metrics are applicable to both corporate and sovereign bonds.

For duration-convexity hedging, we combine the corporate bond with two sovereign bonds (R_1 and R_2) whose sensitivities to interest rates are assumed not to be collinear, i.e., $\begin{pmatrix} \frac{\partial P^{R_1}}{\partial r} & \frac{\partial^2 P^{R_1}}{\partial r^2} \\ \frac{\partial P^{R_2}}{\partial r} & \frac{\partial^2 P^{R_2}}{\partial r^2} \end{pmatrix}$ is invertible, solving the system of equations:

$$N^B \cdot \frac{\partial P^B}{\partial r} + N^{R_1} \cdot \frac{\partial P^{R_1}}{\partial r} + N^{R_2} \cdot \frac{\partial P^{R_2}}{\partial r} = 0$$

$$N^B \cdot \frac{\partial^2 P^B}{\partial r^2} + N^{R_1} \cdot \frac{\partial^2 P^{R_1}}{\partial r^2} + N^{R_2} \cdot \frac{\partial^2 P^{R_2}}{\partial r^2} = 0$$

which is equivalent to:

$$N^B \cdot P^B \cdot MDur^B + N^{R_1} \cdot P^{R_1} \cdot MDur^{R_1} + N^{R_2} \cdot P^{R_2} \cdot MDur^{R_2} = 0$$

$$N^B \cdot P^B \cdot Conv^B + N^{R_1} \cdot P^{R_1} \cdot Conv^{R_1} + N^{R_2} \cdot P^{R_2} \cdot Conv^{R_2} = 0$$

since $\frac{\partial^2 P^B}{\partial r^2} \approx \frac{\partial^2 P^B}{\partial y^2}$ (because of Hypothesis 1) and $\frac{\partial^2 P^R}{\partial r^2} \approx \frac{\partial^2 P^R}{\partial y^2}$ (because $y \approx r$ for sovereign bonds).

This system is solved iteratively for multiple pairs of sovereign bonds, with the optimal pair being the one that minimizes the overall hedging cost. Specifically, the chosen pair is the one that achieves the smallest total absolute notional value, as defined by:

$$(R_1, R_2)^* = \arg\min_{R_1, R_2} \left(|N^{R_1}| + |N^{R_2}| \right)$$

In this study, we evaluate two sovereign bond pairs based on maturity (or "tenor"): a 2–5 year pair and a 5–10 year pair. We select the pair that minimizes the hedging cost according to the above criterion.

5.4.4 Dynamics of the empirical hedge ratio across credit risk and market regimes

Building on the previous discussion, we aim to empirically assess whether the first and second order derivatives of corporate bond prices with respect to yield (i.e. duration and convexity) serve as reliable proxies for their behavior with respect to interest rates under varying credit quality and market conditions. In particular, we aim to answer the following questions: To what extent do these theoretical yield based metrics accurately reflect the observed price movements of corporate bonds in response to changes in interest rates? What adjustments or enhancements can be made to improve the effectiveness of this traditional hedging approach?

To address these questions, we estimate the empirical hedge ratio HR by aligning, for the same underlying company, the P&L of the corporate bond (B), the P&L of the CDS (C), and the P&L of the corresponding sovereign bond pair (R_1 and R_2), sized appropriately by solving the system introduced in Section 5.4.3. This alignment reflects the idea that the combined performance of the corporate bond and CDS should match the performance of the sized sovereign bond position, as the latter absorbs the interest rate risk. According to theory, the following relationship should hold:

$$N^B \cdot \Delta P^B + N^C \cdot \Delta P^C - (N^{R_1} \cdot \Delta P^{R_1} + N^{R_2} \cdot \Delta P^{R_2}) = 0$$

To estimate HR, we reframe this equation as:

$$N^B \cdot \Delta P^B + N^C \cdot \Delta P^C = HR \cdot (N^{R_1} \cdot \Delta P^{R_1} + N^{R_2} \cdot \Delta P^{R_2})$$

$$\iff HR = \frac{N^B \cdot \Delta P^B + N^C \cdot \Delta P^C}{N^{R_1} \cdot \Delta P^{R_1} + N^{R_2} \cdot \Delta P^{R_2}}$$

where HR represents the degree to which it is optimal to use sized sovereign bonds to hedge the corporate bond's interest rate exposure. This allows us to quantify the significance of duration-convexity hedging in taming interest rate movements. Note that we choose N^B and N^C to be equal, meaning that we strive to cover, using CDS, the exact same notional amount that is at risk in the corporate bond position.

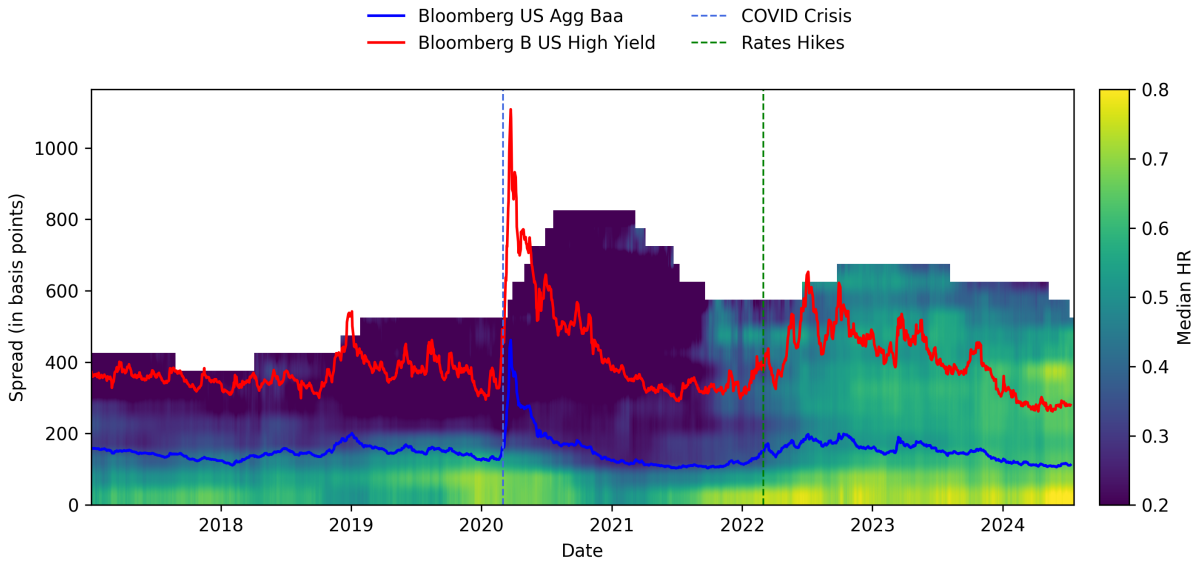

Figure 5.3: Heatmap of the median HR through time and spread bins

Figure 5.3 displays the daily median HR for corporate bonds grouped by spread bins, providing insights into how the HR evolves over time and across different credit spread levels. Moreover, to provide a broader context of credit risk in the market, we present the spreads of two key corporate bond indices. The first is the Bloomberg US Agg Baa Index, which includes bonds with a credit rating close to BBB, offering a perspective on IG securities. The second is the Bloomberg B US High Yield Index, which consists of corporate bonds rated B, highlighting trends in the HY market. This figure illustrates the interplay between market conditions and credit spread dynamics, illustrating their influence on the optimal hedging approach. Specifically, as shown in the figure, the HR declines as credit risk increases. This suggests that as credit risk becomes more prominent, the ability to hedge interest rate movements using yield metrics diminishes, and the corporate bond's exposure to credit risk increasingly dominates. Besides, the HR shows notable variation over time, aligning with the three market regimes discussed in Section 5.3.3. These regimes—pre-COVID, COVID, and post-COVID—exhibit unique patterns, influenced by heightened credit risk during the COVID crisis and significant interest rate increase in the post-COVID period.

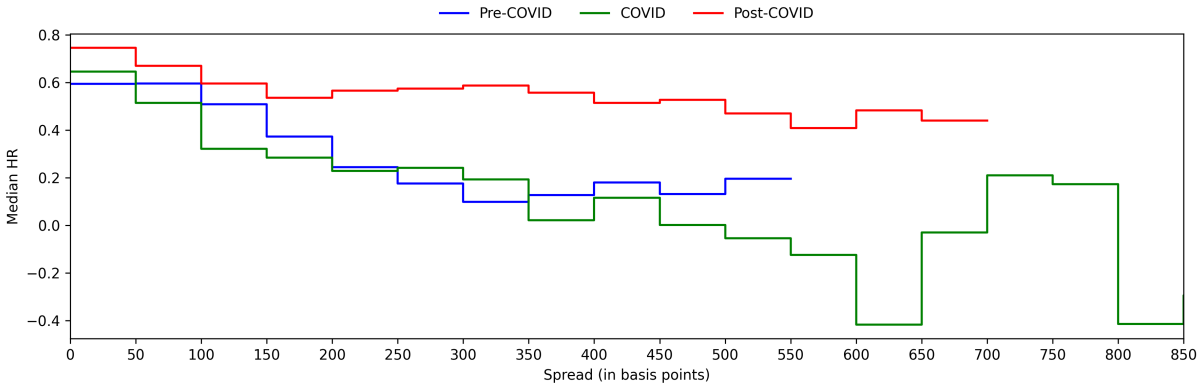

Figure 5.4: Median HR across credit spread bins in three distinct market regimes

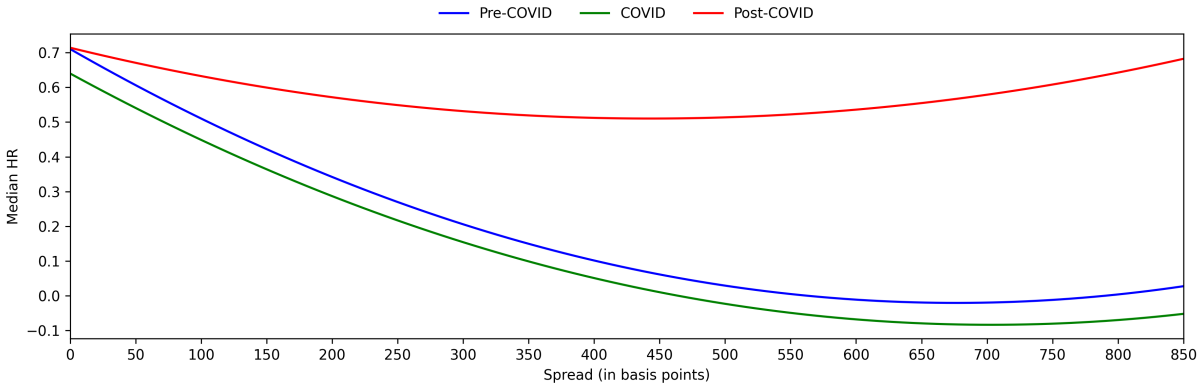

Figure 5.5: Polynomial fit of median HR across credit spread bins in three distinct market regimes

Figure 5.4 displays the evolving median HR for corporate bonds across the three different market regimes introduced in Section 5.3.3. Figure 5.5 displays the polynomial fit of the median HR values from Figure 5.4, providing a smoother representation of the trends across credit spread bins. During periods dominated by credit risk, such as the COVID crisis, the HR consistently moves to a lower level compared to the pre-COVID period. As shown in Figure 5.5, the hedging ratio curve for the COVID period is lower across all credit spread bins, indicating that the optimal hedge level was reduced as credit spreads became the primary driver of corporate bond prices, outweighing the impact of interest rate changes and reducing the effectiveness of duration-convexity hedging. This observation aligns with the findings of [40], which demonstrate that during periods of economic distress, such as the 2008 financial crisis, the interest rate sensitivity of corporate bonds diminishes as credit risk dominates.

Conversely, in the post-COVID period, marked by rising interest rates, the HR curve flattens and shifts to a higher level, reflecting an increased need for interest rate hedging. This is consistent with the conclusions of [40], which indicate that interest rate sensitivity increases during phases of economic normalization, as interest rate changes regain dominance over credit spread dynamics. Additionally, while HR decreases as credit spreads widen in this period, the decline is less pronounced than in the pre-COVID and COVID periods, suggesting that the differentiation in hedging across credit spread bins is less significant. These findings emphasize the dynamic nature of corporate bond hedging and the need

to adjust the HR to reflect prevailing market conditions. A static, one-size-fits-all approach to interest rate hedging is inadequate, as the effectiveness of duration and convexity adjustments varies substantially depending on the corporate bond's credit profile and the broader economic environment. Moreover, in their research, [40] highlight that the sensitivity of corporate bonds to interest rates can shift from positive to negative depending on the phase of the business cycle. This further underscores the critical importance of tailoring hedging strategies to align with evolving market conditions.

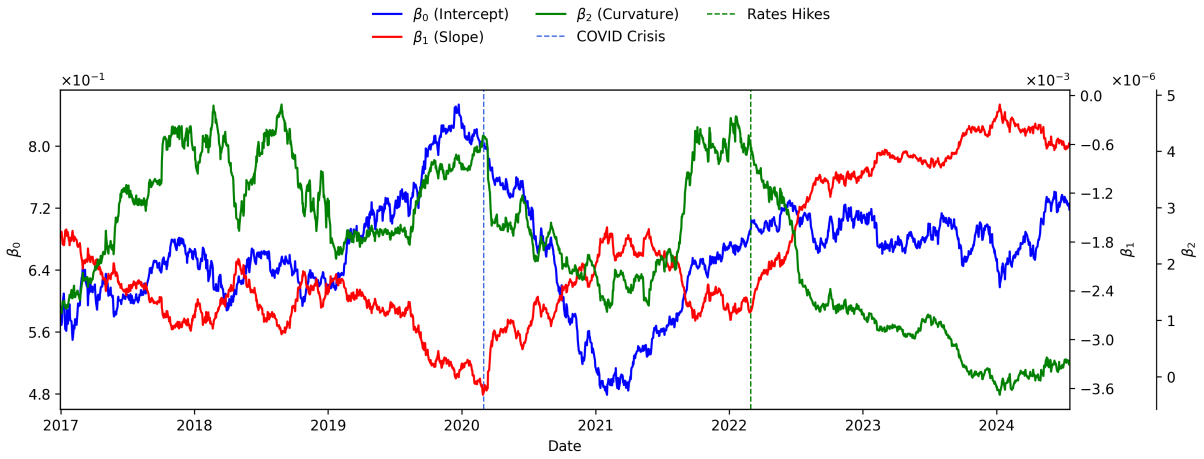

Figure 5.6: Evolution of β_0, β_1, β_2 over time

Following these insights, we characterize the relationship between the HR and credit spread over time, using a one-year rolling linear regression of median HR on spread. Mathematically, the regression model is expressed as:

$$HR_{i,t} = \beta_{t,0} + \beta_{t,1} \cdot s_{i,t} + \beta_{t,2} \cdot s_{i,t}^2 + \epsilon_{i,t} \tag{5.5}$$

where $s_{i,t}$ represents the spread of bin i at time t. $\beta_{t,0}$ is the time-varying intercept, indicating the level of hedging independent of credit spread movements. $\beta_{t,1}$ captures the linear sensitivity of HR to credit spread changes, $\beta_{t,2}$ accounts for non-linear effects of credit spread on HR, and ϵ_t is the residual term. By including the quadratic term, this model captures potential curvature in the relationship.

The results are illustrated in Figure 5.6, where the intercept (β_0), slope (β_1) and curvature (β_2) of the fitted linear regressions are displayed over time. During the pre-COVID period, the level of hedging remains moderate. β_0 is positive, with a value of 0.55, indicating a medium level of hedging. β_1, however, is negative (-0.001), suggesting that optimal hedging decreases as the credit spread increases. Both β_0 and β_1 remain relatively stable throughout this period. This stability in the HR indicates that during calm market conditions, the relationship between the HR and credit spread remains consistent, aligning with previous research ([29]). Besides, β_2, representing the curvature of the relationship between HR and credit spread, is relatively high compared to the post-COVID period but not statistically significant, with values close to zero. Although β_2 is not statistically significant, its relatively higher magnitude during this stable market period indicates a more convex, non-linear relationship between HR and credit spread.

Episode of heightened credit risk, such as the COVID period, reveals notable shifts in both the β_0 and β_1. Specifically, β_0 decreases towards zero, and β_1 decreases in magnitude, weakening the relationship

between credit spread and HR. During times of market turmoil, credit spreads tend to widen significantly, and the relationship between corporate bond prices and interest rate movements diminishes, resulting in a less effective duration-convexity hedge. Moreover, during the COVID period, β_2 remains relatively high compared to the post-COVID period but is not statistically significant. In essence, during times of crisis, the interest rate hedge provided by sovereign bonds becomes less relevant, and the focus shifts toward managing credit risk. This phenomenon is consistent with the flight-to-quality behavior identified in [41], where heightened risk aversion results in a decoupling of credit spreads and interest rate movements, diminishing the role of sovereign bonds in providing effective hedging.

In the post-COVID period, the relationship between the HR and credit spread evolves further. β_0 increases to 0.65, suggesting that a high level of hedging is optimal. This value even exceeds the level observed during the pre-COVID period. Meanwhile, β_1 increases and approaches zero, suggesting that hedging should be performed independently of the credit risk associated with corporate bonds. The reduced curvature in this period, reflected by the near-zero β_2, implies that the relationship between HR and credit spread is more linear. This shift reflects the critical importance of robust hedging strategies in environments characterized by rising interest rates, aligning with empirical observations of capital gains sensitivity in periods of market normalization ([41]).

These findings underscore the importance of adjusting hedging strategies based on prevailing market conditions, as the relative significance of credit risk and interest rate risk can shift dramatically, requiring a dynamic and responsive approach to managing corporate bond exposure.

5.4.5 Comparative performance of hedge ratio adjustment strategies

In this section, we present and evaluate various methods for computing HR to hedge interest rate risk in a portfolio consisting of corporate bonds and CDS, with the assumption that the portfolio primarily exhibits interest rate risk after accounting for credit risk. Our goal is to identify the approach that minimizes performance residuals and volatility, thereby achieving an optimal hedge. By comparing hedging strategies across distinct market regimes—pre-COVID, COVID, and post-COVID—we examine their adaptability and effectiveness in dynamically responding to changing market conditions. This analysis provides insights into which approaches best align with the dual objectives of mitigating interest rate risk and maintaining stability in corporate bond portfolios, ensuring relevance across varying credit risk profiles and economic conditions.

Benchmark approaches

We introduce and benchmark two new approaches aimed at finding the optimal HR. The first is a heuristic method that approximates the HR using the ratio of sovereign bond yield to corporate bond yield. The second approach involves regressing past empirical HR on spread using Equation 5.5. These two methods are benchmarked against classical strategies, including not hedging, fully hedging, or implementing a hedge that decreases with credit rating. We introduce the following methods to compute HR:

- **No hedge (NH):** In this benchmark, no hedging is applied, meaning HR is set to zero:

$$HR_{i,t} = 0$$

  This approach allows us to analyze the raw performance of the portfolio without any interest rate risk mitigation.

- **Fully hedge (FH):** In the fully hedged scenario, the portfolio is entirely hedged against yield movements by setting HR to one:

$$HR_{i,t} = 1$$

  This benchmark provides a view of the maximum theoritical risk reduction.

- **Linear BBB-B credit rating (CR):** This strategy applies a dynamic HR that decreases linearly from 1 for corporate bonds with BBB-rated credit spreads to 0 for corporate bonds with B-rated credit spreads. HR is calculated as follows:

$$HR_{i,t} = \frac{s_{i,t-1} - s_{\text{B Index},t-1}}{s_{\text{BBB Index},t-1} - s_{\text{B Index},t-1}}$$

  This approach gradually reduces the hedge as credit spreads widen, balancing risk exposure across different credit qualities.

- **Yields ratio (YR):** In this benchmark, HR is computed using an exponential weighting scheme, which considers both sovereign bond yields and credit spreads at $t-1$:

$$HR_{i,t} = \frac{\exp(y^R_{i,t-1})}{\exp(y^R_{i,t-1}) + \exp(s_{i,t-1})}$$

  This method allocates more hedging to corporate bonds for which the interest rates—represented by the yield of the corresponding sized sovereign bond position—accounts for a larger portion of the corporate bond's overall yield.

- **Regression model (RM):** This method uses a polynomial function to compute HR at time t. This function, denoted as $\tilde{HR}_{t-1}$, is fitted using a one-year rolling window of data up to time $t-1$ (as shown in Equation 5.5). The fitted polynomial is based on the observed relationship between credit spreads and HR, with the coefficients $\hat{\beta}_{0,t-1}$, $\hat{\beta}_{1,t-1}$ and $\hat{\beta}_{2,t-1}$ representing the estimated parameters at time $t-1$. The model is then evaluated at the credit spread at time $t-1$:

$$HR_{i,t} = \tilde{HR}_{t-1}(s_{i,t-1})$$

  By continuously adjusting to evolving market conditions, it ensures the hedging strategy remains responsive and accurate, capturing the dynamic relationship between corporate bond spreads and interest rates while adapting to changing market environments.

Performances

To enable robust comparison of benchmark strategies over time, we neutralize credit risk disparities and construct a portfolio with constant daily allocation. Specifically, we allocate \$1 daily across the corporate bond universe, ensuring a weighted average credit spread target of 400 basis points. We study the performance of the different hedged portfolios over time. We define the portfolio's P&L at time t as:

$$\Delta P_t^{Portfolio} = \sum_{i \in L_t} w_{i,t} \cdot \left(N_{i,t}^B \cdot \Delta P_{i,t}^B - N_{i,t}^C \cdot \Delta P_{i,t}^C - HR_{i,t} \cdot (N_{i,t}^{R_1} \cdot \Delta P_{i,t}^{R_1} + N_{i,t}^{R_2} \cdot \Delta P_{i,t}^{R_2}) \right) \tag{5.6}$$

where L_t represents the set of corporate bonds available at time t. $w_{i,t}$ represent the allocation weight[4] for hedged corporate bond position i at time t. Here $N_{i,t}^B = N_{i,t}^C = 1 \ \forall \ i, \ t$[5]. $N_{i,t}^{R_1}$ and $N_{i,t}^{R_2}$ are computed solving the hedging system described in Section 5.4.3.

To evaluate the effectiveness of the various hedging approaches in mitigating interest rate risk within corporate bond positions, we analyze portfolio performance alongside several volatility metrics of portfolio P&L, as described in Equation 5.6. Since the portfolio is designed to neutralize both credit and interest rate risks, the residual performance should theoretically be zero, with volatility minimized. Specifically, we assess the mean absolute error (MAE), standard deviation (σ), interquartile range (IQR), value at risk (VaR), and conditional value at risk ($CVaR$) of portfolio performance. These metrics together provide a comprehensive evaluation of the stability and consistency of each hedging strategy over time. The comparative results are summarized in Table 5.2.

[4]Details about the allocation system is presented in Appendix 5.6.3.

[5]With $N_{i,t}^B$ and $N_{i,t}^C$ equal, we strive to cover the exact notional that is at risk in the corporate bond position using CDS.

Period		σ	IQR	VaR	$CVaR$	MAE
pre-COVID	RM	2.209	2.177	-3.056	-4.747	1.518
	YR	**2.208**	2.178	-3.164	-4.728	**1.513**
	CR	2.300	2.261	-3.636	-4.968	1.589
	FH	2.784	2.988	-4.073	-6.372	1.996
	NH	2.209	**2.162**	**-2.904**	**-4.466**	1.548
COVID	RM	4.583	3.104	-4.568	-10.645	2.487
	YR	**4.546**	3.150	**-4.551**	**-10.643**	**2.469**
	CR	5.085	3.554	-5.116	-11.764	2.815
	FH	5.477	3.947	-5.372	-12.506	3.109
	NH	4.556	**3.059**	-4.803	-10.844	2.525
post-COVID	RM	**2.124**	**2.721**	**-3.414**	**-4.678**	**1.645**
	YR	2.195	2.744	-3.770	-5.149	1.667
	CR	2.307	2.834	-3.770	-5.177	1.759
	FH	2.540	3.134	-4.414	-5.691	1.939
	NH	2.846	3.761	-4.741	-6.108	2.206
Overall	RM	**3.004**	**2.530**	**-3.657**	**-6.489**	**1.814**
	YR	3.005	2.576	-3.764	-6.628	**1.814**
	CR	3.276	2.737	-3.959	-7.055	1.967
	FH	3.635	3.226	-4.792	-7.942	2.273
	NH	3.178	2.831	-4.300	-7.020	2.015

Note: Bold font indicates the best metric.
All values are expressed in thousands.

Table 5.2: Evaluation of hedging strategies through daily volatility metrics over different time periods

Based on the results from Table 5.2, we observe distinct performance patterns across the different time periods (pre-COVID, COVID, and post-COVID). During the pre-COVID period, the NH strategy emerges as the most effective approach, followed closely by the YR and RM models. Specifically, the NH strategy shows the lowest IQR (2.162), VaR (-2.904) and $CVaR$ (-4.466) indicating greater stability than other strategies. Besides, the FH strategy proves to be the least effective during this period, exhibiting significantly higher volatility metrics. Indeed, the FH strategy exhibits higher values for both σ (2.784) and IQR (2.988), signaling a less effective hedging strategy during the pre-COVID period.

In the COVID period, the YR strategy performs the best, followed closely by the RM and NH approaches. Notably, the FH strategy once again lags behind, showing the worst performance in terms of volatility metrics. YR achieves the lowest σ (4.546) and IQR (3.150), with RM and NH close behind. FH shows a much higher σ (5.477) and IQR (3.947), reinforcing its ineffectiveness during this period. These results are consistent with the macroeconomic dynamics of the COVID period, where credit risk dominated, and the YR and RM strategies were more adept at navigating these conditions.

In the post-COVID period, we observe a shift in market conditions, with rising interest rates becoming

a key factor. The RM strategy emerges as the best-performing approach ($\sigma = 2.124$, $IQR = 2.721$), closely followed by YR ($\sigma = 2.195$, $IQR = 2.744$). Both strategies outperform NH ($\sigma = 2.846$, $IQR = 3.761$), indicating that they are better suited to handle periods marked by rising interest rates. While the FH strategy shows some improvement during this period, surpassing NH, it still lags behind both RM and YR. This suggests that while NH and FH may provide stability in certain market conditions, they fail to outperform the more adaptive RM and YR strategies during periods of heightened interest rate volatility.

Looking at the overall results, which encompass both the credit crisis and the subsequent post-crisis period, RM and YR consistently outperform other benchmark strategies across nearly all volatility metrics and residual performance level. Both strategies lead in terms of σ, IQR, VaR, $CVaR$ and MAE, showcasing their robustness across varying market conditions. Importantly, the consistent outperformance of RM and YR over the CR strategy demonstrates that focusing solely on credit risk is insufficient for developing a robust hedging approach. Furthermore, the fact that YR, a simpler heuristic-based approach, performs similarly to the more complex RM strategy highlights its practical advantages, as it delivers almost identical results with potentially lower computational costs.

These findings are further corroborated by similar trends observed in the weekly and monthly volatility metrics, as presented in Tables 5.3 and 5.4 in Appendix 5.6.4. This provides additional evidence of the robustness of our strategies across different time horizons, reinforcing the effectiveness of adaptive models in managing interest rate risk under varying market conditions.

5.5 Conclusion

This research introduces a framework for decomposing the equity-corporate bond correlation into its fundamental drivers: credit risk and interest rate risk. Our methodology leverages the equity-CDS correlation to capture credit risk and the equity-sovereign bond correlation to represent interest rate sensitivity. By isolating these components, we provide a perspective on how credit and interest rate risks interact and shape corporate bond performance across different market conditions and credit qualities. This approach empowers investors to make more informed decisions when assessing equity and corporate bond interactions.

For IG bonds, interest rate dynamics remain the primary driver of the equity-corporate bond correlation, reflecting their sensitivity to rate changes due to lower credit spreads and minimal default risk. However, during periods of heightened market stress, such as the COVID crisis, credit risk temporarily takes precedence, as concerns about issuer resilience escalate. Post-crisis, interest rate sensitivity regains dominance, driven by significant interest rate hikes beginning in 2022. In contrast, HY bonds exhibit a persistent reliance on credit risk as the main determinant of their correlation with equities. This influence is particularly pronounced during the COVID crisis, when widening credit spreads and heightened default risk reshape market dynamics. While interest rate factors gain some relevance post-crisis, credit risk remains the dominant force.

Additionally, we highlight the dynamic role of macroeconomic factors, such as inflation and growth

uncertainties, in influencing the equity-corporate bond correlation through both interest rate sensitivity and credit risk. Specifically, inflation uncertainty tends to amplify interest rate sensitivity, particularly in periods of monetary policy tightening, while growth uncertainty predominantly impacts credit risk by altering perceptions of issuer resilience and default probability. These findings illustrate how the equity-corporate bond relationship evolves in response to changing economic conditions, with macroeconomic factors driving the dynamic interplay between interest rate sensitivity and credit risk.

Finally, we propose and validate two adaptive hedging strategies for corporate bond portfolios—based on the yield ratio (YR) and regression model (RM)—aimed at isolating credit risk by neutralizing interest rate exposure. Our empirical results demonstrate that these strategies, which dynamically adjust hedge ratios based on evolving macroeconomic conditions and credit quality, outperform traditional hedging methods across different market regimes. Specifically, we show that the effectiveness of duration-convexity-based hedging diminishes as credit risk becomes more prominent, such as during the COVID crisis, and increases again when interest rates rise in the post-COVID period. Our analysis highlights that a static approach to hedging, relying solely on yield metrics like duration and convexity, is insufficient. Instead, the optimal hedge ratio must reflect the shifting balance between interest rate and credit risks, ensuring more stable and robust performance of corporate bond portfolios in varying market environments.

Future research can refine the framework by incorporating additional macroeconomic variables, such as global liquidity, and fiscal policy, while exploring sector-specific variations in credit and interest rate sensitivity. Additionally, further empirical testing of the adaptive hedging strategies over extended time periods, including crisis events or prolonged low-interest rates, can enhance the model's practical applicability and validate its effectiveness across diverse market conditions and regulatory environments.

5.6 Appendix

5.6.1 Decomposing the equity-corporate bond correlation: credit versus interest rate components

Understanding the relative contributions of the equity-sovereign bond correlation ($\rho_{E,R}$) and the equity-CDS correlation ($\rho_{E,C}$) to the overall equity-corporate bond correlation ($\rho_{E,B}$) is challenging due to the inherent endogeneity between interest rates and credit risk. This endogeneity, captured in the covariance term $\text{cov}(R, C)$ in Equation 5.2, complicates the direct attribution of changes in $\rho_{E,B}$ to either driver. To disentangle these effects, we employ a linear regression framework that decomposes the observed correlation into marginal effects, quantifying the contribution of each driver in isolation. We reformulate Equation 5.2, and model $\rho_{E,B}$ as a function of two components, $a \cdot b$ and $c \cdot d$:

$$\rho_{E,B} = \beta_1 \cdot (a \cdot b) + \beta_2 \cdot (c \cdot d) + \epsilon \tag{5.7}$$

where:

- $a = \frac{\sigma_R}{\sqrt{\sigma_R^2+\sigma_C^2-2\cdot\Sigma_{R,C}}}$, representing the relative influence of sovereign bond volatility
- $b = \rho_{E,R}$, the equity-sovereign bond correlation
- $c = \frac{\sigma_C}{\sqrt{\sigma_R^2+\sigma_C^2-2\cdot\Sigma_{R,C}}}$, representing the relative influence of CDS volatility
- $d = -\rho_{E,C}$, the inverse equity-CDS correlation

The regression coefficients, β_1 and β_2, represent the relative sensitivities of $\rho_{E,B}$ to the components $a \cdot b$ and $c \cdot d$, respectively. Builing on this, we quantify the independent impact of $\rho_{E,R}$ (b) and $\rho_{E,C}$ (d) by computing their marginal effects as:

$$M_b = \beta_1 \cdot \overline{a}, \quad M_d = \beta_2 \cdot \overline{c}$$

where $\overline{a}$ and $\overline{c}$ are the mean values of a and c, respectively. These marginal effects represent the expected change in $\rho_{E,B}$ resulting from a unit change in $\rho_{E,R}$ or $\rho_{E,C}$, while holding all other variables constant.

However, marginal effects alone do not completely capture the contributions of $\rho_{E,R}$ and $\rho_{E,C}$ because they depend on the underlying variability of b and d. To address this, we turn to the variance decomposition framework of [55] to measure the explained variance attributable to each driver. Variance decomposition provides a robust mechanism to attribute the variability of $\rho_{E,B}$ to either $\rho_{E,R}$ or $\rho_{E,C}$. Following the general framework:

$$\mathbb{V}(Y) = \mathbb{V}(\mathbb{E}[Y|X]) + \mathbb{E}[\mathbb{V}(Y|X)]$$

where $\mathbb{V}(\mathbb{E}[Y|X])$ is the explained variance, and $\mathbb{E}[\mathbb{V}(Y|X)]$ is the residual variance. For linear regression, the explained variance is given by:

$$\mathbb{V}(\mathbb{E}[Y|X]) = \beta^\top \operatorname{cov}(X, X)\beta$$

where $\operatorname{cov}(X, X)$ is the covariance matrix of the input features. Assuming that all input features are orthogonal, the contribution λ_i^y of each feature i to the explained variance can then be expressed as:

$$\lambda_i^y = \frac{\mathbb{V}(\mathbb{E}[Y|X_i])}{\sum_{j=1}^n \mathbb{V}(\mathbb{E}[Y|X_j])} \cdot R^2 = \frac{\beta_i^2 \cdot \sigma_{X_i}^2}{\sum_{j=1}^n \beta_j^2 \cdot \sigma_{X_j}^2} \cdot R^2 \tag{5.8}$$

where R^2 is the coefficient of determination of the regression, defined as:

$$R^2 = \frac{\mathbb{V}(\mathbb{E}[Y|X])}{\mathbb{V}(Y)}$$

In our setting, the input features are $a \cdot b$ and $c \cdot d$, but we aim to evaluate the marginal contributions of b and d. To achieve this, we compute the contributions as:

$$\lambda_b^{\rho_{E,B}} = \frac{(\beta_1 \cdot \overline{a})^2 \cdot \sigma_b^2}{(\beta_1 \cdot \overline{a})^2 \cdot \sigma_b^2 + (\beta_2 \cdot \overline{c})^2 \cdot \sigma_d^2} \cdot R^2$$

$$\lambda_d^{\rho_{E,B}} = \frac{(\beta_2 \cdot \overline{c})^2 \cdot \sigma_d^2}{(\beta_1 \cdot \overline{a})^2 \cdot \sigma_b^2 + (\beta_2 \cdot \overline{c})^2 \cdot \sigma_d^2} \cdot R^2$$

This approach isolates the contributions of the equity-sovereign bond correlation ($\lambda_b^{\rho_{E,B}}$) and equity-CDS correlation ($\lambda_d^{\rho_{E,B}}$) to the total explained variance in $\rho_{E,B}$, allowing us to identify which factor dominates under different market conditions.

5.6.2 Macroeconomic drivers of the equity-corporate bond correlation

Another objective is to investigate the relative contributions of the macroeconomic features such as growth volatility (σ_g^2), inflation volatility (σ_π^2), and the growth-inflation correlation ($\rho_{g,\pi}$) in explaining $\rho_{E,R}$ and $\rho_{E,C}$. Additionally, we aim to establish connections between these macroeconomic features and their role in explaining the overall $\rho_{E,B}$.

To achieve this, we denote by $(X_i)_{1\leq i\leq n}$ the selected macroeconomic features. We define $(\alpha_i)_{1\leq i\leq n}$ and $(\gamma_i)_{1\leq i\leq n}$ as the respective coefficients from the linear regressions of b and d with respect to $(X_i)_{1\leq i\leq n}$ such that:

$$b = \rho_{E,R} = \alpha_0 + \sum_{i=1}^{n} \alpha_i X_i + \epsilon_b, \quad d = \rho_{E,C} = \gamma_0 + \sum_{i=1}^{n} \gamma_i X_i + \epsilon_d \tag{5.9}$$

Using the same variance decomposition method of [55] along with the contributions defined in Equation 5.8, we can express the relative contributions of $(X_i)_{1\leq i\leq n}$ in relation to b and d:

$$\lambda_{X_i}^b = \frac{\alpha_i^2 \cdot \sigma_{X_i^2}}{\sum_{j=1}^n \alpha_j^2 \cdot \sigma_{X_j^2}} \cdot R_b^2$$
$$\lambda_{X_i}^d = \frac{\gamma_i^2 \cdot \sigma_{X_i^2}}{\sum_{j=1}^n \gamma_j^2 \cdot \sigma_{X_j^2}} \cdot R_d^2$$

where $\lambda_{X_i}^b$ represents the contribution of the feature X_i in explaining the explained variance of b, and $\lambda_{X_i}^d$ signifies the contribution of the feature X_i in explaining the explained variance of d. R_b^2 and R_d^2 represent the coefficients of determination for the regressions derived from Equation 5.9, corresponding to the dependent variables b and d, respectively.

Building on these results, we can estimate the contributions of the macroeconomic features $(X_i)_{1\leq i\leq n}$ in explaining the R^2 of the linear regression in Equation 5.7, which models $\rho_{(E,B)}$ as a function of $a \cdot b$ and $c \cdot d$. This approach allows us to determine through which channel—interest rate or credit—a given macroeconomic feature transmits its impact, as well as how the strength of this transmission evolves over time. Thus, we define:

$$\lambda_{X_i}^{b\to^{\rho_{E,B}}} = \lambda_{X_i}^b \cdot \lambda_b^{\rho_{E,B}}$$
$$\lambda_{X_i}^{d\to^{\rho_{E,B}}} = \lambda_{X_i}^d \cdot \lambda_d^{\rho_{E,B}}$$

where $\lambda_{X_i}^{b\to^{\rho_{E,B}}}$ represents the contribution of the macroeconomic feature X_i in explaining the total explained variance through b (interest rate channel), and $\lambda_{X_i}^{d\to^{\rho_{E,B}}}$ represents the contribution of the feature X_i in explaining the total explained variance through d (credit risk channel).

The inclusion of the two residuals, ϵ_b and ϵ_d, inevitably reduces the new R^2. Specifically, the revised

total R^2, which captures the explanatory power of the macroeconomic features through $\rho_{E,R}$ and $\rho_{E,C}$, is given by $R^2_{\text{New}} = \sum_{i=1}^{n}(\lambda_{X_i}^{b\to\rho_{E,B}} + \lambda_{X_i}^{d\to\rho_{E,B}})$. By construction, this value is bounded above by the original R^2.

5.6.3 Allocation system

Let B_t be a partition of L_t, representing the available sixteen 50 basis points spread bins at time t.

$$B_t = \{b_{k,t} \mid b_{k,t} = \{i \in L_t \mid 50(k) \leq s_i < 50(k+1)\} \quad \forall k \in [\![0, 15]\!]\}$$

Let H_t be a partition of B_t, separating the spread bins into two groups. The first group is made of spread bin with credit spreads lower than 400 basis points and the other with credit spreads greater than 400 basis points, our target credit spread.

$$H_t = \begin{cases} \ominus_t = \{b \in B_t \mid s_i < 400, \quad \forall i \in b\} \\ \oplus_t = \{b \in B_t \mid s_i \geq 400, \quad \forall i \in b\} \end{cases}$$

We visually describe this decomposition below:

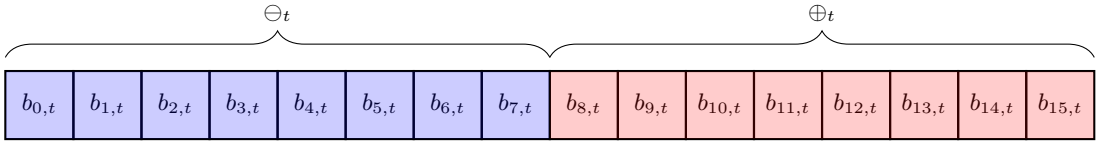

Figure 5.7: L_t partitioning

The allocation is designed to ensure that the weighted average credit spread equals 400 basis points, and that the weighted average allocation equals \$1. To ensure a unique solution, we define two main constraints—the allocation is uniform over credit spread bins in the same half, and it is uniform over corporate bonds in the same spread bin. Mathematically, we aim to solve:

$$\sum_i w_{i,t} \cdot s_{i,t} = 400$$
$$\sum_i w_{i,t} \cdot N^B_{i,t} = 1$$

Subject to the constraints:

$$\sum_{i \in b} w_{i,t} = \sum_{i \in b'} w_{i,t} \quad \forall b, b' \in h, \forall h \in H_t$$
$$w_{i,t} = w_{i',t} \quad \forall i, i' \in b, \forall b \in B_t$$

We can rewrite the constraints as:

$$\sum_{i \in b} w_{i,t} = \frac{\sum_{b \in h} \sum_{i \in b} w_{i,t}}{|h|} \quad \forall b \in h, \forall h \in H_t$$
$$w_{i,t} = \frac{\sum_{i \in b} w_{i,t}}{|b|} = \frac{\sum_{b \in h} \sum_{i \in b} w_{i,t}}{|h| \cdot |b|} \quad \forall i \in b, \forall b \in h, \forall h \in H_t$$

Then, we can rewrite the system as:

$$\sum_{h \in H_t} \sum_{b \in h} \sum_{i \in b} \left(\sum_{k \in \bigcup_{b \in h} b} \frac{w_{k,t}}{|h| \cdot |b|} \right) s_{i,t} = 400$$
$$\sum_{h \in H_t} \sum_{b \in h} \sum_{i \in b} \sum_{k \in \bigcup_{b \in h} b} \frac{w_{k,t}}{|h| \cdot |b|} = 1$$

Simplifying further:

$$\sum_{b \in \ominus_t} \sum_{i \in b} \left(\sum_{k \in \bigcup_{b \in \ominus_t} b} \frac{w_{k,t}}{| \ominus_t | \cdot |b|} \right) \cdot s_{i,t} + \sum_{b \in \oplus_t} \sum_{i \in b} \left(\sum_{k \in \bigcup_{b \in \oplus_t} b} \frac{w_{k,t}}{| \oplus_t | \cdot |b|} \right) \cdot s_{i,t} = 400$$
$$\sum_{b \in \ominus_t} \sum_{i \in b} \sum_{k \in \bigcup_{b \in \ominus_t} b} \frac{w_{k,t}}{| \ominus_t | \cdot |b|} + \sum_{b \in \oplus_t} \sum_{i \in b} \sum_{k \in \bigcup_{b \in \oplus_t} b} \frac{w_{k,t}}{| \oplus_t | \cdot |b|} = 1$$

We now solve this system of two equations and two variables:

$$\sum_{k \in \bigcup_{b \in \ominus_t} b} w_{k,t} \cdot \sum_{b \in \ominus_t} \sum_{i \in b} \frac{s_{i,t}}{| \ominus_t | \cdot |b|} + \sum_{k \in \bigcup_{b \in \oplus_t} b} w_{k,t} \cdot \sum_{b \in \oplus_t} \sum_{i \in b} \frac{s_{i,t}}{| \oplus_t | \cdot |b|} = 400$$
$$\sum_{k \in \bigcup_{b \in \ominus_t} b} w_{k,t} + \sum_{k \in \bigcup_{b \in \oplus_t} b} w_{k,t} = 1$$

5.6.4 Weekly and monthly metrics

Period		σ	IQR	VaR	$CVaR$	$Mean$
pre-COVID	RM	4.891	4.728	**-6.540**	-11.150	3.426
	YR	4.862	4.755	-6.852	-11.161	3.428
	CR	5.044	**4.615**	-6.835	-11.557	**3.409**
	FH	6.168	6.250	-9.002	-15.353	4.151
	NH	**4.839**	5.207	-6.822	**-10.497**	3.580
COVID	RM	12.284	**7.538**	-8.677	-28.176	**6.520**
	YR	12.276	7.810	**-8.624**	-28.414	6.570
	CR	**11.563**	8.665	-9.757	**-27.153**	6.802
	FH	11.875	9.017	-10.765	-29.136	7.285
	NH	12.720	8.272	-9.068	-30.013	6.817
post-COVID	RM	**4.476**	6.680	**-7.761**	**-9.158**	3.589
	YR	4.564	6.569	-8.410	-9.934	**3.581**
	CR	4.915	7.030	-8.853	-9.894	3.897
	FH	4.939	**6.515**	-8.496	-10.592	3.945
	NH	6.527	8.570	-11.958	-13.120	5.153
Overall	RM	7.507	5.931	**-7.961**	**-16.083**	**4.302**
	YR	7.513	**5.595**	-8.221	-16.381	4.309
	CR	**7.325**	6.206	-8.806	-16.351	4.464
	FH	7.807	6.928	-9.630	-18.520	4.933
	NH	8.146	7.411	-9.748	-18.149	4.942

Note: Bold font indicates the best metric.

All values are expressed in thousands.

Table 5.3: Evaluation of hedging strategies through weekly volatility metrics over different time periods

Period		σ	IQR	VaR	$CVaR$	$Mean$
pre-COVID	RM	10.648	**10.822**	**-13.265**	-21.441	8.325
	YR	10.465	11.359	-14.033	-21.263	8.151
	CR	11.019	10.840	-14.170	-23.852	8.596
	FH	13.514	16.089	-22.420	-34.001	10.381
	NH	**10.352**	11.533	-13.307	**-18.681**	**8.126**
COVID	RM	21.555	18.616	**-20.685**	**-47.626**	14.706
	YR	21.531	**15.852**	-23.769	-49.160	14.590
	CR	24.188	24.730	-23.079	-54.393	16.555
	FH	25.945	30.294	-21.063	-57.620	17.467
	NH	**21.018**	16.080	-26.116	-47.991	**14.224**
post-COVID	RM	9.682	13.042	-20.009	**-22.213**	7.393
	YR	10.258	16.528	-21.518	-23.429	7.825
	CR	10.622	16.032	-20.066	-23.757	8.499
	FH	**8.687**	**7.779**	**-19.371**	-22.543	**6.155**
	NH	15.260	23.885	-26.785	-27.775	12.699
Overall	RM	**14.063**	15.375	**-19.929**	**-33.671**	**9.711**
	YR	14.147	**13.552**	-21.493	-34.726	9.745
	CR	15.393	16.025	-20.910	-37.874	10.664
	FH	16.476	13.941	-22.214	-41.472	10.904
	NH	15.367	16.511	-24.762	-35.490	11.191

Note: Bold font indicates the best metric.

All values are expressed in thousands.

Table 5.4: Evaluation of hedging strategies through monthly volatility metrics over different time periods

Chapter 6

General Conclusion

This thesis conducts a comprehensive examination of the challenges and opportunities inherent in credit portfolio strategies, integrating theoretical frameworks, advanced quantitative methodologies, and practical applications. By addressing critical issues in contemporary credit markets, this research sheds light on the complex interactions of corporate actions, asset ranking, and portfolio optimization. The findings offer a dual contribution: advancing the academic understanding of credit markets while providing actionable insights for practitioners, thereby bridging the gap between theory and real-world implementation.

6.1 Key findings and implications for credit trading strategies

Through the exploration of corporate actions, long-short portfolio optimization, and a unified asset ranking and allocation framework, the research provides innovative methodologies and actionable insights for systematic credit trading. Key contributions are summarized as follows:

- **Corporate Actions and Credit Risk Dynamics**: Corporate actions such as stock buybacks, cash dividends, and mergers and acquisitions play a pivotal role in influencing credit risk, as reflected in CDS spreads. Chapter 2 provides a nuanced exploration of these relationships, emphasizing the importance of understanding temporal dynamics. By identifying distinct market responses before, during, and after corporate action announcements, the study demonstrates how these events alter credit perceptions and pricing. The findings highlight the dual influence of wealth transfer and signaling effects in shaping the trajectory of CDS spreads. For instance, stock buybacks generate divergent credit market reactions depending on firm size and market conditions, underscoring the need for tailored strategies when incorporating these events into trading models. Cash dividend changes exhibit pronounced immediate and anticipatory effects, while mergers and acquisitions reveal asymmetrical impacts on buyers and sellers. These results offer critical insights for asset managers to refine trading universe definitions by accounting for periods of volatility linked to corporate actions, thereby enhancing both the precision and reliability of quantitative trading strategies.
- **Enhancing Long-Short Portfolios**: Chapter 3 addresses a significant methodological gap in the

application of learn-to-rank (LTR) algorithms to long-short portfolio strategies. Traditional LTR frameworks often prioritize top-ranked assets, neglecting the accurate ranking of underperforming assets that are crucial for short positions. This inherent asymmetry creates challenges in maintaining market-neutral strategies, where balanced performance across long and short portfolios is essential. The research introduces modifications to widely used LTR loss functions, including ListNet, ListMLE, and ListFold, to address these limitations and ensure a more equitable evaluation of the entire asset ranking spectrum. Empirical validation confirms that these refinements significantly enhance both ranking precision and portfolio performance. The ability to improve returns while maintaining market neutrality underscores the practical value of these innovations, providing a robust framework for long-short strategies in credit markets and beyond.

- **Unified Asset Ranking and Portfolio Weighting**: Chapter 4 introduces a multi-task neural network framework that integrates asset ranking and portfolio allocation. This unified approach challenges the traditional sequential model, where ranking predictions and portfolio construction are treated as distinct steps. By combining these tasks within a single framework, the model captures the interdependencies between ranking accuracy and allocation decisions, producing more coherent and robust investment outcomes. Empirical analysis validates the effectiveness of this integrated approach, demonstrating its ability to outperform traditional methods. The streamlined decision-making process not only improves computational efficiency but also enhances alignment between predictive modeling and real-world trading objectives. For systematic credit trading, this represents a transformative advancement, offering asset managers a tool to optimize credit portfolios with greater precision and adaptability.

- **Decomposing Equity-Corporate Bond Correlation and Adaptive Interest Rate Hedging Strategies**: Chapter 5 introduces a comprehensive framework for decomposing the equity-corporate bond correlation to determine whether interest rate risk or credit risk is the dominant driver over time. Our findings reveal that the underlying drivers of this correlation evolve dynamically, shaped by shifting macroeconomic conditions such as growth and inflation uncertainties. By identifying these structural changes, we provide a deeper understanding of how macroeconomic forces impact the interaction between equity and corporate bond markets. Furthermore, we investigate how corporate bonds can be effectively hedged against interest rate risk, a critical consideration for capital structure arbitrage strategies that rely on isolating credit risk to exploit mispricings between a firm's equity and debt instruments. To address this challenge, we propose two novel hedging approaches that dynamically adjust to changing market conditions. Empirical results confirm that these adaptive strategies significantly outperform traditional methods, offering more robust and responsive risk management solutions for systematic credit trading.

This thesis offers significant contributions to both industry practitioners and academic research. It provides a set of practical tools designed to enhance decision-making and optimize portfolio performance, addressing the challenges of contemporary credit markets. By exploring the impact of corporate actions—such as stock buybacks, dividends, and mergers—on credit risk, the research equips asset managers with actionable insights to better understand market shifts and adjust strategies accordingly. Additionally, the thesis advances asset ranking techniques, particularly through innovations in learn-to-rank (LTR) algorithms, enabling more accurate resource allocation and improved risk-adjusted returns. These

tools are particularly crucial in today's data-driven, electronic markets, where swift, complex decision-making demands sophisticated modeling approaches. Moreover, the thesis introduces novel methodologies for understanding the dynamic relationship between equity and corporate bond markets, including the decomposition of the equity-corporate bond correlation and the development of adaptive interest rate hedging strategies, offering asset managers more effective means for managing interest rate risk within capital structure arbitrage strategies.

From an academic perspective, the thesis fills key gaps in the literature, particularly in understanding the interplay between corporate actions and credit risk, as well as optimizing LTR methodologies for long-short portfolio strategies. It provides a deeper theoretical framework for analyzing how corporate events influence credit risk over time, enriching the academic discourse on this topic. Furthermore, the innovative LTR techniques introduced push the boundaries of asset ranking, offering a more comprehensive and precise framework for evaluating both long and short positions in systematic trading. The thesis also makes a significant theoretical contribution by introducing new methodologies to decompose the equity-corporate bond correlation and develop adaptive interest rate hedging strategies. These advancements deepen our understanding of how macroeconomic factors shape the relationships between equity and credit markets.

In bridging the gap between theory and practice, this research offers both foundational theoretical insights and directly applicable solutions for industry practitioners. Its focus on adaptability and efficiency addresses current challenges in credit trading while ensuring that the methodologies remain relevant and scalable as market conditions evolve. By applying data science to credit portfolio strategies, particularly in an era of rapid technological advancement, this thesis makes a valuable contribution to the ongoing development of both financial markets and academic research.

6.2 Personal insights and challenges on applying data science and machine learning to credit trading

My experience during my PhD and as a quantitative researcher at Drakai Capital has provided a unique vantage point to explore the interplay between finance, data science, and machine learning in understanding credit markets. These fields, when combined, reveal transformative potential for enhancing the predictive power, interpretability, and robustness of credit research. This interdisciplinary approach is particularly crucial in the context of credit trading, a domain undergoing rapid electronification. As traditional trading methods evolve into electronic platforms, the resulting explosion of data offers unparalleled opportunities to identify patterns, manage risk, and uncover alpha through modern analytical techniques.

However, bridging these disciplines is far from straightforward. Financial data is inherently noisy, endogenous, and non-stationary, presenting substantial challenges to extracting meaningful insights. As highlighted in [83], issues such as model overfitting, specification errors, and p-hacking are not abstract academic concerns—they have real consequences, leading to flawed strategies or false discoveries if not addressed with rigor. Mitigating these risks requires disciplined scientific practices, including robust

cross-validation, out-of-sample testing, and stress testing across various market regimes to ensure the reliability and generalizability of findings.

Credit trading itself presents additional complexities. Credit instruments, such as corporate bonds and credit default swaps, often exhibit low liquidity, wide bid-ask spreads, and skewed return distributions. These characteristics challenge both traditional and machine learning-based models, particularly those that assume normality or stationarity in the data. In my research and professional work, I have found that overcoming these hurdles requires blending domain expertise with sophisticated feature engineering and deploying machine learning approaches designed to handle irregular data distributions. Methods such as robust loss functions, Bayesian techniques, and ensemble models can be particularly effective in dealing with these unique challenges.

Another key insight is the value of interpretability in machine learning models for credit trading. The field is influenced by complex microeconomic factors, such as corporate actions, and macroeconomic drivers, such as interest rate movements and credit cycles. Machine learning models that are black-box in nature can obscure valuable insights and hinder their practical adoption in trading environments. Techniques that improve model transparency—such as Shapley values or interpretable algorithms such as decision trees—empower traders to scrutinize predictions, understand underlying drivers, and validate the reliability of model outputs. This is essential for maintaining confidence in strategies, especially in a domain where sudden shifts in sentiment or unexpected market events are common.

The electronification of credit markets, while creating a wealth of data, also highlights the limitations of purely automated systems. My professional experience has shown that fully automating credit trading strategies remains a significant challenge due to the market's inherent complexity and the contextual nuances that require human judgment. Hybrid approaches, combining algorithmic precision with trader expertise, are often more effective in navigating the unpredictable and dynamic nature of credit markets. Monitoring quantitative models and continuously refining their assumptions is crucial to ensuring they remain robust under real-world conditions.

These insights reflect the broader potential of interdisciplinary approaches to transform credit trading while also underscoring the importance of transparency, adaptability, and collaboration between quantitative researchers and domain experts. By leveraging advanced methods alongside human expertise, it becomes possible to build strategies that are not only predictive but also resilient, fostering deeper trust in the use of machine learning within this complex and evolving field.

Chapter 7

Résumé en français

Le marché du crédit a longtemps été caractérisé par son opacité et sa dépendance aux transactions de gré à gré (*over-the-counter* ou OTC), où les échanges se faisaient principalement via des accords bilatéraux et des négociations téléphoniques. Cette structure favorisait un paradigme "Banks-versus-All", dans lequel les banques occupaient une position dominante, contrôlant à la fois le flux d'informations et la liquidité du marché. La crise financière de 2008 a constitué un tournant, les régulateurs du monde entier ayant mis en place des réformes visant à accroître la transparence du marché, promouvoir la standardisation, et accélérer la digitalisation des processus de trading. Ces efforts ont abouti à une "électronification" progressive du marché du crédit, facilitant la transition vers un modèle de trading "All-to-All" qui reflète la structure plus transparente des marchés actions.

Cette évolution structurelle a créé de nouvelles opportunités d'innovation, notamment grâce à l'application des méthodes de recherche quantitative à l'investissement en crédit. Avec la multiplication des plateformes électroniques et la forte augmentation des données disponibles, les acteurs du marché du crédit peuvent désormais s'appuyer sur des technologies avancées afin d'optimiser la prise de décision et élaborer des stratégies plus performantes. Dans ce contexte, l'intégration de l'intelligence artificielle (*artificial intelligence* ou AI) et de l'apprentissage automatique (*machine learning* ou ML) a rendu possibles des approches systématiques d'investissement en crédit, qui étaient auparavant limitées par le manque de données fiables et l'opacité du marché.

Perspectives industrielles

Drakai Capital se place au cœur de cette transformation, tirant parti des avancées technologiques pour répondre aux enjeux majeurs du marché du crédit. Son programme de recherche se concentre plus particulièrement sur trois axes d'innovation clés :

- **Opportunités d'arbitrage *market neutral*** : Drakai Capital cherche à identifier et exploiter des opportunités d'arbitrage sur le marché du crédit grâce à des modèles d'AI. Ces modèles intègrent une grande diversité de données, incluant des indicateurs financiers fondamentaux, des signaux

d'analyse technique, des variables macroéconomiques ainsi que des sources de données alternatives. L'objectif est d'adopter une position *market neutral*, garantissant des rendements non directionnels et décorrélés des fluctuations du marché, minimisant ainsi l'exposition au risque.

- **Diversification de portefeuille et gestion du risque** : Drakai Capital privilégie la construction de portefeuilles de crédit diversifiés, en alliant expertise métier et algorithmes de ML. Ces portefeuilles ont pour objectif d'optimiser le rendement ajusté au risque tout en garantissant stabilité et faible volatilité. Grâce à l'intégration des outils de ML, les portefeuilles peuvent être ajustés en continu, ce qui améliore leur capacité à réagir rapidement aux fluctuations du marché et offre ainsi aux investisseurs des avantages significatifs en matière de diversification.
- **Analyse de la liquidité et exécution optimale** : La liquidité sur le marché du crédit demeure fragmentée, avec une forte hétérogénéité des prix selon les instruments. Drakai Capital mène des analyses approfondies sur la dynamique de la liquidité afin de concevoir des stratégies d'exécution optimales. En comprenant l'évolution structurelle de cette liquidité et l'impact des nouveaux acteurs sur le marché, Drakai Capital est en mesure de limiter les coûts de transaction ainsi que la déviation des prix d'exécution, dans un environnement de trading de plus en plus électronique.

La réussite dans ces domaines nécessite de surmonter plusieurs défis scientifiques et techniques. Drakai Capital a pour ambition de mobiliser des techniques quantitatives de pointe ainsi que des innovations technologiques pour répondre aux enjeux structurels du trading sur le marché du crédit. Au cœur de cette démarche se trouve le développement de modèles statistiques avancés, capables de capturer les dépendances complexes entre actifs financiers, notamment en période de stress de marché. Des épisodes tels que la crise des subprimes en 2008 ou la crise COVID de 2020 ont mis en évidence la nécessité de modéliser les corrélations de manière dynamique, l'interconnexion accrue des actifs ayant fortement contribué à l'amplification des risques systémiques. Drakai Capital accorde par ailleurs une importance particulière à l'intégration de l'AI et du ML dans le processus décisionnel en matière d'investissement, tout en s'attachant à résoudre des problématiques cruciales telles que l'interprétabilité des modèles, leur robustesse et leur utilisation dans des conditions réelles de marché.

Cette thèse, réalisée dans le cadre d'une convention industrielle de formation par la recherche (CIFRE) entre Drakai Capital et le Laboratoire d'Économie d'Orléans (LEO), aborde les principaux défis liés à la conception de stratégies de ML appliquées au trading de produits de crédit. En établissant un lien entre la recherche académique et les applications concrètes du secteur financier, ce travail contribue à l'avancement des approches quantitatives dans l'investissement en crédit, tout en nourrissant la réflexion académique. Les résultats issus de cette recherche ont vocation à éclairer tant les pratiques industrielles que les cadres réglementaires, participant ainsi à la construction d'un marché du crédit plus efficient et plus résilient.

Étapes clés du développement de modèles de trading quantitatif

Le développement de stratégies de trading quantitatives suit généralement un processus itératif. Cependant, chaque gestionnaire de portefeuilles adopte sa propre méthodologie afin de concevoir ses stratégies

systématiques. On peut ainsi dire qu'il existe autant de méthodologies que de gestionnaires, chacune reflétant des philosophies, des outils et des objectifs différents. Cette section décrit, selon moi, l'approche la plus couramment observée dans les sociétés de gestion, afin d'illustrer un processus typique dans le développement de stratégies quantitatives (voir Figure 7.1). Cette thèse vise ensuite à améliorer plusieurs étapes clés de ce processus, notamment la définition de l'univers d'investissement, la conception des modèles quantitatifs et la construction du portefeuille, dans le but d'en renforcer à la fois l'efficacité opérationnelle et le pouvoir prédictif.

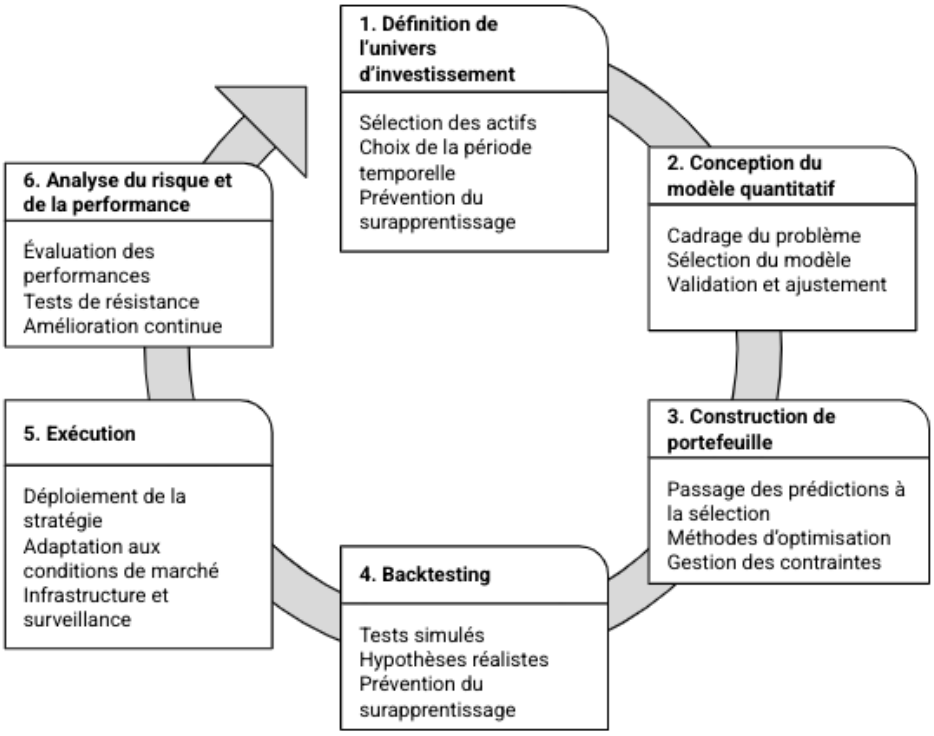

Figure 7.1: Processus de développement d'une stratégie de trading quantitatif

La première étape dans le développement d'une stratégie d'investissement quantitative consiste à définir l'univers de trading. Cela implique de sélectionner les instruments à analyser et à inclure dans le portefeuille, ainsi que de déterminer les périodes temporelles pertinentes au cours desquelles ces actifs seront tradés. Ces choix sont déterminants, car l'univers ainsi défini conditionne directement le jeu de données utilisé pour l'entraînement du modèle quantitatif, influençant ses performances et sa capacité à généraliser sur des données non observées. Par exemple, la sélection d'instruments de crédit peut se concentrer sur certaines maturités de contrats CDS, reconnues pour leur liquidité élevée et leur comportement relativement stable dans le temps. Ce type de choix facilite l'apprentissage du modèle, car des jeux de données homogènes et bien structurés favorisent une meilleure capacité de généralisation. Cependant, un filtrage trop restrictif peut s'avérer contre-productif. En effet, écarter un trop grand nombre de périodes temporelles ou d'actifs peut entraîner un problème de surapprentissage (*overfitting*), ce qui nuit à la capacité du modèle à s'adapter aux conditions réelles de marché. Le surapprentissage constitue un écueil majeur en finance quantitative. Il survient lorsqu'un modèle est trop ajusté aux données historiques, au point de perdre en efficacité lorsqu'il est confronté à des données nouvelles. Il est donc fondamental de définir l'univers de trading avec rigueur, afin de construire un modèle à la fois robuste, flexible et capable

d'être efficient sur une diversité de situations de marché.

La deuxième étape du développement d'une stratégie de trading quantitative consiste à concevoir le modèle quantitatif, en commençant par une définition précise du problème à résoudre. Il s'agit de déterminer la variable à prédire — par exemple, la prévision des rendements d'actifs sur un horizon temporel donné — ainsi que l'objectif du modèle, qu'il s'agisse d'une tâche de classification ou de régression. Une fois le problème posé, le choix de l'algorithme devient crucial. Selon la nature de la tâche, on peut privilégier des méthodes simples comme la régression linéaire, reconnue pour sa facilité d'interprétation, des modèles basés sur des arbres de décision permettant de modéliser des relations non linéaires, ou encore des approches plus complexes, telles que les réseaux de neurones, capables de détecter des liens plus sophistiquées dans les données. La phase suivante vise à entraîner efficacement le modèle. Pour cela, des techniques comme la validation croisée sont souvent employées, afin d'éviter le surapprentissage et garantir un bon compromis entre biais et variance. L'ajustement des hyperparamètres du modèle est également essentiel, renforçant sa robustesse. Ce processus forme le fondement de la conception de modèles prédictifs capables de détecter des tendances et de produire des signaux de trading pertinents.

La troisième étape concerne la construction du portefeuille, qui consiste à traduire les prédictions du modèle en allocations d'actifs concrètes. Dans les sociétés de gestion, cette phase est souvent séparée du processus de modélisation. Dans un premier temps, le modèle entraîné génère des résultats prédictifs, tels que des classements ou des estimations de rendement, qui servent de base aux méthodes de construction du portefeuille. Ces dernières permettent de déterminer une allocation optimale, en cohérence avec les prévisions du modèle. Ce processus s'appuie généralement sur des heuristiques ou des méthodes d'optimisation destinées à trouver un équilibre entre les rendements espérés et le profil de risque du portefeuille. La stratégie d'allocation intègre plusieurs contraintes, telles que les préférences des investisseurs, les limites de risque, ainsi que les exigences réglementaires. Cela garantit que le portefeuille final reflète non seulement les objectifs stratégiques du gestionnaire, mais respecte aussi les cadres normatifs en vigueur.

La quatrième étape est le backtesting, une phase essentielle qui permet d'évaluer à la fois le modèle et la construction du portefeuille avant leur mise en œuvre en conditions réelles. Le backtesting consiste à simuler la stratégie sur des données historiques afin d'en estimer les performances passées. Ce procédé offre des informations précieuses sur les points forts et les limites potentielles de la stratégie, en s'appuyant sur des indicateurs clés tels que les rendements, la volatilité, les pertes maximales (*max drawdowns*) et autres mesures ajustées du risque. De plus, le backtesting permet de tester la robustesse de la stratégie en examinant ses résultats sur différentes périodes, régimes de marché et scénarios de stress, garantissant ainsi sa capacité à rester performante dans diverses conditions. Pour que les résultats soient fiables, il est indispensable d'intégrer des hypothèses réalistes de trading, comme les coûts de transaction et la déviation des prix d'exécution. Ces facteurs peuvent fortement impacter les résultats et donnent une vision plus fidèle du comportement de la stratégie en situation réelle. Néanmoins, un piège courant du backtesting est le risque de surapprentissage : le modèle peut très bien fonctionner sur les données historiques, mais échouer à s'adapter à de nouvelles données futures. Pour limiter ce risque, il est crucial d'utiliser des techniques de validation, telles que les tests hors échantillon (*out-of-sample*) *walk-forward*. Cette méthode permet de vérifier que la performance obtenue n'est pas le fruit d'un simple "*data-mining*" ou d'un surapprentissage, assurant ainsi une évaluation fiable et réaliste du véritable potentiel de la stratégie

en conditions de trading réel.

La cinquième étape dans le développement d'une stratégie de trading quantitative est l'exécution, une phase essentielle où la stratégie passe de la conception théorique à la mise en œuvre concrète sur les marchés. Il s'agit alors de traduire les allocations d'actifs définies en ordres de marché réels, en s'assurant qu'ils soient exécutés de manière efficiente et à un coût maîtrisé. L'exécution est particulièrement déterminante sur les marchés du crédit, où les pratiques de négociation ont connu d'importantes évolutions. L'essor des plateformes de trading électronique a profondément transformé la liquidité, la transparence et l'accessibilité, rendant l'exécution de stratégies systématiques sur les marchés obligataires plus fluide et plus efficace. Dans ce contexte, une exécution réussie ne se limite pas à l'envoi d'ordres : elle repose sur l'intégration de données de marché en temps réel et sur une infrastructure technologique robuste. Des systèmes à faible latence sont essentiels afin de garantir une réactivité optimale. Enfin, un suivi continu de la performance d'exécution s'impose pour mesurer la qualité des transactions effectuées, identifier d'éventuels écarts par rapport aux coûts attendus et évaluer l'impact sur le marché. Ce suivi constitue un mécanisme essentiel d'amélioration continue, permettant de réviser les étapes antérieures de la stratégie et d'en affiner l'exécution au fil du temps, afin de préserver la cohérence avec les objectifs de performance.

La dernière étape dans le développement d'une stratégie de trading quantitative est l'analyse du risque et de la performance. Elle vise principalement à vérifier que les résultats obtenus en conditions réelles sont cohérents avec les performances théoriques observées lors du backtesting. Cette phase permet d'identifier d'éventuels écarts de performance dès le début de la mise en production, et d'évaluer si le modèle s'éloigne de ses résultats attendus. Elle joue un rôle diagnostique essentiel, en aidant à détecter une éventuelle dérive du modèle, des hypothèses de backtesting erronées, ou encore un mauvais fonctionnement du modèle en production. Cette analyse permet également de mettre en lumière certaines faiblesses, telles qu'une sous-performance dans des régimes de marché spécifiques ou une instabilité face à l'évolution des conditions de marché. Un suivi continu est indispensable : il permet de détecter toute dérive du modèle, de recalibrer ses paramètres et d'ajuster les contraintes au fur et à mesure des changements de l'environnement de marché. Cette approche diagnostique garantit que le modèle reste conforme à ses objectifs initiaux, en autorisant des ajustements rapides et ciblés pour préserver sa performance à long terme.

En conclusion, développer une stratégie de trading quantitative robuste et performante est un processus itératif où chaque étape joue un rôle crucial dans le succès global. De la définition de l'univers de trading à la conception des modèles prédictifs, en passant par la construction du portefeuille, le backtesting, l'exécution et l'évaluation des performances, toutes ces phases sont étroitement liées. La moindre faille peut compromettre l'ensemble, soulignant la complexité à bâtir un système à la fois cohérent et efficace. En finance quantitative, le véritable défi n'est pas seulement d'exécuter une stratégie bien conçue, mais surtout de créer un système pertinent, capable de délivrer des résultats solides en backtest comme sur les marchés réels. Cela implique de trouver un équilibre subtil entre pouvoir prédictif et robustesse, pour que la stratégie s'adapte aux fluctuations des marchés tout en générant une performance régulière. Cette exigence est d'autant plus complexe que les marchés sont instables et soumis à de nombreux facteurs imprévisibles. Concevoir un système capable de préserver sa performance dans un tel contexte reste un défi majeur. Si les outils avancés, comme l'apprentissage statistique, offrent de puissantes capacités d'analyse, ils ne garantissent pas le succès. Nombreuses sont les stratégies prometteuses sur le papier

qui échouent en conditions réelles, souvent victimes de surapprentissage ou d'une mauvaise adaptation aux évolutions du marché. C'est pourquoi construire un système durable demande non seulement une compréhension approfondie des mécanismes financiers, mais aussi une grande minutie dans la conception et l'assemblage de chacune des différentes briques du système.

Cette thèse a pour ambition de répondre à plusieurs défis techniques en proposant des solutions innovantes, destinées à renforcer la robustesse et l'adaptabilité des stratégies quantitatives crédit. Elle offre un regard approfondi sur trois aspects clés : la sélection de l'univers d'investissement, la conception des modèles statistiques, ainsi que les techniques de construction de portefeuille. L'objectif est de fournir aux gestionnaires d'actifs des méthodes concrètes et efficaces, leur permettant d'améliorer la fiabilité de leurs systèmes de trading quantitatif et d'assurer leur réussite durable face aux complexités des marchés financiers.

Contributions à l'amélioration des stratégies de trading quantitatives

Cette thèse, organisée en quatre chapitres, analyse les principaux défis auxquels Drakai Capital est confronté dans la pratique. Elle apporte ainsi un éclairage approfondi sur les complexités des marchés de crédit modernes, des enseignements qui pourront également bénéficier aux gestionnaires d'actifs et aux autorités de régulation.

Le chapitre 2, intitulé *« Corporate Actions and Credit Risk Dynamics: Evidence from Stock Buybacks, Cash Dividends and Mergers & Acquisitions »*, a pour objectif d'améliorer la première étape *« Définition de l'univers de trading »* présentée dans la Figure 7.1. Plus précisément, il s'intéresse à l'impact des annonces d'opérations sur titres sur les *spreads* de CDS, qui connaissent souvent des fluctuations anormales autour de ces événements. Ces variations — souvent considérées comme des valeurs aberrantes — peuvent fausser la distribution des données, introduire du bruit et détériorer la performance des modèles d'apprentissage automatique lorsqu'elles sont prises en compte durant l'entraînement. L'étude se concentre sur deux aspects majeurs. D'une part, elle décrit les schémas temporels des ajustements des *spreads* suite à l'annonce d'événements clés tels que les rachats d'actions, les fusions-acquisitions et les versements de dividendes. D'autre part, elle cherche à expliquer les mécanismes sous-jacents à ces réactions, en distinguant les effets de richesse — liés aux changements du bilan financier — des effets de signal, qui reflètent une modification de la perception de la solvabilité de l'entreprise. En examinant ces mécanismes, le chapitre met en évidence les caractéristiques spécifiques des entreprises et des événements qui amplifient ou atténuent les réactions du risque de crédit. Sur le plan pratique, cette analyse a des implications directes pour les système développés par Drakai Capital. Comprendre le comportement anormal des *spreads* de crédit pendant ces plages temporelles permet d'affiner les stratégies de trading de CDS, en évitant les périodes de forte volatilité qui pourraient fausser les performances des modèles. Cela permet notamment aux gestionnaires d'actifs de définir un univers de trading excluant ces périodes instables, garantissant ainsi que les modèles d'apprentissage automatique soient entraînés sur des données plus représentatives et fiables. Par ailleurs, ces connaissances fournissent aux gestionnaires des outils pour mieux anticiper et interpréter les réactions du marché de crédit, améliorant ainsi la prise de décision et

l'exécution autour des annonces d'opérations sur titres. Enfin, les contributions de ce chapitre vont au-delà de l'analyse individuelle de chaque opération : il propose un cadre comparatif permettant d'évaluer différents types d'événements au sein d'une méthodologie unifiée, comblant ainsi certaines lacunes de la littérature existante. Il souligne également l'importance de considérer à la fois le calendrier des événements et les particularités propres à chaque entreprise, mettant en lumière la diversité des réactions face au risque de crédit.

Le chapitre 3 *Enhancing Long-Short Portfolios: A Refined Approach using Learn-to-Rank Algorithms* correspond à l'étape *« 2. Conception du modèle quantitatif »* présentée dans la Figure 7.1. Il s'intéresse aux asymétries inhérentes aux fonctions de perte classiques des méthodes *Learn-to-Rank* (LTR) qui, bien qu'efficaces dans les tâches de recherche d'information, présentent des limites notables lorsqu'elles sont appliquées à la construction de portefeuilles *long-short*. En effet, les approches LTR traditionnelles ont tendance à privilégier la précision dans le haut du classement, en négligeant souvent l'importance d'identifier correctement les actifs sous-performants en bas de classement. Ce déséquilibre constitue un enjeu crucial dans le cadre des stratégies long-short, où les positions longues et courtes jouent un rôle tout aussi essentiel dans la génération de rendement et le maintien de la neutralité de marché. À travers une combinaison d'analyses théoriques et de preuves empiriques, ce chapitre identifie et quantifie ces biais, mettant en lumière leur impact négatif sur la performance des portefeuilles. Sur cette base, nous proposons des modifications adaptées à trois fonctions de perte LTR largement utilisées — ListNet, ListMLE et ListFold — afin de répondre spécifiquement aux exigences duales de classement des stratégies *long-short*. Ces ajustements visent à améliorer la précision du classement sur l'ensemble de la distribution des actifs, garantissant une évaluation plus symétrique et cohérente des actifs surperformants comme sous-performants. La validation empirique, réalisée sur des données hebdomadaires d'actions A-shares chinoises, démontre l'efficacité de ces modifications, avec des améliorations significatives tant en termes de performance de classement que de rendement des portefeuilles. Bien que les tests empiriques de ce chapitre se concentrent sur des données actions dans un cadre académique, les méthodologies et enseignements développés sont directement applicables aux stratégies de trading obligataire de Drakai Capital. Plus précisément, ces fonctions de perte affinées fournissent un cadre robuste pour améliorer la précision des classements et optimiser la construction de portefeuilles *long-short* sur les marchés du crédit.

Le chapitre 4 *Unifying Asset Ranking and Portfolio Weighting through a Multi-Task Neural Network* présente un système novateur qui combine classement des actifs et allocation de portefeuille au sein d'un modèle unique, grâce à un réseau de neurones multitâche. En référence à la Figure 7.1, ce travail vise à fusionner les étapes *« 2. Conception du modèle quantitatif »* et *« 3. Construction du portefeuille »*. Traditionnellement, les gestionnaires quantitatifs de portefeuilles utilisent un processus en deux temps pour construire un portefeuille : d'abord, des modèles de classement prédisent la performance relative des actifs, puis les poids de portefeuille sont déterminés par des règles heuristiques ou des méthodes d'optimisation. Bien que cette approche séquentielle soit efficace, elle ne prend pas toujours en compte les interdépendances entre les prédictions de classement et les décisions d'allocation, ce qui peut limiter la performance globale. En réponse, ce chapitre propose une méthodologie intégrée qui optimise conjointement ces deux tâches au sein d'un même cadre d'apprentissage. En intégrant directement le processus d'optimisation dans le système d'apprentissage, le modèle saisit les interactions entre classement et allocation, conduisant à des décisions d'investissement plus cohérentes et robustes. Même si la validation empirique présentée repose sur des données actions dans un cadre académique, ces résultats sont là aussi

naturellement extensibles aux marchés obligataires. Cette intégration offre à Drakai Capital un outil évolutif et adaptable afin d'optimiser son système de trading, alignant ainsi la modélisation prédictive avec les objectifs concrets de la construction de portefeuilles.

Le chapitre 5 *Understanding the Equity-Corporate Bond Nexus: A Framework for Risk Decomposition and Interest Rate Hedging* correspond à l'étape *« 3. Construction du portefeuille »* dans la Figure 7.1. Il explore les interactions entre les marchés actions et obligataires en décomposant leur corrélation en deux composantes : le risque de crédit et la sensibilité aux taux d'intérêt. Partant du constat que les *spreads* de CDS et les taux des obligations souveraines reflètent des dimensions distinctes du risque obligataire, ce chapitre propose un cadre innovant permettant d'isoler ces deux effets. Plus précisément, il utilise la corrélation entre les actions et les CDS pour capter la sensibilité au risque de crédit, et celle entre les actions et les obligations souveraines pour représenter l'exposition aux taux d'intérêt. Ce cadre permet ainsi de mieux comprendre les moteurs fondamentaux de la corrélation entre actions et obligations d'entreprise selon les régimes de marché. L'analyse empirique révèle des variations marquées de l'importance relative de ces deux composantes selon la qualité de crédit et le contexte économique. En période de stabilité, les obligations *investment grade* sont davantage sensibles aux mouvements de taux d'intérêt, tandis que les titres *high yield* réagissent principalement aux fluctuations du risque de crédit. En revanche, lors de périodes de forte tension — comme la crise du COVID — le risque de crédit devient prépondérant pour l'ensemble des obligations, éclipsant temporairement l'impact des taux. Après la crise, avec le resserrement monétaire des banques centrales, la sensibilité aux taux retrouve de l'importance, en particulier pour les titres bien notés. Du point de vue du trading, ces résultats ont des implications concrètes pour la gestion du risque et les stratégies de couverture de Drakai Capital. Comprendre l'évolution du poids relatif des risques de crédit et de taux dans le temps permet d'ajuster plus finement les techniques de couverture. À cette fin, le chapitre introduit et évalue deux stratégies adaptatives — le *Yield Ratio* et le modèle de régression — qui ajustent dynamiquement les ratios de couverture en fonction des conditions de marché. Ces approches surpassent les méthodes statiques classiques, en particulier dans des environnements marqués par des mouvements rapides des taux ou des *spreads* de crédit. Au-delà de son utilité pour Drakai Capital, ce travail contribue plus largement à la littérature en gestion d'actifs et en gestion des risques, en proposant une approche systématique de la décomposition du risque obligataire et de l'optimisation des stratégies de couverture de taux d'intérêt.

Pris dans leur ensemble, ces contributions soulignent l'importance de combiner une expertise métier approfondie avec des techniques quantitatives rigoureuses afin d'améliorer les stratégies de trading sur les marchés de crédit modernes. En s'attaquant à des enjeux clés liés à la sélection d'actifs, à la conception de modèles statistiques et à la construction de portefeuilles, cette thèse ne se contente pas seulement d'affiner les méthodologies existantes — elle ouvre également la voie à des stratégies systématiques plus adaptatives et résilientes. Ces avancées constituent ainsi une base solide, à la fois pour la recherche académique et pour les applications concrètes, en fournissant à Drakai Capital et à l'ensemble de la communauté de gestionnaires de portefeuilles des outils exploitables pour évoluer dans un environnement financier toujours plus complexe.

Réflexions personnelles et défis liés à l'application de la science des données et de l'apprentissage automatique au trading de crédit

Mon parcours, à la fois durant ma thèse et au sein de Drakai Capital en tant que chercheur quantitatif, m'a permis d'adopter un regard privilégié sur les synergies entre finance, science des données et apprentissage automatique dans l'analyse des marchés du crédit. L'association de ces disciplines ouvre la voie à des avancées majeures en matière de prédiction, d'interprétation et de robustesse des analyses quantitatives. Cette approche interdisciplinaire s'avère d'autant plus pertinente dans un environnement de trading de crédit en pleine mutation, marqué par une transition rapide vers l'"électronification". L'émergence de plateformes électroniques et la croissance exponentielle des données qui en résulte offrent aujourd'hui des opportunités uniques pour détecter des signaux, affiner la gestion des risques et exploiter de nouvelles sources de performance à l'aide de méthodes analytiques modernes.

Toutefois, l'articulation entre finance, data science et apprentissage automatique reste particulièrement complexe. Les données financières se caractérisent par leur bruit, leur endogénéité et leur non-stationnarité, rendant l'extraction d'informations fiables d'autant plus ardue. Comme le rappelle [83], des écueils tels que le surapprentissage, les erreurs de spécification ou le "*p-hacking*" ne relèvent pas de simples considérations théoriques : leurs effets concrets peuvent compromettre la validité d'une stratégie ou conduire à des découvertes trompeuses (faux positifs) si ces risques ne sont pas rigoureusement maîtrisés. Pour y remédier, il est indispensable d'adopter une démarche scientifique exigeante, fondée sur des tests hors échantillon, ainsi que des simulations en conditions de marché variées, afin d'assurer la robustesse et la capacité de généralisation des résultats.

Le trading de crédit comporte des complexités supplémentaires. Les instruments de crédit, comme les obligations d'entreprise et les *credit default swaps*, se distinguent souvent par une faible liquidité, des écarts *bid-ask* importants et des distributions de rendements asymétriques. Ces caractéristiques représentent un défi tant pour les modèles statistiques traditionnels que pour ceux basés sur l'apprentissage automatique, en particulier lorsque ces derniers supposent que les données sont normales ou stationnaires. Au cours de mes recherches et de mon expérience professionnelle, j'ai constaté que surmonter ces difficultés nécessite de conjuguer une expertise financière approfondie avec une ingénierie des variables poussée, ainsi que le recours à des méthodes d'apprentissage automatique spécifiquement adaptées au traitement de distributions de données irrégulières. Des outils tels que les fonctions de perte robustes, les méthodes bayésiennes et les modèles ensemblistes se sont avérées particulièrement efficaces pour relever ces défis spécifiques.

Un autre point essentiel concerne l'importance de l'interprétabilité des modèles d'apprentissage automatique en trading de crédit. Ce domaine est soumis à des influences complexes, tant microéconomiques, comme les opérations sur titres, que macroéconomiques, telles que les variations des taux d'intérêt et les cycles de crédit. Les modèles dits "boîte noire" peuvent cacher des informations cruciales et limiter leur adoption concrète en contexte de trading. Des méthodes favorisant la transparence — comme les valeurs de Shapley ou des algorithmes facilement interprétables, tels que les arbres de décision — offrent aux gestionnaires de portefeuilles la possibilité d'analyser les prédictions, de saisir les facteurs sous-jacents et de vérifier la fiabilité des résultats. Ceci est indispensable pour préserver la confiance dans les stratégies, notamment dans un environnement où les changements soudains de sentiment ou les événements de

marché imprévus sont fréquents.

L'"électronification" des marchés de crédit, bien qu'elle génère un volume important de données, révèle aussi les limites des systèmes entièrement automatisés. Mon expérience professionnelle montre que l'automatisation totale des stratégies de trading de crédit demeure un défi majeur, en raison de la complexité intrinsèque de ce marché et des nuances contextuelles qui nécessitent un jugement humain. Les approches hybrides, combinant la précision des algorithmes à l'expertise des traders, se révèlent souvent plus efficaces pour appréhender la nature imprévisible et dynamique des marchés de crédit. Il est indispensable de surveiller constamment les modèles quantitatifs et de réviser leurs hypothèses afin d'assurer leur robustesse dans des conditions réelles.

Ces enseignements soulignent le fort potentiel des approches interdisciplinaires pour transformer le trading de crédit, tout en insistant sur l'importance de la transparence, de l'adaptabilité et de la collaboration entre chercheurs quantitatifs et experts du secteur. En combinant des méthodes avancées avec l'expertise humaine, on peut concevoir des stratégies à la fois efficaces et robustes, renforçant ainsi la confiance dans l'utilisation de l'apprentissage automatique au sein de ce domaine complexe et en constante évolution.

Bibliography

[1] Viral Acharya and Timothy C. Johnson. Insider trading in credit derivatives. *Journal of Financial Economics*, 84(1):110–141, 2007.

[2] Michael J. Alderson, Joseph T. Halford, and Valeriy Sibilkov. An examination of the wealth effects of share repurchases on bondholders. *Journal of Corporate Finance*, 65(C), 2020.

[3] I. Anagnostou, T. Squartini, D. Kandhai, and D. Garlaschelli. Uncovering the mesoscale structure of the credit default swap market to improve portfolio risk modelling. *Quantitative Finance*, 21(9):1501–1518, 2021.

[4] Christian Andres, André Betzer, and Markus Doumet. Measuring abnormal credit default swap spreads. *Global Finance Journal*, 2021. Forthcoming.

[5] Ivo J.M. Arnold. Internet search volumes of uk banks during the crisis: The role of banking structure and business model. *Global Finance Journal*, 45:100472, 2020.

[6] Paul Asquith, Robert F. Bruner, and David Jr. Mullins. The gains to bidding firms from merger. *Journal of Financial Economics*, 11(1-4):121–139, April 1983.

[7] Sanjai Bhagat, Ming Dong, David Hirshleifer, and Robert Noah. Do tender offers create value? new methods and evidence. *Journal of Financial Economics*, 76(1):3–60, 2005.

[8] Sudipto Bhattacharya. Imperfect information, dividend policy, and "the bird in the hand" fallacy. *The Bell Journal of Economics*, 10:259–270, 1979.

[9] Sudipto Bhattacharya. Nondissipative signaling structures and dividend policy. *Quarterly Journal of Economics*, 95:1–24, 1980.

[10] Matthew T. Billett, Dolly King, and David C. Mauer. Bondholder wealth effects in mergers and acquisitions: New evidence from the 1980s and 1990s. *SPGMI: Compustat Fundamentals (Topic)*, 2002.

[11] Ekkehart Boehmer, Jim Masumeci, and Annette B. Poulsen. Event-study methodology under conditions of event-induced variance. *Journal of Financial Economics*, 30(2):253–272, 1991.

[12] Dion Bongaerts, Frank De Jong, and Joost Driessen. Derivative Pricing with Liquidity Risk: Theory and Evidence from the Credit Default Swap Market. *Journal of Finance*, 66(1):203–240, February 2011.

[13] Alfie Brixton, Jordan Brooks, Peter Hecht, Antti Ilmanen, Thomas Maloney, and Nicholas McQuinn. A changing stock-bond correlation: Drivers and implications. *The Journal of Portfolio Management*, 2023.

[14] Stephen J. Brown and Jerold B. Warner. Measuring security price performance. *Journal of Financial Economics*, 8(3):205–258, September 1980.

[15] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd International Conference on Machine Learning*, ICML '05, page 89–96, New York, NY, USA, 2005. Association for Computing Machinery.

[16] Jeffrey L. Callen, Joshua Livnat, and Dan Segal. The impact of earnings on the pricing of credit default swaps. *The Accounting Review*, 84(5):1363–1394, 2009.

[17] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: From pairwise approach to listwise approach. *Proceedings of the 24th International Conference on Machine Learning*, 227:129–136, June 2007.

[18] CFM. Bond-equity correlations: Are the times a-changin'? *Whitepaper*, 2020.

[19] Fan Chen, Krishnan Ramaya, and Wei Wu. The wealth effects of merger and acquisition announcements on bondholders: New evidence from the over-the-counter market. *Journal of Economics and Business*, 107:105862, 2020.

[20] Guillaume Chevalier, Guillaume Coqueret, and Thomas Raffinot. Supervised portfolios. *Quantitative Finance*, 22:2275–2295, 2022.

[21] William Cready. Aggregate market reaction to earnings announcements. *Journal of Accounting Research*, 48:289–334, May 2010.

[22] Robert T. Daigler and Mark Cooper. A futures duration-convexity hedging method. *Financial Review*, 1998.

[23] L. Dann. Common stock repurchases: an analysis of returns to bondholders and stockholders. *Journal of Financial Economics*, 9:113–138, 1981.

[24] Laurence Daures-Lescourret and Andras Fulop. Standardization, transparency initiatives, and liquidity in the cds market. *Journal of Financial Markets*, 59:100718, 2022.

[25] Alan L. DelFavero. Stock buyback announcements: An examination of abnormal returns in stock price & credit default swaps for s&p100 companies. *Sacred Heart University, Jack Welch College of Business Fairfield, CT.*, 2018.

[26] Upinder S. Dhillon and H. E. Johnson. The effect of dividend changes on stock and bond prices. *Journal of Finance*, 49:281–289, 1994.

[27] Douglas W Diamond. Monitoring and Reputation: The Choice between Bank Loans and Directly Placed Debt. *Journal of Political Economy*, 99(4):689–721, August 1991.

[28] Hitesh Doshi, Jan Ericsson, Kris Jacobs, and Stuart M. Turnbull. Pricing Credit Default Swaps with Observable Covariates. *Review of Financial Studies*, 26(8):2049–2094, 2013.

[29] Brice Dupoyet, Xiaoquan Jiang, and Qianying Zhang. A new take on the relationship between interest rates and credit spreads. *Applied Economics*, 56(5):520–536, January 2024.

[30] Isil Erel, Yeejin Jang, and Michael S. Weisbach. Do Acquisitions Relieve Target Firms' Financial Constraints? *Journal of Finance*, 70(1):289–328, February 2015.

[31] Jan Ericsson, Kris Jacobs, and Rodolfo Oviedo. The determinants of credit default swap premia. *Journal of Financial and Quantitative Analysis*, 44(1):109–132, 2009.

[32] Mara Faccio and Ronald W. Masulis. The choice of payment method in european mergers and acquisitions. *Journal of Finance*, 60:1345–1388, 2005.

[33] Yoav Freund, Raj Iyer, Robert Schapire, and Yoram Singer. An efficient boosting algorithm for combining preferences. *Journal of Machine Learning Research*, 4:933–969, January 2003.

[34] Nils Friewald, Christian Wagner, and Josef Zechner. The cross-section of credit risk premia and equity returns. *The Journal of Finance*, 69, May 2013.

[35] Craig Furfine and R.J. Rosen. Mergers increase default risk. *Journal of Corporate Finance*, 17:832–849, September 2011.

[36] Koresh Galil, Offer Moshe Shapir, Dan Amiram, and Uri Ben-Zion. The determinants of CDS spreads. *Journal of Banking & Finance*, 41(C):271–282, 2014.

[37] Mila Getmansky, Giulio Girardi, and Craig Lewis. Interconnectedness in the cds market. *Financial Analysts Journal*, 72(4):62–82, July 2016.

[38] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*, 2010.

[39] Caitlin Greatrex. The credit default swap market's reaction to earnings announcements. *SSRN Electronic Journal*, 19, March 2008.

[40] Mariya Gubareva and Maria Rosa Borges. Switching interest rate sensitivity regimes of u.s. corporates. *The North American Journal of Economics and Finance*, 54:100888, 2020.

[41] Mariya Gubareva and Ilias Chondrogiannis. Capital Gains Sensitivity of US BBB-Rated Debt to US Treasury Market: Markov-Switching Analyses. *Complexity*, pages 1–13, August 2020.

[42] Sang gyung Jun, Mookwon Jung, and Ralph A. Walkling. Share repurchase, executive options and wealth changes to stockholders and bondholders. *Journal of Corporate Finance*, 15(2):212–229, 2009.

[43] Musa Gün. *The Co-Movement of Credit Default Swaps and Stock Markets in Emerging Economies*, pages 55–69. September 2018.

[44] George Handjinicolaou and Avner Kalay. Wealth redistributions or changes in firm value: An analysis of returns to bondholders and stockholders around dividend announcements. *Journal of Financial Economics*, 13(1):35–63, 1984.

[45] Jarrad Harford and Vahap B. Uysal. Bond market access and investment. *Journal of Financial Economics*, 112(2):147–163, 2014.

[46] Ralf Herbrich, Thore Graepel, and Klaus Obermayer. Large margin rank boundaries for ordinal regression. *Advances in Large Margin Classifiers*, 88, January 2000.

[47] Armen Hovakimian, Tim Opler, and Sheridan Titman. The debt-equity choice. *The Journal of Financial and Quantitative Analysis*, 36(1):1–24, 2001.

[48] Kai Huettermann and Denisa Lleshaj. Debtholder wealth effects in mergers and acquisitions: Evidence from the cds market. *SSRN Electronic Journal*, January 2020.

[49] John Hull, Mirela Predescu, and Alan White. The relationship between credit default swap spreads, bond yields, and credit rating announcements. *Journal of Banking & Finance*, 28(11):2789–2811, 2004.

[50] David Ikenberry, Josef Lakonishok, and Theo Vermaelen. Market underreaction to open market share repurchases. *Journal of Financial Economics*, 39(2):181–208, 1995.

[51] Michael C. Jensen and William H. Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, October 1976.

[52] Dongmei Jing, Mohsen Imeni, Seyyed Ahmad Edalatpanah, Alhanouf Alburaikan, and Hamiden Abd El-Wahed Khalifa. Optimal selection of stock portfolios using multi-criteria decision-making methods. *Mathematics*, 11(2):415, 2023.

[53] M. Johansson and J. Nederberg. Earnings announcements in the credit default swap market - an event study. Dissertation, 2014.

[54] Kose John and Joseph T. Williams. Dividends, dilution, and taxes: A signalling equilibrium. *Journal of Finance*, 40:1053–1070, 1983.

[55] Jeff Johnson and James M. LeBreton. History and use of relative importance indices in organizational research. *Organizational Research Methods*, 2004.

[56] Steven B Kamin and Karsten V Kleist. The evolution and determinants of emerging market credit spreads in the 1990s. *Bank of International Settlements*, Working paper No.68, May 1999.

[57] Simi Kedia and Xing (Alex) Zhou. Informed trading around acquisitions: Evidence from corporate bonds. *Corporate Law: Corporate & Takeover Law*, 2009.

[58] Tong Suk Kim, Jae Won Park, and Yuen Jung Park. Macroeconomic Conditions and Credit Default Swap Spread Changes. *Journal of Futures Markets*, 37(8):766–802, August 2017.

[59] Mark Klock and Katherine Gleason. Bondholder wealth effects from dividend changes. *Corporate Ownership and Control*, 4:42–52, January 2007.

[60] Josef Lakonishok and Theo Vermaelen. Anomalous Price Behavior around Repurchase Tender Offers. *Journal of Finance*, 45(2):455–477, June 1990.

[61] Yanyan Lan, Yadong Zhu, J. Guo, Shuzi Niu, and Xueqi Cheng. Position-aware listmle: A sequential learning process for ranking. In *Conference on Uncertainty in Artificial Intelligence*, 2014.

[62] Hwang Hee Lee and Jung-Soon Hyun. The asymmetric effect of equity volatility on credit default swap spreads. *Journal of Banking & Finance*, 98:125–136, 2019.

[63] Hang Li. Learning to rank for information retrieval and natural language processing. *Synthesis Lectures on Human Language Technologies*, 4:1–113, 2011.

[64] Rose C. Liao. What drives corporate minority acquisitions around the world? The case for financial constraints. *Journal of Corporate Finance*, 26(C):78–95, 2014.

[65] Zilong Liu, Xiaoling Pu, and Xinlei Zhao. What moves the correlation between the equity and credit default swap markets? *The Journal of Fixed Income*, 25:72–87, September 2015.

[66] Yee Cheng Loon and Zhaodong Ken Zhong. The impact of central clearing on counterparty risk, liquidity, and trading: Evidence from the credit default swap market. *Journal of Financial Economics*, 112(1):91–115, 2014.

[67] Tao Ma and Ying Tan. Stock ranking with multi-task learning. *Expert Syst. Appl.*, 199:116886, 2022.

[68] A. Craig Mackinlay. Event studies in economics and finance. *Journal of Economic Literature*, 35(1):13–39, 1997.

[69] Sattar A. Mansi, Yaxuan Qi, and John K. Wald. Bond covenants, bankruptcy risk, and the cost of debt. *Journal of Corporate Finance*, 66:101799, 2021.

[70] V A Marchenko and L A Pastur. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457, April 1967.

[71] Michele Marzano, Gary Dunn, and Nick Constantinou. The relationship between credit default swap spreads and equity prices. *Journal of Risk*, 17:3–28, October 2014.

[72] Miroslav Mateev. Relation between credit default swap spreads and stock prices: A non-linear perspective. *Journal of Economics and Finance*, XX:1–26, December 2017.

[73] William F. Maxwell and Clifford P. Stephens. The Wealth Effects of Repurchases on Bondholders. *Journal of Finance*, 58(2):895–919, April 2003.

[74] Merton H. Miller and Kevin F. Rock. Dividend policy under asymmetric information. *Journal of Finance*, 40:1031–1051, 1985.

[75] Sara B. Moeller, Frederik P. Schlingemann, and René M. Stulz. How Do Diversity of Opinion and Information Asymmetry Affect Acquirer Returns? *Review of Financial Studies*, 20(6):2047–2078, November 2007.

[76] Hans-Jörg Naumer and Burcin Yurtoglu. It is not only what you say, but how you say it: Esg, corporate news, and the impact on cds spreads. *Global Finance Journal*, 52:100571, 2022.

[77] Maureen O'Hara and Xing Alex Zhou. The electronic evolution of corporate bond dealers. *Journal of Financial Economics*, 140(2):368–390, 2021.

[78] Joel Ong and Dorien Herremans. Constructing time-series momentum portfolios with deep multi-task learning. *Expert Syst. Appl.*, 230:120587, 2023.

[79] Maureen O' Hara, Yihui Wang, and Xing (Alex) Zhou. The execution quality of corporate bonds. *Journal of Financial Economics*, 130(2):308–326, 2018.

[80] Heewoo Park and Yuen Jung Park. Stock buybacks and credit default swap spread changes. *Journal of Derivatives and Quantitative Studies*, 31(1):55–75, 2023.

[81] Daniel Poh, Bryan Lim, Stefan Zohren, and Stephen Roberts. Building cross-sectional systematic strategies by learning to rank. *The Journal of Financial Data Science*, 3:jfds.2021.1.060, March 2021.

[82] Lorenzo Portelli and Thierry Roncalli. Stock-bond correlation: Theory & empirical results. May 2024.

[83] Marcos Prado and Vincent Zoonekynd. Why has factor investing failed?: The role of specification errors. *SSRN Electronic Journal*, January 2024.

[84] Jiaping Qiu and Fan Yu. Endogenous liquidity in credit derivatives. *Journal of Financial Economics*, 103(3):611–631, 2012.

[85] L.D.R. Renneboog and P.G. Szilagyi. How do Mergers and Acquisitions Affect Bondholders in Europe? Evidence on the Impact and Spillover of Governance and Legal Standards. Discussion Paper 2006-55, Tilburg University, Center for Economic Research, February 2006.

[86] Katherine Schipper and Abbie Smith. Effects of recontracting on shareholder wealth: The case of voluntary spin-offs. *Journal of Financial Economics*, 12(4):437–467, 1983.

[87] Chengzhu Sun, Shujing Wang, and Chu Zhang. Corporate Payout Policy and Credit Risk: Evidence from Credit Default Swap Markets. *Management Science*, 67(9):5755–5775, September 2021.

[88] Dragon Yongjun Tang and Hongjun Yan. Liquidity and credit default swap spreads. *S&P Global Market Intelligence Research Paper Series*, 2007.

[89] Dragon Yongjun Tang and Hongjun Yan. What moves cds spreads? *Capital Markets: Market Microstructure eJournal*, 2013.

[90] Theo Vermaelen. Common stock repurchases and market signalling: An empirical study. *Journal of Financial Economics*, 9(2):139–183, 1981.

[91] Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.

[92] J. Randall Woolridge. Dividend changes and security prices. *Journal of Finance*, 38:1607–1615, 1983.

[93] Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. Listwise approach to learning to rank - theory and algorithm. *Proceedings of the 25th International Conference on Machine Learning*, pages 1192–1199, January 2008.

[94] Xin Zhang, Lan Wu, and Zhixue Chen. Constructing long-short stock portfolio with a new listwise learn-to-rank algorithm. *Quantitative Finance*, 22:321–331, 2021.

[95] Zihao Zhang, Stefan Zohren, and Stephen J. Roberts. Deep learning for portfolio optimisation. *Capital Markets: Asset Pricing & Valuation eJournal*, 2020.

Mathis Linger

Amélioration des Stratégies de Portefeuille Crédit : Perspectives sur les Opérations sur Titres, les Algorithmes d'Apprentissage du Classement et la Couverture des Taux d'Intérêt

L'évolution du marché du crédit, portée par les réformes réglementaires, l'électronification des passages d'ordre et l'augmentation des données disponibles, a profondément modifié les pratiques des gestionnaires d'actifs. Cette thèse contribue à l'amélioration des stratégies de crédit systématiques en explorant la dynamique du risque de crédit, la conception de modèles d'apprentissage automatique et les techniques de couverture des taux d'intérêt. La première contribution analyse l'impact des annonces d'opérations sur titres—telles que les rachats d'actions, les dividendes et les fusions-acquisitions—sur les spreads des credit default swaps. En quantifiant les mouvements anormaux des spreads autour de ces événements, il a été possible de dissocier l'effet wealth transfer de l'effet signaling, offrant ainsi une meilleure compréhension des mécanismes sous-jacents aux réactions du marché du crédit. La deuxième contribution améliore les modèles de classement d'actifs pour la construction de portefeuilles long-short. Les algorithmes traditionnels d'apprentissage du classement favorisent les actifs les mieux classés au détriment des moins performants, rendant les stratégies neutres au marché moins efficaces. Cette thèse propose des fonctions de perte ajustées afin de renforcer la symétrie du classement sur l'ensemble de la distribution des actifs, améliorant ainsi l'équilibre et la performance des portefeuilles. La troisième contribution propose un modèle de réseau neuronal multitâche unifiant le classement des actifs et la pondération des portefeuilles dans un modèle d'optimisation unique. Contrairement aux méthodes conventionnelles en deux étapes, cette approche produit une construction de portefeuille plus efficiente. Enfin, cette thèse examine la relation entre les marchés actions et des obligations d'entreprise en décomposant leur corrélation en deux composantes : le risque de crédit et la sensibilité aux taux d'intérêt. Puisque ces facteurs varient selon les qualités de crédit et les conditions de marché, deux stratégies adaptatives de couverture des taux d'intérêt sont proposées. Ces stratégies surpassent la couverture traditionnelle duration-convexity en ajustant dynamiquement les ratios de couverture selon les régimes de marché. Les contributions de cette thèse offrent aux gestionnaires d'actifs de nouvelles méthodologies afin d'optimiser les stratégies de gestion sur le marché du crédit, dans un environnement financier toujours plus complexe et axé sur les données.

Mots-clés : Contrat sur le risque de défaut, Obligations d'entreprises, Apprentissage du classement, Opération sur titres, Risque de crédit, Risque de taux d'intérêts

Advancing Credit Portfolio Strategies: Insights on Corporate Actions, Asset Ranking and Interest Rate Hedging

The evolution of credit markets, driven by regulatory reforms, trading electronification, and the increasing availability of data, has reshaped the landscape for asset managers. This thesis contributes to the advancement of systematic credit strategies by investigating credit risk dynamics, machine learning model design, and interest rate hedging techniques. The first contribution examines the impact of corporate action announcements—such as stock buybacks, dividends, and mergers & acquisitions—on credit default swap spreads. Quantifying abnormal spread movements around these events allows for the disentangling of wealth transfer from signaling effects, providing insights into the mechanisms driving credit market reactions. The second contribution focuses on improving asset ranking models for long-short portfolio construction. Traditional learning-to-rank algorithms favor top-ranked assets while neglecting underperformers, creating challenges for market-neutral strategies. This thesis proposes refined loss functions that enhance ranking symmetry across the full asset distribution, thereby improving portfolio balance and performance. The third contribution introduces a multi-task neural network framework that unifies asset ranking and portfolio weighting within a single optimization model. In contrast to conventional two-step approaches, this framework simultaneously optimizes both ranking and allocation, resulting in more efficient investment decisions. Finally, this thesis explores the relationship between equity and corporate bond markets by decomposing their correlation into credit risk and interest rate sensitivity components. Since these drivers vary across credit qualities and market conditions, this thesis proposes two adaptive interest rate hedging strategies. These strategies outperform traditional duration-convexity hedging by dynamically adjusting hedge ratios in response to evolving market regimes. These contributions collectively advance systematic credit investing by integrating financial expertise with advanced quantitative techniques. The findings offer asset managers practical methodologies for improving credit trading strategies in an increasingly data-driven and complex financial environment.

Keywords: Credit Default Swap, Corporate Bond, Learn-to-Rank, Corporate Action, Credit Risk, Interest Rate Risk

LÉO Laboratoire d'Économie d'Orléans

Laboratoire d'Économie d'Orléans, rue de Blois - BP 26739, 45067 ORLEANS Cedex 2

DRAKAI CAPITAL