Essays on Computational Portfolio Management and Asset Pricing

Inaugural Dissertation
submitted in partial fulfillment of the requirement for obtaining the degree of
Doctor rerum oeconomicarum (Dr. rer. oec.)
at the Schumpeter School of Business and Economics
University of Wuppertal

Sebastian Konstantin Schäfer
September 30, 2023

First Supervisor:
Prof. Dr. Uta Pigorsch
Lehrstuhl für Wirtschaftsstatistik und Ökonometrie

Second Supervisor:
Prof. Dr. André Betzer
Lehrstuhl für Finanzwirtschaft und Corporate Governance

Acknowledgement

I am deeply thankful to Prof. Dr. Uta Pigorsch for accepting me as an external doctoral student and for generously supporting me on writing this dissertation. From our very first conversation, she gave me enormous confidence that I would be able to both complete my dissertation and simultaneously start my professional journey in the economy as I could always rely on her unwavering backing. She taught me that doing research necessitates curiosity, endurance and creativity, all of which I believe I gained through working with her and will be highly valuable in the future. Most importantly, she never had to give me particular instructions on what I should be searching for, but always found a way enabling me to make the findings on my own that she was aiming at - a skill that only the best instructors possess.

I also want to thank Prof. Dr. André Betzer for accepting the role as my second supervisor and for giving extremely useful comments on different papers of this dissertation. Moreover, I greatly benefited from the comments from Dr. Markus Doumet, Casimir Carl, Nicola Schumacher and the participants of the 2022 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics, the International Conference on Computational and Financial Econometrics in 2022, the doctoral seminars of Prof. Dr. Uta Pigorsch and Prof. Dr. André Betzer, and the anonymous comments of several referees.

Finally, I am enormously grateful for the continuous support of my family. They are the reason I was able to begin working on my dissertation in the first place and constantly encouraged me to continue when things did not go as planned. I vastly underestimated the amount of effort required to write this dissertation part-time while also beginning my professional career; yet, I could always count on my family to help me find the time that I needed. But, most of all, I have to thank my partner Luisa. She had no choice but to go through every emotional state with me during the last two years, but mastered this with absolute excellence while also sacrificing countless nights and weekends for me to be able to work on this dissertation. I am incredibly grateful for everything she has done for me.

Preface

This cumulative dissertation consists of four individual papers that were created in collaboration with Prof. Dr. Pigorsch. The first paper, *High-Dimensional Stock Portfolio Trading with Deep Reinforcement Learning* is concerned with the empirical application of deep reinforcement learning in portfolio management. I contributed by proposing the ideas and concepts, the empirical analysis, and the initial writing. Prof. Dr. Pigorsch, on the other hand, provided valuable conceptional suggestions that influenced the empirical analysis and made contributions to improving the structure and writing. The paper *High-Dimensional Stock Portfolio Trading with Deep Reinforcement Learning* is published in a shortened version as the proceedings of the 2022 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFEr) (Pigorsch and Schäfer, 2022) where it has undergone a review process with three reviewers.

The second paper, *A Deep Learning Stack for Asset Pricing and Portfolio Management* presents a novel architecture for training deep neural networks for asset pricing. The paper focuses particularly on the computational aspect of deep learning in finance, but also provides valuable insights into the interpretability of neural networks for the task of predicting asset returns. I am the single author of this paper, and did not submit the paper for review as it served as a motivation for the follow-up paper *Anxiety in Returns*.

While the third paper, *Anxiety in Returns*, utilizes concepts of behavioral finance, it ultimately connects the empirical results to financial deep learning. The research idea originated from the second paper during a detailed analysis of the findings and quickly evolved into a separate research project. I developed the concepts and theoretical work presented in this paper as well as the empirical analysis, whereas Prof. Dr. Pigorsch made important suggestions for key empirical tests and greatly contributed to the writing of the paper. I presented the paper on the International Conference on Computational and Financial Econometrics in 2022 (CFE 2022) and Prof. Dr. Pigorsch presented the paper on the annual meeting of the Verein für Socialpolitik in 2023. Furthermore, the paper is published in the Journal of Behavioral Finance (Pigorsch and Schäfer, 2023).

The fourth paper, *Monday Afternoon Reversal*, stems from an empirical analysis of high-frequency data and reveals an asset pricing phenomenon. The distribution of contributions is similar to the third paper, as I developed the research idea, the empirical analysis but got great support from Prof. Dr. Pigorsch in solving empirical issues as well as with writing the paper.

Although the above provides an assessment of the quantifiable contribution to each paper, Prof. Dr. Pigorsch always motivated me to make the crucial "one more discovery", that causes each paper to be significantly more insightful. Furthermore, she shared a rich set of experiences, including writing styles, structure, and strategies for submission to the journal, from which I greatly benefited.

To make reading more convenient, I refer to a paper as a Chapter in the following, which consists of multiple Sections. Each Chapter is closely aligned with the original paper; however, this dissertation contains some additions. Although each Chapter assesses a distinct research question, I discuss important interconnections and joint results in between the Chapters as well as in the Introduction and the Conclusion.

Contents

1 Introduction **1**

2 High-Dimensional Stock Portfolio Trading with Deep Reinforcement Learning **6**

2.1 Introduction 6
2.2 Background and related literature 8
2.2.1 Reinforcement Learning 8
2.2.2 Related work 9
2.3 Deep Q-learning for Portfolio Management 11
2.4 Experiments 14
2.4.1 Methodology 14
2.4.2 Hyperparameters 17
2.4.3 Results 19
2.5 Conclusion 25

3 A Deep Learning Stack for Asset Pricing and Portfolio Management **26**

3.1 Introduction 28
3.2 Neural networks 30
3.3 Data and network designs 32
3.4 Results 36
3.4.1 Long portfolios 36
3.4.2 Long-short portfolios 42
3.4.3 Performance sources 47
3.4.4 Cost Sensitivity 49
3.5 Conclusion 53

4 Anxiety in Returns **54**

4.1 Introduction 55
4.2 Risk aversion, anxiety and investment behavior 59
4.3 Empirical analysis 64
4.3.1 Data 64
4.3.2 Empirical results for the anxiety anomaly 65
4.3.3 Time persistence 68
4.3.4 In-sample predictive power 69
4.3.5 Investor's convex reaction to fear 73
4.3.6 Alternative risk aversion measurement 75

4.3.7 Out-of-sample predictability 76
4.3.8 Potential drivers of the anxiety anomaly 79
4.4 Anxiety anomaly and neural networks 82
4.5 Conclusion 83

5 Monday Afternoon Reversal 86
5.1 Introduction 88
5.2 Monday afternoon reversal 91
5.2.1 Data 91
5.2.2 The Monday afternoon reversal 95
5.2.3 Predictive regression analysis 98
5.2.4 Out-of-sample analysis 99
5.2.5 Overnight versus intraday components 102
5.2.6 Relationship to systematic risk and firm news 102
5.3 U.S. stock market indices and international ETFs 104
5.3.1 Investor disagreement and informed trading 107
5.4 Conclusion 112

6 Summary and conclusions 113

List of Tables

1 List of hyperparameters used for all portfolio setups. 18
2 Out-of-sample cumulative returns at the end of the test period for each level of transaction costs and different portfolios. 20
3 Performance measures for long-only portfolios. 38
4 Baseline OLS regression results for long-only portfolios. 40
5 Fama-French five-factor regression results for long-only portfolios. . . 41
6 Performance measures for market-neutral long-short portfolios. 44
7 Baseline OLS regression results for long-short portfolios. 46
8 Fama-French 5-factor regression results for long-short portfolios. . . . 46
9 Performance measures for portfolios based solely on the directional model. 50
10 Performance measures for long-only portfolios under doubled transaction costs. 51
11 Performance measures for market-neutral long-short portfolios under doubled transaction costs. 52
12 Anxiety anomaly return spreads. 66
13 Returns of stocks with a preceding negative return for factor pairs. . . 72
14 Statistical analysis of the convexity in investors' risk aversion. 74
15 Potential drivers of the anxiety anomaly. 81
16 Summary statistics. 94
17 Rest-of-week returns from a double decile sort on the Monday afternoon return and market capitalization. 96
18 Estimation results for regressions of the rest-of-week return on the Monday afternoon return. 100
19 Overnight and intraday components of the rest-of-week returns of the bottom and top Monday afternoon return decile portfolios. 103
20 Estimation results for Fama and French (2015) and Carhart (1997) 6-factor analysis. 104
21 Estimation results for the regression of the weekly remaining return on the Monday afternoon return for U.S. stock market indices and international ETFs. 106
22 Double decile sort on the Monday afternoon return and volume. . . . 108
23 Double decile sort on the Monday afternoon return and volatility. . . 109
24 Rest-of-week returns based on Monday afternoon return deciles and days-to-cover quintiles. 111

List of Figures

1 Test trading performances of the proposed agent compared to the benchmark strategies including all 500 stocks for the different levels of transaction costs. 19
2 Test trading performances of the proposed agent compared to the benchmark strategies in the respective portfolio setup assuming transaction costs of 1 bps. 21
3 Test trading performances of the proposed agent compared to the benchmark strategies in the respective portfolio setup assuming transaction costs of 5 bps. 22
4 Test trading performances of the proposed agent compared to the benchmark strategies in the respective portfolio setup assuming transaction costs of 10 bps. 23
5 Out-of-sample trading performances for each level of transaction costs while splitting the test period into three phases based on market conditions. 24
6 Rolling window out-of-sample testing approach setup. 34
7 Performance plots for long-only portfolios. 39
8 Fama-French risk factor exposures for long-only portfolios. 42
9 Performance plots for long-short portfolios. 45
10 Fama-French risk factor exposures for long-short portfolios. 47
11 Aggregated feature importances for each model. 48
12 The price of anxiety induced by risk aversion. 75
13 Out-of-sample cumulative log returns. 79
14 Estimated return spread conditional on the preceding return. 84
15 Share of dollar volume per half-hour. 93
16 Value weighted out-of-sample cumulative logarithmic returns of the bottom and top Monday afternoon deciles. 101

List of Acronyms

DRL Deep Reinforcement Learning . . . 2

POMDP Partially Observable Markov Decision Process . . . 8

bps basis points . . . 17

ReLU Rectified Linear Unit . . . 18

Adam Adaptive Moment Estimation . . . 18

CRSP Center For Research in Security Prices . . . 32

ETFs Exchange Traded Funds . . . 86

MSE Mean Squared Error . . . 31

1 Introduction

Understanding the underlying factors of asset prices is one of the most challenging issues in finance. Ever since the development of the capital asset pricing model (Lintner, 1965; Mossin, 1966; Sharpe, 1964), economists seek to enhance the model's ability to describe contemporaneous returns and its ability to predict future returns by proposing hundreds of potential factors that may drive asset returns (Feng et al., 2020; Hou et al., 2015). While artificial intelligence permeates nearly every aspect of contemporary business and economics, researchers utilize machine learning successfully to uncover complex factors that drive asset returns (Almahdi and Yang, 2017; Bianchi et al., 2021; Fischer and Krauss, 2018; Gu et al., 2020; Kraus and Feuerriegel, 2017; Krauss et al., 2017; Patel et al., 2015; Wang et al., 2020; Zhong and Enke, 2017). Due to the high potential economic gains, portfolio management is inherently interested in identifying return predicting factors.

The following two Chapters of this thesis discuss deep learning methods for portfolio management. Deep learning is a subfield of machine learning that focuses on capturing empirical relationships between a set of features and one or more target variables using a variety of statistical models. The most well-known approach is the linear regression including its variants like the Ridge regression or the LASSO. These variants apply parameter shrinkage such that the estimated model becomes less sensitive to noise in the training data, which may ultimately lead to an enhanced out-of-sample performance, i.e., when estimating or predicting previously unseen data. Shrinkage models are also a common choice for predictive asset pricing of cross-sectional return data, which typically come with a low signal-to-noise ratio (Freyberger et al., 2020; Gu et al., 2020; Kozak et al., 2020).

However, the linear regression and its derivatives are, as the name suggests, linear models. Thus, if the data contains valuable non-linear patterns, such as interaction effects between multiple variables, these will remain uncovered when using the latter models. To this end, researchers frequently employ tree-based models to capture non-linear effects in the data, particularly in asset pricing. These models are based on ensembles of decision trees, known as gradient boosting or random forests. Several papers find high out-of-sample performance in predicting asset returns using the latter ensemble models, for instance, see Bianchi et al. (2021); Gu et al. (2020); Krauss et al. (2017).

Deep learning focuses on learning complex non-linear relationships in the data using neural networks. As neural networks may become highly complex and can identify the most spurious patterns in the data, they appear to render the latter statistical approaches for asset pricing obsolete (Bianchi et al., 2021; Chen et al., 2023; Gu et al., 2020; Krauss et al., 2017). However, as pure empirical pricing factors typically lack an economic justification, they may be incapable of predicting previously unseen data. Incorporating computational advances of deep learning into empirical asset pricing while concurrently obtaining economically interpretable models is, therefore, of great value for both academics and practitioners. This dissertation contributes to both of these aspects by combining insights from computational asset pricing, but also by providing a theoretically motivated asset pricing anomaly that significantly predicts equity returns and may explain a fraction of the predictive power of deep learning in asset pricing.

Typically, asset pricing focusses on either explaining or predicting equity returns. Usually, researchers evaluate the predictive power of asset pricing models or factors by evaluating the performance of portfolios that are constructed using the latter predictions. This approach not only reveals the economic value of the predictive power, but also shows its robustness over time in the selection of assets. We combine these two objectives in Chapter 2 by directly training artificial agents to manage portfolios using Deep Reinforcement Learning (DRL). DRL is based on deep learning and is concerned with optimizing the policy of an artificial agent that interacts with its environment by sending actions and receiving rewards. The agent's policy is typically founded on an iteratively optimized neural network. For instance, if the agent must learn to manage a portfolio, it needs to make investment decisions based on a set of features, such as the historical returns of an asset. The policy, i.e., the underlying neural network, transforms these features into a numerical decision, while the reward may be the return of the portfolio in the following period. Consequently, the algorithm learns directly to predict asset returns while simultaneously transforming these predictions into a portfolio.

The novel DRL algorithm we develop in Chapter 2 tackles a number of issues that may limit other algorithms from scaling well with the number of assets in the portfolio. For instance, applying DRL on a fixed set of assets is straightforward, as a single neural network may transform the set of features of these assets into portfolio weights, whereas the size of the feature set and portfolio weights do not vary in time. However, if the latter condition is not fulfilled, common DRL methods cannot be applied to managing portfolios containing a variable number of equities, making the latter method fairly limited in its application. We address this issue, as our agents only interact with one asset at a time, and the rewards are designed such that trading the single asset is evaluated as if the agent were managing a portfolio. Consequently, our approach scales well with arbitrarily large portfolios. As the algorithm also generalizes well, our algorithm outperforms passive and active investment strategies substantially in a variety of out-of-sample tests.

Although DRL comes with the benefit of being able to manage portfolios directly, it also requires high computational efforts. That is, training our algorithm not only involves the optimization of a neural network, which can be computationally costly, but also the simulation of a trading environment that needs to generate rewards on each action taken by the agent. This combination makes DRL hard to parallelize, resulting in relatively long training times. Thus, we develop a novel architecture for asset pricing via deep learning in Chapter 3, which we transform into portfolios heuristically, that is specifically designed to minimize training times and increase the interpretability of the resulting return predictions. To this end, we stack multiple neural networks that predict cross-sectional returns jointly. First, we set up a binary classification task to predict return signs. Second, as it is well known that positive and negative returns behave differently, we train one regression neural network exclusively with positive target returns, while the other is trained exclusively on negative target returns. Finally, we categorize assets as having either positive or negative expected returns and then determine the magnitude of those returns using the corresponding regression models.

Our neural network stack not only yields noteworthy predictive performance in large-scale out-of-sample tests on cross-sectional stock returns but is also computationally highly efficient. As we reduce the complexity of the prediction task by splitting it into multiple subtasks, the neural networks are effective at relatively small sizes. Furthermore, our approach allows to analyze feature importance of each neural network within the stack. We find that, e.g., past return volatility is more than twice as important for predicting positive returns as it is for predicting negative returns. Overall, however, all three models agree that past returns are by far the most important features.

The fourth Chapter presents a novel return anomaly that may be closely related to the importance of past returns, specifically the return sign, for predicting future returns as well as the non-linear factors revealed by neural networks. We find that during periods of high common risk aversion, investors avoid stocks with preceding negative realized returns, making these stocks yield significantly lower returns in the subsequent period than stocks with preceding positive returns. We propose that this return difference is due to investors becoming anxious to invest in losing stocks when they are highly risk-averse, hence, we refer to this phenomenon as the anxiety anomaly. Our findings are based on a newly developed measure of risk aversion but can also be observed using a well-established measure of risk aversion.

We incorporate the anxiety anomaly into a simple linear model and find that it predicts future returns consistently well over multiple decades. In fact, the model yields comparable predictive performance to deep neural networks with hundreds of predictors, as in Gu et al. (2020). Interestingly, neural networks appear to generally prioritize the non-linear transformation of recent returns as well as firm size, closely matching the anxiety anomaly. We train a neural network specifically on the preceding return and the size long-short risk factor spread, which is the underlying factor of our measure of risk aversion, and find that the neural network converges to predictions that are similar to those of the anxiety anomaly. Hence, we conjecture that predictability in returns found by machine learning models may stem from predictability in investors' behavior.

In the fifth Chapter, we describe a strong reversal of returns on Monday afternoons throughout the remaining week's return. Analyzing the high-frequency returns of a U.S. cross-section, indices, and international ETFs, we find that the reversal is consistent among major international equity markets and is not due to systematic or overnight risk. We undertake a variety of robustness checks and find that the reversal is most prevalent among large stocks and is not due to earnings announcements. Conversely, we find no support for the predictive power of the Monday morning return.

On weekends, investors have more time to digest accumulated information in depth without any market interaction, potentially creating heterogeneous beliefs about fundamental values. We measure investor disagreement by volume and volatility on Monday afternoons and find that the reversal exists exclusively among stocks with high volume or volatility on Monday afternoons. Furthermore, we document significant momentum for stocks with low Monday afternoon volatility, strongly suggesting that investor disagreement is the underlying cause of the reversal.

If the Monday afternoon reversal is due to investor disagreement, we anticipate an increased reversal when informed investors strongly disagree with Monday afternoon returns. In order to assess this, we analyze short interest data on Mondays, as short sellers are well-known as informed investors. Indeed, we find that short sellers consistently predict the remaining week's returns on Mondays, and the reversal is strongest when short sellers disagree with Monday afternoon returns, providing further evidence that the reversal is due to investor disagreement.

Ultimately, the final Chapter concludes this dissertation. Here, we provide a summary of the results presented in each Chapter. Furthermore, we discuss the origins of the deep learning's predictive power in asset pricing and, hence, outline the relationships between the consecutive Chapters of this thesis. In particular, as a neural network trained on the features of the anxiety anomaly converges to predictions of the anxiety anomaly, we question whether the predictive power of deep learning may be due to the predictability in investors' behavior. Finally, we close this dissertation with some remarks concerning the future of portfolio management via machine learning.

2 High-Dimensional Stock Portfolio Trading with Deep Reinforcement Learning

Abstract

This Chapter proposes a Deep Reinforcement Learning algorithm for financial portfolio trading based on Deep Q-learning (Mnih et al., 2013). The algorithm is capable of trading high-dimensional portfolios from cross-sectional datasets of any size, which may include data gaps and non-unique history lengths in the assets. We sequentially set up environments by sampling one asset for each environment while rewarding investments with the resulting asset's return and cash reservations with the average return of the set of assets. This forces the agent to strategically assign capital to assets that it predicts will perform above average. We apply our methodology in an out-of-sample analysis to 48 U.S. stock portfolio setups, varying in the number of stocks from ten up to 500, in the selection criteria, and in the level of transaction costs. The algorithm, on average, outperforms all considered passive and active benchmark investment strategies by a large margin using only one hyperparameter setup for all portfolios.

2.1 Introduction

Portfolio management includes the process of analyzing financial assets and estimating future returns and risks. As the amount of available data, especially for stocks, is rising, intelligent systems that automate or even completely control the portfolio management workflow can increase the investor's performance (Gu et al., 2020). Therefore, machine learning is widely used within portfolio management, e.g., for return predictions or risk evaluation as in Fischer and Krauss (2018); Patel et al. (2015); Wolff and Echterling (2020).

DRL is the process of agents finding policies that maximize the cumulative rewards for a specific task resulting from learning through interaction with an environment. As a basis for its decision-making, the agent receives information about the state of the environment for every action it is supposed to take. The agent should become artificially intelligent by using non-linear function approximators like neural networks in order to improve its decision-making.

In the literature, there exist mainly two approaches to applying DRL within the context of portfolio management. Training agents that are capable of trading a single asset, which can include both long and short positions in the respective asset, and training agents to make capital assignments to a portfolio of assets, i.e., learning the optimal weighting policy. However, as we discuss in the following, these approaches suffer from limited flexibility and/or generalization capabilities.

Both methodologies come with substantial drawbacks. Training an agent on a single asset makes the environment, i.e., the sample size, relatively small. For that, considering a stock with a data history of ten years and daily datapoints gives only approx. 2530 samples, whereas, e.g., state-of-the-art neural networks for image recognition are evaluated using datasets that are multiple orders of magnitude larger (He et al., 2016). Small datasets come with the drawback that noise-induced signals are not fit and do not provide sufficient information for a generalizable solution to the problem. Moreover, out-of-sample results for single assets are rather weak in terms of their economic and statistical expressiveness. Finding a hyperparameter setup that performs reasonably well in-sample as well as out-of-sample can be found with tenable computational efforts. This makes solutions found for a single asset fairly restricted to the asset itself and potentially not valid for other assets. These issues motivate us to train agents that are capable of trading asset portfolios. However, when using neural networks to generate output for a vector of asset weights, a complete feature matrix for every period is required. Specifically, each period that contains a missing value cannot be used for training and might also make adjacent periods obsolete as periods are usually interconnected. This requires, e.g., a portfolio consisting of stocks listed for the same time period. As a consequence, the number of available assets to combine into a portfolio is rather limited. We make crucial extensions to Deep Q-learning (Mnih et al., 2013) to simultaneously benefit from large, cross-sectional datasets and portfolios of assets while still using environments with only one asset at a time for training.

Our methodology contains three major extensions. First, we sample one asset from a set of assets and set up a typical DRL environment for trading the single asset. Using asset-specific features in each state, the agent makes periodic decisions about when to invest in the single asset. At each terminal period, the environment is reset with a newly sampled asset. Second, we reward investments in a single asset with its next period's return. Reserving cash, i.e., not investing, is rewarded with the mean return of the set of assets in the next period. Lastly, when evaluating the agent, in every period we build equally weighted portfolios with each asset the agent is willing to invest in. This methodology allows for large, cross-sectional datasets with stronger generalization due to the random sampling of assets, while the reward function motivates the agent to actively avoid investments that the agent predicts will perform below average.

The remainder of the Chapter is structured as follows: Section 2.2 briefly introduces the concepts of DRL in the context of trading financial assets and provides a short review of the related literature. Our algorithm, which uses Deep Q-learning (Mnih et al., 2013), is proposed in Section 2.3 and tested on U.S. stock data in Section 2.4. Section 2.5 concludes and discusses implications for further studies.

2.2 Background and related literature

2.2.1 Reinforcement Learning

We assume that the task of trading a financial asset can be represented as a Partially Observable Markov Decision Process (POMDP) (Spaan, 2012). We do so because there must be some underlying state that describes every possible trading decision made in the markets but is so large that it is practically unobservable. We therefore construct an observable subset of the state consisting of features that are sufficiently relevant. The features are outlined below. We follow the standard definitions for the POMDP and denote the environment by $\mathcal{E}$ in which the agent acts. For each iteration, the agent receives information about the state in time step t, denoted as s_t, and performs an action a_t from the action space $\mathcal{A}$. As a consequence, the environment returns a reward R_t from the reward function $\mathcal{R}$ given the specific state-action pair. Note, that we use R_t for rewards and r_t for financial returns in the following. That is, for every iteration, the agent receives a tuple consisting of the state, action, and reward, i.e., (s_t, a_t, R_t). Based on the tuple, the agent's task is to find a profitable trading strategy given by the policy $\pi(s_t)$.

When trading financial assets, the agent must make periodic decisions. The definition of a period can vary from split seconds to weeks. At each period, the agent receives information about the current state of the environment. In the context of stock trading specifically, this may be the agent's available capital and trading positions as well as asset-specific information such as figures from recent earnings reports and the price history. The latter is frequently represented by technical indicators.

Using the available information, the agent should decide on an action that results in a step in the environment. In a financial trading environment, that means undertaking an investment or adjusting portfolio weights with the available capital. The reward is typically defined by the asset's or portfolio's return, either immediately for each period or cumulatively at the end of multiple periods. Alternatively, the excess return compared to a specific benchmark or the Sharpe ratio can be well suited, too (Almahdi and Yang, 2017; Théate and Ernst, 2021; Yang et al., 2020).

2.2.2 Related work

DRL is widely used for trading and portfolio management. These studies are either concerned with trading assets individually or trading portfolios by learning dynamic asset allocation policies. The following studies are concerned with the trading of single assets.

Huang (2018) demonstrate the effectiveness of DRL for trading currency pairs. The author uses recurrent DRL to trade 12 currency pairs and concludes that using a rather small replay memory, compared to what is used in standard reinforcement learning, can be more effective. Moreover, the author finds that higher trading costs do not necessarily decrease trading performance as learned strategies become more robust. On all pairs, the agent successfully outperforms the benchmark strategies of buy-and-hold as well as sell-and-hold. Moreover, the author provides insights into the trading behavior of the agents. That is, agents maintain higher win rates of approximately 60% with average profits close to zero. These results are in line with our findings. Specifically, we show that the agents are capable of adjusting to increased transaction costs as compared to active benchmark strategies while also yielding higher cumulative returns than a passive investment approach.

The Deep Q-learning algorithm (Mnih et al., 2013) is a popular choice for automated financial asset trading. Théate and Ernst (2021) test trading strategies based on DQN-agents over a cross-section of 30 stocks and stock indices. The agents successfully find trading strategies that, on average, outperform a passive buy-and-hold strategy. The authors train a DQN agent for each asset individually over a period of six years and evaluate their respective performance over a subsequent two-year time period. Overall, agents tend to alternate between a more passive, buy-and-hold strategy and a mean-reversion strategy but cannot outperform a passive investment benchmark for some assets. We find that the capability of alternating and combining multiple investment strategies is crucial as we compare our agents to both passive and active benchmarks whereas none of these is a dominant strategy in the proposed portfolio setups. Li et al. (2019) compare different variants of the Deep Q-learning (vanilla vs. double vs. duelling) and find the best performance when using the vanilla Deep Q-learning on ten U.S. stocks. Additionally, Zhang et al. (2020) compare different DRL algorithms in discrete and continuous action spaces for trading futures contracts. The analysis contains 50 futures contracts from multiple asset classes. The authors compare a variety of DRL methodologies on each asset class by forming portfolios based on the decisions made on each asset.

Overall, Deep Q-learning performs the best. Besides, the authors show robustness in the performance of the algorithms for various levels of transaction costs. Based on these results, we use the vanilla Deep Q-learning algorithm (Mnih et al., 2013) as our baseline algorithm. As our choice of the neural network architecture, we follow Taghian et al. (2020) who compare different feature extraction neural network architectures for the task of DRL for financial asset trading. Based on tests on four assets, the authors find that, overall, a simple multi-layer perceptron architecture based on a Deep Q-learning algorithm performs the best. Furthermore, they find the best results using the time series of raw price data (i.e., open, high, low, and close prices) as inputs while comparing it with hand-crafted inputs from the time series of price data like candlestick patterns. As hand-crafted inputs seem to provide little predictive power, we only transform the time series of past returns using multiple moving averages in combination with stock specific data from quarterly statements to reduce the dimensionality of the input vector.

Besides trading single assets, some studies address the portfolio allocation problem with DRL. Xiong et al. (2018) show the effectiveness of DRL for portfolio trading using a DRL agent to trade a portfolio of 30 assets. The agent achieves an annualized return of 22.24% in comparison to a 15.93% return given by the Dow Jones Industrial. Park et al. (2020) achieve promising results using deep Q-learning for multi-asset trading. The authors use an action mapping function to gain a discrete action space, which is supposed to be more practical. The mapping function selects actions that are close in effect to the agent's chosen action while keeping the tradeover relatively low. With this setup, the authors achieve outperformance over multiple benchmark strategies for an U.S. and Korean portfolio. We follow the choice of a discrete action space and show the agent's capability to yield high cumulative returns even when transaction costs are high.

Jiang et al. (2017) utilize an ensemble of parallelly trained agents to dynamically weight assets. Agents are trained independently and only share the last neural network layer, i.e., the softmax. By applying the structure to the cryptocurrency market on half-hourly returns, the technique produces cumulative returns in multiples of the returns of the benchmark strategies and also strongly outperforms in terms of the Sharpe ratio. Srivastava et al. (2020) use reinforcement learning to solve the problem of finding the optimal, dynamic asset allocation strategy. For that, the authors compare different network structures for the agents, specifically a convolutional neural network, a vanilla recurrent neural network, and a recurrent neural network with long short-term memory cells. Given a portfolio of 24 U.S. stocks, all structures successfully outperform an equally weighted portfolio in terms of total returns and

Sharpe ratio. Additionally, the authors present how feeding in the time series of past weights can stabilize trading and, thus, dramatically reduce the turnover chosen by the agents. Almahdi and Yang (2017) show that recurrent DRL agents perform well when optimizing for risk measures such as the Sharpe or Calmer ratio. The authors validate their findings by comparing out-of-sample trading agents with several benchmarks, all of which were outperformed by the risk-measure-optimized agents.

Overall, these studies demonstrate the potential of DRL in portfolio management applications. We contribute to the existing literature by proposing a DRL algorithm that is highly flexible in the portfolio setup. That is, using our self-regularizing algorithm, one simple hyperparameter setup is sufficient to successfully trade a large variety of stock portfolios. As such, our approach is flexible, i.e., not tailored to individual stocks and, hence, more generally applicable. Furthermore, we completely disregard any hand-crafted, additional trading logic, making the agent's trading strategies fully self-contained.

2.3 Deep Q-learning for Portfolio Management

A (risk-neutral) investor's goal is the maximization of wealth W_T, where T is the terminal period at the end of an investment period, after undertaking a series of investment decisions that manipulate the initial wealth $W_{t=0}$. Starting with an initial wealth of $W_{t=0}$, the investor undertakes investment decisions like buying and selling stocks. In the following, one period corresponds to one day. That is, the agent can change its capital assignments daily. Hence, the investor's final wealth, W_T is the cumulative return resulting from the daily decisions made by the agent, i.e.:

$$W_T = W_{t=0} \prod_{t=1}^{T} (1 + r_t),$$

with r_t denoting the return at the end of day t. To simplify the objective, we consider the natural logarithm of wealth, which is given by:

$$\log(W_T) = \log(W_{t=0}) + \sum_{t=1}^{T} \log(1 + r_t).$$

When the returns r_t are close to zero, which is a reasonable assumption if the returns are measured on a daily frequency, we use the following approximation:

$$\sum_{t=1}^{T} \log(1 + r_t) \approx \sum_{t=1}^{T} r_t,$$

such that the objective becomes to maximize the sum of (daily) returns. However, as investors are myopic (Benartzi and Thaler, 1995a), short-term returns should be more valuable than returns in the distant future. Thus, we discount returns by a factor γ such that the objective changes to maximizing $\sum_{t=1}^{T} \gamma^t r_t$.

As the agent must learn from its periodic rewards, a value should be assigned to each state-action pair. To this end, a state-action function is defined that assigns to each state-action pair a value based on the expected sum of future, discounted rewards, assuming that the agent will follow the policy π from the current timestep k to T, i.e.:

$$\begin{aligned} Q(s,a) &= \mathbb{E}_\pi \left[\sum_{k=0}^{T} \gamma^k R_{t+k} \mid s_t = s, a_t = a \right] \\ &= \mathbb{E}_{a' \sim \pi} \left[R_t + Q(s', a') \mid s_t = s, a_t = a \right], \end{aligned} \tag{1}$$

where $\mathbb{E}_\pi$ is the expected value following policy π, s' is the state that results from executing action a in state s and a' is the action chosen from policy π in state s' (see also Sutton and Barto (2018)). A natural choice for a policy is to always select the action that maximizes the state-action value. However, this requires a model that determines, or, in a stochastic framework, approximates the state-action function $Q(s,a)$. Mnih et al. (2013) suggest to solving this issue by training neural networks (here specifically called Q-networks) to approximate the state-action function. The authors use an agent that has a discrete action space and decides based on the state-action value predictions. The Q-network, in which parameters are randomly initiated, is trained iteratively using stochastic gradient descent. Using the squared loss between predicted state-action values and partially observed state-action values, the neural network is iteratively enhanced. A predicted state-action value is obtained from the Q-network with the state vector as the input layer. Partially observed state-action values, y_i, are obtained based on Equation (1) where the reward R_t is observed and the state-action value is predicted by the Q-network, however, assuming that in the consecutive periods only optimal actions are chosen:

$$y_i = Q^*(s,a) = \mathbb{E}\big[R_t + \max_{a'} Q^*(s', a') | s_t = s, a_t = a\big].$$

If state s is the terminal state, then there are no more actions to take and no state-action value predictions are necessary. Hence, the target is set to the reward that is observed for the state-action pair: $y_i = R_t$. Using the maximum expected state-action value makes the algorithm off-policy, i.e., the policy is updated using actions that are not originated by the policy itself. The squared loss, which should be minimized, results in:

$$L_i(\theta_i) = \mathbb{E}[(y_i - Q(s, a; \theta_i))^2],$$

with the corresponding gradient as:

$$\begin{aligned} \nabla\theta_i L_i(\theta_i) = \mathbb{E}_{s' \sim \mathcal{E}}\big[(R_t + \gamma \max_{a'} Q(s', a'; \theta_{i-1}) \\ - Q(s, a; \theta_i) \nabla\theta_i Q(s, a; \theta_i)\big]. \end{aligned} \tag{2}$$

Samples are generated via an epsilon-greedy strategy. That is, with a probability of ϵ, the agent chooses a random action while selecting the greedy action, i.e., the action with the highest expected state-action value, with a probability of $1 - \epsilon$. This ensures that sufficient exploration is undertaken. However, the concrete value of ϵ depends on the problem itself. Since consecutive samples can be highly correlated, the authors propose using an experience replay memory. While exploring the environment, new samples are stored in the experience replay memory. When updating the Q-network's parameters via stochastic gradient descent, a batch of samples is randomly drawn from the experience replay memory. This furthermore increases the sample efficiency of the algorithm as samples are used more than once (Mnih et al., 2013). The full algorithm is called Deep Q-learning. We use the algorithm as described and include further extensions so that it fits the purpose of trading financial asset portfolios.

Generally, the agent is supposed to be able to train portfolios of assets, and it is also desirable that it can deal with different history lengths of the assets. To deal with that, we differentiate between a training and a trading phase.

When training the agent, we only construct environments with single assets. Here, the agent needs to learn to trade a single asset by either investing in it or holding cash. Our methodology is long-only, which is a product of the reward function that is outlined further below. A single-asset environment is constructed by randomly drawing one of the assets from the selection of available assets with replacement. That is, during training, each asset is shown to the agent potentially multiple times. We define the state as a stack of asset-specific features and a dummy variable that indicates the most recent action the agent has taken. Each time the agent reaches the terminal state of the asset's data history, a new asset is drawn, i.e., the environment is reset.

The process of optimal trading of financial portfolios by an agent can be defined as the maximization of wealth using a policy that assigns capital to a set of assets. We use a discrete action space as motivated by Park et al. (2020). The agent can choose between either $a_t = 1$ with which the agent invests in the asset or holds its long position if it has already invested. If it decides for $a_t = 0$, the agent decides to reserve cash for other investments. That reduces the problem to the following decision made for each asset and period: Will the asset perform above average, or should the capital be reserved for other assets? To promote that decision-making, we reward investments in the asset with the next period's return adjusted for transaction costs C if they occur. The reward for reserving cash is the average return of the set of assets over the next period. Thus, the reward function is defined as:

$$R_t = \begin{cases} r_{t+1,i} - (1 - a_{t-1})C & \text{if } a_t = 1 \\ \frac{1}{N_t} \sum_{j=1}^{N_t} r_{t+1,j} & \text{if } a_t = 0, \end{cases} , \tag{3}$$

where $r_{t+1,i}$ is the next period's return of the currently selected asset in the environment, and N_t is the number of stocks in the set of assets in period t.

This training setup forces the agent to maximize the cumulative return in the training set. However, it is desirable to find a policy that achieves high cumulative returns in out-of-sample datasets. Thus, we prevent the agent from overfitting by using a validation set. We use this set at regular evaluation intervals Ω by computing the out-of-sample cumulative return achieved by the agent's policy. We initialize the cumulative return on the validation set with $CR_v^* = 0$. That is, there will be no solution if the agent yields losses when trading in the validation set. However, this does not occur in our experiments. For brevity, we define the steps without evaluation as ω. We save the Q-network's parameters with which the agent performs best on the validation set in terms of cumulative return. The full algorithm is stated in Algorithm 1.

To compute cumulative portfolio returns, we compute state-action values for the assets in each period for both possible actions, i.e., investing and reserving cash. For that, we construct an equally weighted portfolio from the assets that the agent wants to invest in, i.e., where the state-action values suggest a long position. As the agent can hold assets for multiple periods but the weight of the asset may still change, we account for the costs that occur when the weights need to be increased. These rebalancing costs occur when assets remain in the portfolio for multiple consecutive periods while the number of assets shrinks.

2.4 Experiments

2.4.1 Methodology

We apply our algorithm to a cross-section of U.S. stocks, which we obtain from SimFin Analytics GmbH (2023). To this end, we collect a data history of 500 stocks, including

Algorithm 1 Deep Q-learning for portfolio trading

Extensions to the original Deep Q-learning algorithm from (Mnih et al., 2013) are marked in italics.

Initialize experience replay memory $\mathcal{D}$ to capacity $\mathcal{N}$
Initialize action-value function Q with random weights
Sample random asset and generate environment
Initialize highest cumulative return on validation set to $CR_v^* = 0$
Initialize steps without evaluation to $\omega = 0$
for $t = 1$ to N **do**
  $\omega = \omega + 1$
  if $\omega = \Omega$ ***then***
    Get cumulative return on validation set and store in CR_v
    if $CR_v > CR_v^*$ ***then***
      Update Q-network parameters with current parameters: $\theta^* \leftarrow \theta$
      Set $CR_v^* \leftarrow CR_v$
    end if
    Set $\omega = 0$
  end if
  Get current state s_t
  With probability ϵ select random action a_t
  else select greedy action: $a_t = \text{argmax}_a \; Q(s, a)$
  Update trading position and receive reward R_t *according to Equation* (3)
  Observe next state s_{t+1}
  Store transition (s_t, a_t, r_t, s_{t+1}) in $\mathcal{D}$
  Sample batch of (s_j, a_j, r_j, s_{j+1}) from $\mathcal{D}$
  Set
  $$y_j = \begin{cases} R_j & \text{if state is terminal} \\ R_j + \gamma max_{a'} Q(s_{j+1}, a'; \theta) & \text{otherwise} \end{cases}$$
  Perform gradient step on $(y_j - Q(s_j, a_j; \theta))^2$ according to Equation (2)
  if *state is terminal* ***then***
    Sample new random asset and generate environment
  end if
end for

daily data from January 2010 to June 2021. To evaluate the algorithm's capabilities, we compare cumulative returns over the test period with three benchmarks. First, we use a passive approach via an equally weighted buy-and-hold portfolio using all stocks. Second, we include two active investment strategies following momentum or reversion with a simple rule. The momentum strategy invests in stocks with a positive average return over the last five trading days. In contrast, the reversion strategy buys stocks that have had a negative average return over the last five trading days. We report cumulative returns for all strategies and considered sets of stocks. As the 500 stocks are selected based on their terminal market capitalization, the resulting returns are higher than those of stock indices, like the S&P500.

We construct 15 portfolios from the cross-section. These differ in the number of stocks k and the selection criteria. That is, we form portfolios with k being 10, 25, 50, 100, and 200. For each k, we build portfolios using the biggest and smallest k stocks in terms of market capitalization and one where k stocks are randomly selected from the full cross-section. This gives each portfolio unique characteristics, although they may partly share stocks. Besides, we also test the algorithm using all 500 stocks available. After selecting stocks for each portfolio, we drop for each stock those periods that contain a missing value. However, we keep the period's data for all other stocks, as the algorithm is capable of trading a varying portfolio size and, hence, is not affected by data gaps. We split every portfolio's data history into training, validation, and test sets. Validation starts in January 2019, with testing in January 2020. This yields a test period that is a great challenge for the agent since it includes periods with sharp losses at the beginning of the COVID-19 pandemic coupled with a strong trend change resolving into a market with overall large gains while the agent is trained over a less volatile period. This situation allows access to the true capabilities of the agent to adjust to extreme scenarios, but in early 2020, the agent is also tested during a more tranquil period.

We use both the fundamental data available from quarterly statements and the stock price history. As fundamental data, we use the following indicators: sales per share, gross margin, operating margin, net profit margin, return on equity, return on assets, current ratio, debt ratio, market capitalization, the book to market equity ratio, and the latest close price, which are overall closely related to key factors found in Fama and French (2015). A key feature for trading financial assets via DRL is the price history, which is typically used in the input layer, e.g., in Almahdi and Yang (2017); Jiang et al. (2017); Théate and Ernst (2021); Zhang et al. (2020). To reduce the dimensionality of the input vector, we transform the price history using moving averages of the return history, both arithmetically and exponentially weighted. We hereby use windows of 5, 10, 20, 50, 100, and 200 days. Furthermore, we use rolling standard deviations of returns over a history length of 5, 10, 20, 50, or 100 days. Lastly, the Q-network's input vector contains a dummy variable indicating if the agent has already invested in the currently observed stock. This may be relevant for deciding if a trade is worth the transaction costs. Features are scaled to be normally distributed with a zero mean and unit standard deviation using the training dataset. Afterwards, scaling is only applied to the validation and test datasets.

In our analysis, daily returns are used to calculate the reward for each period. We assume that the agent can enter the market for any amount of shares at the trading day's opening price and close all positions at the closing price. Besides, we analyze the performances for three levels of transaction costs: 1, 5, and 10 basis points (bps). We hereby follow the literature that uses DRL for trading. E.g., Srivastava et al. (2020) assume 5 bps for an U.S. stock portfolio, Théate and Ernst (2021) assume transaction costs from 0 to 20 bps for single stocks, and Zhang et al. (2020) analyze the performance achieved for futures contracts from 1 up to 45 bps.

2.4.2 Hyperparameters

One of the main advantages of using reinforcement learning next to supervised learning for the portfolio management task is the ability of the agents to not only make predictions about the next period but also to place that prediction into the scope of a long-term planning process using discounted rewards. However, a financially reasonable risk-adjusted discount factor, e.g., the daily average rate of return of a market index, is typically very close to one. This means that the agent has to make long-term predictions about the distant future development of the stock. On the other hand, we expect more myopic predictions to be more reliable. Similarly, Théate and Ernst (2021) state that there is a trade-off between long-term orientation and increasing uncertainty in returns far into the future. Still, there are examples in which a discount factor close to one worked out well, e.g., in Huang (2018) or Chen et al. (2021). According to our expectation, we found that when the discount factor is too close to 1, the uncertainty about future returns becomes too high, and the policy breaks down to a constant prediction that chooses the same action in every state. The issue resolves at a discount factor of 0.9, a value that is also used in the Deep Q-learning approach of Park et al. (2020). The same reasoning applies to the exploration rate. We found that a higher exploration rate of 30% enables the agent to regularly find improvements to its policy and avoid getting stuck in local minima. We found that 3,000,000 training steps are sufficient for every setup, and optimal solutions are likely to be found much earlier. As in Mnih et al. (2013) we keep the

Parameter	Symbol	Value
Discount factor	γ	0.9
Exploration probability	ϵ	0.3
Number of iterations	N	3,000,000
Experience memory size	$\mathcal{N}$	300,000
Gradient step interval	–	20
Evaluation interval	Ω	10,000
Optimizer	–	Adam
Batch size	–	1024
Number of hidden layers	–	2
Number of neurons in hidden layers	–	32 or 64 or 128

Table 1: List of hyperparameters used for all portfolio setups.

experience memory size at 10% of the number of training steps, i.e., at 300,000.

Moreover, to increase the training efficiency, we perform gradient steps every 20 iterations and use a large batch size of 1024 to efficiently use the capabilities of GPUs. Alongside, we evaluate the agent 300 times during training, i.e., every 10,000 iterations. It is important to note that the choice of the evaluation interval is associated with a trade-off between the additional time required for computations and the performance that is lost due to infrequent evaluations.

Financial data for return predictions contains a low signal-to-noise ratio. To avoid overfitting, we need to keep the number of parameters in the networks relatively low. We find it more useful to extend the number of layers instead of making the network wider, as, on much larger scales, is also done in He et al. (2016). Hence, we use two hidden layers, activated by the Rectified Linear Unit (ReLU), with either 32, 64, or 128 neurons. Furthermore, we combine predictions from the three network configurations, forming an ensemble of predictors. Krizhevsky et al. (2012) boost accuracy in an image classification task using an ensemble of convolutional neural networks. Xie et al. (2013) show increased performance using ensembles of multiple networks for representation learning. Lastly, to update the weights of the agents' networks, we use the Adaptive Moment Estimation (Adam) optimizer; see also Kingma and Ba (2014). All hyperparameters are summarized in Table 1.

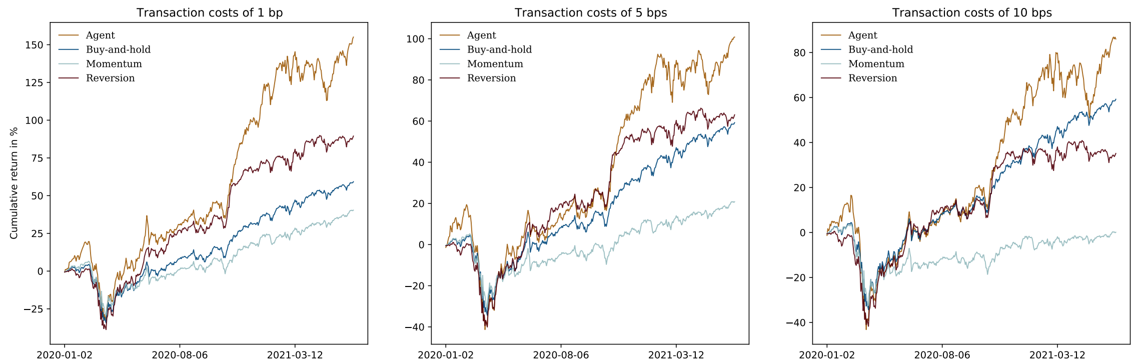

Figure 1: Test trading performances of the proposed agent compared to the benchmark strategies including all 500 stocks for the different levels of transaction costs.

2.4.3 Results

For each of the portfolio and transaction cost setups, we train three agents (one for each Q-network setup) and form the respective ensembles. Then, we evaluate the out-of-sample performance using the cumulative returns achieved in the test period compared to the three benchmark strategies. This gives a total of 148 agents to train and 48 unique portfolio and transaction cost setups.

All out-of-sample cumulative returns are reported in Table 2 (the best performing strategies are indicated in bold) and depicted in detail in Fig. 1 - 4. Figure 1 shows a comparison of the agent's performance compared to the benchmark strategies when trading all 500 stocks for different levels of transaction costs. The performance comparison for smaller portfolios sizes, i.e., from 10 to 200 stocks, are depicted in Figure 2 assuming 1 bp of transaction costs, Figure 3 assuming 5 bps of transaction costs and Figure 4 assuming 10 bps of transaction costs.

Transaction costs	Portfolio size	Portfolio type	Agent	Buy-and-hold	Momentum	Reversion
1 bp	10	big	**222.9%**	129.5%	103.6%	63.8%
		random	**68.1%**	56.4%	50.6%	29.7%
		small	**109.0%**	65.7%	100.7%	108.7%
	25	big	41.4%	**66.2%**	39.4%	58.8%
		random	**78.6%**	48.1%	47.6%	38.2%
		small	**147.8%**	84.0%	108.5%	94.1%
	50	big	**101.0%**	52.7%	44.5%	60.9%
		random	**90.8%**	62.5%	16.9%	79.9%
		small	**119.0%**	66.9%	51.7%	85.9%
	100	big	**126.0%**	56.3%	53.8%	90.1%
		random	**185.9%**	67.8%	17.1%	116.6%
		small	**475.0%**	68.0%	52.2%	72.0%
	200	big	**115.6%**	58.9%	31.0%	111.5%
		random	84.4%	63.6%	17.7%	**106.1%**
		small	**195.0%**	59.4%	41.0%	76.5%
	500	all	**155.1%**	59.2%	40.4%	89.5%
	Mean		***144.7%***	*66.6%*	*51.1%*	*80.1%*
5 bps	10	big	123.2%	**129.5%**	76.6%	44.3%
		random	24.2%	**56.4%**	31.1%	13.1%
		small	**173.7%**	65.7%	73.6%	81.7%
	25	big	11.6%	**66.2%**	20.4%	37.1%
		random	38.5%	**48.1%**	27.4%	19.2%
		small	**134.6%**	84.0%	79.9%	67.5%
	50	big	**76.0%**	52.7%	24.7%	38.8%
		random	**79.2%**	62.5%	0.6%	54.8%
		small	64.3%	**66.9%**	30.7%	60.2%
	100	big	**104.4%**	56.3%	32.5%	63.6%
		random	**107.3%**	67.8%	0.8%	86.5%
		small	**204.1%**	68.0%	30.9%	48.0%
	200	big	**89.0%**	58.9%	12.8%	82.0%
		random	**103.1%**	63.6%	1.3%	77.4%
		small	**99.0%**	59.4%	21.3%	51.8%
	500	all	**100.8%**	59.2%	20.8%	63.1%
	Mean		***95.8%***	*66.6%*	*30.3%*	*55.6%*
10 bps	10	big	127.6%	**129.5%**	47.8%	23.3%
		random	3.5%	**56.4%**	10.2%	-4.7%
		small	**68.3%**	65.7%	44.7%	52.7%
	25	big	46.1%	**66.2%**	0.2%	14.2%
		random	-41.0%	**48.1%**	5.8%	-0.9%
		small	**114.3%**	84.0%	49.5%	39.3%
	50	big	**53.8%**	52.7%	3.6%	15.4%
		random	55.2%	**62.5%**	-16.6%	28.4%
		small	**107.7%**	66.9%	8.5%	32.9%
	100	big	**69.4%**	56.3%	9.9%	35.7%
		random	**131.0%**	67.8%	-16.4%	54.7%
		small	**153.1%**	68.0%	8.5%	22.6%
	200	big	**74.1%**	58.9%	-6.5%	50.8%
		random	**70.2%**	63.6%	-16.0%	47.0%
		small	**92.7%**	59.4%	0.5%	25.8%
	500	all	**86.0%**	59.2%	0.1%	35.1%
	Mean		***75.7%***	*66.6%*	*8.4%*	*29.5%*

Table 2: Out-of-sample cumulative returns at the end of the test period for each level of transaction costs and different portfolios.

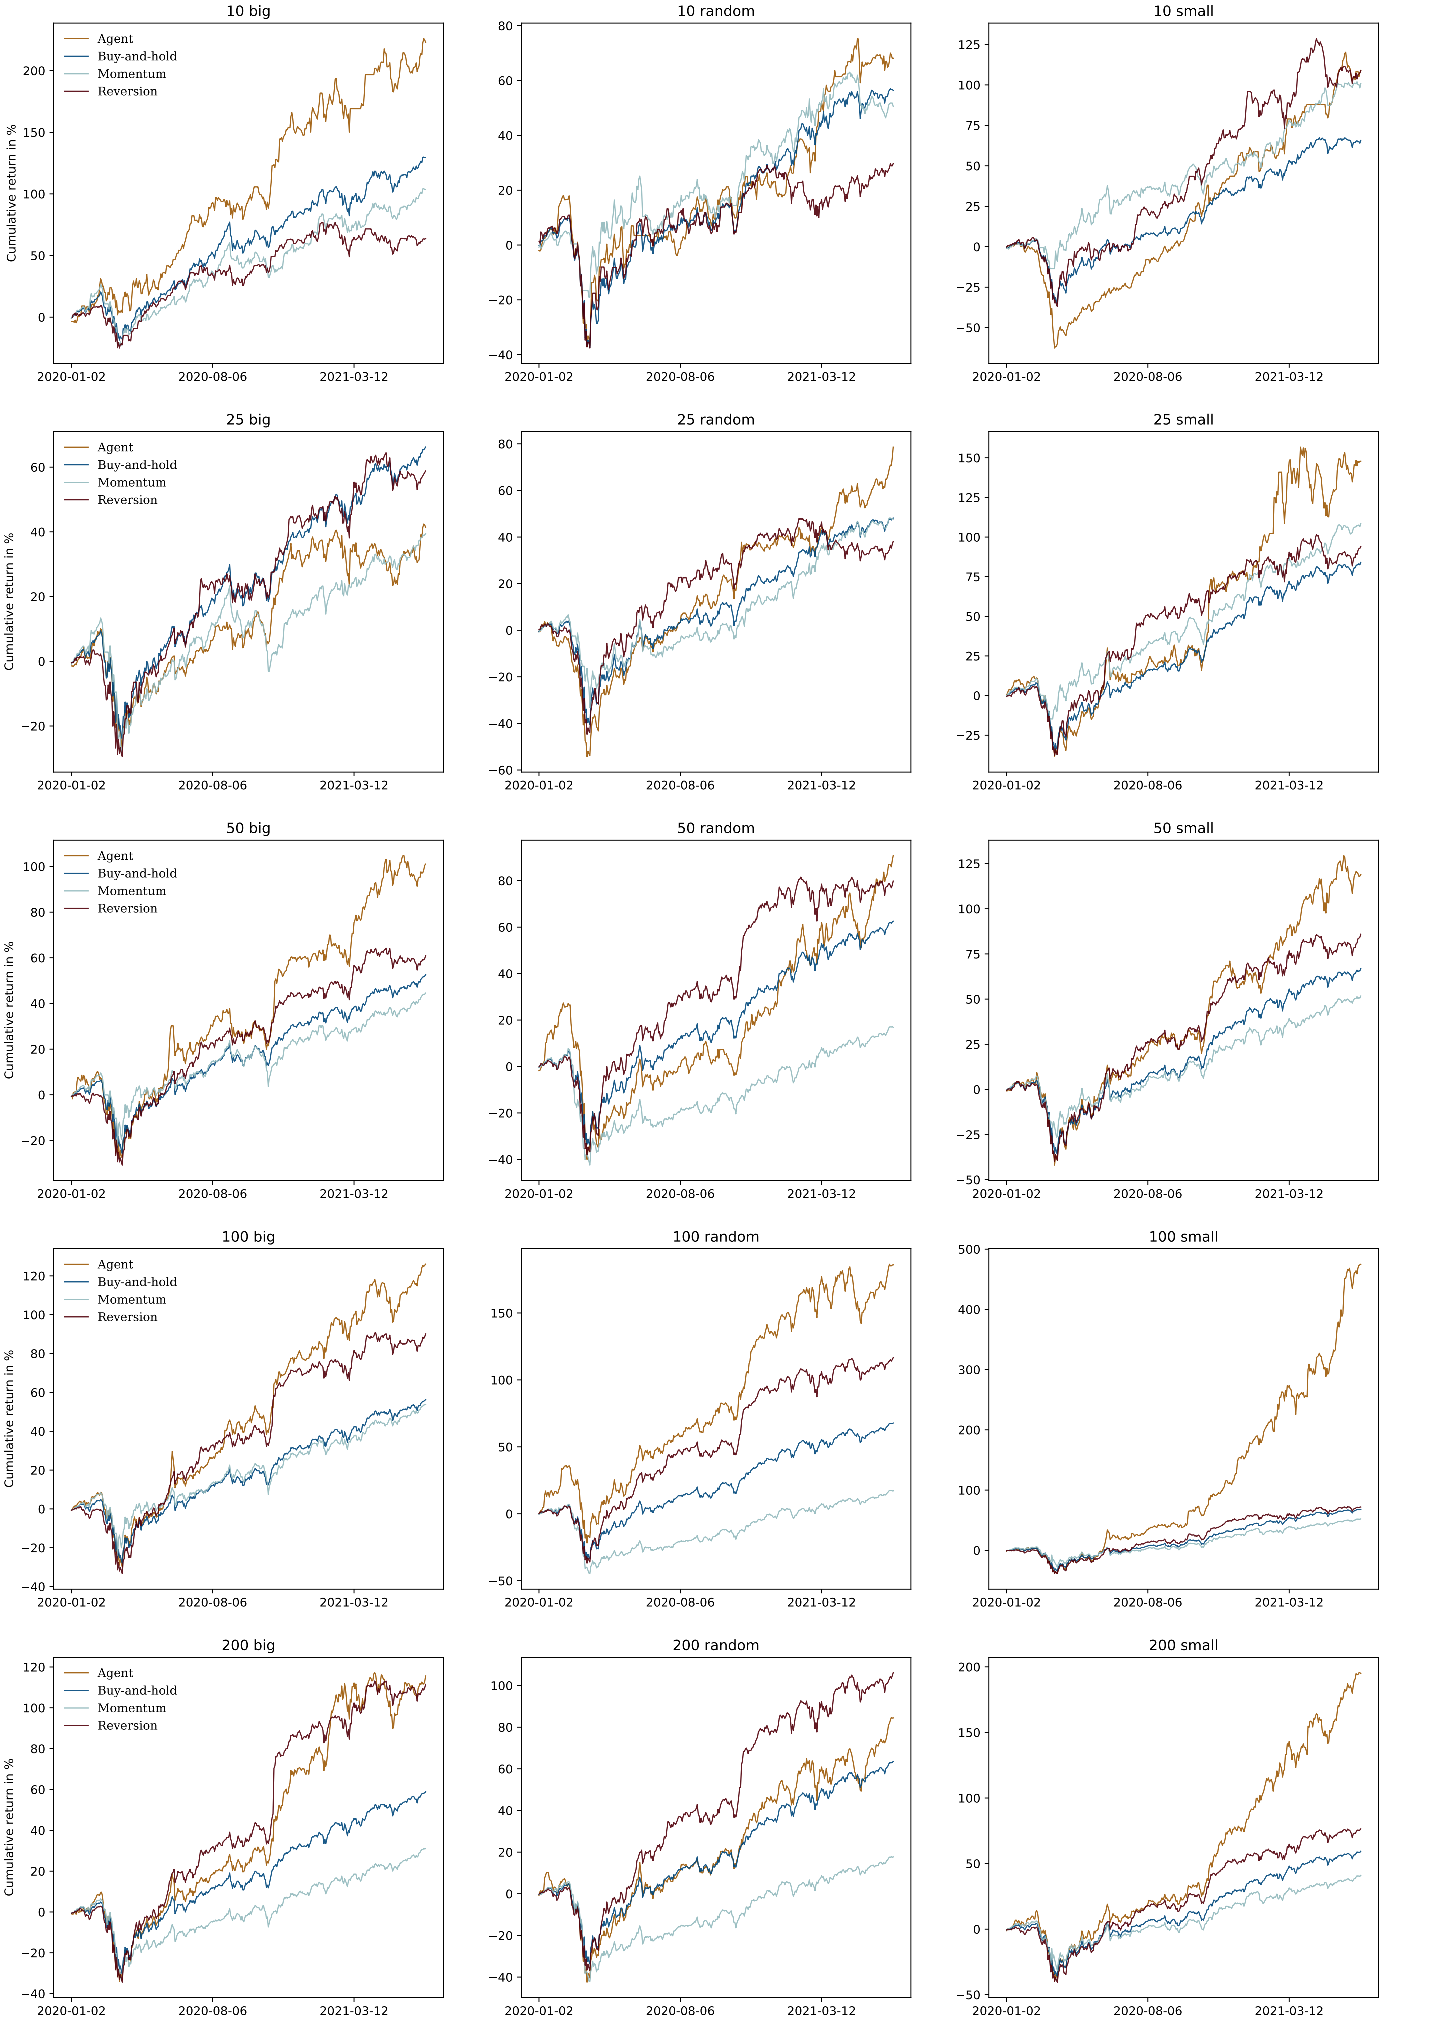

Figure 2: Test trading performances of the proposed agent compared to the benchmark strategies in the respective portfolio setup assuming transaction costs of 1 bps.

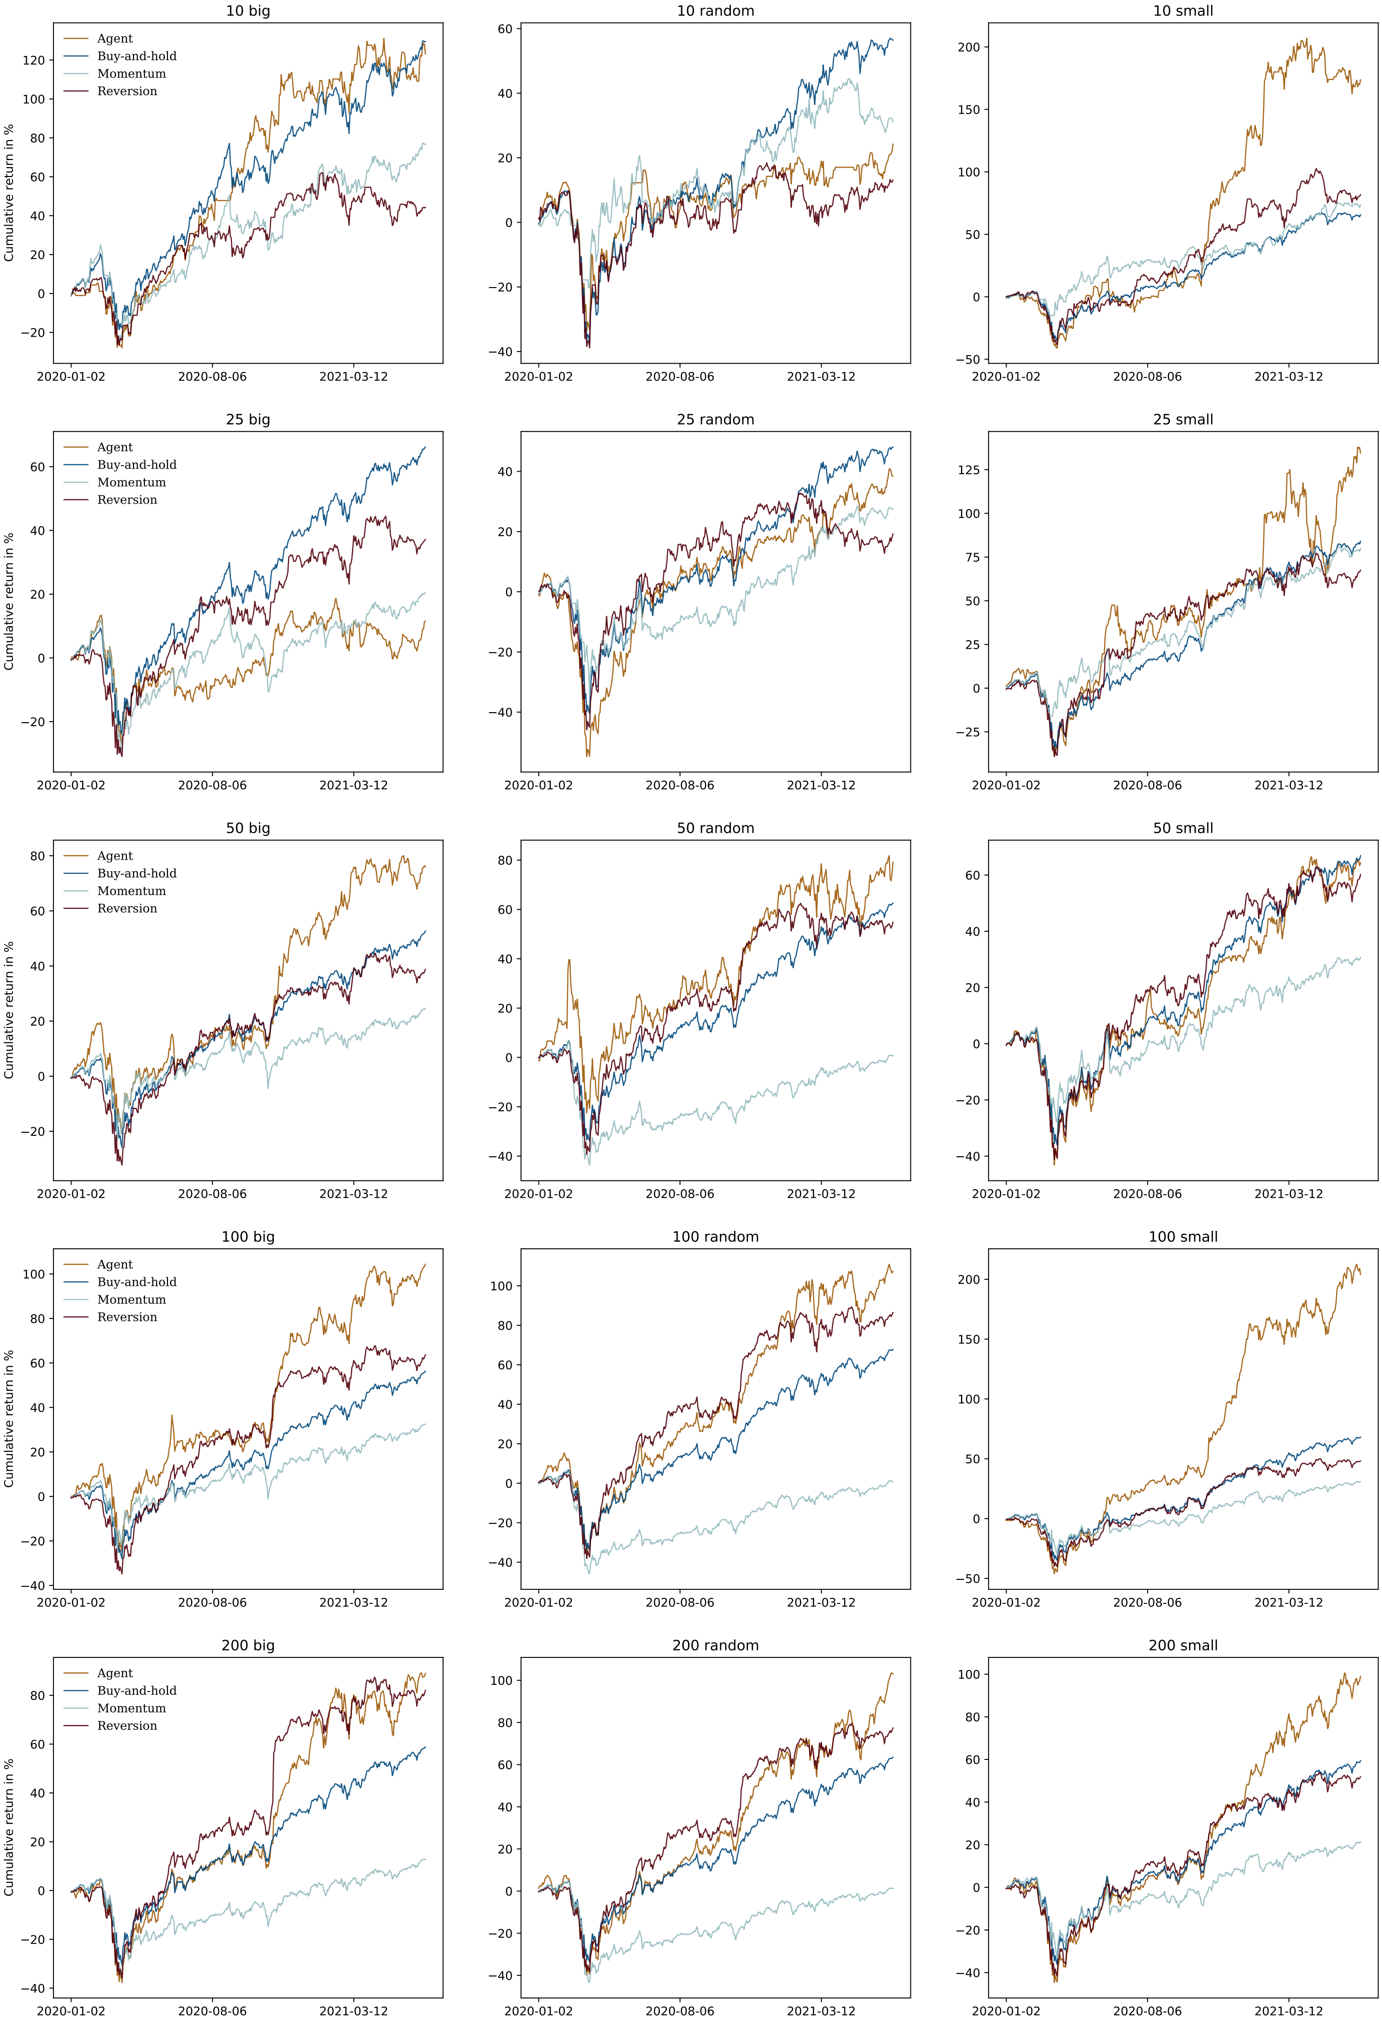

Figure 3: Test trading performances of the proposed agent compared to the benchmark strategies in the respective portfolio setup assuming transaction costs of 5 bps.

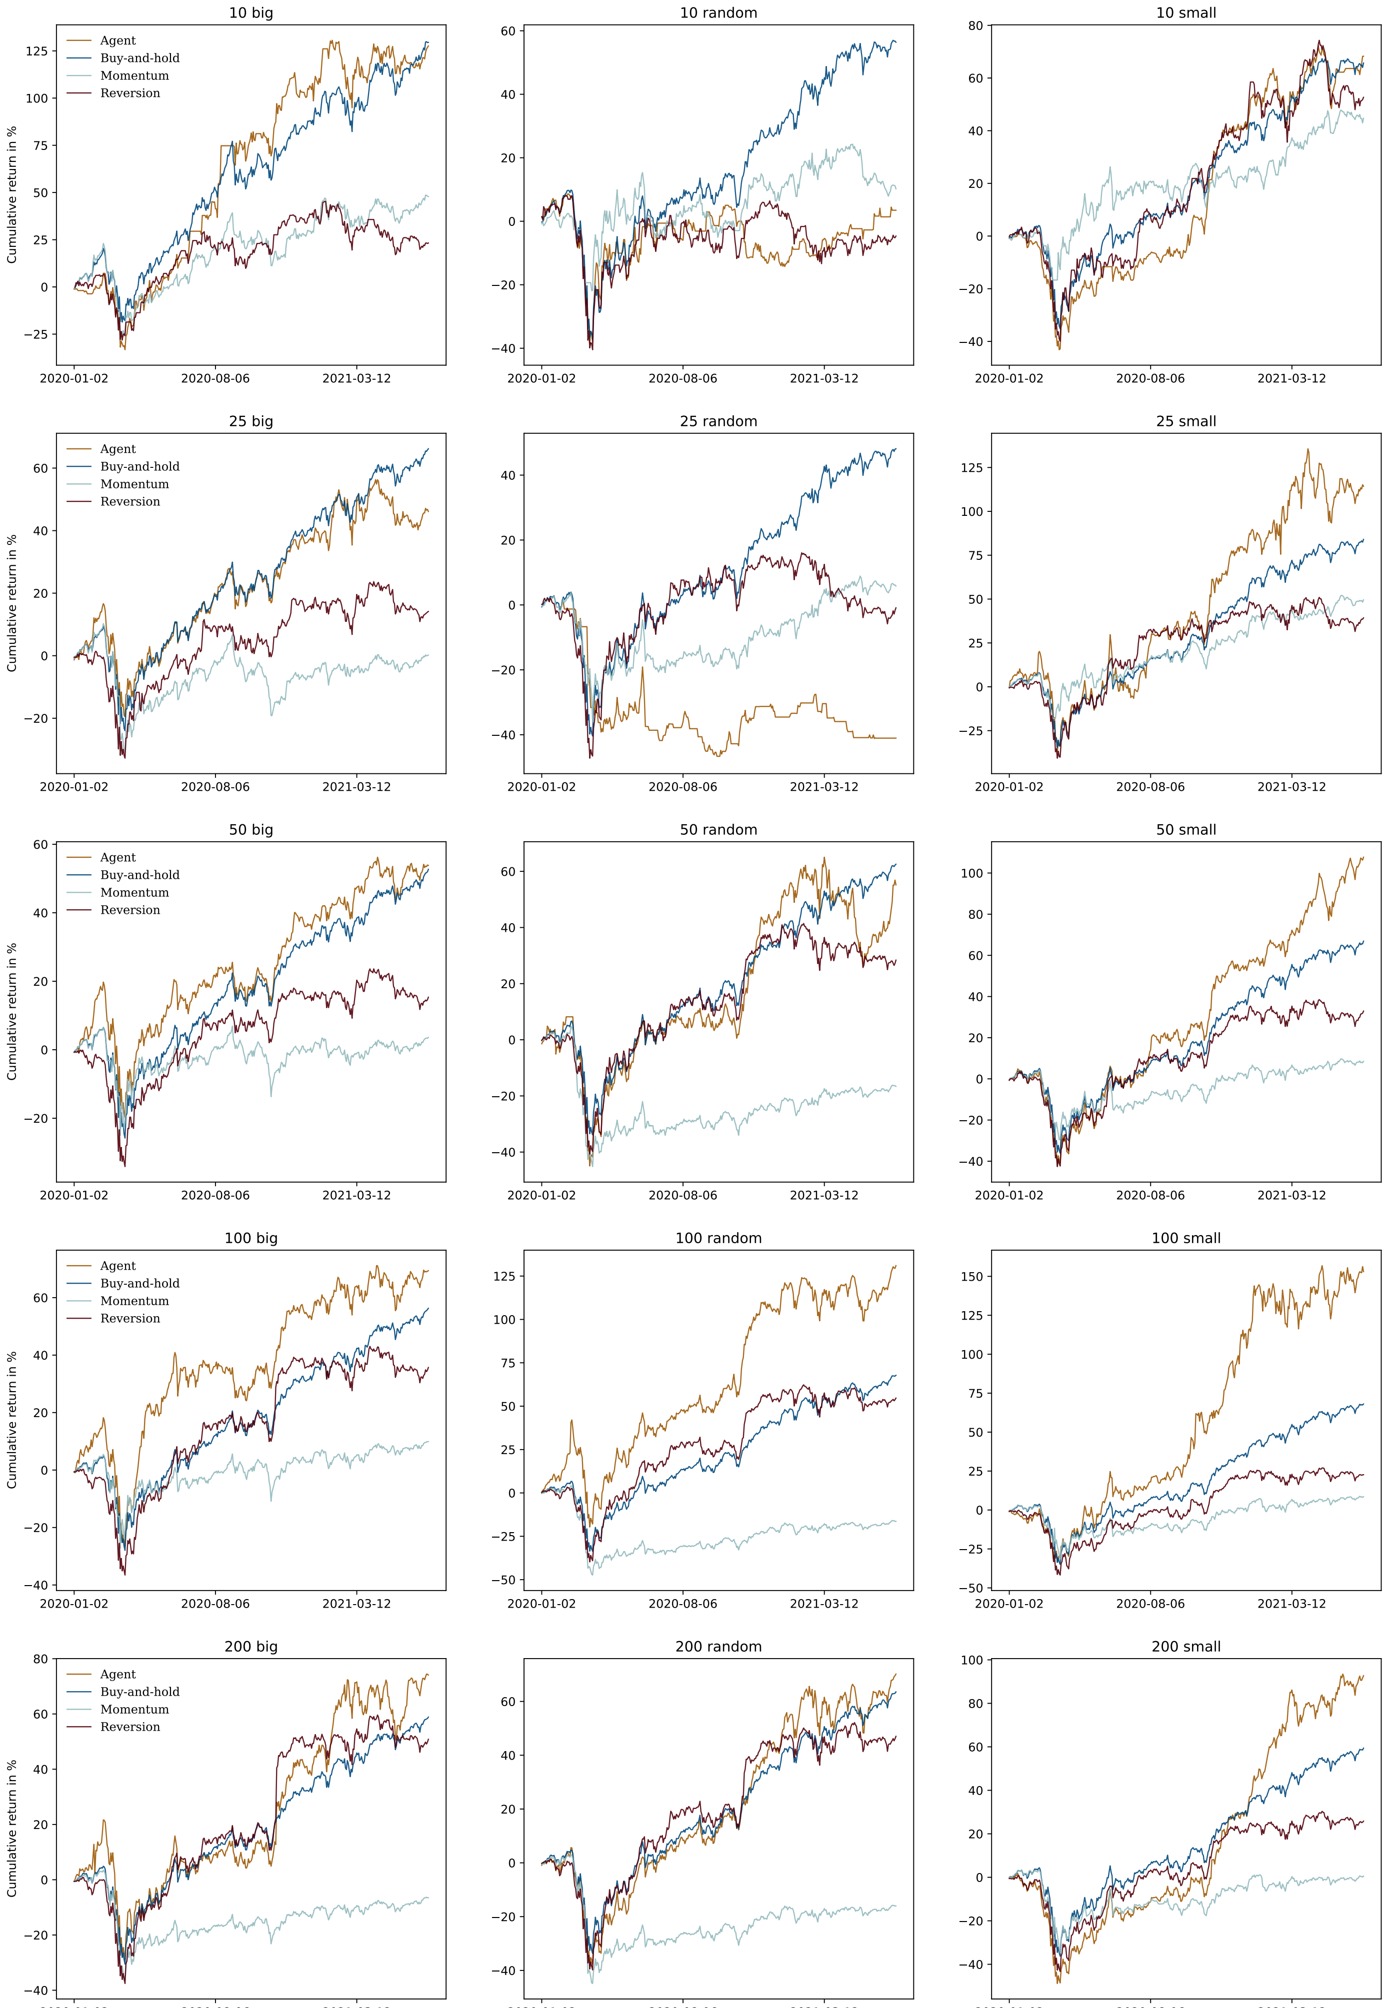

Figure 4: Test trading performances of the proposed agent compared to the benchmark strategies in the respective portfolio setup assuming transaction costs of 10 bps.

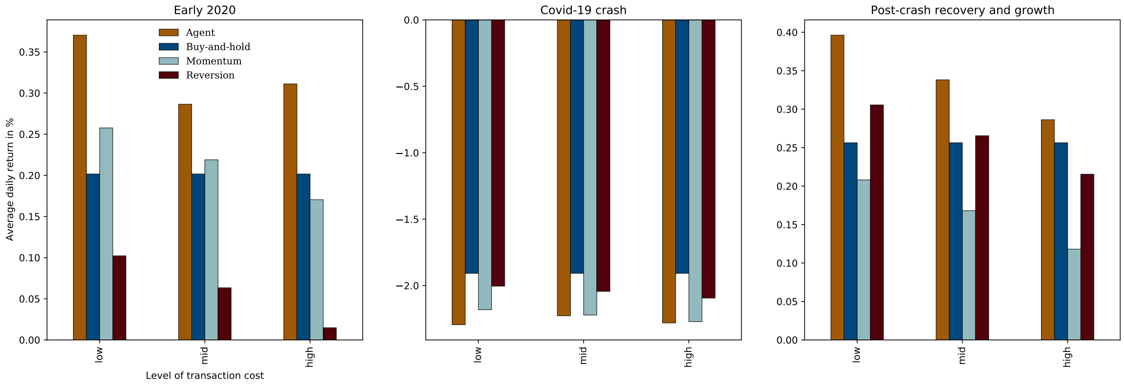

Figure 5: Out-of-sample trading performances for each level of transaction costs while splitting the test period into three phases based on market conditions.

Overall, the agent outperforms all benchmark strategies in 36 out of the 48 cases and achieves the highest mean cumulative returns in every transaction cost setup. Individually compared to each of the benchmark strategies, the agent outperforms in 37, 44, and 44 out of the 48 cases the buy-and-hold, momentum, which buys stocks with recent positive returns, and reversion strategy, which buys stocks with recent negative returns, respectively. Agent outperformance is more consistent and higher when portfolios are larger, i.e., when more than 50 stocks are included. This is to be expected as the training dataset is growing only in the number of stocks and, thus, contains more information when the number of stocks is increased. Moreover, the agent shows the strongest relative outperformance in portfolios with smaller stocks. Both the agent and the active benchmark strategies (momentum and reversion) suffer from increasing transaction costs. However, the agent shows adaptiveness as the loss in performance is notably lower compared to the active benchmark strategies. Although the momentum strategy performs worse than the reversion strategy on average, there are setups in which the momentum strategy is superior to the reversion strategy. This indicates that the agents need to be highly dynamic in terms of the learned strategies for each of the setups.

The test period can be split into three major phases around the crisis, starting in March 2020. Before, the set of assets yielded moderate, positive average returns in an uptrend, followed by sharp losses during the crisis. Afterwards, the stock selection quickly recovers and realizes large gains. The agent's achieved mean returns in each of the three periods are compared to the benchmark strategies in Fig. 5. We find that the proposed algorithm performs especially well at the beginning of the test period and even improves its performance when increasing the transaction costs from 5 to 10 bps. However, the strategy does not outperform the benchmarks during the stock market losses at the beginning of the COVID-19 pandemic. As there is no major crisis in the training set, this result may be expected. In the case of recovering markets and larger gains, we find that the agent outperforms all benchmark strategies on average, which is in line with the results in early 2020.

2.5 Conclusion

This Chapter proposes key extensions to Deep Q-learning (Mnih et al., 2013) to make the algorithm suitable for trading of financial asset portfolios. We generate environments with single assets randomly drawn from a set of assets and put trading returns in the isolated environments into a portfolio-oriented perspective by rewarding cash reservations with the mean return of the set of assets. Furthermore, we use a validation set and reserve the best-performing parameters in the Q-network to prevent overfitting. We use an ensemble of agents, which further decreases overfitting and makes our methodology highly flexible in terms of the number of assets.

With only a single hyperparameter setup, we test our methodology on 16 U.S. stock portfolio configurations, which vary in the number of stocks as well as the selection criteria. We benchmark the performance against a passive buy-and-hold and an active momentum and reversion strategy. Furthermore, we test three levels of transaction costs. The proposed methodology shows promising results, outperforming all benchmarks in 75% of the setups. We find the algorithm to generate excess returns in rising market environments; however, it cannot avoid the sharp losses during the beginning of the pandemic. These findings motivate further extensions, such as the inclusion of short sales and accounting for risk, e.g., using the Sharpe ratio.

Overall, the results are promising in terms of their prospective economic gains. DRL, however, comes with the drawback of high computational costs. That is, training the agents' policies according to our proposed algorithm can hardly be executed in parallel. Thus, the approach may be challenging to incorporate in a, e.g., high-frequency trading setting and should rather be reserved for short- to medium-term investments.

3 A Deep Learning Stack for Asset Pricing and Portfolio Management

In the preceding Chapter, we demonstrated the effectiveness of deep learning and, in particular, of DRL. We find that investors can achieve substantial gains in cumulative returns and Sharpe ratios when benchmarked against passive or even active investment approaches. Although these empirical results are promising, DRL for portfolio management comes with two major drawbacks. First, as the training procedure not only involves the optimization of the policy, i.e., updating the neural network's parameters via gradient descent, the agent also has to take steps in the environment, i.e., generate training samples. As this process is generally hard to parallelize, DRL is usually computationally highly expensive. Second, making the latter point even more important, finding hyperparameters, such as the exploration rate or the experience memory size, that are well suited for the environment and ensure that the policy converges to the optimal policy in a reasonable amount of time may become a taunting task. Moreover, the high sensitivity of the outcome of the final policy to the specified hyperparameters makes DRL a challenging, albeit feasible, methodology for portfolio management.

In the next Chapter, we show how to overcome potential computational constraints and make training significantly more robust, i.e., less sensitive to the choice of the hyperparameters. For that, we break down the task of stock return prediction into smaller and less complex subtasks in a supervised learning environment. This improves the final accuracy of our neural networks and makes them less exposed to the choice of hyperparameters, as we demonstrate that one simple configuration is sufficient for all networks in our sample. Furthermore, we show that this approach is computationally trivial. Lastly, and perhaps most importantly, our two-stage approach gives critical insights into feature importances and enables the investor to identify sources of risk exposure in the final portfolio choices. That is, the investor may discover which features cause the networks to decide on predicted return signs and magnitudes individually, making the overall prediction easier to interpret.

It is important to note that the methodologies utilized in Chapters 2 and 3 are significantly different in their respective targets. Although both approaches share a comparable set of features and data, their training procedures are distinct. That is, DRL updates the network's parameters based on the rewards received from performing (partially random) actions in an environment. Based on the combination of actions and rewards, the neural network may identify general patterns that increase its long-term rewards. In Chapter 3, however, we directly train neural networks to accurately predict return signs and magnitudes, which can lead to substantial differences in the final parameters of the neural network. Investors may then incorporate these predictions into their investment decisions. That is, the training procedure in Chapter 2 already takes into account the investors' preferences in the reward function, whereas these preferences are ignored when training neural networks to predict return signs and magnitudes in the supervised learning framework of Chapter 3. Due to the fact that the approaches differ in their formulation of the target, we do not directly compare them in the subsequent analysis.

Abstract

This Chapter develops a novel neural network-based approach for portfolio management. Our stack of neural networks distinctly predicts the direction and magnitudes of stock returns, making the feature importance clearly allocatable to each prediction subtask. In large-scale out-of-sample tests, we demonstrate the approach's high and robust performance in return predictability and resulting portfolio Sharpe ratios as well as risk-adjusted alphas. Zero-net investment long-short portfolios yield returns that double passive buy-and-hold portfolio performances. Contrary to the reputation of deep learning being computationally heavy, we show that our neural network stacks can be trained in a few seconds to manage vast portfolios by rethinking the training architecture.

3.1 Introduction

Machine learning may become one of the most important advances in recent human research. In the past few years, machine learning has made tremendous progress. The advances are highly prevalent in many subjects, besides many others, such as natural language processing (Brown et al., 2020), image recognition (He et al., 2016) or even a symbiosis of these subdomains (Ramesh et al., 2022). The progress in artificial intelligence also leads to deeper insights into financial markets and asset pricing (Bianchi et al., 2021; Gu et al., 2020). Motivated by the valuable implications of predictive quantitative systems for portfolio management, a large literature is growing that covers portfolio management approaches fully conducted by machine learning forecasts (Almahdi and Yang, 2017; Fischer and Krauss, 2018; Gu et al., 2020; Kraus and Feuerriegel, 2017; Krauss et al., 2017; Patel et al., 2015; Wang et al., 2020; Zhong and Enke, 2017).

Despite the powerful results that the latter research achieves, returns remain notoriously difficult to predict. The vast unpredictability of stock returns makes the financial application of deep learning extraordinarily challenging. That is, the out-of-sample coefficients of determination for return predictions are close to zero (Gu et al., 2020), rarely exceeding even one percent. Based upon this issue, the literature combines multiple predictors, forming an ensemble of machine learning techniques and averaging the resulting predictions, as, e.g., in Krauss et al. (2017). Alternatively, De Prado (2018) proposes to predict return signs by a high-recall model and to train a second model to evaluate the chance of a false-positive classification of the primary model. Although some performance gains can be achieved using these ensemble structures, predicting returns remains tough. To this end, this Chapter develops a network architecture that exploits the predictability of stock return values *after* the sign is resolved. Thus, instead of implementing a secondary model to assess the initial predictions made by the first model, as suggested by De Prado (2018), our objective is to reduce the complexity of the prediction process by breaking it down into several smaller tasks.

Anatolyev and Gospodinov (2010) suggest higher predictability in return directions and return volatility and, hence, predict the sign and absolute values of returns individually and combine these predictions using a copula. We connect to this idea, however, use an alternative decomposition scheme. That is, we also predict the sign of returns as in Anatolyev and Gospodinov (2010) first, but rather than predicting return volatility, we train models for both positive and negative returns individually. We expect that neural networks may identify a different set of predictors to predict positive and negative returns, motivated by the fact that returns behave differently around positive and negative return shocks (Ederington and Guan, 2010). Furthermore, this scheme allows us to analyze potential differences in features importance between the model predicting positive return magnitude and the model predicting negative return magnitudes.

Our approach consists of training three neural networks. Each neural network is supposed to fulfill a specific subtask in the prediction process. Predicting is done via two stages. The first stage consists of one neural network, the "directional model", that predicts the sign of the return. This neural network is trained via the entire training set using binary classified labels, indicating whether the target return is positive or negative. The second stage consists of two more neural networks, the "long model" and the "short model". Whenever the directional predicts a positive (negative) return, we feed the features of the given sample into the long (short) model to obtain the predicted positive (negative) magnitude of the prediction. We train the long (short) model by using target returns of only those samples of the training data, for which we observe a positive (negative) target return.

Using the two-stage approach, the long (short) model achieves an average out-of-sample R^2 of 11.1% (15.5%).[1] Moreover, in large-scale out-of-sample tests, we demonstrate the benefits of this approach: after accounting for transaction costs, the stack yields statistically significant and risk-adjusted alphas in long-only portfolios.

[1] We limit predictions to the minimum and maximum values of the training set in order to account for the dramatic input vector changes during the Covid-19 drawdown in 2020. This way, the mean R^2 is close to the median. This is only done for convenience of interpretation and does not affect any trading results.

The constructed long-short portfolios are market-neutral while still doubling average daily market returns.

While predictive capabilities and performance gains are important, economically understanding high-dimensional and complex models is arguably highly valuable. To uncover the insights of neural networks, which are typically referred to as "black box models", the literature focuses on feature importance measures. For instance, Gu et al. (2020) report individual feature contributions across multiple machine learning techniques such as boosting (Friedman, 2001), random forests (Breiman, 2001) and deep neural networks. As these methods address feature importance from the input side, we contribute to the literature of financial machine learning by revealing the differences in feature importance based on the target formulation. That is, we analyze the feature importances conditional on the target, i.e., based on whether the model predicts return directions or magnitudes. In particular, our stacked neural network approach shows that risk-adjusted alphas and large Sharpe ratios are gained by binary directional return predictions. The follow-up long and short models, however, contribute to considerable gains in returns, i.e., leveraging both risk and returns while keeping the Sharpe ratio constant. Hence, the results suggest that, e.g., investors with low risk aversion may especially benefit from the latter magnitude-based models while risk-averse investors should rather focus on the directional model.

The choice of neural network hyperparameters is crucial for its performance. The training time of neural networks mainly depends on the accuracy of the gradients and the learning rate of the optimizer. Typically, memory constraints limit the training to larger batches. Furthermore, small batches are supposed to have a regularizing effect. The finance literature adopts small batch training schemes, even though the input vector is typically quite small. We argue that neither memory constraints nor regularization issues prevail when training financial predictive neural networks. As a result, our networks train in a few seconds using the GPU alone. This way, we are able to set up a walk-forward backtest consisting of 20 randomly sampled portfolios with 500 stocks each. Although this gives 600 neural networks to train, the computations can be accomplished in several hours, whereas we expect the computational costs to dramatically decrease in the future.

3.2 Neural networks

Due to common factor models as in Fama and French (1993), the linear regression model is the most popular specification of the functional relationship between the cross-section of returns and firm characteristics like the size or value of the company.

However, polynomial specifications or interaction variables are commonly neglected. Many researchers, instead, introduce new factors that may describe and predict the cross-section of stock returns, producing a "factor zoo" (Feng et al., 2020).

Neural networks avoid manual feature engineering and, given sufficient capacities, can approximate any arbitrary function. For that, consider the objective of identifying a parameter vector θ such that the function $\hat{f}(\mathbf{x};\theta)$ approximates a vector of target values $\mathbf{y}$ given a feature vector as precise as possible. The quality of the approximation is defined via a loss function, for instance in a regression task, the mean squared difference between estimates and true values, i.e., the Mean Squared Error (MSE).

A neural network is a directed graph consisting of interconnected layers of neurons, where each neuron represents a non-linear transformation using an activation function on a weighted sum of the input vector, expressed as $f(\mathbf{w_n} \cdot \mathbf{a_n})$, where $\mathbf{a_n}$ is the input vector of neuron n and $\mathbf{w_n}$ is a vector that contains weight parameters[2]. A commonly used activation function is the ReLU, defined as $ReLU(x) = max(x, 0)$. However, many other activation functions, such as the sigmoid-function, are well-suited specifically for classifications tasks to transform outputs to probabilities.

The neural network becomes capable of approximating complex functions when chaining multiple neurons, i.e., non-linear transformations of the inputs. For instance, a simple neural network may consist of three non-linear transformations f_1, f_2, f_3 of a single feature x resolved by a linear output function o giving the final estimation $\hat{y}$:

$$\hat{y} = o(f_3(f_2(f_1(x|w_1)|w_2)|w_3)|w_o) = w_o \cdot ReLU(w_3 \cdot ReLU(w_2 \cdot ReLU(w_1 * x))). \quad (4)$$

The functions f_1, f_2, f_3 are referred to as "hidden layers", while o is the "output layer" and x the "input layer". In this example, the hidden layers contain only one neuron each, however, as the estimation issue becomes more complex, one may include multiple neurons per layer. If multiple neurons per layer are incorporated, w_o, w_1, w_2 and w_3 become vectors and if simultaneously multiple features are given, these latter vectors of weights become matrices of weights. Generally, a network becomes "deep" when adding more hidden layers, while one increases its "width" by adding more neurons to the hidden layers. The set of parameters of the model is typically denoted as θ.

Neural network parameters are optimized via gradient descent. A gradient descent algorithm optimizes the parameters iteratively in small steps using gradients on

[2]Note that we omit the usage of a bias for brevity, however, they should be added to each neuron in the neural network.

a loss function. For instance, in a regression task, the loss function is commonly defined as the squared error. Given the vector of neural network predictions $\hat{\mathbf{y}}(\theta)$, that are a function of the parameter vector θ, and the vector of true values $\mathbf{y}$, we compute the loss as:

$$\mathcal{L}(\theta) = ||\mathbf{y} - \hat{\mathbf{y}}(\theta)||_2^2.$$

More intuitive and interpretable linearly transformed variants of the loss-function are suitable as well, for instance, the mean squared error.

The direction of steepest descent in loss lies in the opposite direction of the loss gradient. The loss gradient is a vector of first-order parameter partial derivatives on the loss function:

$$\nabla\mathcal{L}(\theta) = \begin{bmatrix} \frac{\partial L(\theta)}{\partial w_1} \\ \vdots \\ \frac{\partial L(\theta)}{\partial w_n} \end{bmatrix}.$$

Updates on the parameter vector are straightforward, with $\theta \leftarrow \theta - \eta\nabla\mathcal{L}(\theta)$, where η is the learning rate. The learning rate controls the intensity of parameter changes; however, if it is too high, the parameter vector θ will never reach optimality as it oscillates around the optimum. Researchers proposed several improvements to the algorithm in the recent years, such as including momentum. The Adam (Kingma and Ba, 2014) optimizer includes both momentum in the gradients and a decaying, adaptive learning rate using second-order derivatives. For a more detailed outline of neural networks, see Goodfellow et al. (2016); Hastie et al. (2009).

In order to increase the efficiency or even feasibility of gradient descent optimization, one may not compute the loss and gradient over the entire data set but rather do so in small batches. While this reduces the computational effort, it also introduces uncertainty in the gradients, usually referred to as stochastic gradient descent. However, we outline in Section 3.3 that maximizing the batch size might be favorable for stock return prediction applications.

3.3 Data and network designs

For our empirical analysis, we merge daily price data from Center For Research in Security Prices (CRSP) with fundamental data from Compustat. The sample ranges from 2000 to 2021. From the merged dataset, we select the 2000 stocks that have the fewest data gaps. This gives over 2.6 million data points. Based on this, we generate the following features, which we already discussed in Section 2.4.1: moving averages (both normally and exponentially weighted over 5 to 200 periods) of past returns and return standard deviations as well as the latest values of the following firm characteristics: sales per share, gross profit margin, operating margin, net profit margin, return on equity, return on assets, current ratio, debt ratio, market capitalization, and book to market value. We winsorize the fundamental data at the 0.5%-level at the top and bottom. Lastly, we drop missing values, resulting in 1032 companies and a dataset containing over 2.1 million datapoints. This large loss of stocks and data points is due to a high fraction of missing values in the fundamental data as well as the initialization periods required for the long-term moving averages of the return history.

Predictive models that are used for portfolio management are typically evaluated by their out-of-sample performance. The performance, however, is heavily dependent on the design of the tests. For instance, it is inherently more difficult to provide reasonable performance out-of-sample when considering a large cross-section of stocks rather than just an industry. Many studies consider short out-of-sample periods, small portfolios, or even both. Meanwhile, researchers undertake a complex grid search to optimize lots of hyperparameters. Computational power can make models seem to perform well, even out-of-sample, if the number of test samples is small or transaction costs are neglected, which may result in false conclusions. (Chordia et al., 2020; de Prado, 2020; Harvey and Liu, 2021; Hou et al., 2020; Novy-Marx and Velikov, 2016). Therefore, we aim to design an out-of-sample performance evaluation that is feasible while minimizing the chance that we were just lucky.

To assess the effectiveness of our model design, we want to minimize the chance of obtaining out-of-sample results, that originate from the selection of stocks and the prediction horizon and, hence, do not necessarily unveil the robustness of our approach. Hence, we sample 20 portfolios that originate from the 1032-stock dataset, where each portfolio consists of 500 stocks. Note that the effective number of available stocks may vary each day depending on data gaps. We use a rolling window approach for a small hyperparameter search and for testing, which is illustrated in Figure 6. Validation and testing are done in half-year cycles, i.e., we train on the available history and validate or test it in the following half-year. We search for hyperparameters using the validation set from January 2015 to December 2016 and select those that result in the highest average cumulative return in the validation sets. The algorithm is tested from January 2017 to December 2021.[3] In order to

[3]Computing power is rising exponentially over time. The exponential rise leads to unrealistic results if tested too far into the past, as the computational and transaction costs would have been much higher. Furthermore, there was more limited access to deep learning only a few years ago. For instance, consider the two largest open source deep learning framework releases: Tensorflow in

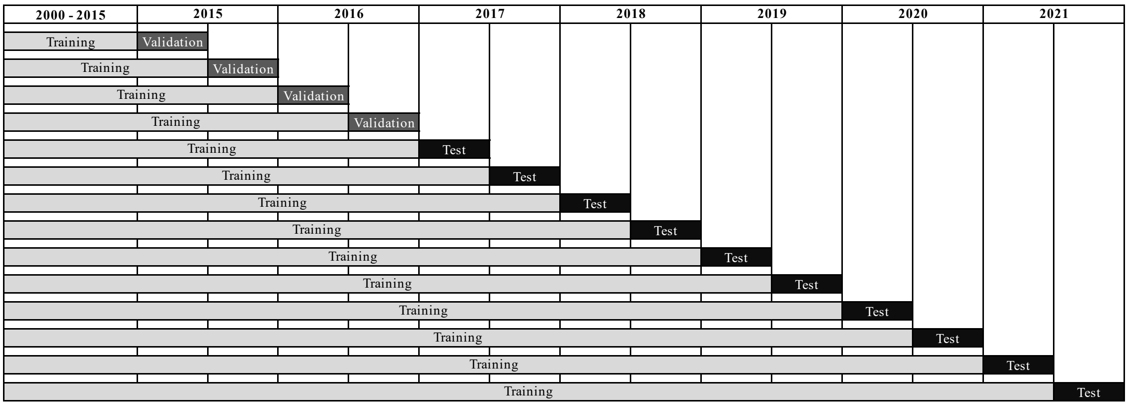

Figure 6: Rolling window out-of-sample testing approach setup.
This figure illustrates the out-of-sample testing scheme we apply for testing our neural network stack. We sample portfolios containing 500 stocks and train, validate and test the performance of the neural network stack according to the illustrated rolling window approach.

avoid any look-ahead bias, we chronologically first choose hyperparameters based on the validation data set and use more recent data, i.e. after December 2016, to test the algorithm. Overall, this gives 1259 test days per portfolio, resulting in 25,180 total test days. The portfolios are evaluated daily, and transaction costs of 5 bps are assumed, which are, however, discussed in greater detail in Section 3.4.4. Therefore, we are confident that our results are not biased by the selection of the out-of-sample test criteria.

For each portfolio and test period, we train three neural networks using the training data which jointly predict returns out-of-sample using the test data set. First, the directional model is trained on the entire training data set, i.e., including positive and negative returns, and learns to predict the next-period return sign. The binary classification model is trained on binary labels based on the sign of the return. That is, the label is assigned a value of one (zero) if $r_{t,i} \geq 0$ ($r_{t,i} < 0$). We use a sigmoid activation in the output layer of the directional model, bounding its predictions within the range of zero to one. Consequently, the model's output can be interpreted as estimated probabilities for a positive return in the subsequent period. The value of returns is predicted via two more regression neural networks, the long and short models. The long (short) model, i.e., the model to predict the magnitude of positive (negative) returns, is trained exclusively on the positive (negative) returns of the training data set. In order to obtain out-of-sample predictions, we first compute predictions of the returns sign using the directional model, and then predict the magnitude using the long (short) model, if the directional model predicts a positive

2015 and PyTorch in 2016.

(negative) return. For out-of-sample predictions on the test data, long and short models are engaged based on the directional model's predictions.

Long-short portfolios are constructed based on a ranking of stocks. We follow the simple approach of Gu et al. (2020) and build zero-net-investment portfolios by buying the top decile and selling the bottom decile stocks in terms of predictions. Therefore, the portfolio return is the return spread between the long and short positions. Consequently, the top decile of stocks is identified by the long model, while the bottom decile is identified by the short model, requiring that half of the stocks' returns are predicted by the long model and the other half by the short model. However, the directional models' training predictions are not centered around 0.5. Depending on which historical years and stocks are included, the directional model may predict significantly more positive return signs than negatives, or vice versa, if we encode any prediction above greater or equal than 0.5 as a positive return sign. Therefore, in order to address this concern, we choose to set the cutoff at the training median prediction for each portfolio. However, the results are qualitatively similar when choosing the threshold at 0.5.

Deep learning requires many computational resources. However, financial datasets usually contain far fewer input features than are required to train models for, e.g., natural language processing or image recognition. Simply put, consider a 100 x 100 pixel image dataset. Each data row has 10,000 features, one for each pixel. In fact, a 100 x 100 image is still at a very low resolution. High-resolution images make training in batches of, e.g., 32 images absolutely necessary as the memory requirements overwhelm typical GPU memories. As our training set only contains 28 features, a modern GPU's memory is sufficient to store the entire dataset at once. Precisely, our largest training dataset requires less than 120 megabytes of memory. Storage is also required by the model itself. Financial models are typically rather small, as we face a low signal-to-noise ratio and may quickly overfit, i.e., show great results in-sample, but perform poorly out-of-sample. Overfitting is generally less prevalent when using smaller models. Here, our models only use about 25 kilobytes of memory. Lastly, forward and backward passes require memory that is close to the model itself. Overall, our neural network stack's memory requirements pose no challenge to the memory available in modern GPUs.

Each model comes with the same hyperparameters. We incorporate a similar structure as in Section 2.4.2 and use two hidden layers, each consisting of 64 neurons, and incorporate the Adam optimizer to train the neural network (Kingma and Ba, 2014). The internal network structure is not fine-tuned and is selected based on the rate of increased training performance per time unit. Fine-tuning hyperparameters can be a daunting task and does not impact the main results of this Chapter as we find negligible sensitivity of the final results to the chosen network structure. Yet, the number of epochs must be limited to effectively prevent overfitting. Using the validation data, we test the performance of the neural network stack using the following number of epochs: 100, 250, 500, 1000, and 2000. While the stack begins to overfit at approximately 2000 epochs, the stack appears to relatively underfit for epochs of 100, 250, and 500 epochs. Training the models for 1000 epochs leads to the best results out of the five candidates on the validation data.

Gradients are typically not computed using the entire training set, but rather estimated by sampling batches from the training data due to memory limitations of the GPU. However, if applicable, computing the exact gradient using the entire training data can be highly favorable out of two reasons. First, training time is split up into two components when training in batches. Each batch needs to be collected from the dataset and transferred from the CPU to the GPU. The GPU forward passes the batch, calculates the gradient, and backward passes the parameter adjustments. Hence, if gradients are computed on the entire training data, the latter step is not required, eliminating a significant portion of required computations.

Second, each batch contributes to the convergence of the neural network. However, input matrices for return predictions contain false signals and noisy data. A high level of noise leads to little or even a negative contribution to the convergence towards the global loss minimum if the batch size is small. Ultimately, we find that using the largest possible batch size mitigates unexpected jumps in the parameters. By utilizing the entire data set for each gradient descent step, the network's predictive ability is guaranteed to increase in the training data, while outliers may have less of an impact. Due to the fact that data transfers between CPU and GPU are not required and the parallel processing of the GPU is maximized, directional models are trained in less than 13 seconds. Moreover, the long and short models are trained in less than seven seconds.[4] The small training time results in a high ratio of backtested portfolios per time spent on computations.

3.4 Results

3.4.1 Long portfolios

This section describes results for long-only portfolios using the directional into long model stack. We report average daily returns, Sharpe ratios and cumulative returns

[4]We use a Nvidia GeForce RTX 3080 for our analysis.

of equally weighted portfolios on the predicted top decile of stocks. Each day, we compute return predictions using the neural network stack. These measures are compared to a passive, equally weighted buy-and-hold portfolio, containing all 500 stocks in each period. All returns are cumulated daily. For each transaction, we assume costs of 5 bps. Detailed results are shown in Table 3. We also provide an estimate on the level of transaction costs that makes the active strategy on par with the passive benchmark in Section 3.4.4.

The model stack outperforms the benchmark portfolio in terms of cumulative returns and Sharpe ratios in every randomly selected portfolio setup.[5] The outperformance is economically meaningful, as the average daily return is more than doubled. While the passive portfolio approach yields average daily returns of 0.08%, managing a portfolio via the proposed model stack results in 0.22% average daily returns. The large difference in average daily returns leads to wide gaps in the cumulative returns at the end of the test period. On average, the initial portfolio values increased by 1,031% after five years using the model stack and only by 148% using a buy-and-hold approach. The difference in Sharpe ratios is less severe. The stack realizes annualized Sharpe ratios of 143% compared to 87% using the passive approach. We also show cumulative returns of our long-only neural network stack along with the buy-and-hold benchmark portfolio in Figure 7 of the best, median, and worst performances generated by the neural network stack.

While the portfolios are similar in average daily returns, there is variation in the annualized Sharpe ratios. Precisely, the minimal annualized Sharpe ratio using the benchmark portfolio is 80.9% compared to the best setup with an annualized Sharpe ratio of 94.9%. This gap of 14 percentage points shows that there is indeed variation across the sampled portfolios, although some stocks overlap. The compositional differences increase the complexity of the test and show weaknesses in our methodology. In fact, in the worst-case portfolio composition, the stacked neural network strategy only generates six excess basis points in returns and nine percentage points in the annualized Sharpe ratio. Although the neural network stack still outperforms in terms of cumulative returns and Sharpe ratios, this portfolio reveals that the economic gains depend on the portfolio composition and that average, large cumulative return and Sharpe ratios relative to the passive approach are not guaranteed.

[5] We dissect the performance sources in Section 3.4.3 in more detail, where we also present a comparative evaluation between our proposed approach and the utilization of a single neural network, specifically the directional model, for predicting returns.

	Mean return		Sharpe ratio		Cumulative return	
	NNs	Benchmark	NNs	Benchmark	NNs	Benchmark
1	0.23%	0.09%	149.59%	91.51%	1138.76%	159.53%
2	0.21%	0.09%	135.52%	90.43%	886.52%	160.99%
3	0.23%	0.08%	144.34%	89.21%	1052.19%	149.99%
4	0.23%	0.08%	146.94%	84.99%	1084.38%	138.89%
5	0.23%	0.09%	147.72%	94.16%	1100.62%	167.96%
6	0.17%	0.08%	117.2%	82.29%	519.47%	134.37%
7	0.24%	0.09%	162.93%	94.9%	1392.67%	165.87%
8	0.25%	0.08%	164.35%	81.8%	1566.37%	131.8%
9	0.21%	0.08%	138.89%	83.23%	853.15%	136.73%
10	0.24%	0.09%	149.62%	92.48%	1231.09%	165.08%
11	0.23%	0.08%	148.7%	85.09%	1080.8%	140.64%
12	0.25%	0.09%	161.9%	87.93%	1484.98%	153.75%
13	0.18%	0.08%	115.37%	84.62%	529.78%	138.24%
14	0.21%	0.08%	135.53%	84.45%	819.45%	140.06%
15	0.14%	0.08%	90.58%	80.9%	298.76%	126.23%
16	0.22%	0.09%	150.42%	94.3%	1050.67%	167.47%
17	0.19%	0.08%	127.15%	87.33%	695.17%	148.98%
18	0.24%	0.08%	154.99%	83.78%	1280.01%	136.0%
19	0.24%	0.08%	162.68%	87.69%	1393.79%	145.38%
20	0.23%	0.08%	151.86%	85.51%	1152.93%	148.63%
Mean	0.22%	0.08%	142.81%	87.33%	1030.58%	147.83%

Table 3: Performance measures for long-only portfolios.
This table shows performance measures for long-only portfolios. Displayed are mean daily returns, annualized Sharpe Ratios, and end-of-period cumulative returns using the stacked neural network structure (NNs) and an equally weighted buy-and-hold strategy (Benchmark). The statistics are based on the out-of-sample period from 2017-01-01 to 2021-12-31. The portfolios are evaluated and cumulated daily. We assume transaction costs of 5 bps per transaction. The third column represents cumulative end-of-period returns, accumulating profits over the period of five years.

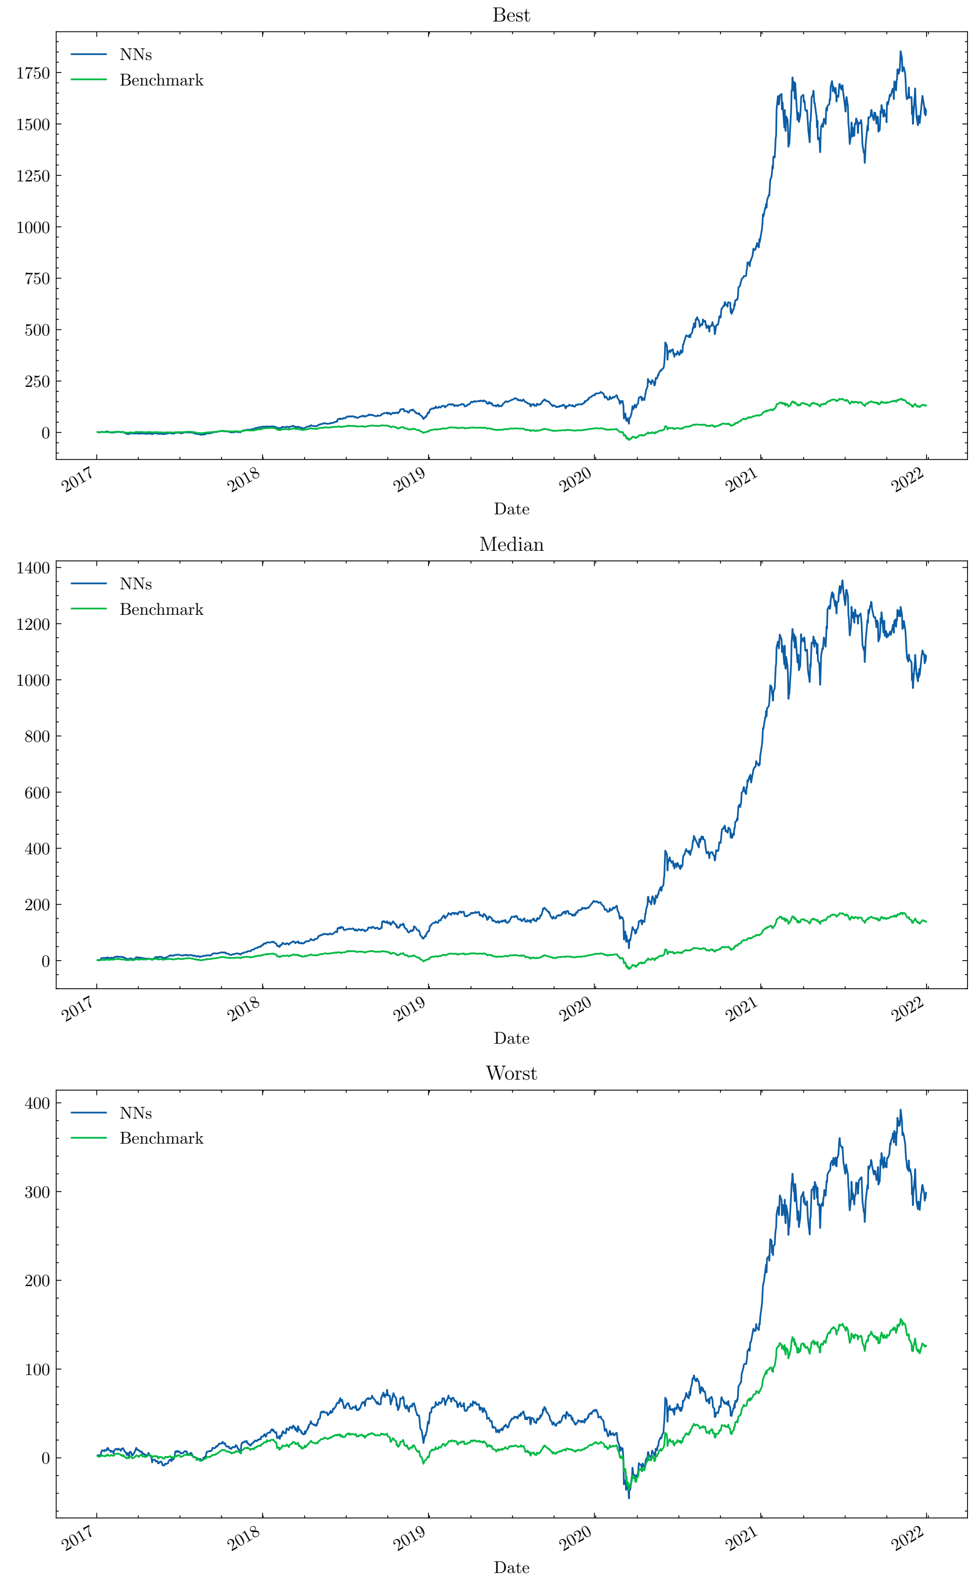

Figure 7: Performance plots for long-only portfolios.
This figure depicts cumulative returns for the long-only neural network stack (NNs) and an equally weighted buy-and-hold portfolio (Benchmark). Displayed are the highest, median, and lowest cumulative returns portfolios. All values are in percentages and cumulated using the full test set of five years. We assume costs of 5 bps per transaction.

	Alpha	Beta	R^2	N Observations
Best	0.00142***	1.43092***	0.79	1259
	(4.459)	(68.696)		
Median	0.00102***	1.47059***	0.79	1259
	(3.162)	(68.565)		
Worst	0.00028	1.48358***	0.81	1259
	(0.909)	(72.327)		

Table 4: Baseline OLS regression results for long-only portfolios.
This table presents results from baseline OLS regressions. We regress the neural network stack (NNs) returns on equally weighted buy-and-hold portfolio returns (Benchmark). Each randomly selected portfolio belongs to an OLS regression. Asterisks indicate the level of significance, i.e., p-values smaller than 10%, 5% or 1%, respectively. t-statistics are reported in parentheses. We assume transaction costs of 5 bps per transaction. Out of the 20 portfolios, we report the best, median, and worst results in terms of alpha. The average alpha is 0.1%, whereas 18 out of 20 portfolios exhibit a statistically significant value of alpha ($p < 5\%$). The average beta is 1.41.

The latter results show that, on average, the neural network stack significantly outperforms a passive, buy-and-hold benchmark in terms of returns and Sharpe ratios. To dissect the performance difference in greater detail, we regress the portfolio returns of the stacked neural network on the equally weighted buy-and-hold portfolio returns. This provides an estimate of how much the larger returns are due to a greater risk loading and how much they are due to the predictive capabilities of the network stack. Table 4 shows the results of a standard OLS regression. For brevity, we report the minimum, median, and maximum results measured in terms of the estimated alpha. Both the highest and median alpha values are 0.142% and 0.102% and significantly different from zero at the 1%-level. In the worst case, the stacked neural network strategy fails to generate a statistically significant alpha at any typical significance level. Still, a significant alpha is generated in 18 out of the 20 cases (5%-level) with an average value of 0.1%. Beta values vary less, with an average value of 1.41. That is, the strategies' return gaps can be explained by two underlying mechanisms. First, the stacked neural network strategy is riskier than the benchmark, i.e., it has leveraging effects on passive, equally weighted returns and, thus, creates larger returns. Second, the strategy predicts returns reasonably well, creating the positive and significant alpha on the passive benchmark, i.e., yielding higher returns without additional loading on risk.

The latter regression shows that the returns of the neural network stack are not entirely explained by the benchmark returns, i.e., returns of an equally weighted buy-and-hold portfolio. Thus, we also test whether the returns may potentially be explained by the well-known risk factors as proposed by Fama and French (2015).

	Alpha	Benchmark-RF	SMB	HML	RMW	CMA	R^2
Best	0.0014***	1.43***	-0.07	-0.01	0.11	0.10	0.79
	(4.41)	(68.60)	(-1.29)	(-0.22)	(1.43)	(1.11)	
Median	0.0011***	1.40***	-0.08	-0.00	0.15*	-0.10	0.76
	(3.09)	(61.84)	(-1.54)	(-0.09)	(1.84)	(-1.02)	
Worst	0.0003	1.48***	-0.03	0.06	-0.01	-0.21**	0.81
	(0.95)	(72.05)	(-0.59)	(1.48)	(-0.11)	(-2.32)	

Table 5: Fama-French five factor regression results for long-only portfolios.
This table presents results from Fama-French five-factor regressions (Fama and French, 1993). We regress the neural network stack (NNs) returns on equally weighted buy-and-hold portfolio returns (Benchmark) and four common risk factors explained in Fama and French (1993). Factors are constructed using six value-weighted portfolios: Size (Small Minus Big stock portfolios, SMB), Value (High Minus Low Value, HML), Profitability (Robust Minus Weak operating profitability, RML) and Investments (Conservative Minus Aggressive, CMA) investment portfolios. Each randomly selected portfolio belongs to one regression. Displayed are the best, median, and worst results based on alpha values. Asterisks indicate the level of significance, i.e., p-values smaller than 10%, 5% or 1%, respectively. t-statistics are reported in parentheses. We assume transaction costs of 5 bps per transaction. The average alpha is 0.1%, whereas 19 out of 20 portfolios come with a statistically significant value of alpha ($p < 5\%$).

Besides market risk, the authors find four risk factors to explain significant portions of the stock return cross-sections: size, value, profitability, and the firm's investment policy.[6] Table 5 presents the regression results. The regression is estimated for each portfolio individually and we report the highest, median, and lowest results in terms of alpha values brevity.

The stacked neural network alpha is not explained by risk factors of Fama and French (2015). Table 5 shows regression results including the equally weighted buy-and-hold portfolio and the risk factors Size, Value, Profitability, and Investments. Including the risk factors does not decrease alpha values or the significant positive relationship to the benchmark strategy, nor does it make the unexplained performance less often significant across the portfolios. The effects are also captured by the mainly unaffected values of R^2 compared to the previous regression including solely benchmark returns. Figure 8 shows the distribution of the estimated regression coefficients of each factor based on the 20 regressions. While median values of the size, value and investments factors being close to zero, there remains a positive relationship of neural network stacks' returns and the operating profitability for some of the portfolios, indicating that the neural network stack prefers high operating

[6] Daily data is publicly available at the authors' website: `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

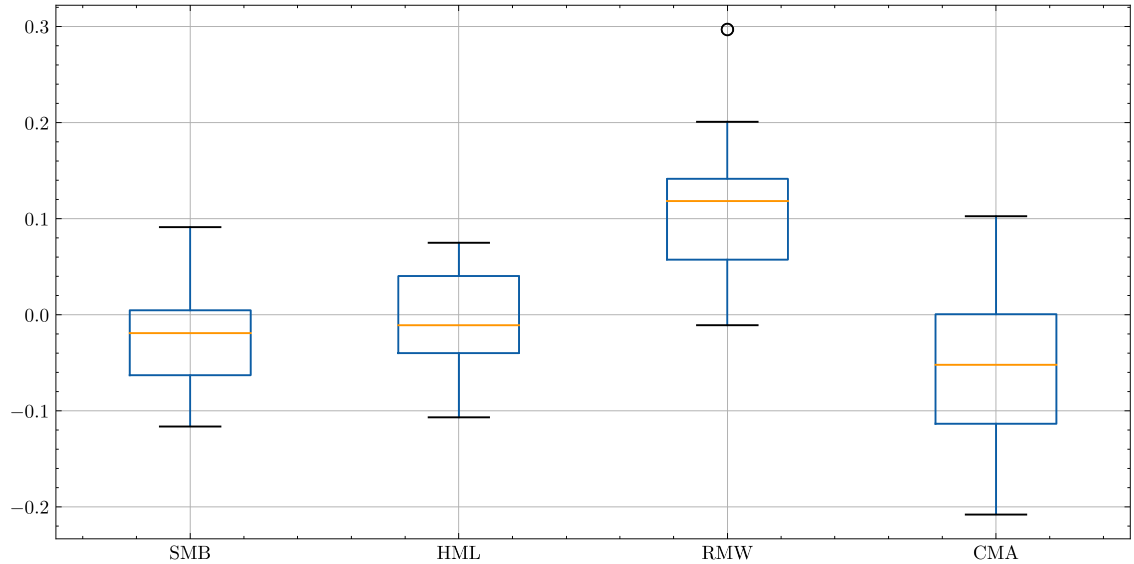

Figure 8: Fama-French risk factor exposures for long-only portfolios.
This figure depicts Fama-French risk factor exposures for out-of-sample long-only portfolios (Fama and French, 1993). For conciseness, we exclude exposure to benchmark returns where all parameter estimates are significantly larger than zero. Boxes show the interquartile range and the median. The farthest points are marked by lines, with outliers (that exceed the box edges by more than 1.5 times the interquartile range) as points.

profitability. Yet, the values of alpha are all significantly different from zero ($p < 5\%$) except for one portfolio. Apart from this one portfolio, the results suggest that the neural network stack predictions are sufficiently accurate to generate out-of-sample abnormal returns that are not explained by common risk factors.

3.4.2 Long-short portfolios

A common practice for algorithmic portfolio trading systems is to build market-neutral long-short portfolios. These portfolios target return spreads between the highest and lowest predicted stock returns. Accurate predictions yield large long-short portfolio returns, i.e., large long-short spreads. While a large spread is desirable, portfolios with a low correlation to market returns may be valuable to as a market-independent portfolio is well suited as, e.g., a diversification device or as an alternative to long-only portfolios in uncertain market phases.

As the neural network stack is designed to explicitly predict both negative and positive returns, it is well suited for market-neutral, i.e., a portfolio that has a low correlation with the benchmark, zero investment portfolios. For that, we combine the long-only portfolio neural network stack with the short counterpart. The short stack combines directional predictions from the same neural network as the long stack and afterwards ranks stocks based on the short model. Both neural network stacks form an isolated portfolio, which we combine as zero investment portfolios: $r_{p,t} = +1 * r_{l,t} - 1 * r_{s,t}$ in every period t. Here, $r_{p,t}$ is the return of the portfolio in period t, $r_{l,t}$ the return of the long stack in period t and $r_{s,t}$ the return of the short stack. The long and short stacks can also be combined using alternative weighting schemes based on the investor's preferences.

Table 6 shows average daily returns, annualized Sharpe ratios, and end-of-period cumulative returns. Returns are based on the combined neural network stacks (NNs) and on equally weighted buy-and-hold portfolios (Benchmark) using the 20 randomly sampled portfolios. On average, the neural network stack still outperforms the benchmark strategy in terms of all performance measures. It generates more than two-fold daily returns (0.17% vs. 0.08%), yielding a 354% increased mean cumulative return at the end of the 5-year period (671.62% vs. 147.83%). The large performance difference is represented in Sharpe ratios, too, given by a difference in Sharpe ratios of 47.23%-points. The neural network stack generates larger cumulative returns, as illustrated in Figure 9, than the benchmark in 18 out of the 20 portfolios. However, in only 15 out of the 20 portfolios, it can beat the annualized benchmark Sharpe ratio. On average, the long-short neural network stacks' performances are larger than a buy-and-hold portfolio even if risk is considered.

Table 7 shows the baseline regression results. It describes the relationship between the long-short neural network stack's returns and the benchmark. We show the best, median, and worst results measured by the value of alpha out of the 20 regressions at the portfolio level. Alpha values are persistently high, with an average value of 0.18%. The highest value of alpha is 0.26% while the lowest is 0.08% and is not significantly different from zero. 18 out of 20 alpha values are significantly different from zero at the 5%-level. Beta values are on average close to zero, with a value of -0.075. These values signal that the generated long-short portfolios are market-neutral. Moreover, the median regression yields an R^2 of only 1%, compared to 79% for the median long-only portfolio. That is, the neural network stacks' returns cannot be well explained by the benchmark returns, giving further support for market neutrality.

The results of the Fama-French five-factor regressions are similar to the long-only results. Table 8 presents the regression results again for the highest, median, and lowest portfolio setups in terms of alpha values. Alpha values remain high, with an average value of 0.17%. 18 out of 20 alpha values are significantly different from zero ($p < 5\%$). Beta values average at -0.073 while only 12 of these are significantly different from zero ($p < 5\%$). These results do not change on the basis of the included risk factors.

The risk factor exposure is larger compared to long-only portfolios. Figure 10

	Mean return		Sharpe ratio		Cumulative return	
	NNs	Benchmark	NNs	Benchmark	NNs	Benchmark
1	0.19%	0.09%	152.55%	91.51%	792.07%	159.53%
2	0.16%	0.09%	119.86%	90.43%	469.23%	160.99%
3	0.17%	0.08%	132.96%	89.21%	556.53%	149.99%
4	0.23%	0.08%	174.57%	84.99%	1228.66%	138.89%
5	0.15%	0.09%	111.56%	94.16%	379.26%	167.96%
6	0.14%	0.08%	107.06%	82.29%	332.7%	134.37%
7	0.17%	0.09%	138.88%	94.9%	536.99%	165.87%
8	0.22%	0.08%	181.38%	81.8%	1191.55%	131.8%
9	0.19%	0.08%	134.7%	83.23%	715.86%	136.73%
10	0.2%	0.09%	158.54%	92.48%	870.46%	165.08%
11	0.16%	0.08%	118.83%	85.09%	460.54%	140.64%
12	0.17%	0.09%	124.76%	87.93%	509.91%	153.75%
13	0.11%	0.08%	92.4%	84.62%	219.28%	138.24%
14	0.19%	0.08%	168.78%	84.45%	811.33%	140.06%
15	0.09%	0.08%	71.4%	80.9%	132.68%	126.23%
16	0.16%	0.09%	114.35%	94.3%	422.31%	167.47%
17	0.1%	0.08%	76.79%	87.33%	177.25%	148.98%
18	0.25%	0.08%	202.48%	83.78%	1801.38%	136.0%
19	0.22%	0.08%	168.84%	87.69%	1113.69%	145.38%
20	0.19%	0.08%	140.57%	85.51%	710.78%	148.63%
Mean	0.17%	0.08%	134.56%	87.33%	671.62%	147.83%

Table 6: Performance measures for market-neutral long-short portfolios.
This table shows performance measures for long-short portfolios. Displayed are mean daily returns, annualized Sharpe-Ratios and end-of-period cumulative returns using the stacked neural network structure (NNs) and an equally weighted buy-and-hold strategy (Benchmark). Long-short portfolios are market-neutral and require zero net investments. Accordingly, the weighting scheme is $w_{short} = -1, w_{long} = +1$. The statistics are based on the out-of-sample period from 2017-01-01 to 2021-12-31. The portfolios are evaluated and cumulated daily. We assume transaction costs of 5 bps per transaction. The third column represents cumulative end-of-period returns, accumulating profits over the period of five years.

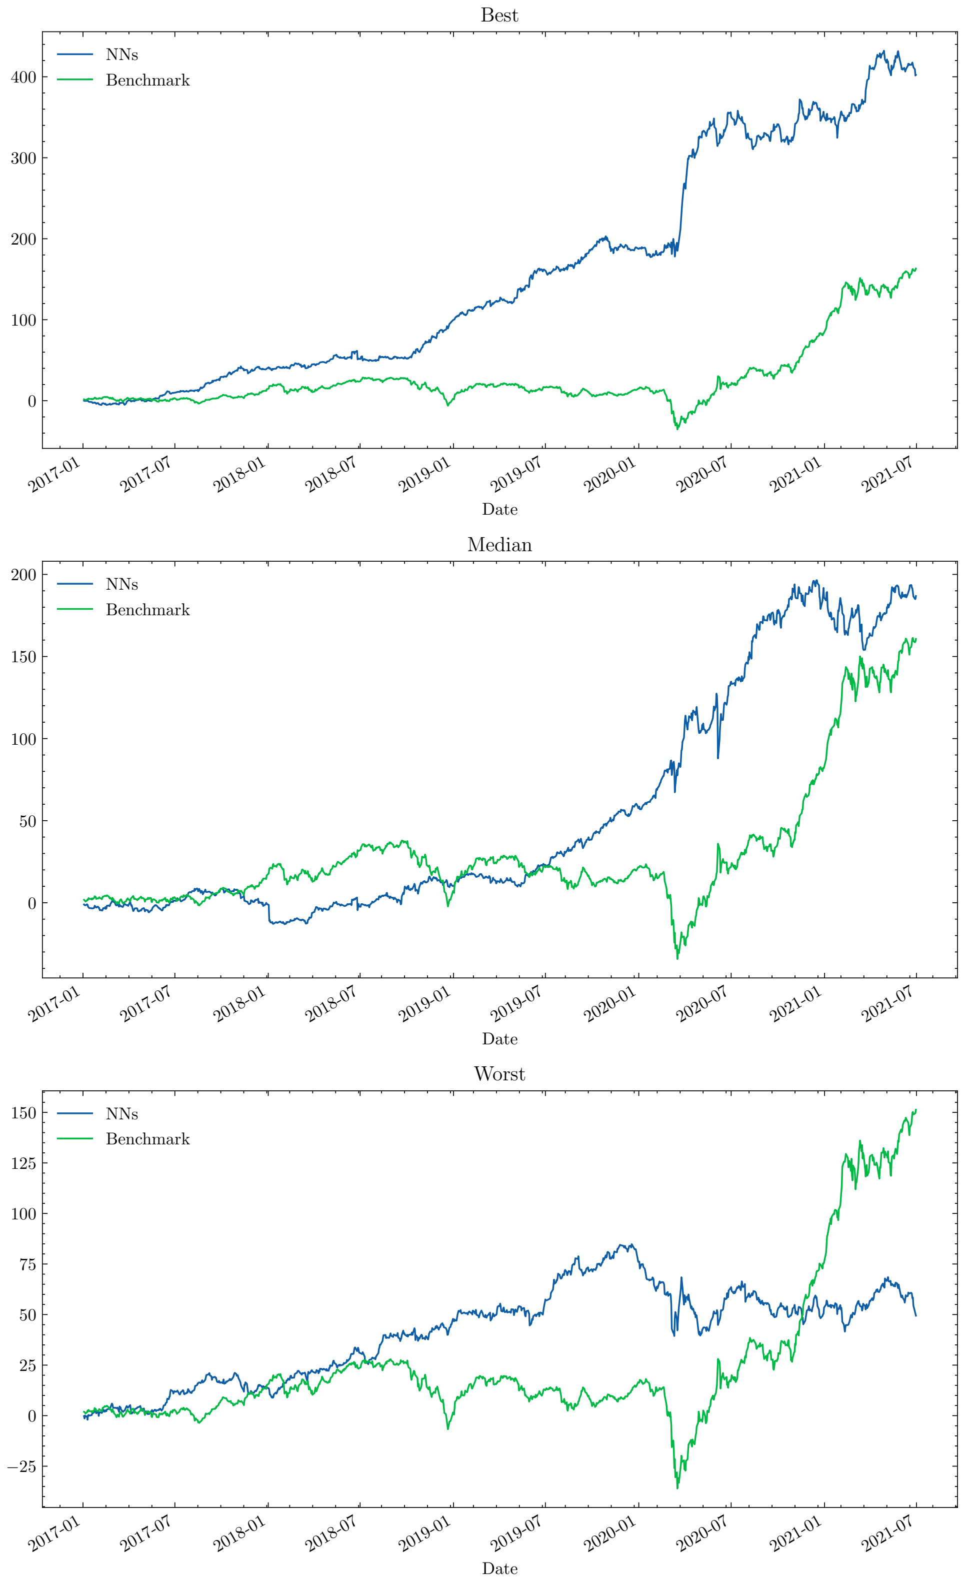

Figure 9: Performance plots for long-short portfolios.
This figure shows cumulative returns for the long-short neural network stack (NNs) and an equally weighted buy-and-hold portfolio (Benchmark). Displayed are the best, median, and worst performances in terms of cumulative returns. All values are in percentages and cumulated using the full test set of five years. We assume costs of 5 bps per transaction.

	Alpha	Beta	R^2	N Observations
Best	0.00262***	-0.09950***	0.01	1259
	(4.672)	(-2.684)		
Median	0.00173***	-0.14765***	0.01	1259
	(2.865)	(-3.729)		
Worst	0.00080	0.06328*	0.0	1259
	(1.505)	(1.771)		

Table 7: Baseline OLS regression results for long-short portfolios.
This table presents results from baseline OLS regressions. We regress the returns of the neural network long-short stack (NNs) on equally weighted buy-and-hold portfolio returns (Benchmark). Each randomly selected portfolio belongs to an OLS regression. Asterisks indicate the level of significance, i.e., p-values smaller than 10%, 5% or 1%, respectively. T-statistics are reported in parentheses. We assume transaction costs of 5 bps per transaction. Out of the 20 portfolios, we report the best, median, and worst results in terms of alpha. The average alpha is 0.18%, whereas 18 out of 20 portfolios have a statistically significant value of alpha ($p < 5\%$). The average beta is -0.075 with 12 out of 20 beta values significantly different from zero.

	Alpha	Benchmark-RF	SMB	HML	RMW	CMA	R^2
Best	0.0025***	-0.10***	-0.03	-0.11	0.27**	0.17	0.01
	(4.45)	(-2.64)	(-0.28)	(-1.41)	(2.09)	(1.03)	
Median	0.0016***	-0.11***	0.02	-0.23***	0.06	0.34*	0.01
	(2.73)	(-2.78)	(0.19)	(-2.72)	(0.40)	(1.96)	
Worst	0.0007	0.07*	-0.09	0.05	0.10	-0.19	0.01
	(1.35)	(1.83)	(-1.04)	(0.63)	(0.79)	(-1.24)	

Table 8: Fama-French 5-factor regression results for long-short portfolios.
This table presents results from Fama-French 5-factor regressions (Fama and French, 1993). We regress the returns of the neural network long-short stack (NNs) on equally weighted buy-and-hold portfolio returns (Benchmark) and returns of four common risk factors explained in Fama and French (1993). Factors are constructed using six value-weighted portfolios: Size (Small Minus Big stock portfolios, SMB), Value (High Minus Low Value, HML), Profitability (Robust Minus Weak operating profitability, RMW), and Investments (Conservative Minus Aggressive, CMA) investment portfolios. Each randomly selected portfolio belongs to one regression. Displayed are the highest, median, and lowest results in terms of alpha values. Asterisks indicate the level of significance, i.e., p-values smaller than 10%, 5% or 1%, respectively. t-statistics are reported in parentheses. We assume transaction costs of 5 bps per transaction. The average alpha is 0.1%, whereas 19 out of 20 portfolios come with a statistically significant value of alpha ($p < 5\%$).

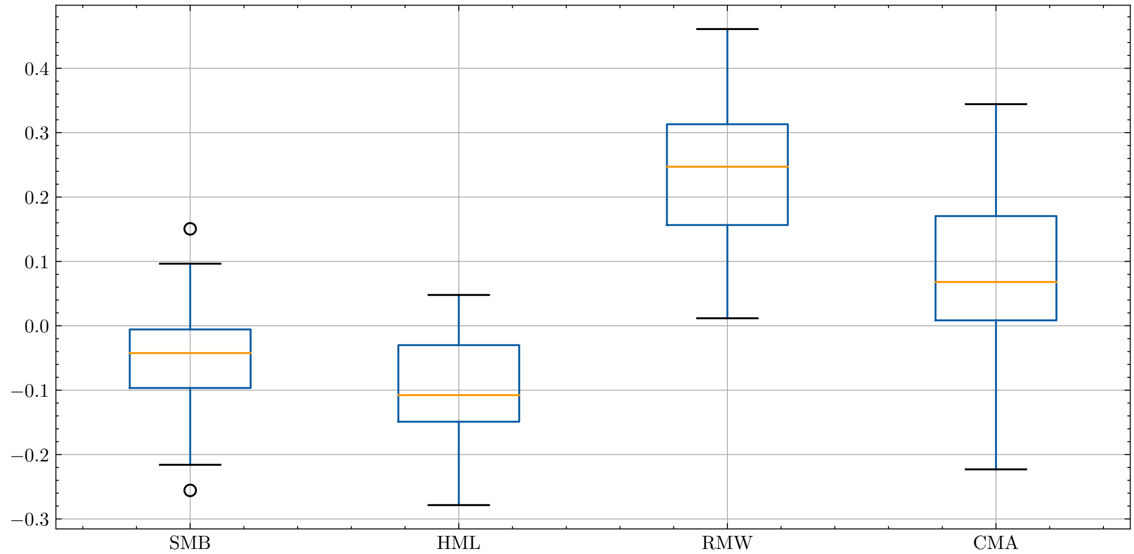

Figure 10: Fama-French risk factor exposures for long-short portfolios.
This figure shows Fama-French risk factor exposures for out-of-sample long-short portfolios (Fama and French, 1993). For conciseness, we exclude exposure to market returns. Boxes show the interquartile range and the median. The farthest points are marked by lines, with outliers (that exceed the box edges by more than 1.5 times the interquartile range) as points.

summarizes regression results for all portfolios using boxplots. Size and value factor relationships are slightly negative, i.e., the neural network stack prefers to buy larger and sell smaller companies, while it prefers value over growth. Analogously to long-only portfolios, operating profitability is the most influential risk factor. The long-short neural network stack buys strong operating profitability and sells weak operating profitability. This is the only risk factor where the neural network stack agrees on the same sign for every portfolio. A conservative investment policy seems to be preferred across the neural network stacks, but there is also a larger interquartile range. Still, these values need to be interpreted cautiously, as the neural network stacks' returns can hardly be explained by the given regression due to a median R^2 of only 1%. Overall, the zero-investment long-short portfolios of the neural network stacks are market-independent and generate statistically significant alpha values even after controlling for empirical risk factors.

3.4.3 Performance sources

Neural networks are considered black-box models. However, as interpreting features may be valuable, the literature provides approximation techniques for feature importance. We follow the approach of Friedman (2001) and derive feature importance based on partial dependence. For each feature, we assign a constant value to each data point for that feature, while leaving the other features unchanged. We select

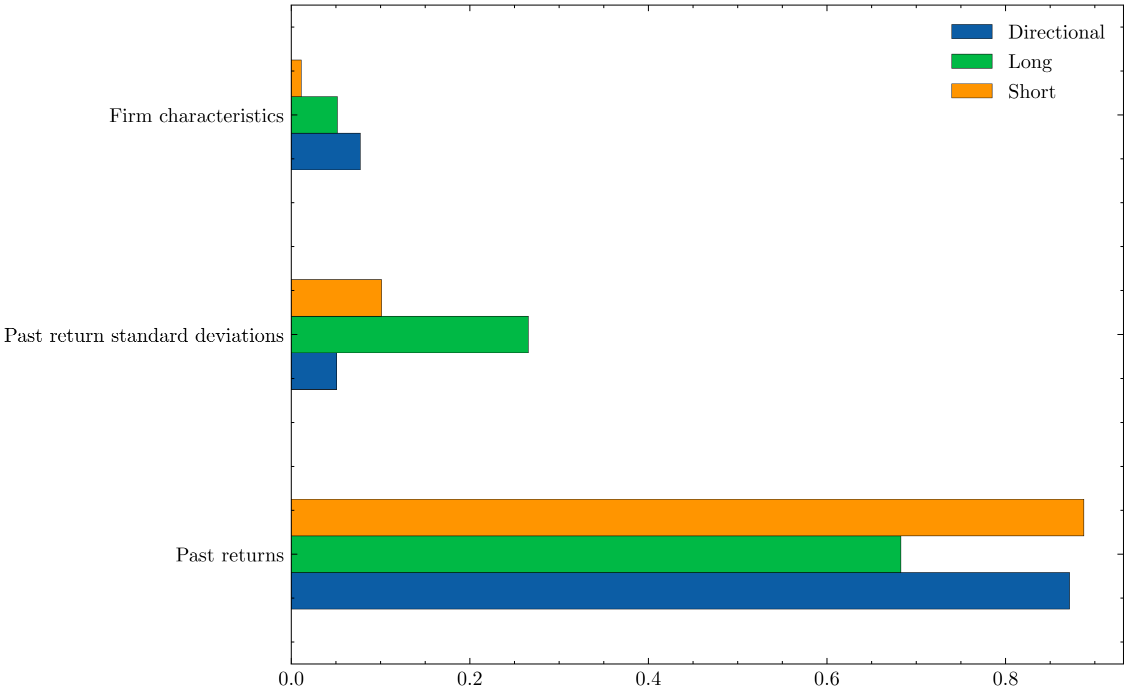

Figure 11: Aggregated feature importances for each model.
This figure shows the importance of features for each of the three models. We aggregate all firm characteristics, return moving averages (Past returns), and return standard deviations. To calculate feature importance, we follow the partial dependence approach of (Friedman, 2001). Feature importances are normalized to sum up to one at the model-level.

ten values between the true 5%-quantile and the 95%-quantile of the feature values. In total, this gives ten artificial data sets. Then, we compute the average prediction for each of the ten artificial data sets. The variance of the ten mean predictions gives the feature importance. The intuition is simple: If changes in the input value lead to relatively large changes in the predictions, then the feature has high importance.

Figure 11 shows feature importance for each of the three models. Past returns are the largest driver for predictions across all models. That is to be expected as the models predict one-day returns, whereas firm-related data is updated at lower frequencies. Accordingly, the short-term task makes firm characteristics less valuable, resulting in a small feature's importance. Firm characteristics have a relatively large impact on the directional models' predictions. The most interesting result is that return standard deviations are more than twice as important for the long model relative to the short model. We conjecture that this is due to the stock market losses starting in March 2020. Here, high volatility is followed by larger stock market gains in the rest of 2020. As a result, the long models pick up this pattern and increase the usage of the return standard deviation.

To better understand the effectiveness of our two-stage neural network structure, we report performance measures for long-only portfolios using solely the directional model, i.e., the first stage. That is, we construct equally weighted portfolios based on the top decile of predictions from the direction model in each period. Table 9 shows average daily returns, Sharpe ratios, and cumulative returns for all sampled portfolios. Although the one-model approach does lead to a slightly higher annualized average Sharpe ratio, the stack yields notably larger average daily returns as well as end-of-period cumulative returns. In fact, average daily returns are 28% larger (a difference of 0.048%-points) and end-of-period cumulative returns are on average 69% larger (a difference of 421%-points). The results are in line with our theoretical expectations. The directional model needs to classify the sign of returns and, hence, does not rely on the concrete return values. The results suggest that while the directional model appears to consistently increase realized Sharpe ratios, the second stage, i.e., the long model, rather exploits risk in order to increase realized returns, keeping the Sharpe ratio approximately constant.

3.4.4 Cost Sensitivity

A central limiting factor of active trading strategies are transaction costs. An active portfolio management strategy must overcome transaction costs while simultaneously beating its benchmark. A common choice for the benchmark is a passive, equally weighted buy-and-hold strategy, as it does not face any transaction costs. To study the effects of transaction costs on our stacked neural network out-of-sample returns, we compute average daily returns, Sharpe ratios, and cumulative returns with doubled transaction costs, i.e., 10 bps. The results give an approximation of the infrastructural requirements to execute the strategy in a profitable way.

Table 10 shows results for long-only portfolios under doubled transaction costs (10 bps), including the passive benchmark that does not face any transaction costs. The results suggest that the strategy is indeed sensitive to increases in transaction costs. That is, the average annualized Sharpe ratio decreases by almost 21%, i.e., a difference of 15.45 percentage points. The average end-of-period cumulative return decreases by roughly 28%. Thus, the average daily return suffers by 0.00472 percentage points per 1 bps-increase in costs. In fact, having transaction costs of 51 bps will result in zero average daily returns. The strategy reaches the benchmark's level at 33 bps. Hence, the neural network-based system for long portfolios is indeed sensitive to transaction costs, yet it is still effective even at higher levels of transaction costs.

Unsurprisingly, the long-short strategy suffers more from increases in costs than the long-only strategy, as shown in Table 11. That is, per basis point increase in transaction costs, the long-short strategy loses 0.0079 percentage points in average daily returns. On average, the long-short neural network yields losses as transaction

	Mean return		Sharpe ratio		Cumulative return	
	Directional NN	NN stack	Directional NN	NN stack	Directional NN	NN stack
1	0.19%	0.23%	169.93%	149.59%	761.0%	1138.76%
2	0.16%	0.21%	132.99%	135.52%	470.31%	886.52%
3	0.17%	0.23%	139.01%	144.34%	543.75%	1052.19%
4	0.18%	0.23%	155.36%	146.94%	653.36%	1084.38%
5	0.17%	0.23%	144.47%	147.72%	550.81%	1100.62%
6	0.15%	0.17%	132.36%	117.2%	441.78%	519.47%
7	0.17%	0.24%	143.44%	162.93%	548.43%	1392.67%
8	0.18%	0.25%	159.04%	164.35%	715.37%	1566.37%
9	0.16%	0.21%	140.51%	138.89%	518.42%	853.15%
10	0.22%	0.24%	178.67%	149.62%	1167.72%	1231.09%
11	0.17%	0.23%	143.0%	148.7%	575.16%	1080.8%
12	0.18%	0.25%	153.97%	161.9%	647.2%	1484.98%
13	0.15%	0.18%	136.24%	115.37%	452.73%	529.78%
14	0.19%	0.21%	161.32%	135.53%	760.65%	819.45%
15	0.12%	0.14%	103.63%	90.58%	275.8%	298.76%
16	0.15%	0.22%	135.15%	150.42%	448.3%	1050.67%
17	0.16%	0.19%	138.64%	127.15%	510.05%	695.17%
18	0.2%	0.24%	170.38%	154.99%	871.23%	1280.01%
19	0.18%	0.24%	157.72%	162.68%	662.68%	1393.79%
20	0.17%	0.23%	142.79%	151.86%	609.97%	1152.93%
Mean	0.17%	0.22%	146.93%	142.81%	609.24%	1030.58%

Table 9: Performance measures for portfolios based solely on the directional model.
This table shows performance measures for the directional model and for long-only portfolios using the stack. Displayed are mean daily returns, annualized Sharpe Ratios, and end-of-period cumulative returns using the directional neural network only (Directional NN) and the long-only neural network stack (NN stack). The statistics are based on the out-of-sample period from January 2017 to December 2021. The portfolios are evaluated and cumulated daily. We assume transaction costs of 5 bps per transaction. The third column represents cumulative end-of-period returns, accumulating profits over the period of five years.

	Mean return		Sharpe ratio		Cumulative return	
	NNs	Benchmark	NNs	Benchmark	NNs	Benchmark
1	0.21%	0.09%	134.04%	91.51%	817.15%	159.53%
2	0.19%	0.09%	120.65%	90.43%	635.16%	160.99%
3	0.2%	0.08%	129.1%	89.21%	754.2%	149.99%
4	0.2%	0.08%	131.41%	84.99%	775.97%	138.89%
5	0.2%	0.09%	132.42%	94.16%	792.89%	167.96%
6	0.15%	0.08%	101.31%	82.29%	361.78%	134.37%
7	0.22%	0.09%	147.26%	94.9%	1012.86%	165.87%
8	0.23%	0.08%	148.78%	81.8%	1131.49%	131.8%
9	0.18%	0.08%	123.24%	83.23%	610.11%	136.73%
10	0.21%	0.09%	134.9%	92.48%	891.7%	165.08%
11	0.2%	0.08%	132.97%	85.09%	775.27%	140.64%
12	0.23%	0.09%	146.6%	87.93%	1078.58%	153.75%
13	0.15%	0.08%	99.92%	84.62%	368.7%	138.24%
14	0.18%	0.08%	120.18%	84.45%	586.42%	140.06%
15	0.12%	0.08%	75.54%	80.9%	197.1%	126.23%
16	0.2%	0.09%	134.33%	94.3%	754.2%	167.47%
17	0.17%	0.08%	111.78%	87.33%	491.92%	148.98%
18	0.21%	0.08%	139.62%	83.78%	925.1%	136.0%
19	0.22%	0.08%	146.76%	87.69%	1007.26%	145.38%
20	0.21%	0.08%	136.48%	85.51%	835.18%	148.63%
Mean	0.19%	0.08%	127.36%	87.33%	740.15%	147.83%

Table 10: Performance measures for long-only portfolios under doubled transaction costs.
This table shows performance measures for long-only portfolios under doubled costs (10 bps) per transaction. Displayed are mean daily returns, annualized Sharpe-Ratios, and end-of-period cumulative returns using the stacked neural network structure (NNs) and an equally weighted buy-and-hold strategy (Benchmark). The statistics are based on the out-of-sample period from January 2017 to December 2021. The portfolios are evaluated and cumulated daily. The third column represents cumulative end-of-period returns, accumulating profits over the period of five years.

costs exceed 27 bps. The strategy loses in Sharpe ratio, respectively. Still, even after doubling transaction costs, zero-net investment portfolios outperform a passive buy-and-hold strategy.

	Mean return		Sharpe ratio		Cumulative return	
	NNs	Benchmark	NNs	Benchmark	NNs	Benchmark
1	0.15%	0.09%	121.2%	91.51%	439.47%	159.53%
2	0.12%	0.09%	90.47%	90.43%	245.44%	160.99%
3	0.13%	0.08%	101.81%	89.21%	297.33%	149.99%
4	0.19%	0.08%	143.49%	84.99%	698.63%	138.89%
5	0.11%	0.09%	81.76%	94.16%	192.71%	167.96%
6	0.1%	0.08%	76.35%	82.29%	162.72%	134.37%
7	0.13%	0.09%	105.16%	94.9%	284.8%	165.87%
8	0.18%	0.08%	148.61%	81.8%	678.95%	131.8%
9	0.15%	0.08%	107.39%	83.23%	396.95%	136.73%
10	0.16%	0.09%	127.48%	92.48%	491.2%	165.08%
11	0.12%	0.08%	89.12%	85.09%	238.41%	140.64%
12	0.13%	0.09%	95.1%	87.93%	270.33%	153.75%
13	0.07%	0.08%	59.12%	84.62%	93.5%	138.24%
14	0.15%	0.08%	134.2%	84.45%	455.38%	140.06%
15	0.05%	0.08%	37.93%	80.9%	40.7%	126.23%
16	0.12%	0.09%	85.15%	94.3%	217.1%	167.47%
17	0.07%	0.08%	47.96%	87.33%	69.14%	148.98%
18	0.21%	0.08%	170.82%	83.78%	1053.59%	136.0%
19	0.18%	0.08%	138.46%	87.69%	637.25%	145.38%
20	0.15%	0.08%	111.69%	85.51%	396.04%	148.63%
Mean	0.13%	0.08%	103.66%	87.33%	367.98%	147.83%

Table 11: Performance measures for market-neutral long-short portfolios under doubled transaction costs.
This table shows performance measures for long-short portfolios under doubled costs (10 bps) per transaction. Displayed are mean daily returns, annualized Sharpe ratios, and end-of-period cumulative returns using the stacked neural network structure (NNs) and an equally weighted buy-and-hold strategy (Benchmark). Long-short portfolios are market-neutral and require zero investments. Accordingly, the weighting scheme is $w_{short} = -1, w_{long} = +1$. The statistics are based on the out-of-sample period from January 2017 to December 2021. The portfolios are evaluated and cumulated daily. The third column represents cumulative end-of-period returns, accumulating profits over the period of five years.

3.5 Conclusion

This Chapter shows how neural network-based portfolio management systems can be improved by using a two-stage prediction approach. By separating sign and magnitude predictions, we show that the former largely contributes to increases in Sharpe ratios and produces statistically significant risk-adjusted alphas compared to an equally weighted buy-and-hold portfolio while accounting for common risk factors. Sign predictions systematically leverage returns and risk simultaneously, keeping the Sharpe ratio constant. Long-short portfolios produced by the neural network stack have low correlations with the equally-weighted buy-and-hold benchmark, yet they more than double average daily benchmark returns. As transaction costs are a central limiting factor of active investment strategies, we outline the maximum theoretical level of transaction costs for the strategy to remain beneficial to the investor. Additionally, we show how financial neural networks can be trained in several seconds by utilizing parallel computing of modern GPUs without the necessity of splitting the training data into multiple batches while training, while our models generalize well.

We first classify returns by their sign and estimate the magnitude separately. There are, however, other potential structures that my be worth exploring. For instance, extreme returns, e.g., the top and bottom decile of returns, may depend on a completely different feature set than the remaining, less extreme returns. Hence, researchers may first estimate the probability of an extreme return and train multiple models to predict the resulting return magnitude depending on the first model's estimate. Although a single neural network may be able to capture such structures given sufficient capacities, the high noise-to-signal ratio in financial data suggests to keep the models' complexity low to avoid potential overfitting. Furthermore, we show that splitting the prediction task in several subtask increases the interpretability of the resulting predictions.

4 Anxiety in Returns

The analysis of the neural network stack's feature importance presented in the previous Chapter reveals that the most important predictor is the history of returns and, also, that features are treated differently by the models that handle predicted positive and negative returns, respectively. Based on these two findings, this Chapter analyzes the importance of the preceding return sign to investors. We show that investors treat positive returns entirely differently than negative returns, a behavior that we theoretically predict from the findings of behavioral finance.

In particular, the prospect theory documents that investors are risk averse and, thus, demand larger risk premiums for stocks that appear to be riskier. As we propose in our analysis, stocks with preceding losses appear riskier to investors than stocks with preceding gains, making investors anxious to buy the former stocks in the following period, which we call the anxiety anomaly. In accordance with our theoretical predictions, we observe significantly larger following returns for stocks with a preceding positive return than stocks with a preceding negative return. For that, we also propose a novel measure of risk aversion that is based on common risk factors' premiums.

Based on our analysis, we further conjecture that most of machine learning's predictive ability in asset pricing may stem from the predictability of human preferences as we find that a simple linear model based on the anxiety anomaly produces an out-of-sample predictive power as large as the predictive power of modern machine learning models with hundreds of predictors (Gu et al., 2020). In these models, the most important predictors are recent returns and a factor approximating risk (Gu et al., 2020), which is in line with our findings in the previous Chapter, i.e., using our neural network stack. These results strongly suggest that the anxiety anomaly may be a potential theoretical explanation for the superior predictive power of machine learning in asset pricing.

The paper of this Chapter is published in the Journal of Behavioral Finance (Pigorsch and Schäfer, 2023).

Abstract

We provide empirical evidence that risk-averse investors become anxious about investments in stocks whose realized losses reveal the downside of risk. Contrary to short-term reversal and in support of convex risk aversion, the latter stocks yield significantly lower returns in the subsequent period. Our findings are based on a novel measure of time-varying risk aversion but can also be observed when using a well-established measure of risk aversion. Moreover, anxiety predicts cross-sectional returns in out-of-sample tests, suggesting that risk-averse investors' preferences drive empirical risk premia.

4.1 Introduction

Does the revelation of downside risk affect the investment behavior of risk-averse investors? The pioneering findings of the prospect theory (Kahneman and Tversky, 1979) state that investors are risk averse, prefer safe over risky prospects, and, thus, demand a premium for risk-taking. In absolute terms, the value of a risky prospect is steeper for losses than gains, whereas the value function is convex in losses and concave in gains. Based on these findings, we propose that a risk-averse investor is anxious to invest in a stock, i.e., a risky investment, if the stock has previously experienced a negative return. The negative return reveals the investment risk to the investor, who responds by requiring a higher rate of return as compensation for taking on risk. This irrational assessment of risk makes stocks with a preceding positive return significantly more appealing to investors, as they appear to be less risky than stocks with a preceding negative return. We provide empirical evidence that the resulting return spread between stocks with a preceding positive and negative return is significantly positive.

Several studies show that risk aversion leads to predictable anomalous return patterns in the cross-section of stock returns. The patterns include the increase of the long-term premium on stock returns by risk aversion (Benartzi and Thaler, 1995b), which may even explain the equity premium puzzle (Mehra and Prescott, 1985). Furthermore, expected returns are determined by investor sentiment (Shen et al., 2017), the sensitivity to innovations in aggregate market risk (Ang et al., 2006), historical peak returns (Bali et al., 2011) and higher moments of the return history (Barberis et al., 2016; Bordalo et al., 2013). However, these studies either concentrate on the medium- to long-term effects of risk aversion on returns or incorporate a large set of economic variables apart from stock price information. In contrast, we find that risk aversion is measured by and simultaneously predicts short-term cross-sectional returns. Specifically, we first motivate our approach to construct a potential measure of risk aversion based on the cross-section of stock returns. Using this measure, we show that the anxiety of risk-averse investors predicts the monthly returns of those stocks that most recently revealed the downside of risk, signaled by a negative return in the preceding month. We find similar results when using a well-established alternative measure of risk aversion (Bollerslev et al., 2009). Additionally, we propose a preference-based yet intuitive explanation for this empirical observation.

To this end, consider an investor whose perception of the firm's fundamental value at period t is:

$$\mathbb{E}[V_t] = \sum_{i=1}^{\infty} \frac{\mathbb{E}[CF_{t+i}]}{(1+\mu)^i}\,, \tag{5}$$

where $\mathbb{E}[CF_{t+i}]$ is the expected cash flow in period $t+i$, and μ is the investor's required rate of return per period.[7] The valuation of the company reflects the investor's expectations regarding the company's ability to generate free cash flows and the investor's perception of the associated uncertainty. This study focuses on the predictive measurement of the latter.

Banz (1981) documents that the returns of small stocks outperform those of large stocks. Small stocks tend to be riskier (Daniel and Titman, 1997; Fama and French, 1993), i.e., have greater uncertainty in expected future cash flows than large stocks. Additionally, the likelihood of default may be more prevailing to investors than in the case of large stocks. Thus, the size premium is, at least partly, a proxy for the investors' required compensation for risk-taking. Therefore, we expect high risk aversion when there is a large size premium, i.e., a large difference in returns between small and large firms. Besides, we demonstrate in Section 4.2 that the approximation of the investor's level of risk aversion by the return spread of a risk factor is in line with standard mean-variance utility. The measurement of risk aversion is also time-varying, a well-known attribute of risk aversion, as we discuss in more detail in the next section. The following provides an intuitive explanation of the risk aversion measure.

Consider a large and positive spread in period t in the size risk factor, i.e., small companies have a higher average realized return than large companies. Closely prior to period t, i.e., at the end of period $t-1$, investors must have avoided small

[7]We do not consider the accounting-implied value of the well-known free cash flow model, as we neither measure expected firm values nor expected cash flows empirically, as, e.g., in Francis et al. (2000). The exact definition of a payback does not change the following considerations and is interchangeable with, for instance, dividends.

companies in favor of larger companies, thereby avoiding uncertainty and being predominantly risk-averse regarding investment decisions for the following period t. Put differently, investors required high compensation for risk during period t, demanding a large expected return for risky investments realized at the end of period t. Consequently, we can observe larger (smaller) returns for smaller (larger) stocks in period t. Hence, the most recent observable magnitude of risk aversion is, at least partly, reflected in the long-short size spread.

The above outlines how the long-short risk factor spread may approximate the degree of risk aversion. Whenever the aggregate risk aversion in the markets is relatively large, we expect a run out of stocks that reveal the downside of risk by preceding realized negative returns. Suppose an investor is already fearful of uncertainty and, as a result, demands low-volatility stocks. In this case, the investor may become anxious to invest in stocks whose most recent return is negative, causing the investor to avoid these stocks. This hypothesis is motivated by the well-documented risk-averse investor's reaction to fear and substantial losses (Cohn et al., 2015; Guiso et al., 2018; Kuhnen and Knutson, 2011; Weber et al., 2013). Furthermore, Walther and Münster (2021) empirically show that investors demand a positive (negative) risk premium for beta risk given a positive (negative) preceding return. On the other hand, when the investor is less risk-averse, we do not expect a significant reaction to the preceding return sign at the firm level. On the other hand, when the investor is less risk-averse, we do not expect a significant reaction to the preceding return sign at the firm level.

In line with our theoretical predictions, we empirically find that in periods with high aggregate risk aversion, stocks with a preceding negative return in period t are indeed excessively avoided in the following period $t + 1$. The effect is statistically and economically significant: in periods of high aggregate risk aversion, stocks with a preceding positive return in t earn in the subsequent period, $t + 1$, a return that is seven times greater than stocks with a preceding negative return. Hereafter, we refer to this spread in returns as the *anxiety anomaly*[8].

The anomaly is also observed for risk aversion derived from alternative risk factors, suggesting that the effect is not specific to the size factor. Importantly, note that we do not claim that the long-short risk factor spread is a perfect measure

[8]We refer to our finding as an anomaly, as the observed investor's behavior is likely to be irrational, i.e., there is no apparent reason not to invest in assets with a preceding negative return despite having a high level of risk aversion, and as it yields high out-of-sample return predictability that results in remarkably strong long-short returns over multiple decades, which is a commonly accepted characteristic of an anomaly; see, for instance Hou et al. (2015).

of risk aversion and therefore assess the robustness of our main results using the risk aversion measure of Bollerslev et al. (2009), which is defined as the difference between implied and realized variance. We find that when using this alternative measure of risk aversion, the large return spread is also present.

Moreover, we confirm convexity in the investors' utility in losses (Kahneman and Tversky, 1979) as we show that the risk-averse investors' avoidance of stocks with a preceding negative return is convex in magnitude. Furthermore, the anomaly counteracts short-term reversal and is not affected by the month of January or any decade we observe.

Further, we assess whether the effect persists in the subsequent period $t + 2$. We expect that the effect is short-term and, at most, reverses in the period $t + 2$. Suppose the anxiety of risk-averse investors triggers the sell-off of stocks with a negative return in period t. In that case, it is disconnected from the company's actual value and should not persist in an arbitrage-free market. Thus, the persistence of the anxious reaction to downside risk revelation stands in conflict with risk aversion being the underlying origin. In line with our expectations, the risk-averse investors' reaction to the revelation of downside risk is myopic. We observe that the anxiety anomaly is short-lived and only exists in the subsequent period $t + 1$. In fact, stocks with a preceding negative return yield above-average returns for the period $t + 2$ after they suffered from the anxiety anomaly.

Our results also contribute to the cross-sectional predictive asset pricing literature. In a multi-decade out-of-sample test with over 3.4 million monthly cross-sectional stock returns, we show that the anxiety anomaly achieves an R^2_{oos} of 0.41%. The performance is comparable to the machine learning-based out-of-sample predictability recorded in the literature. For example, Gu et al. (2020) demonstrate that a neural network predicts the cross-section of stock returns with a R^2_{oos} of 0.4%. We conjecture that most of that predictability may be due to investors' preferences towards risk, captured by the anxiety anomaly.

This Chapter is also related to the literature on the effects of sentiment on investors' risk-taking behavior and resulting differences in the return on high- and low-risk portfolios (Baker and Wurgler, 2006; Bekaert et al., 2010; Campbell and Cochrane, 1999; Kamstra et al., 2003; Shen et al., 2017). From this field, the paper of Shen et al. (2017) is the one that is most closely related to ours. They show empirically that following periods of low (high) levels of investor sentiment (measured by the index developed by Baker and Wurgler (2006)), risk-taking becomes more (less) attractive to investors. In particular, they find that high-risk stocks earn a lower average return than low-risk stocks in the following period. This is similar to our findings, although our approach differs in several ways as we are interested in the effect of downside risk revelation on investors' behavior under time-varying risk aversion. In comparison, Shen et al. (2017) measure the returns of high- and low-risk stocks during periods of high and low investor sentiment, where the stock's riskiness is defined by its correlation to macroeconomic variables. Hence, our approach provides valuable insights into the investor's perception of risk, independent of the investment's objectively measured risk. Moreover, we do not measure investor sentiment but rather develop a novel approach to directly approximate the investor's level of risk aversion.

The remaining sections of the Chapter are organized as follows. In Section 4.2 we provide a theoretical motivation for approximating the investors' degree of risk aversion via the long-short spread of a risk factor and state our predictions regarding the risk-averse investors' reaction to the revelation of downside risk provided by negative returns. Section 4.3 discusses our empirical approach and the corresponding results, including the magnitude of the anomaly, the functional form of investors' preferences, multiple robustness checks, out-of-sample prediction results and potential drivers of the anxiety anomaly. Section 4.4 shows the convergence of a neural network towards the anxiety given the respective features of the anxiety and no a priori knowledge about the functional form. Finally, Section 4.5 concludes.

4.2 Risk aversion, anxiety and investment behavior

In this Chapter, we propose to measure risk aversion using the long-short spread of a risk factor. In addition to our intuitive motivation for this measure provided in the introduction, we outline its plausibility based on mean-variance utility maximization in the following. To this end, consider a capital market consisting of risk-averse investors with homogeneous expectations about returns and variances in available stocks. All investors face identical market frictions, that is, we rule out any personal spreads or price impacts. In addition, each period, all agents review their portfolio and rebalance it according to their utility-maximizing portfolio weights. We further assume that only the first two moments, i.e., expected values and variances of stock returns, affect the investor's utility. At first glance, this may appear to be a restrictive assumption, as it is well known that higher-order moments, such as skewness, may also be relevant components of the utility function, as shown, for example, in Harvey and Siddique (2000); Kane (1982); Kraus and Litzenberger (1976); Mitton and Vorkink (2007). However, as we detail below, our conclusions are unaffected whether we start with a quadratic, mean-variance-based, utility function analogous to the utility used in Markowitz (1952) or initiate the derivation with a cubic form involving, e.g., skewness, as in Mitton and Vorkink (2007).

Standard quadratic utility states that investors generally appreciate returns. At the same time, they dislike volatility, that is, the expected utility from investing in a stock i is increasing for increasing expected returns, μ_i, and decreasing for increasing expected variance, σ_i^2, ceteris paribus, respectively. The most popular representation of this preference function is due to Markowitz (1952) and is given by:

$$\mathbb{E}[U_{t,i}(\theta_{t,i})] = \mu_{t,i} - \gamma_t \, \sigma_{t,i}^2, \tag{6}$$

with $\theta_{t,i} = (\mu_{t,i}, \sigma_{t,i}, \gamma_t)$, t denoting the period, and where we, however, relax the assumption of constant risk-aversion and allow γ_t to be time-varying.[9]

The assumption of a time-varying risk aversion parameter is not novel in the literature. Several studies find that investors change their risk aversion over time. Campbell and Cochrane (1999) and Bekaert et al. (2010) develop models of consumption-based asset pricing that include countercyclical risk aversion. Kamstra et al. (2014) find that risk aversion is multiple times higher during the winter than the summer months. According to Bollerslev et al. (2009), the difference between implied and realized variance significantly predicts the quarterly market return. The authors provide empirical evidence for the predictability not being induced by time-varying systematic risk but show that time-varying risk aversion is likely to explain the results. Closely related, Bekaert and Hoerova (2016) find that risk aversion is highly time-varying and peaks during periods of high financial turmoil in the markets with a time-series model that measures risk-aversion using the implied variance risk premium based on VIX data and the high-frequency realized variance. Therefore, we assume that the risk aversion γ_t is time-varying.

We argue in this Chapter that the spread from long-short risk factors may approximate the short-term aggregate investors' level of risk aversion. To show that this is in line with mean-variance utility maximization, we focus on the size factor, but any other risk factor or portfolio of risk factors can be used as well. Consider the investor's utility from a long-short portfolio that mimics the size factor (Banz, 1981). That is, the investor expects a return $\mu_{t,s}$ of a portfolio of small companies in period t and a return $\mu_{t,l}$ of a portfolio of large companies. As usual, the investor buys the portfolio of small companies and shorts the portfolio of large companies. The corresponding variances are denoted by $\sigma_{t,s}^2$ and $\sigma_{t,l}^2$. Accordingly, the expected

[9] Note that for brevity and based on our assumption of homogeneous investors, we omit a subscript for the investor.

utility U^{Size} defined in Equation (6) for this long-short size portfolio is given by:

$$\mathbb{E}[U_{t,i}^{Size}(\theta_{t,i})] = \mu_{t,s} - \mu_{t,l} - \gamma_t\,(\sigma_{t,s}^2 + \sigma_{t,l}^2 - 2\rho_{t,s,l}\sigma_{t,s}\sigma_{t,l}), \tag{7}$$

with $\mu_{t,s} - \mu_{t,l}$ being the expected return of the long-short size portfolio and $\rho_{t,s,l}$ denoting the expected correlation coefficient between the returns of the small and large portfolio in period t.

The agents only invest in the long-short size portfolio described above if they expect a positive utility from the investment. As the portrayed long-short size portfolio is also a zero-investment portfolio, with the weights in the small and large stock portfolios being $w_s = +1$ and $w_l = -1$, the investor is willing to take the investment at any positive expected utility, no matter the magnitude. However, efficient markets require that prices only allow for marginally positive aggregate investors' expected utility. That is, given a fixed, positive level of utility that is marginally larger than zero, denoted by $\bar{U}^+$, with the investor having some expectation about return (co-)variances, the required long-short size return is given by:

$$\tilde{\mu}_{t,s} - \tilde{\mu}_{t,l} = \bar{U}^+ + \gamma_t\,(\sigma_{t,s}^2 + \sigma_{t,l}^2 - 2\rho_{t,s,l}\sigma_{t,s}\sigma_{t,l}). \tag{8}$$

From this we can infer that to guarantee the positive expected utility level, $\bar{U}^+$, the investor now requires rather than expects a larger (smaller) return of the long-short size portfolio, $\tilde{\mu}_{t,s} - \tilde{\mu}_{t,l}$, for higher (lower) levels of risk aversion, ceteris paribus. This is only not true iff $\sigma_{t,s} = \sigma_{t,l}$ and $\rho_{t,s,l} \neq 1$, which is ruled out as long as the investor mimics the size factor.[10]

The empirical relationship between the time-varying investor's risk aversion and the required size return spread is supported by Baker and Wurgler (2006) and Kamstra et al. (2003). Baker and Wurgler (2006) fuse multiple sentiment measures by their joint first principal component. When this measure of sentiment is above average, the size spread reverses, and small stocks earn a moderately low subsequent return. In contrast, when sentiment is below average, investors favor larger stocks over smaller ones, resulting in a predominantly positive size premium in the subsequent period. These results reinforce the above proposition based on Equation (8) in the sense that a large spread is associated with high risk aversion, i.e., that investors are less attracted to risky investments. Moreover, Kamstra et al. (2003) provide a strong link between the investor's mood and degree of risk aversion. The authors

[10] Arguably, the size long-short spread may be a noisy measure of risk aversion as it does also depend on other factors, which are outside the scope of this Chapter.

find that seasonal affective disorder (a form of mental depression caused by winter's short days) increases investors' aversion to risk. When the days become shorter in the fall, these investors favor safer assets, resulting in higher returns once the days grow longer and their risk aversion subsides. Hence, these studies jointly support our proposition based on Equation (8), i.e., that there seems to be a positive relationship between the required size return spread and the investors' risk aversion.

Previously, we stated that focusing on quadratic, mean-variance utility is not restrictive and that our conclusions also apply for popular utility functions incorporating higher-order moments, such as skewness. To see this, assume that the function $g(r_{t,s}, r_{t,l})$ is a measure of the long-short portfolio skewness. Furthermore, δ_t is the degree to which the investor appreciates skewness in period t. With a cubic utility function like the specification in Mitton and Vorkink (2007), the required long-short spread, which is initially denoted in Equation (8), is given by the following expression:

$$\tilde{\mu}_{t,s} - \tilde{\mu}_{t,l} = \bar{U}^{+} + \gamma_t \left(\sigma_{t,s}^2 + \sigma_{t,l}^2 - 2\rho_{t,s,l}\sigma_{t,s}\sigma_{t,l}\right) - \delta_t \, g(r_{t,s}, r_{t,l}). \tag{9}$$

Obviously, the first-order partial derivative of the long-short spread given in Equation (9) with respect to the risk aversion parameter is independent of the investors' attitude towards skewness:

$$\frac{\partial(\tilde{\mu}_{t,s} - \tilde{\mu}_{t,l})}{\partial \gamma_t} = \sigma_{t,s}^2 + \sigma_{t,l}^2 - 2\rho_{t,s,l}\sigma_{t,s}\sigma_{t,l}. \tag{10}$$

In fact, and more generally, any factor driving the long-short return spread similar to $g(\cdot)$ does not interfere with our results. Furthermore, note that the latter measures marginal positive changes in the level of risk aversion given a marginal change in the required size long-short spread, ceteris paribus. Thus, the approximation can only be applied to short-term returns.

The above outlines how the required long-short risk factor spread, which we approximate by the realized long-short risk factor spread, may measure the market's degree of risk aversion. Now, given that the degree of risk aversion is time-varying, we are particularly interested in the variation of investors' reaction to returns under conditions of high and low risk aversion. That is, in the following we focus on the investors' reaction to the revelation of downside risk under different levels of risk aversion and solely from a behavioral point of view, i.e., irrespective of the investor's underlying economic rationales for investing in the respective assets. We expect a considerable difference in the investors' reaction to negative returns compared to positive returns when we measure a high level of risk aversion. Specifically, whenever the aggregate risk aversion in the markets is relatively high, we expect investors to become anxious about investments in stocks with a preceding negative return. For that, suppose an investor is afraid of uncertainty and, thus, demands low-risk stocks. In that case, the investor should be willing to sacrifice a considerable return to avoid stocks which reveal the downside of risk by realizing a negative preceding return. The negative return reveals the downside risk associated with the uncertain investment and induces anxiety among investors. The anxiety makes the stock appear riskier than those that did not incur losses in the previous period. On the other hand, when investors are less risk-averse, we anticipate no extraordinary reaction to negative firm-level returns.

The sophisticated findings in the field of behavioral finance undermine the investor's reaction to the revelation of the downside risk. For instance, Guiso et al. (2018) and Weber et al. (2013) both measure a significant increase in risk aversion following the financial crisis in 2008. This increase in investors' risk aversion holds true even for those who were only marginally affected financially by the crisis, i.e., those who did not encounter significant losses. Furthermore, the authors find that the effect can be reconstructed by an arbitrary exposure to fear, as they measure the change in risk aversion of subjects before and after they watch a portion of a horror film and find qualitatively similar results. These findings suggest that, in general, the presence of fear hinders investors from taking risky investments in the short term (immediately after the fear has been induced), particularly if the fear is associated with the risky investment itself.

Further evidence on the investors' reaction to downside risk is also provided in Cohn et al. (2015). Using a survey that frames either a boom or a bust market, they compare the investment behavior of financial professionals while keeping the ex-ante success probability of risky investments constant. Those participants who were exposed to a bear market environment exhibited a significantly greater aversion to risk. That is, past returns impact the future decisions of investors, such that displaying a chart with the realization of downside risk in the form of negative returns is sufficient to substantially reduce the investor's willingness to take risk. Closely related, using a controlled experimental market environment and face-reading software, Breaban and Noussair (2018) find that fear among the participants results in lower subsequent prices of the simulated assets.

The latter results are also supported by Kuhnen and Knutson (2011) who show that anxiety significantly increases the reluctance to take risky investments. Moreover, Gilad and Kliger (2008) provide additional evidence that financial professionals' risk-taking is influenced by the subject's priming. For instance, when provided with a positive outcome of a risky investment before decision-making, the participants were more willing to take risks, and vice versa. As a consequence, we anticipate that highly risk-averse investors will avoid stocks that reveal downside risk through a previous negative return, resulting in lower returns in the subsequent period for these stocks compared to stocks that appear to be less risky, i.e., the anxiety anomaly.

On the other hand, Lakonishok et al. (1994) show that investors are much more attracted to stocks with large past growth in fundamentals such as sales or earnings compared to value stocks. However, the authors empirically find that the latter stocks strongly underperform value stocks in the long run. They argue that the preference for growth stocks may hold true for individuals and institutional investors. While investors falsely project past growth rates into the future, individuals may prefer past growth stocks as easy-to-justify selections. If this irrational behavior persists, investors are likely to be more attracted to past winners than past losers.

4.3 Empirical analysis

This section empirically assesses the existence and magnitude of the anxiety anomaly. We first present our baseline results, which show that there is in $t+1$ a large positive difference in the stock returns with a preceding positive and negative return when there is a high level of risk aversion in period t. Next, we provide evidence that the anomaly is short-lived, i.e., the effect vanishes or even reverses. We supplement our analysis by performing a variety of robustness tests, which demonstrate that the anomaly is stable over time, independent of specific stock classes, and not unique to the long-short risk factor spread as a measure of risk aversion. Furthermore, we present additional evidence that risk aversion drives the anomaly as we find a convex reaction to negative foregoing returns, which is in line with the prospect theory (Kahneman and Tversky, 1979). Lastly, we demonstrate that the anomaly's out-of-sample predictability is comparable to that of current machine learning-based predictive models, despite the fact that these models employ a substantial number of additional predictors.

4.3.1 Data

We use standard monthly return data from the CRSP. The data set spans from 1966-07-29 to 2021-07-30 and contains 33,824 companies with an average of 115 monthly returns per company. The mean (median) market capitalization is \$1.87 (\$1.10) billion, while the mean (median) monthly return is 1.24% (0.14%), with a return standard deviation of 1.70%.

We merge this return data with long-short risk factor spreads from Chen and Zimmermann (2022). To increase the ease of reproducibility in the finance literature, the authors successfully replicate 319 risk factors. On their website[11], the authors provide a vast dataset containing monthly long-short factor spreads. We drop factors where we find more than 5% missing values. In order to maintain a high level of comparability across risk factors, we subsequently eliminate the months with missing values. Lastly, the data set includes over 3,900,000 data points and 144 risk factors.

4.3.2 Empirical results for the anxiety anomaly

We hypothesize that when there is a large demand for safety, i.e., high risk aversion, stocks with a negative return in the preceding period are avoided, resulting in subsequent returns that are significantly lower than those of stocks with preceding positive returns. To this end, we build upon the arguments and findings presented in Section 4.2 and use the long-short factor spread to quantify the degree of risk aversion. To remain consistent with the intuitive explanations provided in the previous section and to facilitate interpretation, we focus on the well-known size factor (Banz, 1981), although other factors could also be used, as one of our robustness checks demonstrates. We split the data into two levels of risk aversion. A low (high) level of risk aversion is given when the long-short size factor spread is below (above) the median spread. However, the following results are qualitatively similar when using an alternative threshold, e.g., a spread of zero.

As we want to assess whether risk aversion drives short-term predictability, we measure the following return in month $t + 1$ for each of the four categories: high or low risk aversion paired with a positive or negative preceding stock return, both measured in period t. Our goal is to measure the anxiety anomaly; hence, we are interested in periods with a high demand for safety, i.e., a high long-short risk factor spread. Thus, we compare the mean return of stocks with a positive preceding return to the mean return of stocks with a negative preceding return. Our theoretical outline predicts that the group of stocks with a preceding positive return should have significantly larger mean returns than those with a preceding negative return.

Foremost, we focus on periods with an above-median long-short spread in the size factor. The split results in 313 months with an above-median spread, in which there are in total 891,016 observations for stocks with a preceding negative return. 1,063,132 observations are associated with stocks that have a preceding positive

[11]The data can be downloaded from their website: www.openassetpricing.com.

	Mean return in $t+1$ if $r_{t,i} > 0$	$r_{t,i} < 0$	Spread	t-value of test on equal means
High risk aversion	1.190%	0.174%	1.016%	3.37
	(0.286%)	(0.415%)	(0.301%)	
Low risk aversion	1.433%	1.936%	-0.503%	-1.52
	(0.244%)	(0.442%)	(0.332%)	

Table 12: Anxiety anomaly return spreads.
Reported are the mean returns of stocks with preceding positive and negative returns, their mean difference, corresponding two-way clustered standard errors (in parentheses), and t-test statistics on equal means (i.e. zero spread) for periods of high and low levels of risk aversion. High (low) levels of risk aversion are defined as periods of above (below) median long-short size spread in period t, respectively.

return in months with an above-median spread, i.e., a high level of risk aversion. We compute the corresponding mean, i.e., the equally weighted, return for each of the two stock groups. To account for correlation across stocks within a month and potential within-firm correlation, we compute two-way clustered standard errors (Cameron et al., 2011).[12] The results are reported in the first row of Table 12. In line with our expectations, we find that stocks with a preceding negative return earn a mean return of only 0.17% in the following month, whereas stocks with a preceding positive return earn a seven times larger mean return of 1.19%. The difference of 1.02 percentage points in mean returns is statistically significant.

After splitting the data into periods with a large (above-median) and small (below-median) long-short spread in the size factor, we outlined the effects of the preceding return sign on the following magnitude of returns. However, as the anxiety anomaly suggests that a negative (positive) return sign is associated with a small (large) return in the subsequent period during periods of high risk aversion, we question whether the monthly return data contains a short-term momentum effect. To assess whether momentum drives the results, we measure the return spread between stocks with preceding positive and those with preceding negative returns during periods of below-median long-short spreads in the size factor, i.e., periods of low risk aversion. As we do not expect less risk-sensitive investors to flee from stocks that prevail on the negative realization of a risky investment by a negative return, the analysis of the return spread allows us to control for a momentum effect. More

[12]The clustered standard errors and reported t-statistics are obtained by regressing $r_{t+1,i}$ on appropriately chosen specifications of dummy variables indicating low (high) risk aversion and positive (negative) preceding returns as well as interactions thereof.

precisely, if we find that the mean of the spread is also positive during periods of low risk aversion, it is likely that the observed significant mean spread during periods of high risk aversion is generated by a momentum effect rather than the mechanism underlying the anxiety anomaly.

The corresponding results are presented in the second row of Table 12. The mean return of stocks with a preceding positive return is 1.43%, while it is 1.94% for stocks with a preceding negative return, implying a negative return spread of -0.503%. Hence, the control set contains a noticeable short-term reversal, which is in line with the findings of Jegadeesh (1990). The observed anxiety anomaly is therefore not due to a momentum effect, which is consistent with the findings of Shen et al. (2017): during periods of preceding high sentiment, or in our setup, during periods of a low level of risk aversion, investors demand a significantly lower risk premium.

Given that the data set contains a short-term reversal, we expect that the anxiety anomaly may be even more pronounced than reported in Table 12 as a short-term reversal has a negative effect, opposite to the anxiety anomaly. Therefore, we control for short-term reversal and quantify the effect of the anxiety anomaly by estimating the following regression model:

$$r_{t+1,i} = \beta_0 + \beta_1 RA_t^l + \beta_2 R_{t,i}^+ + \beta_3 RA_t^l \cdot R_{t,i}^+ + \beta_4 r_{t,i} + \varepsilon_{t,i}, \tag{11}$$

where RA_t^l is a dummy variables that is one if we measure a low level of risk aversion in period t and zero else, $R_{t,i}^+$ is a dummy variable that is one if the return of stock i in period t is positive and zero else, and where the inclusion of the return of stock i in period t, $r_{t,i}$, accounts for short-term reversal.

The regression results confirm our expectations that the anxiety anomaly is even stronger when we account for the short-term reversal effect. First, the estimate of β_4 is -0.0249 and statistically significant with a t-statistic of -2.56 based on two-way clustered standard errors, confirming a general short-term reversal. Second, during periods of high risk aversion, stocks with a preceding positive return earn on average $\hat{\beta}_2 = 1.52$ percentage points more than those with a preceding negative return, ceteris paribus. This difference in the returns is statistically significant with a t-statistic of 5.69 and noticeably larger than the one reported in Table 12, where we did not adjust for short-term reversal.

It is well known that the month of January is unique and, thus, requires extraordinary attention regarding asset pricing phenomena (Rozeff and Kinney, 1976). Reinganum (1983) shows that the January effect, especially for small stocks, is consistent with tax savings from the prior year. Stocks that lead to losses in the investors' portfolios are sold to realize the loss and reduce tax expenditures. Using

Equation (11), we estimate for Januaries a return spread of 1.25 percentage points after accounting for short-term reversal. Although statistically insignificant, the magnitude of the spread is closely comparable to that of the remaining months, which suggests that the anxiety anomaly may remain, at least economically, relevant for the month of January.

In summary, the anxiety anomaly occurs when investors experience high risk aversion and, thus, have a high demand for perceived safety. Then, investors tend to avoid stocks that recently showed the negative realization of risk, i.e., have a preceding negative return. The effect is statistically and economically significant as it counteracts short-term reversal and can also be observed for the special month of January.

4.3.3 Time persistence

The large spread in returns strongly supports the hypothesis that investors avoid stocks that reveal risk when investors fear losses intensely. On the other hand, we anticipate that the effect is unique and occurs only immediately after a decline in stock prices. We expect the anxiety anomaly to be a short-term phenomenon, given that the described investors' behavior is unrelated to the firm's true intrinsic value and that stock returns are reversing over the short term. To confirm that the anxiety anomaly does not persist to $t+2$, we calculate the mean returns of each stock group for the subsequent period $t+2$. To this end, we continue to divide the periods into two types based upon whether the size long-short spread is above- or below-median in period t. We maintain the same stock splits within each period according to their return sign in period t. The only difference is that we now focus on the returns for period $t+2$ rather than period $t+1$ as we did previously. Again, we are most interested in the return difference in periods with high risk aversion (high long-short size spread). The results confirm our theoretical predictions. The anxiety anomaly reverses such that stocks with a negative return in period t yield a mean return of 1.32% in period $t+2$ while those with a positive return in period t yield only 1.21% on average in period $t+2$.

We are interested in not only whether the anomaly persists into later periods but also whether it can only be observed during specific market phases in the past. For instance, we may observe that the anomaly is driven by early periods as trading becomes less expensive and information becomes more readily available. We do not believe that the anomaly is a result of market inefficiency. Instead, predictability stems from established human behavior. Even if investors act in accordance with the anxiety anomaly, their capital allocations may still be efficient as long as they continue to maximize their utility. Thus, we anticipate the occurrence of the anomaly every decade.

In order to test whether the anxiety anomaly is present in every decade, we estimate the coefficients of Equation (11) for every decade. That is, we estimate the magnitude of the return difference of stocks with a preceding positive and negative return during periods of high risk aversion for each decade in our sample separately. As our sample starts in July 1966 and ends in July 2021, this gives five complete decades and two incomplete ones.

The estimated coefficient that measures the anxiety anomaly after accounting for short-term reversal, i.e. β_2, is positive in every decade and ranges between 0.49% and 5.55%, except for the first incomplete sample from 1966 to 1970, in which it is insignificantly negative at -0.36%. Note, however, that there are half as many companies in our data set from 1960 to 1970 as there are from 1970 to 1980, with the number of companies increasing over time. The anxiety anomaly is highly statistically significant during the decades starting in 1970, 1980 and 1990 (at the 1% significance level using two-way clustered standard errors). We observe a modest significance for the decades starting in 2000 and 2020 (at the 10% significance level using two-way clustered standard errors), while the coefficient is insignificant from 2010 to 2019. The largest estimated strength of the anxiety anomaly is given from 2020 onwards, as the estimated spread between stocks with preceding positive and negative returns given high risk aversion is 5.55% per month, after accounting for a general short-term reversal. Overall, we conclude that the effect is always present but varies in magnitude. The results are consistent with our hypothesis that the anxiety anomaly stems from human utility and is not the result of mispricing. In the following, we focus on the anomaly's robustness and alternative empirical explanations.

4.3.4 In-sample predictive power

So far, we have only examined equally weighted average returns. Therefore, we cannot yet rule out the possibility that the effect is limited to small stocks. To ensure that this is not the case, we split the data into three groups based on market capitalization and then compare mean returns in periods of high risk aversion. We determine the median market capitalization of each company throughout its listed history and then separate the 20% smallest and largest companies from the remaining 60% of medium-sized companies. For periods of high risk aversion, we determine the difference in returns of stocks with preceding positive and negative returns within each subgroup. The results indicate that the anxiety anomaly is stable across all subgroups. While the mean return differences for small (0.72%) and large (0.77%) firms are nearly identical, they are most pronounced for medium-sized firms (1.19%). For periods with low risk aversion the same mean return differences are negative across all subgroups, i.e., -1.16%, -0.36%, and -0.61% for small-, medium- and large-sized firms, respectively. The results further indicate that the anomaly is not a matter of mispricing but originates from investors' preferences. If the anomaly were due to mispricing, we would expect a larger anxiety anomaly in the subgroup of small stocks, as they are less liquid and therefore offer fewer arbitrage opportunities. However, there is no discernible and significant difference in the magnitude of the anxiety anomaly between subgroups of small and large firms. Consequently, as the anomaly is not related to liquidity, there is little evidence in favor of mispricing.

We fundamentally assume that the sign of returns causes anxiety and a run out of stocks with a preceding negative return, given a high level of risk aversion. A preceding negative return induces fear for the investors, which is leveraged by the high level of risk aversion. Thus, we assume that the sign of the preceding return in conjunction with a large long-short risk factor spread has unique explanatory power for the following period's return in addition to the magnitude of the preceding return and the long-short spread. To test this hypothesis, we estimate the following linear model:

$$r_{t+1,i} = \beta_0 + \beta_1 LSSpread_t + \beta_2 r_{t,i} + \beta_3 AnxietyDummy_{t,i} + \varepsilon_{t,i}, \qquad (12)$$

where $r_{t+1,i}$ and $r_{t,i}$ are the returns for stock i in $t+1$ and t, respectively, and $LSSpread_t$ is the value of the long-short size factor spread in period t. $AnxietyDummy_{t,i}$ signals a possible run out of stock i in period t due to anxiety, i.e., it equals one when there is a large long-short factor spread coupled with a negative return of stock i in period t and zero else.

If the anxiety dummy captures predictive information not yet included in the spread and returns itself, we expect a negative coefficient β_3. We find that the spread has no significant explanatory power while the coefficients of the return ($\hat{\beta}_2 = -0.0260$) and the anxiety anomaly ($\hat{\beta}_3 = -0.0163$) are both significantly negative. In fact, the inclusion of the anxiety dummy leads to a large increase in the R^2 by 12 percentage points, giving an R^2 of 0.17% (the results are reported in more detail in Table 14, specifications (2) and (3)).

So far, we have portrayed the magnitude of the anxiety anomaly using the size factor. Therefore, it is of interest to determine whether a combination of factors strengthens the signal and, hence, to include other risk factors as well. However, we cannot simply combine all available risk factors given by Chen and Zimmermann

(2022), i.e., consider only periods where all risk factors signal a high degree of risk aversion, as there are essentially no periods in which all long-short risk factor spreads are above their respective medians. Consequently, we adhere to the line of research that employs sparse factor models based on ex-ante theoretical considerations and restrict the analysis to each pair of four factors that are closely related to those in the popular models of Hou et al. (2015) and Fama and French (2015).

The number of anomalies discussed in the literature rapidly inflates, but Hou et al. (2015) find that out of 80 regarded anomalies, only half pass the 5% test threshold. In the spirit of a sparse representation, they propose a four-factor model, consisting of the well-known market and size factors of Fama and French (1993) and an investment factor that measures the aggressiveness with which a company undertakes investments. Lastly, they include the firm's profitability and show that these four factors account for the majority of the variance explained by prevalent anomalies. Simultaneously, Fama and French (2015) extend their original three-factor model by adding an investment and profitability factor and conclude that these factors summarize the explanatory power of a value factor. Dropping value, the resulting five-factor model serves as a valuable successor to the primal three-factor model.

Table 13 reports the mean returns of stocks with preceding negative returns in three types of periods; the first is a state of low risk aversion and is determined as none of the factors being above its respective median value. In the second state, the degree of risk aversion is ambiguous and is determined when one factor signals high risk aversion but the other does not. The third is given if both factors jointly signal high risk aversion.

The results indicate that combining two risk factors may improve the measurement of risk aversion, as risk aversion jointly indicated by both factors reduces the following return of stocks with a previous negative return by a remarkable amount in any combination. The two most notable combinations are operating profitability combined with either size or investments. When measuring risk aversion by any of these two combinations of factors, we observe that returns for stocks with preceding negative returns are, on average, over 2.6 percentage points lower than stocks with a preceding positive return during periods of high risk aversion. The results suggest that the combination of factors can improve the measurement of risk aversion and, thus, indicate that the magnitude of the anomaly may be even greater than what we measure by the size factor alone.

		Low risk aversion	Ambiguous	High risk aversion
Size	$r_{t+1,i} \mid r_{t,i} < 0$	1.94%		0.17%
		(0.44%)		(0.42%)
	N	897,907		890,323
BM	$r_{t+1,i} \mid r_{t,i} < 0$	1.42%		0.75%
		(0.42%)		(0.44%)
	N	826,971		961,259
OP	$r_{t+1,i} \mid r_{t,i} < 0$	1.75%		0.49%
		(0.48%)		(0.40%)
	N	806,817		981,413
Investments	$r_{t+1,i} \mid r_{t,i} < 0$	1.74%		0.39%
		(0.46%)		(0.41%)
	N	888,951		899,279
Size + BM	$r_{t+1,i} \mid r_{t,i} < 0$	1.43%	1.36%	-0.10%
		(0.53%)	(0.52%)	(0.52%)
	N	2,624,706	729,536	561,023
Size + OP	$r_{t+1,i} \mid r_{t,i} < 0$	1.48%	1.38%	-1.15%
		(0.91%)	(0.36%)	(0.63%)
	N	2,427,545	1,103,704	384,016
Size + Investments	$r_{t+1,i} \mid r_{t,i} < 0$	1.45%	1.64%	-0.13%
		(0.53%)	(0.59%)	(0.47%)
	N	2,811,764	417,400	686,101
BM + OP	$r_{t+1,i} \mid r_{t,i} < 0$	1.41%	1.26%	0.18%
		(0.68%)	(0.42%)	(0.58%)
	N	2,552,736	782,386	580,143
BM + Investments	$r_{t+1,i} \mid r_{t,i} < 0$	1.31%	2.02%	-0.22%
		(0.53%)	(0.54%)	(0.49%)
	N	2,655,154	659,684	600,427
OP + Investments	$r_{t+1,i} \mid r_{t,i} < 0$	1.39%	1.63%	-1.28%
		(0.98%)	(0.35%)	(0.62%)
	N	2,418,461	1,112,916	383,888

Table 13: Returns of stocks with a preceding negative return for factor pairs.
This table reports mean returns for stocks with preceding negative returns using individual risk factors or combinations thereof to signal the degree of risk aversion. We refer to low risk aversion if the considered risk factors are below their respective medians and to high risk aversion if the considered risk factors are above their median values. Ambiguous periods occur in the case of factor combinations if only one of the two factors is above its median. The considered factors are Size (Banz, 1981), Book-to-market value (BM) (Fama and French, 1992), Operating Profitability (OP) (Fama and French, 2006) and changes in total assets (Investments) (Cooper et al., 2008). Two-way clustered standard errors are given in parenthesis, and the number of samples N in the respective subgroup is displayed in the corresponding third row.

4.3.5 Investor's convex reaction to fear

According to the prospect theory (Kahneman and Tversky, 1979), losses are perceived in a convex manner. Suppose investors dump stocks with preceding negative returns because they are highly risk-averse. In this case, according to prospect theory, the effect should become more pronounced as the preceding return becomes more negative but should stop worsening in the case of extremely negative returns. To test for convexity, we extend the linear regression model in Equation (12) by two interaction terms:

$$\begin{aligned} r_{t+1,i} = \beta_0 + \beta_1\ LSSpread_t + \beta_2\ r_{t,i} + \beta_3\ AnxietyDummy_{t,i} \\ + \beta_4\ r_{t,i}|AnxietyDummy_{t,i} + \beta_5\ r^2_{t,i}|AnxietyDummy_{t,i} + \varepsilon_{t,i}. \end{aligned} \tag{13}$$

The first term, $r_{t,i}|AnxietyDummy_{t,i}$, that interacts returns with the anxiety dummy allows to measure whether the anomaly increases in proportion to the preceding negative return. The second term, $r^2_{t,i}|AnxietyDummy_{t,i}$, that interacts the squared return with the anxiety anomaly, allows to measure whether this increase is convex, i.e., whether the investors' reaction is relatively weakening, the smaller the preceding negative return becomes. If convexity holds, we expect both coefficients β_4 and β_5 to be significantly positive, while the coefficient of the anxiety dummy must remain significantly negative in the first place.

Table 14 reports the corresponding least squares estimation results. Additionally, we present results for alternative specifications in which regressors are gradually omitted. When omitting the regressors that contain the anxiety dummy, the specification yields an in-sample R^2 of only 0.05%. The full specification (5) instead exhibits a R^2 of 0.2% and reveals that not only does the anxiety dummy have a high predictive power for the cross-section of returns, but also the reaction to negative returns during periods of high risk aversion is convex. That is, investors are willing to give up a larger return, the greater the anxiety induced by a negative foregone return of stock i. Specifically, solely a negative sign in the preceding stock return leads, on average, to a 0.87 percentage point decrease in return in the following period. Each percentage point decrease in the preceding negative return decreases the following return even further by 0.14 percentage points on average. However, the willingness to pay for perceived certainty plateaus at substantial monthly foregone losses.

Based on the estimation results presented in Table 14, we can use Equation (13) to estimate the return an investor is willing to forfeit in order to escape stocks that

		β_0	$LSSpread$	$r_{t,i}$	$AD_{t,i}$	$r_{t,i}\|AD_{t,i}$	$r^2_{t,i}\|AD_{t,i}$
	estimate	0.01220***	-0.00099				
(1)	t-statistic	5.44	-1.15				
	R^2	0.04%					
	estimate	0.01232***	-0.00094	-0.01123			
(2)	t-statistic	5.38	-1.11	-1.22			
	R^2	0.05%					
	estimate	0.01603***	-0.00018	−0.02595***	−0.01628***		
(3)	t-statistic	6.90	-0.21	-2.80	-3.94		
	R^2	0.17%					
	estimate	0.01610***	-0.00016	−0.02760***	−0.01464***	0.02059	
(4)	t-statistic	6.94	-0.19	-2.82	-3.97	0.51	
	R^2	0.17%					
	estimate	0.01611***	-0.00016	−0.02760***	−0.00874***	0.14377***	0.31973***
(5)	t-statistic	6.94	-0.19	-2.82	-2.84	3.69	2.77
	R^2	0.20%					

Table 14: Statistical analysis of the convexity in investors' risk aversion.
This table presents the least squares estimation results of the regression model given in Equation (13). The t-test statistics are based on two-way clustered standard errors. The long-short factor spread is measured by the size factor, which also determines the level of risk aversion. Note that AD denotes $AnxietyDummy$ for brevity which is one if a stock's i preceding return is negative under high risk aversion. Significance at the 1%-level is indicated by three stars.

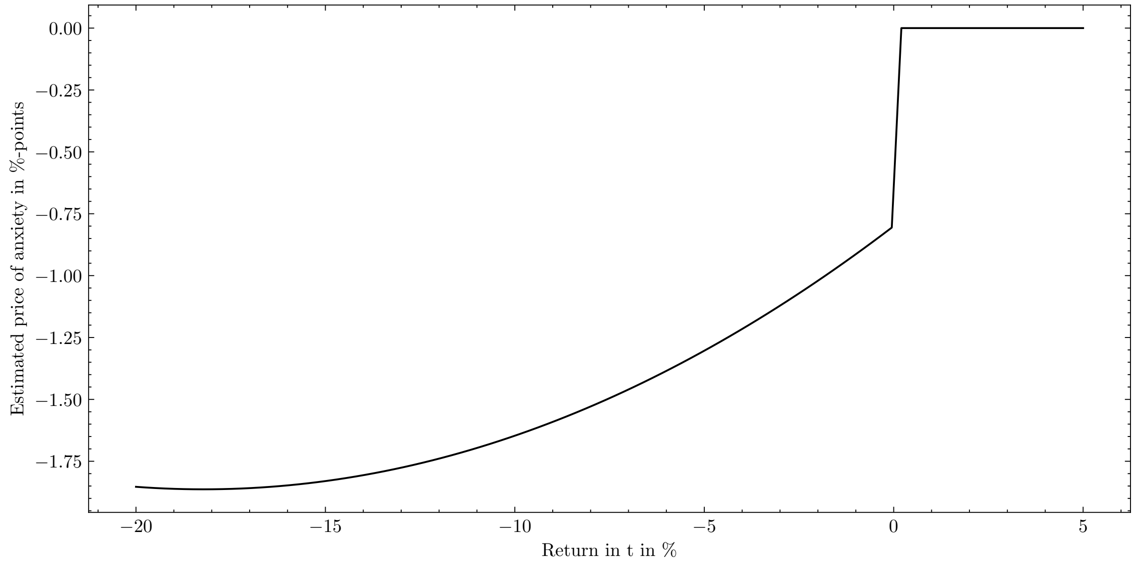

Figure 12: The price of anxiety induced by risk aversion.
This figure depicts the estimated price of anxiety due to risk aversion. Assuming a period t with a high degree of risk aversion (large long-short factor spread), we expect the portrayed return gap in $t+1$ between a stock with a preceding negative return and a preceding return of zero. That is, given the estimates of specification (5) in Table 14, the return gap is given by: $\beta_2\ r_{t,i} + \beta_3\ AnxietyDummy_{t,i} + \beta_4\ r_{t,i}|AnxietyDummy_{t,i} + \beta_5\ r^2_{t,i}|AnxietyDummy_{t,i}$.

show a foregone monthly loss. Figure 12 illustrates the resulting convex preferences of a risk-averse investor. The anxiety anomaly stops increasing in magnitude at a preceding monthly loss of stock i of approximately -15%. At a high level of risk aversion, the maximum expected gap between stocks with a preceding positive and negative return is approximately 1.8 percentage points.

To summarize, the anxiety anomaly is neither captured by the return value nor the long-short risk factor spread. But when we combine them via the proposed dummy variable, we can measure the effects of negative preceding stock returns during times of high risk aversion. Furthermore, investors' preferences in times of high risk aversion are convex in the preceding negative return of stocks. The results favor the hypothesis that risk aversion is a possible explanation for the anxiety anomaly.

4.3.6 Alternative risk aversion measurement

We proposed to measure the degree of risk aversion based on the long-short factor spread of a risk factor, with a particular focus on the size factor, but we also investigated other popular factors like the book-to-market ratio or profitability. Our theoretical considerations suggest that the spread may be a valuable proxy for the level of risk aversion. Nevertheless, as our proxy is novel to the literature and potentially noisy measure, we additionally need to assess whether the anxiety anomaly is specific to the long-short factor or whether it also exists under an alternative measure of risk aversion. To this end, we consider the variance risk premium, as Bollerslev et al. (2009) have shown that it is related to the investors' level of risk aversion. This analysis also addresses the potential concern that the anxiety anomaly may be due to noise in our newly proposed measure of risk aversion.

Following Bollerslev et al. (2009), we define the variance risk premium as the difference between the risk-neutral ex-ante implied variance and the realized variance of the S&P500. As a measure of the implied variance, we employ the end-of-month squared VIX and define realized variance as the sum of squared daily log price differences, i.e., log returns. Using the monthly measurement of the variance risk premium, we measure the anxiety anomaly analogously to our previous definition: an above-median (below-median) variance risk premium signals high (low) risk aversion. We obtain the respective data from the Wharton Research Data Services from January 1990 to July 2021.

Using the variance risk premium to measure the level of risk aversion yields qualitatively similar results. During times of high aggregate risk aversion, the difference in monthly returns between stocks with a preceding positive and a preceding negative return after accounting for short-term reversal is 0.99 percentage points. The difference in returns is statistically significant, with a t-statistic of 2.32 using two-way clustered standard errors. For reference, during the same time horizon, the anxiety anomaly, as measured by the size long-short spread, is 1.28 percentage points (with a t-statistic of 3.43 using two-way clustered standard errors). Although the anomaly is less pronounced when using the variance risk premium as a proxy for risk aversion, our key results remain unchanged. Whenever investors are highly risk-averse, they become anxious to invest in stocks that reveal the downside of risk with a preceding negative return.

4.3.7 Out-of-sample predictability

In the previous sections, we showed that risk aversion reveals a predictable pattern in short-term returns. However, all demonstrated results are in-sample. To measure the out-of-sample predictability of the anxiety anomaly, we analyze Equation (13) in an out-of-sample setting and compare its performance to that of Gu et al. (2020) who employ machine learning methods and report an exceptionally high forecast performance that can currently be interpreted as a kind of benchmark in out-of-sample stock return predictability.

When modeling in-sample returns, the literature usually assumes a linear functional form. However, the linear regression model tends to overfit, especially if many linear predictors, i.e., risk factors, are included, leading to low out-of-sample performance. Gu et al. (2020) confirm this intuition and show that a non-regularized linear regression with hundreds of predictors is incapable of explaining out-of-sample return variation in a large cross-section of stock returns. Besides, it completely discards non-linear predictive information.

Recently, more advanced econometric modeling approaches, such as regularization and shrinkage (Gu et al., 2020; Kozak et al., 2020), principal component analysis (Kelly et al., 2019) or even neural-network-based autoencoders (Gu et al., 2021), are integrated to improve the predictive power. As neural networks allow for non-linearities and regularization, their out-of-sample performance is well suited as a benchmark in cross-sectional asset pricing. However, neural networks exploit deep nonlinear relationships in the data, thereby removing the highly valued interpretability of linear models.

Instead, we build upon our theoretical framework and adopt a simple, i.e., sparse, and simultaneously interpretable setup. In particular, we exploit the fact that returns are a convex function of the magnitude of downside risk revelation. Thus, we specify a linear regression model in the parameters but a non-linear in the predictors, which are related to the anxiety anomaly. Moreover, we non-linearly combine the preceding long-short factor spread with the preceding return. This approach also sheds light on the predictive dependencies found by non-linear machine learning.

Our data set ranges from 1966-07-29 to 2021-07-30. We initially estimate the regression model given in Equation (13) for excess returns using data up to December 1980 and use the estimated model to compute one-step ahead monthly return predictions for the following year, i.e., for January 1981 to December 1981. Each year, we update our coefficient estimates using a forward-rolling approach. This approach gives a total of 42 coefficient updates and years of out-of-sample estimates.

We evaluate the performance using the uncentered, out-of-sample modification of the coefficient of determination used, e.g., in Gu et al. (2020):

$$R^2_{oos} = 1 - \frac{\sum_{t,i}(r_{t+1,i} - \hat{r}_{t+1,i})^2}{\sum_{t,i} r^2_{t+1,i}}, \tag{14}$$

where $r_{t+1,i}$ is the stock i's excess return over the risk-free rate in period t and $\hat{r}_{t+1,i}$ is the corresponding forecast. We determine the metric over all out-of-sample predictions once, i.e., after estimating and predicting yearly.

The specification in Equation (13) yields an R^2_{oos} of 0.42% on the entire cross-section of monthly stock returns. As we are interested in the predictability of the anxiety anomaly, we repeat the procedure using a sparse specification, dropping the values of the preceding return $r_{t,i}$ of stock i and the size long-short factor spread in period t. Surprisingly, R^2_{oos} is only marginally reduced to 0.41%, indicating that the anxiety anomaly may be the underlying source of predictability.

One of the most comprehensive out-of-sample asset pricing studies is arguably Gu et al. (2020), which uses a variety of machine learning models and a variety of stock level characteristics and macroeconomic predictors to forecast monthly excess stock returns. Using the R^2_{oos} metric, the authors evaluate the models' predictive performance with data from 1987 to 2016. In addition, they also update their models annually, and their data contains a number of stocks comparable to ours. We consider their work to be the upper limit of predictability using contemporary machine learning techniques and one of the largest data sets available.

The authors find peak performance when using a neural network with four hidden layers, enabling the learning of non-linear predictive patterns. The model yields an R^2_{oos} of 0.4% when incorporating over 900 predictors, close to the 0.41% we produced with only three variables that measure the anxiety anomaly. The similar performance leads to the following two conjectures. Firstly, a substantial part of the predictability may be attributed to risk aversion. Indeed, all models used in Gu et al. (2020) incorporate the preceding return as the most influential predictor and the size factor as one of the top predictors. The results suggest that investors' preferences are predictable and reflected in the cross-section of returns, resulting in an accurate measurement of short-term risk premia. Secondly, the anxiety anomaly could provide a justification for the superior performance accomplished by non-linear predictors like tree-based learners or neural networks compared to purely linear models.

Last but not least, we are interested in how predictions from annual estimates of Equation (13) translate into portfolio returns. We construct portfolios based on decile sorts and present the top (highest predictions) and bottom (lowest predictions) out-of-sample cumulative returns in Figure 13. More precisely, the figure displays equally weighted cumulative log returns over the out-of-sample period starting in 1980. In addition to the substantial difference in cumulative returns, investors find a nearly threefold greater Sharpe ratio when avoiding stocks that may be affected by the anxiety anomaly (0.941 vs. 0.336). Using value-weighted returns gives a similar picture. The predicted top decile yields an annualized Sharpe ratio of 0.859 vs. 0.518 for the bottom decile. Hence, the given predictions are economically highly valuable. Moreover, the returns are remarkably robust over time, indicating that the results are not driven by external shocks.

Overall, the anxiety anomaly caused by risk aversion and negative returns may

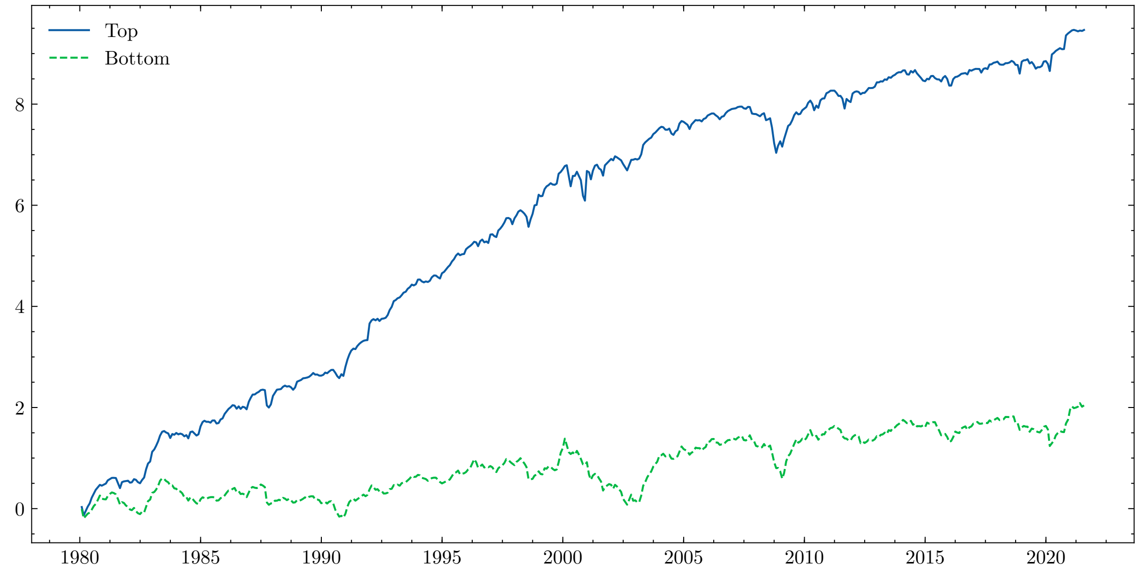

Figure 13: Out-of-sample cumulative log returns.
This figure shows out-of-sample cumulative log returns for the top and bottom deciles based on predictions of the yearly estimated model in Equation (13). Returns are equally weighted. Top (bottom) decile returns yield an annualized Sharpe ratio of 0.941 (0.336).

capture most of the predictability found in the literature. A simple linear specification (Equation (13)) including only one long-short risk factor spread - we use the size factor - and the preceding return yields an R^2_{oos} of 0.42%. The majority of the performance is attributable to the anxiety anomaly alone, indicating that the actual values of the spread and preceding return provide little information for the prediction. Using the resulting predictions based on the simple linear model (Equation (13)), investors can substantially improve realized cumulative returns and Sharpe ratios.

4.3.8 Potential drivers of the anxiety anomaly

In the preceding Section, we discussed the high out-of-sample predictability of the anxiety anomaly, which we further dissect in the following. To this end, we use out-of-sample returns that are generated given the top-minus-bottom decile portfolio based on predictions of the model given in Equation (13). The logarithmic bottom and top returns of this portfolio are also depicted in Figure 13.

We consider four macroeconomic variables that are potentially related to risk aversion and investor sentiment. In particular, Liang et al. (2023) propose that when Monetary Policy Uncertainty (MPU) rises, investors may prefer less risky assets, which may lead to reduced Financial Stress (FS), while Hu et al. (2023) show that MPU significantly affect investor sentiment. We follow their approach and measure MPU via the newspaper-based index developed in Baker et al. (2016)

and the newspaper-based Financial Stress index developed by Püttmann (2018).[13] Additionally, we follow Dai et al. (2023) and incorporate a measure of Crash Risk (CR) into our analysis, which we define as the skewness of the daily returns of stocks in our sample. Acharya et al. (2013) argue that when liquidity is relatively low, investors seek more liquid and less distressed stocks, which is a potential driver of the anxiety anomaly. We follow their approach and define an index for Illiquidity Risk (IR) as the monthly mean of the fraction of the absolute daily return divided by the dollar volume for each stock and day given in our sample.

In addition to these macroeconomic variables that may be related to the risk aversion and investor sentiment, we also control for common risk factors given in Fama and French (2015) and in Carhart (1997). In particular, we incorporate excess returns of a market portfolio (Mkt-RF), returns of a size portfolio (SMB), value portfolio (HML), operating profitability portfolio (RMW), investment strategy portfolio (CMA) and momentum portfolio (MOM) in our analysis.[14]

We merge data on the macroeconomic and common risk factors with the long-short returns of our out-of-sample tests in the preceding section. This results in 249 monthly returns and factor values from December 1984 to November 2016. To make interpretation of the results more convenient, we standardize factors to have zero mean and unit variance. We then regress the anxiety anomaly's long-short return at time t on the contemporaneous values of the factors. Thus, we analyze the characteristics of the realized anxiety anomaly's long-short returns and not the potential predictive power of the macroeconomic and common risk factors on the these returns. Furthermore, as the correlation between the macroeconomic factors is relatively large, we report in addition to the full specification (1), results for individual regressions, i.e., where we regress the anxiety anomaly's long-short return on each macroeconomic and common risk factor individually (specifications (2)-(11)), and consider each macroeconomic factor together with the common risk factors (specifications (12)-(15)).

Table 15 reports the corresponding regression results. Most prominently, the anxiety anomaly generates a positive and significant alpha in every specification, ranging from 1.31% to 1.34%. Unsurprisingly, specification (6) reveals that the anxiety anomaly's long-short returns are negatively correlated with market returns. This behavior is to be expected as the anomaly predicts the underperformance of

[13] Data on MPU and FS is publicly available on `https://www.policyuncertainty.com`.

[14] The data and details on the definitions of the respective factors can be obtained from `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
α	0.0133***	0.0132***	0.0132***	0.0132***	0.0132***	0.0132***	0.0132***	0.0133***	0.0131***	0.0134***	0.0132***	0.0133***	0.0133***	0.0133***	0.0133***
	(4.04)	(3.93)	(3.92)	(3.93)	(3.96)	(3.96)	(3.93)	(3.97)	(3.92)	(4.01)	(3.92)	(4.00)	(3.99)	(4.00)	(4.05)
MPU	0.0016	0.0047										0.0038			
	(0.45)	(1.40)										(1.12)			
FS	-0.0013		0.0024										0.0022		
	(-0.34)		(0.70)										(0.64)		
CR	0.0034			0.0030										0.0028	
	(1.01)			(0.88)										(0.81)	
IR	0.0099***				0.0079**										0.0098***
	(2.72)				(2.38)										(2.92)
Mkt-RF	-0.0078*					-0.0076**						-0.0059	-0.0061	-0.0060	-0.0078*
	(-1.95)					(-2.30)						(-1.47)	(-1.53)	(-1.51)	(-1.97)
SMB	-0.0061*						-0.0031					-0.0046	-0.0044	-0.0046	-0.0059
	(-1.70)						(-0.92)					(-1.27)	(-1.23)	(-1.27)	(-1.64)
HML	0.0053							0.0056				0.0039	0.0037	0.0041	0.0056
	(1.17)							(1.62)				(0.90)	(0.84)	(0.95)	(1.32)
RMW	-0.0001								0.0051			0.0011	0.0010	0.0006	0.0003
	(-0.02)								(1.49)			(0.28)	(0.26)	(0.17)	(0.07)
CMA	0.0016									0.0082**		0.0031	0.0035	0.0034	0.0016
	(0.36)									(2.33)		(0.68)	(0.76)	(0.75)	(0.36)
MOM	0.0053										0.0021	0.0033	0.0033	0.0036	0.0050
	(1.44)										(0.63)	(0.90)	(0.90)	(0.98)	(1.36)

Table 15: Potential drivers of the anxiety anomaly.
This table presents estimation results from regressing out-of-sample long-short returns resulting from top minus bottom decile portfolios based on Equation (13) on macroeconomic and common risk factors. The macroeconomic factors are: Monetary Policy Uncertainty (MPU), Financial Stress (FS), Crash Risk (CR), Illiquidity Risk (IR) whereas the common risk factors are: excess returns on a market portfolio (Mkt-RF), returns on a portfolio based on size (SMB), value (HML), operating profitability (RMW), investment strategy (CMA) and momentum (MOM). All factors are standardized to have zero mean and unit variance for ease of interpretation. The significance levels are given by asterics with *, **, ***, indicating significance at the 10%-, 5%- and 1%-level, respectively. Corresponding t-statistics are given in parentheses.

a set of stocks, thus yielding excess returns when these stocks perform badly. In line with the findings of Acharya et al. (2013), the anxiety anomaly's long-short returns are stronger during periods of high illiquidity risk, i.e., when investors tend to demand less distressed and more liquid assets. Supporting these results, investors significantly prefer firms with more conservative investment policies during these periods. Interestingly, other macroeconomic factors, i.e., monetary policy uncertainty, financial stress, and aggregate crash risk, have no significant explanatory power. Moreover, the estimation results for the momentum factor support our previous finding, that the anxiety anomaly is not due to momentum. Similarly, no other common risk factor appears to explain the anxiety anomaly.

4.4 Anxiety anomaly and neural networks

In the above analysis, we find that the anxiety anomaly yields high out-of-sample predictive performance when compared to neural networks. Interestingly, large networks trained on hundreds of predictors like the one in Gu et al. (2020) incorporate the past return as well as the firm size as two of the most important predictors. Hence, we train a neural network to predict returns using the preceding return as well as the size factor and investigate whether the network converges towards the anxiety anomaly while training.

To this end, we employ a neural network comprising two hidden layers, each consisting of 32 neurons. We use the ReLU activation function, as described in Section 3.2, as well as the Adam optimizer (Kingma and Ba, 2014) for parameter tuning. Given that our primary focus does not lie in evaluating the predictive capabilities of the model, we chose to utilize the complete dataset for training purposes. It is important to acknowledge that we are actively committed to training the model to accurately describe the entire dataset. We do so as we are interested in whether a neural network without any a priori knowledge of the anxiety anomaly will converge towards a parameter set such that it predicts returns in patterns according to the anxiety anomaly given the respective features.

The objective of this analysis is to compare how the neural network maps preceding returns given low and high risk aversion to predictions of future returns at multiple stages of the training procedure. That is, we are interested in the following:

$$\mathbb{E}[r_{t+1}|\theta_n, r_t, LSSpread_h] - \mathbb{E}[r_{t+1}|\theta_n, r_t, LSSpread_l] \ \forall r_t \in \mathcal{R}_t, \qquad (15)$$

where $\mathbb{E}[r_{t+1}|\theta_n, r_t, LSSpread_h]$ is the neural network's prediction of the return r_{t+1} conditional on r_t and $LSSpread$ and given the parameter vector θ_n obtained after n iterations of training, r_t being a fixed value of the preceding return out of a set of values given in $\mathcal{R}_t$, and $LSSpread_h$ ($LSSpread_l$) indicating a high (low) long-short size factor spread. We define the set $\mathcal{R}_t$ as the 40 linearly spaced values between the global minimum and maximum values of r_t within our sample. Given a high (low) level of risk aversion, expressed by $LSSpread_h$, we fix the long-short size factor spread at the 75th percentile (25th percentile) of the size long-short spread of the entire sample.

We evaluate the difference in return predictions given the different levels of risk aversion three times during the training procedure, each after 100 epochs. This results in evaluations with the following in-sample R^2 scores: 0.530%, 0.712%, 0.715%. It is important to note that the network only gains a marginal increase in descriptive power, thus being close to convergence. However, the increase from the first to the second evaluation step is much more noticeable.

Figure 14 shows the corresponding results. The figure reveals multiple interesting insights. First, at an early stage of the training procedure, the neural network identifies that expected returns are low given a high level of risk aversion, while this effect is only marginally larger for low preceding returns than for high preceding returns. Second, the network gradually learns the interconnection between the level of risk aversion and the preceding return. That is, the spread between the expected return of a high and low level of risk aversion is increasing sharply around a preceding return of zero at the second evaluation step (the middle subfigure). This effect is strongly increasing at the third evaluation step.

Third, the spread curve becomes increasingly convex, indicating that the neural network recognizes a convex reaction to negative preceding returns given a high level of risk aversion, similar to our findings depicted in Figure 12. Overall, all of the latter results are in line with the anxiety anomaly. The findings provide further evidence that the anxiety anomaly provides a major source of predictability as training neural networks given these features converges towards the anomaly's characteristics.

4.5 Conclusion

This study investigates the short-term effects of risk aversion on a cross-section of stock returns. We provide a theoretical motivation for measuring investor's risk aversion via the monthly long-short risk factor spreads. Based on our novel short-term measurement of investors' risk aversion, we determine the investors' response to the revelation of downside risk when investors are highly risk averse. We measure the revelation of downside risk by the preceding return sign at the firm level. That is, when we measure high risk aversion, we anticipate that investors will avoid purchasing stocks whose preceding monthly return is negative, as the anxiety caused by the

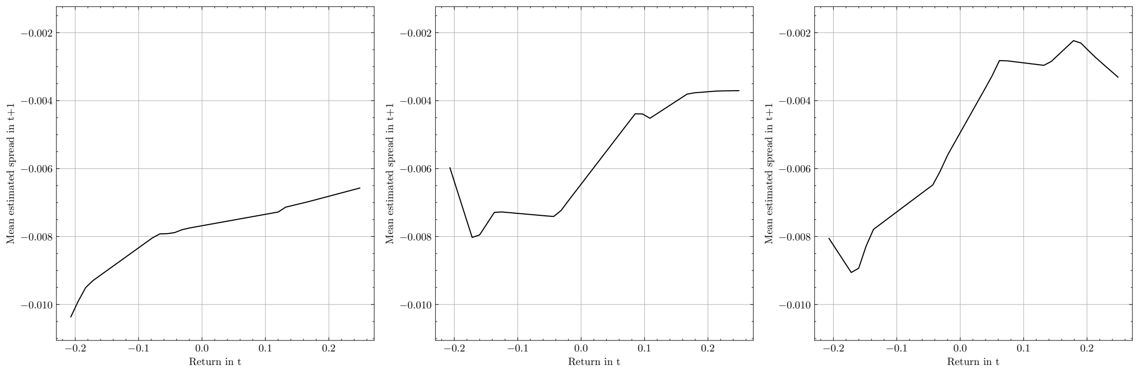

Figure 14: Estimated return spread conditional on the preceding return.
This figure illustrates the spread in estimated returns between a high and low level of risk aversion conditional on the preceding return at progressing stages of the neural network's training procedure, i.e., $\mathbb{E}[r_{t+1}|\theta_n, r_t, LSSpread_h] - \mathbb{E}[r_{t+1}|\theta_n, r_t, LSSpread_l]\ \forall r_t \in \mathcal{R}_t$. The number of completed training epochs increases from left to right by 100, giving parameters vectors θ_{100}, θ_{200}, and θ_{300}, with in-sample R^2 scores of 0.530%, 0.712%, and 0.715%. A high (low) level of risk aversion is defined by fixing the long-short size factor spread to the 75th percentile (25th percentile). We compute the corresponding neural network's predictions of 40 linearly spaced values of r_t within the maximum and minimum values in the entire sample.

negative return prevents the investor from making the investment.

Our expectations are strongly supported by the empirical findings. Stocks that gained in value in the preceding period earn a seven times larger average return than stocks that lost in value during the preceding period when investors are highly risk-averse, which we call the anxiety anomaly. The anomaly is remarkably robust. In particular, it is observed not only for risk aversion measured based on size and other risk factors but also for the well-established measure of risk aversion of Bollerslev et al. (2009). Moreover, the anomaly is present in every decade and stock size.

The anxiety anomaly has its theoretical foundation in risk aversion and fear response. We argue that investors become anxious to invest in a stock that exhibits downside risks by a negative return when investors fear losses the most, i.e., are highly risk-averse. Hence, the reaction to a preceding negative return is supposed to be convex in shape, in line with the prospect theory (Kahneman and Tversky, 1979). Our empirical analysis shows that the convex reaction to preceding negative returns is indeed highly significant. Therefore, the greater the preceding negative return, the greater the investors' reluctance to purchase the stock, resulting in a relatively low return compared to the rest of the market.

We further show that the anxiety anomaly exhibits a remarkably high and stable out-of-sample predictive power that is very close to the recently reported extraordinary performance of neural networks trained on a large data set containing a plethora of different predictors (Gu et al., 2020). Interestingly, according to the neural networks, the preceding return and firm size, i.e., the two variables that determine the anxiety anomaly, belong to the most important features. Moreover, we show that a neural network incorporating the preceding return and the long-short size factor spread converges to similar return estimates as predicted by the anxiety anomaly. As such, the anxiety of risk-averse investors may provide a possible explanation for the non-linear pricing factors revealed by deep learning.

Moreover, analyzing potential drivers of the anxiety anomaly reveals that the anomaly's long-short returns are stronger during periods of high illiquidity risk, i.e., when investors tend to demand less distressed and more liquid assets.

5 Monday Afternoon Reversal

In the first two Chapters, we describe the rise of machine learning in contemporary asset pricing and portfolio management. To this end, we demonstrate that deep learning can generate substantial economic gains for both individuals and institutional investors. In search of an economic rationale for the predictability of returns through machine learning, we dissect the importance and impact of various features and find that the vast majority of short-term predictability stems from the time series of returns itself. Related to these results, we find that the anxiety anomaly, which is based on risk aversion and recent returns, provides substantial out-of-sample predictability of returns.

Motivated by the predictive power of past returns for future returns, we investigate whether other autoregressive patterns at higher frequencies of cross-sectional stock returns exist. By dissecting the half-hourly returns of cross-sectional individual U.S. stocks as well as indices and international Exchange Traded Funds (ETFs), we find that Monday afternoon returns tend to reverse over the remaining week, which we refer to as the Monday afternoon reversal. Similar to the anxiety anomaly, this pattern is remarkably stable over time and cannot be explained by common systematic risk factors. Conversely, we do not find any considerable relationship between Monday morning returns and the remaining week's return.

We aim to solve this puzzling relationship by empirically assessing numerous potential explanations. However, the majority of potential explanations fail to resolve the return pattern: the Monday afternoon reversal is not a result of overnight risk, nor is it driven by small stocks or firm news. The latter results indicate that the reversal is not due to any firm-related characteristics but rather due to the proximity of weekends and a potential disagreement regarding expected returns. To this end, we measure investor disagreement by both volume and volatility on Monday afternoons and find strong support for the latter hypothesis. The reversal is only present among stocks with relatively large volume and volatility on Monday afternoons. In contrast, the returns of stocks with relatively low momentum tend to track their Monday afternoon returns, thereby accumulating return momentum.

With the increasing availability of high-frequency stock market data and the required computational power to process this information, we expect that sophisticated investors will exploit the predictability of patterns such as the Monday afternoon reversal. Thus, we analyze the behavior of a well-studied group of investors - short sellers - on Mondays. Using a large sample of short interest data, we indeed find that when short sellers disagree with Monday afternoon returns, these returns reverse sharply throughout the remainder of the week. Therefore, we presume that knowledgeable investors seek and exploit predictable patterns in the cross-section of return distributions.

Abstract

Using high-frequency data of cross-sectional returns, we document a significant reversal on the Monday afternoon return throughout the week. The effect is stronger for larger firms and cannot be observed for Monday morning returns. Moreover, it emerges during both overnight and intraday trading periods and is not attributable to common risk factors or earnings announcements. Additionally, the reversal also exists in exchange traded funds representing major international equity markets. We find that the reversal is most prevalent when investor disagreement is high, while a significant momentum of returns is observed when investor disagreement is low.

5.1 Introduction

In recent years, we have witnessed a dramatic increase in the availability of detailed stock market data alongside an enormous reduction in computational costs, leading to a substantial rise in the number of studies examining intraday asset pricing anomalies (for instance, see Bogousslavsky (2021); Bollerslev et al. (2020); Gao et al. (2018); Lou et al. (2019)). We contribute to this line of research by revealing a significant reversal on Monday afternoon returns in the cross-section of returns. In particular, we find that stocks exhibiting lower returns on Monday afternoons earn significantly higher returns throughout the remainder of the week than stocks with higher Monday afternoon returns. This observation, which we refer to as the Monday afternoon reversal, is not only evident across multiple U.S. stock market indices but also for international equity ETFs. In addition, we compute the economic gains an investor may accomplish with a long-short portfolio based on Monday afternoon returns, which yields a value-weighted mean annual return of 22.18% and an annualized Sharpe ratio of 1.40.

It has been widely documented that the characteristics of stock returns on Mondays differ from those on other weekdays (Gibbons and Hess, 1981). Dating back to French (1980), who finds particularly low or even negative stock returns on Mondays, several studies have analyzed the effects of the weekend on Monday returns. Although this weekend effect is well documented, relatively little is known about the effect of Monday returns on the remaining week's return. In fact, to the best of our knowledge, there are only two studies that assess the correlation between daily returns after market closures on weekends and subsequent daily returns.

More precisely, Bessembinder and Hertzel (1993) find that Tuesday returns are negatively correlated with the preceding Monday returns within U.S. and Japanese equity indices, U.S. portfolios, Dow 30 firms, as well as futures returns. Although the authors find that this negative correlation is robust to outliers and consistent over time, they offer no explanation for the anomaly. Ülkü and Andonov (2016) analyze daily returns of international equity indices in a more recent sample and find that the reversal did not disappear over time and that the correlation is also present between daily Monday returns and the subsequent four days' return. In this Chapter, however, we use high-frequency data to analyze the timing of the Monday reversal in more detail, as we conjecture that the effects of Monday morning returns may differ from Monday afternoon returns. Our empirical results indeed reveal, among other findings, that the reversal is specific to the Monday afternoon, while we find that the Monday morning appears to be irrelevant.

There are various reasons that motivate our decision to specifically focus our study on the Monday afternoon return. Firstly, Coval and Shumway (2005) find that investors are substantially more willing to take risks when they suffer morning losses. This suggests that if traders are emotionally influenced by their morning trading performance and consequently increase their risk tolerances in the afternoon, we may observe irrational pricing during the afternoon that is subsequently corrected throughout the following days. Our empirical findings indeed support the separation of morning and afternoon returns. While the Monday afternoon return is a significant negative predictor of the week's remaining return, we find no evidence that the Monday morning return has any predictive ability.

The investors' emotional involvement may even be leveraged by a general drop in mood from weekends to Mondays, as individuals tend to be more satisfied with their jobs (Taylor, 2006), are significantly happier and less worried (Helliwell and Wang, 2014) when surveyed on weekends than during the week. Birru (2018) focuses explicitly on the differences in mood levels on Fridays (high mood) and Mondays (poor mood) and links directly these differences to returns from cross-sectional anomalies. Specifically, the author dissects the anomalies' long-short returns into their speculative and non-speculative legs and finds that the relative profitability of the former against the latter leg is significantly larger on Fridays compared to Mondays. Thus, we expect a stronger reversal for stocks that are particularly appealing to speculative investors, i.e., are more volatile, for which we find strong empirical evidence.

Additionally, several studies related to the weekend effect suggest that low Monday returns may only be observable during the early Monday hours (Smirlock and Starks, 1986) or between Friday close and Monday open (Harris, 1986; Rogalski, 1984). Therefore, excluding the early hours of the Monday trading session and focusing instead on the Monday afternoon returns prevents our results from being blurred by a potential weekend effect.

The strong predictive power of the Monday afternoon return on the remaining week's return demands an economic rationale. Thus, this Chapter presents empirical tests to examine various potential explanations for the observed reversal. At first, Lou et al. (2019) document that intraday returns in the cross-section are reversed by the overnight session, and overnight returns are reversed by the intraday session. The authors also attribute popular trading strategies' profits based on a variety of asset pricing factors to one of the respective sessions. We investigate whether the Monday afternoon reversal is induced by either the intraday or overnight session. Although the reversal is stronger overnight than during the intraday session, there is a positive reversal during the intraday session for all remaining weekdays except Friday.

Another potential explanation for the significant reversal is infrequent balancing. Heston et al. (2010) show that each half-hour intraday return exhibits a strong positive autocorrelation between the intraday return of each half-hour and the respective half-hour in the preceding day. This special pattern may be associated with the daily routines of fund managers, who allocate specific time intervals for rebalancing their portfolios. However, we only observe a reversal on the Monday afternoon return by the week's remaining return and no similar pattern for the other business days. As a result, the daily routines of fund managers are unlikely to be the underlying cause of the observed reversal.

Brusa et al. (2011) analyze daily excess returns of four major stock market indices and show that the majority of the weekend effect can be attributed to the market return, the firm size, and the book-to-market equity value. Therefore, we examine whether the reversal's returns are attributable to one of the five factors proposed by Fama and French (2015) or the momentum factor of Carhart (1997). We find no support for any relation to the latter factors. Moreover, we also reject the hypothesis that the reversal is an underreaction with slow adaptation to firm news. This is evident as the reversal remains and is of similar magnitude when we account for firm news announcements in the previous week, on Mondays, or throughout the remaining week (Jiang et al., 2021).

Ülkü and Andonov (2016) propose that the reversal of Monday returns, measured over the whole day, i.e., using daily returns, can be attributed to relatively lower levels of foreign and institutional trading on Mondays combined with a higher proportion of individuals trading, who are supposedly less informed than the former group of traders. As such, Mondays permit noisy prices caused by uninformed individuals, which are subsequently reversed by informed foreign and institutional traders.

However, Ülkü and Andonov (2016) provide only limited empirical evidence for their proposed explanation of the reversal. Foremost, in their sample, the share of individual relative to institutional and foreign trading value is only marginally higher on Mondays than on the remaining business days. This makes it difficult to definitely attribute price discoveries on Mondays to individuals. Additionally, the authors put forward that individuals' net trading flows are negatively correlated with contemporaneous returns. Conversely, the net trading flows of institutions and foreign investors are positively correlated with contemporaneous returns. That is, if individuals consistently trade as liquidity suppliers against contemporaneous returns, it raises doubts about their ability to contribute to price discoveries on Mondays. This is because it would imply the opposite correlation of individuals' net trading flows with contemporaneous returns on Mondays than on other business days.

In contrast, our analysis of high-frequency price data strongly suggests that the reversal is a result of investors' disagreement about fundamental values. We find a particularly strong reversal for stocks with high half-hourly measured volatility over the Monday afternoon, whereas we observe significant return momentum within stocks with low Monday afternoon volatility. Furthermore, using cross-sectional U.S. short interest data, we show that short sellers are informed traders who effectively anticipate the Monday afternoon reversal in both directions.

The remainder of this Chapter proceeds as follows. Section 5.2 quantifies the Monday afternoon reversal in cross-sectional stock returns using high-frequency data and presents several robustness checks. We account for the firm size, the month of January, the week in the month, common systematic risk factors, and firm news announcements. Furthermore, we analyze the out-of-sample predictive power and potential economic gains and dissect the reversal into its overnight and intraday components. Section 5.3 provides additional evidence of the Monday afternoon reversal in U.S. stock market indices and international ETFs. Section 5.4 provides evidence for the hypothesis that the reversal stems from investor disagreement. Finally, Section 5.5 concludes the Chapter.

5.2 Monday afternoon reversal

5.2.1 Data

The empirical analysis is based on high-frequency stock market data from FirstRate-Data (2022). Our data set comprises 5,754 trading days, or 822 trading weeks, and 1,838 companies over a sample period ranging from January 2007 to October 2022. We analyze half-hourly open and close prices along with the corresponding share volumes, with all prices adjusted for stock splits and dividends. As the U.S. market opens at 9:30am (EST) and closes at 4:00pm, this results in a total of 13 observations every trading day (intraday returns), plus an additional 35 half-hour observations per day and stock for after-hours and pre-market trading (overnight returns), whereas we do not consider weeks in which Monday is a holiday. We merge this dataset with earnings report data from Compustat and monthly updated market capitalizations from the Center for Research in Security Prices.

We determine the lowest cross-sectional mean half-hour trading dollar volume during the day when we suspect that most institutional investors take their lunch breaks in order to split the intraday trading session on Mondays into morning and afternoon trading. Figure 15 illustrates the distribution of dollar volume per half-hour. We calculate the dollar volume by multiplying the share volume of a particular stock by the average of its corresponding half-hour open and closing prices and then averaging across all stocks and days. For ease of interpretation, we transform the values into fractions of the total average volume of the day. As previously documented in Hong and Wang (2000), the distribution is U-shaped, with the highest volume occurring at the first and last half-hour periods. We also observe a relatively high volume during the first half-hour after the market closes at 4:00pm. As the volume distribution within the trading day is roughly mirrored at the close of the half-hour period from 12:30pm to 1:00pm, we define the Monday afternoon return $r_{i,w}$ for stock i in week w as the return between the close of the latter period, i.e., at 1:00pm, and the intraday close of the Monday market at 4:00pm. More precisely,

$$r_{i,w} = \frac{c_{i,w,13}}{c_{i,w,7}} - 1, \tag{16}$$

where $c_{i,w,13}$ is the 13th half-hour intraday close of week w, i.e., the price of stock i at the Monday market close at 4:00pm in week w, and $c_{i,w,7}$ is the corresponding seventh half-hour intraday close, i.e., the price of stock i on Monday at 1:00pm in week w.

Our following analysis focuses on determining the impact of the Monday afternoon return on the remaining weekly return measured from Monday market close to Friday market close. If there is a predictive power of the Monday afternoon return on the remaining week's return, one may want to take advantage of it and trade correspondingly, which involves weekly portfolio revisions. To rule out that the return of that strategy is driven by a potential weekend risk premium, we omit the weekend trading time, i.e., Friday close to Monday open, and rather focus on the impact of the Monday afternoon return on the return from Monday close to Friday

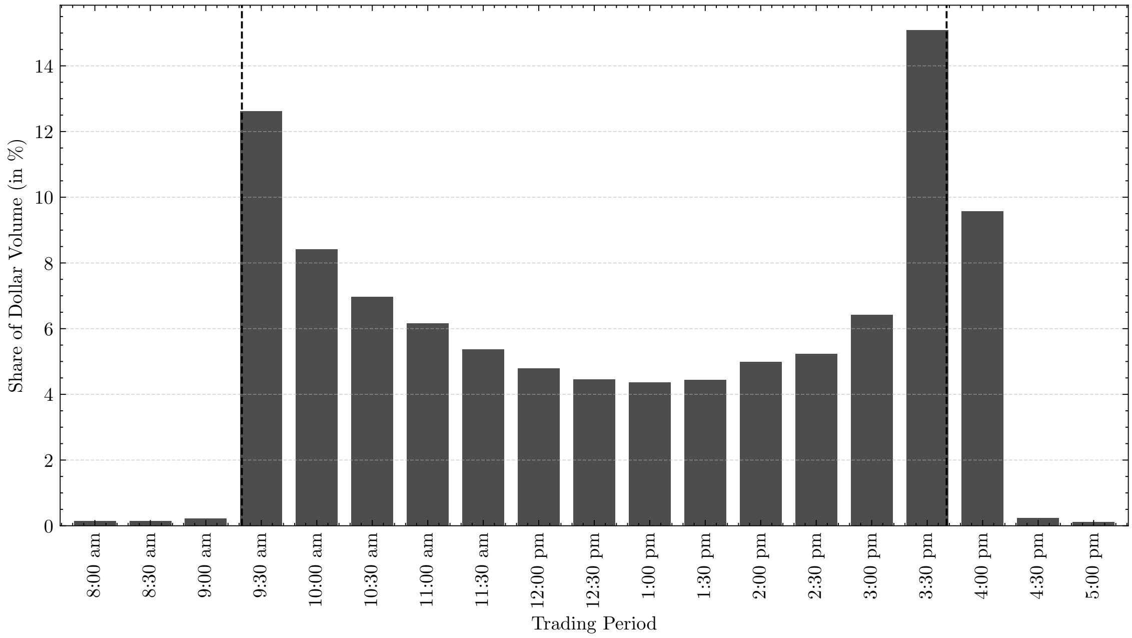

Figure 15: Share of dollar volume per half-hour.
This figure illustrates the dollar volume per half-hour trading period as a share of total daily volume. Only half-hour intervals containing at least 0.1% of the total daily volume are shown. The labels on the horizontal axis give the opening time of the respective half-hour interval, e.g., the first bar on the left represents the period between 8:00am and 8:30am (EST). The two vertical dashed lines indicate the opening and closing times of the U.S. stock market.

close. However, as we show in one of our robustness checks in Section 5.2.2, including the following weekend, i.e., measuring the latter return from Monday close to the next Monday open, even increases the effect.

Hence, for our main analysis, we define the week's remaining return $\tau_{i,w}$ for stock i in week w as the return between the Monday close and the week's close on Friday:

$$\tau_{i,w} = \frac{c_{i,w,65}}{c_{i,w,13}} - 1, \tag{17}$$

where $c_{i,w,65}$ is the 65th half-hour intraday close of week w, i.e., the price of stock i at Friday market close at 04:00pm in week w.

Table 16 presents summary statistics on the number of observations, market capitalizations, and Monday afternoon returns, as well as the week's remaining returns for each year. Clearly, data coverage has vastly improved since 2010. The number of available firms and, hence, observations, increases over time.[15] The data covers periods with negative and positive mean returns and periods of high and low

[15]This also holds true for the year 2022; the relatively small sample size is caused by the data set expiring in October 2022.

	Firms	N	Mean Market Cap.	Median Market Cap.	Mean $r_{i,w}$	Median $r_{i,w}$	Mean $r^2_{i,w}$	Mean $\tau_{i,w}$	Median $\tau_{i,w}$	Mean $\tau^2_{i,w}$
2007	147	6,641	53,715	27,979	-0.0896%	-0.0591%	0.0059%	0.1534%	0.2590%	0.1151%
2008	160	7,077	51,044	26,672	-0.1427%	-0.0602%	0.0290%	-0.4055%	-0.3911%	0.3861%
2009	365	15,745	27,455	12,346	0.0373%	0.0682%	0.0184%	0.8248%	0.2815%	0.3608%
2010	702	31,232	14,787	4,322	-0.0300%	0.0000%	0.0080%	0.1162%	0.1646%	0.1838%
2011	717	31,868	17,831	6,003	-0.0017%	0.0415%	0.0115%	0.4652%	0.3949%	0.2259%
2012	731	32,543	18,118	5,824	0.1000%	0.0921%	0.0045%	0.4636%	0.3095%	0.1236%
2013	804	37,294	18,537	6,266	-0.0346%	-0.0108%	0.0047%	0.5407%	0.5674%	0.0976%
2014	865	39,249	21,130	7,320	-0.0079%	0.0270%	0.0048%	0.3648%	0.3027%	0.1082%
2015	934	40,705	20,751	7,165	-0.0045%	0.0194%	0.0071%	-0.0834%	-0.0580%	0.1512%
2016	1013	44,988	19,235	6,313	-0.0347%	-0.0282%	0.0059%	0.3367%	0.2068%	0.1932%
2017	1056	45,809	22,228	7,531	0.0220%	0.0209%	0.0034%	0.2947%	0.3022%	0.0999%
2018	1143	52,615	23,695	7,748	-0.0969%	-0.0246%	0.0072%	-0.1122%	0.0720%	0.1617%
2019	1203	56,987	22,482	6,966	-0.0089%	-0.0060%	0.0072%	0.5185%	0.5723%	0.1415%
2020	1299	56,022	26,020	7,680	-0.0843%	0.0241%	0.0192%	-0.0040%	0.0410%	0.3848%
2021	1405	63,077	28,420	7,856	-0.0061%	0.0207%	0.0090%	0.3406%	0.3420%	0.1772%
2022	1578	54,114	30,540	8,108	0.1650%	0.0760%	0.0165%	-0.6139%	-0.5482%	0.3141%

Table 16: Summary statistics.
This table presents yearly summary statistics for the utilized data set. The columns represent, from left to right: number of firms, number of observations, mean and median market capitalizations (in million U.S. Dollar), mean and median of the Monday afternoon return $r_{i,w}$ (defined in Equation (16)), Monday afternoon return volatility, mean and median of the rest-of-week return $\tau_{i,w}$ (defined in Equation (17)), and the rest-of-week return volatility, with volatility being the corresponding squared returns. Note that the data is incomplete for 2022, as our sample ends in October 2022.

volatility, measured by the mean squared return. Moreover, the years 2008 and 2020 exhibit the highest levels of volatility, with similar volatility levels for the rest of the week returns. Both mean and median market capitalizations decreased until 2010 and increased again thereafter.

Among others, French (1980) shows that Monday returns for major stock market indices are abnormally low or even negative. However, subsequent studies report that the weekend effect is at most present at Monday openings (Harris, 1986; Rogalski, 1984; Smirlock and Starks, 1986). Doyle and Chen (2009) even suggest that the effect is not necessarily attached to Mondays but rather show that a similar effect can be observed for other weekdays using a sample from 1993 to 2007. Although these results suggest that Monday afternoon returns are unaffected by the weekend, we nevertheless test for negative mean Monday afternoon returns. The mean Monday afternoon returns reported in Table 16 are indeed slightly negative, but statistically insignificant for the full sample period (with a t-statistic of -0.2 based on two-way clustered standard errors) and within each year. Hence, in accordance with the extensive intraday analysis of stock returns in Harris (1986), there is also no indication in our more recent sample of particularly negative returns on Monday afternoons.

5.2.2 The Monday afternoon reversal

We begin our main empirical analysis by sorting the remaining weekly returns based on Monday afternoon return deciles and market capitalization deciles. We perform a double decile sort and include the firm size to ensure that the resulting remaining week's return spread is not entirely driven by small firms. Market capitalization is defined as the firm's average market capitalization of the previous year, and it is updated annually to eliminate any look-ahead bias. Table 17 presents the mean returns for each pair of deciles, with 1 (10) corresponding to the decile with the lowest (highest) Monday afternoon return and market capitalization.

Firstly, we observe that the return of the remaining week cannot be ordered linearly by firm size. More precisely, the medium-sized deciles exhibit the highest mean returns, while the smallest deciles perform the worst. Notably, as we do not report complete consecutive returns for the whole week, this does not necessarily conflict with the broader size premium revealed in Banz (1981). Secondly, the ranking of returns based on market capitalization seems to invert when sorting additionally on the Monday afternoon return. In particular, in the lowest Monday afternoon return decile, medium- to large-sized firms gain higher average returns for the remainder of the week than small-sized firms. In fact, the overall highest average week's remaining return is achieved by stocks in the 8th market capitalization decile

		Monday afternoon return												
	Deciles	1	2	3	4	5	6	7	8	9	10	Mean	Spread (1-10)	t
Market Cap	1	0.391%	0.142%	0.217%	0.192%	0.154%	0.100%	0.208%	0.153%	0.108%	-0.028%	0.164%	0.419%	2.57
	2	0.297%	0.211%	0.095%	0.134%	0.127%	0.194%	0.170%	0.271%	0.108%	0.093%	0.170%	0.204%	1.83
	3	0.346%	0.263%	0.267%	0.301%	0.189%	0.171%	0.100%	0.203%	0.129%	0.066%	0.203%	0.280%	2.65
	4	0.246%	0.246%	0.218%	0.293%	0.285%	0.145%	0.165%	0.195%	0.085%	0.056%	0.193%	0.190%	1.52
	5	0.455%	0.249%	0.234%	0.271%	0.226%	0.073%	0.202%	0.255%	0.191%	-0.013%	0.214%	0.467%	3.57
	6	0.436%	0.371%	0.238%	0.223%	0.240%	0.031%	0.120%	0.188%	0.132%	0.008%	0.199%	0.428%	3.51
	7	0.339%	0.256%	0.221%	0.207%	0.122%	0.155%	0.282%	0.214%	0.062%	-0.003%	0.186%	0.342%	2.53
	8	0.669%	0.256%	0.221%	0.132%	0.301%	0.188%	0.091%	0.106%	0.094%	-0.272%	0.179%	0.941%	2.50
	9	0.504%	0.206%	0.259%	0.257%	0.244%	0.209%	0.113%	0.101%	0.076%	-0.029%	0.194%	0.534%	3.50
	10	0.487%	0.336%	0.288%	0.229%	0.172%	0.170%	0.074%	0.048%	0.070%	-0.086%	0.179%	0.572%	3.09
	Mean	0.417%	0.254%	0.226%	0.224%	0.206%	0.144%	0.152%	0.173%	0.106%	-0.021%	0.188%	0.438%	2.73

Table 17: Rest-of-week returns from a double decile sort on the Monday afternoon return and market capitalization.
This table presents the average rest-of-week returns for decile pairs based on the Monday afternoon return and the market capitalization, along with the lowest-minus-highest spreads and t-statistics based on two-way clustered standard errors for each market capitalization decile. The Monday afternoon return is defined as the intraday return between the closing price of the half-hour trading period from 12:30pm to 1:00pm and market close at 4:00pm. The market capitalization is defined as the mean market value of the firm's equity over the preceding year. The portfolio is evaluated once per week. The reported rest-of-week returns are the returns between the close of the market on Monday at 4:00pm and the close of the market on Friday at 4:00pm.

and the lowest Monday afternoon return decile. In contrast, for stocks in the highest Monday afternoon return decile, medium- to small-sized firms earn generally larger returns than larger firms.

One may conjecture that the observed Monday afternoon reversal is attributable to small stocks, as it is typically challenging to arbitrage price inefficiencies in small stocks due to the potential impact of large orders as well as high transaction costs and lower liquidity compared to large stocks. Hence, markets may permit predictive asset pricing factors that yield positive long-short returns without compromising market efficiency, as these strategies are reliant on the trader's infrastructure and may only be exploitable to a limited degree. Thus, we may expect a greater spread for small firms compared to large firms. But we observe the opposite, i.e., the reversal's long-short spread grows as the firm's size increases. In fact, the lowest-minus-highest Monday afternoon spread is as high as 0.572% for the largest firm size decile, whereas for the smallest firm size decile, it is only 0.419%. Consequently, the Monday afternoon reversal is not primarily attributable to small stocks.

We also consider a one-way decile sort based on the Monday afternoon return and find that the average remaining week return spread of the bottom-minus-top afternoon return decile portfolio is 0.406% and significantly different from zero with a t-statistic of 5.1 based on two-way clustered standard errors. The spread is driven primarily by the long component, which earns a mean remaining week return of 0.400%, compared to the short component, which earns a mean remaining week return of 0.006%. The results are even stronger when we include the weekend return in the remaining week's return, i.e., measure the remaining week's return from Monday close to the next Monday's open, for which we observe a spread of 0.465% (0.474% - 0.009%) with a t-statistic of 6.2.[16]

As the week's structure may produce the Monday afternoon reversal, it is plausible to suppose that the impact is likewise influenced by the year's end. As documented in Reinganum (1983) and Rozeff and Kinney (1976), potential tax savings motivate closing lost holdings towards the end of the year and reinvesting in January, making the latter month distinctive in terms of the cross-sectional return distribution. Therefore, we test whether the reversal on Monday afternoon returns is affected by the speciality of January. However, we find no statistically significant difference in the magnitude of the reversal for Mondays in January than for the rest of the year, suggesting that it is not due to the turn of the year.

[16]The corresponding double decile sort on the Monday afternoon return and the market capitalization are very similar to the results we present in Table 17 and are omitted for brevity.

Closely related, we also investigate whether there are differences in the Monday afternoon reversal depending on the week of the month. Wang et al. (1997) report that the weekend effect occurs particularly during the fourth and fifth weeks of each month, while the weekend effect is insignificant for the remaining weeks. Thus, we undertake a similar analysis for the Monday afternoon reversal. We observe that the reversal is not specific to particular Mondays during a month. More precisely, the reversal is neither increasing nor decreasing when considering the fourth or fifth Monday in a month. Besides, we also find no statistical evidence for the opposite, i.e., that the first, second, or third Monday in the month has a more pronounced reversal effect.

5.2.3 Predictive regression analysis

We observed that the long-short remaining-week-return spread between low and high Monday afternoon return decile stocks is positive for every market capitalization decile and significant for most of them, while being higher for larger firms. To analyze these findings in more detail, we regress the stocks' remaining weekly return on the Monday afternoon return for each market capitalization decile:

$$\tau_{i,w} = \alpha + \beta\, r_{i,w} + \varepsilon_{i,w}, \tag{18}$$

and report the parameter estimates along with their corresponding t-statistics and R^2 in Panel A of Table 18.

The results support the Monday afternoon reversal. The predictive ability of the Monday afternoon return is observable across both small and large firms and, in fact, increases for larger firms. In the largest firm decile, the R^2 is 1.74%, whereas the R^2 for the smallest firm decile is 0.94%, a substantial difference. In addition, each estimate of β is significantly negative and, as expected, becomes more pronounced with increasing firm size, suggesting that the predictive power of the Monday afternoon return on the subsequent remaining week's return is not primarily due to small stocks.

However, although we motivated in the introduction why we focus on the Monday afternoon in our analysis, the choice may still seem arbitary, therefore, we also report results when including the Monday morning return, i.e., from 9:30am to 12:30pm, as a regressor. The Monday morning return is defined as:

$$r_{i,w} = \frac{c_{i,w,7}}{o_{i,w,1}} - 1, \tag{19}$$

where $c_{i,w,7}$ is the 7th half-hour intraday close of week w, i.e., the price of stock i at

12:30pm, and $o_{i,w,1}$ is the opening price of the first half-hour of the week, i.e., the market open price on Monday morning of stock i in week w.

The results are reported in Panel B of Table 18. None of the estimated coefficients of the Monday morning return is statistically significantly different from zero, while all of the coefficients of the Monday afternoon return remain at the same level and are still highly significant. Thus, we find strong evidence for a Monday afternoon reversal, but no indication of any reversal or momentum effect on the Monday morning return.

Heston et al. (2010) suggest that fund managers follow daily routines, which may involve rebalancing fund portfolios at specific times of the day. It is possible that the Monday afternoon reversal is caused by the daily routine of fund managers, such as a daily rebalancing of the portfolio during the afternoon. If this periodic behavior of fund managers causes the afternoon return reversal, it should be observable on other weekdays as well, and we may expect that afternoon returns on other weekdays are also able to predict the remaining week's return. To test for this possible explanation, we estimate modified versions of Equation (18) by separately regressing for each weekday's afternoon return the corresponding remaining return of the week on the respective weekday's afternoon return (e.g., for Tuesday, we regress the weekly return from Tuesday close to the following Friday close on the Tuesday afternoon return). For brevity, we do not split the sample into market capitalization deciles but instead fit the model to the entire data set.[17] The resulting slope coefficients are much smaller (ranging from -0.0937 to -0.0391) and statistically insignificant, indicating that no other weekday's afternoon return meaningfully predicts the remaining week's return. Therefore, the Monday afternoon reversal is unique to the first day of the week and unlikely to be caused by fund manager's daily routines.

5.2.4 Out-of-sample analysis

As the Monday afternoon return is a highly significant predictor for short-term future returns in our in-sample tests, we aim to assess its economic gains in an out-of-sample test. To this end, each week we construct separate value-weighted portfolios of the bottom and top Monday afternoon return deciles and compute the resulting weekly cumulated logarithmic returns. Figure 16 depicts the resulting returns of the two portfolios. The findings reveal the economic significance of the observed effect, as the bottom decile portfolio earns a value-weighted mean weekly return of 0.49% (with a standard error of 0.12%) or 23.67% p.a., whereas the mean return of the top decile

[17]Splitting the sample into market capitalization deciles yields results that are qualitatively similar to those based on the entire dataset.

Panel A: Excluding morning return

Decile	α		$r_{i,w}$		R^2
1	0.0016	(1.54)	-0.4010	(-5.75)	0.94%
2	0.0017	(1.66)	-0.3705	(-4.37)	0.65%
3	0.0020	(1.92)	-0.4294	(-4.63)	0.84%
4	0.0019	(1.86)	-0.4306	(-4.63)	0.84%
5	0.0021	(2.09)	-0.4815	(-4.85)	1.01%
6	0.0020	(1.98)	-0.4865	(-4.65)	1.03%
7	0.0018	(1.88)	-0.5177	(-4.76)	1.12%
8	0.0017	(1.88)	-0.5208	(-6.23)	1.73%
9	0.0019	(2.03)	-0.5397	(-3.92)	1.32%
10	0.0017	(1.96)	-0.6021	(-3.98)	1.74%

Panel B: Including morning return

Decile	α		$r_{i,w}$		$r_{i,w}^{morning}$		R^2
1	0.0022	(2.06)	-0.4192	(-4.51)	0.0217	(0.55)	0.86%
2	0.0018	(1.75)	-0.3758	(-4.46)	-0.0455	(-1.21)	0.71%
3	0.0019	(1.86)	-0.4359	(-4.71)	-0.0184	(-0.46)	0.89%
4	0.0019	(1.88)	-0.4370	(-4.68)	-0.0377	(-1.03)	0.90%
5	0.0021	(2.11)	-0.4924	(-5.01)	-0.0270	(-0.73)	1.09%
6	0.0018	(1.84)	-0.4941	(-4.57)	-0.0095	(-0.25)	1.08%
7	0.0019	(1.95)	-0.5269	(-4.86)	-0.0470	(-1.60)	1.25%
8	0.0017	(1.87)	-0.5347	(-4.42)	-0.0669	(-1.62)	1.37%
9	0.0019	(2.04)	-0.5458	(-3.93)	-0.0369	(-0.70)	1.36%
10	0.0017	(1.96)	-0.6046	(-4.04)	-0.0728	(-1.28)	1.78%

Table 18: Estimation results for regressions of the rest-of-week return on the Monday afternoon return.
This table reports for each firm size decile the results of the regression model (18) where the stock's remaining weekly returns are regressed on the Monday afternoon returns $r_{i,w}$ (Panel A) and additionally on the Monday morning returns $r_{i,w}^{morning}$ (Panel B). Reported are coefficient estimates, corresponding t-statistics (in parenthesis) with two-way clustered standard errors as well as the coefficients of determination R^2.

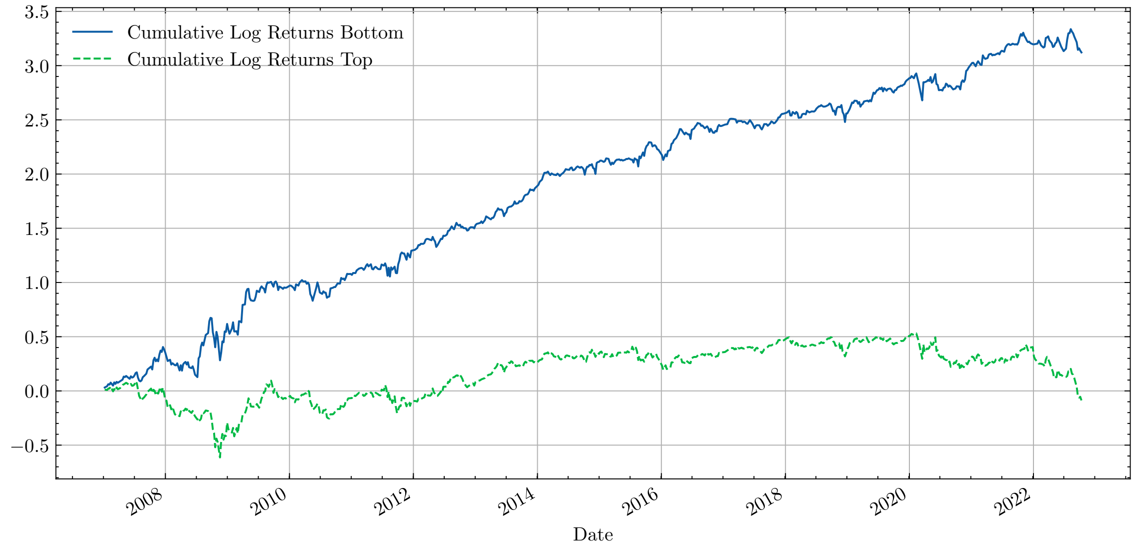

Figure 16: Value-weighted out-of-sample cumulative logarithmic returns of the bottom and top Monday afternoon deciles.
This figure depicts the cumulative logarithmic returns of the bottom and top Monday afternoon return decile portfolios. The portfolios are formed weekly, and the resulting logarithmic returns are value-weighted and cumulated weekly.

portfolio is not significantly different from zero. Hence, investors who purchase the value-weighted portfolio of stocks in the bottom decile of the Monday afternoon return receive a Sharpe ratio of 1.06 on an annualized basis.

Moreover, the investment strategy is remarkably resilient to the 2008 and 2009 financial crises, as well as to the market decline in 2021 and 2022, whereas the value of a portfolio of stocks from the highest Monday afternoon return decile decreases substantially during these time periods. That is, a long-short portfolio that buys the bottom decile and sells the top decile yields an even higher Sharpe ratio of 1.40. This increase in the Sharpe ratio is attributable to a reduction in risk, as the average weekly return of 0.44% per week or 22.18% annually, is slightly lower than that of the long-only portfolio containing stocks from the bottom decile.

It is important to note that we did not consider transaction costs and other market frictions in the latter analysis. Quantifying the economic benefits of the Monday afternoon reversal while including transaction costs is challenging as the investor's particular infrastructure and contractual arrangements with the broker have a non-negligible impact on the effective transaction costs incurring during trading. Nevertheless, to proxy these costs, we compute the average bid-ask-spread of all stocks in our dataset. In general, the mean bid-ask spread is 0.06%, however, it is significantly reduced for stocks in the highest market capitalization decile, having a mean of 0.03%. Thus, as the previously presented long-short returns are value-weighted, the estimated weekly return after accounting for the bid-ask spread of the long-short portfolio based on Monday afternoon return deciles falls within the range of 0.38% and 0.41%, i.e., the trading strategy remains profitable.

5.2.5 Overnight versus intraday components

Lou et al. (2019) demonstrate that overnight and intraday returns exhibit strong return persistence, while returns reverse from the overnight to intraday session and vice versa. In addition, the authors analyze the returns of trading strategies based on common asset pricing factors and conclude that these strategies yield returns either overnight or intraday. Therefore, we investigate whether the bottom-minus-top Monday afternoon return decile spread occurs exclusively overnight or intraday, i.e., whether it is due to the overnight or intraday component.

To this end, we compute the weekly Monday afternoon return deciles as before and decompose the remaining week's returns into their four overnight and four intraday components, i.e., one for each weekday. For each of these components, we compute the respective bottom-minus-top return spreads. Table 19 reports the bottom and top decile mean returns, the return spread, and t-statistics based on two-way clustered standard errors for each time component. The overnight return is defined as the return between the previous day's market close and the respective day's market open. Interestingly, the spread between the bottom and the top deciles exists during both the overnight and intraday trading sessions. The overnight component of the spread is significantly positive on Tuesdays and Fridays, while the intraday component is significantly positive on Tuesdays, Wednesdays, and Thursdays. The Friday intraday session is the only one with a negative spread, albeit of modest magnitude and statistically insignificant. Thus, the Monday afternoon reversal occurs over both the overnight and intraday sessions.

5.2.6 Relationship to systematic risk and firm news

The previous results suggest that a systematic risk factor may be responsible for the high long-short returns caused by the predictive power of the Monday afternoon return. Therefore, in order to determine the relationship between the returns of the Monday afternoon reversal and common systematic risk, we regress the long-short excess returns from the Monday afternoon return sort on the five factors of Fama and French (2015) and the momentum factor of Carhart (1997). The regression provides both an assessment of whether the Monday afternoon reversal is explained by systematic risk, as well as an estimate of the alpha of the long-short portfolio returns based on Monday afternoon returns.

	Overnight			Intraday		
	Bottom	Top	Spread	Bottom	Top	Spread
Tuesday	0.207%	0.082%	0.125%	0.046%	0.000%	0.046%
	(4.72)	(1.96)	(5.47)	(0.87)	(0.01)	(1.35)
Wednesday	0.120%	0.062%	0.058%	0.024%	-0.015%	0.038%
	(3.17)	(1.74)	(3.03)	(0.44)	(-0.29)	(1.13)
Thursday	0.048%	-0.008%	0.055%	0.077%	-0.016%	0.093%
	(1.19)	(-0.20)	(2.83)	(1.39)	(-0.31)	(2.50)
Friday	0.112%	0.069%	0.043%	0.001%	0.017%	-0.016%
	(2.71)	(1.77)	(2.17)	(0.02)	(0.35)	(-0.47)

Table 19: Overnight and intraday components of the rest-of-week returns of the bottom and top Monday afternoon return decile portfolios.
The reported returns are based on equally-weighted portfolios that buy stocks in the bottom or top Monday afternoon return deciles, respectively, and that are evaluated weekly. For each portfolio and week, the remaining weekly return (Monday close to Friday close) is subdivided for each weekday and split into the overnight and intraday components. The overnight return is the return from the previous day's close to the displayed day's open; e.g., the overnight return in row "Tuesday" measures the return between the Monday close and Tuesday open. Spread is the simple difference between the bottom and top returns. The t-statistics (in parentheses) are based on two-way clustered standard errors.

The results are presented in Table 20. As none of the slope coefficients is statistically significantly different from zero, the returns of the bottom-minus-top Monday afternoon returns appear to be unrelated to the considered risk factors. Additionally, the strategy produces a significant alpha of 0.43% per week (with a t-statistic of 5.94 based on weekly clustered standard errors). Interestingly, the long leg of the strategy accounts for the majority of the long-short portfolio returns. When separating the returns into their long (low Monday afternoon returns) and short (high Monday afternoon returns) legs, we observe a significant and positive alpha of 0.45% per week for the long-leg, but no significant alpha for the short leg.

Jiang et al. (2021) find that returns tend to drift in the direction of the initial news response, which may explain a potential mispricing on Mondays that is ultimately reversed throughout the week. In order to test whether the Monday afternoon reversal is related to firm news, we estimate Equation (18) with two additional regressors: the Monday afternoon returns, first, conditional on whether firm earnings announcements occurred during the previous week and, second, conditional on whether they occurred during the week in which we measure the Monday afternoon return. While the unconditional Monday afternoon return remains highly significantly negative (with a

	α	Mkt-RF	SMB	HML	RMW	CMA	MOM	R^2
B-T	0.0043	-0.0533	0.0832	-0.0208	0.1816	0.0579	0.0006	2.532%
	(5.94)	(-1.22)	(0.98)	(-0.23)	(1.52)	(0.40)	(1.09)	
B	0.0045	-0.1724	0.2807	-0.1332	0.3316	-0.1291	-0.0004	3.148%
	(3.66)	(-1.76)	(1.90)	(-0.82)	(1.76)	(-0.50)	(-0.44)	
T	0.0001	-0.1187	0.1981	-0.1129	0.1500	-0.1848	-0.0010	1.726%
	(0.11)	(-1.46)	(1.60)	(-0.79)	(0.95)	(-0.79)	(-1.10)	

Table 20: Estimation results for Fama and French (2015) and Carhart (1997) 6-factor analysis. This table presents estimation results of a regression of the weekly excess return of bottom-minus-top Monday afternoon return portfolios on excess returns of a market portfolio (Mkt-RF), returns of a size portfolio (SMB), value portfolio (HML), a portfolio on operating profitability (RMW), a portfolio on investment strategies (CMA) (Fama and French, 2015) and a portfolio on momentum (Carhart, 1997). Weekly excess returns are measured as the returns from Monday close to Friday close minus the risk-free rate. The t-statistics (in parenthesis) are based on weekly clustered standard errors. Dependent variables are B-T (bottom-minus-top), which are returns from a long-short portfolio on Monday afternoon return deciles, while B and T are the respective long-only portfolios of the bottom and top Monday afternoon return deciles, respectively.

t-statistic of -4.9 based on two-way clustered standard errors), none of the conditional variants are significantly different from zero. Hence, the reversal is not a result of firm earnings announcements durinMonday's surrounding days.

5.3 U.S. stock market indices and international ETFs

Given that the Monday afternoon return is a significant predictor of the remaining week's return across the cross-section of stock returns, we examine whether this relationship also holds true for broad U.S. stock market indices and international equity ETFs. To this end, we analyze the returns of five indices and thirteen international ETFs from January 2004 to October 2022. The selection of ETFs represents major international equity markets and includes specific ETFs for the U.S., Europe, China, Japan, Germany, and the United Kingdom. Specifically, we cover the U.S. indices S&P 500, Dow Jones Industrial Average, Nasdaq Composite, Nasdaq 100, and Russell 3000, as well as the following international ETFs: SPDR S&P 500 ETF Trust, Vanguard Total Stock Market ETF, Invesco QQQ Trust, Vanguard Value ETF, iShares Russell 2000 ETF, Vanguard Total World Stock ETF, Vanguard FTSE Developed Markets ETF, Vanguard FTSE Europe ETF, iShares MSCI Eurozone ETF, SPDR Euro Stoxx 50 ETF, iShares MSCI Germany ETF, iShares MSCI Japan ETF, iShares MSCI United Kingdom ETF, and iShares Trust - China Large-Cap ETF. Note that all ETFs within the selection are traded on U.S. exchanges.

Analyzing the Monday afternoon reversal based on decile sorts is not a reasonable approach due to the relatively small number of indices and ETFs compared to the number of stocks examined in the preceding sections. Therefore, we adhere to the regression analysis based on Equation (18) for each index and ETF. We may anticipate comparable outcomes for indices and ETFs that are typically highly correlated with the broad U.S. stock market. Moreover, if the effect is not unique to the U.S. market, we should observe similar results for other international markets as well.

Table 21 presents the parameter estimates along with corresponding robust t-statistics and values of R^2. First, focusing on the U.S. stock market indices (Panel A) shows clearly that the Monday afternoon return is a highly significant and negative predictor of the week's remaining return for all U.S. indices. Moreover, the estimated coefficients are large in absolute terms; for example, a decrease of 1 percentage point in the S&P500's return on Monday afternoon increases the expected week's remaining return by 0.65 percentage points, all else being equal. Furthermore, with R^2 values as high as 3.07%, the Monday afternoon return adequately explains the variability in the remaining week's return. We subsequently estimate the regression model (18) for all indices jointly, yielding a highly significant and negative estimate of $\beta = -0.63$ with a two-way clustered t-statistic of -4.0, which is comparable to the results for the cross-section of stock returns.

Based on the results for the cross-section of U.S. stocks and for the U.S. stock market indices, we anticipate a significant Monday afternoon reversal among the U.S.-related ETFs. This is obviously supported in Panel B of Table 21, which presents the estimation results for the considered ETFs. However, even more remarkable is the fact that the Monday afternoon reversal is also observable for ETFs that track major international stock markets. In addition, the magnitude of the reversal is remarkably similar across all markets and close to the results observed for the U.S. markets and the cross-section of U.S. stock returns. This implies that the Monday afternoon reversal is not unique to U.S. stocks and that the latter are not causing the effect. We further estimate the regression again jointly for all ETFs, yielding a significantly negative estimate of -0.62 with a t-statistic of -3.7 based on two-way clustered standard errors for the coefficient of the Monday afternoon return. Hence, the magnitude of the Monday afternoon reversal is very similar across international equity markets.

Panel A: U.S. stock market indices

	Ticker	α	t_α	β	t_β	R^2	N
	SPX	0.0013	(1.66)	-0.6473	(-3.55)	2.67%	663
	DJI	0.0015	(1.93)	-0.6367	(-3.08)	2.55%	669
	COMP	0.0019	(2.17)	-0.6687	(-3.45)	2.56%	669
	NDX	0.0021	(2.35)	-0.5162	(-2.53)	1.65%	668
	RUA	0.0014	(1.82)	-0.7023	(-3.67)	3.07%	669

Panel B: ETFs

Market	Ticker	α	t_α	β	t_β	R^2	N
U.S.	SPY	0.0018	(2.55)	-0.5457	(-2.78)	1.68%	783
	VTI	0.0019	(2.61)	-0.5574	(-2.73)	1.65%	766
	QQQ	0.0019	(2.33)	-0.5030	(-2.55)	1.37%	770
	VTV	0.0013	(1.79)	-0.5780	(-2.89)	1.75%	755
	IWM	0.0016	(1.80)	-0.4865	(-2.46)	1.13%	756
Global	VT	0.0017	(1.96)	-0.6476	(-2.38)	1.85%	611
Global excl. U.S.	VEA	0.0012	(1.36)	-0.6924	(-2.16)	1.52%	659
Europe	VGK	0.0019	(2.20)	-0.6430	(-2.27)	1.33%	738
	EZU	0.0016	(1.79)	-0.6533	(-2.62)	1.26%	772
	FEZ	0.0013	(1.47)	-0.8735	(-3.47)	2.41%	712
Germany	EWG	0.0017	(1.91)	-0.6583	(-2.68)	1.15%	790
Japan	EWJ	0.0010	(1.25)	-0.6840	(-2.50)	1.41%	775
UK	EWU	0.0019	(2.42)	-0.6347	(-2.59)	1.30%	798
China	FXI	0.0008	(0.79)	-0.6572	(-2.45)	1.09%	752

Table 21: Estimation results for the regression of the weekly remaining return on the Monday afternoon return for U.S. stock market indices and international ETFs.
Reported are parameter estimates along with corresponding t-statistics (in parenthesis) based on weekly clustered standard errors as well as the R^2 and the available number of observations N. The parameters are estimated for each index and ETF separately. The indices are in row-wise order: S&P500, Dow Jones Industrial Average, Nasdaq Composite, Nasdaq 100, and Russell 3000. The ETFs are in row-wise order: SPDR S&P 500 ETF Trust, Vanguard Total Stock Market ETF, Invesco QQQ Trust, Vanguard Value ETF, iShares Russell 2000 ETF, Vanguard Total World Stock ETF, Vanguard FTSE Developed Markets ETF, Vanguard FTSE Europe ETF, iShares MSCI Eurozone ETF, SPDR Euro Stoxx 50 ETF, iShares MSCI Germany ETF, iShares MSCI Japan ETF, iShares MSCI United Kingdom ETF, and iShares Trust - China Large-Cap ETF.

5.3.1 Investor disagreement and informed trading

Thus far, we have examined several obvious explanations for the Monday afternoon reversal, but none of them have been empirically supported by the data, necessitating further investigation. Recall our findings that the other weekdays' returns have no predictive power for the remaining week's return. Consequently, the Monday is unique, which may be attributable to the weekend. In fact, it is frequently argued that individuals have significantly more time to process information during the weekend than in between weekdays when the stock market is open, which may even apply to institutional investors (Lakonishok and Maberly, 1990). The processing of accumulated information over the weekend, without receiving feedback from markets in the form of returns, may cause investors to hold divergent beliefs regarding expected returns. This, in turn, potentially leads to Monday afternoon returns that are subsequently reversed in the following days, as we discuss below.

Several studies investigate the relationship between heterogeneous beliefs and disagreement regarding expected returns and find a positive relationship between disagreement and stock volume and volatility (Carlin et al., 2014; Cookson and Niessner, 2020). Therefore, we assess the effects of trading volume and volatility during Monday afternoon on the reversal. Hence, if the reversal is more pronounced on Monday afternoons with higher volume or volatility, it indicates that the reversal may be due to disagreement. To this end, we compute the Monday afternoon trading volume as the sum of shares traded times the closing price of the respective half-hour returns. Similarly, we define the Monday afternoon volatility as the sum of squared half-hourly returns for each stock. Afterwards, we compute weekly deciles for each measure.

Tables 22 and 23 report mean remaining week's returns for double decile sorts on the Monday afternoon return and on the Monday afternoon volume and volatility, respectively. Obviously, the bottom-minus-top Monday afternoon return spread is 2.6 times higher in the highest volume decile (0.636%) than in the lowest (0.243%), supporting the hypothesis that disagreement about expected returns drives the reversal. Furthermore, the reversal is only marginally significant at the 10%-level in the smallest volume decile with a t-statistic of 1.78, while it is highly significant in the largest volume decile with a t-statistic of 4.19.

The results are even more pronounced when we consider the Monday afternoon volatility. The bottom-minus-top Monday afternoon return spread is as high as 1.021% for stocks in the highest volatility decile with a t-statistic of 5.5. Although this is the strongest reversal we have observed so far, it is perhaps even more interesting that below-median Monday afternoon volatility stocks do not reverse

		Afternoon return												
	Deciles	1	2	3	4	5	6	7	8	9	10	Mean	Spread (1-10)	t
Volume	1	0.341%	0.242%	0.224%	0.154%	0.212%	0.121%	0.103%	0.164%	0.021%	0.098%	0.168%	0.243%	1.78
	2	0.248%	0.232%	0.183%	0.227%	0.150%	0.068%	0.162%	0.178%	0.067%	-0.007%	0.151%	0.255%	1.86
	3	0.371%	0.256%	0.232%	0.254%	0.231%	0.117%	0.166%	0.188%	0.063%	0.003%	0.188%	0.368%	2.82
	4	0.232%	0.251%	0.203%	0.204%	0.219%	0.194%	0.041%	0.225%	0.200%	0.091%	0.186%	0.141%	1.10
	5	0.372%	0.280%	0.142%	0.150%	0.158%	0.152%	0.221%	0.140%	0.141%	-0.056%	0.170%	0.428%	3.84
	6	0.389%	0.261%	0.259%	0.238%	0.243%	0.149%	0.223%	0.189%	0.114%	-0.024%	0.204%	0.413%	3.42
	7	0.407%	0.208%	0.158%	0.212%	0.184%	0.082%	0.132%	0.207%	0.158%	-0.086%	0.166%	0.493%	3.96
	8	0.447%	0.209%	0.227%	0.265%	0.265%	0.143%	0.228%	0.211%	0.130%	0.110%	0.224%	0.337%	2.68
	9	0.492%	0.265%	0.299%	0.292%	0.236%	0.222%	0.107%	0.148%	0.087%	-0.001%	0.215%	0.493%	3.67
	10	0.539%	0.330%	0.311%	0.276%	0.207%	0.237%	0.136%	0.060%	0.048%	-0.097%	0.205%	0.636%	4.19
	Mean	0.384%	0.253%	0.224%	0.227%	0.211%	0.149%	0.152%	0.171%	0.103%	0.003%	0.188%	0.381%	2.93

Table 22: Double decile sort on the Monday afternoon return and volume.
This table shows the remaining weekly mean returns, bottom-minus-top spreads, and corresponding t-statistics (based on two-way clustered standard errors) resulting from a double decile sort on the Monday afternoon return and trading dollar volume. We define volume as the sum of the trading dollar volume of each half-hour return during the Monday afternoon.

		Afternoon return												
	Deciles	1	2	3	4	5	6	7	8	9	10	Mean	Spread (1-10)	t
Volatility	1	-0.112%	-0.011%	0.064%	0.097%	0.159%	0.105%	0.155%	0.244%	0.415%	0.848%	0.196%	-0.960%	-2.70
	2	-0.188%	-0.054%	0.061%	0.089%	0.159%	0.239%	0.231%	0.241%	0.360%	0.577%	0.171%	-0.766%	-2.57
	3	-0.010%	0.042%	0.132%	0.144%	0.226%	0.124%	0.177%	0.281%	0.266%	0.483%	0.186%	-0.493%	-2.14
	4	0.053%	-0.035%	0.055%	0.231%	0.216%	0.235%	0.120%	0.161%	0.192%	0.255%	0.148%	-0.202%	-0.98
	5	0.057%	0.168%	0.155%	0.306%	0.287%	0.238%	0.219%	0.164%	0.125%	0.098%	0.182%	-0.041%	-0.27
	6	0.130%	0.275%	0.288%	0.290%	0.256%	0.115%	0.193%	0.113%	0.153%	0.117%	0.193%	0.014%	0.11
	7	0.210%	0.333%	0.411%	0.389%	0.219%	0.100%	0.000%	0.080%	0.077%	0.090%	0.191%	0.120%	1.15
	8	0.207%	0.393%	0.501%	0.333%	0.251%	0.153%	0.010%	0.152%	-0.096%	0.084%	0.199%	0.123%	1.22
	9	0.449%	0.614%	0.470%	0.326%	0.244%	0.128%	0.094%	0.046%	-0.026%	-0.100%	0.224%	0.549%	4.44
	10	0.781%	0.566%	0.311%	0.308%	0.070%	-0.257%	0.297%	0.190%	-0.201%	-0.240%	0.183%	1.021%	5.50
	Mean	0.158%	0.229%	0.245%	0.251%	0.209%	0.118%	0.149%	0.167%	0.127%	0.221%	0.187%	-0.064%	0.38

Table 23: Double decile sort on the Monday afternoon return and volatility.
This table shows the remaining weekly mean returns, bottom-minus-top spreads, and corresponding t-statistics (based on two-way clustered standard errors) resulting from a double decile sort on the Monday afternoon return and volatility. Volatility is calculated as the sum of squared half-hour returns during the Monday afternoon.

during the remainder of the week, but rather follow their Monday afternoon return. Specifically, the bottom-minus-top Monday afternoon return spread is strongly negative at -0.960% with a t-statistic of -2.7, indicating strong momentum for low Monday afternoon volatility stocks. Thus, our findings of the Monday afternoon reversal are exclusive to stocks with high Monday afternoon return volatility.

The results indeed strongly suggest that disagreement among investors is responsible for the Monday afternoon reversal. If so, there may be a relatively informed group of traders anticipating a potential over- or undervaluation on Monday afternoons and the subsequent reversal throughout the rest of the week. Therefore, we expect a strong reversal during the rest of the week when returns strongly point in the opposite direction than predicted by informed traders during the Monday afternoon.

A well-studied group of informed traders are short sellers, who accurately anticipate price movements (Boehmer et al., 2008; Cohen et al., 2007; Dechow et al., 2001; Desai et al., 2002; Diether et al., 2009). To test whether short sellers correctly identify potentially over- or undervalued stocks on Monday afternoons, we utilize data about the number of shorted shares available from Compustat. The number of shorted shares per stock is released twice per month. Focusing on Mondays, we only keep observations for this day, totaling 38,450 observations. For our analysis, we use the days-to-cover as a relative measure of short positions in a given stock, which is defined as the number of shorted shares divided by the share volume on the given Monday. This metric is also well-suited for our analysis as it is closely related to short sellers' expected returns. In particular, short sellers generally anticipate high (low) future returns when a stock's days-to-cover is low (high).

As the number of shorted stocks is obtained only twice per month and as these releases do not necessarily coincide with a Monday, the sample size in the following analysis is much smaller than that of Section 5.2. Therefore, we focus on quintiles of the days-to-cover rather than deciles. If short sellers are well-informed and the Monday afternoon reversal is due to investors' disagreement about expected returns, then the reversal should be amplified when short sellers disagree with Monday afternoon prices. That is, we expect significantly higher remaining week's returns for stocks with low days-to-cover (i.e., relatively small amounts of shorted shares) and a low Monday afternoon return, compared to stocks with high days-to-cover (i.e., relatively large amounts of shorted shares) and a high Monday afternoon return.

The results presented in Table 24 confirm this expectation, which reports the remaining week's mean return from a double sort on the Monday afternoon return deciles and the days-to-cover quintiles. According to the results, the reversal is even amplified when short sellers disagree with price movement on the Monday

		Afternoon return									
	Deciles Quintiles	1	2	3	4	5	6	7	8	9	10
Days-to-cover	1	0.748%	-0.217%	-0.122%	0.007%	-0.141%	0.076%	-0.404%	-0.230%	-0.514%	-0.186%
	2	0.440%	0.024%	0.241%	0.066%	0.023%	-0.183%	-0.076%	-0.288%	-0.260%	-0.347%
	3	0.166%	0.269%	-0.074%	-0.013%	-0.120%	-0.217%	-0.124%	-0.039%	-0.435%	-0.596%
	4	0.260%	0.011%	-0.113%	-0.004%	-0.040%	-0.103%	-0.192%	0.057%	-0.466%	-0.991%
	5	-0.238%	-0.254%	0.039%	-0.003%	-0.054%	-0.042%	0.046%	-0.433%	-0.572%	-1.003%
	Mean	0.275%	-0.034%	-0.006%	0.011%	-0.066%	-0.094%	-0.150%	-0.187%	-0.449%	-0.624%

Table 24: Rest-of-week returns based on Monday afternoon return deciles and days-to-cover quintiles.
This table shows the remaining weekly mean returns resulting from a double sort on the Monday afternoon return deciles and days-to-cover quintiles. The days-to-cover are defined as the number of shorted shares divided by the stock's Monday volume.

afternoon return, whereas short sellers appear to anticipate the remaining week's returns on average correctly. The mean remaining week's return for stocks in the lowest days-to-cover quintile and lowest Monday afternoon return decile is 0.75%, while it is -1.00% for stocks in the highest days-to-cover quintile and highest Monday afternoon return decile. The difference of 1.75 percentage points is highly significant, with a t-statistic of 2.8 using two-way clustered standard errors. On the other hand, we find no reversal when short sellers agree with Monday afternoon returns. That is, we find no significant difference between the remaining week's return between stocks in the lowest days-to-cover quintile and highest Monday afternoon return decile and stocks in the highest days-to-cover quintile and lowest Monday afternoon return decile (with a t-statistic of -0.1).

Overall, the results indicate that disagreement about expected returns and ultimately fundamental values may be the underlying cause of our documented reversal. We approximate the disagreement in expected returns by the half-hourly volume and volatility on Monday afternoons and observe the reversal exclusively for stocks with high Monday afternoon volume or volatility. Likewise, we document significant momentum when volatility is low. Moreover, informed traders appear to effectively anticipate reversals in both directions.

5.4 Conclusion

This Chapter shows that the cross-section of returns exhibits a significant reversal on the Monday afternoon return during the remainder of the week. We analyze high-frequency data of 1,838 companies from 2007 to 2022 and find that a bottom-minus-top portfolio based on the Monday afternoon return yields a mean weekly value-weighted return of 0.44% with an annualized Sharpe ratio of 1.40. The majority of the portfolio's performance is attributable to its long position.

We also conduct numerous robustness checks in order to address potential explanations that might drive the observed Monday afternoon reversal. Foremost, we observe that the reversal is not driven by small stocks. We also find no evidence that the Monday afternoon reversal is due to fund managers' daily routines, such as daily rebalancing of the portfolio during the afternoon. In addition, both the overnight and intraday components of weekly returns contribute to the effect. The week's remaining return of the bottom and top Monday afternoon return deciles is not driven by any of the commonly considered risk factors of Fama and French (2015) and Carhart (1997), resulting in a significantly positive risk-adjusted alpha, and it is not due to earnings announcements. Moreover, the results are not attributable to a January, week-of-the-month, or weekend effect. Lastly, the Monday afternoon reversal is present in several U.S. stock market indices and numerous international ETFs.

Hence, the Monday afternoon reversal seems to be a prominent and very robust phenomenon that is not specific to U.S. stocks but can be observed among international equity ETFs. As such, we investigate whether the relatively long market closure on weekends and the subsequent potential disagreement about the effects of information on expected asset returns causes the reversal. In favor of this hypothesis, we find that the reversal only exists for stocks with high Monday afternoon volume or volatility. Furthermore, we show that a well-known group of informed traders, short sellers, correctly anticipates the Monday afternoon reversal in both directions.

6 Summary and conclusions

This dissertation examines several issues in contemporary asset pricing and portfolio management. Initially, we discuss the potential of deep learning for portfolio management and asset pricing. We develop a portfolio management approach based solely on DRL. DRL consists of a set of algorithms that seek to fit policies so that agents who follow these policies receive the greatest possible rewards. An environment, in which we strive to optimally choose a portfolio from a wide range of stocks, provides rewards. The agents receive rewards in the form of returns, and the goal is to maximize portfolio returns.

Our method solves several shortcomings of DRL for portfolio management. Given a large cross-section of n stocks from which one intends to construct a portfolio, the corresponding policy's neural network is typically required to have n output neurons. Furthermore, given p features, each state contains $p \cdot n$ inputs, necessitating a reasonably large sample size. However, a long and continuous, i.e., gap-free, time series of cross-sectional returns at a low frequency (e.g., daily returns) does not exist in finance. Therefore, we propose to sequentially set up trading environments for individual stocks instead of portfolios. To achieve a portfolio-focused policy, we reward investments in each stock with its respective return but no investments with the equally-weighted portfolio return of the remaining available stocks in the respective period.

The latter approach keeps the number of inputs as well as output neurons relatively small, which mitigates common overfitting in financial machine learning and scales well into larger portfolio sizes. We demonstrate these features in numerous out-of-sample tests in which the algorithm shows dominant performance against passive and active benchmark strategies in the vast majority of the setups. Furthermore, as the approach is highly flexible in the number of assets, we obtain strong results with only one configuration of hyperparameters, i.e., the size of the network and the training schedule.

Although our proposed DRL algorithm performs remarkably well, it comes with high computational costs as the algorithm is difficult to parallelize. Thus, we show how the task of algorithmic portfolio management can be conducted highly efficiently with vanilla feed-forward neural networks that are trained in a supervised learning framework. For that, we use a two-stage approach to predicting cross-sectional expected returns. First, we train a binary classification neural network that predicts the sign of returns. Second, we train two more models that predict future return magnitudes, whereas one solely learns from positive returns and the other from negative returns. Ultimately, we combine the models' individual predictions by first predicting the sign of future returns and, depending on this prediction, using the model that is trained on returns with the corresponding sign to predict the magnitude.

Separating the task of predicting cross-sectional returns enables the utilization of smaller models. This is due to the reduced complexity of the subtasks and the consideration of the distinct characteristics exhibited by negative returns compared to positive returns. The smaller models, on the other hand, are less sensitive to overfitting and can be trained with low computational efforts. We rigorously test the performance of the model stack in out-of-sample tests and show that the stack consistently outperforms passive investment benchmarks, even after accounting for common risk factors and transaction costs.

The stack also enables us to conduct a more comprehensive analysis of feature importance. We document large differences in the feature importance between three models, as, e.g., past volatility appears to be significantly more important for predicting the magnitude of positive returns than for negative returns or the direction of returns itself. On the other hand, firm characteristics are most valuable for predicting the direction of returns and least valuable for predicting the magnitude of negative returns. However, past returns are by far the most important predictors among all three models.

Overall, our application of deep learning for portfolio management highlights an important observation: a substantial part of return predictability arises from a non-linear relationship between past and future returns, whereas we suspect that this non-linearity is particularly concentrated around the threshold of zero. Closely related, the findings from behavioral finance suggest that investors exhibit distinct behavior on losses than on gains, as they tend to display a significantly higher level of risk aversion when faced with losses than gains (Kahneman and Tversky, 1979). Thus, we empirically assess in the fourth Chapter whether risk-averse investors avoid stocks with prior realized losses.

To this end, we develop a novel measure of risk aversion that depends on the long-short spread of common risk factors, for which we also establish a strong theoretical motivation. Based on this measure, we find that there are substantial differences in returns between stocks with preceding positive and negative returns during periods of high risk aversion. The disparity in returns stems from unusually low returns from stocks with preceding negative returns during periods of high risk aversion, as we argue that risk-averse investors become anxious to invest in stocks with recent realized losses, a phenomenon we refer to as the anxiety anomaly. Moreover, we find that investors' response to negative returns is convex in losses.

As the investor's reaction to losses appears to be predictable, we construct a straightforward linear model that exploits the anxiety anomaly to its capabilities in an out-of-sample setting. We show that a simple linear model based on the anxiety anomaly yields predictive performance comparable to that of a deep neural network with hundreds of predictors. The striking results suggest that the predictive performance of neural networks in cross-sectional returns may originate from the predictability of investors' preferences. In line with this result, we train a neural network using the preceding return and the long-short size factor spread to predict following returns, which indeed convergences to similar predictions as given by the anxiety anomaly. This could potentially also explain why we find neural networks to appreciate past returns the most among the considered features.

Increasing the frequency of analyzed returns, we document another predictive asset pricing anomaly in the fifth Chapter: the Monday afternoon reversal. We find that returns on Monday afternoons are significantly reversed throughout the remainder of the week. Similar to the anxiety anomaly, we do not find any explanation for this phenomenon in common systematic risk factors or overnight risk. Furthermore, the reversal is not due to limited arbitrage opportunities in small stocks and is attributable to both overnight and intraday returns during the rest of the week. The reversal is even of similar magnitude and significance among international equity ETFs.

We find strong empirical evidence that investors disagree about fundamental values and, thus, expected returns of stocks during the late Monday. Investors have significantly more spare time to process the weekly accumulated information in detail without receiving any market feedback, potentially leading to a heterogeneous evaluation of the information among investors. Thus, we test whether the reversal is stronger for stocks with high investor disagreement, which we measure by relatively high volume and volatility. When volume is low, we find no significant Monday afternoon reversal, while we even find significant momentum when volatility is low. On the other hand, stocks with high volume or volatility on Monday afternoons exhibit the strongest reversals.

Moreover, we provide additional evidence that investor disagreement is likely to be the underlying cause of the Monday afternoon reversal. For that, we analyze short interest data, in particular the days-to-cover of stocks on Mondays. As short sellers are commonly considered well-informed investors, we expect a stronger reversal when this group of informed traders disagrees with returns on Monday afternoons. Indeed, we find the strongest reversals when short sellers disagree with Monday afternoon returns in both directions, further supporting the idea that investor disagreement drives the Monday afternoon reversal.

Overall, this dissertation contributes to the literature in several ways. In the first place, we make contributions to the literature that is concerned with empirical portfolio management using machine learning. To this end, we develop a DRL algorithm that is capable of managing portfolios of any size while mitigating overfitting during training. Furthermore, we show that supervised deep learning can provide significant economic gains to investors without necessitating an advanced computational infrastructure. Our approach also reveals insights into feature importance, showing that past realized returns are the most important features, but with altering effects depending on whether return directions, positive or negative returns, are to be predicted. Motivated by these observations, this thesis additionally analyzes the predictability of past returns and asymmetries given by combining the long-short risk factor spreads and the return sign to explain the cross-section of returns and predict returns out-of-sample. These insights lead us to uncover two novel return anomalies that are both based on past realized returns: the anxiety anomaly and the Monday afternoon reversal. The first builds on behavioral finance, the second focuses explicitly on high frequency return predictability.

We find that investors' behavior is the most plausible explanation for both of the anomalies presented in this dissertation. As empirical models, such as neural networks, appear to capture non-linear patterns in the time series of returns like the anxiety anomaly we presented, are these patterns then merely instrumental variables in predicting long-lasting investors' behavior? This dissertation discusses the interconnections between the domains of artificial intelligence-based portfolio management and behavioral finance, providing evidence of predictable behavioral patterns of investors that are empirically identified by deep learning models.

With increasing computational possibilities and further enhancements in deep learning, it is reasonable to anticipate that these methods will identify a large number of yet unidentified factors that drive asset return. It may, however, be challenging to economically rationalize these empirical factors, leaving them as mere historical artifacts, which are not guaranteed to drive asset returns in the future independent of their historical performance. Consequently, deep learning-based portfolio management carries the additional methodological risk of whether these factors are merely noise in the data or true fundamental drivers of asset returns. We illustrate this risk in the third Chapter, demonstrating that trading performances may greatly vary depending on the portfolio and the neural network's convergence. Therefore, future advancements in deep learning-based portfolio management systems should not result in an exponential increase in the complexity of neural networks, but rather in analyzing their robustness and interpretability, in order to reduce the methodological risk component and build the required investors' trust in these systems.

References

Acharya, V. V., Amihud, Y., and Bharath, S. T. (2013). Liquidity risk of corporate bond returns: conditional approach. *Journal of Financial Economics*, 110(2):358–386.

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

Anatolyev, S. and Gospodinov, N. (2010). Modeling Financial Return Dynamics via Decomposition. *Journal of Business & Economic Statistics*, 28(2).

Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2006). The Cross-Section of Volatility and Expected Returns. *The Journal of Finance*, 61(1):259–299.

Baker, M. and Wurgler, J. (2006). Investor Sentiment and the Cross-Section of Stock Returns. *The Journal of Finance*, 61(4):1645–1680.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty*. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Bali, T. G., Cakici, N., and Whitelaw, R. F. (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99(2):427–446.

Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1):3–18.

Barberis, N., Mukherjee, A., and Wang, B. (2016). Prospect Theory and Stock Returns: An Empirical Test. *Review of Financial Studies*, 29(11):3068–3107.

Bekaert, G., Engstrom, E., and Grenadier, S. R. (2010). Stock and bond returns with Moody Investors. *Journal of Empirical Finance*, 17(5):867–894.

Bekaert, G. and Hoerova, M. (2016). What do asset prices have to say about risk appetite and uncertainty? *Journal of Banking & Finance*, 67:103–118.

Benartzi, S. and Thaler, R. H. (1995a). Myopic loss aversion and the equity premium puzzle. *The quarterly journal of Economics*, 110(1):73–92.

Benartzi, S. and Thaler, R. H. (1995b). Myopic Loss Aversion and the Equity Premium Puzzle. *The Quarterly Journal of Economics*, 110(1):73–92.

Bessembinder, H. and Hertzel, M. G. (1993). Return Autocorrelations around Nontrading Days. *The Review of Financial Studies*, 6(1):155–189.

Bianchi, D., Büchner, M., and Tamoni, A. (2021). Bond risk premiums with machine learning. *The Review of Financial Studies*, 34(2):1046–1089.

Birru, J. (2018). Day of the week and the cross-section of returns. *Journal of Financial Economics*, 130(1):182–214.

Boehmer, E., Jones, C. M., and Zhang, X. (2008). Which Shorts Are Informed? *The Journal of Finance*, 63(2):491–527.

Bogousslavsky, V. (2021). The cross-section of intraday and overnight returns. *Journal of Financial Economics*, 141(1):172–194.

Bollerslev, T., Li, S. Z., and Zhao, B. (2020). Good Volatility, Bad Volatility, and the Cross Section of Stock Returns. *Journal of Financial and Quantitative Analysis*, 55(3):751–781.

Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected Stock Returns and Variance Risk Premia. *The Review of Financial Studies*, 22(11):4463–4492.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and Asset Prices. *American Economic Review*, 103(3):623–628.

Breaban, A. and Noussair, C. N. (2018). Emotional State and Market Behavior. *Review of Finance*, 22(1):279–309.

Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Brusa, J., Lee, W. Y., and Liu, P. (2011). Monday returns and asset pricing. *Journal of Economics and Finance*, 35(3):332–347.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust Inference With Multiway Clustering. *Journal of Business & Economic Statistics*, 29(2):238–249.

Campbell, J. and Cochrane, J. (1999). By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy*, 107(2):205–251.

Carhart, M. M. (1997). On Persistence in Mutual Fund Performance. *The Journal of Finance*, 52(1):57–82.

Carlin, B. I., Longstaff, F. A., and Matoba, K. (2014). Disagreement and asset prices. *Journal of Financial Economics*, 114(2):226–238.

Chen, A. Y. and Zimmermann, T. (2022). Open Source Cross-Sectional Asset Pricing. *Critical Finance Review*, 11(2):207–264. Publisher: Now Publishers, Inc.

Chen, L., Pelger, M., and Zhu, J. (2023). Deep Learning in Asset Pricing. *Management Science*.

Chen, S., Luo, W., and Yu, C. (2021). Reinforcement learning with expert trajectory for quantitative trading. *arXiv preprint arXiv:2105.03844*.

Chordia, T., Goyal, A., and Saretto, A. (2020). Anomalies and false rejections. *The Review of Financial Studies*, 33(5):2134–2179.

Cohen, L., Diether, K. B., and Malloy, C. J. (2007). Supply and Demand Shifts in the Shorting Market. *The Journal of Finance*, 62(5):2061–2096.

Cohn, A., Engelmann, J., Fehr, E., and Maréchal, M. A. (2015). Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals. *American Economic Review*, 105(2):860–885.

Cookson, J. A. and Niessner, M. (2020). Why Don't We Agree? Evidence from a Social Network of Investors. *The Journal of Finance*, 75(1):173–228.

Cooper, M. J., Gulen, H., and Schill, M. J. (2008). Asset Growth and the Cross-Section of Stock Returns. *The Journal of Finance*, 63(4):1609–1651.

Coval, J. D. and Shumway, T. (2005). Do Behavioral Biases Affect Prices? *The Journal of Finance*, 60(1):1–34.

Dai, P.-F., Goodell, J. W., Huynh, L. D. T., Liu, Z., and Corbet, S. (2023). Understanding the transmission of crash risk between cryptocurrency and equity markets. *Financial Review*, 58(3):539–573.

Daniel, K. and Titman, S. (1997). Evidence on the Characteristics of Cross Sectional Variation in Stock Returns. *The Journal of Finance*, 52(1):1–33.

De Prado, M. L. (2018). *Advances in financial machine learning*. John Wiley & Sons.

de Prado, M. M. L. (2020). *Machine learning for asset managers.* Cambridge University Press.

Dechow, P. M., Hutton, A. P., Meulbroek, L., and Sloan, R. G. (2001). Short-sellers, fundamental analysis, and stock returns. *Journal of Financial Economics*, 61(1):77–106.

Desai, H., Ramesh, K., Thiagarajan, S. R., and Balachandran, B. V. (2002). An Investigation of the Informational Role of Short Interest in the Nasdaq Market. *The Journal of Finance*, 57(5):2263–2287.

Diether, K. B., Lee, K.-H., and Werner, I. M. (2009). Short-Sale Strategies and Return Predictability. *Review of Financial Studies*, 22(2):575–607.

Doyle, J. R. and Chen, C. H. (2009). The wandering weekday effect in major stock markets. *Journal of Banking & Finance*, 33(8):1388–1399.

Ederington, L. H. and Guan, W. (2010). How asymmetric is U.S. stock market volatility? *Journal of Financial Markets*, 13(2):225–248.

Fama, E. F. and French, K. R. (1992). The Cross-Section of Expected Stock Returns. *The Journal of Finance*, 47(2):427–465.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

Fama, E. F. and French, K. R. (2006). Profitability, investment and average returns. *Journal of Financial Economics*, 82(3):491–518.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.

Feng, G., Giglio, S., and Xiu, D. (2020). Taming the factor zoo: A test of new factors. *The Journal of Finance*, 75(3):1327–1370.

FirstRateData (2022). *High resolution financial data provider.* `https://firstratedata.com` [Accessed: 2022-11-30].

Fischer, T. and Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2):654–669.

Francis, J., Olsson, P., and Oswald, D. R. (2000). Comparing the Accuracy and Explainability of Dividend, Free Cash Flow, and Abnormal Earnings Equity Value Estimates. *Journal of Accounting Research*, 38(1):45–70.

French, K. R. (1980). Stock returns and the weekend effect. *Journal of Financial Economics*, 8(1):55–69.

Freyberger, J., Neuhierl, A., and Weber, M. (2020). Dissecting Characteristics Nonparametrically. *The Review of Financial Studies*, 33(5):2326–2377.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.

Gao, L., Han, Y., Zhengzi Li, S., and Zhou, G. (2018). Market intraday momentum. *Journal of Financial Economics*, 129(2):394–414.

Gibbons, M. R. and Hess, P. (1981). Day of the Week Effects and Asset Returns. *The Journal of Business*, 54(4):579–596.

Gilad, D. and Kliger, D. (2008). Priming the Risk Attitudes of Professionals in Financial Decision Making*. *Review of Finance*, 12(3):567–586.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT press.

Gu, S., Kelly, B., and Xiu, D. (2020). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies*, 33(5):2223–2273.

Gu, S., Kelly, B., and Xiu, D. (2021). Autoencoder asset pricing models. *Journal of Econometrics*, 222(1, Part B):429–450.

Guiso, L., Sapienza, P., and Zingales, L. (2018). Time varying risk aversion. *Journal of Financial Economics*, 128(3):403–421.

Harris, L. (1986). A transaction data study of weekly and intradaily patterns in stock returns. *Journal of Financial Economics*, 16(1):99–117.

Harvey, C. R. and Liu, Y. (2021). Lucky factors. *Journal of Financial Economics*, 141(2):413–435.

Harvey, C. R. and Siddique, A. (2000). Conditional Skewness in Asset Pricing Tests. *The Journal of Finance*, 55(3):1263–1295.

Hastie, T., Tibshirani, R., Friedman, J. H., and Friedman, J. H. (2009). *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

Helliwell, J. F. and Wang, S. (2014). Weekends and Subjective Well-Being. *Social Indicators Research*, 116(2):389–407.

Heston, S. L., Korajczyk, R. A., and Sadka, R. (2010). Intraday Patterns in the Cross-section of Stock Returns. *The Journal of Finance*, 65(4):1369–1407.

Hong, H. and Wang, J. (2000). Trading and Returns under Periodic Market Closures. *The Journal of Finance*, 55(1):297–354.

Hou, K., Xue, C., and Zhang, L. (2015). Digesting Anomalies: An Investment Approach. *Review of Financial Studies*, 28(3):650–705.

Hou, K., Xue, C., and Zhang, L. (2020). Replicating anomalies. *The Review of Financial Studies*, 33(5):2019–2133.

Hu, Y., Hong, Y., Feng, K., and Wang, J. (2023). Evaluating the Importance of Monetary Policy Uncertainty: The Long- and Short-Term Effects and Responses. *Evaluation Review*, 47(2):264–286. Publisher: SAGE Publications Inc.

Huang, C. Y. (2018). Financial trading as a game: A deep reinforcement learning approach. *arXiv preprint arXiv:1807.02787*.

Jegadeesh, N. (1990). Evidence of Predictable Behavior of Security Returns. *The Journal of Finance*, 45(3):881–898.

Jensen, M. C. (1968). The performance of mutual funds in the period 1945-1964. *The Journal of finance*, 23(2):389–416.

Jiang, H., Li, S. Z., and Wang, H. (2021). Pervasive underreaction: Evidence from high-frequency data. *Journal of Financial Economics*, 141(2):573–599.

Jiang, Z., Xu, D., and Liang, J. (2017). A deep reinforcement learning framework for the financial portfolio management problem. *arXiv preprint arXiv:1706.10059*.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–291.

Kamstra, M. J., Kramer, L. A., and Levi, M. D. (2003). Winter Blues: A SAD Stock Market Cycle. *American Economic Review*, 93(1):324–343.

Kamstra, M. J., Kramer, L. A., Levi, M. D., and Wang, T. (2014). Seasonally Varying Preferences: Theoretical Foundations for an Empirical Regularity. *The Review of Asset Pricing Studies*, 4(1):39–77.

Kane, A. (1982). Skewness Preference and Portfolio Choice. *Journal of Financial and Quantitative Analysis*, 17(1):15–25.

Kelly, B. T., Pruitt, S., and Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, 134(3):501–524.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kozak, S., Nagel, S., and Santosh, S. (2020). Shrinking the cross-section. *Journal of Financial Economics*, 135(2):271–292.

Kraus, A. and Litzenberger, R. H. (1976). Skewness Preference and the Valuation of Risk Assets. *The Journal of Finance*, 31(4):1085–1100.

Kraus, M. and Feuerriegel, S. (2017). Decision support from financial disclosures with deep neural networks and transfer learning. *Decision Support Systems*, 104:38–48.

Krauss, C., Do, X. A., and Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the s&p 500. *European Journal of Operational Research*, 259(2):689–702.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105.

Kuhnen, C. M. and Knutson, B. (2011). The Influence of Affect on Beliefs, Preferences, and Financial Decisions. *Journal of Financial and Quantitative Analysis*, 46(3):605–626.

Lakonishok, J. and Maberly, E. (1990). The Weekend Effect: Trading Patterns of Individual and Institutional Investors. *The Journal of Finance*, 45(1):231–243.

Lakonishok, J., Shleifer, A., and Vishny, R. W. (1994). Contrarian Investment, Extrapolation, and Risk. *The Journal of Finance*, 49(5):1541–1578.

Li, Y., Ni, P., and Chang, V. (2019). Application of deep reinforcement learning in stock trading strategies and stock forecasting. *Computing*, pages 1–18.

Liang, C., Hong, Y., Huynh, L. D. T., and Ma, F. (2023). Asymmetric dynamic risk transmission between financial stress and monetary policy uncertainty: thinking in the post-covid-19 world. *Review of Quantitative Finance and Accounting*, 60(4):1543–1567.

Lintner, J. (1965). The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *The Review of Economics and Statistics*, 47(1):13–37.

Lou, D., Polk, C., and Skouras, S. (2019). A tug of war: Overnight versus intraday expected returns. *Journal of Financial Economics*, 134(1):192–213.

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1):77–91.

Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.

Mitton, T. and Vorkink, K. (2007). Equilibrium Underdiversification and the Preference for Skewness. *The Review of Financial Studies*, 20(4):1255–1288.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.

Mossin, J. (1966). Equilibrium in a Capital Asset Market. *Econometrica*, 34(4):768–783.

Novy-Marx, R. and Velikov, M. (2016). A taxonomy of anomalies and their trading costs. *The Review of Financial Studies*, 29(1):104–147.

Park, H., Sim, M. K., and Choi, D. G. (2020). An intelligent financial portfolio trading strategy using deep q-learning. *Expert Systems with Applications*, 158.

Patel, J., Shah, S., Thakkar, P., and Kotecha, K. (2015). Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques. *Expert systems with applications*, 42(1):259–268.

Pigorsch, U. and Schäfer, S. (2022). High-Dimensional Stock Portfolio Trading with Deep Reinforcement Learning. In *2022 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFEr)*. IEEE.

Pigorsch, U. and Schäfer, S. (2023). Anxiety in Returns. *Journal of Behavioral Finance*.

Püttmann, L. (2018). Patterns of Panic: Financial Crisis Language in Historical Newspapers. SSRN Scholarly Paper 3156287, Social Science Research Network, Rochester, NY.

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.

Reinganum, M. R. (1983). The anomalous stock market behavior of small firms in January: Empirical tests for tax-loss selling effects. *Journal of Financial Economics*, 12(1):89–104.

Rogalski, R. J. (1984). New Findings Regarding Day-of-the-Week Returns over Trading and Non-Trading Periods: A Note. *The Journal of Finance*, 39(5):1603–1614.

Rozeff, M. S. and Kinney, W. R. (1976). Capital market seasonality: The case of stock returns. *Journal of Financial Economics*, 3(4):379–402.

Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk*. *The Journal of Finance*, 19(3):425–442.

Shen, J., Yu, J., and Zhao, S. (2017). Investor sentiment and economic forces. *Journal of Monetary Economics*, 86:1–21.

SimFin Analytics GmbH (2023). https://www.simfin.com/de/. Accessed: 2024-02-15.

Smirlock, M. and Starks, L. (1986). Day-of-the-week and intraday effects in stock returns. *Journal of Financial Economics*, 17(1):197–210.

Spaan, M. T. (2012). Partially observable markov decision processes. In *Reinforcement Learning*, pages 387–414. Springer.

Srivastava, S. et al. (2020). Deep reinforcement learning for asset allocation in us equities. *arXiv preprint arXiv:2010.04404*.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.

Taghian, M., Asadi, A., and Safabakhsh, R. (2020). Learning financial asset-specific trading rules via deep reinforcement learning. *arXiv preprint arXiv:2010.14194*.

Taylor, M. P. (2006). Tell me why I don't like Mondays: investigating day of the week effects on job satisfaction and psychological well-being. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 169(1):127–142.

Théate, T. and Ernst, D. (2021). An application of deep reinforcement learning to algorithmic trading. *Expert Systems with Applications*, 173.

Ülkü, N. and Andonov, K. (2016). Reversal of Monday returns. *Quantitative Finance*, 16(4):649–665.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention Is All You Need. Technical Report arXiv:1706.03762, arXiv.

Walther, M. and Münster, M. (2021). Conditional Risk Premiums and the Value Function of Prospect Theory. *Journal of Behavioral Finance*, 22(1):74–83.

Wang, K., Li, Y., and Erickson, J. (1997). A New Look at the Monday Effect. *The Journal of Finance*, 52(5):2171–2186.

Wang, W., Li, W., Zhang, N., and Liu, K. (2020). Portfolio formation with preselection using deep learning from long-term financial data. *Expert Systems with Applications*, 143:113042.

Weber, M., Weber, E. U., and Nosić, A. (2013). Who takes Risks When and Why: Determinants of Changes in Investor Risk Taking. *Review of Finance*, 17(3):847–883.

Wolff, D. and Echterling, F. (2020). Stock picking with machine learning. *Available at SSRN 3607845.*

Xie, J., Xu, B., and Chuang, Z. (2013). Horizontal and vertical ensemble with deep representation for classification. *arXiv preprint arXiv:1306.2759.*

Xiong, Z., Liu, X.-Y., Zhong, S., Yang, H., and Walid, A. (2018). Practical deep reinforcement learning approach for stock trading. *arXiv preprint arXiv:1811.07522.*

Yang, H., Liu, X.-Y., Zhong, S., and Walid, A. (2020). Deep reinforcement learning for automated stock trading: An ensemble strategy. *Available at SSRN.*

Zhang, Z., Zohren, S., and Roberts, S. (2020). Deep reinforcement learning for trading. *The Journal of Financial Data Science*, 2(2):25–40.

Zhong, X. and Enke, D. (2017). Forecasting daily stock market return using dimensionality reduction. *Expert Systems with Applications*, 67:126–139.