TECHNISCHE UNIVERSITÄT MÜNCHEN

TUM School of Management

Essays on Empirical Asset Pricing

Tobias Kalsbach

Vollständiger Abdruck der von der TUM School of Management der Technischen Universität München zur Erlangung eines

Doktors der Wirtschafts- und Sozialwissenschaften
(Dr. rer. pol.)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Reiner Braun
Prüfer der Dissertation:
1. Prof. Dr. Christoph Kaserer
2. Prof. Dr. Laurens Swinkels

Die Dissertation wurde am 09.03.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Management am 15.06.2023 angenommen.

Advisor: Prof. Dr. Christoph Kaserer Tobias Kalsbach

Essays on Empirical Asset Pricing

Abstract

This dissertation examines three research questions on empirical asset pricing. First, I[1] study how market news, firm-specific news, and noise diffuse among firms and how they affect stock returns in global networks. Market and firm-specific news as well as noise, are estimated through a structural vector auto-regression and the global network is based on analyst co-coverage. I show that investors show a categorical learning behavior by being able to differentiate between news and noise and rather processing market wide than firm-specific news. This underreaction of investors is driven by limited attention to the diffusing firm-specific news. Second, I compare various machine learning models to predict the cross-section of emerging market stock returns. I document that allowing for non-linearities and interactions leads to economically and statistically superior out-of-sample returns compared to traditional linear models. Furthermore, significant net returns can be achieved when accounting for transaction costs, short-selling constraints, and limiting the investment universe to big stocks only. Third, I investigate the anchoring effect as an explanation for investor underreaction to global firm-specific news. The anchoring effect refers to the tendency of investors to stick to their initial beliefs about a stock, even when facing new information. My results provide evidence of investors' distorted belief updating process and show that the anchoring of investors induced by the 52-week high impacts the processing of the firm-specific news. Regression analyses decompose stock returns into three independent components and reveal that the interaction effect between the firm-specific news return and the nearness to the 52-week high are related to a significant risk-adjusted return.

[1] In this dissertation, I use the term "I" in the introduction and conclusion. It does not necessarily refer to me directly since the second and third essay are based on joint work with my co-authors.

Betreuer: Prof. Dr. Christoph Kaserer Tobias Kalsbach

Aufsätze zu empirischer Kapitalmarktforschung

Kurzfassung

In dieser Dissertation werden drei Forschungsfragen zur empirischen Kapitalmarktforschung untersucht. Zunächst erforsche ich wie sich Markt- und firmenspezifische Nachrichten sowie Rauschen unter Unternehmen verbreitet und wie diese die Aktienrenditen in globalen Netzwerken beeinflussen. Die einzelnen Komponenten werden durch eine strukturelle Vektor-Autoregression bestimmt, und das globale Netzwerk basiert auf den Abschätzungen von Analysten. Ich zeige, dass Investoren ein kategorisches Lernverhalten zeigen, da sie zwischen Nachrichten und Rauschen unterscheiden und eher marktweite als firmenspezifische Nachrichten verarbeiten. Diese Unterreaktion der Anleger ist darauf zurückzuführen, dass sie den sich verbreitenden firmenspezifischen Nachrichten nur begrenzte Aufmerksamkeit schenken. Des Weiteren vergleiche ich verschiedene Modelle des maschinellen Lernens zur Vorhersage der Aktienrenditen von Schwellenländern. Ich belege, dass die Berücksichtigung von Nichtlinearitäten und Wechselwirkungen zu wirtschaftlich und statistisch besseren Renditen führt als bei traditionellen linearen Modellen. Darüber hinaus können erhebliche Nettorenditen erzielt werden, wenn Transaktionskosten und Leerverkaufsbeschränkungen berücksichtigt werden und das Anlageuniversum auf große Aktien beschränkt wird. Drittens untersuche ich den Verankerungseffekt als Erklärung für die Unterreaktion der Anleger auf globale unternehmensspezifische Nachrichten. Der Verankerungseffekt bezieht sich auf die Tendenz von Anlegern, an ihren ursprünglichen Überzeugungen über eine Aktie festzuhalten, selbst wenn sie mit neuen Informationen konfrontiert werden. Die Ergebnisse zeigen hierbei, getrieben durch die Nähe zum 52-Wochen-Hoch, dass Anleger ihre Ansichten über Aktien nicht anpassen, wenn neue Nachrichten erscheinen. Des Weiteren zerlege ich in einer Regressionsanalysen die Aktienrenditen in drei unabhängige Komponenten und zeigen, dass lediglich der Interaktionseffekt zwischen den neuen Nachrichten und der Nähe zum 52-Wochen-Hoch mit einer signifikanten risikobereinigten Rendite verbunden ist.

Overview

Contents

List of Tables

List of Figures

Contribution to Essays

Essay 1: What Diffuses in Stock Prices? The Roles of News and Noise in Global Networks

Authors: Tobias Kalsbach

Tobias Kalsbach reviewed the literature, developed the research design, collected all the data, conducted all analyses, interpreted the results, and prepared and revised the manuscript.

Tobias Kalsbach

Essay 2: Machine Learning and the Cross-Section of Emerging Market Stock Returns

Authors: Matthias Xaver Hanauer, Tobias Kalsbach

Tobias Kalsbach conducted all analyses and drafted the manuscript. Matthias Xaver Hanauer and Tobias Kalsbach were jointly involved in formulating the research question, reviewing the literature, developing the research design, collecting the data, constructing the sample, interpreting the empirical analyses, and revising the manuscript.

Matthias Xaver Hanauer

Tobias Kalsbach

Essay 3: Anchoring and Global Underreaction to Firm-Specific News

Authors: Tobias Kalsbach, Steffen Windmüller

Tobias Kalsbach conducted all analyses and drafted the manuscript. Steffen Windmüller and Tobias Kalsbach were jointly involved in formulating the research question, reviewing the literature, developing the research design, collecting the data, constructing the sample, interpreting the empirical analyses, and revising the manuscript.

Tobias Kalsbach   Steffen Windmüller

To my parents, my siblings, and Christina.

I am very grateful to my advisor, Christoph Kaserer, for his continuous support and exceptional guidance throughout my dissertation. Moreover, I am thankful to him for giving me the opportunity to carry out a joint research project with the Erasmus University Rotterdam and to complete a research stay abroad.

I am very grateful to Laurens Swinkels for hosting me at Erasmus University Rotterdam. During my research visit, I received very helpful comments on my work.

I would also like to express my gratitude to my co-authors Matthias Xaver Hanauer and Steffen Windmüller for their commitment and dedication throughout our collaborative work. I also thank my co-authors on other projects: Frederik Muskens and Laurens Swinkels.

Further, I thank my colleagues Vitor Azevedo, Daniel Bias, Minghui Chen, Christoph Großmann, Michael Haimann, Robert Heigermoser, Lisa Knauer, Daniel Schmidt, and Victoria Tressel for insightful discussions and encouragement.

Lastly, I am especially grateful to Reiner Braun for his comments and advice. I also thank all the participants of the conferences and seminars that I attended.

0 Introduction

The movement of asset prices is a complex and dynamic process that reflects a wide range of forward-looking information. Understanding how information is incorporated into asset prices and the factors influencing information diffusion and mispricing is crucial for investors, regulators, and policymakers in making informed decisions. It is essential to ensure that markets function efficiently, investors make informed decisions, and regulators and policymakers can protect investors and maintain market integrity.

Over the years, researchers have devoted significant attention to understanding how information is incorporated into asset prices. The efficient market hypothesis asserts that prices reflect all available information immediately and that any new information is instantly reflected in asset prices (Fama, 1970). However, this view has been challenged by studies showing that information diffusion in asset prices can be slow, resulting in the mispricing of underlying assets (Hong and Stein, 1999). The concept of information is complex and multifaceted, and it can be of various kinds, come from various sources, and vary among levels. Information can be public or private, firm-specific or market-wide, explicit or implicit, and originate at the firm level or among economically linked firms. Accurately assessing and incorporating information is crucial to making sound investment decisions and ensuring market efficiency. Different explanations exist for the slow diffusion of information resulting in the mispricing of the underlying asset. On the one hand, frictions in the market can prevent arbitrage from fully eliminating mispricing. These frictions can include transaction costs, liquidity constraints, and other barriers to the rapid and efficient transfer of assets between buyers and sellers. On the other hand, behavioral biases can also impact the processing of information and the resulting mispricing of assets. For instance, limited attention to information arrival or the source of information, the anchoring on previous prices, and the missing awareness of the interaction between information components can influence how investors process new information.

This dissertation consists of three essays examining empirical asset pricing research questions. In the first essay, I study the slow diffusion of firm-specific news, market news, and noise from fundamentally related firms into the focal firms' stock prices. In the second essay, I compare various machine learning models to predict the cross-section of emerging market stock returns. In the third essay, I investigate how the anchoring bias distorts investors' belief updating process after the arrival of global firm-specific news.

0.1 Research questions

In each essay, I utilize a specific methodology and data set to examine the respective research question. I outline the three research questions and designs in the following subsections.

0.1.1 What Diffuses in Stock Prices? The Roles of News and Noise in Global Networks

Understanding what drives the diffusion of information among fundamentally linked firms is crucial to get a better sense of the behavior of investors and how they process information. While neoclassic theory states that markets should be informationally efficient and incorporate all available information about future values, the slow information diffusion among firms presents contrary evidence. A behavioral explanation of this empirical artifact is that investors have limited access to a scarce cognitive resource—attention—which limits investors in their information processing capabilities. This results in partial processing of a firm's publicly available and relevant information environment leading to a delayed stock return reaction. This limited attention and, therefore, limited processing capabilities cause a categorial-learning behavior that lets investors rather digest market and sector-wide information than firm-specific information.

My sample for the empirical analysis covers the period from January 1992 to December 2021 and is determined by the availability of the international stock market and analyst coverage data. The underlying sample consists of 42,789 stocks from 49 equity markets. I limit myself to countries included in the MSCI Developed or Emerging Markets Index in the respective year and stocks linked to at least one other stock. For non-U.S. countries, the stock market data is from Refinitiv Datastream, and the accounting data is retrieved from Refinitiv Worldscope. In the case of U.S. stocks, data is obtained from the Center for Research on Security Prices (CRSP) and the Compustat database. Further, I include analyst-related data from Institutional Brokers' Estimate System (I/B/E/S) and institutional ownership data from the FactSet Ownership database. The final sample consists of 3,901,237 stock-month observations.

I investigate the relation between the slow diffusion of different information components from linked firms and the reaction of the focal firms' stock prices. I decompose a firm's stock return into its firm-specific news, market news, and noise component using a structural vector auto-regression to determine the different information components. To estimate the slow information diffusion, I further aggregate the stock information components into three different spillover measures. I use portfolio sorts and cross-sectional regressions to determine which spillover component affects stock returns, how long it takes till they fully diffuse into a firm's stock price, and what causes the slow diffusion?

I present robust evidence that firm-specific news is the key driver of information diffusion across fundamentally linked firms. The cross-sectional return difference between firms exposed to negative news from linked firms and those exposed to positive news amounts to approximately 7% per year. For large firms, the firm-specific news diffuses into the stock price on average within one month, whereas market news is directly incorporated. When looking at small firms, the noise component and the market news are predictive for one month, while the firm-specific news takes about three months to be fully incorporated into the stock price. I further prove that investors underreact to firm-specific news because of limited attention. The results imply that investors can differentiate between news and noise and further rather process market than firm-specific news, which is consistent with categorical learning behavior.

0.1.2 Machine Learning and the Cross-Section of Emerging Market Stock Returns

Machine learning algorithms have been available for a long time. However, due to increased computing power and data availability, decreased data storage costs, and algorithmic innovations in recent years, machine learning methods have seen increasing popularity in research fields such as economics, finance, and accounting. This paper compares various machine learning models to predict the cross-section of emerging market stock returns. More specifically, I analyze the predictive power of nine algorithms: ordinary least squares regression and elastic net as examples for traditional linear models; tree-based models such gradient boosted regression trees and random forest; and neural networks with one to five layers. Furthermore, I investigate the performance of an ensemble comprising the five different neural networks and an ensemble of methods that allow for non-linearities and interactions, i.e., the two tree-based models and the ensemble of neural networks.

The sample comprises data from emerging stock markets as classified by Morgan Stanley Capital International (MSCI). The accounting data is from Refinitiv Worldscope, and the stock market data is from Refinitiv Datastream. The sample period starts in July 1995 and ends in December 2021. The result is a comprehensive dataset spanning 15.152 unique stocks from 32 emerging market countries with more than 1.42 million stock-month observations and the 36 firm-level characteristics falling into categories such as value, past returns, investment, profitability, intangibles, and trading frictions.

This paper compares various machine learning models to predict the cross-section of emerging market stock returns. More specifically, I analyze the predictive power of nine algorithms: ordinary least squares regression and elastic net as examples for traditional linear models; tree-based models such gradient boosted regression trees and random forest; and neural networks with one to five layers. Furthermore, I investigate the performance of an ensemble comprising the five different neural networks and an ensemble of methods that allow for non-linearities and interactions, i.e., the two tree-based models and the ensemble of neural networks.

The main findings can be summarized as follows. First, I document that the different prediction algorithms pick up similar characteristics. However, I observe that tree-based methods and neural networks also identify non-linearities and interactions of characteristics. In contrast, linear methods are restricted to linear relationships and do not allow for interactions among characteristics. Second, return forecasts based on machine learning models lead to economically and statistically superior out-of-sample long-short returns compared to traditional linear models. Furthermore, a factor model can only partly explain these long-short returns, and their alphas remain highly significant. These findings are robust to several methodological choices and for emerging market subregions. Finally, I document that machine learning forecasts beat linear models consistently over my sample period, and I cannot observe a decline in predictability over time. Third, developed market long-short returns based on machine learning forecasts derived in the same way as their emerging market counterparts cannot explain emerging market out-of-sample returns. However, models estimated solely on developed markets data also predict emerging market stock returns. These findings indicate that similar relationships between firm characteristics and future stock returns exist for developed and emerging markets but that the pricing of these characteristics is not fully integrated between developed and emerging markets. Fourth, the high returns of the machine learning strategies in emerging markets do not primarily stem from higher-risk months and do not revert quickly, suggesting that an underreaction explanation is more likely than a risk-based explanation. Furthermore, both linear and machine learning models show higher predictability for stocks associated with higher limits to arbitrage. However, I also document that this effect is less pronounced for machine learning forecasts than for linear regression forecasts, indicating that the superiority of machine learning models in emerging markets does not stem from limits to arbitrage. Finally, accounting for transaction costs, short-selling constraints, and limiting my investment universe to big stocks, I document that machine learning-based return forecasts can lead to significant net outperformance over the market and net alphas, at least when efficient trading rules are applied.

0.1.3 Anchoring and Global Underreaction to Firm-Specific News

Investor underreaction to the arrival of news has been a long-standing topic in the finance literature. A large body of empirical and theoretical evidence argues that firms' stock prices respond slowly due to investors' behavioral biases. Theoretical literature often suggests that investors' limited attention results in underreaction to the news. At the same time, empirical evidence supports this limited attention hypothesis by showing that firms' stock prices respond slowly to the arrival of new information. I aim to test a novel psychological explanation, the anchoring effect, as an additional explanation for investor underreaction to global firm-specific news measured through the nearness to the 52-week high price. The anchoring effect refers to the tendency of investors to stick to their initial beliefs about a stock, even when facing new information. This psychological barrier can be enforced when investors use the 52-week high as an anchor when making investment decisions. For example, investors influenced by the anchoring effect will not fully adjust their beliefs if the firm experiences the arrival of positive news (negative news) and if the stock price is close to (far from) the 52-week high, leading to a slow stock price response.

My sample for the empirical analysis covers the period from January 2004 to December 2021 and is determined by the broad availability of firm-specific news data. The underlying sample consists of 24,337 stocks from 23 equity markets. I limit myself to countries included in the MSCI Developed Markets Index in the respective year and stocks experiencing a minimum of one firm-specific news event in the previous month. The U.S. and international equities analyses are based on a global sample comprising stock market data from Refinitiv Datastream and accounting data from Worldscope. To combine the stock market data with the firm-specific news, I follow a multi-step procedure to find all corresponding news articles for a corresponding firm on a trading day. Additionally, I include analyst and institutional ownership data for the stock data. All analyst-related data is collected from Institutional Brokers' Estimate System (I/B/E/S), whereas the Institutional ownership data is from the FactSet Ownership database. The final sample consists of 1,417,250 stock-month observations.

The main objective is to investigate the impact of firm-specific news in conjunction with the proximity to the 52-week high on investor behavior. I begin by forming independent, country-neutrally double-sorted quintile portfolios using the last month's firm-specific news return and the nearness to the 52-week high at the previous month-end as sorting criteria. Afterward, I utilize a return decomposition methodology to disentangle the stock return predictability into three components. The first component measures the pure firm-specific news return, the second the pure effect resulting from the stock price nearness to its 52-week high, and the third component the interaction effect between the firm-specific news return and the nearness to its 52-week high.

The main findings can be summarized as follows. First, the interaction effect yields an average Fama-French-Carhart (1997) four-factor alpha of 1.47% (t=4.67). In contrast, the pure firm-specific news effect and pure 52-week high effect are insignificant, excluding the interaction effects from the regression results in two positive and significant pure effects. These results allow me to conclude that the investors' underreaction to the firm-specific news is partially explained by the anchoring bias induced by the nearness to the 52-week high. Second, I investigate the role of a stock's limits to arbitrage in causing mispricing. My results provide evidence that the induced underreaction of investors is indeed driven by firms with high limits to arbitrage. The effect exists among stocks that are smaller in market capitalization, have lower institutional ownership or analyst coverage, have higher idiosyncratic volatility, and have higher transaction costs. Third, Restricting the company-specific news solely to earnings announcement days decreases the risk-adjusted return of the interaction effect, causing it to lose significance. However, when these days are excluded, the global four-factor alpha increases to 1.68% (t=5.17) per month. The interaction effect in the U.S. stock market, which is the most efficient, becomes insignificant when a slower information diffusion process is modeled. Moreover, a positive and significant risk-adjusted return is observed globally after eliminating macroeconomic announcements and the predictable component from daily returns. Lastly, I explore how the nearness to the 52-week high distorts the belief-updating process leading to an underreaction. I use analyst recommendation changes as a direct proxy to observe the belief updating process in financial markets. My results suggest that analysts are indeed influenced by firm-specific news as they change their recommendations after the arrival of news. However, the upgrade (downgrade) is less likely if positive (negative) news arrives at the firm and the underlying stock price is near (far from) the 52-week high. The findings provide evidence for the hypothesis that the belief updating process is distorted and influenced by stock prices' nearness to the 52-week high and the arrival of firm-specific news.

0.2 Contributions

The three essays of this dissertation contribute to multiple strands of the literature. I briefly summarize the main contribution of each essay.

In the *first* essay, I investigate the slow diffusion of firm-specific news, market news, and noise into the focal firms' stock prices. The firm-specific news component is the key driver behind the slow information diffusion around the globe. Firms exposed to most negative firm-specific news from linked firms earn approximately 7% more per year than firms exposed to positive firm-specific news. For large firms, the firm-specific news diffuses into the stock price on average within one month, whereas market news is directly incorporated. When looking at small firms, the noise component and the market news are predictive for one month, while firm-specific news takes about three months to diffuse into the stock price fully. The results imply that investors can differentiate between news and noise and rather process market than firm-specific news, consistent with categorical learning behavior. Further, the underreaction to firm-specific news is driven by investors' limited attention.

The findings contribute to the literature on the role of slow and gradual information diffusion among fundamentally linked firms for asset prices (Moskowitz and Grinblatt, 1999; Menzly and Ozbas, 2010; Cohen and Lou, 2012; Lee et al., 2019; Parsons, Sabbatucci and Titman, 2020; Ali and Hirshleifer, 2020). Disentangling firm-specific news, market news, and noise from each other allows for unveiling the component of the market inefficiency causing the cross-firm predictability (Burt and Hrdlicka, 2021). Further, I account for the investors' categorical learning behavior, which sheds light on the mechanism leading to the slow information diffusion among fundamentally linked firms (Peng and Xiong, 2006; Huang, Lin and Xiang, 2021; Huang et al., 2022). Lastly, the global sample contributes to understanding stock price formation (Jacobs and Müller, 2020).

In the *second* essay, I compare various machine learning models to predict the cross-section of emerging market stock returns. I analyze the predictive power of nine algorithms. The algorithms cover traditional linear models, tree-based models, and neural networks. Furthermore, I include two ensemble methods that allow for non-linearities and interactions. I train the algorithms using 36 firm-level characteristics for 32 emerging market countries from July 1995 to December 2021, while the 20-year out-of-sample period is from January 2002 to December 2021. I document that return forecasts from machine learning methods lead to superior out-of-sample returns in emerging markets. Interestingly, investors already applying such a strategy in developed markets seem to enjoy potential diversification benefits when applying them also in emerging markets. I further investigate the source of the predictability and conclude that it rather stems from mispricing than higher risk. Still, the superiority of machine learning models in emerging markets does not stem from limits to arbitrage. Finally, significant net returns can be achieved when accounting for transaction costs, short-selling constraints, and limiting my investment universe to big stocks only.

This essay contributes to the literature in at least three aspects. I contribute to the rapidly expanding literature on predicting the cross-section of stock returns with machine learning methods. To this point, there exists only evidence that more complex machine learning models are superior to linear models in developed markets (Rasekhschaffe and Jones, 2019; Freyberger, Neuhierl and Weber, 2020; Gu, Kelly and Xiu, 2020; Tobek and Hronec, 2020; Drobetz and Otto, 2021). Under the hypothesis that developed markets are integrated, the same risk factors should apply to these markets. Therefore, similar results within developed markets are unsurprising, and emerging markets provide an attractive alternative for out-of-sample tests in independent and new samples. Further, I add to the literature on the drivers of emerging market stock returns and market integration (Bekaert and Harvey, 1995; Harvey, 1995). The machine learning models allow me to consider non-linearities and interactions next to linear relationships. Lastly, I contribute to understanding the source of return predictability from machine learning forecasts (Avramov, Cheng and Metzker, 2022; Leung et al., 2021; Cakici et al., 2022a). I provide evidence that machine learning models show higher predictability for stocks associated with higher limits to arbitrage. A positive and significant outperformance can be achieved even when accounting for transaction costs, short-selling constraints, and limiting the investment universe to big stocks only.

In the *third* essay, I examine investor underreaction to global firm-specific news and seek to test the anchoring effect as an explanation. The anchoring effect refers to the tendency of investors to cling to their initial beliefs even when facing new information, as reinforced by their use of the 52-week high as an anchor. The paper tests the central hypothesis that the anchoring effect distorts the investor's belief updating process after the arrival of firm-specific news, resulting in the predictability of future stock returns. The sample for the empirical analysis covers stocks from developed markets across 23 countries from January 2004 to December 2021. Using a novel return decomposition methodology allows me to conclude that the investors' underreaction to the firm-specific news is partially explained by the anchoring bias induced by the nearness to the 52-week high. Furthermore, I provide evidence that the stock's limits to arbitrage are causing the mispricing and inducing investors' underreaction. In addition, I show that the anchoring bias effect on investors' underreaction over the subsequent month is driven by unscheduled, firm-specific news. Finally, I show that analysts react to firm-specific news but are less likely to change their recommendation if the stock price is near the 52-week high.

This study adds to understanding investor underreaction in an international asset pricing context in at least four aspects. First, I contribute to a better understanding investor underreaction by explicitly using firm-specific news (Jiang, Li and Wang, 2021) instead of proxying news with economically-linked, past-month firm momentum (Huang, Lin and Xiang, 2021). I provide insights into investor underreaction by showing that limits to arbitrage amplify the underreaction potential. Second, I reveal a crucial economic mechanism behind investor underreaction in global equity markets. I rely on the anchoring and adjustment hypothesis by showing that professional forecasters (Campbell and Sharpe, 2009; Cen, Hilary and Wei, 2013) include the firm-specific news in their recommendation but are affected by the anchoring bias if the stock is near (far from) the 52-week high and positive (negative) news arrives. Third, I show that unscheduled, firm-specific news drives the anchoring bias effect on investors' underreaction over the subsequent month. Empirical evidence suggests that investors' underreaction is driven by scheduled news (Birru, 2013; George, Hwang and Li, 2014). My results on the investors' distorted belief updating process provide strong evidence of a longer-dated, monthly investor underreaction to unscheduled news, indicating that unscheduled news items require more time to be reflected within stock prices. Fourth, I contribute to the literature on empirical asset pricing for global equity markets by using an international sample and extended metrics. Most literature on news-induced momentum (Chan, 2003; Gutierrez and Kelly, 2008; Hillert, Jacobs and Müller, 2014; Jiang, Li and Wang, 2021) concentrates solely on the U.S. stock market. Therefore, I add to the ongoing discussion about the investor underreaction hypothesis and its economic channels by providing non-U.S. out-of-sample evidence (Hou, Xue and Zhang, 2018) for the anchoring bias and investor underreaction to firm-specific news.

0.3 Outline

The remainder of this dissertation is structured as follows. In Chapter 1, I analyze the relation between the slow diffusion of information components from linked firms and the reaction of the focal firms' stock prices. In Chapter 2, I compare machine learning models to predict the cross-section of emerging market stock returns. In Chapter 3, I study how the anchoring bias distorts investors' belief updating process, leading to an underreaction to global firm-specific news. Finally, in Chapter 4, I briefly summarize the main results and highlight their contributions and implications.

1 What Diffuses in Stock Prices? The Roles of News and Noise in Global Networks

Abstract

This paper studies how market news, firm-specific news, and noise diffuse among firms and how they affect stock returns in global networks. Market and firm-specific news as well as noise, are estimated through a structural vector auto-regression and the global network is based on analyst co-coverage. I show that investors show a categorical learning behavior by being able to differentiate between news and noise and rather processing market wide than firm-specific news. It takes up to 3 months for investors to adjust the prices of stocks to the information contained in the firm-specific news component. This underreaction of investors is driven by limited attention to the diffusing firm-specific news.

Key words:	Asset Pricing, Information and Market Efficiency, Limited Attention, Analysts, Stock-Returns
JEL Codes:	G12, G14, G17, G41, L14

Authors:	Tobias Kalsbach
First Author:	Tobias Kalsbach

Presentations:	36^{th} International Conference of the French Finance Association 2021 TUM School of Management 2020 Robeco
Current Status:	Working Paper

1.1 Introduction[1]

Having an understanding of what drives the diffusion of information among fundamentally linked firms is crucial to get a better sense of the behavior of investors and how they process information.[2] Recent empirical work mostly focuses on identifying new fundamental links (Moskowitz and Grinblatt, 1999; Hou, 2007; Cohen and Lou, 2012; Cohen and Frazzini, 2008; Menzly and Ozbas, 2010; Hoberg and Phillips, 2018; Lee et al., 2019; Ali and Hirshleifer, 2020; Parsons, Sabbatucci and Titman, 2020; Ying, 2020; Lee et al., 2022), or the psychological explanation of the delayed price response (Huang, Lin and Xiang, 2021; Huang et al., 2022). However, those literature streams fail to identify which type of information nor how fast the information diffuses (Burt and Hrdlicka, 2021).

While neoclassic theory states that markets should be informational efficient and therefore incorporate all available information about future values (Fama, 1970), the slow information diffusion among firms presents contrary evidence to this theory. A behavioral explanation of this empirical artifact is that investors have limited access to a scarce cognitive resource—attention (Kahneman, 1973)—which limits investors in their information processing capabilities (Simon, 1955; Jensen and Heckling, 1995). This results in partial processing of a firm's publicly available and relevant information environment leading to a delayed stock return reaction (Hirshleifer and Teoh, 2003; Hung, Li and Wang, 2015). This limited attention and, therefore, limited processing capabilities cause a categorial-learning behavior that lets investors rather digest market and sector-wide information than firm-specific information (Peng and Xiong, 2006).

This paper investigates the relation between the slow diffusion of different information components and focal firms' prices and derives insights on the following questions. Which return components diffuse among fundamentally linked firms? How long does the diffusion take place? What causes the slow diffusion? I utilize the return decomposition model of Brogaard et al. (2022), which allows me to decompose a stock return into firm-specific news, market news, and a noise component. To model information flow, I rely on the

[1] I thank David Blitz, Matthias Hanauer, Christoph Kaserer, Lisa Knauer, Michael Weber, Sebastian Müller, Cameron Peng, Laurens Swinkels, Milan Vidojevic, Florian Weigert, Steffen Windmüller, and the conference and seminar participants at the 36th International Conference of the French Finance Association (AFFI), TUM School of Management, and Robeco for helpful comments.

[2] This chapter is based on Kalsbach (2023).

expertise of analysts (Lee and So, 2017) and use the individual analyst's firm co-coverage to identify the link between firms (Ali and Hirshleifer, 2020; Kaustia and Rantala, 2020). I use portfolio sorts and Fama and MacBeth (1973) regressions to examine which, how, and why information diffuses into a firm's stock prices. I present robust evidence that firm-specific news is the key driver of information diffusion across fundamentally linked firms. The cross-sectional return difference between firms exposed to negative news from linked firms and those with exposure to positive news amounts to approximately 7% per year. For large firms, the firm-specific news diffuses into the stock price on average within one month, whereas market news is directly incorporated. When looking at small firms, the noise component and the market news are predictive for one month, while the firm-specific news takes about 3 months to be fully incorporated into the stock price. I further provide evidence that investors underreact to firm-specific news because of limited attention. The results imply that investors can differentiate between news and noise and further rather process market than firm-specific news, which is consistent with categorical learning behavior.

The closest paper to my research are Burt and Hrdlicka (2021), Huang, Lin and Xiang (2021), and Huang et al. (2022). Burt and Hrdlicka (2021) decomposes the diffusing return into an idiosyncratic and predictable return component. The predictable component is the driver for long-term, cross-firm predictability, which indicates that slow information diffusion is not the only source of cross-firm predictability. Huang, Lin and Xiang (2021) argue that investors do not update their beliefs about a firm due to the anchoring effect causing an underreaction to the diffusing news. Huang et al. (2022) show that investors underreact to continuous information while discrete information is quickly absorbed into price. All three papers focus on a U.S.-only setting, while I use a network of global firms.

This paper contributes to the literature on momentum spillovers originating from fundamentally related firms by providing an understanding of the role of news diffusion among fundamentally linked firms for asset prices. I am the first to study the diffusion of news among fundamentally linked firms in a large sample covering both developed as well as emerging markets. The final sample contains 30 years of data on 42,789 firms that are located in 49 countries. Further, I contribute to the literature by differentiating between firm-specific news, market news and noise using the return decomposition model of Brogaard et al. (2022). These individual components allow me to uncover the return component causing the cross-firm predictability. To my knowledge, I am the first to combine the categorical learning behavior by Peng and Xiong (2006) with the diffusion of market news and firm-specific news among firms and further relate this slow information diffusion to investors' underreaction due to limited attention.

The remainder of this work is structured as follows: Section 1.2 reviews related literature. Section 1.3 describes the underlying return decomposition into firm-specific news, market news, and noise and the construction of the analyst co-coverage network. Section 1.4 introduces the dataset. Section presents evidence from portfolio sorts and cross-sectional regression and relates the news diffusion to underreaction and limited attention. Section 1.6 concludes.

1.2 Related literature

1.2.1 Information diffusion in networks

The first strand of literature related to this paper focuses on the slow and gradual information diffusion between fundamentally linked firms. The first set of studies focuses on the slow adjustment of stock prices to industry information (Moskowitz and Grinblatt, 1999; Ramnath, 2002; Hong, Torous and Valkanov, 2007; Hou, 2007; Hoberg and Phillips, 2018). Further, less complex firms adjust their prices faster to this industry-wide information (Cohen and Lou, 2012). Other sources of relevant industry information are the industries connected through the supply chain (Menzly and Ozbas, 2010), and industries that occur along complementarity networks (Lee et al., 2022). Besides the cross-industry supply chains, direct firm-to-firm costumer-supplier relationships are crucial when studying how information diffuses (Cohen and Frazzini, 2008). Another way to identify firms with exposure to new information is the usage of related technologies (Lee et al., 2019), companies with headquarters located in the same city (Parsons, Sabbatucci and Titman, 2020), and common institutional ownership (Ying, 2020). A network that unifies most of these information channels is based on analyst co-coverage (Ali and Hirshleifer, 2020), which I use as the underlying network in this paper to study which news is diffusing among firms. In addition to the existing literature, this paper takes a global view on the role of information diffusion among fundamentally linked firms for asset returns by combining global individual analyst coverage with international stock market data. This novel dataset allows me to model the diffusion of information within and across 38 industries as well as within and across 49 countries.

1.2.2 News decomposition measures

The second strand of literature related to this paper explores the various news decomposition measures proposed in the literature. Prominent news measures are cash-flow and discount-rate news which can be estimated using different methodologies like the vector auto-regression (Campbell and Shiller, 1988; Campbell, 1991; Vuolteenaho, 2002), implied cost of capital (Chen, Da and Zhao, 2013), or direct estimation of the cash-flow news (Easton and Monahan, 2005; Da and Warachka, 2009; Da, Liu and Schaumburg, 2014). An alternative measure decomposes the return into market-wide news and firm-specific news using regressions of stock returns on market return (Roll, 1988; Morck, Yeung and Yu, 2000). Another measure differentiates return into the news- and non-news-driven components Jiang, Li and Wang (2021). Brogaard et al. (2022) propose an alternative that allows distinguishing between market news, firm-specific news, and noise. While Brogaard et al. (2022) tries to understand better the roles of information and noise for the focal firm itself, I add to this literature by testing the diffusion of market news, firm-specific news, and noise in an empirical asset pricing context. I provide evidence that firm-specific news is the actual return component that diffuses slowly among fundamentally linked firms.

1.2.3 Underreaction and limited attention

The third strand focuses on the growing literature on investors' limited attention. Traditional asset pricing models often fail to explain the existence of certain profitable trading strategies. These strategies often arise due to the recognition of information caused by limited attention. Investors cannot fully process all information driven by the scarcity of attention (Kahneman, 1973). Therefore investors avoid stocks they are unfamiliar with (Merton, 1987). Processing the information signals of a stock is not a binary decision, as investors neglect publicly available accounting information in the case the expected cost of processing outweighs the expected benefits (Hirshleifer and Teoh, 2003). Further, investor neglect information of the latest earning release (Hirshleifer, Lim and Teoh, 2011). Due to their limited attention, investors further show a categorical learning behavior which makes them prone to process market- and sector-wide than firm-specific information (Peng and Xiong, 2006). Broad empirical literature also provides evidence on investors' limited attention. This behavioral bias influences investors' trading behavior as they trade on market-wide attention-grabbing events (Yuan, 2015), whereas only individual investors are impacted by these events (Barber and Odean, 2008). On the other hand, institutional investors can process complex information faster (Cohen and Frazzini, 2008). There exists a variety of proxies that model the attention of investors, including trading volume and down market periods (Hou, Xiong and Peng, 2009), media attention (Drake, Guest and Twedt, 2014; Twedt, 2016), past return patterns (Huang et al., 2022), Fridays (DellaVigna and Polle, 2009), non-trading hours (Francis, Pagach and Stephan, 1992), analyst coverage (Hong, Torous and Valkanov, 2007), and many more. Due to limited attention, investors are not able to fully process the complex information environment (Chen, Da and Zhao, 2013), resulting in the situation that complex relations between firms might be overlooked (Kovacs, 2016). I contribute to this literature by providing empirical evidence on the categorical-learning behavior and by testing investors' limited attention when firm-specific news diffuses.

1.3 Empirical measurement

1.3.1 News and noise

The stock price of a firm is driven by various components. The measure of Brogaard et al. (2022) explicitly differentiates between market-wide information, firm-specific information, and a noise component. It differs from other proposed measures in the literature by not requiring data that varies on annual frequency (Campbell and Shiller, 1988; Campbell, 1991; Vuolteenaho, 2002), using biased analyst estimates (Easton and Monahan, 2005; Da and Warachka, 2009; Da, Liu and Schaumburg, 2014; Chen, Da and Zhao, 2013), leaving out the noise component (Roll, 1988; Morck, Yeung and Yu, 2000), or relying on third-party data (Jiang, Li and Wang, 2021).

In the first step, Brogaard et al. (2022) divide the stock return into the discount rate (μ), new information about the stock's fundamentals which follow a random walk (w_t), and a pricing error (Δs_t).[3] The new information is then further categorized in market-wide information ($\theta_{r_m}\varepsilon_{r_m,t}$), firm-specific private information revealed through trading ($\theta_x\varepsilon_{x,t}$), and firm-specific public information not captured by trading ($\theta_r\varepsilon_{r,t}$) as described in Equation (1.1).

$$\begin{aligned} r_t &= \mu + w_t + \Delta s_t \\ w_t &= \theta_{r_m}\varepsilon_{r_m,t} + \theta_x\varepsilon_{x,t} + \theta_r\varepsilon_{r,t} \end{aligned} \tag{1.1}$$

To estimate the individual components of Equation (1.1), a structural vector autoregression (VAR) with five lags to allow for a full week of serial correlation and lagged effects is applied in Equation (1.2).[4] It includes an individual regression for the market return ($r_{m,t}$), the signed dollar volume of trading in the given stock (x_t)[5], and the stock return (r_t). I estimate the VAR on a monthly rolling base requiring a minimum of 20 daily observations while the stock needs to trade in the last 12 months.

$$\begin{aligned} r_{m,t} &= a_0 + \sum_{l=1}^{5} a_{1,l} r_{m,t-l} + \sum_{l=1}^{5} a_{2,l} x_{t-l} + \sum_{l=1}^{5} a_{3,l} r_{t-l} + \varepsilon_{r_m,t} \\ x_t &= b_0 + \sum_{l=1}^{5} b_{1,l} r_{m,t-l} + \sum_{l=1}^{5} b_{2,l} x_{t-l} + \sum_{l=1}^{5} b_{3,l} r_{t-l} + \varepsilon_{x,t} \\ r_t &= c_0 + \sum_{l=1}^{5} c_{1,l} r_{m,t-l} + \sum_{l=1}^{5} c_{2,l} x_{t-l} + \sum_{l=1}^{5} c_{3,l} r_{t-l} + \varepsilon_{r,t} \end{aligned} \tag{1.2}$$

Next, the permanent long-run effects (θ_{r_m}, θ_x, θ_r) have to be estimated.[6] They are

[3] Note: The drift in the efficient price (μ) is a constant in this model. New information about a stock is unpredictable; therefore, it holds $\mathrm{E}_{t-1}[w_t] = 0$.

[4] According to Brogaard et al. (2022), the lag structure of the VAR accounts for short-term momentum as well as reversals, persistence in order flow, first-order serial correlation in market returns due to non-synchronous trading and delayed stock-price reactions to market-wide information.

[5] Following Pástor and Stambaugh (2003), this proxy has minimal data requirements, whereby positive values indicate a net buying and negative values a net selling behavior.

[6] θ_{r_m} is estimated through the unit shock of $\varepsilon_{r_m,t} = 1$, θ_x follows by unit shock $\varepsilon_{x,t} = 1$, and θ_r is derived

inferred from the cumulative impulse response function of the returns (r_t) at $t = 15$ using the respective unit shock. This is equal to Brogaard et al. (2022) and allows the return to stabilize after the arrival of a shock.

To derive the firm-specific, the market information and the noise component, I aggregate the individual daily residuum in the structural VAR (ε_t) and the respective long-run effects (θ) of the last month using the model as outlined in Equation (1.3). The return attributable to the noise component equals the realized return that is not attributable to either information or the discount rate.

$$\begin{aligned} FS_t &= ((\prod_{d \in t} 1 + (\theta_x \varepsilon_{x,d}) + \prod_{d \in t} 1 + (\theta_r \varepsilon_{r,d})) - 1) \times 100 \\ Mkt_t &= ((\prod_{d \in t} 1 + (\theta_{r_m} \varepsilon_{r_m,d})) - 1) \times 100 \\ Noise_t &= ((\prod_{d \in t} 1 + (r_d - a_0 - \theta_x \varepsilon_{x,d} - \theta_r \varepsilon_{r,d} - \theta_{r_m} \varepsilon_{r_m,d})) - 1) \times 100 \end{aligned} \tag{1.3}$$

I do not differentiate between private and public firm-specific information as Peng and Xiong (2006) do not include such differentiation in their model.

1.3.2 Fundamentally linked firms

A growing body of empirical literature explores the gradual diffusion of information between fundamentally linked firms. Ali and Hirshleifer (2020) provide evidence that a network based on analyst co-coverage is superior to previous studies using supply-chain relationships (Cohen and Frazzini, 2008; Menzly and Ozbas, 2010), industry affiliation (Moskowitz and Grinblatt, 1999; Hoberg and Phillips, 2018; Cohen and Lou, 2012), technology expertise (Lee et al., 2019), or geographic location (Parsons, Sabbatucci and Titman, 2020) to model the link between firms. Ali and Hirshleifer (2020) specify two firms as fundamentally linked if an analyst has issued an earnings estimate on either firm within the last twelve months. They define the fundamentally linked firms' spillover (r^{FL}) on the focal firm as the weighted returns of the linked firms (J) by the number of unique analysts

by setting $\varepsilon_{r,t} = 1$.

with an estimate (n_j) as described in Equation (1.4).

$$r^{FL} = \frac{1}{\sum_{j=1}^{J} n_j} \sum_{j=1}^{J} n_j \times r_j \tag{1.4}$$

The authors show that fundamentally linked firms can predict future returns as analysts and investors react with delay to new information. There are two reasons why this definition of fundamentally linked firms is a good proxy for information spillover. First, analysts gather firm environment-related information such as firm-specific news, information about competitors, or customer-supplier relationships to form earning-per-share estimates. Estimates from the same analyst can cover firms from the same information environment (Lee and So, 2017; Kaustia and Rantala, 2020) with a high probability; hence there is an information spillover. Second, the two firms become more fundamentally linked when more analysts provide estimates for both firms (Ali and Hirshleifer, 2020).

To account for the different return components defined by Brogaard et al. (2022), I decompose the return from fundamentally linked firms' spillover into its firm-specific news (FS^{FL}), market news (Mkt^{FL}) and noise ($Noise^{FL}$) components, using the same methodology.

1.4 Data and descriptive statistics

1.4.1 Stock data and controls

The data I analyze in this paper is collected from various sources. The main sample consists of 42,789 stocks for 49 equity markets, covering the period from January 1992 to December 2021. Countries are only part of the sample in years in which they are included in the MSCI Developed or Emerging Markets Index.[7] The sample covers the following countries: Argentina, Australia, Austria, Belgium, Brazil, Canada, Chile, China, Czechia, Denmark, Egypt, Finland, France, Germany, Greece, Hong Kong, India, Indonesia, Ireland, Israel, Italy, Japan, Jordan, Korea, Kuwait, Malaysia, Mexico, Morocco, Netherlands, New Zealand, Norway, Pakistan, Philippines, Poland, Portugal, Russia, Saudi Arabia, Singapore, South Africa, Spain, Sri Lanka, Sweden, Switzerland, Taiwan, Thailand,

[7] See https://www.msci.com/market-classification for details.

Turkey, UAE, United Kingdom, and the United States.

For non-U.S. countries, the accounting data is retrieved from Refinitiv Worldscope, and the stock market data is from Refinitiv Datastream. I apply several static and dynamic screens to ensure that the sample comprises exclusively common stocks and provides the highest data quality. First, stocks are identified using Refinitiv Datastream constituent lists, particularly Refinitiv Worldscope lists, research lists, and - to eliminate survivorship bias - dead lists. Second, following Ince and Porter (2006), Griffin, Kelly and Nardari (2010), Schmidt et al. (2017), and Hanauer (2020), non-common equity stocks are eliminated through generic and country-specific static screens. Furthermore, several dynamic screens are applied to stock returns and prices to exclude erroneous and illiquid observations. Appendix A provides a detailed description of the utilized constituent lists and the associated static and dynamic screens.

The U.S. stock market data comes from the Center for Research on Security Prices (CRSP), whereas accounting data is retrieved from Compustat. I include all common equity stocks traded on NYSE, NYSE MKT (formerly: AMEX), or NASDAQ. I exclude all stocks with a CRSP share code (SHRCD) different than 10 or 11. [8]

Additionally, I include analyst and institutional ownership data for the stock data. All analyst related data is collected from Institutional Brokers' Estimate System (I/B/E/S). To identify fundamentally linked firms, I use the unadjusted detail file from I/B/E/S, which contains all estimates a permanently identified analyst made from January 1991 to December 2021. For analyst coverage, I utilize the unadjusted analyst consensus earnings estimates.[9] Institutional ownership data is from the FactSet Ownership database (formerly LionShares).

Furthermore, I require stocks to have market capitalization data for the previous month and valid news and noise spillover. Finally, a country is only part of the final sample in those months for which at least 30 stock-month observations are available after filters.[10] I end up with a total of 3,901,237 stock-month observations. Table 1.1 shows the country descriptive statistics for the stocks in the final sample.

[8] For a more detailed description of the share codes, see https://wrds-www.wharton.upenn.edu/data-dictionary/form_metadata/crsp_a_stock_msf_identifyinginformation/shrcd/.

[9] The earnings forecasts are based on the I/B/E/S unadjusted files, as they are not affected by share splits after their publication date, which could distort the results (Diether, Malloy and Scherbina, 2002).

[10] I thereby ensure that each portfolio contains more than five stocks on average.

Table 1.1
Summary statistics by country

The table presents summary statistics for the 49 countries of our global stock sample. Panel A shows the global summuary, Panel B covers all developed markets whereas Panel C covers all emerging markets. Column 1 reports the country names, and Columns 2, 3, 4, and 5 report the total, minimum, mean, and maximum number of firms per country. Columns 6 and 7 state the average mean and median size per country-month. Column 8 shows the average total size per country-month and column 9 reports these values in percentage of the respective total across countries. Size is measured as market capitalization in million USD. The last columns report the actual start date in which each country is included in our sample. The sample consists all stocks for the period between January 1992 and December 2021, with minium one analyst co-covering the stocks, and minimum 30 valid country-month observations.

	Number of firms				Size			
	Total	Min	Mean	Max	Mean	Median	Total	%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Global								
Global	42789	4613	10831	13829	3024	455	35041841	100.0
Panel B: Developed Markets								
Australia	1507	62	324	508	1870	267	685524	1.79
Austria	134	30	40	67	1937	1074	73960	0.18
Belgium	173	43	72	95	3167	537	234240	0.57
Canada	2085	169	463	706	1736	237	896202	2.42
Denmark	210	36	68	105	2891	766	168228	0.40
Finland	227	30	90	129	2045	329	187960	0.44
France	1092	177	333	417	4012	338	1400390	4.06
Germany	1015	119	289	421	3585	302	1065504	3.19
Hong Kong	593	61	174	328	4032	957	712381	1.90
Ireland	52	30	31	34	1737	188	55113	0.02
Israel	85	30	36	53	2947	1127	107253	0.06
Italy	527	31	153	212	2951	472	465782	1.33
Japan	4248	367	1297	2255	2684	690	3317449	11.05
Netherlands	243	56	91	144	4341	783	348111	1.07
New Zealand	150	30	51	71	743	302	40380	0.08
Norway	390	30	100	147	1469	224	162807	0.37
Portugal	76	30	33	39	1378	387	45942	0.05
Singapore	583	45	123	209	1767	342	209707	0.62
Spain	239	76	92	111	5298	1244	482634	1.42
Sweden	700	38	182	341	1683	236	336640	0.90
Switzerland	297	76	136	164	5973	787	863168	2.33
United Kingdom	3034	379	750	985	2819	258	2173528	7.13
United States	12367	2596	3282	4457	4907	617	15241280	45.50
DM	17660	1812	4810	6080	2730	390	13770672	86.88

Continued on next page

Table 1.1 continued

	Number of firms				Size			
	Total	Min	Mean	Max	Mean	Median	Total	%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel C: Emerging Markets								
Argentina	57	30	33	43	793	220	26393	0.02
Brazil	188	30	95	116	4319	1403	398128	0.40
Chile	122	30	41	60	2611	1301	102754	0.21
China	3131	35	1058	1902	2945	1095	3524168	4.04
Czechia	54	30	39	44	257	68	10050	0.00
Egypt	97	30	45	61	972	473	40903	0.04
Greece	264	30	81	155	863	300	63437	0.17
India	1611	34	389	775	1753	396	857295	1.60
Indonesia	327	30	83	138	1639	494	159232	0.37
Jordan	201	40	104	156	272	41	26356	0.01
Korea	1857	59	326	688	1806	473	526332	1.27
Kuwait	32	30	30	30	2399	913	71980	0.00
Malaysia	842	67	213	306	927	233	205270	0.64
Mexico	154	31	51	69	3245	1338	170212	0.38
Morocco	81	30	42	63	1318	414	54924	0.05
Pakistan	112	30	48	68	686	340	33579	0.03
Philippines	152	30	51	72	1761	915	98685	0.19
Poland	371	30	91	186	1063	261	98793	0.16
Russia	149	32	61	80	8378	1815	505952	0.39
Saudi Arabia	118	30	74	89	7169	1826	545645	0.39
South Africa	416	43	110	190	2377	884	258570	0.75
Sri Lanka	77	30	42	54	241	134	10086	0.00
Taiwan	1376	43	344	739	1539	410	537575	1.20
Thailand	620	37	150	224	1031	281	175937	0.41
Turkey	291	32	91	196	1211	373	110395	0.28
UAE	62	30	35	41	4155	1533	147613	0.12
EM	12762	178	2739	5460	1690	463	6029888	13.12

Besides the different spillover variables, I construct a set of controls according to standard definitions in the literature. I use balance-sheet data from December in year $t-1$ for the stock returns from July of year t to June of year $t+1$ as in Fama and French (1993). Following Lewellen and Nagel (2006), I calculate beta ($Beta$) as the sum of the regression coefficients of daily excess returns on the local market excess return and one lag of the local market excess return for the past 12 months. I require at least 126 observations for valid beta estimates, as in Welch (2020). Size (MV_{ln}) is a stock's log market capitalization at the end of the previous month and measured in USD, as in Fama and French (1992). Book-to-market (BM) is the ratio of the book value of equity to the market value of equity, following Rosenberg, Reid and Lanstein (1985) and Davis, Fama and

French (2000). I define the book value of equity as common equity plus deferred taxes. If no deferred taxes are given, the book value of equity equals common equity. The market value of equity is as of December $t-1$. Momentum (MOM) is the cumulative return from month $t-12$ to $t-2$ as in Fama and French (1996). Short-term reversal ($STREV$) is the lagged one-month return as in Jegadeesh (1990). Analyst coverage (ANA) is the number of sell-side analysts forecasting annual firm earnings in each month $t-1$ as in Hong, Torous and Valkanov (2007). I calculate illiquidity ($ILLIQ$) according to Amihud (2002). It is defined as the arithmetic mean for the past month of the absolute daily return divided by the product of the end-of-day stock price and the daily trading volume. Institutional ownership ($OWNER$) denotes the fraction of firm shares outstanding owned by institutional investors end of the previous month as in Hirshleifer, Hsu and Li (2013).

To derive the risk-adjusted returns, I calculate the following six-factor returns: market ($RMRF$), size (SMB), value (HML), momentum (WML), and liquidity (LIQ). I closely follow Fama and French (2017), Carhart (1997), Pástor and Stambaugh (2003), and Jensen, Kelly and Pedersen (2021). In the case of the market, I define the market factor as the value-weighted returns of all available stocks in excess of the risk-free rate. For the remaining factors, I first assign stocks into different size groups (micro, small, and large) separately for each country and month following Fama and French (2008, 2012, 2017). Large stocks are the biggest stocks that account for 90% of a country's aggregated market capitalization. Small stocks are defined as those stocks that comprise the next 7% of aggregated market capitalization (so that big and small stocks together account for 97% of the aggregated market size of a country). Microcaps comprise the remaining 3%. Next, for each factor's underlying characteristic, I sort each stock into its tercile using the country-specific 30% and 70% percentile breakpoints based on large stocks in that country. The underlying characteristics for each factor are MV for size, BM for value, MOM for momentum, and $ILLIQ$ for liquidity. Besides the size factor, each factor return is then defined as the average of the big-high-tercile and small-high-tercile minus the average of the big-low-tercile and small-low-tercile. In the case of the size factor, I subtract the average of the three value terciles of the large stocks from the average of the respective small stocks. For each of the corresponding portfolios, I define the portfolio return as the value-weighted return using the market capitalization of each stock.

1.4.2 Descriptive statistics

Table 1.2 reports the time-series averages of the number of observations, the minimum, 25th percentile, mean, median, 75th percentile, maximum, and standard deviation in the cross-section.

In Panel A, I report the descriptive statistics of the three main variables of interest, firms-specific and market news, and the noise component. The three key explanatory variables all display a negative minimum and a positive maximum, pointing towards both positive and negative spillovers for each information source. The average firm-specific news return is -1.02, and the median is -0.90, indicating, on average, a negative firm-specific news spillover. In the case of market news, the mean and median are close to each other, with positive values of 1.54 and 1.52, respectively. Similar to the firm-specific news, the average noise spillover is also negative but less pronounced, with a value of -0.24 and a median of -0.20. Panel B summarizes the stock-level characteristics of focal firms, which I later use as controls. The average beta is close to 1, with a mean of 0.96 and a median of 0.92. The median log market capitalization is 6.00, which is equal to a median market capitalization of 403 million U.S. dollars. The average book-to-market ratio is 0.75, indicating that the sample is rather composed of overvalued firms. On average, 6.29 analysts issue a firm's earnings per share forecast, with a minimum of 0 analysts. This indicates that some of the stocks are not covered over the full period of 12 months which are used to construct the network. In Panel C, I display the three variables which are used as limited-attention proxies. As institutional ownership data is only available as of April 2000 and March 2020, I limit the sample in this panel to this period. In this period, on average, 5.76 analysts covered a firm, the average firm size is 3.3 billion U.S. dollars and 22.12% of a firm a held by institutional investors. The last Panel, Panel D, includes statistics on the analyst co-coverage network. By construction, at least one analyst is co-covering two firms, while on average, 7.56 analysts are linking two firms. On the other hand, one firm has, on average, 58.53 direct peers, reaching up to 201.05 peers.

Table 1.2
Descriptive statistics

This table reports the descriptive statistics of our main variables, controls, and the global network. The panels reports the time-series average of the cross-sectional mean, standard deviation, and quantiles of each variable. Panel A describes the news and noise spillover variables. Panel B summarizes the stock-level characteristics of focal firms. In Panel C, different proxies for investors' limited attention are presented. Panel D covers general statistics on the co-coverage network. In the case of Panel A, Panel B and Panel D, the sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. For Panel C the sample starts in April 2000 and ends in March 2020 due to the availability of ownership data. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage. $Beta$ is the sum of the regression coefficients of daily excess returns on the local market excess return and one lag of the local market excess return for the past 12 months. MV_{ln} is the log of the product of the closing price and the number of shares outstanding. BM is the book-to-market ratio in June of year t, which is computed as the ratio of the book value of common equity in fiscal year $t-1$ to the market value of equity in December of year $t-1$. MOM is the cumulative return from month x$t-12$ to month $t-1$ for a given month t. $STREV$ is the short-term reversal of the focal firm for a given month t. ANA is the analyst coverage, which is the number of sell-side analysts forecasting annual firm earnings in each month t. $ILLIQ$ is the illiquidity measure of Amihud (2002), which is the average daily ratio of the absolute stock return to the dollar trading volume in month t. RET^{FL}_{t+h} is the monthly contemporaneous returns from stocks that are connected through shared analyst coverage at time $t+h$. MV is the product of the closing price and the number of shares outstanding. $OWNER$ are the holdings by all institutional investors as a fraction of the market capitalization. N is the average number of co-covering analysts per stock releasing an EPS forecast in the last 12 months. C is the average number of firms that are connected through shared analyst coverage per stock.

	N	Mean	Std	Min	P25	P50	P75	Max
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Main Variables								
FS^{FL}	10836.77	-1.02	3.96	-10.95	-3.38	-0.90	1.42	8.35
Mkt^{FL}	10836.77	1.54	1.40	-1.46	0.61	1.52	2.43	4.76
$Noise^{FL}$	10836.77	-0.24	0.67	-2.06	-0.60	-0.20	0.17	1.28
Panel B: Control Variables								
$Beta$	10836.77	0.96	0.50	-0.03	0.64	0.92	1.23	2.48
MV_{ln}	10836.77	6.08	1.85	2.09	4.79	6.00	7.27	10.65
BM	10836.77	0.75	1.63	-0.16	0.32	0.57	0.95	3.74
MOM	10836.77	15.21	84.41	-72.03	-17.42	5.26	32.40	231.72
$STREV$	10836.77	1.05	14.14	-30.34	-6.02	0.18	6.78	44.13
ANA	10836.77	6.29	6.77	0.00	1.34	3.70	8.82	29.87
$ILLIQ$	10836.77	2.17	176.08	0.00	0.00	0.00	0.03	6.58
RET^{FL}	10836.77	1.01	6.10	-13.24	-2.40	0.88	4.09	17.93
Panel C: Limited-attention Variables								
ANA	11880.07	5.76	6.49	0.00	1.01	3.16	8.12	28.81
MV	11880.07	3299.66	14312.05	8.45	142.11	493.39	1723.01	51543.57
$OWNER$	11880.07	15.01	22.12	0.00	0.10	5.82	19.34	96.80
Panel D: Spillover Variables								
N	10836.77	7.56	7.99	1.00	2.00	4.68	10.32	36.86
C	10836.77	58.53	45.81	1.91	22.71	47.31	83.36	201.05

1.5 Empirical results

Without further specifying which news component is diffusing among stocks that are fundamentally linked, Ali and Hirshleifer (2020) provide evidence that investors are sluggish in impounding news, measured as return, from peers in the network. I hypothesize that the different types of news and noise have different diffusion patterns. The noise component should not be predictive if investors can differentiate between information and noise. Further, if investors tend to process market rather than firm-specific information following the limited attention model of Peng and Xiong (2006), they tend to underreact to the firm-specific information, which leads to cross-firm predictability. I make use of the return decomposition model of Brogaard et al. (2022) to test this hypothesis, whereby I predict future stock returns using two approaches commonly used in the asset-pricing literature: portfolio sorts and the Fama and MacBeth (1973) regression.

1.5.1 Diffusion of news and noise

1.5.1.1 Portfolio sorts

Portfolio tests provide a way of using cross-sectional data to test asset pricing predictions. At the beginning of each month t different portfolio strategies are implemented, I sort each firm into country-neutral quintiles based on their individual firm-specific news, market news, and noise components from fundamentally linked firms. The first quintile ('Low') contains all firms with the lowest country-specific spillover measure, while the last quintile ('High') encompasses firms with the highest spillover from analyst co-covered firms. To analyze which return measures are positively associated with future stock returns, I compare the returns of two extreme portfolios by combining them with a short position in the 'Low' quintile and a long position in the 'High' quintile. Afterwards, I regress the return differential on different asset pricing models. All of the presented t-statistics are adjusted for serial auto-correlation using Newey and West (1987) standard errors with 12 lags following Huang, Lin and Xiang (2021) and Huang et al. (2022).

The columns 'Low', two till four, and 'High' of Table 1.3 report the excess returns over the risk-free rate of the equal-weighted, Panel A, and value-weighted, Panel B, portfolios sorted on firm-specific news, market news, and noise, respectively. The last column shows the difference in return between the fifth and the first portfolio.

Table 1.3
Portfolio sorts

This table reports the performance of the stock portfolios sorted by the firm-specific and market news as well as the noise spillover from analyst co-covered stocks. At the beginning of each month the stocks are sorted into 5 country-neutral portfolios. In Panel A reports the equal-weighted and Panel B the value-weighted monthly excess of each the of portfolios and the long minus short portfolio. All standard-errors are adjusted using Newey and West (1987). The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage.

	Low	2	3	4	High	High-Low
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Equal-Weighted						
FS^{FL}	0.225	0.648	0.873	1.067	1.333	1.108
	(0.63)	(2.12)	(2.97)	(3.64)	(4.30)	(6.84)
Mkt^{FL}	0.642	0.766	0.814	0.883	0.944	0.303
	(1.93)	(2.50)	(2.67)	(2.87)	(2.67)	(1.77)
$Noise^{FL}$	0.817	0.788	0.806	0.786	0.733	-0.084
	(2.45)	(2.56)	(2.73)	(2.69)	(2.33)	(-0.94)
Panel B: Value-Weighted						
FS^{FL}	0.225	0.524	0.704	0.734	0.798	0.574
	(0.75)	(2.11)	(3.01)	(3.20)	(3.02)	(3.42)
Mkt^{FL}	0.399	0.559	0.620	0.776	0.615	0.216
	(1.40)	(2.26)	(2.52)	(3.25)	(2.29)	(1.07)
$Noise^{FL}$	0.519	0.640	0.641	0.689	0.504	-0.015
	(1.76)	(2.57)	(2.67)	(2.98)	(1.89)	(-0.10)

The univariate sorts show that the relation between the different news measures and the portfolio returns increases monotonically across the portfolios in most cases. In the case of the equal-weighted portfolios, the firm-specific news spillover yields an average monthly excess return of 0.225% (t=0.63) in the lowest portfolio and increases to 1.333% (t=4.30) in the highest portfolio. For market news, the portfolio containing stocks with the lowest market news spillover, the excess return is 0.642% (t=1.93), monotonically increasing to 0.944% (t=2.67). The excess return of the lowest portfolio sorted on the peers' noise component is 0.817% (t=2.45) and decreases to 0.7433% (t=2.33) in a non-monotonical way. The value-weighted portfolios in Panel B yield a similar but less clear pattern in terms of monotonicity. When sorting on the firm-specific news spillover, the portfolios' excess return increases from 0.225% (t=0.75) in the 'Low' to 0.798% (t=3.02) in the 'High' portfolio. Value-weighting changes the return pattern of the market news portfolios. The results do not reveal a continuous monotonous increase as the average monthly excess return decreases from the fourth portfolio with 0.776% (t=3.25) to 0.615% (t=2.29) in the 'High' portfolio. Value-weighting noise-based portfolios results in a reverse smile return function, starting with 0.519% (t=1.76), increasing to 0.689% (t=2.98), and ending with the lowest monthly excess return among the five portfolios of 0.504% (t=1.89).

In the last column of Table 1.3, I test the hypothesis that the different return components can predict the cross-section of returns by forming a hypothetical portfolio strategy that goes short in the lowest quintile and long in the highest quintile. Only the long-short positions in the firm-specific news spillover exhibit significant positive excess returns for both the equal-weighted and the value-weighted portfolios. The equal-weighted portfolio shows a monthly excess return of 1.108% (t=6.84), whereas the value-weighted portfolio yields half the magnitude in excess returns, summing up to 0.574% (t=3.42). Therefore, both firm-specific news long-short portfolios survive the proposed t-statistics hurdle rate of 3.0 proposed by Harvey, Liu and Zhu (2016). Neither the noise component nor the market-news portfolio yield significant returns. This is consistent with the view that investors are aware of the respective information contained in the market news and noise portfolio, but they do not trade on the firms-specific news. Further, the difference in effect size and statistical significance when using the value-weighted portfolios points towards that the firm-specific news diffusion effect is stronger among small stocks.

Next, I perform time-series regressions to risk-adjust the different news portfolios' monthly excess returns with well-known asset pricing factor models. In Table 1.4, I first solely include the market risk premium as in the capital asset pricing model by Sharpe (1966) ($CAPM$). For the second factor model, I include the two additional factors, size and value, as proposed by Fama and French (1993) ($FF3$). Next, as in Carhart (1997), I augment the previous factor model with the momentum factor ($FF4$). Lastly, I additionally add the liquidity (LIQ) factor by Pástor and Stambaugh (2003) ($FF4 + LIQ$) to the time-series regression. The first two rows of Panel A and Panel B show the risk-adjusted return of the firm-specific news spillover portfolio remains significant when different asset pricing factor models are employed. Due to the significance levels of the market news and noise portfolios in Table 1.3, I concentrate my main analysis in Table 1.4 on the firm-specific news spillover. [11]

In both Panel A, presenting the equal-weighted results, and Panel B, containing the value-weighted results, the risk-adjusted returns are significant, independent of the underlying factor model. In the case of equal-weighting the long-short portfolio and employing the four-factor model augmented with liquidity, the risk-adjusted return is 1.067% (t=7.71), while the value-weighted portfolio approach results in a monthly risk-adjusted return of 0.558% (t=3.56). The different weighting methodologies indicate that rather small firms are driving the slow firm-specific news diffusion. This pattern is stable among the other three-factor models which yield similar risk-adjusted returns. By further investigating the long and short sides of the equal- and value-weighted portfolios, I observe that both sides contribute towards the portfolio strategy's performance. While the low firm-specific news portfolio yields statistically negative alphas (underperformance), the long side yields statistically positive alphas (overperformance). I further analyze if the portfolio strategy is exposed to certain risk factors. For brevity reasons, I stick to the most comprehensive factor model as presented in columns 10 to 12. Both the equal-weighted as well as the value-weighted portfolios load significantly on the market, indicating that the portfolios are not well diversified concerning market risk. For the equal-weighting approach, I can further identify a positive relation between the firm-specific portfolio and two risk factors, value and momentum. In the case of value-weighting, none of the other risk factors yield significant exposure.

[11] For completeness, I attach the risk-adjusted results of the two other return components in Table A. 8 and Table A. 9, yielding similar results as in Table 1.3.

Table 1.4
Portfolio time-series regression of FS^{FL}

This table reports the Jensen's alpha of the four different regression models and long, short, and long–short portfolio based on quintiles of FS^{FL} using country-neutral break points. Each regression model includes a different set of tradable common risk factors $MKTRF$, SMB, HML, UMD, and LIQ (Fama and French, 1993; Carhart, 1997; Pástor and Stambaugh, 2003) Panel A shows the results for the equal- weighted portfolios and Panel B shows the results for the value-weighted counterpart. t-statistics are reported in parentheses. All standard-errors are adjusted using Newey and West (1987). The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage.

	$CAPM$			$FF3$			$FF4$			$FF4+LIQ$		
	Low	High	H-L	Low	High	H-L	Low	High	H-L	Low	High	H-L
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: Equal-Weighted												
α	-0.450	0.753	1.203	-0.502	0.695	1.198	-0.356	0.709	1.066	-0.355	0.712	1.067
	(-3.22)	(4.89)	(8.00)	(-4.42)	(6.01)	(7.04)	(-3.64)	(5.95)	(7.61)	(-3.69)	(6.41)	(7.71)
$Mkt-RF$	1.241	1.065	-0.176	1.153	1.002	-0.151	1.111	0.998	-0.113	1.095	0.968	-0.127
	(27.83)	(31.41)	(-5.02)	(36.57)	(40.61)	(-4.65)	(35.62)	(45.03)	(-3.76)	(31.47)	(41.63)	(-3.54)
SMB				-1.010	-0.780	0.230	-0.973	-0.776	0.197	-0.993	-0.815	0.178
				(-9.26)	(-8.91)	(1.68)	(-10.16)	(-8.44)	(1.55)	(-10.35)	(-9.28)	(1.30)
HML				0.100	0.143	0.043	0.003	0.133	0.130	0.039	0.202	0.163
				(1.06)	(2.66)	(0.32)	(0.05)	(2.69)	(1.62)	(0.51)	(4.94)	(2.01)
MOM							-0.211	-0.020	0.191	-0.213	-0.023	0.189
							(-5.24)	(-0.31)	(2.16)	(-5.00)	(-0.39)	(2.26)
LIQ										0.097	0.185	0.088
										(1.08)	(2.92)	(0.79)

Continued on next page

Table 1.3 continued

	$CAPM$			$FF3$			$FF4$			$FF4+LIQ$		
	Low	High	H-L	Low	High	H-L	Low	High	H-L	Low	High	H-L
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel B: Value-Weighted												
α	-0.394	0.265	0.659	-0.378	0.266	0.645	-0.314	0.243	0.557	-0.314	0.244	0.558
	(-4.19)	(3.10)	(4.00)	(-3.84)	(2.77)	(3.60)	(-3.13)	(3.26)	(3.58)	(-3.13)	(3.32)	(3.56)
$Mkt-RF$	1.138	0.980	-0.157	1.128	0.988	-0.140	1.110	0.995	-0.115	1.103	0.982	-0.120
	(39.30)	(41.28)	(-3.60)	(43.22)	(45.49)	(-3.53)	(40.24)	(39.35)	(-2.63)	(32.83)	(34.35)	(-2.24)
SMB				-0.051	0.079	0.130	-0.035	0.073	0.108	-0.044	0.057	0.102
				(-0.68)	(1.18)	(0.96)	(-0.47)	(1.02)	(0.78)	(-0.54)	(0.71)	(0.64)
HML				-0.065	0.002	0.067	-0.107	0.018	0.125	-0.090	0.047	0.137
				(-0.78)	(0.03)	(0.47)	(-1.55)	(0.38)	(1.22)	(-1.17)	(0.95)	(1.20)
MOM							-0.092	0.034	0.126	-0.093	0.032	0.126
							(-1.99)	(0.49)	(1.17)	(-1.99)	(0.49)	(1.19)
LIQ										0.046	0.077	0.032
										(0.50)	(0.85)	(0.19)

Next, I investigate the long-run effect of the different spillover portfolios. To analyze the long-run buy-and-hold returns, I follow the methodology of Smajlbegovic (2019). First, I identify the stocks used for constructing the different long-short portfolios, and afterward, I calculate their equal and value-weighted excess returns in the subsequent months $t + k$, where $k \in \{0, ..., 24\}$. Then, I run a time-series regression for each holding period month k on the four-factor model augmented with liquidity. The corresponding intercept of the regression is then the alpha of the buy-and-hold strategy k months after the portfolio formation. Lastly, I add up these intercepts, which results in the cumulative alpha in month k - the variable of interest.

Figure 1.1 presents the equal-weighted (Panel A) and the value-weighted (Panel B) cumulative five-factor alpha over a holding period of 24 months.

The graphical representation in Panel A reveals that only the firm-specific news portfolio does not revert in the long run. I can identify a strong increase in the risk-adjusted return in the first 6 months after portfolio construction, yielding an alpha of 2.065%. In the subsequent 18 months, the alpha slightly increases to 2.970%. Both the market news as well as the noise component do not yield any positive long-term risk-adjusted return. Looking at the value-weighted results as presented in Panel b, only the firm-specific news portfolio yields a positive long-term risk-adjusted return. Comparing the pattern of the individual news and noise portfolios reveals that the news measure of Ali and Hirshleifer (2020) is only similar to the firm-specific news in the first two months. Afterward, the news measure of Ali and Hirshleifer (2020) rather correlates with the market news and the noise component.

1.5.1.2 Cross-sectional regressions

A common way to test return predictability is the Fama and MacBeth (1973) procedure. The regression method is frequently used to control for time effects as it accounts for dependence in the time dimension (Petersen, 2009) and other potential drivers that could affect the return predictability.

For each month, I run a cross-sectional regression including all stock with news and

Figure 1.1

Long-horizon performance of news and noise spillover

This figure shows the average cumulative risk-adjusted return of the different spillover long-short portfolios. Panel A shows equal-weighted risk-adjusted returns while Panel B shows value-weighted risk-adjusted returns. First, I obtain the return of the portfolio formed at the end of month t for month t+k, where k $\in \{1, ..., 24\}$. Second, I run a time-series regression with the full five-factor model, including $MKTRF$, SMB, HML, UMD, and LIQ (Fama and French, 1993; Carhart, 1997; Pástor and Stambaugh, 2003) for the corresponding months. The regression intercept is defined as the average risk-adjusted portfolio return for the long–short portfolio at month t + k. In the final step, I compute the average holding period (cumulative) risk-adjusted return for the next k months since formation as the sum over the previous k months.

Panel A: Equal-Weighted

Panel B: Value-Weighted

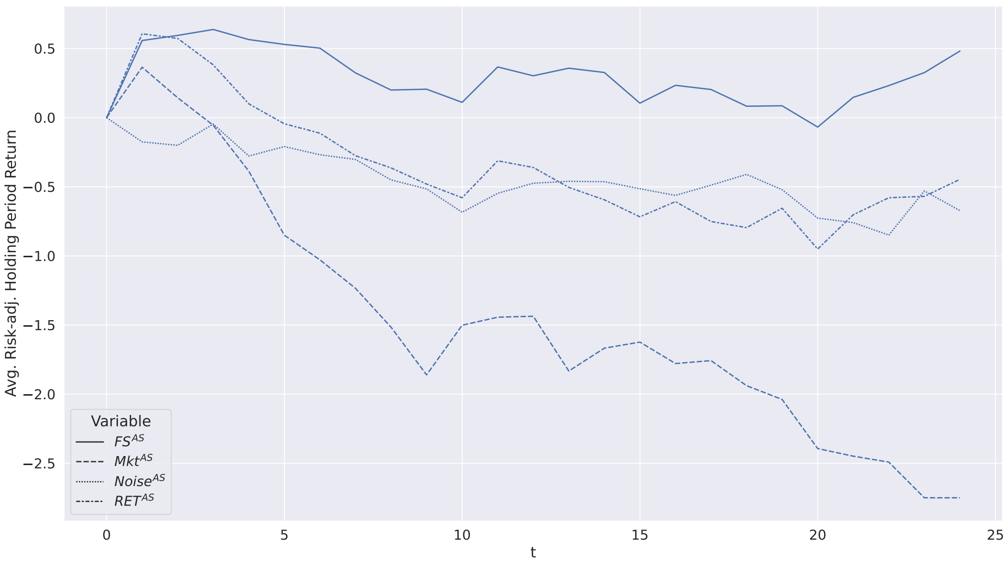

noise spillover and with future monthly excess returns as the dependent variable:

$$RET_{i,t+h} = a_{t+h} + \lambda_{FS} FS^{FL}_{i,t-1} + \lambda_{Mkt} Mkt^{FL}_{i,t-1} + \lambda_{Noise} Noise^{FL}_{i,t-1} + \lambda_c \boldsymbol{C} + \epsilon_{i,t+h}, \quad (1.5)$$

where $FS^{FL}_{i,t-1}$, $Mkt^{FL}_{i,t-1}$, and $Noise^{FL}_{i,t-1}$ are the firm.specific news, market news, and noise spillover components of the previous month, and $\boldsymbol{C}$ represents the vector of control variables. The reported regression coefficients are determined by calculating the time-series averages. Further, the corresponding t-statistics are adjusted using the Newey and West (1987) procedure with 12 lags. In Table 1.5, I report the regression results for different setups. In columns 1-3, I regress on the next month's excess return (h=0), whereas the dependent variable in columns 4-5 and 6-7 is the second (h=1) or third (h=2) month after the spillover, respectively.

For the results presented in column 1, I only include $FS^{FL}_{i,t-1}$, $Mkt^{FL}_{i,t-1}$, and $Noise^{FL}_{i,t-1}$ as well as industry and country dummies as independent variables. The industry dummies are included to mitigate concerns that the diffusion of industry news documented by Hou (2007) explains the news and noise spillover. Additionally, country dummies are included to ensure that specific countries do not drive the results. Similar to the portfolio results, also in the most simple setup without any further controls, only the firm-specific news from peers is positive and significant, yielding a regression coefficient of 0.121 (t=9.76). The market news and the noise spillover do not yield a statistically significant coefficient.

I include a standard set of controls for the second specification as presented in column 2. The controls are focal stocks' beta, log market capitalization, book-to-market ratio, momentum, short-term reversal, the number of analyst forecasts in the consensus of the current fiscal year, and the Amihud illiquidity measure. Including this set of controls allows me to measure the spillover effects more precisely, resulting in a small increase of the λ_{FS} to 0.129 (t=10.57). I further document that the regression coefficient of the market news stays positive and additionally turns significant, yielding a λ_{Mkt} of 0.117 (t=2.65). Besides the market beta, all other controls yield significant coefficients with the expected direction. This is in line with previous findings in the literature, which document that the market beta has only marginal predictive power for returns.

Following Burt and Hrdlicka (2021), I add the contemporaneous return of the funda-

Table 1.5
Peers news and noise and stock returns

This table reports the estimated regression coefficients and Newey-West t-statistics (in parentheses) from Fama-MacBeth cross-sectional regressions predicting one-month ahead excess stock returns. The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage. $Beta$ is the sum of the regression coefficients of daily excess returns on the local market excess return and one lag of the local market excess return for the past 12 months. MV_{ln} is the log of the product of the closing price and the number of shares outstanding. BM is the book-to-market ratio in June of year t, which is computed as the ratio of the book value of common equity in fiscal year $t-1$ to the market value of equity in December of year $t-1$. MOM is the cumulative return from month $t-12$ to month $t-1$ for a given month t. $STREV$ is the short-term reversal of the focal firm for a given month t. ANA is the analyst coverage, which is the number of sell-side analysts forecasting annual firm earnings in each month t. $ILLIQ$ is the illiquidity measure of Amihud (2002), which is the average daily ratio of the absolute stock return to the dollar trading volume in month t. RET^{FL}_{t+h} is the monthly contemporaneous returns from stocks that are connected through shared analyst coverage at time $t+h$.

	$R_{i,t}$			$R_{i,t+1}$		$R_{i,t+2}$	
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
FS^{FL}	0.121	0.129	0.089	0.046	0.027	0.035	0.011
	(9.76)	(10.57)	(12.18)	(6.36)	(5.76)	(4.53)	(2.77)
Mkt^{FL}	0.073	0.117	0.081	0.038	0.020	0.021	-0.005
	(1.07)	(2.65)	(4.09)	(0.85)	(0.73)	(0.39)	(-0.17)
$Noise^{FL}$	-0.015	0.044	0.091	0.057	0.027	0.019	-0.002
	(-0.39)	(1.48)	(4.73)	(1.86)	(1.32)	(0.50)	(-0.06)
RET^{FL}_{t+h}			0.536		0.538		0.536
			(47.12)		(47.29)		(46.20)
$Beta$		-0.206	-0.186	-0.235	-0.213	-0.182	-0.177
		(-1.24)	(-1.22)	(-1.47)	(-1.50)	(-1.21)	(-1.31)
MV_{ln}		-0.122	-0.119	-0.071	-0.070	-0.055	-0.056
		(-2.38)	(-2.78)	(-1.46)	(-1.77)	(-1.14)	(-1.43)
BM		0.276	0.266	0.266	0.257	0.277	0.264
		(4.80)	(5.69)	(4.84)	(5.84)	(4.93)	(5.91)
MOM		0.007	0.006	0.006	0.005	0.004	0.003
		(3.10)	(2.98)	(2.78)	(2.70)	(1.75)	(1.66)
$STREV$		-0.029	-0.032	0.004	0.004	0.014	0.013
		(-6.19)	(-6.56)	(1.49)	(1.19)	(6.19)	(6.27)
ANA		0.027	0.026	0.019	0.018	0.017	0.016
		(4.57)	(4.99)	(3.45)	(3.81)	(3.27)	(3.65)
$ILLIQ$		-0.084	-0.082	0.043	0.043	0.072	0.075
		(-3.27)	(-3.20)	(1.47)	(1.62)	(2.32)	(2.64)
FF-38	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2 (%)	14.51	16.50	18.67	16.09	18.28	15.94	18.12
Avg. Obs	10837	10837	10837	10805	10804	10769	10769

mentally linked stocks in column 3 to the regression setup. This proxy for the common momentum component helps to isolate the actual delay in information diffusion from the different spillover components. By including the contemporaneous return, the coefficients of the two news components of the spillover decrease by about 31% to 0.089 (t=12.18) in the case of the firm-specific news. The respective coefficients of the market news spillover decrease to 0.081 (t=4.09). Adding this proxy to the specification nearly doubles the magnitude of the noise coefficient to 0.091 (t=4.73) and turns the λ_{Noise} significant.

In columns 4 and 5, I test if the different news and noise components are still diffusing even two months after they initially occurred. The difference between columns 4 and 5 is that I further control for the common momentum components. In both setups, the λ_{Mkt} and λ_{Noise} lose their predictive power. The coefficients amount to 0.038 (t=0.85) and 0.057 (t=1.86) in column 4 and 0.020 (t=0.73) and 0.027 (t=1.32) when including the contemporaneous return. In the case of firm-specific news, the coefficient stays positive and significant in both specifications. Comparing the coefficient in both specifications, I document a decrease from 0.046 (t=6.36) to 0.027 (t=5.76) when including the contemporaneous return.

In the last two columns, 6 and 7, I regress the news and noise components on the three-month ahead monthly excess return. Similar to the two-month-ahead return, both market news and noise are insignificant. While in column 6, the firm-specific news component takes a value of 0.035 (t=4.35), and it decreases to 0.011 (t=2.77) in column 7, not passing the proposed t-statistics hurdle rate of 3.0 by Harvey, Liu and Zhu (2016).

These results are partially in line with the categorical learning behavior and the inattention of investors as proposed by Peng and Xiong (2006). At the one-month horizon, both market and firm-specific news are positive and significant, indicating that investors do not directly process both components, resulting in delayed information diffusion. At the two-month horizon, in line with the model of Peng and Xiong (2006), the market news turns insignificant, proving that investors tend to process more market news than firm-specific news. By controlling for the delay in information diffusion at the three-month horizon, the firm-specific news turns insignificant, indicating that all information from the peers is fully diffused in the focal firms' stock price.

Hou, Xue and Zhang (2018) advocates for the use of weighed-least squares by using the market capitalization of each stock as the weight to estimate the regression coefficients. This helps in controlling for the different behavior of small stocks compared to large stocks, as the high number of small firms would be the driver of the coefficient estimates. To further investigate the different diffusion patterns among small and large stocks, I estimate two sets of regressions, one for large capitalization stocks and one for small capitalization stocks. The results are presented in Table 1.6. Panel A covers the regression results of the large stocks sample, and Panel B of the small stocks. These sample splits by size additionally proxy for investor's attention (Hirshleifer and Teoh, 2003) and allows me to investigate further the information processing model of Peng and Xiong (2006). If investors tend to process more market than firm-specific information and if they have to decide on which firms they spend processing capabilities on, I expect different behaviors among the two samples. An investor should first spend his capabilities on large firms, which in the first place would allow him to differentiate between news and noise and further directly process the diffusing market news. If the investor then has remaining capabilities left, he will subsequently spend his attention on the small stocks.

In columns 1 and 2, I regress the linked firm news and noise components on the next month's return. I can already identify different diffusion patterns among the two samples in these two columns. While for the large stocks, only the firm-specific news is relevant when including the contemporaneous return, in the case of small stocks, also the market news diffuses into the stock price. This is in line with the mechanism proposed above. In the second month, after the news occurred, as reported in columns 3 and 4, the firm-specific news coefficient of the large stocks sample also lost its predictive power, indicating that investors now also have processed this type of news. In the case of small stocks, only the firm-specific news is still predictive, indicating that the firm-specific news is not yet fully diffused. The last two columns, 5 and 6, of Table 1.6 report the coefficients three months after the different news arrived at the peer stock. In this regression setting, all news and noise coefficients of the large stock sample are still insignificant, while the firm-specific news is still diffusing into the stock price of the small stocks.

These results underline that, on the one hand, investors process different news components differently, and on the other hand, this diffusion process varies among stocks to which investors pay attention.

Table 1.6
Peers news and noise diffusion by size

This table reports the estimated regression coefficients and Newey-West t-statistics (in parentheses) from Fama-MacBeth cross-sectional regressions predicting one-month, two-month, and three-month ahead excess stock returns. Panel A estimates the coefficients for large stocks whereas Panel B focuses on small stocks. I follow Fama and French (2012, 2017) to calculate the country specific breakpoints. Large stocks are defined as the largest stocks which together account for 90% of a country's aggregated market capitalization. Small stocks are defined as those stocks that comprise the remaining 10% of aggregated market capitalization. The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage. RET^{FL}_{t+h} is the monthly contemporaneous returns from stocks that are connected through shared analyst coverage at time $t+h$.

	$R_{i,t}$		$R_{i,t+1}$		$R_{i,t+2}$	
Variables	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Large Stocks						
FS^{FL}	0.099	0.055	0.030	0.010	0.026	-0.011
	(8.18)	(8.32)	(3.10)	(1.51)	(2.64)	(-1.84)
Mkt^{FL}	0.129	0.061	0.023	-0.032	-0.018	-0.063
	(2.31)	(1.79)	(0.42)	(-0.77)	(-0.31)	(-1.78)
$Noise^{FL}$	0.028	0.071	0.043	0.007	0.034	0.001
	(0.51)	(2.06)	(0.82)	(0.21)	(0.60)	(0.02)
RET^{FL}_{t+h}		0.747		0.746		0.745
		(62.30)		(64.01)		(65.85)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
FF-38	Yes	Yes	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2 (%)	27.14	31.83	26.80	31.48	26.70	31.36
Avg. Obs	4135	4135	4127	4127	4117	4117
Panel B: Small Stocks						
FS^{FL}	0.136	0.103	0.054	0.037	0.041	0.023
	(11.11)	(12.51)	(7.21)	(7.20)	(4.77)	(4.22)
Mkt^{FL}	0.103	0.078	0.026	0.022	0.039	0.020
	(2.46)	(3.40)	(0.59)	(0.76)	(0.68)	(0.53)
$Noise^{FL}$	0.065	0.103	0.066	0.045	0.012	-0.004
	(1.81)	(3.53)	(1.66)	(1.39)	(0.32)	(-0.12)
RET^{FL}_{t+h}		0.405		0.408		0.406
		(40.96)		(39.85)		(37.68)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
FF-38	Yes	Yes	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2 (%)	13.58	14.76	13.11	14.30	12.91	14.09
Avg. Obs	6702	6702	6678	6678	6652	6652

1.5.1.3 Robustness

I run a series of robustness tests as presented in Table 1.7 to examine the consistency of the relation between news and noise diffusion and stock returns. First, I exclude all stocks with a price below the country's 10th percentile following Landis and Skouras (2021) to ensure that the results are not driven by small illiquid stocks. Second, I limit the sample period to the time between April 2000 and March 2020. On the hand, this period is in line with the limited attention sample period and on the other hand, allows to evaluate if investor behavior has changed over time, e.g. if they process different return components faster. Third, the sample is limited to focal firms from the United States, similar to the sample of Ali and Hirshleifer (2020). Lastly, I apply the method of Kaustia and Rantala (2020) to account for random connections between stocks.[12]

In Panel A, I report the coefficients of the one-month and two-month-ahead returns. Excluding the smallest and most illiquid stocks from the global sample has the largest implications of all robustness checks on the noise coefficients. While in column 3 of Table 1.5, the noise component turned significant by including the peers' contemporaneous returns, the λ_{Noise} now stays insignificant, indicating that the noise component only diffuses into the stock price of rather small stocks. The market and firm-specific news coefficients are similar to the baseline setup and robust to the price filter. Shortening the sample period to the period between April 2000 and March 2020 reduces the magnitude of coefficients. One reason for this could be that the spillover has become less important in recent years, or investors pay more attention to it. Nevertheless for both return horizons, the λ_{FS}, λ_{Mkt}, and λ_{Noise} are significant for the same setups as in Table 1.5. Next, I explicitly limit the sample to stocks from the United States. Without including the contemporaneous returns, only the firm-specific news stay significant. Similar to the previous robustness test, this could be again driven by two different reasons. First, the market news is not important when predicting the future return of a U.S. stock. The second and the more likely reason is that investors in the U.S are aware of the market news, and therefore, they directly diffuse into the focal firms' stock price. Lastly, I apply the methodology of Kaustia and Rantala (2020). In doing so, I can account for randomness in the links between firms and further

[12] Kaustia and Rantala (2020) state that unrelated firms may have common analysts just by chance because of analysts who cover several industries or groups of unrelated firm.

Table 1.7
Robustness tests

This table reports the estimated regression coefficients and Newey-West t-statistics (in parentheses) from Fama-MacBeth cross-sectional regressions predicting one-month ahead excess stock returns. Robustness checks include: (1), (2), prices above the countries 25th percentile following Landis and Skouras (2021). (3), (4), sample period between April 2000 and ends in March 2020. (5), (6), only stocks from the U.S. (7), (8), reduced spillover sample following Kaustia and Rantala (2020). The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage. RET^{FL}_{t+h} is the monthly contemporaneous returns from stocks that are connected through shared analyst coverage at time $t+h$.

	PRC		Apr 00 - Mar 20		U.S.		$-C$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: $R_{i,t}$								
FS^{FL}	0.123	0.080	0.099	0.074	0.165	0.112	0.074	0.072
	(9.46)	(10.60)	(9.62)	(10.49)	(8.06)	(8.57)	(9.58)	(11.05)
Mkt^{FL}	0.152	0.114	0.116	0.058	0.089	0.050	0.081	0.079
	(3.37)	(4.29)	(2.11)	(2.59)	(1.19)	(1.12)	(3.12)	(4.65)
$Noise^{FL}$	-0.041	0.018	0.029	0.089	0.093	0.139	0.024	0.048
	(-1.20)	(0.80)	(0.83)	(3.98)	(1.11)	(2.35)	(1.28)	(3.04)
RET^{FL}_{t+h}		0.578		0.532		0.551		0.395
		(53.37)		(49.36)		(38.24)		(58.56)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FF-38	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Adj. R^2 (%)	18.89	21.50	15.89	18.07	9.01	10.66	19.22	22.53
Avg. Obs	9042	9042	11880	11880	3282	3282	6313	6313
Panel B: $R_{i,t+1}$								
FS^{FL}	0.048	0.025	0.037	0.019	0.068	0.041	0.021	0.015
	(5.74)	(4.88)	(4.56)	(4.57)	(5.03)	(4.80)	(5.11)	(4.50)
Mkt^{FL}	0.045	0.027	-0.028	-0.007	-0.001	-0.020	0.010	0.009
	(0.97)	(0.95)	(-0.55)	(-0.28)	(-0.01)	(-0.44)	(0.37)	(0.44)
$Noise^{FL}$	0.063	0.028	0.049	0.021	0.016	-0.012	0.038	0.023
	(2.16)	(1.20)	(1.26)	(0.88)	(0.16)	(-0.18)	(2.21)	(1.90)
RET^{FL}_{t+h}		0.568		0.534		0.552		0.393
		(55.88)		(52.09)		(37.40)		(56.50)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FF-38	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Adj. R^2 (%)	18.11	20.78	15.56	17.75	8.50	10.18	18.79	22.09
Avg. Obs	9018	9018	11845	11845	3264	3264	6299	6297

confirm that the baseline results are not driven by stocks that are randomly co-covered by the analysts. Indeed, the results as presented in columns 7 and 8, yield similar coefficients as in column 2 and 3 of Table 1.5.

Panel B includes the regression results of the different robustness measures on the two-month ahead returns. It underlines the robustness of the results of Table 1.5. Only the diffusing firm-specific news yields a positive and significant coefficient in every setting.

Overall the robustness tests indicate that the previously reported results are not driven by illiquid stocks, the selected sample period, the underlying countries, or the links between the firms themselves.

1.5.2 Underreaction and limited attention

The current literature on the news diffusion from fundamentally linked firms poses that investors underreact due to limited attention (Cohen and Frazzini, 2008; Lee et al., 2019; Ali and Hirshleifer, 2020). In this section, I will test the firm-specific and market news as well as the noise component regarding this behavior.

Cohen and Frazzini (2008) propose a setup to test whether the focal firms underreact or overreact to the diffusing news. Similar to the approach presented in Table 1.3, I form long-short country-neutral calendar-time portfolios based on the diffusing news and noise component in month $t-1$. RET_{t-1} represents the return of the focal firm in the same month in which the shock arrived at the peers in month $t-1$, and $RET_{t,t+h}$ is the focal firm cumulative return over the subsequent h months. In the last row of each panel, I report the underreaction coefficient ($URC_{t-1,t+h}$), which is a measure of the focal firms' initial price response to a given shock as a fraction of the subsequent cumulative return. As in Cohen and Frazzini (2008), I define $URC_{t-1,t+h}$ as:

$$URC_{t-1,t+h} = \frac{RET_{t-1}}{RET_{t-1} + RET_{t,t+h}}. \tag{1.6}$$

If the market reacts efficiently to the news, the underreaction coefficient would equal one. If the focal firms underreact to the shock, the fraction is larger than zero and smaller than one. On the other hand, a value larger than one indicates that the focal firm overreacts to the shock. In Table 1.8, I report the value-weighted underreaction coefficients and the corresponding portfolio returns. [13] To be not biased by the holding periods, I vary the time frame of the holding periods between 3, 6, and 9 months. Based on the previous results, I first test in Panel A of Table 1.8 the full sample and afterward in Panel B all large stocks and Panel C all small stocks on the existence of underreaction.

The results in Table 1.8 indicate that the focal firm is underreacting to the arrival of firm-specific news. In Panel A, the underreaction coefficient for a holding period of six months is 0.818 (t=2.90), which is equal to an underreaction to firm-specific news from linked firms by roughly 18.2%. The coefficient drops even further by limiting the sample to small stocks. Depending on the holding period, the underreaction coefficient varies between 0.785 (t=6.28) and 0.674 (t=5.29). In the case of market news, I cannot identify the underreaction behavior of the focal firm, whereas, in the case of the noise component, only small stocks tend to underreact, yielding a coefficient between 0.557 (t=7.66) and 0.843 (t=4.68).

Lastly, I test if this underreaction is driven by the limited attention of investors. Prior literature mentions three different measures for investor attention. For example, Cohen and Frazzini (2008), Menzly and Ozbas (2010), Hirshleifer, Hsu and Li (2013), and Jiang, Qian and Yao (2016) use analyst stock coverage, firm size, and institutional ownership as a proxy for investor attention for an underlying stock. Due to data availability of institutional ownership data, the sample period starts in April 2000 and ends in March 2020. If the diffusion of firm-specific news is related to limited attention, I should be able to measure a more substantial effect for firms that attract less investor attention (Lee et al., 2019). To test this hypothesis, I interact the firm-specific news with different dummy variables, which proxy if the underlying stock is exposed to limited attention. The dummy variable takes the value of one if the underlying stock characteristic is above the cross-sectional median, and zero otherwise. In the first three regression specifications, I use analyst coverage (ANA), size (MV), and institutional ownership ($OWNER$) as limited attention proxies. In the fourth regression specification, I construct a composite measure that is equal to one if the cross-sectional average of the other dummies is larger than 0.5. In case of the existence of limited investor attention, I expect a negative sign on the different interaction terms and, therefore, a lower predictive power of the underlying news

[13] In Appendix A. 9, I report the results when equal-weighting the portfolios.

Table 1.8
Underreaction coefficients

This table shows returns on the firm-specific and market news as well as noise spillover portfolio and the corresponding underreaction coefficients. At the beginning of every month, stocks are ranked in ascending order based on the corresponding spillover at the end of the previous month. At the beginning of each month the stocks are sorted into 5 country-neutral portfolios. All stocks are value-weighted within a given portfolio, and the portfolios are rebalanced every calendar month to maintain value weights Panel A bases the analysis on all stocks, Panel B focuses on large stocks whereas Panel c focuses on small stocks. I follow Fama and French (2012, 2017) to calculate the country specific breakpoints. Large stocks are defined as the largest stocks which together account for 90% of a country's aggregated market capitalization. Small stocks are defined as those stocks that comprise the remaining 10% of aggregated market capitalization. Each panel reports the average cumulative returns on the long–short portfolio formed on the respective spillover in month t. RET_t is the focal firm stock return in month t. $RET_{t+1,t+h}$ is the cumulative return over the subsequent h, for $h \in \{3, 6, 9\}$, months. URC (underreaction coefficient) is defined as the fraction of total returns from month t to month $t+h$ that occurs in month t ($URC = RET_t/(RET_t + RET_{t+1,t+h})$). t-statistics are shown below the coefficient estimates. In the case of URC the t-statistics represent the distance of the coefficient from one, which is the case of no underreaction. The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage.

	FS^{FL}			Mkt^{FL}			$Noise^{FL}$		
	$h=3$	$h=6$	$h=9$	$h=3$	$h=6$	$h=9$	$h=3$	$h=6$	$h=9$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: All Stocks									
RET_t	6.866	6.862	6.859	2.590	2.612	2.633	0.465	0.459	0.472
	(46.88)	(46.53)	(46.21)	(15.35)	(15.41)	(15.44)	(3.59)	(3.52)	(3.60)
$RET_{t+1,t+h}$	1.074	1.530	1.649	0.291	0.072	0.059	0.017	-0.035	-0.266
	(4.62)	(4.79)	(4.23)	(0.98)	(0.17)	(0.11)	(0.09)	(-0.13)	(-0.81)
URC	0.865	0.818	0.806	0.899	0.973	0.978	0.964	1.083	2.300
	(2.82)	(2.90)	(3.02)	(1.59)	(0.39)	(0.34)	(0.59)	(1.76)	(31.84)
Panel B: Large Stocks									
RET_t	6.760	6.755	6.752	2.566	2.592	2.613	0.443	0.437	0.451
	(46.29)	(45.94)	(45.64)	(15.16)	(15.28)	(15.31)	(3.33)	(3.27)	(3.36)
$RET_{t+1,t+h}$	0.965	1.430	1.503	0.295	0.049	0.011	0.045	-0.053	-0.346
	(4.09)	(4.33)	(3.77)	(1.00)	(0.11)	(0.02)	(0.22)	(-0.20)	(-1.01)
URC	0.875	0.825	0.818	0.897	0.982	0.996	0.908	1.137	4.290
	(2.37)	(2.30)	(2.62)	(1.78)	(0.27)	(0.06)	(1.57)	(3.05)	(74.99)
Panel C: Small Stocks									
RET_t	7.436	7.434	7.426	2.810	2.844	2.867	0.290	0.271	0.271
	(44.80)	(44.45)	(44.12)	(13.59)	(13.71)	(13.74)	(2.31)	(2.15)	(2.14)
$RET_{t+1,t+h}$	2.033	2.805	3.587	0.172	-0.095	0.012	0.231	0.201	0.051
	(8.45)	(8.36)	(8.30)	(0.46)	(-0.19)	(0.02)	(1.17)	(0.73)	(0.15)
URC	0.785	0.726	0.674	0.942	1.034	0.996	0.557	0.575	0.843
	(6.28)	(5.41)	(5.29)	(1.09)	(0.66)	(0.08)	(7.66)	(9.06)	(4.68)

components. Specifically, the level of investor attention rises if the focal firm is covered by many analysts, is big, or has a high share of institutional ownership which should lead to accelerated news incorporation and, therefore, lower predictive power. In Table 1.9, I present the results of the test for limited investor attention on the diffusion of firm-specific news.[14]

Table 1.9
Limited attention

This table reports the estimated regression coefficients and Newey-West t-statistics (in parentheses) from Fama-MacBeth cross-sectional regressions predicting one-month ahead excess stock returns with limited attention proxies. I interact the firm-specific news spillover with limited attention dummies. The indicator variables that take the value of one if the underlying variable is above the median in the cross-section, and zero otherwise. The sample consists of all stocks for the period between April 2000 and ends in March 2020 , with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage. RET^{FL}_{t+h} is the monthly contemporaneous returns from stocks that are connected through shared analyst coverage at time $t+h$. ANA is the analyst coverage, which is the number of sell-side analysts forecasting annual firm earnings in each month t. MV is the product of the closing price and the number of shares outstanding. $OWNER$ are the holdings by all institutional investors as a fraction of the market capitalization. $COMP$ is equal to 1 if cross-sectional average of the dummies of ANA, MV, and $OWNER$ is larger than 0.5.

	ANA		MV		$OWNER$		$COMP$	
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$LA \times FS^{FL}$	-0.051	-0.042	-0.085	-0.070	-0.006	-0.004	-0.061	-0.051
	(-3.86)	(-3.93)	(-7.28)	(-7.39)	(-0.59)	(-0.47)	(-4.33)	(-4.35)
FS^{FL}	0.116	0.088	0.133	0.101	0.102	0.076	0.122	0.093
	(12.57)	(13.66)	(12.76)	(14.82)	(10.60)	(10.61)	(13.14)	(14.57)
Mkt^{FL}	0.115	0.058	0.115	0.057	0.115	0.057	0.111	0.054
	(2.09)	(2.57)	(2.13)	(2.56)	(2.12)	(2.57)	(2.05)	(2.38)
$Noise^{FL}$	0.031	0.090	0.019	0.080	0.028	0.088	0.009	0.072
	(0.86)	(4.00)	(0.52)	(3.62)	(0.79)	(3.85)	(0.24)	(3.21)
LA	0.177	0.163	-0.098	-0.105	-0.050	-0.063	0.056	0.031
	(2.82)	(2.81)	(-0.83)	(-1.01)	(-1.42)	(-1.84)	(0.48)	(0.30)
RET^{FL}_{t+h}		0.531		0.533		0.531		0.537
		(49.45)		(49.39)		(49.31)		(50.00)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FF-38	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2 (%)	15.93	18.09	15.85	18.03	15.94	18.10	15.74	17.96
Avg. Obs	11880	11880	11880	11880	11880	11880	11880	11880

Columns 1 and 2 include analyst coverage as the first limited attention proxy. In

[14] In Appendix A. 10, I report the results when interacting the limited attention proxies with each return component individually.

both specifications, with and without the contemporaneous returns, the coefficient of the interaction effect is negative and significant, amounting to -0.051 (t=-3.86) and -0.042 (t=-3.93), respectively. The next pair of results in columns 3 and 4 use the firms' market capitalization as a proxy for investors' limited attention. Similar to the first two regression specifications is the interaction coefficient negative and significant. The coefficient without controlling for common momentum equals -0.085 (t=-7.28) and is reduced to -0.070 (t=-7.39) when including the contemporaneous return. By using institutional ownership as a proxy, in columns 5 and 6, the coefficient is negative but not statistically significant. This is in line with the findings of Cohen and Frazzini (2008), which provide evidence that institutional investors can process diffusing information faster. In the last two columns, 7 and 8, I use the composite measure as a proxy for limited attention. The interaction coefficient of these two last specifications amounts to -0.061 (t=-4.33) and -0.051 (t=-4.35). Overall, the regressions on limited attention indicate that investors cannot process firm-specific news directly. The documented underreaction is therefore driven by the limited attention of investors, which leads to slow information diffusion.

1.6 Conclusion

There exists a well-documented effect that investors underreact to information that diffuses among fundamentally linked firms. Understanding what drives this slow diffusion of information helps better understand investors' behavior. This paper studies which type of news or noise is slowly diffusing into the firm's stock price leading to an investors' underreaction and with which time lag the information is incorporated into the stock price. I estimate the different return components using a structural vector auto-regression, which allows me to decompose a stock's return into firm-specific news, market news, and a noise component. To model fundamental links between firms, I rely on the characteristic of analysts to cover stocks that are connected to each other.

For a global sample that consists of 42,789 stocks for 49 equity markets and spans 30 years, I present robust evidence that firm-specific news is the driver of slow information diffusion. The cross-sectional return difference between firms exposed to negative news from linked firms and those exposed to positive news amounts to approximately 7% per year. Further, it takes up to three months for an investor to incorporate the diffusing firm-specific news. Another finding of this paper is that investors can differentiate between noise and news and tend to process first market-wide news and afterward firm-specific news. This underreaction to the arrival of news is caused by the investors' limited attention to firm-specific news.

Overall, my results imply that investors are not fully limited in attention. They show a strong categorial learning behavior that enables them to incorporate at least partially relevant information diffusing among fundamentally linked firms.

2 Machine Learning and the Cross-Section of Emerging Market Stock Returns

Abstract

This paper compares various machine learning models to predict the cross-section of emerging market stock returns. We document that allowing for non-linearities and interactions leads to economically and statistically superior out-of-sample returns compared to traditional linear models. Although we find that both linear and machine learning models show higher predictability for stocks associated with higher limits to arbitrage, we also show that this effect is less pronounced for non-linear models. Furthermore, significant net returns can be achieved when accounting for transaction costs, short-selling constraints, and limiting our investment universe to big stocks only.

Key words:	Machine Learning, Return Prediction, Cross-Section of Stock Returns, Emerging Markets, Random Forest, Gradient Boosting, Neural Networks
JEL Codes:	C14, C52, C58, G11, G12, G14, G15, G17

Authors:	Matthias Xaver Hanauer, Tobias Kalsbach
First Author:	Tobias Kalsbach

Presentations:	2022 Munich Finance Day 2022 Robeco 2022 TUM School of Management
Current Status:	Working Paper

Note: In this chapter I use the first-person plural narrative.

2.1 Introduction[1]

Machine learning algorithms have been available for a long time.[2] However, due to increased computing power and data availability, decreased data storage costs, and algorithmic innovations in recent years (cf., Rasekhschaffe and Jones, 2019), machine learning methods see increasing popularity in research fields such as economics, finance, and accounting.[3]

This paper compares various machine learning models to predict the cross-section of emerging market stock returns. More specifically, we analyze the predictive power of nine algorithms: ordinary least squares regression and elastic net as examples for traditional linear models; tree-based models such as gradient-boosted regression trees and random forest; and neural networks with one to five layers. Furthermore, we investigate the performance of an ensemble comprising the five different neural networks and an ensemble of methods that allow for non-linearities and interactions, i.e., the two tree-based models and the ensemble of neural networks. In the remainder of the paper, we often use the term 'machine learning' only for the two tree-based methods, the neural networks, and the two ensembles. Our data set contains stocks from 32 emerging market countries and the 36 firm-level characteristics from Kelly, Pruitt and Su (2019) and Windmüller (2022) falling into categories such as value, past returns, investment, profitability, intangibles, and trading frictions. The data sample covers the sample period from July 1995 to December

[1] We thank David Blitz, Nicole Branger, Clint Howard, Christoph Kaserer, Tim Kroencke, Markus Leippold, Tizian Otto, Steffen Windmüller, and seminar participants at the Munich Finance Day 2022, TUM School of Management, and Robeco for their helpful comments and suggestions. Any remaining errors are our own.
Disclosures: Hanauer is also employed by Robeco. The views expressed in this paper are those of the authors and not necessarily shared by Robeco.

[2] This chapter is based on Hanauer and Kalsbach (2022).

[3] For instance, machine learning methods are applied to predict stock returns in Moritz and Zimmermann (2016), Rasekhschaffe and Jones (2019), Freyberger, Neuhierl and Weber (2020), Gu, Kelly and Xiu (2020), Tobek and Hronec (2020), Chen, Pelger and Zhu (2021), Drobetz and Otto (2021), Leippold, Wang and Zhou (2022), Azevedo et al. (2022), Cakici et al. (2022a), and Rubesam (2022), stock market betas in Drobetz et al. (2021), country stock returns in Cakici and Zaremba (2022), industry stock returns in Rapach et al. (2019), option returns in Bali et al. (2022a), corporate bond returns in Kaufmann, Messow and Vogt (2021) and Bali et al. (2022b), the equity premium in Rossi (2018), Treasury bond returns in Bianchi, Büchner and Tamoni (2021) and Bianchi et al. (2021), commodity returns in Struck and Cheng (2020), short-term bitcoin returns in Jaquart, Dann and Weinhardt (2021), cryptocurrency returns in Cakici et al. (2022b), (changes) in future company profitability in Anand et al. (2019), Van Binsbergen, Han and Lopez-Lira (2020) and Chen et al. (2022), peer-implied market capitalizations in Hanauer, Kononova and Rapp (2022), mutual fund selection in Kaniel et al. (2022), hedge fund selection in Wu et al. (2021), mortgage risk in Sadhwani, Giesecke and Sirignano (2021), or corporate directors in Erel et al. (2021).

2021, while our 20-year out-of-sample period is from January 2002 to December 2021.

Our main findings can be summarized as follows. First, we document that the different prediction algorithms pick up similar characteristics. However, we observe that tree-based methods and neural networks also identify non-linearities and interactions of characteristics. In contrast, linear methods are restricted to linear relationships and do not allow for interactions among characteristics.

Second, return forecasts based on machine learning models lead to economically and statistically superior out-of-sample long-short returns compared to traditional linear models. Furthermore, the Fama and French (2018) six-factor model can only partly explain these long-short returns, and their alphas remain highly significant. These findings are robust to several methodological choices and for emerging market subregions. Finally, we document that machine learning forecasts beat linear models consistently over our sample period, and we cannot observe a decline in predictability over time.

Third, developed market long-short returns based on machine learning forecasts derived in the same way as their emerging market counterparts cannot explain emerging market out-of-sample returns. However, models estimated solely on developed markets data also predict emerging market stock returns. These findings indicate that similar relationships between firm characteristics and future stock returns exist for developed and emerging markets but that the pricing of these characteristics is not fully integrated between developed and emerging markets.

Fourth, the high returns of the machine learning strategies in emerging markets do not primarily stem from higher-risk months and do not revert quickly, suggesting that an underreaction explanation is more likely than a risk-based explanation. Furthermore, both linear and machine learning models show higher predictability for stocks associated with higher limits to arbitrage. However, we also document that this effect is less pronounced for machine learning forecasts than for linear regression forecasts, indicating that the superiority of machine learning models in emerging markets does not stem from limits to arbitrage.

Finally, accounting for transaction costs, short-selling constraints, and limiting our investment universe to big stocks only, we document that machine learning-based return forecasts can lead to significant net outperformance over the market and net alphas, at least when efficient trading rules are applied.

This paper contributes to the literature in at least three aspects. First, we contribute to the rapidly expanding literature on predicting the cross-section of stock returns with machine learning methods. Rasekhschaffe and Jones (2019), Freyberger, Neuhierl and Weber (2020), and Gu, Kelly and Xiu (2020) document that more complex machine learning models are superior to linear models for the U.S. Tobek and Hronec (2020) and Drobetz and Otto (2021) find similar evidence for developed markets and Europe, respectively. However, none of the studies mentioned above investigates emerging markets. Emerging markets are important as they account for around 58% of the global gross domestic product (GDP), which is forecasted to rise to 61% by 2026.[4] Furthermore, the same risk factors should apply to these markets under the hypothesis that developed markets are integrated. Therefore, similar results within developed markets are not surprising, and emerging markets provide an attractive alternative for out-of-sample tests in terms of independent and new samples.

Two contemporaneously written papers, Azevedo et al. (2022) and Cakici et al. (2022a), also include emerging markets in their analysis next to developed markets. While Azevedo et al. (2022) also find that most machine learning models outperform a linear combination of anomalies, their results do not discriminate between emerging and other markets. Therefore, their results are mainly driven by developed markets. In contrast to our study, Cakici et al. (2022a) do not find superior forecasts for machine learning models compared to linear models. A potential reason for this difference might be that they train their models for each country separately while we train our models on a pooled sample of countries. However, more data might be necessary for more complex models to robustly identify non-linearities and interactions in the data.[5] We provide some supportive evidence for this claim by documenting that models trained on emerging market subregions underperform models trained on the pooled sample of emerging market subregions and that the performance loss is more pronounced for machine learning models and smaller subregions.

[4] See, IMF, World Economic Outlook database, April 2022, https://www.imf.org/en/Publications/WEO/weo-database/2022/April.

[5] While a linear model asks for a single parameter for each predictor, in the case of non-linear models, the number of parameters to estimate rapidly expands even with a moderate number of predictors (cf., Gu, Kelly and Xiu, 2020; Hanauer, Kononova and Rapp, 2022). As such, pooling data across countries will arguably improve the observations-to-parameters ratio.

Finally, Leippold, Wang and Zhou (2022) show that machine learning models dominate linear models for Chinese A-shares. In contrast, our sample purposely excludes Chinese A-shares to represent an international investor's investable emerging market universe: for the majority of our sample period, the China A-share market was only accessible to local investors and only gradually opened up to international investors (cf., Jansen, Swinkels and Zhou, 2021).

Second, we add to the literature on the drivers of emerging market stock returns. Bekaert and Harvey (1995) and Harvey (1995) were among the first to investigate emerging market country returns and their market integration. Early studies on the cross-section of emerging market stocks, such as Rouwenhorst (1999), van der Hart, Slagter and van Dijk (2003), van der Hart, de Zwart and van Dijk (2005), Griffin, Kelly and Nardari (2010), Cakici, Fabozzi and Tan (2013), and Hanauer and Linhart (2015) mainly focus on size, value, and momentum. Later studies such as Zaremba and Czapkiewicz (2017) and Hanauer and Lauterbach (2019) also investigate firm characteristics belonging to categories such as profitability, investment, intangibles, and trading frictions. Our study includes characteristics from all these groups, but machine learning models can also take non-linearities and interactions into account next to linear relationships.

Finally, our paper also contributes to the understanding of the source of return predictability from machine learning forecasts. Avramov, Cheng and Metzker (2022) show that return forecasts from deep learning models for the U.S. extract their profitability mainly from difficult-to-arbitrage stocks and during high limits-to-arbitrage market states. The authors also argue that the performance of machine learning forecasts further deteriorates when microcaps are excluded and when reasonable transaction costs are considered. Similarly, Leung et al. (2021) find that the economic gains of a gradient boosting machine model for developed market stocks tend to be more limited and critically dependent on the ability to take risk and implement trades efficiently. Furthermore, Cakici et al. (2022a) document that machine learning strategies work best for small stocks, as well as in countries with many listed firms and high idiosyncratic risk. In our paper, we follow Hou, Xue and Zhang (2018) and exclude microcaps from our analysis. While we also find that both linear and machine learning models show higher predictability for stocks associated with higher limits to arbitrage, we also document that this effect is less pronounced for machine learning models. Furthermore, we also provide evidence that a positive and significant outperformance and six-factor alpha can be achieved even when accounting for transaction costs, short-selling constraints, and limiting the investment universe to big stocks only.

The remainder of the paper is structured as follows: Section 2.2 describes the data sources, sample composition, and utilized firm-level characteristics. Section 2.3 outlines our methodology for predicting returns with machine learning algorithms, portfolio construction, and benchmark models. Section 2.4 presents evidence of the superiority of more complex machine learning models, while Section 2.5 strives to understand the source of this superiority better. We provide our conclusions in Section 2.6.

2.2 Data

2.2.1 Stock market data

Our sample comprises data from emerging stock markets as classified by Morgan Stanley Capital International (MSCI). The accounting data is from Refinitiv Worldscope, and the stock market data is from Refinitiv Datastream. The sample period starts in July 1990 and ends in December 2021. Countries are included in the sample only in those years in which they are part of the MSCI Emerging Markets Index.[6] Furthermore, countries are only part of the final sample in those months for which at least 10 stock-month observations are available after applying screens. The following 32 countries meet these criteria: Argentina, Brazil, Chile, China, Colombia, Czechia, Egypt, Greece, Hungary, India, Indonesia, Israel, Jordan, Korea, Kuwait, Malaysia, Mexico, Morocco, Pakistan, Peru, Philippines, Poland, Portugal, Qatar, Russia, Saudi Arabia, South Africa, Sri Lanka, Taiwan, Thailand, Turkey, and the UAE.[7]

We apply several static and dynamic screens to ensure that our sample comprises exclusively of common stocks and provides the highest data quality. First, we identify stocks using Refinitiv Datastream constituent lists, particularly Refinitiv Worldscope lists, re-

[6] See https://www.msci.com/market-classification for details.

[7] The Chinese sample includes only non "A"-shares to proxy the investment universe for an international investor, as the China A-share market was only accessible to local investors for the majority of our sample period (cf., Jansen, Swinkels and Zhou, 2021).

search lists, and dead lists (to eliminate survivorship bias). Following Ince and Porter (2006), Griffin, Kelly and Nardari (2010), Schmidt et al. (2019), and Hanauer (2020), we eliminate non-common equity stocks through generic and country-specific static screens. Furthermore, we apply several dynamic screens to stock returns and prices to exclude erroneous and illiquid observations. Appendix B provides a detailed description of the constituent lists and the associated static and dynamic screens. Furthermore, we require stocks to have market capitalization data for the previous month.

We follow the size group methodology of Fama and French (2008, 2012, 2017) and Hanauer and Lauterbach (2019) and assign stocks into three size groups (micro, small, and big) for each country and month. Big stocks are the largest stocks, which together account for 90% of a country's aggregated market capitalization. Small stocks comprise the next 7% of aggregated market capitalization (so that big and small stocks together account for 97% of the aggregated market size of a country). Microcaps comprise the remaining 3%.[8] Although micro stocks represent only 3% of the total market capitalization of our emerging market universe, they account for 67% of the number of stocks, which is similar to the proportion reported in Fama and French (2008) and Hanauer (2020) for the U.S. and developed markets, respectively. To prevent our results from being driven by microcaps, we follow Hou, Xue and Zhang (2018) and Hanauer and Lauterbach (2019) and exclude them. Finally, we cap the market capitalization of each stock within each month by its 99% percentile to avoid our results being driven by erroneous data and a few mega-caps.

We calculate returns from the total return index in USD. Following Jacobs (2016) and Hanauer and Lauterbach (2019), we winsorize all returns each month within a country at 0.1% and 99.9% to eliminate potential errors. To calculate the excess returns, we obtain the risk-free rate from Kenneth R. French's homepage.[9]

The result is a comprehensive dataset spanning 15.152 unique stocks and more than 1.42 million stock-month observations. Table 2.1 depicts the descriptive statistics for the final sample.

[8] To distinguish between these size groups, Fama and French (2008) use the 20^{th} and 50^{th} percentiles of end-of-June market cap on NYSE stocks as size breakpoints for the U.S. market, which on average are bigger than AMEX or NASDAQ stocks. However, these breakpoints are applied to all (NYSE, AMEX, and NASDAQ) stocks. For international markets, Fama and French (2012, 2017) propose to calculate breakpoints based on aggregated market capitalization, as we do.

[9] See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Table 2.1
Summary statistics by country

The table presents summary statistics for the 32 countries of our sample. Columns (1), (2), (3), and (4) report the total, minimum, mean, and maximum number of firms per country. Columns (5) and (6) state the average mean and median size per country-month. Column (7) shows the average total size per country-month and column (8) reports these values in percentage of the respective total across countries. Size is measured as market capitalization in million USD. The last two columns, Column (9) and Column (10), report the actual beginning and ending dates during which each country is included in our sample.

	Number of firms				Size				Date	
	Total	Min	Mean	Max	Mean	Median	Total	%	Start	End
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Argentina	96	11	30	45	832	376	25467	1.12	91-05	21-12
Brazil	289	17	65	154	3121	1500	247409	4.08	94-09	21-12
Chile	201	53	74	102	1740	826	125716	4.26	90-07	21-12
China	28	10	16	24	2772	1292	45063	0.10	16-05	21-12
Colombia	50	14	19	25	2791	2178	53710	0.95	94-07	21-12
Czechia	89	10	36	77	662	234	13541	0.24	97-07	05-08
Egypt	199	51	81	123	577	213	46194	0.59	01-07	21-12
Greece	334	37	90	224	471	178	39442	1.56	90-07	21-12
Hungary	41	10	12	22	1921	486	22053	0.44	97-07	21-12
India	2238	356	593	893	1242	334	788478	13.09	94-07	21-12
Indonesia	649	35	150	296	1019	324	181003	4.50	90-07	21-12
Israel	634	173	245	331	270	60	62659	1.31	95-07	10-06
Jordan	161	10	98	119	328	70	32467	0.08	06-04	09-06
Korea	2972	394	803	1343	622	134	572232	12.87	92-07	21-12
Kuwait	81	73	75	78	1504	396	113350	0.02	21-07	21-12
Malaysia	1173	177	389	534	641	154	242228	9.97	90-07	21-12
Mexico	181	25	54	70	2692	1316	156068	4.92	90-07	21-12
Morocco	57	24	31	37	1323	658	42758	0.41	01-07	14-06
Pakistan	362	73	139	205	190	68	28089	0.50	94-07	21-12
Peru	103	17	28	40	1100	643	31521	0.61	94-07	21-12
Philippines	270	23	80	113	1087	433	101378	2.70	90-07	21-12
Poland	591	26	135	232	653	143	100667	1.74	95-07	21-12
Portugal	99	35	51	60	394	155	18182	0.84	90-07	98-06
Qatar	32	25	27	29	4969	2880	134437	0.42	14-07	21-12
Russia	232	10	54	102	4711	1958	285849	3.77	98-07	21-12
Saudi Arabia	89	34	52	84	8189	4936	398893	0.36	19-07	21-12
South Africa	517	80	131	240	2445	1101	274534	6.31	95-07	21-12
Sri Lanka	150	88	98	105	15	7	1530	0.03	94-07	01-06
Taiwan	1912	339	743	978	772	223	596406	11.36	97-07	21-12
Thailand	823	133	237	387	728	188	194393	6.48	90-07	21-12
Turkey	430	50	134	237	862	243	125222	3.77	90-07	21-12
UAE	69	33	43	49	4703	1653	201327	0.61	14-07	21-12
Global	15152	594	3763	5690	901	209	3909693	100.00	90-07	21-12

2.2.2 Firm-level characteristics

The 36 firm-level characteristics in this study are analogous to those in Kelly, Pruitt and Su (2019) and Windmüller (2022) and constructed using data from Refinitiv Datastream and Worldscope. Appendix B outlines the detailed construction of the characteristics. We follow Windmüller (2022) and substitute the daily bid-ask spreads with the daily version of Amihud (2002) illiquidity as a proxy for trading frictions. As shown by Fong, Holden and Trzcinka (2017), the Amihud (2002) illiquidity measure increases the number of observations in the cross-section and is the best daily cost-per-dollar-volume proxy for international data.

The 36 characteristics are: assets-to-market (`A2ME`), total assets (`AT`), sales-to-assets (`ATO`), book-to-market (`BEME`), market beta (`Beta`), cash-and-short-term-investment-to-assets (`C`), capital turnover (`CTO`), capital intensity (`D2A`), leverage (`Debt2P`), ratio of change in property, plants, and equipment to change in total assets (`DPI2A`), earnings-to-price (`E2P`), fixed costs-to-sales (`FC2Y`), cash flow-to-book (`FreeCF`), idiosyncratic volatility (`Idiovol`), investment (`INV`), market capitalization (`LME`), turnover (`LTurnover`), net operating assets (`NOA`), operating accruals (`OA`), operating leverage (`OL`), price relative to its 52-week high (`P2P52WH`), price-to-cost margin (`PCM`), profit margin (`PM`), gross profitability (`Prof`), Tobin's Q (`Q`), momentum ($\mathtt{r}_{12-2}$), intermediate momentum ($\mathtt{r}_{12-7}$), short-term reversal ($\mathtt{r}_{2-1}$), long-term reversal ($\mathtt{r}_{36-13}$), return on net operating assets (`RNA`), return on assets (`ROA`), return on equity (`ROE`), sales-to-price (`S2P`), the ratio of sales and general administrative costs to sales (`SGA2S`), unexplained volume (`SUV`), and Amihud (2002) illiquidity (`Illiqu`).

Moreover, in a robustness check, we add the following four characteristics that have been shown to be strong return predictors for emerging markets (Hanauer and Lauterbach, 2019): monthly updated book-to-market ($\mathtt{BEME}_\mathtt{m}$, Asness and Frazzini, 2013), composite equity issuance (`CEI`, Daniel and Titman, 2006), cash flow-to-price (`CF2P`, Lakonishok, Shleifer and Vishny, 1994), and gross profitability-to-assets (`GP2A`, Novy-Marx, 2013).

We do not exclude financial firms but set the following characteristics as missing as they are not meaningfully defined for financials: `ATO`, `C`, `D2A`, `DPI2A`, `FC2Y`, `FreeCF`, `CF2P`, `GP2A`, `OA`, `PCM`, `PM`, `Prof`, `RNA`, `SGA2S`, and `NOA`.

Following Freyberger, Neuhierl and Weber (2020), Gu, Kelly and Xiu (2020), and Leippold, Wang and Zhou (2022), we rank all stock characteristics cross-sectionally for each month and country into the [-1,1] interval to limit the effect of outliers. These country-based ranks aim to address the impact of different accounting standards across countries, particularly in the earlier part of the sample period, and thus account for cross-country differences in characteristics. In case of missing characteristics, we replace them with a 0 to ensure broad cross-sectional coverage. Balance sheet data from the fiscal year ending in calendar year t-1 is used from end-of-June in year t to end-of-May in year t+1 to predict stock returns from July in year t to end-of-June in year t+1.

2.3 Methodology

2.3.1 Return prediction using machine learning

Rasekhschaffe and Jones (2019) stress that domain knowledge is essential to structure the forecasting problem in a way that increases the signal-to-noise ratio. As we are interested in the cross-section of stock returns and rank stocks in portfolio sorts later in a country-neutral manner, we aim to forecast the outperformance of a stock relative to its country market return. Therefore, we define the abnormal return of a stock i, $i = 1, ..., N$ in month t, $t = 1, ..., T$ in the country c, $c = 1, ..., C$ as

$$r^{abn}_{i,t,c} = r_{i,t,c} - Mkt_{t,c}, \tag{2.1}$$

where $r_{i,t,c}$ is the return of stock i in month t of country c and $Mkt_{t,c}$ is the value-weighted market return in month t of country c.

Following Gu, Kelly and Xiu (2020), we employ a general additive prediction model to describe the one-month-ahead abnormal return of a stock $r^{abn}_{i,t+1,c}$, which can be written as

$$r^{abn}_{i,t+1,c} = E_t[r^{abn}_{i,t+1,c}|x_{i,t}] + \epsilon_{i,t+1,c}, \tag{2.2}$$

where $E_t[r^{abn}_{i,t+1,c}|x_{i,t}]$ is the conditional expected abnormal return of stock i in month t for month $t + 1$ given a vector of stock-specific p characteristics known at month t, $x_{i,t} \in \mathbb{R}^p$, and $\epsilon_{i,t+1,c}$ is the prediction error term. Our objective is to estimate the expected abnormal return by using an unknown function f^*, $f^* : \mathbb{R}^p \rightarrow \mathbb{R}$, which estimates the expected returns independently of any other information besides the vector of p stock-specific characteristics available in month t:

$$E_t[r^{abn}_{i,t+1,c}|x_{i,t}] = f^*(x_{i,t}). \tag{2.3}$$

In the case of supervised machine learning, the unknown function $f^*(x)$ is approximated by some function $f(x, \theta, \rho)$, which is parameterized by a vector of coefficients θ and a set of hyperparameters ρ. While θ is directly derived from the underlying training data with respect to ρ and a specific loss function L, ρ itself depends on the user input but is optimized concerning L based on available data. The exact functional form of f depends on the family and can be either linear or non-linear, parametric or non-parametric.

For this paper, we build on Rasekhschaffe and Jones (2019), Gu, Kelly and Xiu (2020), Tobek and Hronec (2020), Drobetz and Otto (2021), and Leippold, Wang and Zhou (2022) to select a representative amount of machine learning models from the finance literature. We analyze the predictive power of nine different algorithms: ordinary least squares (OLS) regression, elastic net ($ENet$), gradient-boosted regression trees ($GBRT$), random forest (RF), and neural networks with one to five layers (NN_1, NN_2, NN_3, NN_4, NN_5). We also investigate the performance of an ensemble of the five different neural networks (NN_{1-5}) and the average combination of the more advanced machine learning methods (ENS): $GBRT$, RF, and NN_{1-5}. We provide a more detailed description of the models in Appendix B.

Besides the model selection, we also follow the standard approach in the literature (Gu, Kelly and Xiu, 2020; Leippold, Wang and Zhou, 2022) for selecting the hyperparameter range, the training of the models, and the performance evaluation. One of the most crucial things when estimating the different machine learning models is to avoid data leakage. This happens when information exceeding the training dataset is used to create the model. Therefore, we divide our data into three disjoint periods, which always maintain the temporal ordering: the training, validation, and testing samples. We first estimate the models for a range of hyperparameters based on the training data. Next, we determine the respective loss of each hyperparameter set and model in the validation sample. The optimal hyperparameter set minimizes the individual model's respective loss function. Afterward, we retrain the model with the optimal hyperparameter set on the combined training and validation data. Next, the models are used to predict the monthly returns for the test dataset. We describe an example of this procedure for the first two years in our sample: we first estimate the models for a range of hyperparameters based on the training data from July 1990 to December 1995. Afterward, we determine the best hyperparameters through the validation sample from January 1996 to December 2001. Finally, the model is retrained with the optimal hyperparameter using the data from July 1990 to December 2001 and evaluated in the testing sample using data from January 2002 to December 2002. To test our models from January 2003 to December 2003, we extend the training sample by one year (July 1990 to December 1996) and roll the validation sample forward by one year (January 1997 to December 2002). This procedure ensures that no future information is leaked from a previous period. Since machine learning models are computationally intensive, we retrain them only once at the end of every year but do the prediction every month using the latest model and data. Appendix B summarizes the hyperparameter tuning schemes for each model.

2.3.2 Machine learning portfolios

We mainly rely on portfolio performance analysis to evaluate the predictive performance of the different machine learning models. For a given machine learning model, we follow the following approach: At the end of each month t, we predict the next month's abnormal return ($\hat{r}^{abn}_{i,t,c}$), which we use for sorting stock into quintiles. To avoid that small stocks or certain countries dominate our results, we estimate the quintile breakpoints for each country separately based on big stocks as recommended in Hou, Xue and Zhang (2018) and applied in Hanauer and Lauterbach (2019). Furthermore, the machine-learning-based signals should not only contain information on the return predictability in equal-weighted sorts, which may be driven by smaller stocks, but also in value-weighted sorts, which are dominated by larger stocks. Finally, we construct a zero-net investment portfolio (long-short) that goes long in the highest quintile portfolio and short in the lowest quintile portfolio. We reassign and rebalance all portfolios at the end of each month.

2.3.3 Benchmark factor models

To benchmark the results of the different machine learning portfolio sorts, we consider the Fama and French (2018) six-factor model, i.e., the Fama and French (2015) five-factor model with a cash-based profitability factor and augmented with the Carhart (1997) momentum factor. The corresponding six factors are market ($RMRF$), size (SMB, small minus big), value (HML, high minus low), profitability (RMW, robust minus weak), investment (CMA, conservative minus aggressive), and momentum (WML, winners minus losers). These factors are based on the same stock sample as the machine learning portfolios, i.e., we also exclude microcaps. Furthermore, we use regional versions of the factors for studying emerging market regions. Appendix B provides a detailed description of how the factors are constructed.

2.4 Empirical results

This section presents evidence on the application of various machine learning models in emerging markets. We begin by analyzing the out-of-sample R^2_{OOS} of individual stock returns. Subsequently, we evaluate the importance of different characteristics, the sensitivity of the predicted returns to various characteristics, and the sensitivity to the interaction effects of different characteristics. Next, we employ portfolio sorts to assess the economic gains of using different machine learning models. Finally, we investigate the impact of various methodological changes and the robustness of our findings in emerging market subregions.

2.4.1 Prediction performance

Table 2.2 presents the out-of-sample R^2_{OOS} for our set of machine learning models, which measures the predictive power on the individual stock level. In Panel B, we show the Newey and West (1987) adjusted Diebold and Mariano (1995) test statistics to compare the out-of-sample stock-level prediction performance between each machine learning model. We measure the pooled out-of-sample R^2_{OOS} in Panel A as:

$$R^2_{OOS} = 1 - \frac{\sum_t^T \sum_i^N (r^{abn}_{i,t,c} - \hat{r}^{abn}_{i,t,c})^2}{\sum_t^T \sum_i^N (r^{abn}_{i,t,c})^2}. \tag{2.4}$$

Table 2.2
Monthly out-of-sample stock-level prediction performance

This table summarizes the monthly out-of-sample stock-level prediction performance using OLS (OLS), elastic net ($ENet$), random forest (RF), gradient boosted regression trees ($GBRT$), neural networks with 1 to 5 layers (NN_1–NN_5), an ensemble of the different neural networks (NN_{1-5}), and an ensemble of the different non-linear machine learning algorithms (ENS). Panel A reports the monthly R^2_{OOS} statistics for the full sample and within subsamples that include only large stocks or small stocks. Panel B reports pairwise Newey and West (1987) adjusted Diebold-Mariano test statistics comparing the out-of-sample stock-level prediction performance among each machine learning model. Positive numbers indicate the column model outperforms the row model. Bold font indicates the difference is significant at 1% level or better for individual tests, and an asterisk indicates significance at the 1% level for 10-way comparisons via our conservative Bonferroni adjustment. The out-of-sample period is from January 2002 to December 2021.

	OLS	$ENet$	RF	$GBRT$	NN_1	NN_2	NN_3	NN_4	NN_5	NN_{1-5}	ENS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Panel A: Percentage R^2_{OOS}											
Full Sample	0.29	0.18	0.40	0.52	0.49	0.53	0.53	0.55	0.54	0.60	0.60
Large firms	0.12	-0.01	0.25	0.30	0.19	0.24	0.27	0.31	0.31	0.34	0.38
Small firms	0.40	0.29	0.49	0.66	0.67	0.70	0.68	0.70	0.68	0.75	0.73
Panel B: Between-model comparison of predictive performance											
OLS		-2.13	**3.45***	**6.35***	**5.57***	**5.42***	**5.64***	**6.05***	**6.65***	**7.34***	**8.25***
$ENet$			**4.53***	**6.50***	**5.91***	**6.08***	**6.29***	**7.21***	**7.01***	**7.68***	**8.19***
RF				**6.80***	**2.96**	**3.96***	**4.38***	**5.27***	**4.57***	**6.59***	**12.65***
$GBRT$					-0.72	0.71	0.68	1.74	1.00	**3.49***	**7.59***
NN_1						2.51	2.12	**3.24**	**2.85**	**9.19***	**4.38***
NN_2							-0.23	1.69	0.25	**6.01***	2.37
NN_3								1.82	0.39	**5.86***	**2.82**
NN_4									-1.42	**3.12**	1.45
NN_5										**5.30***	**2.60**
NN_{1-5}											-0.45

The first row in Panel A of Table 2.2 reports the R^2_{OOS} of the full sample. The OLS yields a benchmark R^2_{OOS} of 0.29%, which all other models improve except for the $ENet$ (R^2_{OOS} of 0.18%). Since the $ENet$ shrinks certain coefficients towards zero but does not consider interactions or non-linearities, it seems that this regularization does not increase the predictability. The RF and $GBRT$ are superior to the OLS, producing fits of 0.40% and 0.52%, respectively. Only the NN_1 underperforms the $GBRT$ but outperforms all other linear and non-parametric models and yields a R^2_{OOS} of 0.49%. The NN_2 to NN_5

show R^2_{OOS} between 0.53% and 0.55%, with the NN_4 performing the best. Creating an ensemble of neural networks (NN_{1-5}) and an ensemble of the non-linear machine learning models (ENS) produces fits for both models of 0.60%.

A closer look at the second and third rows in Panel A of Table 2.2 reveals an interesting pattern: in all the cases, the predictive performance is better for small firms than for large firms. The ensemble of neural networks (NN_{1-5}) and the ensemble of non-linear machine learning models (ENS) yield a R^2_{OOS} of 0.34% and 0.38% for large firms and 0.75% and 0.73% for small firms, respectively.

Whereas Panel A measures the individual predictive performance of the different machine learning models, Panel B assesses the statistical significance of differences among the models using the Newey and West (1987) adjusted Diebold and Mariano (1995) test statistics (DM_{kj}) comparing a column model (k) versus a row model (j). We compute the Newey-West adjusted Diebold-Mariano test statistics as:

$$\begin{aligned} MSFE_t^m &= \frac{1}{N_t}\sum_{i=1}^{N_t}(r_{i,t,c}^{abn} - \hat{r}_{i,t,c,m}^{abn})^2 \\ d_{kj,t} &= MSFE_t^k - MSFE_t^j \\ \bar{d}_{kj} &= \frac{1}{T}\sum_{T}^{t=1} d_{kj,t} \\ DM_{kj} &= \frac{\bar{d}_{kj}}{\hat{\sigma}_{d_{kj},NW(4)}}, \end{aligned} \tag{2.5}$$

where $\hat{\sigma}_{d_{kj},NW(4)}$ is the Newey and West (1987) standard error of $d_{kj,t}$ with four lags. The Diebold-Mariano test statistic is normally distributed with a mean of 0 and a standard deviation of 1 ($\mathcal{N}(0, 1)$) with the null hypothesis that there exists no difference between the models, which allows us to map the magnitudes of the test statistic to p-values. Bold numbers indicate a significant difference between the models at the 1% level ($DM \geq 2.60$). An asterisk indicates statistical significance at the 1% level for 10-way comparisons via the conservative Bonferroni adjustment, which increases the critical value to 3.33.

Except for the $ENet$, all machine learning models outperform the OLS and exceed the Bonferroni adjusted critical value of 3.33. The comparison between the RF and all other non-linear models yields a similar result, with the $GBRT$ and all other neural networks, except NN_1, outperforming the RF. For the $GBRT$, only the two machine-learning ensembles significantly improve prediction performance, as evidenced by a DM statistic of 3.49 and 7.59, respectively. The different neural networks with one to five layers do not differ much in their prediction performance. In the case of the NN_1, the neural networks with four and five layers are superior. The two best-performing machine learning models are the two ensembles. While the ensemble of neural networks (NN_{1-5}) significantly outperforms all other machine learning models, the ensemble of the trees and neural networks (ENS) exhibits statistically significant outperformance when compared to the OLS, $ENet$, RF, $GBRT$, NN_1, NN_3, and NN_5.

2.4.2 Characteristics importance and marginal relationships

Next, we examine the importance of individual characteristics in predicting abnormal returns and the model-implied marginal impact of individual characteristics on expected abnormal returns.

We determine the importance of each characteristic for each model by measuring the average reduction in R^2_{oos} by setting each value of the particular characteristic to zero and keeping the remaining model estimates fixed. Figure 2.1 visualizes the sum over the cross-sectional ranked characteristics for the different machine learning models.[10] A darker color in the figure indicates higher importance of the characteristic for the individual model, while a lighter color indicates lower importance for the R^2_{oos}.

The most influential characteristics are similar among the different machine learning models, with turnover ($LTurnover$), idiosyncratic volatility ($Idiovol$), price relative to its 52-week high ($P2P52WH$), Amihud (2002) illiquidity ($Illiqu$), total assets (AT), market capitalization (LME), and market beta ($Beta$) from the trading frictions category; momentum (r_{12-2}), short-term reversal (r_{2-1}), and intermediate momentum (r_{12-7}) from the past returns category; and assets-to-market ($A2ME$), Tobin's Q (Q), book-to-market ($BEME$), and leverage ($Debt2P$) from the value category all being among the top 15 characteristics. However, characteristics from the profitability and intangibles categories, except for return on asset (ROA), are not present among the top 15.

[10] In addition, we present the most influential characteristics per model and the corresponding normalized importance in Figure B. 1 in the appendix.

Figure 2.1
Characteristic importance by model

This figure shows the ranked characteristic importance for the variables in each model. Characteristic importance is an average over all training samples and importance within each model is normalized to sum to one.

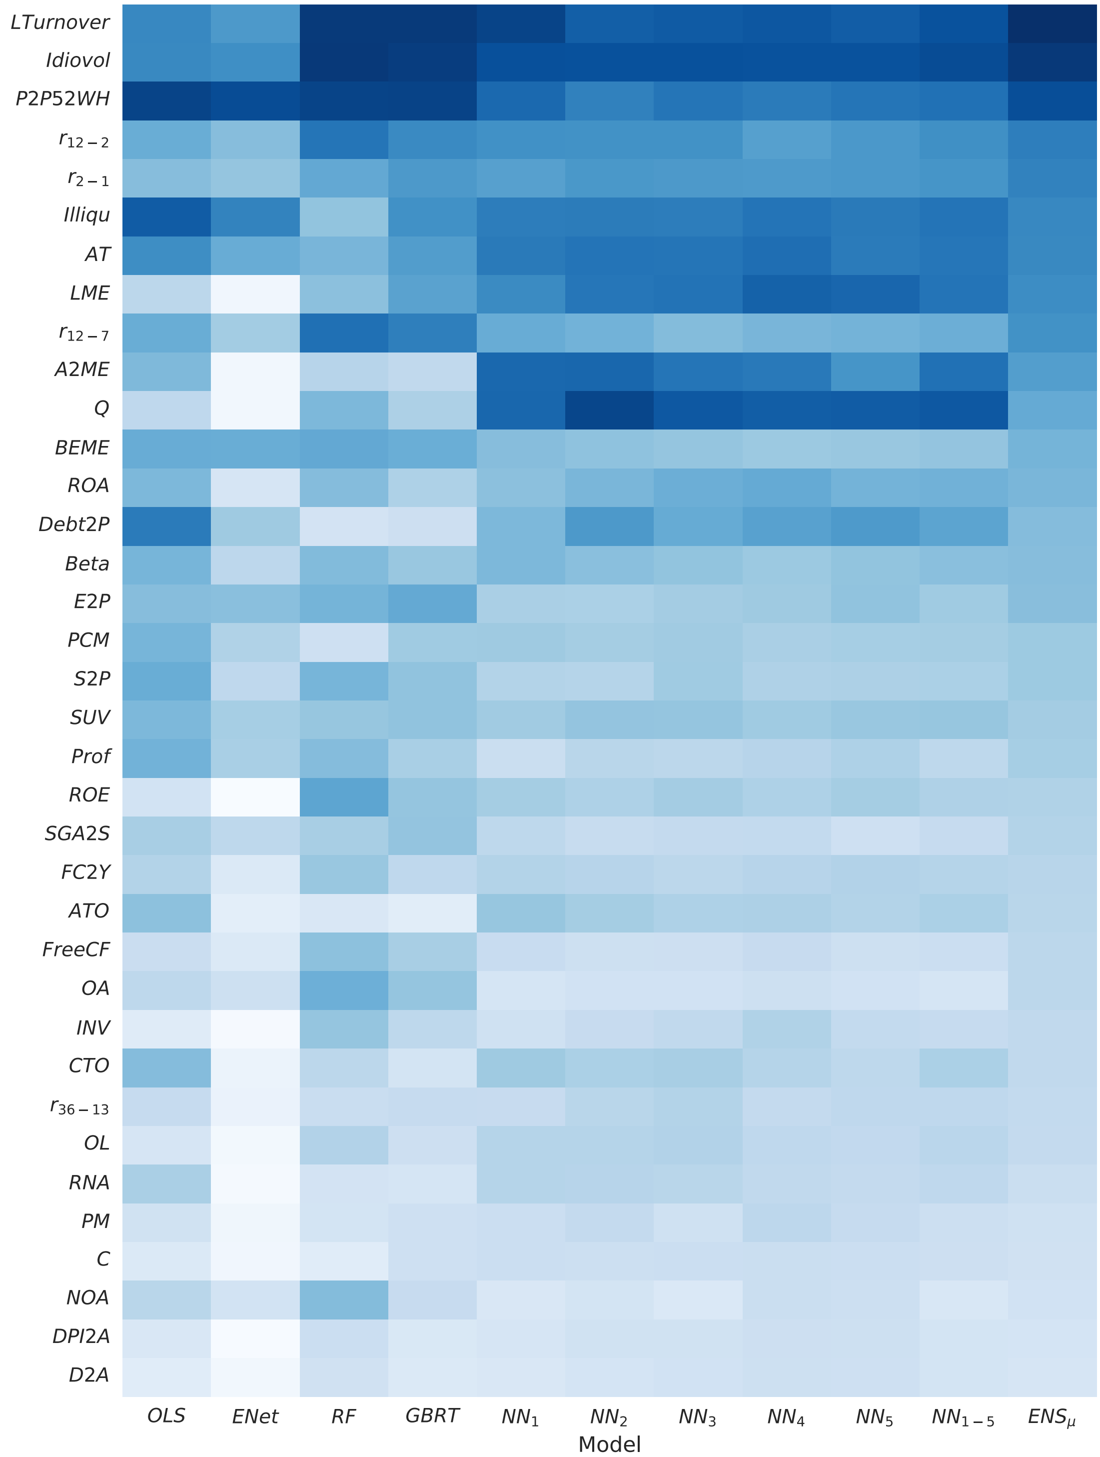

Figure 2.2 visualizes the marginal impact of individual characteristics on expected abnormal returns for the OLS, $ENet$, RF, $GBRT$, and NN_{1-5}. We predict the returns for each model and characteristic by iterating over the (-1,1) interval and holding all other characteristics fixed at zero. We do this for each time period and model individually and calculate the average predicted return among the different machine learning models. We select short-term reversal (r_{2-1}), idiosyncratic volatility ($Idiovol$), turnover ($LTurnover$), and operating leverage (OL) as examples to visualize how the different machine learning models associate the underlying characteristic with the expected abnormal returns.

Figure 2.2

Marginal association between expected returns and characteristics

The figure shows the sensitivity of expected returns (vertical axis) to the four following individual characteristics (holding all other covariates fixed at their median values): short-term reversal (r_{2-1}, top-left), idiosyncratic volatility ($Idiovol$, top-right), turnover($LTurnover$, bottom-left), and operating leverage (OL, bottom-right).

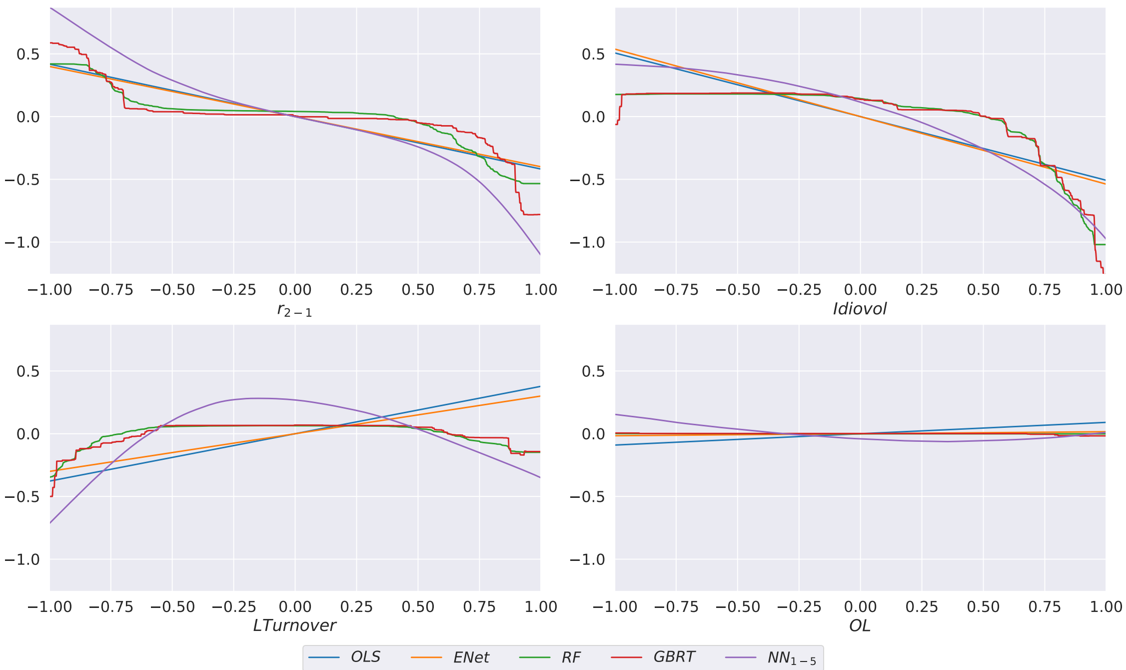

Inspecting the relationships in Figure 2.2, we observe that all methods identify the well-known negative relationship between expected returns with short-term reversal (r_{2-1}, top-left) or idiosyncratic volatility ($Idiovol$, top-right). While the two linear models are, per definition, restricted to linear relationships, we see that tree-based methods and neural networks identify more pronounced short-term reversal patterns in the extremes.[11] Similarly, these methods detect a relatively flat relationship for low and medium levels of idiosyncratic volatility ($Idiovol$) but an increasingly negative relationship for high idiosyncratic volatility, echoing the empirical results in Ang et al. (2006). The differences are even more pronounced for turnover ($LTurnover$, bottom-left). While both OLS and $ENet$ find a positive slope, the two tree-based models, RF and $GBRT$, and the neural network ensemble, NN_{1-5}, identify an inverted U-shape pattern: extreme positive and negative values of $LTurnover$ are associated with lower expected return than the middle region in the interval, echoing the pattern documented in Freyberger, Neuhierl and Weber (2020). Such differences in marginal relationships can partly explain the divergence in the performance of linear and non-linear methods. However, we also observe that all methods agree on a nearly zero relationship between operating leverage (OL, bottom-right) and expected returns.

A significant advantage of the tree-based models and the different neural networks is that they can model complex interactions among the various characteristics. In Figure 2.3, we illustrate how the NN_{1-5} can model complex interactions between characteristics. Specifically, we show the sensitivity of the expected returns to pairwise interaction effects for Amihud (2002) illiquidity ($Illiqu$) and idiosyncratic volatility ($Idiovol$) with short-term reversal (r_{2-1}) and market capitalization (LME) by varying both pairs of characteristics while holding the other predictors fixed. We choose $Illiqu$ and $Idiovol$ as they are prominent hard-to-value proxies (cf., Kumar, 2009) and r_{2-1} and LME as they are two main control characteristics in the asset pricing literature.

The upper-left figure illustrates that the difference between high and low previous month returns is the most substantial for very illiquid stocks (purple line). In contrast, the lines remain mostly parallel for other values of $Illiqu$. This finding is consistent with the empirical observation for the interaction between short-term reversal with turnover reported in Medhat and Schmeling (2022). The upper-right figure depicts the interactions between a stock's market capitalization and the Amihud (2002) illiquidity measure. For liquid firms (blue and orange line), the expected return increases with market capitalization.

[11] This finding is consistent with the empirical pattern for short-term reversal deciles that can be found on Kenneth R. French's homepage.

Figure 2.3
Expected returns and characteristic interactions (NN_{1-5})

The figure shows the sensitivity of the expected returns (vertical axis) to interactions effects for four selected combinations in model NN_{1-5} (holding all other characteristics fixed at their median values of 0): Amihud (2002) illiquidity (*Illiqu*) and short-term reversal (r_{2-1}) (top-left), Amihud (2002) illiquidity and market capitalization (LME) (top-right), idiosyncratic volatility (*Idiovol*) and short-term reversal (bottom-left), and idiosyncratic volatility and market capitalization (bottom-right).

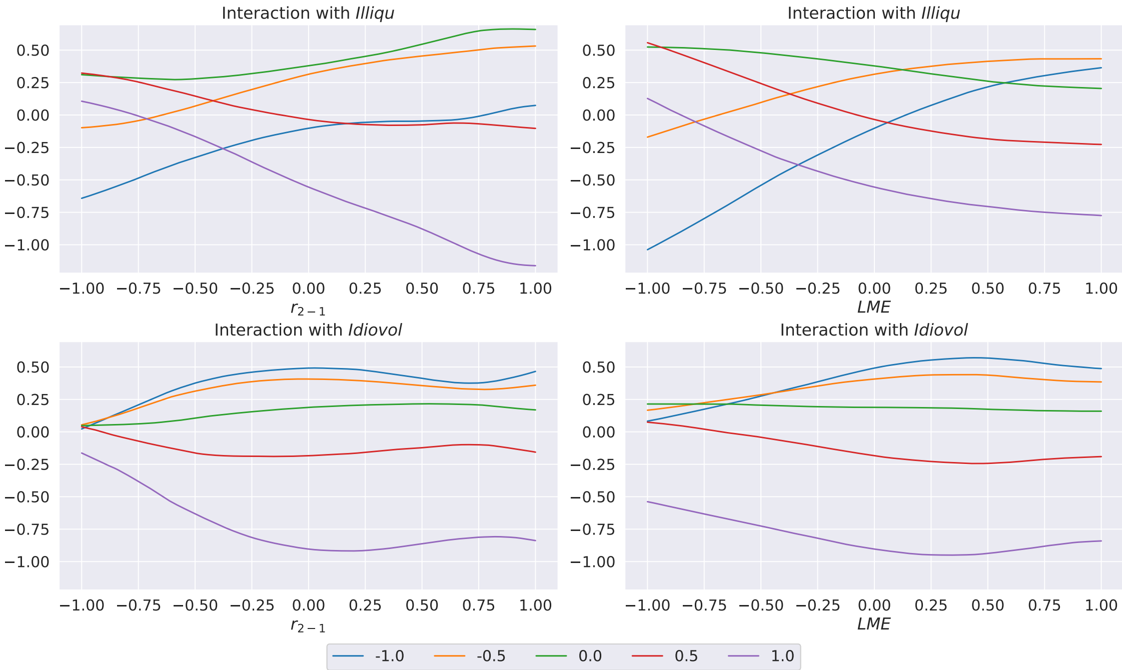

In contrast, the relationship is reversed for illiquid firms (red and purple line), implying that expected returns decrease for larger firms. The bottom-left figure reveals that the short-term reversal effect is most pronounced and S-shaped for risky stocks (purple line). In contrast, the reversal effect is concave for less risky stocks (blue and orange line), yielding significantly lower returns when the prior month's returns are high. Finally, the bottom-right figure indicates that no strong interaction effects exist between $Idiovol$ and LME.

2.4.3 Portfolio performance

Following our analysis of the predictive ability of the different machine learning methods for individual stock returns, we will now proceed with a general overview of the profitability of machine learning signal-based portfolios.

Table 2.3 displays the results of our analysis on equal- and value-weighted country-neutral quintile portfolio sorts using big-stock breakpoints. In Panel A and Panel D, we report the predicted monthly returns for the long-short quintile (Pred), the average monthly return for the long-short quintile (Avg), Newey and West (1987) adjusted t-statistics with four lags (t-stat), monthly standard deviations (SD), and Sharpe ratios (SR). Panel B and Panel E show the alphas (α), corresponding Newey and West (1987) adjusted t-statistics with four lags (t-stat$_\alpha$), and R^2 with respect to the Fama and French (2018) six-factor model:

$$\begin{aligned} r_{p,t,ML} - r_{f,t} = \alpha + \beta_1 RMRF_t + \beta_2 SMB_t + \beta_3 HML_t \\ + \beta_4 RMW_t + \beta_5 CMA_t + \beta_6 WML_t + \epsilon_t. \end{aligned} \tag{2.6}$$

We additionally provide detailed results on every quintile in Table B. 1 in the appendix.[12]

Panel C and Panel F describe the maximum drawdowns (Max DD), the most negative monthly return (Max 1M Loss), and the average monthly percentage change in holdings (TO) of different machine learning-based long-short portfolios. We define maximum drawdowns as

$$\text{Max DD} = \max_{0 \leq t_1 \leq t_2 \leq T} (Y_{t_1} - Y_{t_2}), \tag{2.7}$$

where Y_t is the cumulative log return from date 0 through t. The strategy's average monthly turnover is defined as

$$\text{TO} = \frac{1}{T} \sum_{t=1}^{T} \left(\sum_{i=1}^{N_t} \left| w_{i,t+1} - \frac{w_{i,t}(1 + r_{i,t+1})}{1 + \sum_{j=1}^{N_t} w_{j,t} r_{j,t+1}} \right| \right), \tag{2.8}$$

where $w_{i,t}$ is the weight of stock i in the portfolio at time t.

We start by analyzing the equal-weighted long-minus-short quintile returns in Panel A of Table 2.3. All machine learning models yield positive and highly significant long-short returns. The order is similar to the monthly out-of-sample stock-level prediction perfor-

[12] We also include the performance of a strategy that uses the equal-weighted (1/N) average of all standardized characteristics ($\mu_{sign(c)}$) in this table. Thereby, characteristics are sorted in such a way that higher values correspond to higher expected returns. The performance of this simple linear combination is slightly worse (similar) than that of the other two linear strategies for equal-weighted (value-weighted) portfolios.

Table 2.3

Drawdowns, turnover, and risk-adjusted performance of machine learning portfolios

This table reports the out-of-sample performance of the different machine learning long-short portfolios. Stocks are sorted into country-neutral quintiles portfolios based on their predicted returns for the next month. The sorting breakpoints are based on big stocks only, which are in the top 90% of a country's aggregated market capitalization. Panel A (Panel D) summarizes the quintile sort results from equal-weighting (value-weighting) and provides the predicted monthly returns for the long-short quintile (Pred), the average monthly returns of the long-short quintile (Avg), Newey and West (1987) adjusted t-statistics with 4 lags (t-stat), their standard deviations (SD), and annualized Sharpe ratios (SR), respectively. Panel B (Panel E) reports the average Fama and French (2018) six-factor model alphas (α_{FF6}), corresponding Newey and West (1987) adjusted t-statistics with 4 lags (t-stat$_\alpha$), and corresponding R^2 using equal-weighting (value-weighting). Panel C (Panel F) describes the maximum drawdowns (Max DD), the most negative monthly return (Max 1M Loss), and the average monthly turnover in % of the equal-weighted (value-weighted) long-short portfolio. The sample period is from January 2002 to December 2021.

	OLS	$ENet$	RF	$GBRT$	NN_1	NN_2	NN_3	NN_4	NN_5	NN_{1-5}	ENS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Panel A: Quntile sorts performance - Equal-weighted											
Pred	1.93	1.97	1.50	1.80	2.61	2.60	2.41	2.29	2.25	2.30	1.71
Avg	1.38	1.20	1.60	1.82	1.89	1.91	1.84	1.86	1.85	1.88	1.86
t-stat	7.82	6.83	9.33	11.57	14.01	15.75	14.81	13.82	13.50	13.50	11.79
SD	2.04	2.12	2.04	1.88	1.68	1.58	1.60	1.66	1.69	1.72	1.87
SR	2.34	1.96	2.71	3.35	3.91	4.21	4.00	3.89	3.78	3.79	3.44
Panel B: Risk-adjusted performance - Equal-weighted											
α_{FF6}	0.97	0.83	1.19	1.40	1.47	1.55	1.49	1.48	1.44	1.46	1.43
t-stat$_\alpha$	8.02	6.94	14.10	15.65	15.67	19.02	16.93	16.72	15.79	15.81	15.66
R^2	62.42	55.79	59.81	60.89	53.21	48.65	52.70	54.66	56.15	55.67	58.17
Panel C: Drawdowns and turnover - Equal-weighted											
Max DD (%)	26.35	26.23	21.69	18.84	16.70	13.45	16.00	16.07	17.61	17.82	19.04
Max 1M loss (%)	13.97	12.96	10.53	10.70	9.37	7.65	10.20	10.17	10.25	10.75	10.68
Turnover (%)	89.27	96.38	89.61	97.39	101.87	102.02	100.80	99.21	99.50	99.72	95.77
Panel D: Quntile sorts performance - Value-weighted											
Pred	1.85	1.89	1.39	1.61	2.30	2.21	2.04	1.94	1.93	1.97	1.52
Avg	0.84	0.73	0.99	1.06	1.04	1.12	1.12	1.20	1.17	1.21	1.21
t-stat	4.64	4.01	5.28	6.14	7.00	9.47	7.91	8.35	8.17	8.55	7.04
SD	2.22	2.36	2.32	2.17	1.95	1.75	2.01	1.97	1.87	1.98	2.20
SR	1.31	1.07	1.48	1.69	1.85	2.23	1.93	2.11	2.17	2.12	1.91
Panel E: Risk-adjusted performance - Value-weighted											
α_{FF6}	0.28	0.27	0.47	0.57	0.57	0.71	0.66	0.73	0.71	0.72	0.67
t-stat$_\alpha$	2.72	2.28	5.24	6.73	4.83	9.16	6.55	7.76	8.21	8.26	8.29
R^2	68.25	56.39	67.61	68.50	52.97	48.50	51.79	58.50	59.39	56.86	67.17
Panel F: Drawdowns and turnover - Value-weighted											
Max DD (%)	30.49	31.45	31.07	26.54	23.63	15.44	23.19	21.81	20.81	20.36	25.28
Max 1M loss (%)	16.60	17.82	14.46	14.73	16.45	9.58	17.36	15.90	12.83	14.81	14.81
Turnover (%)	91.28	97.18	90.46	101.10	103.81	106.35	106.02	104.35	104.43	101.46	96.85

mance in Table 2.2. The linear methods OLS and $ENet$ yield a monthly return of 1.38% (t-stat 7.82) and 1.20% (t-stat 6.83), respectively. However, the tree-based methods RF and $GBRT$ exhibit even higher long-short returns of 1.60% (t-stat 9.33) and 1.80% (t-stat 11.57), which themselves are outperformed by the neural networks with returns between 1.84% (NN_3) and 1.91% (NN_2) and t-statistics between 13.82 (NN_4) and 15.75 (NN_2). The ensemble of the different neural networks (NN_{1-5}) yields a similar performance as NN_5, and the ensemble of the tree-based methods and neural networks has a performance similar to the $GBRT$.

The risk-adjusted performance displayed in Panel B leads to the same order as the raw long-short returns. However, the increase in the six-factor alpha for the machine learning models compared to the linear models is even more pronounced as the six-factor model has less explanatory power. Furthermore, Panel C reveals that the neural network portfolios exhibit a smaller maximum drawdown and maximum one-month loss than the linear and tree-based models. The maximum drawdown (worst one-month return) in the case of the ensemble of neural networks is 17.82% (10.75%), whereas this number is 26.35% (13.97%) for the OLS. The superior performance of the machine learning models comes at the cost of a somewhat higher turnover. However, compared to the performance gains, this turnover increase from 89.27% for OLS to values between 89.61% for RF and 102.02 for NN_2 is relatively small.

Turning to the results for value-weighted portfolios in Panels D to E of Table 2.3 reveals identical qualitative conclusions, but the return spreads, t-statistics, and Sharpe ratios are substantially lower. Although the return forecasts derived from linear models already lead to economically and statistically significant long-short mean returns and six-factor alphas, the tree-based methods and neural networks do even better. Again, the neural network with two layers exhibits the highest t-statistics and Sharpe ratios while suffering from the mildest drawdowns. Comparing the ensemble of machine learning methods (ENS) with the linear OLS regressions shows performance gains of roughly 50% for the raw quintile returns and even higher for the risk-adjusted performance. In sum, allowing for non-linearities and interactions also leads to economically superior out-of-sample returns compared to traditional linear models, as summarized in Figure 2.4.

Figure 2.5 illustrates the results of Table 2.3 by plotting the equal-weighted and value-

Figure 2.4
Long-horizon performance of machine learning forecasts

This figure shows the Fama and French (2018) six-factor models alphas for various machine learning long-short portfolios. Stocks are sorted into country-neutral and value-weighted quintiles based on their predicted returns for the next month. The sorting breakpoints are based on big stocks only, which are in the top 90% of a country's aggregated market capitalization. The sample period is from January 2002 to December 2021.

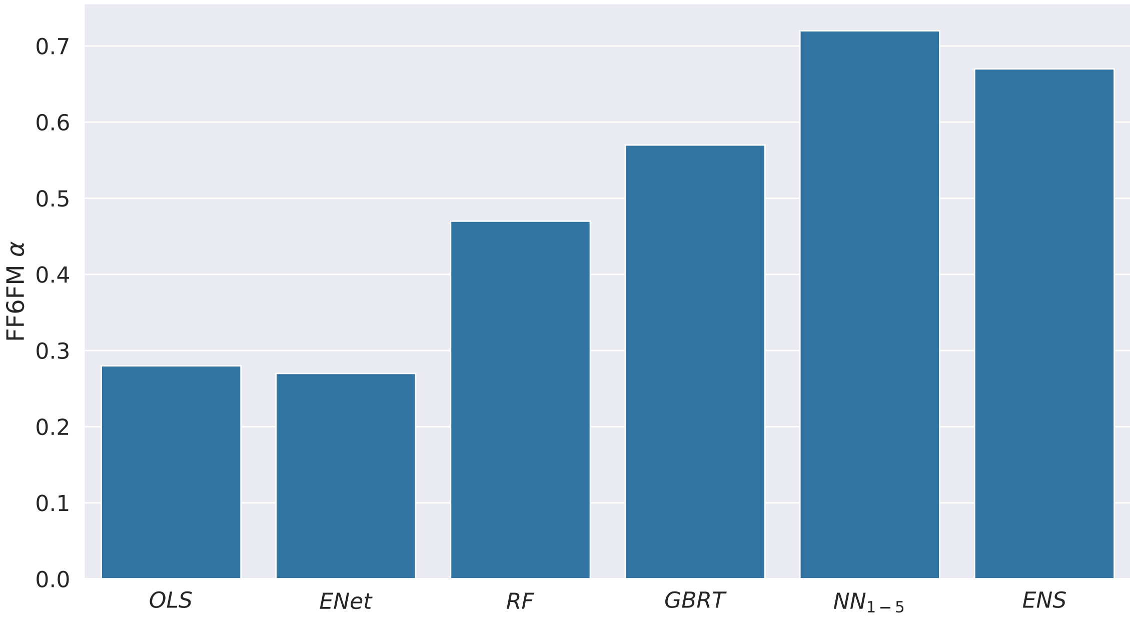

weighted cumulative performance of selected long-short strategies. We additionally include the cumulative performance for the long and short sides for select strategies in Appendix B. 2. Notably, the performance of our strategies does not predominantly stem from the short side, which would raise investability concerns due to shorting frictions.

Using a value-weighted portfolio strategy, RF initially dominates the other methods, while the outperformance of $GBRT$ and NN_{1-5} mainly stems from the period after 2009. As the ENS comprises all three methods, we observe a rather consistent outperformance versus OLS that is not driven by a particular period. In the case of equal-weighted portfolios, there are only small differences between the portfolio returns of $GBRT$, NN_{1-5}, and ENS till 2021. As for the value-weighted portfolios, the machine learning methods outperform the linear approaches consistently over time. The model with the lowest cumulative return is the $ENet$, whereby the underperformance versus the OLS is mainly driven by the first years of the sample period. Besides a sharp drawdown in 2009, there

Figure 2.5
Cumulative return of machine learning portfolios

The figure shows the cumulative log returns of long-short quintile portfolios sorted on the out-of-sample machine learning return forecasts. Panel A shows equal-weighted returns, while Panel B shows value-weighted returns. The sample period is from January 2002 to December 2021.

Panel A: Equal-Weighted

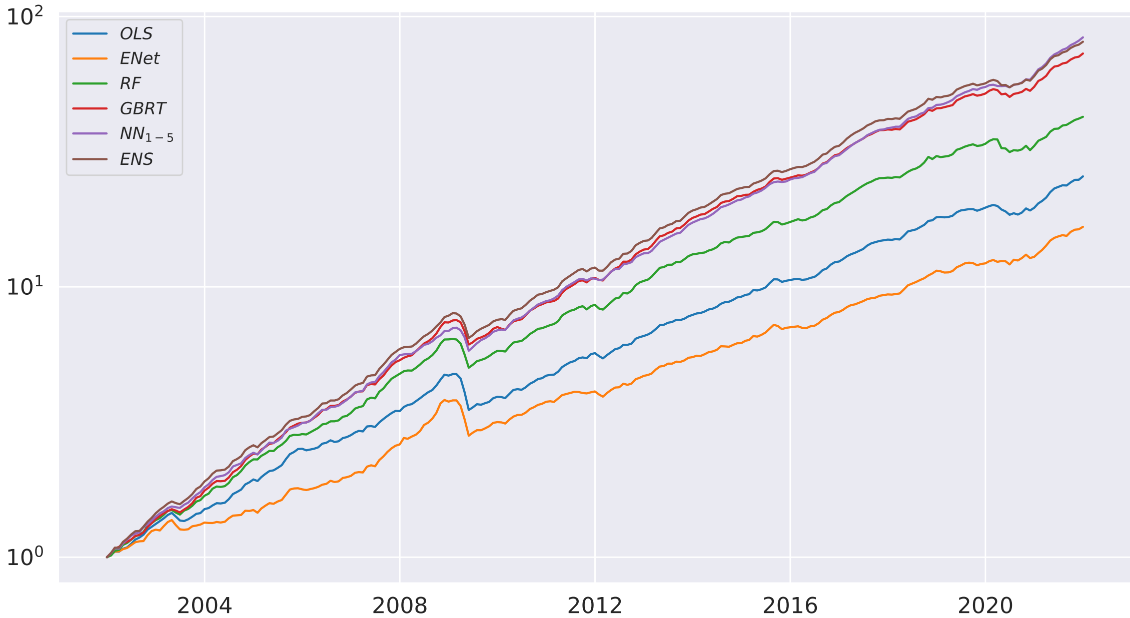

Panel B: Value-Weighted

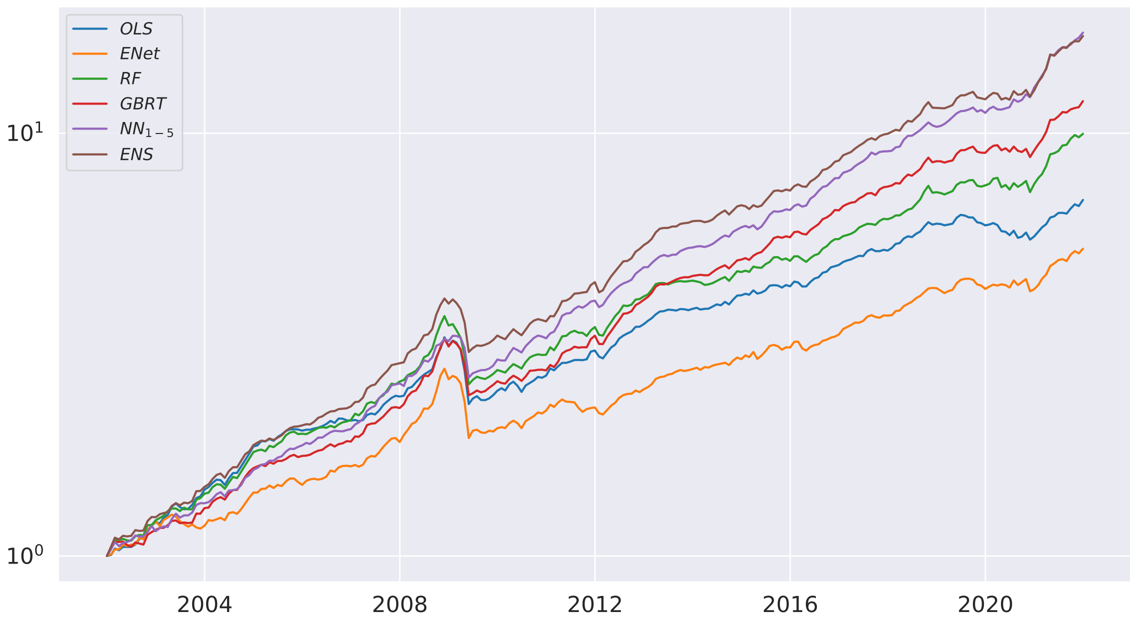

are no other notable downturns for all approaches. The drawdown in 2009 probably stems from the models' exposure to momentum that exhibited a momentum crash at that time (Daniel and Moskowitz, 2016; **?**). The recent global shock due to the COVID-19 pandemic in early 2020 did not lead to a significant portfolio-level downturn.

2.4.4 Robustness

To check the robustness of the results presented above, we will investigate (i) the impact of various methodological changes and (ii) the robustness within emerging market subregions.

Table 2.4 summarizes the robustness tests for methodological changes. We include various performance indicators for our equal-weighted and value-weighted machine learning portfolio strategies. However, we will only compare the performance of the benchmark OLS model to the ENS model. We select the ENS to be not driven by a look-ahead bias regarding the model selection and its portfolio performance. Besides the individual long-short return and the six-factor model of the two machine learning models, we include the results of the following two regressions in the last two rows of each panel:

$$\begin{aligned} r_{LS,t,ENS} &= \alpha + \beta_{OLS} r_{LS,t,OLS} + \epsilon_t, \\ r_{LS,t,OLS} &= \alpha + \beta_{ENS} r_{LS,t,ENS} + \epsilon_t. \end{aligned} \tag{2.9}$$

A positive and significant alpha indicates that the returns of the strategy on the right-hand side cannot fully explain the portfolio returns on the left-hand side.

The first two rows in Panel A show again our baseline result for OLS and ENS, as previously shown in Table 2.3. In addition, the last two rows of Panel A demonstrate that the ENS long-short portfolio spans the OLS long-short portfolio for both equal- and value-weighted portfolios, but the OLS portfolios cannot span ENS portfolios.

In Panel B, we construct our long-short trading strategy using decile instead of quintile sorts. By focusing on more extreme predicted abnormal returns and due to the monotonic increase among the portfolios, the equal-weighted and value-weighted long-short returns of the OLS increase to 1.84% (t-stat 10.09) and 1.18% (t-stat 5.91), whereas the Fama and French (2018) six-factor alpha increases to 1.41% (t-stat 11.30) and 0.55% (t-stat 4.66). The ENS shows an increase in the return to 2.50% (t-stat 13.93) and 1.66% (t-stat 8.12)

Table 2.4
Robustness

This table reports robustness tests for the out-of-sample performance of equal- and value-weighted long-short portfolios. All stocks are sorted into country-neutral quintile portfolios based on their predicted returns for the next month. We investigate predictions from a linear OLS model and an ensemble (ENS) of non-linear machine learning models (RF, $GBRT$, and NN_{1-5}). The sorting breakpoints are based on big stocks only, which are in the top 90% of a country's aggregated market capitalization. Panel A summarizes the baseline results as presented in Table 2.3. Panel B reports results on using decile sorts. Panel C uses an extended feature set following Hanauer and Lauterbach (2019). Panel D applies a feature selection before training the machine learning algorithms. Panel E uses predictions stemming from machine learning algorithms only trained on developed market data. Panel F excludes the high-turnover characteristics `Idiovol`, `LTurnover`, r_{2-1}, `SUV`, `Illiqu` from the feature set. Panel G shows the results for models trained on emerging market subregions. The first two rows of each panel provide the average monthly returns of the long-short quintile (Avg), corresponding t-statistics (t), the average Fama and French (2018) six-factor alphas (α), corresponding t-statistics (t_α), and R^2. The next two rows show spanning alphas (α), corresponding t-statistics (t_α), and R^2 when regressing the long-short ENS returns on OLS returns and vice versa. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	Equal-weighted					Value-weighted				
	Avg	t	α	t_α	R^2	Avg	t	α	t_α	R^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Baseline										
OLS	1.38	7.82	0.97	8.02	62.42	0.84	4.64	0.28	2.72	68.25
ENS	1.86	11.79	1.43	15.66	58.17	1.21	7.04	0.67	8.29	67.17
$ENS \sim OLS$			0.73	9.01	80.28			0.49	7.83	74.66
$OLS \sim ENS$			-0.44	-2.10	80.28			-0.22	-1.50	74.66
Panel B: Decile sorts										
OLS	1.84	10.09	1.41	11.30	52.39	1.18	5.91	0.55	4.66	62.37
ENS	2.50	13.93	2.02	18.31	54.16	1.66	8.12	1.10	10.30	57.37
$ENS \sim OLS$			1.01	7.10	71.88			0.78	6.00	56.06
$OLS \sim ENS$			-0.38	-2.76	71.88			-0.07	-0.36	56.06
Panel C: Extended feature set										
OLS	1.51	9.51	1.14	11.05	52.93	0.87	5.22	0.36	3.17	55.86
ENS	1.96	13.14	1.57	19.57	56.22	1.22	6.80	0.71	7.55	63.06
$ENS \sim OLS$			0.69	9.72	80.98			0.45	5.37	73.16
$OLS \sim ENS$			-0.38	-2.43	80.98			-0.13	-1.26	73.16
Panel D: Feature selection										
OLS	1.36	8.03	0.97	8.94	61.23	0.82	4.51	0.28	2.61	65.34
ENS	1.83	12.31	1.43	17.52	58.83	1.23	7.36	0.73	8.71	63.59
$ENS \sim OLS$			0.75	8.96	80.53			0.60	9.09	69.83
$OLS \sim ENS$			-0.50	-2.34	80.53			-0.29	-1.60	69.83
Panel E: Trained on developed markets										
OLS	1.29	6.71	0.93	6.76	62.40	0.89	4.67	0.38	3.37	68.17
ENS	1.67	10.55	1.23	10.12	59.97	1.20	6.64	0.62	5.31	61.15
$ENS \sim OLS$			0.72	6.99	79.14			0.43	4.75	74.18
$OLS \sim ENS$			-0.50	-3.84	79.14			-0.15	-1.20	74.18
Panel F: Excluding short-term feature set										
OLS	1.36	7.12	0.89	8.85	63.66	0.77	4.25	0.24	3.37	74.39
ENS	1.59	9.13	1.17	11.62	59.27	1.00	5.56	0.46	5.91	70.92
$ENS \sim OLS$			0.45	7.31	88.62			0.32	3.93	82.31
$OLS \sim ENS$			-0.32	-4.99	88.62			-0.16	-2.18	82.31
Panel G: Subregional training										
OLS	1.09	6.29	0.77	6.23	53.92	0.78	4.42	0.29	2.57	56.92
ENS	1.35	8.88	0.97	9.75	57.22	0.97	5.95	0.44	4.59	58.10
$ENS \sim OLS$			0.49	7.78	77.56			0.28	4.46	75.77
$OLS \sim ENS$			-0.23	-1.76	77.56			-0.06	-0.51	75.77

as well as in the risk-adjusted return to 2.02% (t-stat 18.31) and 1.10% (t-stat 10.30). Therefore, both OLS and ENS show stronger results when using decile sorts. Still, the increase in returns of the ENS is higher than the OLS, resulting in a larger α when regressing the ENS on the OLS compared to Panel A.

The feature set for the robustness test presented in Panel C includes the additional predictive characteristics described in Hanauer and Lauterbach (2019). The OLS particularly benefits from this extended feature set. The average equal-weighted and value-weighted long-short return increase by 9% and 4%, while only the equal-weighted return of ENS increases by 5%. In the case of the value-weighted risk-adjusted return, the OLS alpha increases by 28% and the ENS alpha increases by 6%.

Reducing the number of characteristics by applying a lasso regression, i.e., feature selection, before training the machine learning models reduces the equal-weighted and value-weighted long-short returns as well as the equal-weighted risk-adjusted returns of both machine learning models. Still, it increases the value-weighted alpha of the ensemble, as presented in Panel D.

In Panel E, we utilize machine learning models, which were never trained on emerging market stock returns; instead, the models are trained on developed markets (as defined by MSCI).[13] Although the models were solely trained on developed markets, we surprisingly do not observe a big performance loss but actually very similar returns. Furthermore, models that allow for non-linearities and interactions (ENS) still significantly outperform linear models (OLS). This indicates that machine learning models can create significant results even if they are evaluated on data from a totally different region.

For the robustness test in Panel F, we exclude the high-turnover characteristics, namely, `Idiovol`, `LTurnover`, $\mathtt{r}_{2-1}$, `SUV`, `Illiqu`, from the feature set. While the risk-adjusted returns of the OLS decrease by 8% and 15%, the ensemble is even more affected as the alphas are reduced by 19% and 31%. This indicates that these high-turnover features are relatively more important for more complex methods. However, even after excluding these characteristics, the long-short portfolios based on the ENS can span the long-short

[13] For the construction of the developed market sample, we follow the same procedure as for the emerging market sample. I.e., we use country-specific constituent lists, apply static and dynamic screens as outlined in Appendix B, compute the same set of features as for emerging markets, and estimate the machine learning models in the same way as for emerging markets.

portfolios constructed based on the OLS, while the converse is not the case.

In Panel G, we do not train our models on a pooled sample of all countries but separately for each of the following subregions: Central and Latin America (Americas); Asia; and Europe, Middle East, and Africa (EMEA). On the one hand, this allows the models to capture potential region-specific effects. On the other hand, each model is now trained on fewer data, which might be a drawback, especially for identifying non-linearities and interactions. We document that subregional training leads to inferior return forecasts than training models on pooled data from all subregions. This finding indicates that region-specific effects play a minor role compared to more data for out-of-sample returns. Furthermore, we find that the performance decay is more pronounced for the machine learning ensemble (ENS), i.e., indicating that more data is better for robustly identifying non-linearities and interactions.[14] Nevertheless, the OLS long-short portfolios cannot span the ENS long-short portfolio, but the ENS spans the OLS.

Finally, we assess if the superior performance of the machine learning return forecasts is robust across emerging market regions in Table 2.5. Therefore, we divide the countries of our full sample into three regions: Central and Latin America (Americas); Asia; and Europe, Middle East, and Africa (EMEA).

Overall, the results are robust for the different sub-regions Americas, Asia, and EMEA. Both OLS and ENS yield positive and significant long-short returns and alphas for both weighting schemes, but ENS exhibit higher returns and t-statistics. Furthermore, significant and positive alphas remain in the spanning regression of ENS on OLS for all sub-regions, but no positive spanning alphas remain when regressing OLS on ENS. Comparing the results across sub-regions, we find the strongest results for Asia and EMEA and a bit weaker but still highly significant results for Americas.

[14] Table B. 8 in the Appendix shows that the performance decline is most pronounced for the smaller regions Americas and EMEA while smaller for Asia. Furthermore, we compare the performance of models trained on pooled data with those trained solely on local country data. We restrict this analysis to stocks from the seven countries (Chile, Indonesia, Mexico, Malaysia, Philippines, Thailand, and Turkey) that are in our sample throughout the entire sample period. The results of this analysis are presented in Table B. 9. Consistent with our findings for models trained on subregional data, models trained on individual country data underperform global models.

Table 2.5
Regional performance

This table reports the out-of-sample performance of equal- and value-weighted long-short portfolios for emerging market subregions. All stocks are sorted into country-neutral quintile portfolios based on their predicted returns for the next month. We investigate predictions from a linear OLS model and an ensemble (ENS) of non-linear machine learning models (RF, $GBRT$, and NN_{1-5}). The sorting breakpoints are based on big stocks only, which are in the top 90% of the country's aggregated market capitalization. Panel A summarizes the baseline results as presented in Table 2.3, and Panel B shows the result for all countries being part of emerging Americas, Panel C combines all emerging Asian countries, and Panel D reports results for emerging countries from Europe, the Middle East, and Africa. The first two rows of each panel provide the average monthly returns of the long-short quintile (Avg), corresponding t-statistics (t), the average Fama and French (2018) six-factor alphas (α), corresponding t-statistics (t_α), and R^2. The next two rows show spanning alphas (α), corresponding t-statistic (t_α), and R^2 when regressing the long-short ENS returns on OLS returns and vice versa. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	Equal-weighted					Value-weighted				
	Avg	t	α	t_α	R^2	Avg	t	α	t_α	R^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Emerging Markets										
OLS	1.38	7.82	0.97	8.02	62.42	0.84	4.64	0.28	2.72	68.25
ENS	1.86	11.79	1.43	15.66	58.17	1.21	7.04	0.67	8.29	67.17
$ENS \sim OLS$			0.73	9.01	80.28			0.49	7.83	74.66
$OLS \sim ENS$			-0.44	-2.10	80.28			-0.22	-1.50	74.66
Panel B: Americas										
OLS	0.70	2.73	0.51	2.56	39.51	0.75	2.83	0.37	1.70	39.83
ENS	0.88	4.06	0.69	3.90	25.45	0.85	3.20	0.57	2.73	33.47
$ENS \sim OLS$			0.45	3.45	46.98			0.33	1.95	48.58
$OLS \sim ENS$			0.03	0.15	46.98			0.16	0.82	48.58
Panel C: Asia										
OLS	1.46	7.59	1.13	9.34	61.82	0.84	3.95	0.37	2.93	66.71
ENS	1.98	11.18	1.63	17.32	60.28	1.34	7.02	0.87	8.81	67.14
$ENS \sim OLS$			0.74	7.97	79.70			0.62	6.85	75.00
$OLS \sim ENS$			-0.40	-1.83	79.70			-0.33	-1.64	75.00
Panel D: Europe, the Middle East and Africa										
OLS	1.12	6.46	0.96	6.23	18.96	0.82	4.00	0.37	2.00	26.18
ENS	1.57	10.27	1.32	9.33	16.23	1.13	5.54	0.59	3.38	29.44
$ENS \sim OLS$			0.83	7.42	51.63			0.57	4.21	46.30
$OLS \sim ENS$			-0.11	-0.58	51.63			0.05	0.38	46.30

2.5 Understanding the sources of return predictability

The results so far provide evidence that return forecasts based on machine learning models lead to economically and statistically superior out-of-sample long-short returns compared to traditional linear models. To further understand the source of return predictability, we first investigate the performance of the two models in higher- versus lower-risk months. Second, we explore to what extent developed markets' long-short returns can explain emerging markets' long-short returns. Third, we turn to the time-series properties of the long-short machine learning portfolios over the next 36 months after portfolio formation. Fourth, we link the profitability of the machine learning models to several proxies for limits to arbitrage. Finally, we investigate the performance of an investment strategy that considers real-life investment frictions such as short-selling restrictions and transaction costs.

2.5.1 Performance in higher-risk versus lower-risk months

The profitability of return forecasts based on machine learning models may reflect risks not captured by the standard risk factors we control so far. Hence, we examine the performance of *OLS* and *ENS* forecasts during higher- versus lower-risk months. As proxies for risk, we apply whether (i) emerging markets as a whole go up or down, (ii) the rate on long-term U.S. government bonds is going up or down, (iii) the TED spread is below or above its median value, and (iv) the time-varying risk aversion index (RAbex) proposed by Bekaert, Engstrom and Xu (2022) is below or above its median value.[15] Splitting the sample period into up and down markets is done, for example, by Chan, Karceski and Lakonishok (1998), van der Hart, de Zwart and van Dijk (2005), or Asness, Frazzini and Pedersen (2019). The change in the U.S. government bond rate as a proxy for risk is motivated by the substantial financial instability experienced by emerging markets during the 'taper tantrum' in 2013 when U.S. yields surprisingly surged (cf., Estrada, Park

[15] The TED spread is defined as the difference between the LIBOR rate and the 3-month U.S. T-bill rate. The 10-year constant maturity U.S. Treasury rate (item DGS10) and the TED spread (item TEDRATE) are from the FRED database of the Federal Reserve Bank of St. Louis, and the time-varying risk aversion index RAbex is from Nancy Xu's website: https://www.nancyxu.net/risk-aversion-index.

and Ramayandi, 2016). According to Frazzini and Pedersen (2014), the TED spread is a gauge of funding conditions. Lastly, Bianchi, Büchner and Tamoni (2021) employ RAbex to investigate the link between time-varying risk aversion and excess bond returns.

Table 2.6

Higher-risk versus lower-risk periods

This table reports the out-of-sample performance of long-short portfolios in higher- versus lower-risk months. All stocks are sorted into country-neutral quintile portfolios based on their predicted returns for the next month. The sorting breakpoints are based on big stocks only, which are in the top 90% of the country's aggregated market capitalization. Risk proxies are whether emerging markets as a whole go up or down (Mkt), the rate on long-term U.S. government bonds is going up or down ($\Delta Yield$), whether the TED spread is below or above its median value (TED), and whether the time-varying risk aversion index proposed by Bekaert, Engstrom and Xu (2022) (RAbex) is below or above its median value ($RAbex$). Panel A (Panel B) summarizes the results from equal-weighting (value-weighting). The first two rows of each panel provide the average monthly long-short returns and corresponding t-statistics. The next two rows show spanning alphas and corresponding t-statistic when regressing the long-short ENS returns on OLS returns and vice versa. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	Mkt_{up}	Mkt_{down}	$\Delta Yield_{up}$	$\Delta Yield_{down}$	TED_{high}	TED_{low}	$RAbex_{high}$	$RAbex_{low}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Equal-Weighted								
OLS	1.07	1.82	1.60	1.17	1.19	1.55	1.13	1.63
	(3.82)	(11.81)	(7.25)	(6.20)	(3.60)	(9.93)	(3.55)	(12.48)
ENS	1.73	2.06	2.00	1.74	1.68	2.02	1.70	2.02
	(7.01)	(14.26)	(9.52)	(9.85)	(6.07)	(12.72)	(6.15)	(14.41)
$ENS \sim OLS$	0.84	0.47	0.63	0.81	0.73	0.68	0.80	0.56
	(7.60)	(5.43)	(6.31)	(6.60)	(7.08)	(4.95)	(8.07)	(4.25)
$OLS \sim ENS$	-0.64	0.03	-0.34	-0.52	-0.58	-0.13	-0.66	0.05
	(-2.56)	(0.23)	(-1.46)	(-2.13)	(-2.38)	(-0.81)	(-2.93)	(0.53)
Panel B: Value-Weighted								
OLS	0.64	1.12	1.06	0.63	0.69	0.96	0.74	0.93
	(2.13)	(5.51)	(4.69)	(3.31)	(2.03)	(6.49)	(2.29)	(6.30)
ENS	1.12	1.34	1.40	1.03	1.12	1.29	1.07	1.35
	(4.12)	(5.84)	(5.53)	(6.40)	(3.55)	(7.95)	(3.65)	(8.38)
$ENS \sim OLS$	0.57	0.41	0.47	0.52	0.53	0.45	0.45	0.48
	(6.47)	(3.88)	(4.36)	(6.37)	(6.84)	(5.03)	(4.94)	(5.27)
$OLS \sim ENS$	-0.37	0.08	-0.13	-0.30	-0.44	0.10	-0.27	-0.00
	(-2.01)	(0.54)	(-0.71)	(-1.78)	(-3.35)	(0.75)	(-1.55)	(-0.00)

Table 2.6 summarizes the top-bottom quintile returns for OLS and ENS for the different subsamples. For both equal- and value-weighted portfolios, we observe that the performance is somewhat higher in down-market months and months with rising bond yields. However, we also document higher returns when the TED spread is low, i.e., when funding conditions are better and in months with below-median risk aversion. Nevertheless, the quintile spreads are statistically significant and positive for all subsamples, prediction models, and weighting schemes. Furthermore, the difference between the subsamples is less pronounced for ENS than for OLS, and significant and positive alphas remain in the spanning regression of ENS on OLS. At the same time, the converse is not the case. This evidence suggests that the superiority of machine learning models compared to linear models in our sample does not stem solely from higher-risk months, at least for the definitions considered here.

2.5.2 Market integration

The robustness tests in Table 2.4 reveal an interesting finding: models trained solely on developed markets data perform similarly to models trained on emerging markets data in predicting emerging market stock returns. This result suggests that the pricing between developed and emerging markets could be more integrated, as indicated by the results for value and momentum returns in Cakici, Fabozzi and Tan (2013) and Hanauer and Linhart (2015). If developed and emerging markets are integrated, the developed market machine learning long-short portfolio returns would be able to explain the machine learning long-short portfolio returns for emerging markets, i.e., resulting in an insignificant α in the following regression for the global and regional emerging market samples:

$$r_{LS,Region_{EM},t,ENS} = \alpha + \beta_1 r_{LS,Global_{Dev},t,ENS} + \epsilon_t. \tag{2.10}$$

However, the results in Table 2.7 reveal that this is not the case. All alphas remain highly statistically significant for both the equal- and value-weighted portfolios. For the equal-weighted factor, Asia has the highest alpha of 1.47% (t-stat 8.66), followed by EMEA with 1.24% (t-stat 7.39). The value-weighted factor construction yields the highest alpha for Asia with 0.98% (t-stat 4.92), followed by EMEA with an alpha of 0.87% (t-stat 4.36).

Furthermore, the developed market long-short portfolio returns can only explain between 10% and 33% of the variation in emerging market long-short portfolio returns, which corresponds to correlations between 32% and 57%.[16]

Table 2.7
Market integration

This table reports summary statistics for regressions of emerging market regions' long-short returns on developed market's long-short returns. The long-short returns are based on ensemble (ENS) return forecasts of non-linear machine learning models (RF, $GBRT$, and NN_{1-5}) and are separately estimated for emerging and developed markets. All stocks are sorted into country-neutral quintile portfolios based on their predicted returns for the next month. The sorting breakpoints are based on big stocks only, which are in the top 90% of a country's aggregated market capitalization. Panel A (Panel B) summarizes the results of equal-weighting (value-weighting) of the prediction-sorted portfolios based on the different regional subsets. Each Panel provides the average monthly return of the long-short quintile (Avg), the alphas (α), betas (β), their corresponding t-statistics, and R^2 with respect to the developed market ensemble machine-learning factor. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	Avg	t	α	t_α	β	t_β	R^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Equal-weighted							
$Global_{EM}$	1.86	11.79	1.39	9.24	0.50	7.01	32.34
AME_{EM}	0.88	4.06	0.35	1.73	0.56	5.57	19.44
$ASIA_{EM}$	1.98	11.18	1.47	8.66	0.55	6.71	28.42
$EMEA_{EM}$	1.57	10.27	1.24	7.39	0.36	4.64	12.78
Panel B: Value-weighted							
$Global_{EM}$	1.21	7.04	0.89	5.73	0.49	4.95	28.17
AME_{EM}	0.85	3.20	0.53	2.30	0.49	3.68	14.37
$ASIA_{EM}$	1.34	7.02	0.98	4.92	0.54	3.95	22.47
$EMEA_{EM}$	1.13	5.54	0.87	4.36	0.40	4.55	10.73

Our interpretation of these results is that, although similar relationships between firm characteristics and future stock returns exist in both developed and emerging markets, the pricing of these characteristics is still not fully integrated. Furthermore, our results suggest that investors already applying machine learning strategies in developed markets may benefit from potential diversification benefits when applying such a strategy also in

[16] For the construction of the developed market sample, we follow the same procedure as for the emerging market sample. I.e., we use country-specific constituent lists, apply static and dynamic screens as outlined in Appendix B, compute the same set of features as for emerging markets, and estimate the machine learning models in the same way as for emerging markets, but now on developed market data. Furthermore, we report descriptive statistics of the developed market long-short machine learning portfolio in Table B. 10 of the appendix. The returns for the different developed market strategies are roughly half compared to their emerging market counterparts. However, we also document that models that allow for non-linearities and interactions (ENS) also significantly outperform linear models (OLS) in developed markets.

emerging markets.

2.5.3 Performance for longer holding periods

Is the profitability of the machine learning forecasts the result of temporary or permanent price changes? We analyze the long-run buy-and-hold returns following the methodology in Smajlbegovic (2019) and **?** to answer this question. First, we identify stocks used for constructing the long-short machine learning portfolios and calculate their value-weighted raw monthly returns for each month $t+k$, where $k \in \{1, ..., 36\}$. Second, we run a time-series regression of the the six-factor model for each holding period month k of the machine learning long-short factor. The corresponding average six-factor alpha for month k is the intercept (a_k) of the following regression:

$$r_{t+k,ML} - r_{f,t+k} = \alpha_k + \sum_{i}^{|f|} \beta_{i,k} f_{i,t+k} + \epsilon_{t+k}, \tag{2.11}$$

where $r_{t+k,ML} - r_{f,t+k}$ is the raw long-short return in month $t+k$ of stocks used for construction of the long-short machine learning factor in month t and $f_{i,t+k}$ indicates the individual factor returns of the six-factor model in month $t+k$: $RMRF_{t+k}$, SMB_{t+k}, HML_{t+k}, RMW_{t+k}, CMA_{t+k}, and WML_{t+k}. The intercept of the regression (α_k) is the alpha of the buy-and-hold strategy k months after portfolio formation, which is used to form the cumulative alpha in month k denoted as ACR_k:

$$ACR_k = \sum_{t=1}^{k} \alpha_t. \tag{2.12}$$

Figure 2.6 illustrates the value-weighted cumulative six-factor alpha of both OLS and ENS over a 36-month holding period. The figure reveals that both OLS and ENS can predict long-term returns and that their performance does not revert quickly. Together with the fact that standard risk factors cannot explain the performance of the strategies and the consistent performance over calendar time, we conclude that an underreaction explanation is more likely than an overreaction explanation. We further document that the superior performance of ENS compared to OLS is mainly driven by the first six months. Later both lines show a relatively parallel trend. This observation is unsurprising as the models are trained on one-month ahead returns and not longer periods.

Figure 2.6
Fama and French (2018) six-factor model alphas

This figure shows the average cumulative risk-adjusted return of the different machine learning long-short portfolios. First, we obtain the return of the portfolio formed at the end of month t for month t+k, where $k \in \{1, ..., 36\}$. Second, we run a time-series regression with the Fama and French (2018) six-factor model for the corresponding months. The regression intercept is defined as the average risk-adjusted portfolio return for the long–short portfolio at month t + k. In the final step, we compute the average holding period (cumulative) risk-adjusted return for the next k months since formation as the sum over the previous k months.

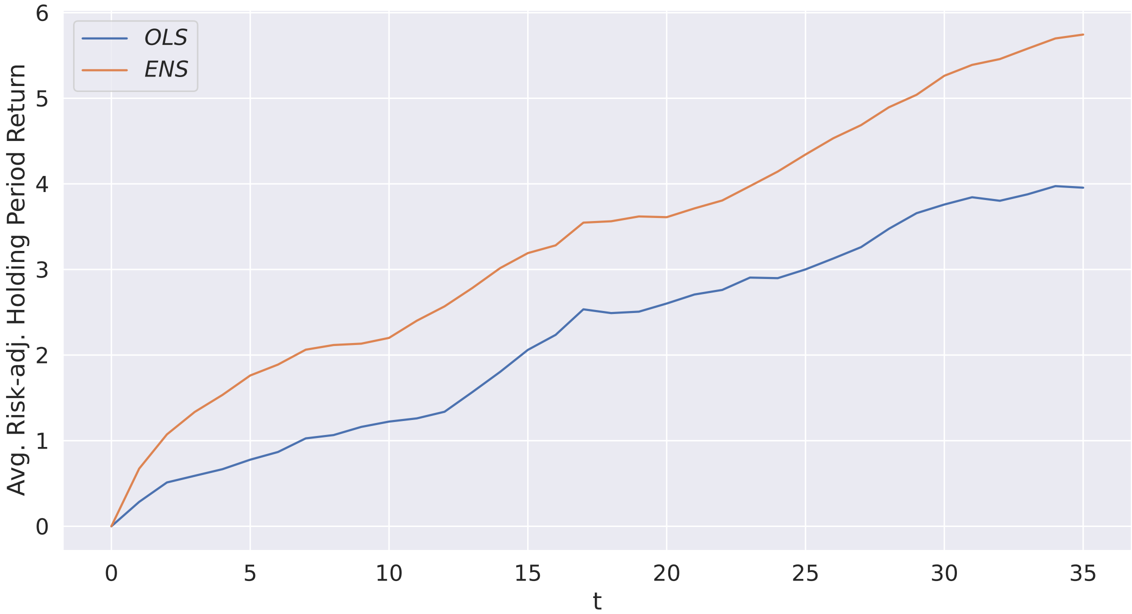

2.5.4 Limits to arbitrage

Our results thus far suggest that the high returns of the machine learning strategies in emerging markets cannot be explained by standard risk factors such as the factors of the Fama and French (2018) six-factor model and are consistent over time. Furthermore, the high returns do not primarily stem from higher-risk months and do not revert quickly. Therefore, a simple question arises: Why do investors not arbitrage away these abnormal returns? If limits to arbitrage hinder investors from doing that, we would expect that the predictability of the machine learning forecasts is concentrated in stocks with the highest limits to arbitrage.

To test whether the predictability of machine learning methods arises, at least in part, from such frictions, we interact the predicted returns of the machine learning models with different proxies for limits to arbitrage within a Fama and MacBeth (1973) regression. We additionally include both parts of the interaction term as controls as well as country dummies to account for any country effect yielding the following regression framework:

$$r_{i,t+1} - r_{f,t+1} = \alpha + \beta_1 ML_{i,t} + \beta_2 LTA_{i,t} + \beta_3 ML_{i,t} \times LTA_{i,t} + \beta_4 X_i + \epsilon_{i,t+1}, \quad (2.13)$$

where $LTA_{i,t}$ denotes the cross-sectional and country-neutral standardized variable measuring the limits to arbitrage of stock i while $ML_{i,t}$ is the predicted return based on the underlying machine learning model.

The coefficient β_3 is most relevant for this analysis as it indicates if the predictability of the different machine learning models is increasing with higher limits to arbitrage. We include three different variables that are closely related to limits to arbitrage and commonly used in the literature: size as a measure of information ambiguity (Zhang, 2006), idiosyncratic volatility as a proxy for arbitrage risk (Pontiff, 2006; Stambaugh, Yu and Yuan, 2015), and Amihud (2002) illiquidity as a potential proxy for transaction costs. If limits to arbitrage are important for the persistence of mispricing, we expect that predictability is the strongest for smaller stocks with high idiosyncratic volatility and low liquidity. Therefore, we additionally include the average of these three variables.

The results of this analysis are reported in Table 2.8. We first examine firm size's role in predicting future returns. Most small firms are less diversified and less fundamental information is available. In the case of fixed information acquisition costs, small firms are less attractive. The results in Columns (1) and (2) underline this hypothesis. The smaller the stock, the higher the return predictability for both methods.

The coefficient β_3 is most relevant for this analysis as it indicates if the predictability of the different machine learning models is increasing with higher limits to arbitrage. We include three commonly used variables that are closely related to limits to arbitrage: size as a measure of information ambiguity (Zhang, 2006), idiosyncratic volatility as a proxy for arbitrage risk (Pontiff, 2006; Stambaugh, Yu and Yuan, 2015), and Amihud (2002) illiquidity as a potential proxy for transaction costs. If limits to arbitrage are important for the persistence of mispricing, we expect that predictability is the strongest for smaller stocks with high idiosyncratic volatility and low liquidity. Therefore, we additionally include the average of these three variables ($COMBO$).[17]

The results of this analysis are reported in Table 2.8. We first examine firm size's role in predicting future returns. Most small firms are less diversified and less fundamental information is available. In the case of fixed information acquisition costs, small firms are less attractive. The results in Columns (1) and (2) support this hypothesis. The smaller the stock, the higher the return predictability for both methods.

Table 2.8
Limits to abitrage

This table reports the results of a Fama and MacBeth (1973) regression of next month's returns on machine learning return forecasts (ML), proxies for limits to arbitrage, and their interaction term. Each month, we conduct cross-sectional regressions of excess stock returns in month t + 1 on firms' ML return forecasts, limits-to-arbitrage scores, and their interaction terms, all from the end of the previous month t. The proxies for limits to arbitrage are: -1 × market capitalization ($SIZE$), idiosyncratic volatility ($IVOL$), Amihud illiquidity ($ILLIQ$), and a combination of the different proxies. All proxies for limits to arbitrage are ranked into the [-1,1] interval for each month and country. The t-statistics in parentheses are the corresponding Newey and West (1987) adjusted t-statistics with 4 lags. The sample period is from January 2002 to December 2021.

	$SIZE$		$IVOL$		$ILLIQ$		$COMBO$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
OLS	0.72		0.72		0.74		0.74	
	(9.43)		(9.16)		(9.55)		(9.27)	
ENS		1.11		1.13		1.15		1.13
		(13.81)		(14.31)		(14.58)		(12.88)
$LTA \times ML$	0.27	0.19	0.48	0.21	0.01	-0.06	0.52	0.27
	(5.36)	(3.27)	(9.65)	(3.85)	(0.23)	(-0.90)	(7.36)	(2.80)
LTA	0.15	0.12	0.06	0.12	0.21	0.19	0.27	0.26
	(2.42)	(1.93)	(0.68)	(1.45)	(2.78)	(2.48)	(3.06)	(2.94)
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2 (%)	15.00	15.22	15.10	15.30	15.10	15.34	15.01	15.23
Avg. Obs	4419	4419	4419	4419	4419	4419	4419	4419

In the second specification, we study how arbitrage risk affects the link between ma-

[17] We also report the average size, idiosyncratic volatility, illiquidity, and limit to arbitrage combination scores of the OLS and ENS portfolios in Table B. 11 of the appendix. All proxies for limits to arbitrage are ranked into the [-1,1] interval for each month and country, where higher values indicate higher limits to arbitrage. We find relatively similar exposures for OLS and ENS portfolios. The most pronounced differences are as follows: First, the short leg of the OLS strategy invests on average in below-average market capitalization stocks, while the short leg of the ENS strategy invests in above-average market capitalization stocks. Second, the long leg of the OLS strategy invests on average in stocks with above-average liquidity, while the long leg of the ENS strategy does this to a lower extent. In combination, this leads to a below-average exposure to limits of arbitrage for the long leg of the OLS strategy, while the long leg of the ENS strategy has only a slightly below-average exposure.

chine learning-based prediction and future stock returns. According to Pontiff (2006), arbitrageurs prefer to hold fewer stocks with higher idiosyncratic stock return volatility. Columns (3) and (4) provide empirical evidence that stocks with higher $IVOL$ exhibit larger predictable returns than less volatile stocks.

Next, we test how stock illiquidity relates to our previous findings. The intuition behind this proxy is based on the tradeability of the stock. The more illiquid the stock, the slower and more costly it should be to trade on the market. However, we are not able to provide empirical evidence that the return predictability of the machine learning models is driven by transaction costs.

Finally, we combine all three limits-to-arbitrage proxies to measure their mutual influence on the effect of future return predictability. Columns (7) and (8) provide evidence that stocks associated with more substantial limits-to-arbitrage characteristics exhibit stronger predictability independent of the underlying machine learning model.

However, we also find that the higher predictability for stocks with higher limits of arbitrage is less pronounced for the machine learning ensemble ENS than for the linear OLS regression, indicating the superiority of machine learning models in emerging markets does not stem from limits to arbitrage.

2.5.5 Further investment frictions

A common feature of the results presented above is that they are based on theoretical "zero-investment" long-short portfolio returns. However, it is questionable whether these returns can be realized in practice, as short-selling constraints may prevent the implementation of long-short strategies, and transaction costs may erode the strategy's profits. These constraints are particularly relevant for emerging markets (see, e.g., Roon, Nijman and Werker, 2001). Therefore, in this subsection, we limit ourselves to long-only portfolios of big stocks (i.e., also remove small stocks) and consider reasonable transaction costs. To estimate transaction costs, we compute the efficient discrete generalized estimator (EDGE) of the bid-ask spread for each stock and month, recently proposed by Ardia, Guidotti and Kroencke (2022). These bid-ask spread estimates vary considerably across time and stocks (cf., Figure B. 3 in the appendix) and therefore provide a more sophisticated estimate than the flat 100 basis points per single-trip used in van der Hart, Slagter and van Dijk (2003) and Hanauer and Lauterbach (2019).[18] The transaction cost per single-trip is one-half of the estimated bid-ask spread, and we define the transaction cost of portfolio L as:

$$\text{T-Cost}_{L,t} = \left(\sum_{i=1}^{N_{L,t-1 \cup t}} \left| w_{i,t} - \frac{w_{i,t-1}(1 + r_{i,t})}{1 + \sum_{j=1}^{N_{L,t}} w_{j,t-1} r_{j,t}} \right| \times \frac{S_{i,t}}{2} \right), \tag{2.14}$$

where $w_{i,t}$ is the weight of stock i at the end of month t, $r_{i,t}$ is the total return of stock i in month t, and $S_{i,t}$ is the estimated bid-ask spread. Furthermore, the net portfolio returns are defined as:

$$r_{L,t,net,ML} = r_{L,t,gross,ML} - \text{T-Cost}_{L,t}. \tag{2.15}$$

In the final step, we calculate the Fama and French (2018) six-factor model alpha return as:

$$r_{L,t,net,ML} - r_{f,t} = \alpha_{net} + \sum_{i}^{|f|} \beta_i f_{i,t,net} + \epsilon_t, \tag{2.16}$$

where $f_{i,t,net}$ is the risk factor return after transaction cost.

Furthermore, we also consider trading cost mitigation rules following Novy-Marx and Velikov (2016) and Blitz et al. (2022), which are common among practitioners. Such buy/hold strategies consist of the stocks that currently belong to the top X% plus the stocks selected in previous months that are still among the top Y% of stocks. In Table 2.9, we compare the quintile long-only strategy (20%/20%) with the transaction-cost-mitigation strategy buying the top 10% and holding them in our portfolio as long as they belong to the top 30% (10%/30%).

Table 2.9 reports the strategies' average gross excess over the market, their turnover and transaction costs, as well as the resulting net outperformance.[19] Limiting the investment universe to long-only portfolios of big stocks, we still see positive and significant gross

[18] Table B. 8 in the appendix also provides the results for transaction cost estimates of 100 basis points per single-trip.

[19] Our market portfolio is a reasonable proxy for the MSCI Emerging Market Index, a popular reference index for passive investment strategies. The mean return and standard deviation for the two time-series are 1.04% and 5.88%, and 1.00% and 5.99%, respectively. Furthermore, the correlation of the two time-series is 0.97. The small differences in the returns of our market portfolio and the MSCI Emerging Market Index may be due to MSCI's use of free float-adjusted market capitalization weighting and the inclusion of Chinese A-shares since 2018.

Table 2.9
Further investment frictions

This table reports the performance of different buy/hold long-only strategies before and after transaction-cost. The investment universe is limited to big stocks. We investigate predictions from a linear OLS model and an ensemble (ENS) of non-linear machine learning models (RF, $GBRT$, and NN_{1-5}). Every month the portfolio consists of the stocks that currently exhibit the highest X% forecasted returns per country plus those selected in previous months whose forecasted returns have not deteriorated beyond the top Y%. The first number in the column row names represents X, while the second represents Y. We report the strategies' gross returns in excess of the market, average turnover, transaction costs, net returns in excess of the market, and net Fama and French (2018) six-factor models alphas. We estimate one-way transaction costs as one-half of a stock's bid-ask spread, estimated as in Ardia, Guidotti and Kroencke (2022). All t-statistics are Newey and West (1987) adjusted with 4 lags. Panel A summarizes results from equal-weighting, while Panel B shows results from value-weighting. The sample period is from January 2002 to December 2021.

	OLS		ENS	
	20%/20%	10%/30%	20%/20%	10%/30%
	(1)	(2)	(3)	(4)
Panel A: Equal-weighted				
$r^e_{gross} - Mkt$	0.49	0.46	0.78	0.79
	(5.46)	(5.33)	(8.07)	(7.42)
TO (in %)	44.29	24.86	45.20	27.53
T-cost (in %)	0.31	0.18	0.32	0.20
$r^e_{net} - Mkt$	0.19	0.28	0.46	0.59
	(2.11)	(3.30)	(4.88)	(5.63)
α^{FF6}_{net}	0.29	0.39	0.55	0.67
	(4.91)	(6.19)	(7.66)	(8.28)
Panel B: Value-weighted				
$r^e_{gross} - Mkt$	0.32	0.32	0.47	0.48
	(3.39)	(3.47)	(4.88)	(4.66)
TO (in %)	44.48	22.16	45.51	23.40
T-cost (in %)	0.25	0.13	0.26	0.14
$r^e_{net} - Mkt$	0.07	0.19	0.21	0.34
	(0.75)	(2.08)	(2.20)	(3.31)
α^{FF6}_{net}	0.06	0.17	0.22	0.34
	(1.26)	(3.15)	(3.91)	(5.18)

outperformance for the top quintile portfolio (20%/20%) for both OLS and ENS and both weighting schemes. We observe similar gross outperformance when switching to the transaction cost mitigation strategies (10%/30%). However, the turnover and transactions are reduced by roughly 40%. This reduction in transaction costs substantially positively affects the net performance. For the equal-weighted strategies in Panel A, the net outperformance for OLS increases from 0.19% (t-stat 2.11) to 0.28% (t-stat 3.30). The net outperformance for ENS of 0.46% (t-stat 4.88) is also significant for the standard top quintile approach but also increases to 0.59% (t-stat 5.63) when applying a more efficient portfolio construction. Value-weighting the returns in Panel B leads to more challenging results. In this setup, the top OLS quintile yields only an insignificant net return of 0.07% (t-stat 0.75). Applying the trading-cost mitigation strategy increases the net returns to 0.19% (t-stat 2.08) for OLS and even to 0.34% (t-stat 3.31) for ENS. Similar results can be derived by comparing the Fama and French (2018) net alphas for which the turnover-reducing strategy for ENS exhibits again the highest net alpha of 0.34 (t-stat 5.18).[20] Therefore, we conclude that machine learning-based return forecasts can lead to significant net outperformance and net alphas, at least when efficient trading rules are applied.

2.6 Conclusion

This paper compares the out-of-sample predictive power of various machine learning models for a broad sample of 32 emerging market countries and a 20-year out-of-sample period. More specifically, we use both linear and more complex algorithms that allow for non-linearities and interactions.

We document that the different prediction algorithms identify similar characteristics. However, we also observe that tree-based methods and neural networks identify non-linearities and interactions of characteristics. Furthermore, return forecasts based on machine learning models lead to economically and statistically superior out-of-sample long-short returns compared to traditional linear models. This finding is robust to several methodological choices and for emerging market subregions.

We also find that developed market long-short returns based on machine learning forecasts derived in the same way as their emerging market counterparts cannot explain emerging market out-of-sample returns. However, models estimated solely on developed markets data can also predict emerging market stock returns. This finding indicates that similar relationships between firm characteristics and future stock returns exist for developed and emerging markets but that the pricing of these characteristics is not fully integrated

[20] When applying the more conservative transaction cost estimates of 100 basis points per single-trip, only the machine learning ensemble in combination with transaction cost mitigation exhibits significant net returns and alphas of 0.23% and 0.25%, respectively.

between developed and emerging markets.

Furthermore, we also document that the high returns of the machine learning strategies in emerging do not primarily stem from higher-risk months and do not revert quickly, suggesting that an underreaction explanation is more likely than a risk-based explanation. Although both linear and machine learning models show higher predictability for stocks associated with higher limits to arbitrage, we also document that this effect is less pronounced for machine learning forecasts than for linear regression forecasts. This finding indicates that the superiority of machine learning models in emerging markets does not stem from limits to arbitrage. Finally, accounting for transaction costs, short-selling constraints, and limiting our investment universe to big stocks only, we document that machine learning-based return forecasts can lead to significant net outperformance over the market and net alphas, at least when efficient trading rules are applied.

3 Anchoring and Global Underreaction to Firm-Specific News

Abstract

This paper investigates the anchoring effect as an explanation for investor underreaction to global firm-specific news. The anchoring effect refers to the tendency of investors to stick to their initial beliefs about a stock, even when facing new information. The paper adopts a novel high-frequency methodology to identify news events and covers stocks from 23 developed countries from January 2004 to December 2021. Our results provide evidence of investors' distorted belief updating process and show that the anchoring of investors induced by the 52-week high impacts the processing of the firm-specific news. Regression analyses decompose stock returns into three independent components and reveal that the interaction effect between the firm-specific news return and the nearness to the 52-week high are related to a significant risk-adjusted return.

Key words:	Portfolio, Choice Asset Pricing, Information and Market Efficiency, Belief, Stock-Returns
JEL Codes:	G11, G12, G14, D83, L14
Authors:	Tobias Kalsbach, Steffen Windmüller
First Author:	Tobias Kalsbach
Presentations:	2022 TUM School of Management
Current Status:	Working Paper
Note:	In this chapter I use the first-person plural narrative.

3.1 Introduction[1]

Investor underreaction to the arrival of news has been a long-standing topic in the finance literature. A large body of empirical and theoretical evidence argues that firms' stock prices respond slowly due to investors' behavioral biases. Theoretical literature often suggests that investors' limited attention results in underreaction to the arrival of news.[2] At the same time, empirical evidence supports this limited attention hypothesis by showing that firms' stock prices respond slowly to the arrival of new information.[3] This paper aims to test a novel psychological explanation, the anchoring effect, as an additional explanation for investor underreaction to global firm-specific news measured through the nearness to the 52-week high price.[4] The anchoring effect, as introduced by Tversky and Kahneman (1974), refers to the tendency of investors to stick to their initial beliefs about a stock, even when facing new information, as suggested by Coibion and Gorodnichenko (2015). This psychological barrier can be enforced when investors use the 52-week high as an anchor when making investment decisions. For example, investors influenced by the anchoring effect will not fully adjust their beliefs if the firm experiences the arrival of positive news (negative news) and if the stock price is close to (far from) the 52-week high, leading to a slow stock price response. Following this, our central hypothesis is closely related to the main argument of the anchoring bias literature, namely that investors do not fully incorporate the new information into their beliefs due to the anchoring effect resulting in the predictability of future stock returns (George and Hwang, 2004; Hong, Jordan and Liu, 2015; Huang, Lin and Xiang, 2021) and earnings surprises (Birru, 2013).

---

[1] We thank Matthias Hanauer, Christoph Kaserer, Lisa Knauer, and Laurens Swinkels for their helpful comments and suggestions. Any remaining errors are our own.

[2] Several behavioral theories that focus on investors' underreaction to public news have been proposed (Daniel, Hirshleifer and Subrahmanyam, 1998; Hong and Stein, 1999; Hirshleifer, Lim and Teoh, 2011; Peng and Xiong, 2006).

[3] Investors limited attention causes an underreact to different information types. The first type is the release of earnings information (cf., Ben-Rephael, Da and Israelsen, 2017). The second type covers general news from peer economically linked firms (cf., Ali and Hirshleifer, 2020). The last type covers the underreaction to news-driven returns (cf., Jiang, Li and Wang, 2021).

[4] The nearness of the firm's stock price or market to its 52-week high comes with a change in investor trading behavior (cf., Huddart, Lang and Yetman, 2009). The 52-week is often associated with two different investor biases. The disposition bias causes the sale (buying) of stocks trading at a historical high (low) (Heath, Huddart and Lang, 1999; Poteshman and Serbin, 2003). The second bias the nearness to the 52-week high is associated with is the anchoring bias. This bias leads to the investor's underreaction to news (cf., Huang, Lin and Xiang, 2021), and helps to explain the stocks momentum anomaly (cf., Hung, Lin and Yang, 2022).

We adopt the high-frequency methodology of Jiang, Li and Wang (2021) to identify all unscheduled and scheduled news events, which allows us to estimate monthly firm-specific news returns. The sample for the empirical analysis is limited to stocks from the developed markets covering 23 global countries from January 2004 to December 2021. We restrict the sample to this period and markets due to the increase in global news coverage starting in 2004 and the technical requirement of the portfolio sort analysis of having a firm-specific news event in the previous month. We begin by forming independent, country-neutrally double-sorted quintile portfolios using the last month's firm-specific news return (FN) and the nearness to the 52-week high ($NEAR$) at the previous month-end as sorting criteria.

The firm-specific news return and the nearness to the 52-week high yield positive and significant risk-adjusted returns, providing out-of-sample evidence for additional countries and an extended sample period. Next, we follow a portfolio strategy utilizing the anchoring bias in combination with the arrival of firm-specific news. The long (short) leg of the strategy incorporates stocks near (far from) their 52-week high and experiencing extremely positive (negative) firm-specific news. The short leg of the strategy yields a monthly average Fama-French-Carhart (1997) four-factor alpha of -0.30% (t=-2.37), whereas the long side earns an alpha of 1.14% (t=9.60). The combined long-short strategy returns an alpha of 1.44% (t=9.37), which provides the first evidence of the investors' distorted belief updating process. It is important to denote that the portfolio with bad news but a near 52-week high as well as the portfolio with very good news but far from the 52-week high both do not yield any significant returns, further underlining our hypothesis that investors only underreact to good (bad) news if the stock price is near (far from) its 52-week high.

To further analyze how the anchoring of investors induced by the 52-week high impacts the processing of the firm-specific news, we follow the innovative decomposition methodology by George, Hwang and Li (2014). We perform a similar Fama and MacBeth (1973) cross-sectional regression analysis to decompose the returns of the double-sorted portfolios into three independent components. The first component captures the interaction effect between the firm-specific news return and the nearness to the 52-week high and measures the degree to which the 52-week high effect causes the underreaction to the firm-specific news. The second component focuses solely on the pure firm-specific news effect, and the last part yields the return attributable to the pure nearness to the 52-week high effect.

The interaction effect yields an average Fama-French-Carhart (1997) four-factor alpha of 1.47% (t=4.67). In contrast, the pure firm-specific news effect and pure 52-week high effect are insignificant, earning a risk-adjusted return of -0.13% (t=-0.66) and 0.10% (t=0.61), respectively. Excluding the interaction effects from the regression results in two positive and significant pure effects with an alpha of 0.68% (t=10.20) in the case of the pure firm-specific news and an alpha of 0.55% (t=3.88) in the case of the pure 52-week high. These results allow us to conclude that the investors' underreaction to the firm-specific news is partially explained by the anchoring bias induced by the nearness to the 52-week high.

Next, we investigate the role of a stock's limits to arbitrage in causing mispricing. Our results provide evidence that firms indeed drive the induced underreaction of investors with high limits to arbitrage. The effect exists among stocks that are smaller in market capitalization, have lower institutional ownership or analyst coverage, have higher idiosyncratic volatility, and have higher transaction costs.

In several robustness tests, we underline the persistence of our results. By applying different factor models, the risk-adjusted returns of the interaction effect, the pure firm-specific news effect, and the 52-week high effect do not change. Additionally, we provide results on six variations of our firm-specific news measure. By limiting the firm-specific news to earnings announcement days, we find that the risk-adjusted return of the interaction effect is reduced and loses its significance, yielding a global alpha of 1.21% (t=1.60) per month. These results are robust across U.S. and non-U.S. firms, indicating that investors quickly incorporate scheduled news into the stock prices. By excluding earnings announcement days, we find that the global four-factor alpha increases to 1.68% (t=5.17) per month. By modeling a slower information diffusion process, the interaction effect in the most efficient stock market, the U.S., becomes insignificant, yielding a monthly return of 0.52% (t=1.19) per month. Further exclusion of macroeconomic announcements and the predictable component from daily returns resulted in a global risk-adjusted return of interaction effect of 1.38% (t=3.86) and 1.32% (t=4.30), respectively. In the last robustness test, we tackle the concern that the 52-week high is just a replacement of the stock's momentum (MOM). Thus, we replace the nearness to the 52-week high with momentum and run a placebo test. In this case, the risk-adjusted return in the Fama-French-Carhart (1997) model is negative and insignificant. The pure firm-specific news and momentum effects are positive and significant, independent of including or excluding the interaction effect. This provides evidence that the 52-week high is not just another sort on momentum and that the nearness to the 52-week high proxies the investor's underreaction to firm-specific news.

Lastly, we explore how the nearness to the 52-week high distorts the belief-updating process leading to an underreaction. We use analyst recommendation changes as a direct proxy to observe the belief updating process in financial markets. Our results suggest that analysts are indeed influenced by firm-specific news as they change their recommendations after the arrival of news. However, the upgrade (downgrade) is less likely if positive (negative) news arrives at the firm and the underlying stock price is near (far from) the 52-week high. The findings provide evidence for the hypothesis that the belief updating process is distorted and influenced by stock prices' nearness to the 52-week high and the arrival of firm-specific news.

This study adds to understanding investor underreaction in at least four aspects in an international asset pricing context.

First, we contribute to a better understanding of investor underreaction by explicitly using firm-specific news as the cause. Using the novel high-frequency approach introduced in Jiang, Li and Wang (2021), we show that investors not only underreact to focal firm news in the U.S. but also in non-U.S. equity markets. In contrast to Huang, Lin and Xiang (2021), who investigate investor underreaction by proxying news with economically-linked, past-month firm momentum, our news measure utilizes firm-specific news of the focal firm. We provide insights into investor underreaction by showing that limits to arbitrage amplify the underreaction potential, i.e., investor underreaction increases with higher limits to arbitrage.

Second, our paper reveals a crucial economic mechanism behind investor underreaction in global equity markets. We show that the underreaction to firm-specific news disappears when controlling for its interaction with investors' anchoring bias. Prior theoretical studies suggest that investors' underreaction to new information, such as earnings news, can be attributed to different psychological biases (Barberis, Shleifer and Vishny, 1998; Daniel, Hirshleifer and Subrahmanyam, 1998) and limits to arbitrage (Shleifer and Vishny, 1997). Empirically, the mispricing caused by the anchoring bias can be partially explained by the firms' exposure to limits to arbitrage (Byun, Goh and Kim, 2020). We explore the economic mechanism causing the underreaction. We, therefore, rely on the anchoring and adjustment hypothesis by showing that professional forecasters (Campbell and Sharpe, 2009; Cen, Hilary and Wei, 2013) include the firm-specific news in their recommendation but are affected by the anchoring bias if the stock is near (far from) the 52-week high and positive (negative) news arrives.

Third, we show that unscheduled, firm-specific news drives the anchoring bias effect on investors' underreaction over the subsequent month. Empirical evidence so far suggested that limits to arbitrage are an important driving force behind investors' underreaction to new information on the earnings announcement date (Hung, Li and Wang, 2015). Moreover, Birru (2013) and George, Hwang and Li (2014) find that investors' underreaction is driven by scheduled news, respectively, earnings announcements when quantified by price changes over the subsequent days. In contrast, we consider all news releases over the previous month to measure their pricing impact within the current month. Our results on the investors' distorted belief updating process provide strong evidence on a longer-dated, monthly investor underreaction to unscheduled news, indicating that unscheduled news items require more time to be reflected within stock prices.

Fourth, we contribute to the literature on empirical asset pricing for global equity markets by using an international sample and extended metrics. According to Karolyi (2016), most of today's published studies in top finance journals focused on the United States. In this regard, most literature on news-induced momentum (Chan, 2003; Gutierrez and Kelly, 2008; Hillert, Jacobs and Müller, 2014; Jiang, Li and Wang, 2021) concentrates solely on the U.S. stock market. Therefore, we add to the ongoing discussion about the investor underreaction hypothesis and its economic channels by providing non-U.S. out-of-sample evidence (Hou, Xue and Zhang, 2018) for the anchoring bias and investor underreaction to firm-specific news.

The remainder of the paper is structured as follows: Section 3.2 defines the main variables and the return decomposition methodology, which allows us to derive the interaction effect and the two pure effects of the firm-specific news and 52-week high. Section 3.3 introduces the global dataset. Section 3.4 presents the empirical results on the anchoring effect and the underreaction to firm-specific news. Section 3.5 concludes.

3.2 Empirical Strategy

In this section, we present our empirical strategy for measuring how the nearness to the 52-week high distorts the belief-updating process of an investor after the arrival of firm-specific news and explains investor underreaction. In the first subsection, we present the underlying methodology to construct the two required signals, firm-specific news and the nearness to the 52-week high. The second subsection focuses on the decomposition methodology, which allows us to differentiate between the pure effect of firm-specific news, the pure effect of the nearness to the 52-week high, and their interaction effect.

3.2.1 Firm-specific news and nearness to the 52-week high

This paper adapts the high-frequency decomposition methodology by Jiang, Li and Wang (2021) to identify scheduled and unscheduled firm-specific news.[5] We then relate the investor underreaction to firm-specific news to investors' distorted belief updating process, similar to George and Hwang (2004), Birru (2013), and Huang, Lin and Xiang (2021). This distortion is driven by the psychological barrier imposed by the nearness to the 52-week high price. If the current stock price is near the 52-week high and positive firm-specific news arrives at the firm or the price is far from the 52-week high, and negative firm-specific news arrives, investors are not willing to update their beliefs about the firm's fundamentals due to the anchoring effect.

Our measure of firm-specific news combines daily stock returns with firm-specific news events to decompose daily stock returns into news-driven and non-news-driven returns based on market reactions to firm-specific news releases.[6] To calculate the daily firm-specific news returns, we rely on the regular trading hours of the individual stock exchanges a stock is traded on. If the news is released within regular trading hours of day t, the news

[5] The method of Jiang, Li and Wang (2021) has several advantages over other low-frequency news types. One regularly used news type differentiates between cash-flow news and discount-rate news estimated through vector auto-regression (Campbell and Shiller, 1988; Campbell, 1991; Vuolteenaho, 2002), implied cost of capital (Chen, Da and Zhao, 2013), and analyst estimations (Easton and Monahan, 2005; Da and Warachka, 2009; Da, Liu and Schaumburg, 2014). Other regression-based news types differentiate between market-wide and firm-specific news (Roll, 1988; Morck, Yeung and Yu, 2000) and additionally noise (Brogaard et al., 2022).

[6] In the study of Jiang, Li and Wang (2021), the authors use high-frequency, intraday data for U.S. stocks. Due to global data non-availability, we are restricted to daily stock returns. However, within robustness tests, Jiang, Li and Wang (2021) show that their high-frequency-based results hold when using daily instead of intraday data to identify firm-specific news returns.

return equals the respective daily return. For news occurring after the closing of a stock's main stock market (i.e., an overnight release), at the weekend, or on a holiday, the news is incorporated into the return of the next trading day $t+1$. If no firm-specific news occurs, we declare the return of this day as non-news. We further aggregate the daily firm-specific news returns to a monthly level, similar to the intraday news aggregation to a daily level in Jiang, Li and Wang (2021). Suppose there are M trading days per month. Let $fn_{i,t,n}$ be the mth daily news-driven return for stock i in month t, where m = 1,2,...,M, we can compute the monthly firm-specific focal news return ($FN_{i,t}$) as follows:

$$FN_{i,t} = (\prod_{m=1}^{M} 1 + fn_{i,t,m}) - 1 \times 100 \tag{3.1}$$

To determine the impact of the psychological barrier and the distortion in the belief updating, we need to derive the stock's nearness to the 52-week high. We, therefore, follow George and Hwang (2004) and Windmüller (2022) and define:

$$NEAR_{i,t} = \frac{UP_{i,t}}{\max\limits_{0 \leq d \leq 52} UP_{i,t-d}}, \tag{3.2}$$

where $UP_{i,t}$ is the unadjusted stock price of stock i at end of the previous week t.

3.2.2 Decomposition methodology

To shed light on the distortion in the belief-updating process, we follow the methodology proposed by George, Hwang and Li (2014) and Huang, Lin and Xiang (2021). Therefore, we first sort all stocks that experienced a firm-specific news arrival based on their nearness to the 52-week high and their firm-specific news return into two independent country-neutral quintile portfolios. Afterward, we utilize two different Fama and MacBeth (1973) regressions to decompose the returns of the double-sorted portfolios. In the first regression, we run a monthly stock-level Fama and MacBeth (1973) regression to estimate the two pure effects of firm-specific news and the nearness to the 52-week high as well as the interaction effect of both across the 5x5 = 25 portfolios. The regression model is specified as follows:

$$\begin{aligned} R_{i,t+1} = b_0 &+ b_1 FN^5_{i,t} + b_2 FN^4_{i,t} + b_3 FN^2_{i,t} + b_4 FN^1_{i,t} + b_5 NEAR^5_{i,t} + b_6 NEAR^1_{i,t-1} \\ &+ b_7 FN^5_{i,t} \times NEAR^5_{i,t} + b_8 FN^4_{i,t} \times NEAR^5_{i,t} + b_9 FN^2_{i,t} \times NEAR^5_{i,t} \\ &+ b_{10} FN^1_{i,t} \times NEAR^5_{i,t} + b_{11} FN^5_{i,t} \times NEAR^1_{i,t} + b_{12} FN^4_{i,t} \times NEAR^1_{i,t} \\ &+ b_{13} FN^2_{i,t} \times NEAR^1_{i,t} + b_{14} FN^1_{i,t} \times NEAR^1_{i,t} + \epsilon, \end{aligned} \tag{3.3}$$

where $R_{i,t+1}$ is the stock return of firm i in the next month $t+1$, and right-hand-side variables are dummies indicating the quintile ranking of firm i at the end of the month t for FN and $NEAR$. In the second regression, we exclude the interaction effect from the model, leaving us only with the estimation of the two pure effects of firm-specific news and the nearness to the 52-week high:

$$R_{i,t+1} = b_0 + b_1 FN^5_{i,t} + b_2 FN^4_{i,t} + b_3 FN^2_{i,t} + b_4 FN^1_{i,t} + b_5 NEAR^5_{i,t} + b_6 NEAR^1_{i,t-1} + \epsilon \tag{3.4}$$

In Table 3.1, we describe the methodology by George, Hwang and Li (2014), and Huang, Lin and Xiang (2021) on how the individual average portfolio return in each of the 5×5 portfolios sorted by the firm-specific news return and the nearness to the 52-week high is decomposed by the regression parameters and the return components. The lowest nearness to the 52-week high (firm-specific news return) quintile is defined as $NEAR1$ ($FN1$), while the highest nearness to the 52-week high (firm-specific news return) quintile is specified as $NEAR5$ ($FN5$). Similar to Huang, Lin and Xiang (2021), we merge the $NEAR2$, $NEAR3$, and $NEAR4$ quintiles into one group (referred to as $NEAR2 \sim 4$ in Table 3.1) since it is assumed that the nearness to the 52-week high only exists in the two most extreme $NEAR$ portfolios.

Table 3.1

Specification of return decompostition

This table describes the specification of the return decomposition as in George, Hwang and Li (2014) and Huang, Lin and Xiang (2021) by regression parameter and return component for the double-sorted firm portfolios by firm-specific news returns (FN) and nearness to the 52-week high ($NEAR$). To form the double-sorting portfolios we sort each month all firms which experienced firm-specific news arrival into independent and country-neutral 5×5 portfolios based on FN in the previous month and $NEAR$ at the previous month-end. Each cell represents a group of stocks with a particular $NEAR$ and FN ranking. Portfolios with $NEAR$ ranked in the $NEAR2$, $NEAR3$, and $NEAR4$ quintiles are combined into one group in the return decomposition. In Panel A (Panel B), we show how the respective portfolio return can be decomposed using the regression parameters from the monthly stock-level Fama and MacBeth (1973) regression as specified in Equation 3.3 (Equation 3.4). In Panel C (Panel D), we show how the respective portfolio return can be decomposed into different return components. The return components can be disentangled into the benchmark return (μ), the returns associated with the 52-week high (H), the returns attributable to the firm-specific news (N), and the returns associated with the interaction between the firm-specific news and nearness of the stock price to the 52-week high (I). μ reflects the average return of stocks in the portfolio with neither extreme firm-specific news returns nor an extreme nearness to the 52-week high. H reflects the returns associated with being near (n), middle (m), or far (f) from the 52-week high, regardless of the FN ranking. N reflects the returns associated with having extremely good (gg), good (g), bad (b), or extremely bad (bb) firm-specific news about the firms, regardless of the $NEAR$ ranking. I reflects the returns associated with having both good (bad) firm-specific news about the firm and stock prices near (far from) the 52-week high.

	$FN1$	$FN2$	$FN3$	$FN4$	$FN5$
	(1)	(2)	(3)	(4)	(5)
Panel A: Decomposition by regression parameter *including* the interactions effect					
$NEAR1$	$b_0+b_4+b_6+b_{14}$	$b_0+b_3+b_6+b_{13}$	b_0+b_6	$b_0+b_2+b_6+b_{12}$	$b_0+b_1+b_6+b_{11}$
$NEAR2 \sim 4$	b_0+b_4	b_0+b_3	b_0	b_0+b_2	b_0+b_1
$NEAR5$	$b_0+b_4+b_5+b_{10}$	$b_0+b_3+b_5+b_9$	b_0+b_5	$b_0+b_2+b_5+b_8$	$b_0+b_1+b_5+b_7$
Panel B: Decomposition by regression parameter *excluding* the interactions effect					
$NEAR1$	$b_0+b_4+b_6$	$b_0+b_3+b_6$	b_0+b_6	$b_0+b_2+b_6$	$b_0+b_1+b_6$
$NEAR2 \sim 4$	b_0+b_4	b_0+b_3	b_0	b_0+b_2	b_0+b_1
$NEAR5$	$b_0+b_4+b_5$	$b_0+b_3+b_5$	b_0+b_5	$b_0+b_2+b_5$	$b_0+b_1+b_5$
Panel C: Decomposition by return component *including* the interactions effect					
$NEAR1$	$\mu+H_f+N_{bb}+I_{bb,f}$	$\mu+H_f+N_b+I_{b,f}$	$\mu+H_f$	$\mu+H_f+N_g$	$\mu+H_f+N_{gg}$
$NEAR2 \sim 4$	$\mu+N_{bb}+I_{bb,m}$	$\mu+N_b+I_{b,m}$	μ	$\mu+N_g+I_{g,m}$	$\mu+N_{gg}+I_{gg,m}$
$NEAR5$	$\mu+H_n+N_{bb}$	$\mu+H_n+N_b$	$\mu+H_n$	$\mu+H_n+N_g+I_{g,n}$	$\mu+H_n+N_{gg}+I_{gg,n}$

Continued on next page

Table 3.1 continued

	$FN1$	$FN2$	$FN3$	$FN4$	$FN5$
	(1)	(2)	(3)	(4)	(5)
Panel D: Decomposition by return component *excluding* the interactions effect					
$NEAR1$	$\mu + H_f + N_{bb}$	$\mu + H_f + N_b$	$\mu + H_f$	$\mu + H_f + N_g$	$\mu + H_f + N_{gg}$
$NEAR2 \sim 4$	$\mu + N_{bb}$	$\mu + N_b$	μ	$\mu + N_g$	$\mu + N_{gg}$
$NEAR5$	$\mu + H_n + N_{bb}$	$\mu + H_n + N_b$	$\mu + H_n$	$\mu + H_n + N_g$	$\mu + H_n + N_{gg}$

In Panel A and Panel B of Table 3.1, we present how the different estimated parameters of Equation 3.3 and Equation 3.4 can be combined to derive the respective average portfolio return in each of the portfolios. We further show how the respective portfolio return can be decomposed into four different return components in Panel C and D. The return components are the benchmark return (μ), the returns associated with the 52-week high (H), the returns attributable to the firm-specific news (N), and the returns associated with the interaction between the firm-specific news and nearness of the stock price to the 52-week high (I). The first return component reflects the benchmark portfolio. It is the average return of the stocks in the portfolio with neither extreme firm-specific news returns nor an extreme nearness to the 52-week high. The second return component is solely driven by the stock's nearness to the 52-week high, regardless of the firm-specific news return ranking. Sorting the stocks into quintiles based on their nearness to the 52-week high results in a return component common among the stocks in the same portfolio. Stocks that are far (f) away from the 52-week high are denoted as H_f and are expected to have a negative return, while stocks that are near (n) the 52-week high are denoted as H_n and are expected to have a positive return. To derive the pure 52-week high effect, we build a long-short strategy that relies solely on the return predictability of the nearness to the 52-week high. We, therefore, define the pure 52-week high effect as:

$$\text{Pure 52-week High Effect} = H_n - H_f = b_5 - b_6. \tag{3.5}$$

The third return component is solely driven by the firm-specific news return, regardless of the firm-specific news return ranking. Sorting the stocks into quintiles based on their firm-specific news return results in a common return component among the stocks in the same portfolio. Following Jiang, Li and Wang (2021), do positive firm-specific news returns predict higher future stock returns, and therefore the firm-specific news component increases from the $FN1$ to the $FN5$ quintile. Stocks with extremely bad (bb) firm-specific news returns are denoted as N_{bb} and bad (b) firm-specific news returns are denoted as N_b, whereas good (g) firm-specific news return are denoted as N_g and extremely good (gg) firm-specific news return are denoted as N_{gg}. While extremely bad firm-specific news returns are associated with negative news momentum and therefore expected to have negative returns in the future, are the extremely good firm-specific news return related to positive future returns. To derive the pure firm-specific news return effect, we build a long-short strategy that relies solely on the return predictability of the firm-specific news return. Depending on the assumption that the 52-week high effect moderates the market underreaction to firm-specific news or not, we define pure firm-specific news as:

$$\text{Pure Firm-specific News Effect} = N_{gg} - N_{bb} = (b_1 + b_{11}) - (b_4 + b_{10}), \text{and} \tag{3.6}$$

$$= b_1 - b_4. \tag{3.7}$$

The fourth and last return component is associated with having, on the one hand, good firm-specific news about the firm and a stock price near the 52-week high and, on the other hand, experiencing bad firm-specific news while having a stock price that is far from the 52-week high. While the underreaction to the firm-specific news due to the nearness to the 52-week high could also be driven by the less extreme quintiles (e.g., the $FN2$ and $FN4$ quintile) but with a smaller magnitude, we focus our analysis on the most extreme FN and $NEAR$ quintiles. Stocks with extremely bad firm-specific news returns far from the 52-week high are denoted as $I_{bb,f}$, whereas stocks with extremely good firm-specific news returns near the 52-week high are represented as $I_{gg,n}$. Hence, the interaction effect is defined as:

$$\text{Interaction Effect} = I_{gg,n} - I_{bb,f} = (b_7 - b_{11}) - (b_{14} - b_{10}) \tag{3.8}$$

If investors don't show any issues with their belief updating process after the arrival of good (bad) firm-specific news while having a stock price that is near (far) its 52-week high, the interaction effect's long-short strategy will not yield any additional significant return component. This would point towards the hypothesis that the portfolio returns are entirely attributable to the pure firm-specific news effect and the pure 52-week high effect. On the other hand, if the return of the interaction effect long-short strategy is positive and significant, this would induce that investors are not willing to update their beliefs and hence are underreacting to the good (bad) news if the stock price is near (far from) its 52-week high.

Finally, the time-series average of the pure firm-specific news effect, the pure 52-week high impact, and the interaction effect are computed. The alphas are calculated by regressing the return components on different asset pricing models to further account for risk factors.[7] To account for serial auto-correlation, we adjust the t-statistics using Newey and West (1987) standard errors with 12 lags.

3.3 Data and descriptive statistics

This section describes the data sources used to create the data sets for our empirical analyses and the sample selection procedure. Afterward, we summarize the characteristics of the underlying data set.

3.3.1 Data

To extract firm-specific news of a firm, we use the RavenPack news database similar to Jiang, Li and Wang (2021). This database structures all relevant information on news articles from thousands of providers, including Dow Jones Newswires, the Wall Street Journal, and MarketWatch, Barron's, into machine-readable measures.[8] We rely on a comprehensive global sample from the most relevant sources from different news providers and their archives for our analysis.[9] To rank a firm-specific news story about a given firm, we use two relevance scores provided by RavenPack, which range between 0 and 100, and the novelty score, which ranges from 0 to 365 days. The first score is entity relevance and captures how strongly the underlying news refers to a specific company. A value of 0 (100) means that the company is only mentioned passively (actively). The second score relates to the event relevance and indicates where the underlying event is mentioned

[7] To benchmark the results of the portfolio sorts, we consider various factor models compromised of the following factors: market ($RMRF$), size (SMB), value (HML), profitability (RMW), investment (CMA), momentum (WML), and liquidity (LIQ). Appendix C provides a detailed description of the factor construction.

[8] Recent studies using this data set comprise Jiang and Sun (2014), Kelley and Tetlock (2017), Ke, Kelly and Xiu (2020), and Jiang, Li and Wang (2021).

[9] We use every news provider, namely Alliance News, Benzinga Pro, Dow Jones Newswires, Dow Jones Third Party, EDGAR SEC Filings, The Fly, FX Street News and FX Street Economic Calendar, Lexis-Nexis News and Social Media, MT Newswires, and Factset Transcripts. In a subsequent step, we filter out certain unreliable sources for each news provider by relying on the source rank. The highest source rank is 1, classified as 'Fully accountable, reputable and balanced,' followed by rank 2, described as 'Official, reliable and honest.' and rank 3, classified as 'Acknowledged, formal, and credible.' To include only the most reliable sources, we filter out every source ranked below 2.

the first time. A value greater or equal to 90 suggest that the event is prominently placed in the title or headline within a news feed. Last, we filter the news on its event novelty score. The measure indicates how new the information contained in the event is compared to previous news. This score specifies how many days have passed since the same event for the given entity was published. For our final sample, we require news stories to have an entity relevance score of 100, an event relevance of 90, and a minimum event novelty score of 1. These filters guarantee that our news sample covers only economically or fundamentally relevant, non-repeated, and, therefore, undisclosed information about a company. We include only firm-specific news, i.e., mergers and acquisitions, analyst ratings, assets, bankruptcy, credit, credit ratings, dividends, earnings, equity actions, labor issues, product services, and revenue from 29 newsgroups. Applying these filters does not introduce any look-ahead bias, as RavenPack assesses all news articles within milliseconds of receipt and immediately sends the resulting data to the users. All information is thus available at the time of news release.

The U.S. and international equities analyses are based on a global sample comprising stock market data from Refinitiv Datastream and accounting data from Worldscope. Several static and dynamic screens are applied to ensure that our sample comprises exclusively of common stocks and provides the highest data quality. First, stocks are identified using Refinitiv Datastream constituent lists, particularly Worldscope lists, research lists, and — to eliminate survivorship bias — dead lists. Following Ince and Porter (2006), Griffin, Kelly and Nardari (2010), and Schmidt et al. (2017), non-common equity stocks are eliminated through generic and country-specific static screens. Furthermore, several dynamic screens are applied to stock returns and prices to exclude erroneous and illiquid observations. Appendix C. 3 and C. 4 provide a detailed description of the static and dynamic screens. Finally, stocks must have a market capitalization greater than zero for the previous month, positive book equity, and a return. We limit our-self to countries that are constituents of the MSCI Developed Markets Index in the respective year.[10] To calculate excess returns, we obtain the risk-free rate from Kenneth R. French's homepage.[11]

To combine the stock market data with the firm-specific news, we follow a multi-step

[10] See https://www.msci.com/market-classification for details.
[11] See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

procedure to match all corresponding news articles of a corresponding firm to a trading day. In the first step, we determine a firm's Datastream identifier, which corresponds to the RavenPack entity identifier, using the provided ISIN and firm name. In the second step, we map the opening hours of the underlying stock exchange to the merged dataset. Lastly, we allocate the identified firm-specific news based on the opening hours of the respective trading day.

Additionally, we include analyst and institutional ownership data for the stock data. All analyst-related data is collected from Institutional Brokers' Estimate System (I/B/E/S), whereas the Institutional ownership data is from the FactSet Ownership database (formerly LionShares). We merge the I/B/E/S data to our stock sample using the provided I/B/E/S ticker and the FactSet data using the provided ISIN of the firm.

In Table 3.2, we summarize our sample selection procedure, allowing us to assess whether investors update their beliefs about a stock after the arrival of firm-specific news between January 2004 and December 2021. After applying the different static and dynamic

Table 3.2
Sample selection

This table presents the sample selection process for U.S. and Ex-U.S. firms. Columns (1) and (2) cover the number of stock-month observations for the firms located in the U.S. and Ex-U.S. Columns (3) and (4) cover the unique stocks, and column (5) the number of unique countries.

	Observation		Firm		Country
	U.S	Ex-U.S.	U.S	Ex-U.S.	Global
	(1)	(2)	(3)	(4)	(5)
... after static and dynamic screens	2157734	6438126	18818	52295	53
... from developed markets	2157734	3290151	18818	27378	24
... with daily stock return	2040066	3270483	18593	27275	24
... with firm-specific news after debut	744563	2597811	6306	18825	23
... with firm-specific news in previous month	549547	965680	6284	18555	23
... with minimum price	539091	891575	6254	18135	23
... with at least 25 stocks per country-month	539091	878159	6254	18083	23
Sample	**539091**	**878159**	**6254**	**18083**	**23**

screens, the original sample covers 8.60 million stock-month observations based on 71.113 unique stocks from 53 countries. Due to data availability and quality, we focus our analysis on developed markets reducing the main sample to 24 unique countries. We limit the sample to stocks with daily returns in the previous month to compute the daily news returns. This reduces the sample to 5.31 million stock-month observations. After mapping the firm-specific news events on the daily returns and applying the monthly news aggregation methodology, we end up with 3.34 million stock-month observations, of which 22.3% are from the United States, and 77.7% are from 23 other countries. To analyze investor behavior after the arrival of firm-specific news, we further require the arrival of firm-specific news during the last month. This leads us to a sample size of 1.52 million observations covering 23 countries. To ensure that small and illiquid stocks do not drive our results, we exclude stocks with a market capitalization below its country's 10% quantile each month, in line with Landis and Skouras (2021). We end up with 1.43 million observations. We require a minimum of 25 stocks for each country-month combination to limit the role of idiosyncratic stock price movements and to ensure a minimum level of stock market coverage within each portfolio.Our final sample includes 1.42 million stock-month observations representing 24.337 unique stocks and 23 countries.

3.3.2 Descriptive statistics

Table 3.3 provides the summary statistics by country, averaged over time. We provide detailed summary statistics of the developed market countries such as Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Hong Kong, Israel, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, United Kingdom, and the United States.

Table 3.3
Summary statistics by country

The table presents summary statistics for each of the 22 countries of our sample. Columns (1), (2), (3), and (4) report the total, minimum, mean, and maximum number of firms per country. Columns (5) and (6) state the average mean and median size per country month. Column (7) shows the average total size per country month and column (8) reports these values in percentage of the respective total across countries. The last two columns (9) and (10) report the actual beginning and ending dates during which each country is included in my sample. Size is measured as market capitalization in million USD. The sample period starts in January 2004 and ends in December 2021.

	Number of firms				Size			Date	
	Total	Min	Mean	Max	Mean	Median	%	Start	End
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Australia	2057	78	233	596	3353	396	2.63	04-01	21-12
Austria	74	25	31	42	2571	1498	0.09	04-06	21-12

Continued on next page

Table 3.2 continued

	Number of firms				Size			Date	
	Total	Min	Mean	Max	Mean	Median	%	Start	End
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Belgium	146	25	41	72	5566	864	0.68	04-01	21-12
Canada	3270	348	752	1100	1414	64	3.90	04-01	21-12
Denmark	211	25	49	107	4667	739	0.58	04-03	21-12
Finland	172	25	58	93	3055	574	0.63	04-01	21-12
France	834	90	184	307	7227	770	4.73	04-01	21-12
Germany	844	84	193	335	6257	681	4.17	04-01	21-12
Hong Kong	680	50	132	287	4283	720	1.95	04-01	21-12
Israel	162	25	41	70	2333	656	0.11	10-07	21-12
Italy	444	29	93	196	3437	940	1.07	04-01	21-12
Japan	4503	336	1325	2673	2899	279	11.55	04-01	21-12
Netherlands	147	25	39	69	7798	2187	0.94	04-01	21-12
New Zealand	134	25	34	48	1228	513	0.04	04-03	21-12
Norway	298	26	70	136	2902	383	0.66	04-03	21-12
Portugal	44	26	27	30	2050	408	0.00	06-04	16-06
Singapore	668	31	102	249	2331	365	0.79	04-01	21-12
Spain	241	25	53	99	7843	2303	1.41	04-01	21-12
Sweden	521	37	114	276	2655	378	1.01	04-01	21-12
Switzerland	258	33	68	123	8201	1712	1.95	04-01	21-12
United Kingdom	2128	351	519	721	3435	253	6.90	04-01	21-12
United States	6254	1904	2495	3027	6221	897	54.09	04-01	21-12
Global	24337	4054	6561	8948	4295	424	100.00	04-01	21-12

On average, we can identify 6.561 stocks per month with a market size of 4.3 billion USD that experience a firm-specific news arrival. The largest market in terms of the number of stocks as well as the market size is the United States, with an average of 2,495 stocks per month and a market size that represent 54.09% of the total market size. The second largest market is Japan, with a maximum of 2,673 stocks per month and coverage of 11.55% of the total market size. While the U.S. market is the largest country in terms of total market size due to its high number of stocks, it is topped with regards to the median size of companies within a country by Austria, Italy, Netherlands, Spain, and Switzerland. We select January 2004 as the start of the sample period due to the broad coverage of firm-specific news events. But several countries like Austria, Denmark, Israel, New Zealand, Norway, and Portugal join the sample at a later stage. In Table 3.4, we depict the descriptive statistics of the main variables for our final sample. Since our interest is in the behavior of investors after the arrival of firm-specific news in combination with the nearness to the 52-week high, we determine the time-series average of the mean, standard deviation, and quantile breakpoints of the cross-section of the two main variables. We additionally include the share of trading days of observations with a minimum of one firm-specific news story in the month the firm-specific news arrives at the firm. $NEAR$

Table 3.4
Variable descriptives

The table reports the time-series average of the cross-sectional mean, standard deviation, and quantiles of each variable for the sample of firm-month observations from January 2004 to December 2021. $NEAR$ is the ratio of the unadjusted stock price at the end of the previous month to the past 52-weeks high, as in George and Hwang (2004). FN is the previous monthly firm-specific news from the firm that is based on decomposed daily returns and the RavenPack news database, as in Jiang, Li and Wang (2021). $FN_{\%}$ is the average share of news days of the firm in the previous month.

	N	Mean	Std	Min	P1	P25	P50	P75	P99	Max
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$NEAR$	6561.34	0.77	0.19	0.18	0.26	0.67	0.82	0.91	1.00	1.00
FN	6561.34	0.81	9.31	-20.47	-15.55	-2.37	0.18	3.14	20.95	29.44
$FN_{\%}$	6561.34	0.11	0.09	0.05	0.05	0.05	0.09	0.14	0.39	0.48

has a mean of 0.77 and a standard deviation of 0.19, indicating that most firms have stock prices close to the 52-week high. The distribution of $NEAR$ is close to symmetric, with a median of 0.82 and a minimum and maximum of 0.18 and 1.00, respectively. FN, the monthly firm-specific news-driven return, has a mean of 0.81 and a standard deviation of 9.31, indicating significant variation in the firm-specific news return. The distribution of FN is positively skewed, with a median of 0.18 and a minimum and maximum of -20.47 and 29.44, respectively. $FN_{\%}$ has a mean of 0.11 and a standard deviation of 0.09. This implies that if a firm is experiencing a firm-specific news arrival in the month, it is, on average, in the news on two days. The distribution of $FN_{\%}$ is also positively skewed, as the median of 0.09 and a range from 0.05 to 0.48 indicate.

3.4 Empirical Results

The main objective of this research is to investigate the impact of firm-specific news in conjunction with the proximity to the 52-week high on investor behavior. To achieve this, we utilize firm-specific news returns as the foundation of our analysis. Our methodology involves sorting the firms independently, first by their news returns and then by their proximity to the 52-week high. To verify the validity of the independent double-sort technique, we demonstrate the lack of correlation between the two sorting variables. We then analyze the cross-sectional return patterns among the double-sorted portfolios and examine how the return predictability of news returns is affected by the firm's proximity to the 52-week high.

3.4.1 Portfolio characteristics

We create double-sorted portfolios by categorizing stocks into country-neutral quintile portfolios based on their firm-specific news returns (FN) and proximity to their 52-week high ($NEAR$) at the previous month-end. This process results in 25 portfolios which are held for one month. In Panels A and B of Table 3.5, we report the average firm-specific news return and the respective average nearness to the 52-week high by each of the 25 portfolios. In both Panels, $NEAR1$ represents the lowest quintile of $NEAR$, while $NEAR5$ represents the highest quintile of $NEAR$. Analogously, $FN1$ represents the lowest quintile of FN, and $FN5$ represents the highest quintile of FN. In Panel C of Table 3.5, the correlation between the two main variables is shown. We identify a smaller variation within the $FN1$ quintile among the five $NEAR$ quintiles in the case of firm-specific news returns. The average news return increases from -9.09% in $NEAR1$ to -5.21% in $NEAR5$. For the other four quintiles ($FN2$, $FN3$, $FN4$, $FN5$), the average firm-specific news return does not vary much within the respective FN portfolio and among the different $NEAR$ portfolios. In the case of the highest FN quintile, the average FN in the $NEAR1$ portfolio is 12.03%, whereas the average FN in the $NEAR5$ portfolio is 11.90%. In Panel B, we identify a very similar pattern. Within the lowest $NEAR$ quintile, the average nearness to the 52-week high varies between 0.58 and 0.57 among the five FN quintiles. Among the higher $NEAR$ portfolios, the average nearness to the 52-week high increases to 0.74 in the case of the $NEAR2$ portfolio, 0.83 ($NEAR3$), and 0.90 ($NEAR4$). In the highest $NEAR$ portfolio, the average value varies between 0.96 for the $FN1$ and 0.97 for the $FN5$ portfolio. The correlation statistic in Panel C further reduces the concern that sorting by $NEAR$ could also be a sort by FN.

In the next step, we analyze the return patterns across the 25 portfolios sorted by the firm-specific news returns and the nearness to the 52-week high. We calculate each portfolio's average risk-adjusted monthly equal-weighted returns and report them in Table 3.6.

Table 3.5
Portfolio characteristics

This table reports the characteristics of the firm portfolios sorted by their firm-specific news return (FN) and their nearness to the 52-week high ($NEAR$). FN is the previous monthly firm-specific news return from the firm that is based on decomposed daily returns and the RavenPack news database, as in Jiang, Li and Wang (2021). $NEAR$ is the ratio of the unadjusted stock price at the end of the previous month to the past 52-weeks high, as in George and Hwang (2004). To form the double-sorting portfolios, in each month, the firms are independently sorted into 5 × 5 country-neutral portfolios based on the FN in the previous month and $NEAR$ at the previous month-end. The portfolios are held for one month. Panels A and B report the average FN and $NEAR$ (sorting variables) for each portfolio, respectively. Mean FN in Panel A is shown in percent. Panel C reports the average correlation between $NEAR$ and FN each month. The sample period is from January 2004 to December 2021.

	(1)	(2)	(3)	(4)	(5)
Panel A: mean FN					
	$FN1$	$FN2$	$FN3$	$FN4$	$FN5$
$NEAR1$	-9.09	-1.65	0.28	2.57	12.03
$NEAR2$	-7.37	-1.64	0.28	2.53	11.09
$NEAR3$	-6.34	-1.61	0.29	2.49	10.57
$NEAR4$	-5.51	-1.57	0.30	2.49	10.32
$NEAR5$	-5.21	-1.46	0.32	2.52	11.90
Panel B: mean $NEAR$					
	$FN1$	$FN2$	$FN3$	$FN4$	$FN5$
$NEAR1$	0.57	0.58	0.58	0.58	0.57
$NEAR2$	0.74	0.74	0.74	0.74	0.74
$NEAR3$	0.83	0.83	0.83	0.83	0.83
$NEAR4$	0.90	0.90	0.90	0.90	0.90
$NEAR5$	0.96	0.97	0.97	0.97	0.97
Panel C: Correlation between FN and $NEAR$					
	FN	PRC			
FN	100.00	12.90			
$NEAR$	12.90	100.00			

Panel A reports the excess return; Panel B reports the $CAPM$ alpha, Panel C focuses on the Fama and French (1993) three-factor alpha ($FF3$), and Panel D uses the Carhart (1997) four-factor model ($FFC4$) as underlying.

Table 3.6
Portfolio returns

This table reports the performance of the firm portfolios sorted by firm-specific news return (FN) and nearness to the 52-week high ($NEAR$). To form the double-sorting portfolios, each month, the firms are independently sorted into 5×5 country-neutral portfolios based on FN and $NEAR$. Additionally, we utilize the country-neutral quintile breakpoints to replicate the equal-weighted portfolio returns of the original (*Orig.*) studies (George and Hwang, 2004; Jiang, Li and Wang, 2021). The equal-weighted portfolios are held for one month and rebalanced every month. This table reports the average monthly excess and risk-adjusted returns of the portfolios. Panel A shows the excess returns. Risk-adjusted returns in Panels A, B, and C are the intercept estimates from the time-series regressions of the monthly excess portfolio returns on market excess return (CAPM), Fama and French (1993) three factors (FF3), and Carhart (1997) four factors (FFC4), respectively. We report the portfolio holding period returns from January 2004 to December 2021. We compute the original firm-specific news returns and nearness to the 52-week high strategy as follows. In each month, we sort the firms into quintiles based on their previous month's firm-specific news returns or their nearness to the 52-week high. We long firms in the highest quintile and short firms in the lowest quintile and hold the portfolios for one month. We track the equal-weighted portfolio returns in the holding period. Alphas in this table are reported in percent. All standard errors are adjusted using Newey and West (1987). t-statistics are in parentheses.

	Orig.	$FN1$	$FN2$	$FN3$	$FN4$	$FN5$	$FN5-1$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Excess Return							
Orig.		0.86	1.03	1.12	1.12	1.50	0.64
		(2.89)	(3.83)	(4.20)	(3.97)	(4.88)	(14.33)
$NEAR1$	1.01	0.71	1.04	1.28	0.97	1.07	0.35
	(2.78)	(1.42)	(2.21)	(2.49)	(2.00)	(2.01)	(3.08)
$NEAR2$	1.00	0.56	0.92	1.07	1.04	1.39	0.84
	(3.16)	(1.30)	(2.28)	(2.57)	(2.40)	(2.88)	(6.21)
$NEAR3$	1.08	0.88	1.02	1.05	1.00	1.46	0.58
	(3.91)	(2.25)	(2.83)	(3.03)	(2.68)	(3.39)	(5.83)
$NEAR4$	1.24	1.03	1.09	1.16	1.28	1.65	0.62
	(4.91)	(2.93)	(3.25)	(3.74)	(3.79)	(4.12)	(5.82)
$NEAR5$	1.31	1.13	1.09	1.07	1.29	1.95	0.82
	(5.59)	(2.83)	(3.48)	(3.71)	(4.22)	(5.53)	(6.19)
$NEAR5-1$	0.29	0.42	0.04	-0.21	0.32	0.89	1.24
	(3.01)	(1.87)	(0.18)	(-0.69)	(1.27)	(3.35)	(5.20)
Panel B: CAPM alpha							
Orig.		-0.04	0.20	0.34	0.27	0.60	0.64
		(-0.41)	(2.30)	(3.34)	(2.92)	(4.96)	(14.15)
$NEAR1$	-0.04	-0.36	-0.00	0.27	-0.08	-0.03	0.33
	(-0.25)	(-1.80)	(-0.01)	(1.13)	(-0.35)	(-0.12)	(3.10)
$NEAR2$	0.05	-0.42	-0.01	0.17	0.10	0.42	0.83
	(0.47)	(-3.02)	(-0.04)	(1.11)	(0.62)	(1.95)	(6.34)
$NEAR3$	0.24	-0.00	0.22	0.27	0.16	0.56	0.56
	(2.70)	(-0.01)	(1.86)	(2.27)	(1.49)	(3.63)	(5.94)
$NEAR4$	0.48	0.24	0.34	0.48	0.52	0.81	0.57
	(6.06)	(2.04)	(3.59)	(4.73)	(5.14)	(5.58)	(5.86)
$NEAR5$	0.63	0.34	0.44	0.49	0.64	1.23	0.89
	(6.99)	(1.98)	(3.30)	(4.34)	(6.28)	(8.51)	(6.86)

Continued on next page

Table 3.5 continued

	$Orig.$	$FN1$	$FN2$	$FN3$	$FN4$	$FN5$	$FN5-1$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$NEAR5-1$	0.66	0.70	0.44	0.22	0.71	1.25	1.59
	(8.72)	(3.55)	(2.28)	(0.95)	(3.47)	(6.10)	(8.30)
Panel C: FF3 alpha							
$Orig.$		-0.03	0.22	0.36	0.29	0.61	0.64
		(-0.37)	(3.77)	(5.12)	(4.75)	(6.64)	(14.08)
$NEAR1$	-0.02	-0.34	0.03	0.30	-0.05	-0.02	0.32
	(-0.16)	(-2.61)	(0.19)	(1.71)	(-0.33)	(-0.10)	(3.02)
$NEAR2$	0.08	-0.39	0.02	0.20	0.14	0.43	0.82
	(1.14)	(-4.35)	(0.24)	(1.82)	(1.40)	(2.75)	(6.47)
$NEAR3$	0.27	0.01	0.25	0.30	0.19	0.59	0.58
	(4.37)	(0.14)	(3.00)	(3.90)	(2.45)	(4.71)	(5.98)
$NEAR4$	0.49	0.25	0.36	0.50	0.54	0.82	0.57
	(8.24)	(2.53)	(4.63)	(6.37)	(6.96)	(6.81)	(5.98)
$NEAR5$	0.63	0.34	0.45	0.50	0.65	1.23	0.89
	(8.05)	(2.19)	(3.85)	(5.15)	(7.24)	(9.58)	(6.78)
$NEAR5-1$	0.65	0.68	0.42	0.20	0.70	1.24	1.57
	(9.82)	(3.80)	(2.41)	(1.04)	(4.15)	(6.67)	(8.96)
Panel D: FFC4 alpha							
$Orig.$		-0.05	0.21	0.36	0.28	0.58	0.63
		(-0.70)	(3.50)	(4.96)	(4.47)	(6.36)	(13.74)
$NEAR1$	0.02	-0.30	0.06	0.34	-0.01	0.02	0.32
	(0.21)	(-2.37)	(0.44)	(1.99)	(-0.04)	(0.10)	(2.93)
$NEAR2$	0.09	-0.39	0.04	0.24	0.16	0.41	0.80
	(1.29)	(-4.36)	(0.44)	(2.15)	(1.53)	(2.63)	(6.54)
$NEAR3$	0.25	-0.01	0.24	0.29	0.18	0.54	0.55
	(4.08)	(-0.07)	(2.87)	(3.77)	(2.37)	(4.35)	(5.66)
$NEAR4$	0.45	0.20	0.32	0.47	0.49	0.77	0.57
	(8.01)	(2.10)	(4.32)	(6.30)	(6.59)	(6.42)	(5.80)
$NEAR5$	0.56	0.25	0.38	0.44	0.57	1.14	0.89
	(7.93)	(1.65)	(3.26)	(4.99)	(6.91)	(9.60)	(6.83)
$NEAR5-1$	0.53	0.55	0.32	0.10	0.57	1.12	1.44
	(8.82)	(3.42)	(1.92)	(0.56)	(4.16)	(6.79)	(9.37)

We will focus our discussion of Table 3.6 on Panel D, respectively the $FFC4$ alpha, because the other risk-adjusted results are comparable. We first investigate the risk-adjusted portfolio returns of the two original settings ($Orig.$) following George and Hwang (2004) and Jiang, Li and Wang (2021). The risk-adjusted portfolio returns increase monotonically from -0.05% (t=-0.70) in the lowest FN to 0.58% (t=6.36) in the highest FN portfolio. A long-short portfolio results in a significant monthly alpha of 0.63% (t=13.74). For the original nearness to the 52-week portfolios, the alpha in the lowest portfolio ($NEAR1$) is equal to 0.02% (t=0.21) and increases to 0.56% (t=7.93) in the $NEAR5$ portfolio. A long-short strategy using the nearness to the 52-week high returns a monthly alpha of 0.53% (t=8.82). Similar to the empirical results of George, Hwang and Li (2014) and Huang, Lin and Xiang (2021), we also discover that the risk-adjusted portfolio return increases within the respective news portfolio when the $NEAR$ portfolio ranking increases. In the case of the lowest FN portfolio with an average alpha of -0.05% (t=-0.70), the $NEAR1$ portfolio earns a monthly alpha of -0.30% (t=-2.37), which increases to 0.25% (t=1.65) for the $NEAR5$ delivering a long-short return of 0.55% (t=3.42). Moving to the portfolios with high firm-specific news ($FN5$) in the previous month. The firms which are far from the 52-week high ($NEAR1$) earn a monthly risk-adjusted alpha of 0.02% (t=0.10), which increases to 1.14% (t=9.60) for firms close to their 52-week high ($NEAR5$). A long-short portfolio strategy that builds on investors' underreaction due to their belief updating bias earns a monthly risk-adjusted alpha of 1.44% (t=9.37). In the case of the long position, investors do not update their beliefs after the arrival of very good news due to the nearness of the stock to the 52-week high. For the short position, investors are unwilling to update their beliefs after the arrival of very negative news as the stock price is already far-away from its 52-week high. This suggests that very positive stock returns are only predicted by very positive firm-specific news. We find a similar pattern when the stock prices are close to their 52-week high. Low stock returns are predicted by very negative firm-specific news returns when the stock prices are far from their 52-week high.

3.4.2 Baseline return decomposition results

Next, we apply the return decomposition methodology described in Section 3.2.2 to disentangle the portfolio returns into the pure effect of the nearness to the 52-week high, the pure firm-specific news effect and their interaction effect and report the results in Table 3.7. In Panel A of Table 3.7, we follow Equation 3.3 and include the interaction effects, whereas, in Panel B, we follow Equation 3.4 and exclude the interaction effects. If the interaction effect is positive and significant, this indicates that a large part of the portfolio formed on the firm-specific news return is driven by having a stock price close or far from the 52-week high. The interaction effect in Panel A of Table 3.7 is positive and significant, independent of which factor model is used to calculate the risk-adjusted return

Table 3.7
Return decomposition results

This table reports the estimates of the monthly averages for the pure firm-specific news return effect, the pure 52-week high effect, and the interaction effect in the firm-specific news setting. The return decomposition methodology is described in 3.2.2, and the specifications of the return decomposition are shown in Table 3.1 and based on the Equation 3.3 and Equation 3.4. The pure firm-specific news effect is computed as $N_{gg} - N_{bb}$, where N_{gg} (N_{bb}) is the return associated with having extremely good (bad) firm-specific news regardless of the nearness to the 52-week high. The pure 52-week high effect is computed as $H_n - H_f$, where H_n (H_f) is the return attributable to having stock prices near (far from) the 52-week high regardless of news about the customer firms. The interaction effect is computed as $I_{gg,n} - I_{bb,f}$, where $I_{gg,n}$ ($I_{bb,f}$) is the return associated with having both very good (very bad) firm-specific news and stock prices near (far from) the 52-week high. Panel A reports return decomposition in which interaction effects are included. Panel B reports return decomposition in which interaction effects are excluded. Average monthly $CAPM$ alpha, $FF3$ alpha, and $FFC4$ alpha are the intercepts from time-series regressions of monthly estimates of each effect (e.g., the pure firm-specific news return effect) on market excess returns, Fama and French (1993) three factors, and Carhart (1997) four factors, respectively. The sample period is from January 2004 to December 2021. Alphas in this table are reported in percent. All standard errors are adjusted using Newey and West (1987). t-statistics are in parentheses.

	Alpha		
	CAPM	FF3	FFC4
	(1)	(2)	(3)
Panel A: Interaction effect included			
Interaction	1.51	1.51	1.47
	(4.90)	(4.83)	(4.67)
Pure Firm-specific News	-0.14	-0.15	-0.13
	(-0.75)	(-0.78)	(-0.66)
Pure 52-week High	0.22	0.20	0.10
	(1.01)	(1.10)	(0.61)
Panel B: Interaction effect excluded			
Pure Firm-specific News	0.70	0.69	0.68
	(10.54)	(10.49)	(10.20)
Pure 52-week High	0.68	0.66	0.55
	(3.42)	(3.93)	(3.88)

component. In column (1), the interaction effect generates an $CAPM$ alpha of 1.51% (t=4.90) per month. By using the Fama and French (1993) three-factor model in column (2), the risk-adjusted return is also equal to 1.34% (t=4.52) per month, and by additionally including the momentum factor by Carhart (1997) in column (3), the monthly alpha is reduced to 1.47% (t=4.67) per month. In the three previously mentioned setups, the return component driven by the pure firm-specific news effect is negative but insignificant. In the case of the $CAPM$, the risk-adjusted alpha is -0.14% (t=-0.75) per month; for the $FF3$, the alpha amounts to -0.15% (t=-0.78) per month, and in the case of $FFC4$ the alpha is equal to -0.13% (t=-0.66).

The pure firm-specific news effect turns positive and significant by excluding the interaction effect in Panel B. In column (1), the effect amounts to 0.70% (t=10.54) per month; in column (2), the effect is very similar by using $FF3$ as a factor model resulting in an alpha of 0.69% (t=10.49) per month. In column (3), the risk-adjusted return of the pure firm-specific news component amounts to 0.68% (t=10.20) per month when regressing the monthly returns on the Carhart (1997) four-factor model.

The comparison of Panels A and B indicates that positive (negative) firm-specific news results predict high (low) future returns for a company only when the stock prices are close to (far from) the 52-week high. These findings suggest that the nearness to the 52-week high causes investors to react inadequately to firm-specific news, significantly contributing to the firm-specific news phenomenon.

3.4.3 Results by information environment

Our results suggest that investors cannot update their beliefs about the fair value of a stock after the arrival of good (bad) firm-specific news if the stock price is near (far from) its 52-week high resulting in the mispricing of the stock. Prior studies suggest that investors' underreaction to new information, like earnings news, as well as the anchoring bias, can be partially attributed to the firms' exposure to limits to arbitrage (Shleifer and Vishny, 1997; Hung, Li and Wang, 2015; Byun, Goh and Kim, 2020). We, therefore, split in Table 3.8 our primary analysis in chapter 3.4.2 into two different sub-samples. Panel A covers all stocks with high exposure to limits to arbitrage, while Panel B contains the stock with low limits to arbitrage. We include five different variables that are closely related to limits to arbitrage and are commonly used in the literature (Lam and Wei, 2011). The first two proxies are the stock market capitalization ($Size$) and analyst coverage ($Coverage$), which are a measure of information uncertainty (Hong, Torous and Valkanov, 2007; Gleason and Lee, 2003; Zhang, 2006). The third proxy is the share of institutional ownership (IO), indicating low short-sale constraints (Nagel, 2005), and the fourth proxy measures through idiosyncratic volatility ($Risk$) potential arbitrage costs (Pontiff, 1996; Wurgler and Zhuravskaya, 2002; Mashruwala, Rajgopal and Shevlin, 2006; Pontiff, 2006; Duan, Hu and McLean, 2010; McLean, 2010; Stambaugh, Yu and Yuan, 2015). The last individual variable is the efficient discrete generalized estimator (TC) as a proxy for potential transaction costs (Ardia, Guidotti and Kroencke, 2022).[12] Similar to Smajlbegovic (2019), we add the limits to arbitrage index (LTA) using a linear combination of the ranks of negative market capitalization, negative institutional ownership, negative analysts coverage, idiosyncratic volatility, and transaction costs.

To be able to investigate how the limits to arbitrage affect the three return components, we sort the stocks each month into three country-neutral portfolios based on the underlying limits to arbitrage proxy. Next, due to the reduced number of stocks in each portfolio, we return the decomposition methodology described in Appendix C using 3×3 country-neutral portfolio sorts. The results in Table 3.8 provide further empirical evidence on the belief updating process of investors. The mispricing effect of investors not being able to update their beliefs about a stock after the arrival of good (bad) firm-specific news and in the case the stock is near (far from) its 52-week high is partially driven by exposure of the stocks to high limits to arbitrage. Each of the six subsamples yields a similar pattern that only the interaction effect in Panel A covering the stocks with high limits to arbitrage is positive and significant. Splitting the sample by size yields an interaction effect for small stocks of 0.96% (t=2.53) per month, whereas large stocks have a monthly alpha of only 0.16% (t=1.02). Using institutional ownership as the underlying splitting criterion to proxy for the information environment, the alpha of low IO stocks is 0.78% (t=2.27) per month, and of high IO stocks, it is equal to 0.07% (t=0.38). In the case of low analyst coverage, the monthly risk-adjusted return of the interaction effect is 0.87% (t=2.40), whereas, for stocks with high coverage, the alpha is 0.14% (t=0.51). Dividing the sample by idiosyncratic volatility yields a monthly $FF4C$ alpha of 0.60% (t=1.67) for high-risk stocks and 0.24% (t=1.09) for low-risk stocks. The sample split by transaction costs results in a monthly risk-adjusted return of 1.06% (t=3.31) for stocks with high transaction costs and a risk-adjusted return of -0.05% (t=-0.28) per month for stocks with low transaction costs. The linear combination of size, institutional ownership, analyst coverage, risk, and transaction cost underlines the previous results by yielding a monthly alpha of 0.67% (t=2.14) for high limits to arbitrage stocks. In contrast, the low limits to arbitrage stocks are associated with a $FF4C$ alpha of 0.00% (t=0.02).

[12] Ardia, Guidotti and Kroencke (2022) show in their paper that the efficient discrete generalized estimator (EDGE) is superior to other proxies for transaction costs estimators from Roll (1988), Corwin and Schultz (2012), as well as Abdi and Ranaldo (2017) or the Amihud (2002) illiquidity measure.

Table 3.8

Return decomposition results: subsamples by information environment.

This table reports the results of the return decomposition in subsamples classified by the firms' information environment. The subsamples are generated as follows. Each month firms are sorted equally into country-neutral tercile portfolios based on the underlying firm characteristic available in the previous month. Panel A covers firms in the high limits to arbitrage tercile, whereas Panel B covers firms in the low limits to arbitrage tercile. In column (1), firms are sorted into groups based on market capitalization. In column (2), firms are sorted into groups based on institutional ownership. Institutional ownership is the holdings by all institutional investors as a fraction of the market capitalization. Firms not covered by FactSet are assumed to have zero institutional ownership. In column (3), firms are sorted into groups based on analyst coverage. Analyst coverage is the number of distinct analysts who make fiscal year one earnings forecasts. Firms not covered by I/B/E/S are assumed to have zero analyst coverage. In column (4), firms are sorted into groups based on idiosyncratic volatility. we define idiosyncratic volatility as the standard deviation of the residuals from a regression of excess returns on a local Fama and French (1993) three-factor model. We use one month of daily data and require at least fifteen non-missing observations. In column (5), firms are sorted into groups based on transaction cost. To estimate transaction cost, we compute for each stock and month, the efficient discrete generalized estimator (EDGE) of the bid-ask spread proposed by Ardia, Guidotti and Kroencke (2022). In column (6), firms are sorted into groups based on the ranked average among the five information environment variables. Within each subsample, we sort the firms into country-neutral 3 × 3 portfolios based on FN and $NEAR$ and conduct a return decomposition using the methodology described in Appendix B. The sample period is from January 2004 to December 2021. FFC4 Alpha in this table is reported in percent. All standard errors are adjusted using Newey and West (1987). t-statistics are in parentheses.

	FF4C Alpha					
	Size	IO	Coverage	Risk	TC	LTA
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: High limits to abitrage						
	Small	Low	Low	High	High	High
Interaction	0.96	0.78	0.87	0.60	1.06	0.67
	(2.53)	(2.27)	(2.40)	(1.67)	(3.31)	(2.14)
Pure Firm-Specific News	0.47	0.18	0.05	0.48	0.29	0.53
	(1.79)	(0.71)	(0.22)	(2.29)	(1.34)	(2.32)
Pure 52-week High	0.22	0.23	0.37	0.52	0.08	0.26
	(1.06)	(1.19)	(1.94)	(2.75)	(0.36)	(1.26)
Panel B: Low limits to arbitrage						
	Large	High	High	Low	Low	Low
Interaction	0.16	0.07	0.14	0.24	-0.05	0.00
	(1.02)	(0.38)	(0.51)	(1.09)	(-0.28)	(0.02)
Pure Firm-Specific News	0.05	0.21	-0.08	-0.00	0.35	0.08
	(0.50)	(1.90)	(-0.42)	(-0.03)	(2.24)	(0.85)
Pure 52-week High	0.29	0.15	0.08	-0.07	0.36	0.11
	(2.63)	(1.11)	(0.65)	(-0.54)	(3.86)	(1.07)

3.4.4 Robustness checks

Next, we perform various robustness checks and additional tests to support the findings of the return decomposition. First, we provide further evidence of investor behavior in an out-of-sample application. Second, we use a variety of factor models. Third, we use different definitions of firm-specific news. The last robustness test focuses on a placebo test using the return decomposition based on MOM.

Table 3.9 splits the sample into two sub-samples. The first sub-sample focuses on the return decomposition in the U.S. and is similar to the analysis in George, Hwang and Li (2014), Huang, Lin and Xiang (2021), and Jiang, Li and Wang (2021). The second sub-sample focuses on a true out-of-sample analysis by excluding the U.S. Even after focusing

Table 3.9
Return decomposition results: U.S vs Ex-U.S.

This table reports the results of the return decomposition, including the interaction effects based on different country subsamples. Panel A, reports the results of firms located within the U.S., and Panel B of firms located outside of the U.S. Within each subsample, we sort the firms into country-neutral 5 × 5 portfolios based on FN and PRC and conduct a return decomposition using the methodology described in Appendix B. Average monthly FFC4 Alpha are the intercepts from time-series regressions of monthly estimates of each effect (e.g., the pure firm-specific news return effect) on the Carhart (1997) four factors. The sample period is from January 2004 to December 2021. FFC4 Alpha in this table is reported in percent. All standard errors are adjusted using Newey and West (1987). t-statistics are in parentheses.

	FF4C Alpha			
	U.S.		Ex-U.S.	
	(1)	(2)	(3)	(4)
Interaction	0.61		1.87	
	(2.53)		(4.16)	
Pure Firm-Specific News	0.11	0.54	-0.23	0.78
	(0.64)	(5.54)	(-0.83)	(11.19)
Pure 52-week High	0.13	0.31	0.14	0.71
	(0.70)	(1.66)	(0.63)	(4.82)

the return decomposition on the most efficient and mature market, the interaction effect in column (1) remains positive and statistically significant, yielding a risk-adjusted return of 0.61% (t=2.53) per month. Further, excluding the interaction term from the Fama and MacBeth (1973) in column (2) yields positive and statistically significant pure effects. The results of the out-of-sample test in column (3) and column (4) underline the robustness of our results. By comparing the four-factor alpha of the interaction term in column (3) to column (1), we can identify an increase of 1.26 percentage points to 1.87% (t=4.16) per month. The exclusion of the interaction effect in column (4) still yields positive and significant pure effects.

In Table 3.10, we use a variety of factor models to determine whether the employed factor drives the risk-adjusted returns of the interaction term. More specifically, we extend the Fama and French (1993) three-factor model with the profitability and investment factor as proposed by Fama and French (2015), resulting in the proposed five-factor model ($FF5$). Further, similar to the model by Carhart (1997), we add momentum to the five-factor model ($FF5C$). To control for liquidity constraints, we employ the four-factor model proposed by Pástor and Stambaugh (2003) by adding the liquidity factor to the Fama and French (1993) three-factor model (PS). Next, we combine the five-factor model by Fama and French (2015) with the liquidity factor resulting in a six-factor model ($FF5 + LIQ$). The last model augments the $FF5$ with the momentum and liquidity factor resulting in a seven-factor model ($FF5C + LIQ$). In column (1) of Table 3.10, the risk-adjusted return

Table 3.10
Return decomposition results: other risk-adjustment methods.

This table reports the return decomposition results using various risk-adjustment methods. In the FF5 column, risk-adjusted returns are estimated from time-series regressions of monthly return components (effects) on Fama and French (2015) five factors. In the FF5C column, we augment FF5 factors with the momentum factor (UMD) by Carhart (1997) in the time-series regression. In the PS column, risk-adjusted returns are estimated from time-series regressions of monthly return components (effects) on Pástor and Stambaugh (2003) four factors. In the FF5+LIQ column, we augment FF5 factors with the liquidity (LIQ) factor by Pástor and Stambaugh (2003) in the time-series regression. In the FF5C+LIQ column, we augment FF5 factors with the liquidity (LIQ) factor by Pástor and Stambaugh (2003) and the momentum factor (UMD) by Carhart (1997) in the time-series regression. The sample period is from January 2004 to December 2021. Alphas in this table are reported in percent. All standard errors are adjusted using Newey and West (1987). t-statistics are in parentheses.

	Alpha				
	FF5	FF5C	PS	FF5+LIQ	FF5C+LIQ
	(1)	(2)	(3)	(4)	(5)
Interaction	1.61	1.55	1.53	1.64	1.58
	(5.30)	(5.11)	(4.79)	(5.28)	(5.11)
Pure Firm-Specific News	-0.17	-0.13	-0.17	-0.21	-0.17
	(-0.83)	(-0.66)	(-0.85)	(-0.99)	(-0.82)
Pure 52-week High	0.02	-0.13	0.23	0.05	-0.10
	(0.09)	(-0.72)	(1.21)	(0.24)	(-0.53)

of the interaction effect slightly increases to 1.61% (t=5.30) per month compared to the three-factor model in Table 3.7 by adding profitability and investment factor to the factor model. A similar effect can be identified when adding the two additional factors to the Carhart (1997) four-factor model resulting in a monthly alpha of 1.55% (t=5.11). Adding the liquidity factor to the three-factor and five-factor model in column (3) and column (4) yields a monthly risk-adjusted return of 1.53% (t=4.79) and 1.64% (t=5.28), respectively. After controlling for the largest factor model in column (5), the interaction effect remains positive, economically, and statistically significant.

For our third robustness check in Table 3.11, we add several variations of the previously defined measure of firm-specific news for three samples. The first sample in Panel A uses the entire sample, Panel B focuses on the firms located in the U.S., and Panel C limits the sample to all firms outside of the U.S. The first two alternative measures are related to earnings announcement days (EAD), as these scheduled news events are well-known and followed by investors. To identify these earnings announcement days, we rely on the methodology of Engelberg, McLean and Pontiff (2018) by identifying the earnings announcement day as the day with the highest volume within a three-day window around the reported announcement day. To be as precise as possible about the impact of EADs, we define two measures that deviate slightly from our firm-specific news measure. The first measure only includes days with an earnings announcement as a firm-specific news day, while all other days are classified as non-firm-specific news days. The second measure investigates the incremental value of the underlying firm-specific news provider, as we exclude all EADs from the firm-specific news days. The third measure tries to model a slower information diffusion of firm-specific news. If a firm-specific news event occurred on the day t, we additionally classify the next day $t + 1$ as a firm-specific news day. For the next measure, we classify additional events as firm-specific. [13] We exclude the days from the firm-specific news measure on which relevant macroeconomic information is released to exclude the possibility that our results are driven by macroeconomic news. To identify all relevant macro-economic news, we follow Savor and Wilson (2013) by using only the macro announcements that have statistically and economically significant impacts on an individual country's market risk premium.[14] For the last measure, we follow

[13] These events are 'partnerships,' 'indexes,' 'marketing,' 'regulatory,' 'permits,' 'exploration,' 'commodity-prices,' 'industrial-accidents,' 'business-operations,' 'credit-default-swap,' 'privacy,' and 'ownership.'

[14] Due to the availability we limit our analysis to the following countries: Australia, Canada, France, Germany, Italy, Japan, Netherlands, Norway, New Zealand, Spain, Sweden, Switzerland, United Kingdom, and the United States.

Burt and Hrdlicka (2021) to extract the idiosyncratic news part from daily returns. To decompose the returns into a predictable and unpredictable (idiosyncratic) component, we use an asset pricing model derived from the daily returns of the last 12 months ($t-1$ till $t-12$) and the factor realizations at time t. The estimated parameters enable us to derive the predictable component (ϵ_t), equal to the daily return minus the non-idiosyncratic component. In the final step, we differentiate between idiosyncratic firm-specific and non-firm-specific news returns. For the last measure, we follow Burt and Hrdlicka (2021) to extract the idiosyncratic news part from daily returns. Limiting the firm-specific news to the earnings announcement days in column (1) reduces the risk-adjusted return of the interaction effect, loses its significance, and yields an alpha of 1.21% (t=1.60) per month. These results are robust by limiting the sample to only firms located in the U.S. or outside of the U.S. This indicates that investors are paying a lot of attention to these earnings announcement days. Therefore, the news component diffuses very fast into the stock price. A less likely alternative explanation could be that the anchoring effect is not persistent these days, contrary to the results of George, Hwang and Li (2014). Excluding the earnings announcement days from the firm-specific news return estimation in column (2) increases the four-factor alpha to 1.68% (t=5.17) per month. In column (3), we model a slower information diffusion, resulting in a lower global monthly risk-adjusted return of 1.33% (t=4.90). Limiting the sample to the most efficient stock market, the U.S., the interaction effect even becomes insignificant, yielding an alpha of 0.52% (t=1.19) per month. Including more events in the firm-specific news detection in column (4) further decreases the monthly four-factor alpha to 1.22% (t=3.58). We can identify a similar pattern as in column (3), in which the stocks from outside the U.S. yield a positive and significant alpha. In contrast, the alpha of the U.S. sample is insignificant. This underlines the importance of the event selection by Jiang, Li and Wang (2021). Column (5) excludes all the days important macroeconomic announcements are released. The risk-adjusted return of the interaction effect is unaffected by this correction, yielding a global monthly alpha of 1.38% (t=3.86). Similar to column (5), we try to measure the firm-specific news component more exactly by excluding the predictable part from the daily return in column (6). By aggregating the daily idiosyncratic and firm-specific returns, the global risk-adjusted return of interaction effect amounts to 1.32% (t=4.30).

Table 3.11
Return decomposition results: robustness measures.

This table reports the return decomposition results using different firm-specific news measures. Panel A covers all firms with firm-specific news, Panel B covers firms from the U.S., and Panel C excludes all U.S. firms. In the EAD column, only earnings announcement dates are used to identify firm-specific news. We follow Engelberg, McLean and Pontiff (2018) by identifying the earnings announcement day as the day with the highest volume within a the three-day window around the reported announcement day in I/B/E/S. In the $-EAD$ column, we exclude all earnings announcement dates from the firm-specific news measure. In the $t, t+1$ column, we model a slower information diffusion by tagging the next day after news occurrence as a firm-specific news day. In the $+Events$ column, we add further essential events to the firm-specific news measure. Additional events include: partnerships, indexes, marketing, regulatory, permits, exploration, commodity-prices, industrial-accidents, business-operations, credit-default-swap, privacy, ownership. In the $-Macro$ column, we exclude macro-news days. We follow Savor and Wilson (2013) by excluding macro announcement days that have statistically and economically significant impacts on an individual country's market risk premium. In the ϵ column, we extract the idiosyncratic news part from daily returns. We follow Burt and Hrdlicka (2021) by decomposing the daily return into a predictable and idiosyncratic component using the data and asset pricing model from the previous twelve month (t-2 till t-13). Average monthly FFC4 Alpha are the intercepts from time-series regressions of monthly estimates of each effect (e.g., the pure firm-specific news return effect) on the Carhart (1997) four factors. The sample period is from January 2004 to December 2021. FFC4 alpha in this table is reported in percent. All standard errors are adjusted using Newey and West (1987). t-statistics are in parentheses.

	FF4C Alpha					
	EAD	$-EAD$	$t, t+1$	$+Events$	$-Macro$	ϵ
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Global						
Interaction	1.21	1.68	1.33	1.22	1.38	1.32
	(1.60)	(5.17)	(4.90)	(3.58)	(3.86)	(4.30)
Pure Firm-Specific News	0.14	-0.37	-0.23	0.02	-0.08	-0.15
	(0.25)	(-2.03)	(-1.51)	(0.10)	(-0.45)	(-0.77)
Pure 52-week High	0.03	0.05	0.21	0.17	0.25	0.11
	(0.16)	(0.24)	(1.31)	(0.95)	(1.47)	(0.65)
Panel B: U.S.						
Interaction	-0.29	0.72	0.52	0.11	0.73	0.87
	(-0.30)	(2.21)	(1.19)	(0.30)	(2.65)	(2.89)
Pure Firm-Specific News	0.84	-0.10	0.01	0.37	0.07	-0.12
	(1.26)	(-0.43)	(0.05)	(1.27)	(0.33)	(-0.55)
Pure 52-week High	0.46	0.14	0.22	0.27	0.13	0.02
	(1.51)	(0.73)	(1.04)	(1.33)	(0.58)	(0.10)
Panel C: Ex-U.S.						
Interaction	1.72	2.21	1.74	1.83	1.81	1.54
	(1.86)	(4.92)	(6.87)	(3.98)	(3.12)	(3.96)
Pure Firm-Specific News	-0.07	-0.48	-0.36	-0.17	-0.11	-0.19
	(-0.10)	(-1.96)	(-2.20)	(-0.61)	(-0.32)	(-0.75)
Pure 52-week High	-0.04	0.03	0.26	0.17	0.34	0.20
	(-0.15)	(0.14)	(1.51)	(0.72)	(1.42)	(0.92)

Similar to Huang, Lin and Xiang (2021), our findings may be driven by the momentum effect instead of nearness to the 52-week high since MOM, and $NEAR$ are potentially positively correlated. To rule out this possibility, we address this concern by performing a placebo return decomposition based on MOM instead of $NEAR$ in Table 3.12. Independent of the underlying factor model is the risk-adjusted return of interaction effect between the firm-specific news and the momentum effect negative and not significant. In column (5), we use the four-factor model as underlying to estimate the alpha of the interaction effect, which amounts to -0.19% (t=-0.62) per month. In contrast, the pure firm-specific news and momentum effects stay significant, yielding a monthly risk-adjusted return of 0.82% (t=4.70) and 0.46% (t=4.62), respectively. Excluding the interaction effect from the regression in column (6) results in a monthly alpha of 0.77% (t=11.10) and 0.39% (t=4.79), when including only the pure firm-specific news effect and the pure momentum effect. This placebo test highlights the uniqueness of the nearness to the 52-week high in explaining the underreaction to the arrival of firm-specific news.

Table 3.12
Return decomposition results: placebo test.

This table reports the estimates of the monthly averages for the pure firm-specific news return effect, the pure momentum effect, and the interaction effect. The return decomposition methodology is described in 3.2.2, and the specifications of the return decomposition are shown in Table 3.1 and based on the Equation 3.3 and Equation 3.4, where we replace $NEAR$ with MOM. The pure firm-specific news effect is computed as $N_{gg} - N_{bb}$, where N_{gg} (N_{bb}) is the return associated with having extremely good (bad) firm-specific news regardless of the stock's momentum. The pure momentum effect is computed as $H_n - H_f$, where H_n (H_f) is the return attributable to having high (low) stock momentum regardless of firm-specific news about the firms. The interaction effect is computed as $I_{gg,n} - I_{bb,f}$, where $I_{gg,n}$ ($I_{bb,f}$) is the return associated with both very good (very bad) firm-specific news and high (low) momentum. Average monthly $CAPM$ alpha, $FF3$ alpha, and $FFC4$ alpha are the intercepts from time-series regressions of monthly estimates of each effect (e.g., the pure firm-specific news return effect) on market excess returns, Fama and French (1993) three factors, and Carhart (1997) four factors, respectively. The sample period is from January 2004 to December 2021. Alphas in this table are reported in percent. All standard errors are adjusted using Newey and West (1987). t-statistics are in parentheses.

	Alpha					
	CAPM		FF3		FFC4	
	(1)	(2)	(3)	(4)	(5)	(6)
Interaction	-0.12		-0.15		-0.19	
	(-0.42)		(-0.49)		(-0.62)	
Pure Firm-Specific News	0.83	0.81	0.83	0.80	0.82	0.77
	(4.72)	(11.51)	(4.79)	(11.04)	(4.70)	(11.10)
Pure Momentum	0.73	0.68	0.66	0.60	0.46	0.39
	(4.14)	(4.38)	(4.45)	(3.94)	(4.62)	(4.79)

3.4.5 Analysis of the economic mechanism

In this section, we further investigate the economic mechanism behind the distortion of the belief updating process by combining analyst recommendations revisions and the arrival of firm-specific news. Similar to Huang, Lin and Xiang (2021), we examine analysts' recommendation changes as they provide a direct proxy to observe the belief-updating process of essential information intermediaries in financial markets (Campbell and Sharpe, 2009; Cen, Hilary and Wei, 2013). We perform two types of regressions to examine the impact of the nearness to the 52-week high on analyst reactions to the arrival of firm-specific news. The first set of regressions uses an ordered logit, whereas the other set uses an ordinary least squares regression. Each of the regressions uses a binary indicator if the analyst changed his recommendation after the arrival of firm-specific news, the associated firm-specific news return, the nearness to the 52-week high, an interaction of both, as well as several controls resulting in the following equation:

$$RecChange_{i,j} = \beta_1 FN_{i,j} + \beta_2 NEAR_{i,j} + \beta_3 FN_{i,j} \times NEAR_{i,j} + \beta_{1c}\mathbf{C} + \epsilon_{i,j}, \qquad (3.9)$$

where $RecChange_{i,j}$ is recommendation revision event j of firm i. Based on the different Panels in Table 3.13, the revision event can take different values. In Panel A, the recommendation revision event is defined as $RecChange$. It takes a value of one if the analyst revised his stock recommendation upwards, zero if it is unchanged, and minus one in the case of a downgrade. In Panel B, we regress the independent variables on the dummy $Upgrade$, which equals one for a positive revision and otherwise zero. In Panel C, the dummy $Downgrade$ is defined as one in the case of a negative revision and otherwise zero. The regression includes three fundamental variables to understand further analysts' distorted belief updating process. The first variable, FN, is the cumulative firm-specific news return in the 21 trading days before the day of the recommendation change event. The second variable, $NEAR$, is the nearness to the 52-week high at the end of the trading day before the recommendation change, and the last variable, $FN \times NEAR$, is the interaction term between FN and $NEAR$. We include similar control variables as in Huang, Lin and Xiang (2021), determined at the previous month-end before the analyst revision events. The controls cover analyst-based variables like the number of earnings forecast revisions, analyst dispersion, analyst coverage, and standardized unexpected earnings, and further firm-specific controls like firm size, book-to-market ratio, asset growth, and accruals, as well as return-driven controls such as momentum, short-term reversal, and idiosyncratic volatility. We further include industry, year, and country fixed effects in the regression and cluster the standard errors by each firm.

Table 3.13
Analyst recommendation revision.

This table reports the predictive effects of firm-specific news returns, nearness to the 52-week high, and their interaction on the direction of subsequent analyst recommendation revisions. The analysis is conducted using analyst recommendation revisions on firms with firm-specific news from January 2004 to December 2021. In columns (1–3) of Panel A, we estimate an ordered logit regression model as in Eq. (1), where the dependent the variable takes a value of one when the analyst recommendation revision on a firm is an upgrade, zero when the revision is a reiteration and a negative one when the revision is a downgrade. The independent variable FN is the cumulative firm-specific news returns in the 21 trading days before the recommendation revision days. $NEAR$ is the nearness to the 52-week high of the firm on the trading day before the announcement days. $FN \times NEAR$ is the interaction term between FN and $NEAR$. The control variables are supplier firm characteristics, including analyst dispersion, analyst coverage, standardized unexpected earnings (SUE), market capitalization, book-to-market ratio, past 12-month cumulative returns, idiosyncratic volatility, asset growth, and accruals as of the month-end before the recommendation announcement date. Fama-French 48-industry, year, month, and country fixed effects are included in the regressions. In columns (4–6), we re-perform the above regressions in OLS regressions. Z-statistics in parentheses of columns (1–3) or t-statistics in parentheses of columns (4–6) are computed based on standard errors clustered by firm. In Panel B, the dependent variable is replaced by Upgrade, which is a dummy variable that equals one if the revision is an upgrade and zeroes otherwise. In Panel C, the dependent variable is replaced by Downgrade, which is a dummy variable that equals one if the revision is a downgrade and zeroes otherwise. In Panels B and C, we estimate logit regression models in columns (1–3) and OLS regression models in columns (4–6).

	Ordered Logit			OLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: ***RecChange*** **as the dependent variable**						
FN	0.010	0.030	0.036	0.005	0.013	0.016
	(21.66)	(22.16)	(23.77)	(21.46)	(21.93)	(23.58)
$NEAR$		-0.079	-0.229		-0.064	-0.091
		(-5.74)	(-11.14)		(-8.98)	(-8.47)
$FN \times NEAR$		-0.029	-0.029		-0.013	-0.013
		(-15.13)	(-13.70)		(-14.47)	(-13.25)
Controls	No	No	Yes	No	No	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs	684678	684678	545959	684678	684678	545959
Pseudo/Adj. R^2	0.001	0.001	0.004	0.004	0.004	0.011

Continued on next page

Table 3.13 continued

	Ordered Logit			OLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B: *Upgrade* **as the dependent variable**						
FN	0.009	0.018	0.024	0.002	0.004	0.005
	(19.16)	(12.95)	(15.46)	(18.75)	(12.87)	(15.47)
$NEAR$		-0.142	-0.319		-0.050	-0.068
		(-9.40)	(-13.63)		(-12.80)	(-11.55)
$FN \times NEAR$		-0.011	-0.012		-0.002	-0.003
		(-5.86)	(-5.44)		(-5.48)	(-5.33)
Controls	No	No	Yes	No	No	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs	684678	684678	545959	684678	684678	545959
Pseudo/Adj. R^2	0.001	0.001	0.006	0.007	0.008	0.013
Panel C: *Downgrade* **as the dependent variable**						
FN	-0.011	-0.039	-0.044	-0.003	-0.009	-0.010
	(-21.59)	(-27.28)	(-27.88)	(-21.68)	(-27.42)	(-28.05)
$NEAR$		0.027	0.146		0.014	0.023
		(1.69)	(6.33)		(3.43)	(3.91)
$FN \times NEAR$		0.043	0.042		0.010	0.010
		(21.22)	(19.05)		(20.95)	(18.98)
Controls	No	No	Yes	No	No	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs	684678	684678	545959	684678	684678	545959
Pseudo/Adj. R^2	0.001	0.002	0.007	0.004	0.005	0.011

We will focus our discussion of Table 3.13 on columns (1) to (3) of each Panel. Starting with Panel A, the results in column (1) suggest that analysts are more inclined to change their recommendations on a stock in the direction of the firm-specific news event, as evidenced by the positive and significant FN coefficient. This indicates that analysts pay attention to the news and incorporate them into their recommendations. The negative and significant coefficients of the two interaction terms in column (2) and column (3) further indicate that analysts are less likely to upgrade (downgrade) the stock recommendation in response to positive (negative) firm-specific news when the stock price is near (far from) the 52-week high. The results of Panel B and Panel C of Table 3.13 underline our results by replacing the recommendation change with the two dummy variables *Upgrade* and *Downgarde*. In the case of Panel B, the coefficient of the firm-specific news return is still positive and significant, and the interaction term is negative and significant. These coefficients provide evidence that analysts upgrade their recommendations as soon as positive news arrives at the firm, while they are less likely to do this if the stock price is near its 52-week high. For Panel C, the results are the same as the coefficients are flipped. The negative and significant coefficient of the firm-specific news indicates that if negative firm-specific news arrives, analysts tend to downgrade their recommendations, but due to the positive and significant coefficient on the interaction term are less likely to do this if the stock price is far away from its 52-week high. The results in columns (4) to column (6) of Panel A, B, and C are robust to replacing the ordered logit regression with an ordinary least squares regression. These results provide evidence for our main hypothesis that the distortion of the belief updating process causes an underreaction influenced by the stock prices' nearness to the 52-week high and the arrival of firm-specific news.

3.5 Conclusion

The paper examines investor underreaction to firm-specific news in global equity markets and tests the anchoring effect as an economic mechanism. The anchoring effect refers to the tendency of investors to cling to their initial beliefs even when facing new information, as reinforced by their use of the 52-week high as an anchor. The paper investigates the central hypothesis that the anchoring effect distorts the investor's belief updating process after the arrival of firm-specific news, resulting in the predictability of future stock returns. The sample for the empirical analysis covers stocks from developed markets across 23 countries from January 2004 to December 2021.

A return decomposition methodology allows us to disentangle the stock return predictability into three components. The first component measures the pure firm-specific news return, the second the pure effect resulting from the stock price nearness to its 52-week high, and the third component the interaction effect between the firm-specific news return and the nearness to its 52-week high. By including all three effects, the interaction effect is positive and significant. In contrast, the pure firm-specific news return turns insignificant compared to the configuration in which only the two pure effects are included in the regression. Our results show that the investors' underreaction to the firm-specific news is at least partially explained by the anchoring bias induced by the nearness to the 52-week high. We also explore how the nearness to the 52-week high distorts the belief-updating process by utilizing analyst recommendation changes, leading to an underreaction. Analysts react to firm-specific news but are less likely to change their recommendation if the stock price is near the 52-week high. Finally, we show that unscheduled, firm-specific news drives the anchoring bias effect on investors' underreaction over the subsequent month. This contrasts previous studies, claiming that investors underreact to scheduled news, such as earnings announcements, over the subsequent days. Our study indicates that investors tend to underreact to unscheduled firm-specific news due to the psychological barrier created by the 52-week high in global equity markets. This offers a fresh perspective on investor underreaction, often attributed solely to investor inattention in the existing literature.

The insights in this paper give rise to future research in at least three dimensions. First, while this study investigates the general underreaction to firm-specific news, one potential avenue for further research could be to explore which news categories are causing the underreaction or by investigating macroeconomic news. Second, while this study focuses on stocks, it could be extended by investigating investors' underreaction within corporate bonds. Third, instead of using the nearness in terms of price, the news decay over time could play another critical role in explaining the underreaction.

4 Conclusion

Empirical and theoretical evidence exists that asset prices are not solely driven by their systematic risk compared to the market but rather by a large number of factors. The release of new information and the diffusion among markets and firms helps explain the variation in asset prices. Further, previous empirical asset pricing research focuses prominently on linear models, but research on machine learning methods that help model non-linearities is comparatively sparse.

The three essays in this dissertation aim to fill this gap and explore research questions on empirical asset pricing in global stock markets. In the first essay, I study the slow diffusion of firm-specific news, market news, and noise from fundamentally related firms into the focal firms' stock prices. In the second essay, I compare various machine learning models to predict the cross-section of emerging market stock returns. In the third essay, I test the anchoring effect on investor underreaction to global firm-specific news. In this chapter, I briefly summarize the main results of each essay and highlight the contributions and implications.

In the *first* essay, I study how firm-specific news, market news, and the noise of fundamentally linked firms diffuse into the stock prices of firms in a global network. Firm-specific news, market news, and noise are estimated at the stock-level using a structural vector auto-regression on daily data. To model fundamental links between firms, I rely on the characteristic of analysts to cover stocks that are connected.

For a global sample that consists of 42,789 stocks for 49 equity markets and spans 30 years, I present robust evidence that firm-specific news is the driver of slow information diffusion. The cross-sectional return difference between firms exposed to negative news from linked firms and those exposed to positive news amounts to approximately 7% per year. Further, it takes up to three months for an investor to incorporate the diffusing firm-specific news. Another finding of this paper is that investors can differentiate between noise and news and tend to process first market-wide news and afterward firm-specific news. This underreaction to the arrival of news is caused by the investors' limited attention to firm-specific news.

Overall, the results imply that investors are not fully limited in attention. They show a strong categorial learning behavior that enables them to incorporate at least partially relevant information diffusing among fundamentally linked firms.

The *first* essay contributes to the literature on momentum spillovers originating from fundamentally related firms by providing an understanding of the role of news diffusion among fundamentally linked firms for asset prices. I am the first to study the diffusion of news among fundamentally linked firms in a large sample covering both developed and emerging markets. Further, I contribute to the literature by differentiating between firm-specific news, market news and noise using the return decomposition model of Brogaard et al. (2022). These individual components allow me to uncover the return component causing the cross-firm predictability. To my knowledge, I am the first to combine the categorical learning behavior by Peng and Xiong (2006) with the diffusion of market news and firm-specific news among firms and further relate this slow information diffusion to investors' underreaction due to limited attention.

In the *second* essay, I compare various machine learning models to predict the cross-section of emerging market stock returns from 32 emerging market countries between January 2002 to December 2021. More specifically, I analyze the predictive power of nine algorithms: ordinary least squares regression and elastic net as examples for traditional linear models; tree-based models such gradient boosted regression trees and random forest; and neural networks with one to five layers. Furthermore, we investigate the performance of an ensemble comprising the five different neural networks and an ensemble of methods that allow for non-linearities and interactions, i.e., the two tree-based models and the ensemble of neural networks.

I document that return forecasts from machine learning methods lead to superior out-of-sample returns in emerging markets. Interestingly, investors already applying such a strategy in developed markets seem to enjoy potential diversification benefits when applying them also in emerging markets. I further investigate the source of the predictability and conclude that it rather stems from mispricing than higher risk. Still, the superiority of machine learning models in emerging markets does not stem from limits to arbitrage. Finally, significant net returns can be achieved when accounting for transaction costs, short-selling constraints, and limiting our investment universe to big stocks only.

Altogether, the findings of the second essay suggest that predicting the cross-section of emerging market stock returns and allowing for non-linearities and interactions leads to economically and statistically superior out-of-sample returns compared to traditional linear models.

The *second* essay contributes to the literature in at least three aspects. It contributes to the rapidly expanding literature on predicting the cross-section of stock returns with machine learning methods. There is only evidence that more complex machine learning models are superior to linear models in developed markets (Rasekhschaffe and Jones, 2019; Freyberger, Neuhierl and Weber, 2020; Gu, Kelly and Xiu, 2020; Tobek and Hronec, 2020; Drobetz and Otto, 2021). Under the hypothesis that developed markets are integrated, the same risk factors should apply to these markets. Therefore, similar results within developed markets are unsurprising, and emerging markets provide an attractive alternative for out-of-sample tests in independent and new samples. Further, it adds to the literature on the drivers of emerging market stock returns and market integration (Bekaert and Harvey, 1995; Harvey, 1995). The machine learning models allow taking non-linearities and interactions into account next to linear relationships. Lastly, it contributes to the understanding of the source of return predictability from machine learning forecasts (Avramov, Cheng and Metzker, 2022; Leung et al., 2021; Cakici et al., 2022a). It provides evidence that machine learning models show higher predictability for stocks associated with higher limits to arbitrage. A positive and significant outperformance can be achieved even when accounting for transaction costs, short-selling constraints, and limiting the investment universe to big stocks only.

In the *third* essay, I investigate the impact of firm-specific news in conjunction with the proximity to the 52-week high on investor behavior. I utilize a return decomposition methodology to disentangle the stock return predictability into three components. The first component measures the pure firm-specific news return, the second the pure effect resulting from the stock price nearness to its 52-week high, and the third component the interaction effect between the firm-specific news return and the nearness to its 52-week high.

My sample includes 1.42 million stock-month observations representing 24.337 unique stocks and 23 countries. This allows me to conclude that the investors' underreaction to the firm-specific news is partially explained by the anchoring bias induced by the nearness to the 52-week high. The interaction effect yields an average Fama-French-Carhart (1997) four-factor alpha of 1.47% (t=4.67), whereas the two pure effects are insignificant. Further, my results provide evidence that firms drive the induced underreaction of investors with high limits to arbitrage and unscheduled firm-specific news. Lastly, I provide evidence that analysts react to firm-specific news but are less likely to change their recommendation if the stock price is near the 52-week high.

The third essay indicates that investors tend to underreact to firm-specific news due to the psychological barrier created by the 52-week high. This offers a fresh perspective on the market underreaction, often attributed to investor inattention in the existing literature.

This study adds to understanding investor underreaction in at least four aspects in an international asset pricing context. First, we contribute to a better understanding of investor underreaction by explicitly using firm-specific news (Jiang, Li and Wang, 2021) instead of proxying news with economically-linked, past-month firm momentum (Huang, Lin and Xiang, 2021). I provide insights into investor underreaction by showing that limits to arbitrage amplify the underreaction potential. Second, our paper reveals a crucial economic mechanism behind investor underreaction in global equity markets. I rely on the anchoring and adjustment hypothesis by showing that professional forecasters (Campbell and Sharpe, 2009; Cen, Hilary and Wei, 2013) include the firm-specific news in their recommendation but are affected by the anchoring bias if the stock is near (far from) the 52-week high and positive (negative) news arrives. Third, I show that unscheduled, firm-specific news drives the anchoring bias effect on investors' underreaction over the subsequent month. Empirical evidence so far suggests that investors' underreaction is driven by scheduled news (Birru, 2013; George, Hwang and Li, 2014). My results on the investors' distorted belief updating process provide strong evidence of a longer-dated, monthly investor underreaction to unscheduled news, indicating that unscheduled news items require more time to be reflected within stock prices. Fourth, I contribute to the literature on empirical asset pricing for global equity markets by using an international sample and extended metrics. Most literature on news-induced momentum (Chan, 2003; Gutierrez and Kelly, 2008; Hillert, Jacobs and Müller, 2014; Jiang, Li and Wang, 2021) concentrates solely on the U.S. stock market. Therefore, I add to the ongoing discussion about the investor underreaction hypothesis and its economic channels by providing non-U.S. out-of-sample evidence (Hou, Xue and Zhang, 2018) for the anchoring bias and investor underreaction to firm-specific news.

To conclude, the three essays in this dissertation examine research questions on empirical asset pricing in global stock markets. The findings motivate several avenues for future research.

The *first* essay sheds light on investors' underreaction to the arrival of firm-specific news from fundamentally linked firms. Given that investors are limited in their attention, it seems relevant to investigate further the actual behavior of private and institutional investors after the arrival of firm-specific news.

The *second* essay suggests that machine learning models help better to predict the cross-section of emerging market stock returns. Due to the increasing number of anomalies with explanatory power and the importance of macroeconomic predictors, it is relevant to analyze the inclusion of these features within the different machine learning models to further improve the model performance.

The *third* essay provides empirical evidence on how the anchoring effect distorts the investor's belief updating process after the arrival of firm-specific news. Future research could investigate the specific news categories and timing that cause underreaction to firm-specific news across multiple asset classes and explore the role of time perspective in explaining underreaction.

A Chapter 1

Datastream sample definition

Constituent lists

Datastream comprises three types of constituent lists: (1) research lists, (2) Worldscope lists, and (3) dead lists. By using dead lists, we ensure that any survivorship bias is obviated. For each country, we use the union of all available lists and eliminate any duplicates. As a result, one list remains for each country to be used in the subsequent static filter process. Table A. 1 and Table A. 2 provide an overview of the constituent lists for developed and emerging markets that are used in this study.

Table A. 1
Constituent lists developed markets

The table contains the research lists, Worldscope lists and dead lists of developed markets countries in my sample.

Country	List	Country	List	Country	List
Australia	DEADAU	**Hong Kong**	DEADHK	**Spain**	DEADES
	FAUALL		FHKALL		WSCOPEES
	WSCOPEAU		WSCOPEHK		FESALL
Austria	WSCOPEOE	**Ireland**	WSCOPEIR		FSPDOM
	DEADAT		FIEALL		FSPNQ
	FATALL		DEADIE	**Sweden**	WSCOPESD
	FOSTDCT	**Israel**	DEADIL		FSEALL
	FOSTOM		WSCOPEIS		FXSTOALL
Belgium	FBEALL		FILALL		DEADSE
	WSCOPEBG	**Italy**	FITALL	**Switzerland**	WSCOPESW
	DEADBE		DEADIT		FCHALLP
Canada	DEADCA1		WSCOPEIT		DEADCH
	...	**Japan**	WSCOPEJP	**United Kingdom**	DEADGB
	DEADCA6		FJPALL		...
	WSCOPECN		FJPCONS		DEADGB7
	FXTSEALL		FTOKYO		FGBALL
	FCAALL		FXTKSALL		WSCOPEUK
Denmark	FDKALL		DEADJP	**United States**	WSUS1
	WSCOPEDK	**Netherlands**	DEADNL		...
	DEADDK		FNLALL		WSUS26
Finland	FFIALL		WSCOPENL		FUSALL1
	WSCOPEFN	**New**	WSCOPENZ		...
	DEADFI	**Zealand**	FNZALL		FUSALL7
France	DEADFR		DEADNZ		FUSALLA
	WSCOPEFR	**Norway**	DEADNO		...
	FFRALL		FNOALL		FUSALLZ
Germany	DEADDE1		WSCOPENW		DEADUS1
	...	**Portugal**	WSCOPEPT		...
	DEADDE9		FPTALL		DEADUS12
	FGKURS		DEADPT		
	FDEALLP	**Singapore**	DEADSG		
	WSCOPEBD		FSGALL		
			FXSESM		
			WSCOPESG		

Table A. 2
Constituent lists emerging markets

The table contains the research lists, Worldscope lists and dead lists of emerging markets countries in my sample.

Country	List	Country	List	Country	List
Argentina	DEADAR	**Jordan**	DEADJO	**Russia**	DEADRU
	FARALL		FJOALL		FRUSXALL
	WSCOPEAR		WSCOPEJO		WSCOPERS
Brazil	DEADBR	**Korea**	DEADKR	**Saudi Arabia**	DEADSA
	FBRALL		FKRALL		FSAALL
	WSCOPEBR		WSCOPEKO		WSCOPESI
Chile	DEADCL	**Kuwait**	DEADKW	**South Africa**	DEADZA
	FCLALL		FKWALL		FZAALL
	WSCOPECL		WSCOPEKW		WSCOPESA
China	DEADCN	**Malaysia**	DEADMY	**Sri Lanka**	DEADLK
	FCNALL		FACE		FLKALL
	WSCOPECH		FMYALL		WSCOPECY
Czechia	DEADCZ		WSCOPEMY	**Taiwan**	DEADTW
	FCZALL	**Mexico**	DEADMX		FROCOALL
	WSCOPECZ		FMXALL		FTWALL
Egypt	DEADEG		WSCOPEMX		WSCOPETA
	FEGALL	**Morocco**	DEADMA	**Thailand**	DEADTH
	WSCOPEEY		FMAALL		FTHALL
Greece	DEADGR		WSCOPEMC		WSCOPETH
	FGRALL	**Pakistan**	DEADPK	**Turkey**	DEADTR
	WSCOPEGR		FPKALL		FTRALL
India	DEADIN		WSCOPEPK		WSCOPETK
	FINALL	**Philippines**	DEADPH	**UAE**	DEADAE
	FINCONS		FPHALL		FAEALL
	FXBOMALL		WSCOPEPH		FXADSALL
	FXNSEALL	**Poland**	DEADPL		FXDFMALL
	WSCOPEIN		FPLALL		WSCOPEAE
Indonesia	DEADID		FPOLCM		
	FIDALL		WSCOPEPO		
	WSCOPEID				

Static screens

I restrict the sample to common equity stocks by applying several static screens, as shown in Table A. 3. Screens (1) to (7) are straightforward to apply and common in the literature.

Table A. 3

Static screens

The table displays the static screens applied in our study, mainly following Ince and Porter (2006), Schmidt et al. (2017) and Griffin, Kelly and Nardari (2010). Column 3 lists the Datastream items involved (on the left of the equals sign) and the values which we set them to in the filter process (to the right of the equals sign). Column 4 indicates the source of the screens.

Nr.	Description	Datastream item(s) involved	Source
(1)	For firms with more than one security, only the one with the biggest market capitalization and liquidity is used.	MAJOR = Y	Schmidt et al. (2017)
(2)	The type of security must be equity.	TYPE = EQ	Ince and Porter (2006)
(3)	Only the primary quotations of a security are analyzed.	ISINID = P	Fong, Holden and Trzcinka (2017)
(4)	Firms are located in the respective domestic country.	GEOGN = country shortcut	Ince and Porter (2006)
(5)	Securities are listed in the respective domestic country.	GEOLN = country shortcut	Griffin, Kelly and Nardari (2010)
(6)	Securities whose quoted currency is different to the one of the associated country are disregarded.[a]	PCUR = currency shortcut of the country	Griffin, Kelly and Nardari (2010)
(7)	Securities whose ISIN country code is different to the one of the associated country are disregarded.[b]	GGISN = country shortcut	Annaert, Ceuster and Verstegen (2013)
(8)	Securities whose name fields indicate non-common stock affiliation are disregarded.	NAME, ENAME, ECNAME	Ince and Porter (2006), Campbell, Cowan and Salotti (2010), Griffin, Kelly and Nardari (2010) and Karolyi, Lee and van Dijk (2012)

[a] In this filter rule, the respective pre-euro currencies are also accepted for countries within the euro-zone. Moreover, in Russia 'USD' is accepted as currency, in addition to 'RUB'.

[b] In Hong Kong, ISIN country codes equal to 'BM' or 'KY' and in the Czech Republic ISIN country codes equal to 'CS' are also accepted.

Screen (8) relates to, among others, to work by the following: Ince and Porter (2006), Campbell, Cowan and Salotti (2010), Griffin, Kelly and Nardari (2010), Karolyi, Lee and van Dijk (2012). The authors provide generic filter rules to exclude non-common equity securities from Refinitiv Datastream. we apply the identified keywords and match them with the security names provided by Datastream. A security is excluded from the sample in the event that a keyword coincides with part of the security name. The following three Datastream items store security names and are applied to the keyword filters: 'NAME', 'ENAME', and 'ECNAME'. Table A. 4 gives an overview of the keywords used.

Table A. 4

Generic keyword deletions

The table reports generic keywords searched for in the names of all stocks of all countries. If a harmful keyword is detected as part of the name of a stock, the respective stock is removed from the sample.

Non-common equity	**Keywords**
Duplicates	1000DUPL, DULP, DUP, DUPE, DUPL, DUPLI, DUPLICATE, XSQ, XETa
Depository receipts	ADR, GDR
Preferred stock	PF, 'PF', PFD, PREF, PREFERRED, PRF
Warrants	WARR, WARRANT, WARRANTS, WARRT, WTS, WTS2
Debt	%, DB, DCB, DEB, DEBENTURE, DEBENTURES, DEBT
Unit trusts	.IT, .ITb, TST, INVESTMENT TRUST, RLST IT, TRUST, TRUST UNIT, TRUST UNITS, TST, TST UNIT, TST UNITS, UNIT, UNIT TRUST, UNITS, UNT, UNT TST, UT
ETFs	AMUNDI, ETF, INAV, ISHARES, JUNGE, LYXOR, X-TR
Expired securities	EXPD, EXPIRED, EXPIRY, EXPY
Miscellaneous (mainly taken from Ince and Porter (2006))	ADS, BOND, CAP.SHS, CONV, DEFER, DEP, DEPY, ELKS, FD, FUND, GW.FD, HI.YIELD, HIGH INCOME, IDX, INC.&GROWTH, INC.&GW, INDEX, LP, MIPS, MITS, MITT, MPS, NIKKEI, NOTE, OPCVM, ORTF, PARTNER, PERQS, PFC, PFCL, PINES, PRTF, PTNS, PTSHP, QUIBS, QUIDS, RATE, RCPTS, REAL EST, RECEIPTS, REIT, RESPT, RETUR, RIGHTS, RST, RTN.INC, RTS, SBVTG, SCORE, SPDR, STRYPES, TOPRS, UTS, VCT, VTG.SAS, XXXXX, YIELD, YLD

In addition, Griffin, Kelly and Nardari (2010) introduce specific keywords for individual countries. The keywords are thus applied to the security names of single countries only. For example, German security names are parsed to contain the word 'GENUSSSCHEINE', which declares the security to be a non-common equity. In Table A. 5, we give an overview of country-specific keyword deletions conducted in our study.

Table A. 5

Country-specific keyword deletions

The table reports country-specific keywords searched for in the names of all stocks of the respective countries. If a harmful keyword is detected as part of the name of a stock, the respective stock is removed from the sample.

Country	Keywords
Australia	PART PAID, RTS DEF, DEF SETT, CDI
Austria	PC, PARTICIPATION CERTIFICATE, GENUSSSCHEINE, GENUSSCHEINE
Belgium	VVPR, CONVERSION, STRIP
Brazil	PN, PNA, PNB, PNC, PND, PNE, PNF, PNG, RCSA, RCTB
Canada	EXCHANGEABLE, SPLIT, SPLITSHARE, VTG\\., SBVTG\\., VOTING, SUB VTG, SERIES
Denmark	\\)CSE\\)
Finland	USE
France	ADP, CI, SICAV, \\)SICAV\\), SICAV-
Germany	GENUSSCHEINE
Greece	PR
Indonesia	FB DEAD, FOREIGN BOARD
Israel	P1, 1, 5
Italy	RNC, RP, PRIVILEGIES
Korea	1P
Malaysia	'A'
Mexico	'L', 'C'
Netherlands	CERTIFICATE, CERTIFICATES, CERTIFICATES\\), CERT, CERTS, STK\\.
New Zealand	RTS, RIGHTS
Philippines	PDR
South Africa	N', OPTS\\., CPF\\., CUMULATIVE PREFERENCE
Sweden	CONVERTED INTO, USE, CONVERTED-, CONVERTED - SEE
Switzerland	CONVERTED INTO, CONVERSION, CONVERSION SEE
United Kingdom	PAID, CONVERSION TO, NON VOTING, CONVERSION 'A'

Dynamic screens

For the securities remaining from the static screens above, we obtained return and market capitalization data from Datastream and accounting data from Worldscope. Several dynamic screens that are common in the literature were installed in order to account for data errors, mainly within return characteristics. The dynamic screens are shown in Table A. 6.

Table A. 6
Dynamic screens

The table displays the dynamic screens applied to the data in our study, following Ince and Porter (2006), Griffin, Kelly and Nardari (2010), Jacobs (2016) and Schmidt et al. (2017). Column 3 lists the respective Datastream items. Column 4 refers to the source of the screens.

Nr.	Description	Datastream item(s) involved	Source
(1)	We delete the zero returns at the end of the return time-series that exist because in the case of a delisting, Datastream displays stale prices from the date of delisting until the end of the respective time-series. We also delete the associated market capitalizations.	RI, MV	Ince and Porter (2006)
(2)	We delete the associated returns and market capitalizations in case of abnormal prices (unadjusted prices > 1000000).	RI, MV, UP	The screen originally stems from Schmidt et al. (2017), however we employ it on unadjusted price.
(3)	We delete monthly (daily) returns and the associated market capitalizations if returns exceed 990% (200%).	RI, MV	Griffin, Kelly and Nardari (2010); Schmidt et al. (2017)
(4)	We delete monthly returns and the associated market capitalizations in the case of strong return reversals, defined as $(1+r_{t-1})(1+r_t)-1 < 0.5$ given that either r_{t-1} or $r_t \geq 3.0$.	RI, MV	Ince and Porter (2006)
(5)	We delete daily returns and the associated market capitalizations in the case of strong return reversals, defined as $(1+r_{t-1})(1+r_t)-1 < 0.2$ with r_{t-1} or $r_t \geq 1.0$.	RI, MV	Griffin, Kelly and Nardari (2010); Jacobs (2016)

Tables

Table A. 7
Portfolio time-series regression of Mkt^{FL}

This table reports the Jensen's alpha of the four different regression models and long, short, and long–short portfolio based on quintiles of Mkt^{FL} using country-neutral break points. Each regression model includes a different set of tradable common risk factors $MKTRF$, SMB, HML, UMD, and LIQ (Fama and French, 1993; Carhart, 1997; Pástor and Stambaugh, 2003) Panel A shows the results for the equal- weighted portfolios and Panel B shows the results for the value-weighted counterpart. t-statistics are reported in parentheses. All standard-errors are adjusted using Newey and West (1987). The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage.

	$CAPM$			$FF3$			$FF4$			$FF4+LIQ$		
Variable	Low	High	H-L	Low	High	H-L	Low	High	H-L	Low	High	H-L
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: Equal-Weighted												
α	0.021	0.297	0.276	-0.027	0.274	0.302	0.002	0.405	0.403	0.005	0.407	0.402
	(0.18)	(1.51)	(1.45)	(-0.29)	(1.82)	(1.58)	(0.02)	(2.53)	(2.08)	(0.05)	(2.63)	(2.06)
$Mkt-RF$	1.140	1.191	0.050	1.059	1.110	0.051	1.050	1.072	0.022	1.013	1.050	0.036
	(20.45)	(26.36)	(0.69)	(27.16)	(24.94)	(0.76)	(26.49)	(24.99)	(0.31)	(23.60)	(23.55)	(0.49)
SMB				-0.940	-0.869	0.071	-0.933	-0.836	0.097	-0.981	-0.865	0.116
				(-7.63)	(-8.45)	(0.42)	(-7.46)	(-8.83)	(0.60)	(-8.06)	(-7.97)	(0.62)
HML				0.092	0.002	-0.091	0.073	-0.085	-0.158	0.159	-0.034	-0.193
				(1.81)	(0.02)	(-1.05)	(1.35)	(-1.03)	(-1.45)	(2.92)	(-0.38)	(-1.69)
MOM							-0.042	-0.189	-0.147	-0.046	-0.192	-0.145
							(-0.95)	(-2.63)	(-1.49)	(-1.11)	(-2.68)	(-1.46)
LIQ										0.230	0.137	-0.093
										(2.74)	(0.99)	(-0.44)

Continued on next page

Table 1.3 continued

Variable	$CAPM$			$FF3$			$FF4$			$FF4+LIQ$		
	Low	High	H-L	Low	High	H-L	Low	High	H-L	Low	High	H-L
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel B: Value-Weighted												
α	0.021	0.297	0.276	-0.027	0.274	0.302	0.002	0.405	0.403	0.005	0.407	0.402
	(0.18)	(1.51)	(1.45)	(-0.29)	(1.82)	(1.58)	(0.02)	(2.53)	(2.08)	(0.05)	(2.63)	(2.06)
$Mkt-RF$	1.140	1.191	0.050	1.059	1.110	0.051	1.050	1.072	0.022	1.013	1.050	0.036
	(20.45)	(26.36)	(0.69)	(27.16)	(24.94)	(0.76)	(26.49)	(24.99)	(0.31)	(23.60)	(23.55)	(0.49)
SMB				-0.940	-0.869	0.071	-0.933	-0.836	0.097	-0.981	-0.865	0.116
				(-7.63)	(-8.45)	(0.42)	(-7.46)	(-8.83)	(0.60)	(-8.06)	(-7.97)	(0.62)
HML				0.092	0.002	-0.091	0.073	-0.085	-0.158	0.159	-0.034	-0.193
				(1.81)	(0.02)	(-1.05)	(1.35)	(-1.03)	(-1.45)	(2.92)	(-0.38)	(-1.69)
MOM							-0.042	-0.189	-0.147	-0.046	-0.192	-0.145
							(-0.95)	(-2.63)	(-1.49)	(-1.11)	(-2.68)	(-1.46)
LIQ										0.230	0.137	-0.093
										(2.74)	(0.99)	(-0.44)

Table A. 8

Portfolio time-series regression of $Noise^{FL}$

This table reports the Jensen's alpha of the four different regression models and long, short, and long–short portfolio based on quintiles of $Noise^{FL}$ using country-neutral break points. Each regression model includes a different set of tradable common risk factors $MKTRF$, SMB, HML, UMD, and LIQ (Fama and French, 1993; Carhart, 1997; Pástor and Stambaugh, 2003) Panel A shows the results for the equal- weighted portfolios and Panel B shows the results for the value-weighted counterpart. t-statistics are reported in parentheses. All standard-errors are adjusted using Newey and West (1987). The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage.

	$CAPM$			$FF3$			$FF4$			$FF4+LIQ$		
Variable	Low	High	H-L	Low	High	H-L	Low	High	H-L	Low	High	H-L
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: Equal-Weighted												
α	0.176	0.139	-0.037	0.158	0.052	-0.106	0.283	0.105	-0.179	0.287	0.105	-0.182
	(1.25)	(0.99)	(-0.39)	(1.66)	(0.70)	(-1.44)	(2.70)	(1.35)	(-2.23)	(3.17)	(1.39)	(-2.32)
$Mkt-RF$	1.178	1.091	-0.087	1.081	1.040	-0.041	1.045	1.025	-0.020	1.003	1.020	0.017
	(27.45)	(28.26)	(-2.16)	(32.46)	(41.93)	(-1.21)	(38.23)	(37.32)	(-0.71)	(37.93)	(35.35)	(0.55)
SMB				-1.019	-0.731	0.288	-0.988	-0.718	0.270	-1.042	-0.725	0.317
				(-11.47)	(-9.81)	(3.89)	(-10.72)	(-10.39)	(3.18)	(-13.49)	(-10.50)	(4.03)
HML				-0.033	0.255	0.288	-0.116	0.221	0.336	-0.019	0.234	0.252
				(-0.50)	(6.84)	(4.25)	(-1.97)	(4.63)	(5.75)	(-0.28)	(4.85)	(3.80)
MOM							-0.181	-0.075	0.105	-0.186	-0.076	0.109
							(-5.22)	(-2.80)	(3.34)	(-5.27)	(-2.77)	(3.65)
LIQ										0.260	0.034	-0.225
										(4.80)	(0.66)	(-4.70)

Continued on next page

Table 1.3 continued

Variable	$CAPM$			$FF3$			$FF4$			$FF4+LIQ$		
	Low	High	H-L	Low	High	H-L	Low	High	H-L	Low	High	H-L
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel B: Value-Weighted												
α	0.176	0.139	-0.037	0.158	0.052	-0.106	0.283	0.105	-0.179	0.287	0.105	-0.182
	(1.25)	(0.99)	(-0.39)	(1.66)	(0.70)	(-1.44)	(2.70)	(1.35)	(-2.23)	(3.17)	(1.39)	(-2.32)
$Mkt-RF$	1.178	1.091	-0.087	1.081	1.040	-0.041	1.045	1.025	-0.020	1.003	1.020	0.017
	(27.45)	(28.26)	(-2.16)	(32.46)	(41.93)	(-1.21)	(38.23)	(37.32)	(-0.71)	(37.93)	(35.35)	(0.55)
SMB				-1.019	-0.731	0.288	-0.988	-0.718	0.270	-1.042	-0.725	0.317
				(-11.47)	(-9.81)	(3.89)	(-10.72)	(-10.39)	(3.18)	(-13.49)	(-10.50)	(4.03)
HML				-0.033	0.255	0.288	-0.116	0.221	0.336	-0.019	0.234	0.252
				(-0.50)	(6.84)	(4.25)	(-1.97)	(4.63)	(5.75)	(-0.28)	(4.85)	(3.80)
MOM							-0.181	-0.075	0.105	-0.186	-0.076	0.109
							(-5.22)	(-2.80)	(3.34)	(-5.27)	(-2.77)	(3.65)
LIQ										0.260	0.034	-0.225
										(4.80)	(0.66)	(-4.70)

Table A. 9
Underreaction coefficients

This table shows returns on the firm-specific and market news as well as noise spillover portfolio and the corresponding underreaction coefficients. At the beginning of every month, stocks are ranked in ascending order based on the corresponding spillover at the end of the previous month. At the beginning of each month the stocks are sorted into 5 country-neutral portfolios. All stocks are equal-weighted within a given portfolio, and the portfolios are rebalanced every calendar month to maintain value weights Panel A bases the analysis on all stocks, Panel B focuses on large stocks whereas Panel c focuses on small stocks. I follow Fama and French (2012, 2017) to calculate the country specific breakpoints. Large stocks are defined as the largest stocks which together account for 90% of a country's aggregated market capitalization. Small stocks are defined as those stocks that comprise the remaining 10% of aggregated market capitalization. Each panel reports the average cumulative returns on the long–short portfolio formed on the respective spillover in month t. RET_t is the focal firm stock return in month t. $RET_{t+1,t+h}$ is the cumulative return over the subsequent h, for $h \in \{3, 6, 9\}$, months. URC (underreaction coefficient) is defined as the fraction of total returns from month t to month $t+h$ that occurs in month t ($URC = RET_t/(RET_t + RET_{t+1,t+h})$). t-statistics are shown below the coefficient estimates. In the case of URC the t-statistics represent the distance of the coefficient from one, which is the case of no underreaction. The sample consists of all stocks for the period between January 1992 and December 2021, with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage.

	FS^{FL}			Mkt^{FL}			$Noise^{FL}$		
	$h=3$	$h=6$	$h=9$	$h=3$	$h=6$	$h=9$	$h=3$	$h=6$	$h=9$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: All Stocks									
RET_t	6.866	6.862	6.859	2.590	2.612	2.633	0.465	0.459	0.472
	(46.88)	(46.53)	(46.21)	(15.35)	(15.41)	(15.44)	(3.59)	(3.52)	(3.60)
$RET_{t+1,t+h}$	1.074	1.530	1.649	0.291	0.072	0.059	0.017	-0.035	-0.266
	(4.62)	(4.79)	(4.23)	(0.98)	(0.17)	(0.11)	(0.09)	(-0.13)	(-0.81)
URC	0.865	0.818	0.806	0.899	0.973	0.978	0.964	1.083	2.300
	(2.82)	(2.90)	(3.02)	(1.59)	(0.39)	(0.34)	(0.59)	(1.76)	(31.84)
Panel B: Large Stocks									
RET_t	6.760	6.755	6.752	2.566	2.592	2.613	0.443	0.437	0.451
	(46.29)	(45.94)	(45.64)	(15.16)	(15.28)	(15.31)	(3.33)	(3.27)	(3.36)
$RET_{t+1,t+h}$	0.965	1.430	1.503	0.295	0.049	0.011	0.045	-0.053	-0.346
	(4.09)	(4.33)	(3.77)	(1.00)	(0.11)	(0.02)	(0.22)	(-0.20)	(-1.01)
URC	0.875	0.825	0.818	0.897	0.982	0.996	0.908	1.137	4.290
	(2.37)	(2.30)	(2.62)	(1.78)	(0.27)	(0.06)	(1.57)	(3.05)	(74.99)
Panel C: Small Stocks									
RET_t	7.436	7.434	7.426	2.810	2.844	2.867	0.290	0.271	0.271
	(44.80)	(44.45)	(44.12)	(13.59)	(13.71)	(13.74)	(2.31)	(2.15)	(2.14)
$RET_{t+1,t+h}$	2.033	2.805	3.587	0.172	-0.095	0.012	0.231	0.201	0.051
	(8.45)	(8.36)	(8.30)	(0.46)	(-0.19)	(0.02)	(1.17)	(0.73)	(0.15)
URC	0.785	0.726	0.674	0.942	1.034	0.996	0.557	0.575	0.843
	(6.28)	(5.41)	(5.29)	(1.09)	(0.66)	(0.08)	(7.66)	(9.06)	(4.68)

Table A. 10
Limited attention

This table reports the estimated regression coefficients and Newey-West t-statistics (in parentheses) from Fama-MacBeth cross-sectional regressions predicting one-month ahead excess stock returns with limited attention proxies. I interact the firm-specific and market news as well as noise spillover individually with limited attention dummies. The indicator variables that take the value of one if the underlying variable is above the median in the cross-section, and zero otherwise. The sample consists of all stocks for the period between April 2000 and ends in March 2020 , with a minimum of one analyst co-covering the stocks, and a minimum of 30 valid country-month observations. FS^{FL} is the weighted average VAR-based firm-specific news in the previous month of stocks that are connected through shared analyst coverage. Mkt^{FL} is the weighted average VAR-based market news in the previous month of stocks that are connected through shared analyst coverage. $Noise^{FL}$ is the weighted average VAR-based noise in the previous month of stocks that are connected through shared analyst coverage. RET^{FL}_{t+h} is the monthly contemporaneous returns from stocks that are connected through shared analyst coverage at time $t+h$. ANA is the analyst coverage, which is the number of sell-side analysts forecasting annual firm earnings in each month t. MV is the product of the closing price and the number of shares outstanding. $OWNER$ are the holdings by all institutional investors as a fraction of the market capitalization. $COMP$ is equal to 1 if cross-sectional average of the dummies of ANA, MV, and $OWNER$ is larger than 0.5.

	ANA		MV		$OWNER$		$COMP$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$LA \times FS^{FL}$	-0.049	-0.041	-0.086	-0.073	-0.005	-0.003	-0.061	-0.051
	(-3.64)	(-3.82)	(-7.20)	(-7.53)	(-0.53)	(-0.38)	(-4.23)	(-4.34)
$LA \times Mkt^{FL}$	-0.080	-0.101	-0.103	-0.114	-0.017	-0.022	-0.101	-0.112
	(-1.55)	(-2.12)	(-2.06)	(-2.64)	(-0.49)	(-0.75)	(-1.91)	(-2.26)
$LA \times Noise^{FL}$	-0.102	-0.094	-0.044	-0.052	0.050	0.055	-0.072	-0.072
	(-1.65)	(-1.90)	(-0.59)	(-0.85)	(1.11)	(1.46)	(-1.31)	(-1.54)
FS^{FL}	0.116	0.088	0.134	0.103	0.102	0.076	0.123	0.094
	(13.16)	(14.02)	(13.37)	(15.13)	(10.53)	(10.41)	(13.76)	(14.85)
Mkt^{FL}	0.143	0.092	0.159	0.104	0.125	0.068	0.149	0.093
	(2.76)	(3.72)	(2.68)	(3.70)	(2.43)	(3.10)	(2.75)	(3.48)
$Noise^{FL}$	0.060	0.117	0.034	0.098	0.010	0.068	0.033	0.096
	(1.82)	(5.09)	(0.80)	(3.10)	(0.26)	(2.67)	(0.84)	(3.59)
LA	0.200	0.183	0.105	0.097	-0.168	-0.175	0.126	0.097
	(1.92)	(1.97)	(0.78)	(0.80)	(-1.90)	(-2.36)	(0.95)	(0.87)
RET^{FL}_{t+h}		0.528		0.530		0.530		0.534
		(49.04)		(48.96)		(49.50)		(49.56)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FF-38	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2 (%)	16.00	18.13	15.95	18.10	15.97	18.12	15.82	18.02
Avg. Obs	11880	11880	11880	11880	11880	11880	11880	11880

B Chapter 2

Datastream sample definition

Constituent lists

Datastream comprises three types of constituent lists: (1) research lists, (2) Worldscope lists, and (3) dead lists. By using dead lists, we ensure that any survivorship bias is obviated. For each country, we use the union of all available lists and eliminate any duplicates. As a result, one list remains for each country to be used in the subsequent static filter process. Table A. 1 provides an overview of the constituent lists for emerging markets that are used in our study.

Static screens

We restrict our sample to common equity stocks by applying several static screens, as shown in Table A. 2. Screens (1) to (7) are straightforward to apply and common in the literature.

Screen (8) relates to, among others, to work by the following: Ince and Porter (2006), Campbell, Cowan and Salotti (2010), Griffin, Kelly and Nardari (2010), Karolyi, Lee and van Dijk (2012). The authors provide generic filter rules to exclude non-common equity securities from Refinitiv Datastream. We apply the identified keywords and match them with the security names provided by Datastream. A security is excluded from the sample in the event that a keyword coincides with part of the security name. The following three Datastream items store security names and are applied to the keyword filters: 'NAME', 'ENAME', and 'ECNAME'. Table A. 3 gives an overview of the keywords used.

In addition, Griffin, Kelly and Nardari (2010) introduce specific keywords for individual countries. The keywords are thus applied to the security names of single countries only.

Table B. 1
Constituent lists emerging markets

The table contains the research lists, Worldscope lists and dead lists of emerging markets countries in our sample.

Country	List	Country	List	Country	List
Argentina	DEADAR	**Israel**	DEADIL	**Portugal**	WSCOPEPT
	FARALL		WSCOPEIS		FPTALL
	WSCOPEAR		FILALL		DEADPT
Brazil	DEADBR	**Jordan**	DEADJO	**Qatar**	DEADQA
	FBRALL		FJOALL		FQAALL
	WSCOPEBR		WSCOPEJO		WSCOPEQA
Chile	DEADCL	**Korea**	DEADKR	**Russia**	DEADRU
	FCLALL		FKRALL		FRUSXALL
	WSCOPECL		WSCOPEKO		WSCOPERS
China	DEADCN	**Kuwait**	DEADKW	**Saudi Arabia**	DEADSA
	FCNALL		FKWALL		FSAALL
	WSCOPECH		WSCOPEKW		WSCOPESI
Colombia	DEADCO	**Malaysia**	DEADMY	**South Africa**	DEADZA
	FCOALL		FACE		FZAALL
	WSCOPECB		FMYALL		WSCOPESA
Czechia	DEADCZ		WSCOPEMY	**Sri Lanka**	DEADLK
	FCZALL	**Mexico**	DEADMX		FLKALL
	WSCOPECZ		FMXALL		WSCOPECY
Egypt	DEADEG		WSCOPEMX	**Taiwan**	DEADTW
	FEGALL	**Morocco**	DEADMA		FROCOALL
	WSCOPEEY		FMAALL		FTWALL
Greece	DEADGR		WSCOPEMC		WSCOPETA
	FGRALL	**Pakistan**	DEADPK	**Thailand**	DEADTH
	WSCOPEGR		FPKALL		FTHALL
Hungary	DEADHU		WSCOPEPK		WSCOPETH
	FHUALL	**Peru**	DEADPE	**Turkey**	DEADTR
	WSCOPEHN		FPEALL		FTRALL
India	DEADIN		WSCOPEPE		WSCOPETK
	FINALL	**Philippines**	DEADPH	**UAE**	DEADAE
	FINCONS		FPHALL		FAEALL
	FXBOMALL		WSCOPEPH		FXADSALL
	FXNSEALL	**Poland**	DEADPL		FXDFMALL
	WSCOPEIN		FPLALL		WSCOPEAE
Indonesia	DEADID		FPOLCM		
	FIDALL		WSCOPEPO		
	WSCOPEID				

For example, German security names are parsed to contain the word 'GENUSSSCHEINE', which declares the security to be a non-common equity. In Table A. 4, we give an overview of country-specific keyword deletions conducted in our study.

Dynamic screens

For the securities remaining from the static screens above, we obtained return and market capitalization data from Datastream and accounting data from Worldscope. Several dynamic screens that are common in the literature were installed in order to account for data errors, mainly within return characteristics. The dynamic screens are shown in Table A. 5.

Table B. 2
Static screens

The table displays the static screens applied in our study, mainly following Ince and Porter (2006), Schmidt et al. (2017) and Griffin, Kelly and Nardari (2010). Column 3 lists the Datastream items involved (on the left of the equals sign) and the values which we set them to in the filter process (to the right of the equals sign). Column 4 indicates the source of the screens.

Nr.	Description	Datastream item(s) involved	Source
(1)	For firms with more than one security, only the one with the biggest market capitalization and liquidity is used.	MAJOR = Y	Schmidt et al. (2017)
(2)	The type of security must be equity.	TYPE = EQ	Ince and Porter (2006)
(3)	Only the primary quotations of a security are analyzed.	ISINID = P	Fong, Holden and Trzcinka (2017)
(4)	Firms are located in the respective domestic country.	GEOGN = country shortcut	Ince and Porter (2006)
(5)	Securities are listed in the respective domestic country.	GEOLN = country shortcut	Griffin, Kelly and Nardari (2010)
(6)	Securities whose quoted currency is different to the one of the associated country are disregarded.[a]	PCUR = currency shortcut of the country	Griffin, Kelly and Nardari (2010)
(7)	Securities whose ISIN country code is different to the one of the associated country are disregarded.[b]	GGISN = country shortcut	Annaert, Ceuster and Verstegen (2013)
(8)	Securities whose name fields indicate non-common stock affiliation are disregarded.	NAME, ENAME, ECNAME	Ince and Porter (2006), Campbell, Cowan and Salotti (2010), Griffin, Kelly and Nardari (2010) and Karolyi, Lee and van Dijk (2012)

[a] In this filter rule, the respective pre-euro currencies are also accepted for countries within the euro-zone. Moreover, in Russia 'USD' is accepted as currency, in addition to 'RUB'.
[b] In Hong Kong, ISIN country codes equal to 'BM' or 'KY' and in the Czech Republic ISIN country codes equal to 'CS' are also accepted.

Table B. 3

Generic keyword deletions

The table reports generic keywords searched for in the names of all stocks of all countries. If a harmful keyword is detected as part of the name of a stock, the respective stock is removed from the sample.

Non-common equity	Keywords
Duplicates	1000DUPL, DULP, DUP, DUPE, DUPL, DUPLI, DUPLICATE, XSQ, XETa
Depository receipts	ADR, GDR
Preferred stock	PF, 'PF', PFD, PREF, PREFERRED, PRF
Warrants	WARR, WARRANT, WARRANTS, WARRT, WTS, WTS2
Debt	%, DB, DCB, DEB, DEBENTURE, DEBENTURES, DEBT
Unit trusts	.IT, .ITb, TST, INVESTMENT TRUST, RLST IT, TRUST, TRUST UNIT, TRUST UNITS, TST, TST UNIT, TST UNITS, UNIT, UNIT TRUST, UNITS, UNT, UNT TST, UT
ETFs	AMUNDI, ETF, INAV, ISHARES, JUNGE, LYXOR, X-TR
Expired securities	EXPD, EXPIRED, EXPIRY, EXPY
Miscellaneous (mainly taken from Ince and Porter (2006))	ADS, BOND, CAP.SHS, CONV, DEFER, DEP, DEPY, ELKS, FD, FUND, GW.FD, HI.YIELD, HIGH INCOME, IDX, INC.&GROWTH, INC.&GW, INDEX, LP, MIPS, MITS, MITT, MPS, NIKKEI, NOTE, OPCVM, ORTF, PARTNER, PERQS, PFC, PFCL, PINES, PRTF, PTNS, PTSHP, QUIBS, QUIDS, RATE, RCPTS, REAL EST, RECEIPTS, REIT, RESPT, RETUR, RIGHTS, RST, RTN.INC, RTS, SBVTG, SCORE, SPDR, STRYPES, TOPRS, UTS, VCT, VTG.SAS, XXXXX, YIELD, YLD

Table B. 4

Country-specific keyword deletions

The table reports country-specific keywords searched for in the names of all stocks of the respective countries. If a harmful keyword is detected as part of the name of a stock, the respective stock is removed from the sample.

Country	Keywords
Brazil	PN, PNA, PNB, PNC, PND, PNE, PNF, PNG, RCSA, RCTB
Greece	PR
Indonesia	FB DEAD, FOREIGN BOARD
Israel	P1, 1, 5
Korea	1P
Mexico	'L', 'C'
Peru	INVERSION, INVN, INV
Philippines	PDR
South Africa	N', OPTS\\., CPF\\., CUMULATIVE PREFERENCE

Table B. 5

Dynamic screens

The table displays the dynamic screens applied to the data in our study, following Ince and Porter (2006), Griffin, Kelly and Nardari (2010), Jacobs (2016) and Schmidt et al. (2017). Column 3 lists the respective Datastream items. Column 4 refers to the source of the screens.

Nr.	Description	Datastream item(s) involved	Source
(1)	We delete the zero returns at the end of the return time-series that exist because in the case of a delisting, Datastream displays stale prices from the date of delisting until the end of the respective time-series. We also delete the associated market capitalizations.	RI, MV	Ince and Porter (2006)
(2)	We delete the associated returns and market capitalizations in case of abnormal prices (unadjusted prices $>$ 1000000).	RI, MV, UP	The screen originally stems from Schmidt et al. (2017), however we employ it on unadjusted price.
(3)	We delete monthly (daily) returns and the associated market capitalizations if returns exceed 990% (200%).	RI, MV	Griffin, Kelly and Nardari (2010); Schmidt et al. (2017)
(4)	We delete monthly returns and the associated market capitalizations in the case of strong return reversals, defined as $(1+r_{t-1})(1+r_t)-1 < 0.5$ given that either r_{t-1} or $r_t \geq 3.0$.	RI, MV	Ince and Porter (2006)
(5)	We delete daily returns and the associated market capitalizations in the case of strong return reversals, defined as $(1+r_{t-1})(1+r_t)-1 < 0.2$ with r_{t-1} or $r_t \geq 1.0$.	RI, MV	Griffin, Kelly and Nardari (2010); Jacobs (2016)
(6)	We delete observations of stocks that show non-zero price changes in less than 50% of the traded months in the previous 12 months.	RI, MV	Griffin, Hirschey and Kelly (2011)
(7)	We delete observations of stocks in the lowest 3% of a country's aggregated market capitalization.	MV	Hanauer and Lauterbach (2019)

Characteristics definition

This section outlines the construction of characteristic variables that we use in the paper. For each characteristic, we give the respective Datastream and Worldscope items in parentheses, the category (past returns, investment, profitability, intangibles, value, or trading frictions) and frequency (monthly vs. yearly), plus the relevant reference. As described in Section 2.2, we use balance-sheet data from December in year t-1 for the stock returns from July of year t to June of year t + 1 as in Fama and French (1993).

A2ME (assets-to-market), Value, Yearly Assets-to-market cap is the ratio of total assets (WC02999) to market capitalization as at December t-1, as in Bhandari (1988).

AT (total assets), Trading Frictions, Yearly Total assets measured in USD (WC02999) as in Gandhi and Lustig (2015).

ATO (sales-to-assets), Profitability, Yearly As in Soliman (2008), we calculate net sales (WC01001) over lagged net operating assets. Net operating assets are defined following Hirshleifer et al. (2004) and are explained in the construction of NOA.

BEME (book-to-market), Value, Yearly Book-to-market is the ratio of book value of equity to market value of equity. We define the book value of equity as common equity (WC03501) plus deferred taxes (WC03263). If no deferred taxes are given, the book value of equity equals common equity (WC03501). The market value of equity is as of December t-1. See Rosenberg, Reid and Lanstein (1985) and Davis, Fama and French (2000).

$\mathbf{BEME_m}$ (monthly updated book-to-market), Value, Monthly Monthly updated book-to-market is the ratio of book value of equity to the most recent market value of equity. Book value of equity is defined as for $BEME$. The most recent market value of equity is of the end of month t to predict returns of month t+1 as in Asness (2011).

Beta (market beta), Trading Frictions, Monthly Following Lewellen and Nagel (2006), we calculate beta daily as the sum of the regression coefficients of daily excess returns on the local market excess return and one lag of the local market excess return for the past 12 months. We require at least 126 observations for valid beta estimates, as in Welch (2020).

C (cash-and-short-term-investment-to-assets), Value, Yearly The ratio of cash and short-term investments (WC02001) to total assets (WC02999), as in Palazzo (2012).

CbOPtA (cash-based operating profits-to-asset), Profitability, Yearly As in Ball et al. (2016), cash-based operating profits-to-asset is operating profits converted to a cash basis divided by total assets (WC02999). Following Ball et al. (2015), operating profits is net sales or revenues (WC01001) minus cost of goods sold (WC01501) minus selling, general, and administrative expenses (WC01101), excluding research and development expense (WC01201). The cash-based adjustment is the year-on-year change in deferred income (WC03262), plus change in accounts payable (WC03040), plus change in accrued expenses (WC03054 + WC03069), minus change in accounts receivable (WC02051), minus change in inventory (WC02101), minus prepaid expenses (WC02140), all divided by total assets. All changes are set to zero if missing.

CEI (composite equity issuance), Intangibles, Monthly Similar to Daniel and Titman (2006), we define composite equity issuance as the growth rate in the market capitalization not attributable to the total stock return R: $log(MC_{t-1}/MC_{t-13})-R_{(t-13,t-1)}$. To predict the returns of month t, $R_{(t-13,t-1)}$ is the cumulative log return (calculated via the total return index, Datastream item RI) from month $t-13$ to month $t-1$ and MC_{t-1} is the market capitalization (Datastream item MV) from the end of month $t-1$.

CF2P (cash flow-to-price), Value, Yearly Cash flow to price is the ratio of net cash flow from operating activities (WC04860) to the market capitalization as at December t-1, as in Lakonishok, Shleifer and Vishny (1994).

CTO (capital turnover), Profitability, Yearly We define capital turnover as the ratio of net sales (WC01001) to lagged total assets (WC02999), as in Haugen and Baker (1996).

D2A (capital intensity), Intangibles, Yearly Capital intensity is the ratio of depreciation and amortization (WC01151) over total assets (WC02999), as in Gorodnichenko and Weber (2016).

Debt2P (leverage), Value, Yearly Following Litzenberger and Ramaswamy (1979), debt to price is the ratio of total assets (WC02999) minus common equity (WC03501) to the market capitalization as of December t-1.

DPI2A (ratio of change in property, plants & equipment to total assets), Investment, Yearly Following Lyandres, Sun and Zhang (2007), we define the changes in PP&E and inventory as the annual change in gross property, plant, and equipment (WC02301) plus the annual change in inventory (WC02101) over lagged total assets (WC02999).

E2P (earnings-to-price), Value, Yearly Earnings to price is the ratio of income before extraordinary items (WC01551) to the market capitalization as at December t-1, as in Basu (1983).

FC2Y (fixed costs-to-sales), Profitability, Yearly As in Gorodnichenko and Weber (2016), fixed costs to sales is the sum of selling, general and administrative expenditures (WC01101) and research and development expenses (WC01201) over net sales (WC01001).

FreeCF (cash flow-to-book), Value, Yearly Following Hou, Karolyi and Kho (2011), we define cash flow to book as free cash flow to book value of equity. Free cash flow is calculated as net income (WC01551) plus depreciation and amortization (WC01151) minus changes in working capital minus capital expenditure (WC04601). The book value of equity is defined in the construction of BEME.

GP2A (gross profits-to-assets), Profitability, Yearly Gross profits-to-assets is net sales (WC01001) minus costs of goods sold (WC01051) divided by total assets (WC02999), as in Novy-Marx (2013).

Idiovol (idiosyncratic volatility with respect to the Fama and French (1993) three-factor model), Trading Frictions, Monthly As in Ang et al. (2006), we define idiosyncratic volatility as the standard deviation of the residuals from a regression of excess returns on a local Fama and French (1993) three-factor model. We use one month of daily data and require at least fifteen non-missing observations.

INV (investment), Investment, Yearly Investment is the percentage year-to-year growth rate of total assets (WC02999) following Cooper, Gulen and Schill (2008).

LME (market capitalization), Trading Frictions, Monthly Size is a stock's market capitalization at the end of the previous month and measured in USD, as in Fama and French (1992).

LTurnover (turnover), Trading Frictions, Monthly Turnover is a stock's trading volume (VO) divided by its shares outstanding (NOSH) during the last month, as in Datar, Y. Naik and Radcliffe (1998).

NOA (net operating assets), Investment, Yearly Following Hirshleifer et al. (2004), net operating assets are defined as the difference between operating assets and operating liabilities, scaled by lagged total assets. Operating assets are total assets (WC02999) minus cash and short-term investments (WC02001). Operating liabilities are total assets (WC02999), minus total debt (WC03255), minus minority interest (WC03426), minus preferred stock and common equity (WC03995).

OA (operating accruals), Intangibles, Yearly Following Sloan (1996), operating accruals are calculated as changes in working capital minus depreciation (WC01151) scaled by lagged total assets (WC02999). Changes in operating working capital are changes in current assets (WC02201) minus changes in cash and short-term investments (WC02001) minus changes in current liabilities (WC03101), plus changes in debt in current liabilities (WC03051) plus changes in income taxes payable (WC03063).

OL (operating leverage), Intangibles, Yearly We define operating leverage as the sum of costs of goods sold (WC01051) and selling, general, and administrative expenses (WC01101) over total assets (WC02999), as in Novy-Marx (2010).

P2P52WH (price relative to its 52-week high), Trading Frictions, Monthly Rel to high price is the ratio of the unadjusted stock price (UP) at the end of the previous calendar month to the past 52-weeks high, as in George and Hwang (2004).

PCM (price-to-cost margin), Profitability, Yearly As in Gorodnichenko and Weber (2016) and D'Acunto et al. (2018), the price-to-cost margin is net sales (WC01001) minus costs of goods sold (WC01051), divided by net sales (WC01001).

PM (profit margin), Profitability, Yearly As in Soliman (2008), we calculate the profit margin as operating income after depreciation or EBIT (WC18191) over sales (WC01001).

Prof (gross profitability), Profitability, Yearly Profitability is net sales (WC01001) minus costs of goods sold (WC01051) divided by the book value of equity, following Ball et al. (2015). The book value of equity is defined in the construction of BEME.

Q (Tobin's Q), Value, Yearly As in Freyberger, Neuhierl and Weber (2020), we define Tobin's Q as total assets (WC02999) plus the market capitalization as of December t-1 minus cash and short-term investments (WC02001) and minus deferred taxes (WC03263), scaled by total assets (WC02999).

$r_{12\text{-}2}$ (momentum), Past Returns, Monthly Momentum is the cumulative return from month t-12 to t-2 as in Fama and French (1996).

$r_{12\text{-}7}$ (intermediate momentum), Past Returns, Monthly Intermediate momentum is the cumulative return from t-12 to t-7 as in Novy-Marx (2012).

$r_{2\text{-}1}$ (short-term reversal), Past Returns, Monthly Short-term reversal is the lagged one-month return as in Jegadeesh (1990).

$r_{36\text{-}13}$ (long-term reversal), Past Returns, Monthly Long-term reversal is the cumulative return from t-36 to t-13 as in De Bondt and Thaler (1985).

RNA (return on net operating assets), Profitability, Yearly As in Soliman (2008), we calculate the return on net operating assets as the ratio of operating income after depreciation or EBIT (WC18191) to lagged net operating assets. Net operating assets are defined following Hirshleifer et al. (2004) and explained in the construction of NOA.

ROA (return on assets), Profitability, Yearly Following Balakrishnan, Bartov and Faurel (2010), return-on-assets is the ratio of earnings before extraordinary items (WC01551) to lagged total assets (WC02999).

ROE (return on equity), Profitability, Yearly Following Haugen and Baker (1996), return-on-equity are earnings before extraordinary items (WC01551) to lagged book equity. The book value of equity is defined in the construction of BEME.

S2P (sales-to-price), Value, Yearly Following Lewellen (2015), sales-to-price is the ratio of net sales (WC01001) to the market capitalization as of December t-1.

SGA2S (sales and general administrative costs to sales), Intangibles, Yearly As in Freyberger, Neuhierl and Weber (2020), we define SG&A to sales as the ratio of selling, general and administrative expenses (WC01101) to net sales (WC01001).

Illiqu (Amihud (2002) illiquidity), Trading Frictions, Monthly We calculate illiquidity according to Amihud (2002) as the arithmetic mean of the following ratio for the past month: the daily absolute return divided by the product of the end-of-day stock price (UP) and the daily trading volume (VO).

SUV (unexplained volume), Trading Frictions, Monthly Following Garfinkel (2009), standard unexplained volume is the difference between actual volume and predicted volume in the previous month. Predicted volume comes from a regression of daily volume on a constant and the absolute values of positive and negative returns. We use two months of data to estimate the model parameters (data from t-2 and t-1) and estimate the predicted volume using data from the previous month (t-1). I require at least fifteen daily observations in the previous month. Unexplained volume is standardized by the standard deviation of the residuals from the regression.

Methodology

Simple linear regression

The least complex method in our analysis and most widely used in the context of empirical asset pricing is the simple linear regression model estimated via the ordinary least squares (OLS) method. We will use it as a benchmark to compare the more complex machine learning models to it. In the case of the simple linear regression, the conditional expectations $f^*(x)$ can be modeled using the following linear model:

$$f(x_{i,t,c}, \theta) = \theta^T x_{i,t,c}, \tag{B.1}$$

where θ, $\theta^T = (\theta_1, \theta_2, ..., \theta_p) \in \mathbb{R}^p$, is the column vector of coefficients that can be estimated with OLS by minimizing the loss function:

$$L_{MSE}(\theta) = \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} (r^{abn}_{i,t+1,c} - f(x_{i,t,c}, \theta)^2, \tag{B.2}$$

which is also known as the Mean Squared Error (MSE). The OLS has the big advantage that it does not require any hyperparameter input from the user. Further, by minimizing the loss function L_{MSE} a unique analytical solution can be extracted, which is easy to interpret as the coefficients, θ directly describe how a change in the stock characteristics affects the expected return. Additionally, if the number of observations in the underlying dataset is larger than the number of coefficients that need to be estimated, the OLS yields an efficient and unbiased estimator according to Wooldridge (2001). But if the number of characteristics approaches the number of observations in the dataset, the OLS has issues distinguishing between signal and noise. While the signal is the portion we can understand, model, and predict, noise consists of the unpredictable component of price movements. In the case of a small sample or a large number of characteristics, the OLS starts with overfitting the coefficients to noise rather than extracting the signal. This is of particular importance in the field of asset pricing, which empirically relies on a low signal-to-noise ratio. This overfitting yields a higher in-sample performance but a poor out-of-sample performance. Further, multicollinearity between the different characteristics can lead to a fallacious interpretation of test statistics as well as misleading coefficients. Lastly, the OLS does not model or evaluate any non-linearities of the characteristics nor any potential interactions between them. Any non-linearity would have to be imputed by the user.

Regularized regression

To avoid overfitting in the case of empirical asset pricing, the user could increase the training sample, reduce the number of characteristics used to predict future returns, or utilize regularized regression techniques that identify which characteristics are informative and omits those that are not. Classical regularized regression techniques are ridge regression, lasso regression, or elastic net. To limit the number of machine learning methods, we concentrate on the elastic net, which is a combination of ridge and lasso regression. While the different regularized regression models have the same linear functional form as the simple linear regression, they differ with respect to the loss function by adding a penalty term ($\phi_{ENet}(\theta, \lambda, \alpha)$) to it:

$$L_{ENet}(\theta, \lambda, \alpha) = L_{MSE}(\theta) + \phi_{ENet}(\theta, \lambda, \alpha). \tag{B.3}$$

This penalty term reduces the model's in-sample performance and increases its out-of-sample stability by shrinking the coefficients of noisy characteristics, improving the signal-to-noise ratio. The penalty function of the elastic net is defined as:

$$\phi_{ENet}(\theta, \lambda, \alpha) = (1-\alpha)\lambda \sum_{j=1}^{P} \mid \theta_j \mid + \frac{1}{2}\alpha\lambda \sum_{j=1}^{P} \theta_j^2, \tag{B.4}$$

where λ, $\lambda \in \mathbb{R}^+$ defines the magnitude of shrinkage and α, $\alpha \in \{0, ..., 1\}$ which determines the relative weight between the two penalty components of the ridge and lasso regression. In the case of $\lambda = 0$ the regularized regression models yield a simple linear regression model. The coefficients are shrunk towards zero by setting $\lambda > 0$. As these two hyperparameters have to be set by the user, we utilize our validation sample to find the optimal in-sample λ and α in the first run. We determine the optimal θ in the second run using the full training and validation sample.

Tree-based regression

Tree-based models represent the first non-parametric regression model as their structure is decided by the training data. For our return prediction, we will utilize two tree-based methods: the random forest, as well as the gradient boosted regression tree. Compared to the linear methods, one advantage of these tree methods is that the user does not have to manually add any potential non-linearities or interactions to the data as the tree methods build these by construction.

Regression trees follow the idea of sequentially partitioning the underlying data into groups that behave similarly to each other based on a selected characteristic with regard to the future return. By sequentially separating the data, the tree ”grows” and new ”branches” are created each time the data is split into new groups. The tree can grow to a depth of D based on the user input. At each new branch, the characteristic is picked that causes the biggest separation in the data based on an optimized cut-off value.[1] As soon as the data can not be split into subgroups or the depth D is reached, a ”leave” is created. In asset pricing, the tree yields a return that is clustered by the underlying characteristics.

The following equation describes a tree with a depth of D and K leaves:

$$\begin{aligned} f(x_{i,t,c}, \theta, D, K) &= \sum_{k=1}^{K} \theta_k 1_{\{x_{i,t,c} \in C_k(D)\}} \\ \theta_k &= \frac{1}{N_k} \sum_{x_{i,t,c} \in C_k(D)} r^{abn}_{i,t+1,c}, \end{aligned} \tag{B.5}$$

where D is the depth of the tree measured as the maximum number of separations following the longest branch, $C_k(D)$ indicates the k-th separation of the characteristics, θ_k is average abnormal return within the partition, and $1_{\{x_{i,t,c} \in C_k(D)\}}$ indicates if $x_{i,t,c}$ is part of $C_k(D)$. Following this methodology, a tree of depth D can capture up to $D-1$ interactions. To avoid overfitting, the tree must be regularized. We follow two different approaches in our analysis.

The first regularization approach uses bootstrap aggregation, or “bagging,” developed

[1] In our case, for each separation, the characteristic is selected that minimizes the MSE.

by Breiman (2001). In this approach, each of the T trees starts with a share of B bootstrap samples from the data and fits an individual regression tree to the bootstrapped data. Afterward, the forecasts from the individual trees are averaged. This reduces the variation in the prediction and stabilizes the prediction performance. In the case of the random forest, the trees additionally use random subsets R of characteristics to grow the branches. This reduces the impact of certain dominant return characteristics and creates de-correlated trees.

The second regularization approach is ”boosting.” It starts by training a weak and shallow regression tree on the full training data. In the next step, a second regression tree with the same depth D is trained on the residuals of the first tree. The prediction of these two trees is then averaged while the contribution of the second tree is shrunken by a factor LR (learning rate), $LR \in (0, 1)$ to avoid the model overfitting the residuals. At each new step b, till the model reaches a total of B trees, a new shallow tree is fitted to the residual, which is based on the $b-1$-th model and added to it with a shrinkage weight of LR.

Both regression trees share the two main hyperparameters: the number of trees in the forest T, $T \in \mathbb{Z}^+$ and the maximum depth D, $D \in \mathbb{Z}^+$. While the random forest additionally requires the share of the bootstrapped samples B, $0 > B \leq 1$, the gradient boosted regression tree requires a certain learning rate LR, $0 > B \leq 1$. These hyperparameters are optimized through the validation step. Additionally, we can provide the share R, $0 > R \leq 1$, of randomly selected characteristics that are used in each tree of the random forest.

Neural networks

Neural networks are another highly flexible but opposed to the regression trees, a parametric model. While these models can have various forms, we focus on the standard structure of a feed-forward neural network. A feed-forward neural network consists of an ”input” layer of input characteristics and the intercept, at least one ”hidden” layer compromising activation functions, and an ”output” layer that aggregates the outcome of the last hidden layer into a return prediction.

A feedforward neural network consists of several subsequent layers l, $l = 0, 1, ..., L$, one input layer ($l = 0$), $L - 1$ hidden layers ($l = 1, 2, ..., L - 1$) and one output layer $l = L$.

Each layer l contains n^l nodes. In the case of the input layer, the number of nodes is equal to the number of characteristics, including an intercept, while the output layer contains due to the regression setting one node. In the case of the hidden layer, we consider an architecture of up to five hidden layers while the first hidden layer contains 32 nodes and each additional hidden layer divides the number of nodes by two compared to the previous layer following the geometric pyramid rule according to Masters (1993). This results in the following number of nodes per layer:

$$\begin{aligned} n^0 &= p + 1, \\ n^1 &= 32, \\ n^l &= \frac{n^{l-1}}{2} \forall l \in \{2, ..., L-1\}, \\ n^L &= 1. \end{aligned} \tag{B.6}$$

Each of the nodes in the hidden layer contains an activation function. In our case we follow Gu, Kelly and Xiu (2020) and Leippold, Wang and Zhou (2022) and choose the rectified linear unit defined as:

$$\text{ReLU}(x) = \max(0, x), \tag{B.7}$$

As in De Nard, Hediger and Leippold (2022), we adopt the Adam optimization algorithm (Kingma and Ba, 2014), early stopping, batch normalization (Ioffe and Szegedy, 2015), ten ensembles with individual seeds (Hansen and Salamon, 1990; Dietterich, 2000) and dropout (Srivastava et al., 2014) when training our models.

Hyperparameters

We will use the following hyperparameters based on the hyperparameter range in Gu, Kelly and Xiu (2020), Tobek and Hronec (2020), Drobetz and Otto (2021), and Leippold, Wang and Zhou (2022):

- Elastic net
  - λ: $[1 \times 10^{-5}, 2 \times 10^{-5}, ..., 1 \times 10^{-2}]$
  - α: $[0, 0.01, ..., 1]$

- Random forest
  - R: $[0.01, 0.02, ..., 1]$
  - B: 1
  - T: $[100, 102, ..., 600]$
  - D: $[1, 2, ..., 8]$
- Gradient-boosted regression tree
  - LR: $[0.01, 0.02, ..., 0.1]$
  - T: $[50, 52, ..., 500]$
  - D: $[1, 2, ..., 8]$
- Neural networks
  - l_1: $[0.00001, ..., 0.001]$
  - LR: $[0.001, 0.1]$
  - Batch Size: 10000
  - Epochs: 100

Factor construction

This section outlines the construction of the factors of the Fama and French (2018) six-factor model, using the same stock sample as for the machine learning portfolios described in Section 2.2.1.

The market factor, RMRF, is the value-weighted return of all stocks minus the risk-free rate. The remaining factors are constructed using a 2 x 3 sorting approach commonly employed for international markets (Fama and French, 2012, 2017). The portfolio breakpoints for the value, profitability, investment, and momentum factors are the 30% and 70% percentiles of the underlying characteristic of the big-stock sample per country. In the case of the value factor, we use the book-to-market ratio to form Growth (G), Neutral (N), and Value (V) portfolios. For profitability, we use cash-based operating profitability to sort the stocks into the extreme portfolios Weak (W) and Robust (R). We use asset growth for the investment factor, which yields Conservative (C) and Aggressive (A) portfolios. For the momentum factor, we sort stocks into the Winner (W) and Loser (L) portfolios based on a stock's momentum. Finally, we classify stocks into the two size groups big (B) and small (S) as described in Section 2.2.1. The final factor calculation is based on the intersection of the different portfolios, while the portfolio returns are value-weighted,

$$\begin{aligned} SMB &= (SV + SN + SG)/3 - (BV + BN + BG)/3, \\ HML &= (BV + SV)/2 - (BG + SG)/2, \\ RMW &= (BR + SR)/2 - (BW + SW)/2, \\ CMA &= (BC + SC)/2 - (BA + SA)/2, \\ WML &= (BW + SW)/2 - (BL + SL)/2. \end{aligned} \tag{B.8}$$

Table B. 6 presents summary statistics for the monthly factor returns.

Table B. 6
Summary statistics for the factors of the Fama and French (2018) six-factor model

This table presents the average monthly return, standard deviation, and corresponding t-statistic for the following set of factors: the market ($RMRF$), size (SMB, small minus big), value (HML, high minus low), profitability (RMW, robust minus weak based on cash-based operating profitability), investment (CMA, conservative minus aggressive), and momentum (WML, winners minus losers). The t-statistics are Newey-West adjusted with 4 lags. The sample period for the analysis is January 2002 to December 2021.

	$RMRF$	SMB	HML	CMA	RMW	WML
	(1)	(2)	(3)	(4)	(5)	(6)
mean	0.94	-0.0	0.52	0.07	0.49	0.79
std. dev.	5.88	1.36	1.89	2.24	2.65	4.16
t-stat	2.05	-0.05	3.40	0.38	2.45	2.54

Figures

Figure B. 1
Variable importance by model

This figure shows the importance for Individual characteristics in each model. Characteristics importance is an average over all training samples. Variable importance within each model is normalized to sum to one.

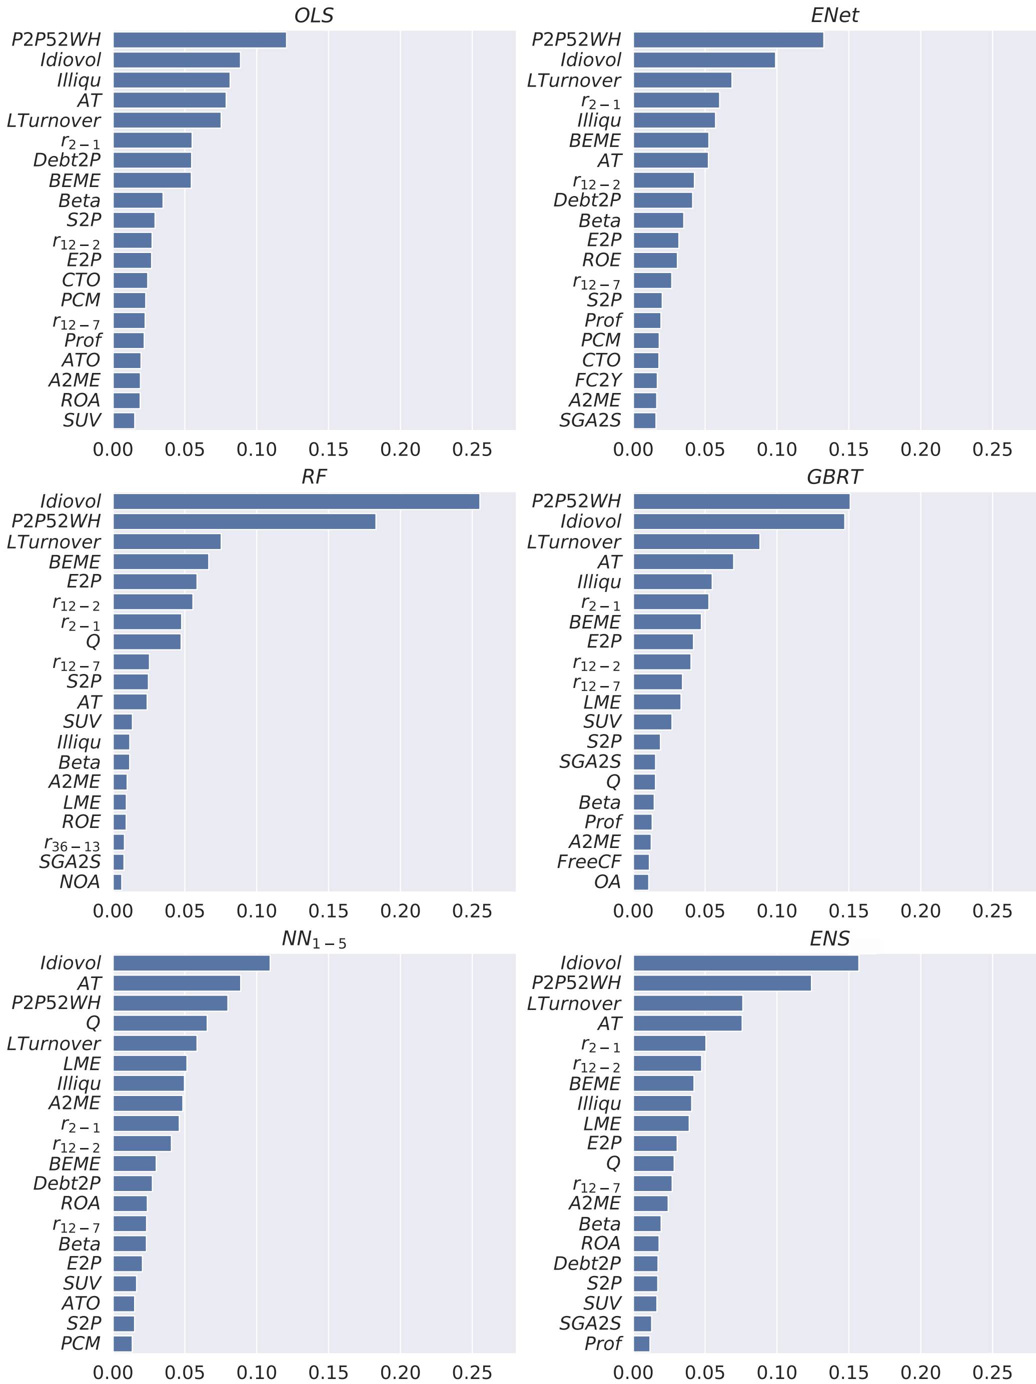

Figure B. 2
Cumulative return of machine learning portfolios

The figure shows the cumulative log returns in excess of the market of portfolios sorted on out-of-sample machine learning return forecasts. The solid and dashed lines represent long (top quintile) and short (bottom quintile) portfolios, respectively. In Panel A equal-weighted cumulative log returns are shown while in Panel B the long and short portfolios are value-weighted. The sample period is from January 2002 to December 2021.

Panel A: Equal-weighted

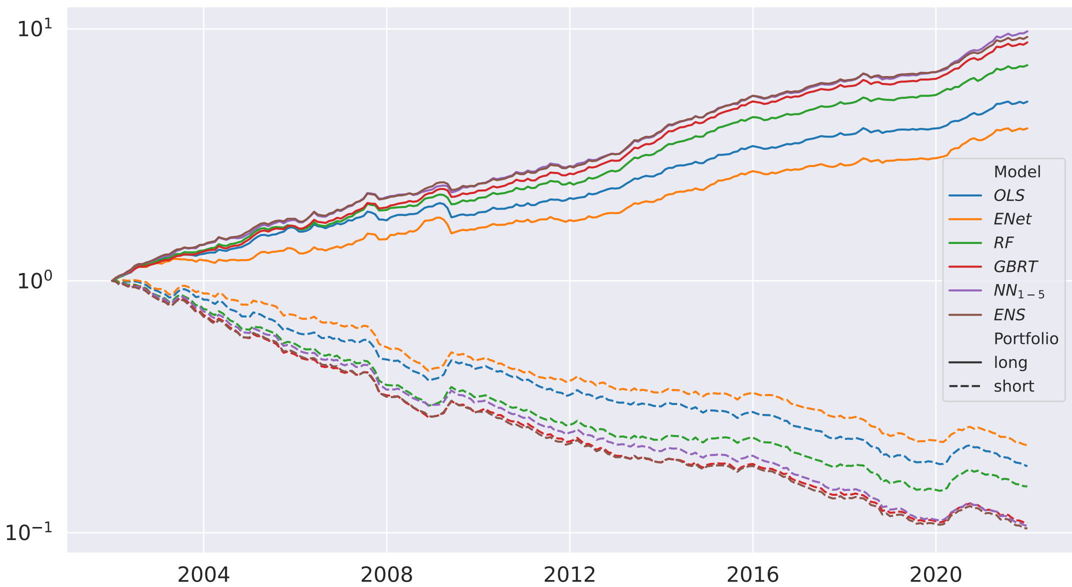

Panel B: Value-weighted

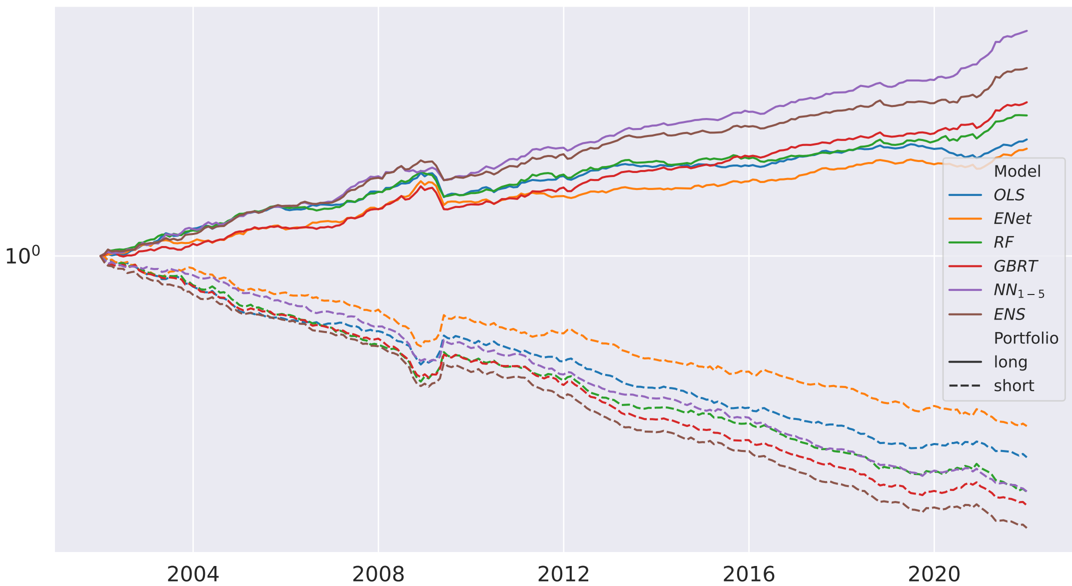

Figure B. 3

Estimated bid-ask spreads based on the EDGE estimator

This figure shows the cross-sectional distribution of estimated bid-ask spreads for big stocks in emerging markets. Thereby, big stocks are defined as the biggest stocks, which together account for 90% of a country's aggregated market capitalization. For each stock and month, we compute the efficient discrete generalized estimator (EDGE) of the bid-ask spread, proposed in Ardia, Guidotti and Kroencke (2022). The estimators are based on daily prices using a monthly estimation window. Following Novy-Marx and Velikov (2016), we replace zero estimates with the non-zero estimate of the stock of the same country with the shortest Euclidean distance in size and characteristic volatility rank space. The sample period is from January 2002 to December 2021.

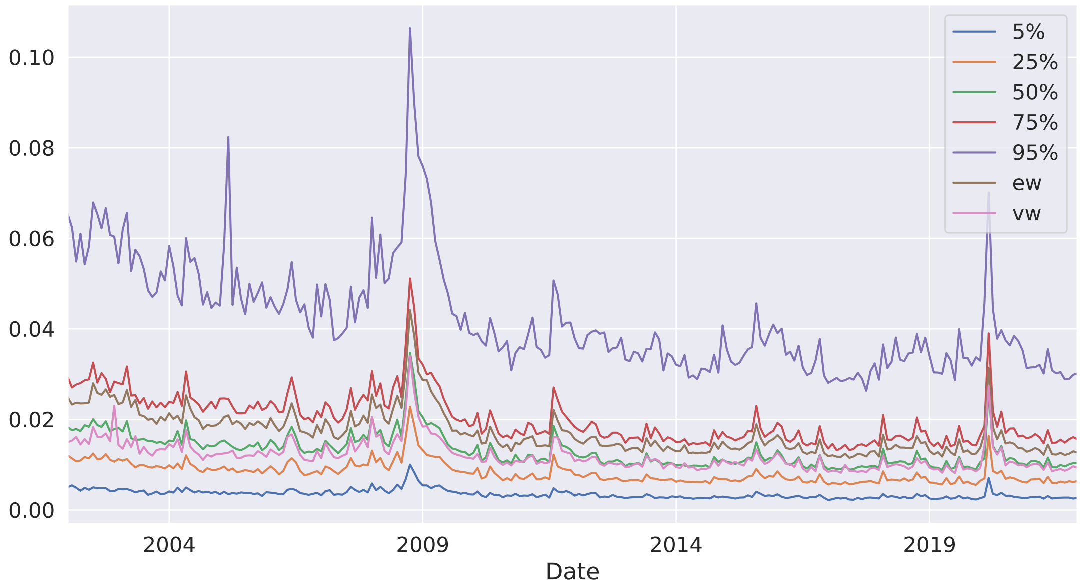

Tables

Table B. 7
Detail performance of the machine learning portfolios

This table reports the out-of-sample performance of the different machine learning quintile portfolios. Stocks are sorted into country-neutral quintiles based on their predicted returns for the next month. The sorting breakpoints are based on big stocks only, which are in the top 90% of a country's aggregated market capitalization. Each Panel provides the predicted monthly returns (Pred), the average monthly excess returns (Avg), corresponding t-statistics (t), the Fama and French (2018) six-factor model alpha (α), and corresponding t-statistics. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	Equal-weighted					Value-weighted				
	Pred	Avg	t	α	t_α	Pred	Avg	t	α	t_α
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: *OLS*										
Low (L)	-1.07	0.26	0.50	-0.39	-3.43	-0.99	0.43	0.83	-0.23	-3.40
2	-0.38	0.90	1.85	0.11	1.16	-0.37	0.82	1.70	-0.00	-0.05
3	0.00	1.15	2.51	0.26	3.08	0.01	0.95	2.09	0.09	2.55
4	0.36	1.34	3.03	0.41	4.91	0.37	1.13	2.55	0.17	3.45
High (H)	0.86	1.64	3.73	0.58	6.39	0.87	1.25	2.82	0.06	1.05
H-L	1.93	1.38	7.76	0.97	8.02	1.85	0.83	4.57	0.28	2.72
Panel B: *ENet*										
Low (L)	-1.07	0.34	0.65	-0.33	-3.05	-0.99	0.51	0.98	-0.18	-2.33
2	-0.37	0.92	1.86	0.12	1.32	-0.36	0.81	1.70	0.02	0.32
3	0.02	1.16	2.48	0.27	3.19	0.03	0.96	2.06	0.09	1.95
4	0.39	1.32	2.97	0.38	4.52	0.40	1.09	2.43	0.11	2.21
High (H)	0.90	1.53	3.61	0.50	4.99	0.91	1.23	2.86	0.09	1.56
H-L	1.97	1.20	6.79	0.83	6.94	1.90	0.72	3.95	0.27	2.28
Panel C: *RF*										
Low (L)	-0.79	0.18	0.35	-0.47	-4.05	-0.69	0.34	0.66	-0.30	-4.37
2	-0.21	0.82	1.74	0.07	0.81	-0.21	0.86	1.82	0.09	1.77
3	0.09	1.12	2.41	0.23	2.62	0.09	0.94	2.06	0.03	0.61
4	0.37	1.35	2.92	0.40	5.16	0.37	1.15	2.51	0.09	1.45
High (H)	0.71	1.78	3.96	0.72	8.34	0.70	1.32	2.99	0.17	3.38
H-L	1.50	1.60	9.29	1.19	14.10	1.39	0.99	5.24	0.47	5.24
Panel D: *GBRT*										
Low (L)	-0.87	0.04	0.08	-0.59	-5.24	-0.74	0.30	0.58	-0.38	-5.57
2	-0.18	0.87	1.82	0.11	1.21	-0.17	0.81	1.71	0.09	1.68
3	0.13	1.13	2.42	0.25	3.25	0.13	0.94	2.04	0.03	0.73
4	0.43	1.36	2.98	0.39	4.76	0.42	1.12	2.47	0.11	1.62
High (H)	0.93	1.86	4.12	0.81	8.77	0.86	1.35	3.09	0.19	4.06
H-L	1.80	1.82	11.49	1.40	15.65	1.61	1.05	6.06	0.57	6.73

Continued on next page

Table B. 6 continued

	Equal-weighted					Value-weighted				
	Pred	Avg	t	α	t_α	Pred	Avg	t	α	t_α
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel E: NN_1										
Low (L)	-1.27	0.03	0.06	-0.62	-5.23	-1.05	0.41	0.81	-0.26	-2.74
2	-0.28	0.80	1.71	0.04	0.54	-0.26	0.76	1.68	-0.01	-0.28
3	0.16	1.07	2.31	0.21	2.60	0.16	0.94	2.08	0.05	1.15
4	0.58	1.34	2.93	0.40	4.92	0.57	1.07	2.35	0.05	0.78
High (H)	1.34	1.91	4.09	0.86	8.93	1.25	1.45	3.09	0.30	5.86
H-L	2.61	1.88	13.92	1.47	15.67	2.30	1.04	6.94	0.57	4.83
Panel F: NN_2										
Low (L)	-1.27	-0.01	-0.03	-0.68	-6.33	-1.01	0.37	0.74	-0.35	-5.42
2	-0.23	0.82	1.74	0.06	0.70	-0.21	0.80	1.72	0.03	0.52
3	0.17	1.01	2.18	0.16	2.10	0.18	0.93	2.08	0.03	0.65
4	0.56	1.35	3.06	0.42	4.76	0.55	1.08	2.42	0.06	0.87
High (H)	1.32	1.90	3.99	0.86	8.87	1.21	1.49	3.08	0.36	6.84
H-L	2.60	1.91	15.67	1.55	19.02	2.21	1.11	9.40	0.71	9.16
Panel G: NN_3										
Low (L)	-1.20	0.02	0.05	-0.66	-5.87	-0.94	0.38	0.74	-0.35	-4.59
2	-0.18	0.82	1.75	0.06	0.74	-0.16	0.73	1.60	-0.01	-0.34
3	0.17	1.08	2.32	0.24	3.47	0.17	0.97	2.17	0.09	2.29
4	0.51	1.31	2.92	0.39	4.42	0.50	1.08	2.36	0.06	1.13
High (H)	1.22	1.86	3.99	0.83	8.34	1.11	1.49	3.17	0.32	5.34
H-L	2.41	1.84	14.72	1.49	16.93	2.04	1.12	7.85	0.66	6.55
Panel H: NN_4										
Low (L)	-1.14	0.04	0.07	-0.63	-5.57	-0.90	0.32	0.62	-0.38	-5.18
2	-0.17	0.79	1.67	0.03	0.34	-0.14	0.73	1.59	-0.02	-0.50
3	0.16	1.09	2.37	0.26	3.34	0.17	0.95	2.14	0.08	1.87
4	0.48	1.31	2.91	0.35	4.39	0.47	1.11	2.43	0.05	0.81
High (H)	1.15	1.89	4.07	0.86	8.26	1.04	1.52	3.25	0.35	6.35
H-L	2.29	1.86	13.75	1.48	16.72	1.94	1.20	8.30	0.73	7.76
Panel I: NN_5										
Low (L)	-1.12	0.04	0.08	-0.62	-5.43	-0.90	0.36	0.68	-0.37	-4.83
2	-0.17	0.82	1.74	0.04	0.46	-0.15	0.71	1.56	0.01	0.29
3	0.18	1.10	2.37	0.27	3.68	0.18	0.91	1.97	0.02	0.50
4	0.50	1.32	2.96	0.38	4.59	0.49	1.13	2.54	0.08	1.47
High (H)	1.14	1.88	4.04	0.82	8.38	1.04	1.52	3.26	0.34	7.61
H-L	2.26	1.84	13.42	1.44	15.79	1.93	1.17	8.11	0.71	8.21

Continued on next page

Table B. 6 continued

	Equal-weighted					Value-weighted				
	Pred	Avg	t	α	t_α	Pred	Avg	t	α	t_α
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel J: NN_{1-5}										
Low (L)	-1.13	0.03	0.06	-0.62	-5.41	-0.91	0.34	0.65	-0.36	-4.82
2	-0.19	0.77	1.65	-0.00	-0.00	-0.16	0.73	1.62	0.02	0.42
3	0.17	1.10	2.38	0.27	3.57	0.17	0.95	2.10	0.06	1.35
4	0.51	1.31	2.92	0.39	4.50	0.50	1.07	2.39	0.03	0.54
High (H)	1.17	1.90	4.05	0.84	8.70	1.07	1.54	3.22	0.36	7.02
H-L	2.30	1.87	13.42	1.46	15.81	1.97	1.21	8.48	0.72	8.26
Panel K: ENS										
Low (L)	-0.85	0.02	0.04	-0.61	-5.32	-0.71	0.24	0.47	-0.42	-5.89
2	-0.17	0.85	1.78	0.09	0.99	-0.16	0.80	1.71	0.09	1.85
3	0.13	1.13	2.46	0.29	3.44	0.13	0.90	1.98	0.01	0.29
4	0.41	1.38	3.02	0.40	5.17	0.41	1.15	2.59	0.12	1.95
High (H)	0.87	1.88	4.12	0.82	9.02	0.81	1.45	3.17	0.25	5.82
H-L	1.71	1.86	11.73	1.43	15.66	1.52	1.20	6.97	0.67	8.29
Panel L: $\mu_{sign(c)}$										
Low (L)	-0.20	0.22	0.42	-0.41	-3.16	-0.19	0.48	0.93	-0.22	-3.42
2	-0.08	0.79	1.57	0.00	0.02	-0.08	0.90	1.84	0.10	2.10
3	-0.01	1.05	2.20	0.18	2.19	-0.01	0.96	2.10	0.08	1.75
4	0.06	1.25	2.77	0.32	3.72	0.06	1.08	2.45	0.05	1.06
High (H)	0.15	1.56	3.67	0.57	7.27	0.15	1.24	2.91	0.11	2.18
H-L	0.35	1.34	6.48	0.98	7.26	0.34	0.77	4.16	0.32	3.32

Table B. 8

Robustness - Models trained on subregional data

This table reports the out-of-sample performance of equal-weighted and value-weighted long-short portfolios sorted on forecasts derived from models trained on subregional data. All stocks are sorted into country-neutral portfolios based on their predicted returns for the next month. The sorting breakpoints are based on big stocks only, which are in the top 90% of the country's aggregated market capitalization. Panel A shows the results pooled emerging markets, Panel B for all countries being part of emerging Americas, Panel C combines all emerging Asian countries, and Panel D reports results for emerging countries from Europe, the Middle East, and Africa. The first two rows of each panel provide the average monthly return of the long-short quintile (Avg), corresponding t-statistics (t), the average Fama and French (2018) six-factor alpha (α), corresponding t-statistics (t_α), and R^2. The next two rows show spanning alpha (α), corresponding t-statistic (t_α), and R^2 when regressing the long-short ENS returns on OLS returns and vice versa. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	Equal-weighted					Value-weighted				
	Avg	t	α	t_α	R^2	Avg	t	α	t_α	R^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Emerging Markets										
OLS	1.09	6.29	0.77	6.23	53.92	0.78	4.42	0.29	2.57	56.92
ENS	1.35	8.88	0.97	9.75	57.22	0.97	5.95	0.44	4.59	58.10
$ENS \sim OLS$			0.49	7.78	77.56			0.28	4.46	75.77
$OLS \sim ENS$			-0.23	-1.76	77.56			-0.06	-0.51	75.77
Panel B: Americas										
OLS	0.76	3.23	0.36	1.82	46.02	0.60	2.67	0.04	0.22	49.64
ENS	0.70	3.37	0.41	2.52	36.93	0.64	3.12	0.20	1.14	35.55
$ENS \sim OLS$			0.19	1.56	56.53			0.22	1.66	49.56
$OLS \sim ENS$			0.17	1.02	56.53			0.14	0.83	49.56
Panel C: Asia										
OLS	1.43	7.73	1.12	9.62	60.08	0.81	4.03	0.41	3.22	63.14
ENS	1.95	11.16	1.61	17.97	60.73	1.27	6.62	0.83	8.24	67.10
$ENS \sim OLS$			0.65	8.64	83.54			0.52	5.07	71.55
$OLS \sim ENS$			-0.36	-2.10	83.54			-0.17	-1.11	71.55
Panel D: Europe, the Middle East and Africa										
OLS	1.06	5.97	0.91	5.67	19.35	0.92	4.43	0.47	2.39	26.25
ENS	1.39	8.11	1.06	6.59	20.58	0.99	4.75	0.34	1.97	37.32
$ENS \sim OLS$			0.61	5.47	56.62			0.25	2.01	59.45
$OLS \sim ENS$			-0.01	-0.06	56.62			0.19	1.45	59.45

Table B. 9

Robustness - Models trained on pooled versus individual countries

This table reports the out-of-sample performance of value-weighted long-short portfolios sorted on local and pooled model forecasts. Local model forecasts are based on machine learning models trained separately for each country on local country data. In contrast, the pooled model forecasts are from our baseline machine learning models trained on pooled emerging market data. Both local and pooled strategies include only stocks from the following seven countries that are in our sample throughout the entire sample period: Chile, Indonesia, Mexico, Malaysia, Philippines, Thailand, and Turkey. All stocks are sorted into country-neutral quintile portfolios based on predicted returns for the next month. The sorting breakpoints are based on big stocks only. The first two rows of each panel provide the average monthly excess returns (Avg), corresponding t-statistics (t), the average Fama and French (2018) six-factor alphas (α), corresponding t-statistics (t_α), and R^2. The next two rows show spanning alphas (α), corresponding t-statistics (t_α), and R^2 when regressing the long-short *local* returns on *pooled* returns and vice versa. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	OLS					*ENS*				
	Avg	t	α	t_α	R^2	Avg	t	α	t_α	R^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
local	0.61	3.68	0.41	2.83	17.13	0.92	5.93	0.63	4.70	23.16
pooled	0.80	4.37	0.41	2.26	34.41	1.18	6.67	0.81	5.66	27.08
local $\sim$ *pooled*			0.17	1.42	38.67			0.25	1.83	40.26
pooled $\sim$ *local*			0.37	2.15	38.67			0.53	3.15	40.26

Table B. 10

Performance of machine learning portfolios for developed markets

This table reports the out-of-sample performance of long-short portfolios in developed markets. Stocks are sorted into country-neutral quintile portfolios based on the predicted returns from machine learning models trained with developed market data. The sorting breakpoints are based on big stocks only. The first two rows of each panel provide the average monthly returns of the long-short quintiles (Avg), corresponding t-statistics (t), the average Fama and French (2018) six-factor alphas (α), corresponding t-statistics (t_α), and R^2. The next two rows show spanning alphas (α), corresponding t-statistics (t_α), and R^2 when regressing the long-short *ENS* returns on *OLS* returns and vice versa. All t-statistics are calculated using Newey and West (1987) adjusted standard errors with 4 lags. The sample period is from January 2002 to December 2021.

	Equal-Weighted					Value-Weighted				
	Avg	t	α	t_α	R^2	Avg	t	α	t_α	R^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	0.68	2.83	0.42	2.32	55.74	0.43	1.85	0.12	0.74	53.00
ENS	0.93	4.76	0.65	4.15	51.69	0.66	3.22	0.32	2.04	48.20
ENS $\sim$ *OLS*			0.42	5.82	89.49			0.31	4.01	84.55
OLS $\sim$ *ENS*			-0.43	-5.45	89.49			-0.26	-3.21	84.55

Table B. 11
Limits to arbitrage: Summary statistics

This table reports the summary statistics of limits-to-arbitrage proxies of different machine learning quintile portfolios. All stocks are sorted into country-neutral quintile portfolios based on their predicted returns for the next month. The sorting breakpoints are based on big stocks only, which are in the top 90% of the country's aggregated market capitalization. We compute for each quintile portfolio the average monthly value of each of three proxies for limits to arbitrage: -1 $\times$ market capitalization ($SIZE$), idiosyncratic volatility ($IVOL$), Amihud illiquidity ($ILLIQ$), and a combination of the different proxies ($COMBO$). All proxies for limits to arbitrage are ranked into the [-1,1] interval for each month and country and higher values indicate higher limits to arbitrage. Afterward, we report the time-series average. The sample period is from January 2002 to December 2021.

	$SIZE$		$IVOL$		$ILLIQ$		$COMBO$	
	OLS	ENS	OLS	ENS	OLS	ENS	OLS	ENS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Low (L)	0.10	-0.11	0.03	0.02	0.33	0.33	0.15	0.15
2	0.01	0.06	-0.02	-0.03	0.11	-0.06	0.03	-0.05
3	-0.03	0.10	-0.03	-0.05	-0.04	-0.14	-0.03	-0.10
4	-0.06	0.09	-0.03	-0.03	-0.17	-0.18	-0.09	-0.10
High (H)	-0.05	-0.04	0.03	0.04	-0.36	-0.13	-0.13	-0.02

Table B. 12
Further investment frictions

This table reports the performance of different buy/hold long-only strategies before and after transaction costs. The investment universe is limited to big stocks. We investigate predictions from a linear OLS model and an ensemble (ENS) of non-linear machine learning models (RF, $GBRT$, and NN_{1-5}). Every month the portfolio consists of the stocks that currently exhibit the highest X% forecasted returns per country plus those selected in previous months whose forecasted returns have not deteriorated beyond the top Y%. The first number in the column row names represents X, while the second represents Y. We report the strategies' gross returns in excess of the market, average two-way turnover, transaction costs, net returns in excess of the market, and net Fama and French (2018) six-factor models alphas. We assume one-way transaction costs of 100 basis points. All t-statistics are Newey and West (1987) adjusted with 4 lags. Panel A summarizes results from equal-weighting while Panel B shows results from value-weighting. The sample period is from January 2002 to December 2021.

	OLS		*ENS*	
	20%/20%	10%/30%	20%/20%	10%/30%
	(1)	(2)	(3)	(4)
Panel A: Equal-weighted				
$r^e_{gross} - Mkt$	0.47	0.44	0.78	0.78
	(5.31)	(5.07)	(7.88)	(7.37)
TO (in %)	44.14	24.68	45.13	27.38
T-cost (in %)	0.44	0.25	0.45	0.27
$r^e_{net} - Mkt$	0.03	0.19	0.33	0.51
	(0.37)	(2.26)	(3.35)	(4.82)
α^{FF6}_{net}	0.17	0.33	0.43	0.62
	(2.77)	(5.17)	(5.84)	(7.35)
Panel B: Value-weighted				
$r^e_{gross} - Mkt$	0.27	0.29	0.45	0.46
	(2.99)	(3.13)	(4.67)	(4.53)
TO (in %)	44.09	21.86	45.46	23.28
T-cost (in %)	0.44	0.22	0.45	0.23
$r^e_{net} - Mkt$	-0.17	0.07	-0.01	0.23
	(-1.88)	(0.74)	(-0.08)	(2.26)
α^{FF6}_{net}	-0.16	0.06	0.02	0.25
	(-3.27)	(1.14)	(0.36)	(3.72)

C Chapter 3

Decomposition methodology of the 3×3

In the case of a reduced number of stocks, we sort all stocks which experienced a firm-specific news arrival based on their nearness to the 52-week high and their firm-specific news return into two independent country-neutral tercile portfolios. Similar to the main decomposition methodology, we utilize two different Fama and MacBeth (1973) regressions to decompose the returns of the double-sorted portfolios into the two pure effects of firm-specific news and the nearness to the 52-week high as well as the interaction effect of both across the nine portfolios. The regression model, including the interactions, is specified as follows:

$$\begin{aligned} R_{i,t+1} = b_0 &+ b_1 FN^3_{i,t} + b_2 FN^1_{i,t} + b_3 NEAR^3_{i,t} + b_4 NEAR^1_{i,t-1} \\ &+ b_5 FN^3_{i,t} \times NEAR^3_{i,t} + b_6 FN^1_{i,t} \times NEAR^3_{i,t} \\ &+ b_7 FN^3_{i,t} \times NEAR^1_{i,t} + b_8 FN^1_{i,t} \times NEAR^1_{i,t} + \epsilon, \end{aligned} \tag{C.1}$$

where $R_{i,t+1}$ is the stock return of firm i in the next month $t+1$, and right-hand-side variables are dummies indicating the tercile ranking of firm i at the end of the month t for FN and $NEAR$. In the second regression, we exclude the interaction effect from the model:

$$R_{i,t+1} = b_0 + b_1 FN^3_{i,t} + b_2 FN^1_{i,t} + b_3 NEAR^3_{i,t} + b_4 NEAR^1_{i,t-1} + \epsilon \tag{C.2}$$

In Table C. 1, we describe how the individual average portfolio return in each of the 3×3 portfolios sorted by the firm-specific news return and the nearness to the 52-week high is decomposed by using the regression parameters and the return components. The lowest nearness to the 52-week high (firm-specific news return) tercile is defined as $NEAR1$ ($FN1$), while the highest nearness to the 52-week high (firm-specific news return) tercile is specified as $NEAR3$ ($FN3$).

Table C. 1
Specification of return decomposition 3×3

This table describes the specification of the return decomposition as in George, Hwang and Li (2014) and Huang, Lin and Xiang (2021) by regression parameter and return component for the double-sorted firm portfolios by firm-specific news returns (FN) and nearness to the 52-week high ($NEAR$). To form the double-sorting portfolios we sort each month all firms which experienced firm-specific news arrival into independent and country-neutral 3×3 portfolios based on FN in the previous month and $NEAR$ at the previous month-end. Each cell represents a group of stocks with a particular $NEAR$ and FN ranking. In Panel A (Panel B), we show how the respective portfolio return can be decomposed using the regression parameters from the monthly stock-level Fama and MacBeth (1973) regression as specified in Equation C.1 (Equation C.2). In Panel C (Panel D), we show how the respective portfolio return can be decomposed into different return components. The return components can be disentangled into the benchmark return (μ), the returns associated with the 52-week high (H), the returns attributable to the firm-specific news (N), and the returns associated with the interaction between the firm-specific news and nearness of the stock price to the 52-week high (I). μ reflects the average return of stocks in the portfolio with neither extreme firm-specific news returns nor an extreme nearness to the 52-week high. H reflects the returns associated with being near (n), middle (m), or far (f) from the 52-week high, regardless of the FN ranking. N reflects the returns associated with having good (g) or bad (b) firm-specific news about the firms, regardless of the $NEAR$ ranking. I reflects the returns associated with having both good (bad) firm-specific news about the firm and stock prices near (far from) the 52-week high.

	$FN1$	$FN2$	$FN3$
	(1)	(2)	(3)
Panel A: Decomposition by regression parameter *including* the interactions effect			
$NEAR1$	$b_0 + b_2 + b_4$	$b_0 + b_4$	$b_0 + b_1 + b_4 + b_7$
$NEAR2$	$b_0 + b_2$	b_0	$b_0 + b_1$
$NEAR3$	$b_0 + b_2 + b_3 + b_6$	$b_0 + b_3$	$b_0 + b_1 + b_3 + b_5$
Panel B: Decomposition by regression parameter *excluding* the interactions effect			
$NEAR1$	$b_0 + b_2 + b_4$	$b_0 + b_4$	$b_0 + b_1 + b_4$
$NEAR2$	$b_0 + b_2$	b_0	$b_0 + b_1$
$NEAR3$	$b_0 + b_2 + b_3$	$b_0 + b_3$	$b_0 + b_1 + b_3$
Panel C: Decomposition by return component *including* the interactions effect			
$NEAR1$	$\mu + H_f + N_b + I_{b,f}$	$\mu + H_f$	$\mu + H_f + N_g$
$NEAR2$	$\mu + N_b + I_{b,m}$	μ	$\mu + N_g + I_{g,m}$
$NEAR3$	$\mu + H_n + N_b$	$\mu + H_n$	$\mu + H_n + N_g + I_{g,n}$
Panel D: Decomposition by return component *excluding* the interactions effect			
$NEAR1$	$\mu + H_f + N_b$	$\mu + H_f$	$\mu + H_f + N_g$
$NEAR2$	$\mu + N_b$	μ	$\mu + N_g$
$NEAR3$	$\mu + H_n + N_b$	$\mu + H_n$	$\mu + H_n + N_g$

In Panel A and Panel B of Table C. 1, we present how the different estimated parameters of Equation C.1 and Equation C.2 can be combined to derive the respective average portfolio return in each of the portfolios. In Panel C and Panel D, we further show how the respective portfolio return can be decomposed into four different return components. The return components are the benchmark return (μ), the returns associated with the 52-week high (H), the returns attributable to the firm-specific news (N), and the returns associated with the interaction between the firm-specific news and nearness of the stock price to the 52-week high (I). The first return component reflects the benchmark portfolio. It is the average return of the stocks in the portfolio with neither extreme firm-specific news returns nor an extreme nearness to the 52-week high. The second return component is solely driven by the stocks nearness to the 52-week high, regardless of the firm-specific news return ranking. Sorting the stocks into terciles based on their nearness to the 52-week high results in a return component which is common among the stocks in the same portfolio. Stocks that are far (f) away from the 52-week high are denoted as H_f and are expected to have a negative return, while stocks that are near (n) the 52-week high are denoted as H_n and are expected to have a positive return. To derive the pure 52-week high effect, we build a long-short strategy that relies solely on the return predictability of the nearness to the 52-week high. We, therefore, define the pure 52-week high effect as:

$$\text{Pure 52-week High Effect} = H_n - H_f = b_3 - b_4. \tag{C.3}$$

The third return component is solely driven by the firm-specific news return, regardless of the firm-specific news return ranking. Sorting the stocks into terciles based on their firm-specific news return results in a return component that is common among the stocks in the same portfolio. Similiar to Jiang, Li and Wang (2021), do positive firm-specific news returns predict higher future stock returns, and therefore the firm-specific news component increases from the $FN1$ tercile to the $FN3$ tercile. Stocks with bad (b) firm-specific news returns are denoted as N_b, whereas good (g) firm-specific news return are denoted as N_g. While bad firm-specific news returns are associated with negative news momentum and therefore expected to have negative returns in the future, are the good firm-specific news return associated with positive future returns. To derive the pure firm-specific news return effect we build a long-short strategy that relies solely on the return predictability of the firm-specific news return. Depending on the assumption that the 52-week high effect moderates the market underreaction to firm-specific news or not we define pure firm-specific news as:

$$\text{Pure Firm-Specific News Effect} = N_g - N_b = (b_1 + b_5) - (b_2 + b_{10}), \text{and} \tag{C.4}$$

$$= b_1 - b_4. \tag{C.5}$$

The fourth and last return component is associated with having, on the one hand, good firm-specific news about the firm and a stock price near the 52-week high and, on the other hand, experiencing bad firm-specific news while having a stock price that is far from the 52-week high. While the underreaction to the firm-specific news due to the nearness to the 52-week high could also be driven by the less extreme quintiles (e.g., the $FN2$ and $FN4$ quintile) but with a smaller magnitude, we focus our analysis on the most extreme FN and $NEAR$ quintiles. Stocks with extremely bad firm-specific news returns which are far from the 52-week high are denoted as $I_{b,f}$ whereas stocks with extremely good firm-specific news returns that are near the 52-week high are denoted as $I_{g,n}$. Hence, the interaction effect is defined as:

$$\text{Interaction Effect} = I_{g,n} - I_{b,f} = (b_7 - b_{11}) - (b_{14} - b_{10}) \tag{C.6}$$

If investors had a non-distorted belief updating process after the arrival of good (bad) firm-specific news while having a stock price that is near (far) its 52-week high, the interaction effect's long-short strategy would not yield a significant coefficient. In this case, the portfolio returns could still be fully attributable to the single components of the interaction, namely the pure firm-specific news effect and the pure 52-week high effect. On the other hand, if the coefficient of the interaction effect for the long-short strategy is positive and significant, one potential implication is that investors are not willing to update their beliefs and hence are underreacting to the good (bad) news if the stock price is near (far from) its 52-week high.

Datastream sample definition

Constituent lists

Datastream comprises three types of constituent lists: (1) research lists, (2) Worldscope lists, and (3) dead lists. By using dead lists, we ensure that any survivorship bias is obviated. For each country, we use the union of all available lists and eliminate any duplicates. As a result, one list remains for each country to be used in the subsequent static filter process. Table C. 2 provides an overview of the constituent lists for developed markets that are used in this study.

Static screens

I restrict the sample to common equity stocks by applying several static screens, as shown in Table C. 3. Screens (1) to (7) are straightforward to apply and common in the literature.

Screen (8) relates to, among others, to work by the following: Ince and Porter (2006), Campbell, Cowan and Salotti (2010), Griffin, Kelly and Nardari (2010), Karolyi, Lee and van Dijk (2012). The authors provide generic filter rules to exclude non-common equity securities from Refinitiv Datastream. we apply the identified keywords and match them with the security names provided by Datastream. A security is excluded from the sample in the event that a keyword coincides with part of the security name. The following three Datastream items store security names and are applied to the keyword filters: 'NAME', 'ENAME', and 'ECNAME'. Table C. 4 gives an overview of the keywords used.

In addition, Griffin, Kelly and Nardari (2010) introduce specific keywords for individual countries. The keywords are thus applied to the security names of single countries only. For example, German security names are parsed to contain the word 'GENUSSSCHEINE', which declares the security to be a non-common equity. In Table C. 5, we give an overview of country-specific keyword deletions conducted in our study.

Dynamic screens

For the securities remaining from the static screens above, we obtained return and market capitalization data from Datastream and accounting data from Worldscope. Several dynamic screens that are common in the literature were installed in order to account for data

Table C. 2
Constituent lists developed markets

The table contains the research lists, Worldscope lists and dead lists of developed markets countries in my sample.

Country	List	Country	List	Country	List
Australia	DEADAU	**Hong Kong**	DEADHK	**Spain**	DEADES
	FAUALL		FHKALL		WSCOPEES
	WSCOPEAU		WSCOPEHK		FESALL
Austria	WSCOPEOE	**Ireland**	WSCOPEIR		FSPDOM
	DEADAT		FIEALL		FSPNQ
	FATALL		DEADIE	**Sweden**	WSCOPESD
	FOSTDCT	**Israel**	DEADIL		FSEALL
	FOSTOM		WSCOPEIS		FXSTOALL
Belgium	FBEALL		FILALL		DEADSE
	WSCOPEBG	**Italy**	FITALL	**Switzerland**	WSCOPESW
	DEADBE		DEADIT		FCHALLP
Canada	DEADCA1		WSCOPEIT		DEADCH
	...	**Japan**	WSCOPEJP	**United Kingdom**	DEADGB
	DEADCA6		FJPALL		...
	WSCOPECN		FJPCONS		DEADGB7
	FXTSEALL		FTOKYO		FGBALL
	FCAALL		FXTKSALL		WSCOPEUK
Denmark	FDKALL		DEADJP	**United States**	WSUS1
	WSCOPEDK	**Netherlands**	DEADNL		...
	DEADDK		FNLALL		WSUS26
Finland	FFIALL		WSCOPENL		FUSALL1
	WSCOPEFN	**New Zealand**	WSCOPENZ		...
	DEADFI		FNZALL		FUSALL7
France	DEADFR		DEADNZ		FUSALLA
	WSCOPEFR	**Norway**	DEADNO		...
	FFRALL		FNOALL		FUSALLZ
Germany	DEADDE1		WSCOPENW		DEADUS1
	...	**Portugal**	WSCOPEPT		...
	DEADDE9		FPTALL		DEADUS12
	FGKURS		DEADPT		
	FDEALLP	**Singapore**	DEADSG		
	WSCOPEBD		FSGALL		
			FXSESM		
			WSCOPESG		

errors, mainly within return characteristics. The dynamic screens are shown in Table C. 6.

Table C. 3
Static screens

The table displays the static screens applied in our study, mainly following Ince and Porter (2006), Schmidt et al. (2017) and Griffin, Kelly and Nardari (2010). Column 3 lists the Datastream items involved (on the left of the equals sign) and the values which we set them to in the filter process (to the right of the equals sign). Column 4 indicates the source of the screens.

Nr.	Description	Datastream item(s) involved	Source
(1)	For firms with more than one security, only the one with the biggest market capitalization and liquidity is used.	MAJOR = Y	Schmidt et al. (2017)
(2)	The type of security must be equity.	TYPE = EQ	Ince and Porter (2006)
(3)	Only the primary quotations of a security are analyzed.	ISINID = P	Fong, Holden and Trzcinka (2017)
(4)	Firms are located in the respective domestic country.	GEOGN = country shortcut	Ince and Porter (2006)
(5)	Securities are listed in the respective domestic country.	GEOLN = country shortcut	Griffin, Kelly and Nardari (2010)
(6)	Securities whose quoted currency is different to the one of the associated country are disregarded.[a]	PCUR = currency shortcut of the country	Griffin, Kelly and Nardari (2010)
(7)	Securities whose ISIN country code is different to the one of the associated country are disregarded.[b]	GGISN = country shortcut	Annaert, Ceuster and Verstegen (2013)
(8)	Securities whose name fields indicate non-common stock affiliation are disregarded.	NAME, ENAME, ECNAME	Ince and Porter (2006), Campbell, Cowan and Salotti (2010), Griffin, Kelly and Nardari (2010) and Karolyi, Lee and van Dijk (2012)

[a] In this filter rule, the respective pre-euro currencies are also accepted for countries within the euro-zone. Moreover, in Russia ‘USD’ is accepted as currency, in addition to ‘RUB’.
[b] In Hong Kong, ISIN country codes equal to ‘BM’ or ‘KY’ and in the Czech Republic ISIN country codes equal to ‘CS’ are also accepted.

Table C. 4

Generic keyword deletions

The table reports generic keywords searched for in the names of all stocks of all countries. If a harmful keyword is detected as part of the name of a stock, the respective stock is removed from the sample.

Non-common equity	Keywords
Duplicates	1000DUPL, DULP, DUP, DUPE, DUPL, DUPLI, DUPLICATE, XSQ, XETa
Depository receipts	ADR, GDR
Preferred stock	PF, 'PF', PFD, PREF, PREFERRED, PRF
Warrants	WARR, WARRANT, WARRANTS, WARRT, WTS, WTS2
Debt	%, DB, DCB, DEB, DEBENTURE, DEBENTURES, DEBT
Unit trusts	.IT, .ITb, TST, INVESTMENT TRUST, RLST IT, TRUST, TRUST UNIT, TRUST UNITS, TST, TST UNIT, TST UNITS, UNIT, UNIT TRUST, UNITS, UNT, UNT TST, UT
ETFs	AMUNDI, ETF, INAV, ISHARES, JUNGE, LYXOR, X-TR
Expired securities	EXPD, EXPIRED, EXPIRY, EXPY
Miscellaneous (mainly taken from Ince and Porter (2006))	ADS, BOND, CAP.SHS, CONV, DEFER, DEP, DEPY, ELKS, FD, FUND, GW.FD, HI.YIELD, HIGH INCOME, IDX, INC.&GROWTH, INC.&GW, INDEX, LP, MIPS, MITS, MITT, MPS, NIKKEI, NOTE, OPCVM, ORTF, PARTNER, PERQS, PFC, PFCL, PINES, PRTF, PTNS, PTSHP, QUIBS, QUIDS, RATE, RCPTS, REAL EST, RECEIPTS, REIT, RESPT, RETUR, RIGHTS, RST, RTN.INC, RTS, SBVTG, SCORE, SPDR, STRYPES, TOPRS, UTS, VCT, VTG.SAS, XXXXX, YIELD, YLD

Table C. 5

Country-specific keyword deletions

The table reports country-specific keywords searched for in the names of all stocks of the respective countries. If a harmful keyword is detected as part of the name of a stock, the respective stock is removed from the sample.

Country	Keywords
Australia	PART PAID, RTS DEF, DEF SETT, CDI
Austria	PC, PARTICIPATION CERTIFICATE, GENUSSSCHEINE, GENUSSCHEINE
Belgium	VVPR, CONVERSION, STRIP
Canada	EXCHANGEABLE, SPLIT, SPLITSHARE, VTG\\., SBVTG\\., VOTING, SUB VTG, SERIES
Denmark	\\)CSE\\)
Finland	USE
France	ADP, CI, SICAV, \\)SICAV\\), SICAV-
Germany	GENUSSCHEINE
Israel	P1, 1, 5
Italy	RNC, RP, PRIVILEGIES
Netherlands	CERTIFICATE, CERTIFICATES, CERTIFICATES\\), CERT, CERTS, STK\\.
New Zealand	RTS, RIGHTS
Sweden	CONVERTED INTO, USE, CONVERTED-, CONVERTED - SEE
Switzerland	CONVERTED INTO, CONVERSION, CONVERSION SEE
United Kingdom	PAID, CONVERSION TO, NON VOTING, CONVERSION 'A'

Table C. 6
Dynamic screens

The table displays the dynamic screens applied to the data in our study, following Ince and Porter (2006), Griffin, Kelly and Nardari (2010), Jacobs (2016) and Schmidt et al. (2017). Column 3 lists the respective Datastream items. Column 4 refers to the source of the screens.

Nr.	Description	Datastream item(s) involved	Source
(1)	We delete the zero returns at the end of the return time-series that exist because in the case of a delisting, Datastream displays stale prices from the date of delisting until the end of the respective time-series. We also delete the associated market capitalizations.	RI, MV	Ince and Porter (2006)
(2)	We delete the associated returns and market capitalizations in case of abnormal prices (unadjusted prices $>$ 1000000).	RI, MV, UP	The screen originally stems from Schmidt et al. (2017), however we employ it on unadjusted price.
(3)	We delete monthly (daily) returns and the associated market capitalizations if returns exceed 990% (200%).	RI, MV	Griffin, Kelly and Nardari (2010); Schmidt et al. (2017)
(4)	We delete monthly returns and the associated market capitalizations in the case of strong return reversals, defined as $(1+r_{t-1})(1+r_t)-1 < 0.5$ given that either r_{t-1} or $r_t \geq 3.0$.	RI, MV	Ince and Porter (2006)
(5)	We delete daily returns and the associated market capitalizations in the case of strong return reversals, defined as $(1+r_{t-1})(1+r_t)-1 < 0.2$ with r_{t-1} or $r_t \geq 1.0$.	RI, MV	Griffin, Kelly and Nardari (2010); Jacobs (2016)

Factor construction

We calculate the market factor as the value-weighted returns of all available stocks in excess of the risk-free rate. For the factors value, profitability, investment, and momentum, we estimate the portfolio breakpoints using the country-specific 30% and 70% percentile of the underlying characteristic using only the big-stock sample. In the case of the value stocks, we use the book-to-market ratio to categorize the stocks as Growth (G), Neutral (N), and Value (V). For profitability, we use the cash-based profitability as an underlying characteristic which enables us to sort the stocks into the extreme portfolios Weak (W) and Robust (R). In the case of the investment factor, we base the sorting on the stock's asset growth, which yields a Conservative (C) and Aggressive (A) portfolio. The next factor is based on the stock's momentum and sorts the stocks into the Winner (W) and Loser (L) portfolios. The last factor is based on the stock's Amihud (2002) illiquidity and sorts the stocks into the liquid (AL) and illiquid (AI) portfolios. We follow the size group methodology of Fama and French (2008, 2012, 2017) and assign stocks into three size groups (micro, small, and big) separately for each country and month. Big stocks are defined as the biggest stocks, which together account for 90% of a country's aggregated market capitalization. Small stocks are defined as those stocks that comprise the next 7% of aggregated market capitalization (so that big and small stocks together account for 97% of the aggregated market size of a country). Microcaps comprise the remaining 3%.[1] The final factor calculation is based on the intersection of the different portfolios, while

[1] To distinguish between these size groups, Fama and French (2008) use the 20^{th} and 50^{th} percentiles of end-of-June market cap on NYSE stocks as size breakpoints for the U.S. market, which on average are bigger than AMEX or NASDAQ stocks. However, these breakpoints are applied to all (NYSE, AMEX, and NASDAQ) stocks. For international markets, Fama and French (2012, 2017) propose to calculate breakpoints based on aggregated market capitalization, as we do.

the portfolio returns are value-weighted,

$$\begin{aligned}
SMB &= (SV + SN + SG)/3 - (BV + BN + BG)/3, \\
HML &= (BV + SV)/2 - (BG + SG)/2, \\
RMW &= (BR + SR)/2 - (BW + SW)/2, \\
CMA &= (BC + SC)/2 - (BA + SA)/2, \\
MOM &= (BW + SW)/2 - (BL + SL)/2, \\
LIQ &= (BAL + SAL)/2 - (BAI + SAI)/2.
\end{aligned} \tag{C.7}$$

Bibliography

Abdi, F., Ranaldo, A., 2017. A Simple Estimation of Bid-Ask Spreads from Daily Close, High, and Low Prices. The Review of Financial Studies 30 (12), 4437–4480.

Ali, U., Hirshleifer, D., 2020. Shared analyst coverage: Unifying momentum spillover effects. Journal of Financial Economics 136 (3), 649–675.

Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5 (1), 31–56.

Anand, V., Brunner, R., Ikegwu, K., Sougiannis, T., 2019. Predicting profitability using machine learning. Working Paper.

Ang, A., Hodrick, R. J., Xing, Y., Zhang, X., 2006. The cross-section of volatility and expected returns. The Journal of Finance 61 (1), 259–299.

Annaert, J., Ceuster, M. D., Verstegen, K., 2013. Are extreme returns priced in the stock market? European evidence. Journal of Banking & Finance 37 (9), 3401–3411.

Ardia, D., Guidotti, E., Kroencke, T. A., 2022. Efficient estimation of bid-ask spreads from open, high, low, and close prices. Working Paper.

Asness, C. S., 2011. Momentum in Japan: The exception that proves the rule. Journal of Portfolio Management 37 (4), 67–75.

Asness, C. S., Frazzini, A., 2013. The devil in HML's details. Journal of Portfolio Management 39 (4), 49–68.

Asness, C. S., Frazzini, A., Pedersen, L. H., 2019. Quality minus junk. Review of Accounting Studies 24 (1), 34–112.

Avramov, D., Cheng, S., Metzker, L., 2022. Machine learning vs. economic restrictions: Evidence from stock return predictability. Management Science , forthcoming (0).

Azevedo, V., Kaiser, G. S., Kaiser, S., Müller, S., 2022. Stock market anomalies and machine learning across the globe. Working Paper.

Balakrishnan, K., Bartov, E., Faurel, L., 2010. Post loss/profit announcement drift. Journal of Accounting and Economics 50 (1), 20–41.

Bali, T. G., Beckmeyer, H., Moerke, M., Weigert, F., 2022a. Option return predictability with machine learning and big data. Working Paper.

Bali, T. G., Goyal, A., Huang, D., Jiang, F., Wen, Q., 2022b. Predicting corporate bond returns: Merton meets machine learning. Working Paper.

Ball, R., Gerakos, J., Linnainmaa, J. T., Nikolaev, V. V., 2015. Deflating profitability. Journal of Financial Economics 117 (2), 225–248.

Ball, R., Gerakos, J., Linnainmaa, J. T., Nikolaev, V. V., 2016. Accruals, cash flows, and operating profitability in the cross section of stock returns. Journal of Financial Economics 121 (1), 28–45.

Barber, B. M., Odean, T., 2008. All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. Review of Financial Studies 21 (2), 785–818.

Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. Journal of Financial Economics 49 (3), 307–343.

Basu, S., 1983. The relationship between earnings yield, market value and return for nyse common stocks : Further evidence. Journal of Financial Economics 12 (1), 129–156.

Bekaert, G., Engstrom, E. C., Xu, N. R., 2022. The time variation in risk appetite and uncertainty. Management Science 68 (6), 3975–4004.

Bekaert, G., Harvey, C. R., 1995. Time-varying world market integration. The Journal of Finance 50 (2), 403–444.

Ben-Rephael, A., Da, Z., Israelsen, R. D., 2017. It depends on where you search: Institutional investor attention and underreaction to news. The Review of Financial Studies 30 (9), 3009–3047.

Bhandari, L. C., 1988. Debt/equity ratio and expected common stock returns: Empirical evidence. The Journal of Finance 43 (2), 507–528.

Bianchi, D., Büchner, M., Hoogteijling, T., Tamoni, A., 2021. Corrigendum: Bond risk premiums with machine learning. The Review of Financial Studies 34 (2), 1090–1103.

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

Birru, J., 2013. Psychological Barriers, Expectational Errors, and Underreaction to News. SSRN Electronic Journal.

Blitz, D., Hanauer, M. X., Honarvar, I., Huisman, R., van Vliet, P., 2022. Beyond Fama-French factors: Alpha from short-term signals. Working Paper.

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

Brogaard, J., Nguyen, T. H., Putnins, T. J., Wu, E., 2022. What Moves Stock Prices? The Roles of News, Noise, and Information. The Review of Financial Studies 35 (9), 4341–4386.

Burt, A., Hrdlicka, C., 2021. Where Does the Predictability from Sorting on Returns of Economically Linked Firms Come From? Journal of Financial and Quantitative Analysis 56 (8), 2634–2658.

Byun, S.-J., Goh, J., Kim, D.-H., 2020. The role of psychological barriers in lottery-related anomalies. Journal of Banking Finance 114, 105786.

Cakici, N., Fabozzi, F. J., Tan, S., 2013. Size, value, and momentum in emerging market stock returns. Emerging Markets Review 16, 46–65.

Cakici, N., Fieberg, C., Metko, D., Zaremba, A., 2022a. Machine learning goes global: Cross-sectional return predictability in international stock markets. Working Paper.

Cakici, N., Shahzad, S. J. H., Bedowska-Sojka, B., Zaremba, A., 2022b. Machine learning and the cross-section of cryptocurrency returns. SSRN Working Paper no. 4295427.

Cakici, N., Zaremba, A., 2022. Empirical asset pricing via machine learning: The global edition. Working Paper.

Campbell, C. J., Cowan, A. R., Salotti, V., 2010. Multi-country event-study methods. Journal of Banking & Finance 34 (12), 3078–3090.

Campbell, J. Y., 1991. A Variance Decomposition for Stock Returns. The Economic Journal 101 (405), 157.

Campbell, J. Y., Shiller, R. J., 1988. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. The Review of Financial Studies 1 (3), 195–228.

Campbell, S. D., Sharpe, S. A., 2009. Anchoring Bias in Consensus Forecasts and Its Effect on Market Prices. Journal of Financial and Quantitative Analysis 44 (2), 369–390.

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

Cen, L., Hilary, G., Wei, K. C. J., 2013. The Role of Anchoring Bias in the Equity Market: Evidence from Analysts' Earnings Forecasts and Stock Returns. Journal of Financial and Quantitative Analysis 48 (1), 47–76.

Chan, L. K. C., Karceski, J., Lakonishok, J., 1998. The risk and return from factors. Journal of Financial and Quantitative Analysis 33, 159–188.

Chan, W. S., 2003. Stock price reaction to news and no-news: drift and reversal after headlines. Journal of Financial Economics 70 (2), 223–260.

Chen, L., Da, Z., Zhao, X., 2013. What drives stock price movements? The Review of Financial Studies 26 (4), 841–876.

Chen, L., Pelger, M., Zhu, J., 2021. Deep learning in asset pricing. Working Paper.

Chen, X., Cho, Y. H. T., Dou, Y., Lev, B., 2022. Predicting future earnings changes using machine learning and detailed financial data. Journal of Accounting Research 60 (2), 467–515.

Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. The Journal of Finance 63 (4), 1977–2011.

Cohen, L., Lou, D., 2012. Complicated Firms. Journal of Financial Economics 104 (2), 383–400.

Coibion, O., Gorodnichenko, Y., 2015. Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. American Economic Review 105 (8), 2644–2678.

Cooper, M. J., Gulen, H., Schill, M. J., 2008. Asset growth and the cross-section of stock returns. The Journal of Finance 63 (4), 1609–1651.

Corwin, S. A., Schultz, P., 2012. A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices. The Journal of Finance 67 (2), 719–760.

Da, Z., Liu, Q., Schaumburg, E., 2014. A closer look at the short-term return reversal. Management Science 60 (3), 658–674.

Da, Z., Warachka, M. C., 2009. Cashflow risk, systematic earnings revisions, and the cross-section of stock returns. Journal of Financial Economics 94 (3), 448–468.

Daniel, K., Hirshleifer, D., Subrahmanyam, A., 1998. Investor Psychology and Security Market Under- and Overreactions. The Journal of Finance 53 (6), 1839–1885.

Daniel, K., Moskowitz, T. J., 2016. Momentum crashes. Journal of Financial Economics 122 (122), 221–247.

Daniel, K., Titman, S., 2006. Market reactions to tangible and intangible information. The Journal of Finance 61 (4), 1605–1643.

Datar, V. T., Y. Naik, N., Radcliffe, R., 1998. Liquidity and stock returns: An alternative test. Journal of Financial Markets 1 (2), 203–219.

Davis, J. L., Fama, E. F., French, K. R., 2000. Characteristics, covariances, and average returns: 1929 to 1997. The Journal of Finance 55 (1), 389–406.

De Bondt, W. F. M., Thaler, R., 1985. Does the stock market overreact? The Journal of Finance 40 (3), 793–805.

De Nard, G., Hediger, S., Leippold, M., 2022. Subsampled factor models for asset pricing: The rise of vasa. Journal of Forecasting 41 (6), 1217–1247.

DellaVigna, S., Polle, J. M., 2009. Investor Inattention and Friday Earnings Announcements. The Journal of Finance 64 (2), 709–749.

Diebold, F. X., Mariano, R. S., 1995. Comparing predictive accuracy. Journal of Business & Economic Statistics 13 (3), 134–144.

Diether, K. B., Malloy, C. J., Scherbina, A., 2002. Differences of opinion and the cross section of stock returns. The Journal of Finance 57 (5), 2113–2141.

Dietterich, T. G., 2000. Ensemble methods in machine learning. In: Multiple Classifier

Systems. Springer Berlin Heidelberg, Berlin, Heidelberg, p. 1.

Drake, M. S., Guest, N. M., Twedt, B. J., 2014. The media and mispricing: The role of the business press in the pricing of accounting information. Accounting Review 89 (5), 1673–1701.

Drobetz, W., Hollstein, F., Otto, T., Prokopczuk, M., 2021. Estimating security betas via machine learning. Working Paper.

Drobetz, W., Otto, T., 2021. Empirical asset pricing via machine learning: evidence from the European stock market. Journal of Asset Management 22 (7), 507–538.

Duan, Y., Hu, G., McLean, R. D., 2010. Costly arbitrage and idiosyncratic risk: Evidence from short sellers. Journal of Financial Intermediation 19 (4), 564–579.

D'Acunto, F., Liu, R., Pflueger, C., Weber, M., 2018. Flexible prices and leverage. Journal of Financial Economics 129 (1), 46–68.

Easton, P. D., Monahan, S. J., 2005. An evaluation of accounting-based measures of expected returns. Accounting Review 80 (2), 501–538.

Engelberg, J., McLean, D., Pontiff, J., 2018. Anomalies and News. The Journal of Finance 73 (5), 1971–2001.

Erel, I., Stern, L. H., Tan, C., Weisbach, M. S., 2021. Selecting directors using machine learning. The Review of Financial Studies 34 (7), 3226–3264.

Estrada, G. B., Park, D., Ramayandi, A., 2016. Taper tantrum and emerging equity market slumps. Emerging Markets Finance and Trade 52 (5), 1060–1071.

Fama, E., 1970. Efficient Capital Markets. A Review of Theory and Empirical Work. The Journal of Finance 25 (2), 383–417.

Fama, E. F., French, K. R., 1992. The cross-section of expected stock returns. The Journal of Finance 47 (2), 427–465.

Fama, E. F., 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., French, K. R., 1996. Multifactor explanations of asset pricing anomalies. The Journal of Finance 51 (1), 55–84.

Fama, E. F., French, K. R., 2008. Dissecting anomalies. The Journal of Finance 63 (4), 1653–1678.

Fama, E. F., French, K. R., 2012. Size, value, and momentum in international stock returns. Journal of Financial Economics 105 (3), 457–472.

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

Fama, E. F., French, K. R., 2017. International tests of a five-factor asset pricing model. Journal of Financial Economics 123 (3), 441–463.

Fama, E. F., French, K. R., 2018. Choosing factors. Journal of Financial Economics 128 (2), 234–252.

Fama, E. F., MacBeth, J. D., 1973. Risk, return, and equilibrium: Empirical tests. Journal of Political Economy 81 (3), 607–636.

Fong, K. Y. L., Holden, C. W., Trzcinka, C. A., 2017. What are the best liquidity proxies for global research? Review of Finance 21 (4), 1355–1401.

Francis, J., Pagach, D., Stephan, J., 1992. The Stock Market Response to Earnings Announcements Released During Trading versus Nontrading Periods. Journal of Accounting Research 30 (2), 165.

Frazzini, A., Pedersen, L. H., 2014. Betting against beta. Journal of Financial Economics

111 (1), 1 – 25.

Freyberger, J., Neuhierl, A., Weber, M., 2020. Dissecting characteristics nonparametrically. The Review of Financial Studies 33 (5), 2326–2377.

Gandhi, P., Lustig, H., 2015. Size anomalies in U.S. bank stock returns. The Journal of Finance 70 (2), 733–768.

Garfinkel, J. A., 2009. Measuring investors' opinion divergence. Journal of Accounting Research 47 (5), 1317–1348.

George, T. J., Hwang, C.-Y., 2004. The 52-week high and momentum investing. The Journal of Finance 59 (5), 2145–2176.

George, T. J., Hwang, C.-Y., Li, Y., 2014. Anchoring, the 52-Week High and Post Earnings Announcement Drift. SSRN Electronic Journal.

Gleason, C. A., Lee, C. M. C., 2003. Analyst Forecast Revisions and Market Price Discovery. The Accounting Review 78 (1), 193–225.

Gorodnichenko, Y., Weber, M., 2016. Are sticky prices costly? Evidence from the stock market. American Economic Review 106 (1), 165–99.

Griffin, J. M., Hirschey, N. H., Kelly, P. J., 2011. How important is the financial media in global markets? The Review of Financial Studies 24 (12), 3941–3992.

Griffin, J. M., Kelly, P. J., Nardari, F., 2010. Do market efficiency measures yield correct inferences? A comparison of developed and emerging markets. The Review of Financial Studies 23 (8), 3225–3277.

Gu, S., Kelly, B. T., Xiu, D., 2020. Empirical asset pricing via machine learning. The Review of Financial Studies 33 (5), 2223–2273.

Gutierrez, R. C., Kelly, E. K., 2008. The Long-Lasting Momentum in Weekly Returns.

The Journal of Finance 63 (1), 415–447.

Hanauer, M., 2020. A comparison of global factor models. Working Paper.

Hanauer, M. X., Kalsbach, T., 2022. Machine Learning and The Cross-Section of Emerging Market Stock Returns. SSRN Electronic Journal.

Hanauer, M. X., Kononova, M., Rapp, M. S., 2022. Boosting agnostic fundamental analysis: Using machine learning to identify mispricing in european stock markets. Finance Research Letters 48, 102856.

Hanauer, M. X., Lauterbach, J. G., 2019. The cross-section of emerging market stock returns. Emerging Markets Review 38, 265–286.

Hanauer, M. X., Linhart, M., 2015. Size, value, and momentum in emerging market stock returns: Integrated or segmented pricing? Asia-Pacific Journal of Financial Studies 44 (2), 175–214.

Hansen, L., Salamon, P., 1990. Neural network ensembles. IEEE Transactions on Pattern Analysis and Machine Intelligence 12 (10), 993–1001.

Harvey, C. R., 1995. Predictable risk and returns in emerging markets. The Review of Financial Studies 8 (3), 773–816.

Harvey, C. R., Liu, Y., Zhu, H., 2016. … and the Cross-Section of Expected Returns. The Review of Financial Studies 29 (1), 5–68.

Haugen, R. A., Baker, N. L., 1996. Commonality in the determinants of expected stock returns. Journal of Financial Economics 41 (3), 401–439.

Heath, C., Huddart, S., Lang, M., 1999. Psychological Factors and Stock Option Exercise. The Quarterly Journal of Economics 114 (2), 601–627.

Hillert, A., Jacobs, H., Müller, S., 2014. Media makes momentum. The Review of Financial

Studies 27 (12), 3467–3501.

Hirshleifer, D., Hsu, P. H., Li, D., 2013. Innovative efficiency and stock returns. Journal of Financial Economics 107 (3), 632–654.

Hirshleifer, D., Kewei Hou, Teoh, S. H., Yinglei Zhang, 2004. Do investors overvalue firms with bloated balance sheets? Journal of Accounting and Economics 38, 297–331.

Hirshleifer, D., Lim, S. S., Teoh, S. H., 2011. Limited Investor Attention and Stock Market Misreactions to Accounting Information. Review of Asset Pricing Studies 1 (1), 35–73.

Hirshleifer, D., Teoh, S. H., 2003. Limited attention, information disclosure, and financial reporting. Journal of Accounting and Economics 36 (1-3 SPEC. ISS.), 337–386.

Hoberg, G., Phillips, G. M., 2018. Text-Based Industry Momentum. Journal of Financial and Quantitative Analysis 53 (6), 2355–2388.

Hong, H., Stein, J. C., 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. The Journal of Finance 54 (6), 2143–2184.

Hong, H., Torous, W., Valkanov, R., 2007. Do industries lead stock markets? Journal of Financial Economics 83 (2), 367–396.

Hong, X., Jordan, B. D., Liu, M. H., 2015. Industry information and the 52-week high effect. Pacific-Basin Finance Journal 32, 111–130.

Hou, K., 2007. Industry information diffusion and the lead-lag effect in stock returns. The Review of Financial Studies 20 (4), 1113–1138.

Hou, K., Karolyi, G. A., Kho, B.-C., 2011. What factors drive global stock returns? The Review of Financial Studies 24 (8), 2527–2574.

Hou, K., Xiong, W., Peng, L., 2009. A Tale of Two Anomalies: The Implications of Investor Attention for Price and Earnings Momentum. Working Paper.

Hou, K., Xue, C., Zhang, L., 2018. Replicating anomalies. The Review of Financial Studies 33 (5), 2019–2133.

Huang, S., Lee, C. M., Song, Y., Xiang, H., 2022. A frog in every pan: Information discreteness and the lead-lag returns puzzle. Journal of Financial Economics 145 (2), 83–102.

Huang, S., Lin, T.-C., Xiang, H., 2021. Psychological barrier and cross-firm return predictability. Journal of Financial Economics 142 (1), 338–356.

Huddart, S., Lang, M., Yetman, M. H., 2009. Volume and Price Patterns Around a Stock's 52-Week Highs and Lows: Theory and Evidence. Management Science 55 (1), 16–31.

Hung, M., Li, X., Wang, S., 2015. Post-Earnings-Announcement Drift in Global Markets: Evidence from an Information Shock. The Review of Financial Studies 28 (4), 1242–1283.

Hung, W., Lin, C.-T., Yang, J. J., 2022. Aggregate 52-week high, limited attention, and time-varying momentum profits. Journal of Banking Finance 141, 106531.

Ince, O. S., Porter, R. B., 2006. Individual equity return data from Thomson Datastream: Handle with care! Journal of Financial Research 29 (4), 463–479.

Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift.

Jacobs, H., 2016. Market maturity and mispricing. Journal of Financial Economics 122 (2), 270–287.

Jacobs, H., Müller, S., 2020. Anomalies across the globe: Once public, no longer existent? Journal of Financial Economics 135 (1), 213–230.

Jansen, M., Swinkels, L., Zhou, W., 2021. Anomalies in the china a-share market. Pacific-Basin Finance Journal 68, 101607.

Jaquart, P., Dann, D., Weinhardt, C., 2021. Short-term bitcoin market prediction via machine learning. The Journal of Finance and Data Science 7, 45–66.

Jegadeesh, N., 1990. Evidence of predictable behavior of security returns. The Journal of Finance 45 (3), 881–898.

Jensen, M. C., Heckling, W. H., 1995. Specific and General Knowledge, and Organizational Structure. Journal of Applied Corporate Finance 8 (2), 4–18.

Jensen, T. I., Kelly, B. T., Pedersen, L. H., 2021. Is There a Replication Crisis in Finance? Working Paper.

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

Jiang, H., Sun, Z., 2014. News and Corporate Bond Liquidity. SSRN Electronic Journal.

Jiang, Y., Qian, Y., Yao, T., 2016. R&D Spillover and Predictable Returns. Review of Finance 20 (5), 1769–1797.

Kahneman, D., 1973. Attention and Effort. Prentice Hall, New Jersey.

Kalsbach, T., 2023. What Diffuses in Stock Prices? The Role of News and Noise in Global Networks. SSRN Electronic Journal.

Kaniel, R., Lin, Z., Pelger, M., Van Nieuwerburgh, S., 2022. Machine-learning the skill of mutual fund managers. Working Paper.

Karolyi, A. G., 2016. Home bias, an academic puzzle. Review of Finance 20 (6), 2049–2078.

Karolyi, G. A., Lee, K.-H., van Dijk, M. A., 2012. Understanding commonality in liquidity around the world. Journal of Financial Economics 105 (1), 82–112.

Kaufmann, H., Messow, P., Vogt, J., 2021. Boosting the equity momentum factor in credit. Financial Analysts Journal 77 (4), 83–103.

Kaustia, M., Rantala, V., 2020. Common Analysts: Method for Defining Peer Firms. Journal of Financial and Quantitative Analysis, 1–32.

Ke, Z., Kelly, B. T., Xiu, D., 2020. Predicting Returns with Text Data. SSRN Electronic Journal.

Kelley, E. K., Tetlock, P. C., 2017. Retail Short Selling and Stock Prices. The Review of Financial Studies 30 (3), 801–834.

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

Kingma, D. P., Ba, J., 2014. Adam: A method for stochastic optimization.

Kovacs, T., 2016. Intra-Industry Information Transfers and the Post-Earnings Announcement Drift. Contemporary Accounting Research 33 (4), 1549–1575.

Kumar, A., 2009. Hard-to-value stocks, behavioral biases, and informed trading. Journal of Financial and Quantitative Analysis 44 (6), 1375–1401.

Lakonishok, J., Shleifer, A., Vishny, R. W., 1994. Contrarian investment, extrapolation, and risk. The Journal of Finance 49 (5), 1541–1578.

Lam, F. E. C., Wei, K. J., 2011. Limits-to-arbitrage, investment frictions, and the asset growth anomaly. Journal of Financial Economics 102 (1), 127–149.

Landis, C., Skouras, S., 2021. Guidelines for asset pricing research using international equity data from Thomson Reuters Datastream. Journal of Banking Finance 130, 106128.

Lee, C. M., Shi, T., Sun, S. T., Zhang, R., 2022. Production Complementarity and Momentum Spillover Across Industries. Working Paper.

Lee, C. M., So, E. C., 2017. Uncovering expected returns: Information in analyst coverage proxies. Journal of Financial Economics 124 (2), 331–348.

Lee, C. M., Sun, S. T., Wang, R., Zhang, R., 2019. Technological Links and Predictable Returns. Journal of Financial Economics 132 (3), 76–96.

Leippold, M., Wang, Q., Zhou, W., 2022. Machine learning in the chinese stock market. Journal of Financial Economics 145 (2, Part A), 64–82.

Leung, E., Lohre, H., Mischlich, D., Shea, Y., Stroh, M., 2021. The promises and pitfalls of machine learning for predicting stock returns. The Journal of Financial Data Science 3 (2), 21–50.

Lewellen, J., 2015. The cross-section of expected stock returns. Critical Finance Review 4 (1), 1–44.

Lewellen, J., Nagel, S., 2006. The conditional capm does not explain asset-pricing anomalies. Journal of Financial Economics 82 (2), 289 – 314.

Litzenberger, R. H., Ramaswamy, K., 1979. The effect of personal taxes and dividends on capital asset prices: Theory and empirical evidence. Journal of Financial Economics 7 (2), 163–195.

Lyandres, E., Sun, L., Zhang, L., 2007. The New Issues Puzzle: Testing the Investment-Based Explanation. The Review of Financial Studies 21 (6), 2825–2855.

Mashruwala, C., Rajgopal, S., Shevlin, T., 2006. Why is the accrual anomaly not arbitraged away? The role of idiosyncratic risk and transaction costs. Journal of Accounting and Economics 42 (1-2), 3–33.

Masters, T., 1993. Practical Neural Network Recipes in C++. Academic Press Professional, Inc., USA.

McLean, R. D., 2010. Idiosyncratic Risk, Long-Term Reversal, and Momentum. Journal of Financial and Quantitative Analysis 45 (4), 883–906.

Medhat, M., Schmeling, M., 2022. Short-term Momentum. The Review of Financial Studies 35 (3), 1480–1526.

Menzly, L., Ozbas, O., 2010. Market Segmentation and Cross-predictability of Returns. The Journal of Finance 65 (4), 1555–1580.

Merton, R. C., 1987. A simple model of capital market equilibrium with incomplete information. The Journal of Finance 42 (3), 483–510.

Morck, R., Yeung, B., Yu, W., 2000. The information content of stock markets: why do emerging markets have synchronous stock price movements? Journal of Financial Economics 58 (1-2), 215–260.

Moritz, B., Zimmermann, T., 2016. Tree-based conditional portfolio sorts: The relation between past and future stock returns. Working Paper.

Moskowitz, T. J., Grinblatt, M., 1999. Do Industries Explain Momentum? The Journal of Finance 54 (4), 1249–1290.

Nagel, S., 2005. Short sales, institutional investors and the cross-section of stock returns. Journal of Financial Economics 78 (2), 277–309.

Newey, W. K., West, K. D., 1987. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. Econometrica 55 (3), 703–708.

Novy-Marx, R., 2010. Operating leverage. Review of Finance 15 (1), 103–134.

Novy-Marx, R., 2012. Is momentum really momentum? Journal of Financial Economics 103 (3), 429–453.

Novy-Marx, R., 2013. The other side of value: The gross profitability premium. Journal of Financial Economics 108 (1), 1–28.

Novy-Marx, R., Velikov, M., 2016. A taxonomy of anomalies and their trading costs. The

Review of Financial Studies 29 (1), 104–147.

Palazzo, B., 2012. Cash holdings, risk, and expected returns. Journal of Financial Economics 104 (1), 162–185.

Parsons, C. A., Sabbatucci, R., Titman, S., 2020. Geographic Lead-Lag Effects. The Review of Financial Studies 33 (10), 4721–4770.

Pástor, L., Stambaugh, R. F., 2003. Liquidity Risk and Expected Stock Returns. Journal of Political Economy 111 (3), 642–685.

Peng, L., Xiong, W., 2006. Investor attention, overconfidence and category learning. Journal of Financial Economics 80 (3), 563–602.

Petersen, M. A., 2009. Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches. The Review of Financial Studies 22 (1), 435–480.

Pontiff, J., 1996. Costly Arbitrage: Evidence from Closed-End Funds. The Quarterly Journal of Economics 111 (4), 1135–1151.

Pontiff, J., 2006. Costly arbitrage and the myth of idiosyncratic risk. Journal of Accounting and Economics 42 (1), 35–52, conference Issue on Implications of Changing Financial Reporting Standards.

Poteshman, A. M., Serbin, V., 2003. Clearly Irrational Financial Market Behavior: Evidence from the Early Exercise of Exchange Traded Stock Options. The Journal of Finance 58 (1), 37–70.

Ramnath, S., 2002. Investor and Analyst Reactions to Earnings Announcements of Related Firms: An Empirical Analysis. Journal of Accounting Research 40 (5), 1351–1376.

Rapach, D. E., Strauss, J. K., Tu, J., Zhou, G., 2019. Industry return predictability: A machine learning approach. The Journal of Financial Data Science 1 (3), 9–28.

Rasekhschaffe, K. C., Jones, R. C., 2019. Machine learning for stock selection. Financial Analysts Journal 75 (3), 70–88.

Roll, R., 1988. R2. The Journal of Finance 43 (3), 541–566.

Roon, F. A. D., Nijman, T. E., Werker, B. J. M., 2001. Testing for mean-variance spanning with short sales constraints and transaction costs: The case of emerging markets. The Journal of Finance 56 (2), 721–742.

Rosenberg, B., Reid, K., Lanstein, R., 1985. Persuasive evidence of market inefficiency. Journal of Portfolio Management 11 (3), 9–16.

Rossi, A. G., 2018. Predicting stock market returns with machine learning. Working paper.

Rouwenhorst, K. G., 1999. Local return factors and turnover in emerging stock markets. The Journal of Finance 54 (4), 1439–1464.

Rubesam, A., 2022. Machine learning portfolios with equal risk contributions: Evidence from the brazilian market. Emerging Markets Review 51, 100891.

Sadhwani, A., Giesecke, K., Sirignano, J., 2021. Deep learning for mortgage risk. Journal of Financial Econometrics 19 (2), 313–368.

Savor, P., Wilson, M., 2013. How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements. Journal of Financial and Quantitative Analysis 48 (2), 343–375.

Schmidt, P. S., Von Arx, U., Schrimpf, A., Wagner, A. F., Ziegler, A., 2017. On the construction of common size, value and momentum factors in international stock markets: A guide with applications. Swiss Finance Institute Research Paper 10.

Schmidt, P. S., von Arx, U., Schrimpf, A., Wagner, A. F., Ziegler, A., 2019. Common risk factors in international stock markets. Financial Markets and Portfolio Management 33 (3), 213–241.

Sharpe, W. F., 1966. Mutual Fund Performance. The Journal of Business 39 (S1), 119.

Shleifer, A., Vishny, R. W., 1997. The Limits of Arbitrage. The Journal of Finance 52 (1), 35–55.

Simon, H. A., 1955. A Behavioral Model of Rational Choice. The Quarterly Journal of Economics 69 (1), 99.

Sloan, R. G., 1996. Do stock prices fully reflect information in accruals and cash flows about future earnings? The Accounting Review 71 (3), 289–315.

Smajlbegovic, E., 2019. Regional economic activity and stock returns. Journal of Financial and Quantitative Analysis 54 (3), 1051—1082.

Soliman, M. T., 2008. The use of DuPont analysis by market participants. The Accounting Review 83 (3), 823–853.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research 15 (1), 1929–1958.

Stambaugh, R. F., Yu, J., Yuan, Y., 2015. Arbitrage asymmetry and the idiosyncratic volatility puzzle. The Journal of Finance 70 (5), 1903–1948.

Struck, C., Cheng, E., 2020. The cross section of commodity returns: A nonparametric approach. The Journal of Financial Data Science 2 (3), 86–103.

Tobek, O., Hronec, M., 2020. Does it pay to follow anomalies research? Machine learning approach with international evidence. Journal of Financial Markets, 100588.

Tversky, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. Science 185 (4157), 1124–1131.

Twedt, B., 2016. Spreading the Word: Price discovery and newswire dissemination of management earnings guidance. Accounting Review 91 (1), 317–346.

Van Binsbergen, J. H., Han, X., Lopez-Lira, A., 2020. Man vs. machine learning: The term structure of earnings expectations and conditional biases. NBER Working Paper No. 27843.

van der Hart, J., de Zwart, G., van Dijk, D., 2005. The success of stock selection strategies in emerging markets: Is it risk or behavioral bias? Emerging Markets Review 6 (3), 238–262.

van der Hart, J., Slagter, E., van Dijk, D., 2003. Stock selection strategies in emerging markets. Journal of Empirical Finance 10 (1-2), 105–132.

Vuolteenaho, T., 2002. What drives firm-level stock returns? The Journal of Finance 57 (1), 233–264.

Welch, I., 2020. Simpler better market betas. Working Paper.

Windmüller, S., 2022. Firm characteristics and global stock returns: A conditional asset pricing model. The Review of Asset Pricing Studies 12 (2), 447–499.

Wooldridge, J. M., December 2001. Applications of generalized method of moments estimation. Journal of Economic Perspectives 15 (4), 87–100.

Wu, W., Chen, J., Yang, Z. B., Tindall, M. L., 2021. A cross-sectional machine learning approach for hedge fund return prediction and selection. Management Science 67 (7), 4577–4601.

Wurgler, J., Zhuravskaya, E., 2002. Does Arbitrage Flatten Demand Curves for Stocks? The Journal of Business 75 (4), 583–608.

Ying, J., 2020. Institutional Trading on Information Diffusion across Fundamentally Related Firms. Working Paper.

Yuan, Y., 2015. Market-wide attention, trading, and stock returns. Journal of Financial Economics 116 (3), 548–564.

Zaremba, A., Czapkiewicz, A., 2017. Digesting anomalies in emerging European markets: A comparison of factor pricing models. Emerging Markets Review 31, 1–15.

Zhang, X., 2006. Specification tests of international asset pricing models. Journal of International Money and Finance 25 (2), 275–307.