From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2023

Essays on Investor Heterogeneity in Equity Markets

Noh, Dong Woo

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp018336h5169

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Essays on Investor Heterogeneity in Equity Markets

Dong Woo Noh

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Economics
Advisers: Motohiro Yogo and Wei Xiong

May 2023

© Copyright by Dong Woo Noh, 2023.

All Rights Reserved

Abstract

This dissertation examines asset pricing implications of investor heterogeneity in equity markets. Chapters 1 and 2 study differences between retail investors and institutional investors, while Chapter 3 focuses on heterogeneity among different institutions.

Chapter 1 studies retail investor trading behavior around stock-level trading pauses in Korea, which are triggered after extreme intraday price movements. The stated goals of this market intervention tool are investor protection and price stability. However, I find no significant improvement in measures of price stability, and instead find unintended negative consequences on retail investors. Because trading pauses are salient events, retail investors trade excessively once pauses are triggered, and they trade in the wrong direction by betting on price continuation. Trading pauses become a window of wealth transfer from retail to institutional investors, and estimates for the wealth transfer and trading costs together amount to approximately USD 350 million per annum.

Chapter 2 (co-authored with Stella Park and Alfred Qi Fan) studies the "overnight-intraday return gap." In both developed and emerging stock markets, overnight (close-to-open) returns are higher relative to intraday (open-to-close) returns. We propose a noise trader-based explanation for this overnight-intraday return gap. We use exact and exhaustive retail investor trade flow data available for the Korean stock market to establish the relationship between the overnight-intraday return gap and retail investor intensity. Furthermore, we use the pandemic-induced retail trading frenzy as a shock to retail trading intensity to address endogeneity concerns.

Chapter 3 (co-authored with Sangmin Oh and Jihong Song) shifts the focus to the US market and studies the heterogeneity in investor demand for sustainable equity investing. We measure firm-level sustainability across three dimensions: third-party environment scores, emissions, and green patents. Separately estimated institutional investor demands are sensitive to scores and emissions, but not to green patents. We then aggregate these heterogeneous demands in an equilibrium framework to draw implications for the effectiveness of sustainable investing: (i) price-elastic investors do not "undo" effects of sustainable investors, (ii) investor pressure for sustainability only weakly predicts future improvements in firm sustainability, and (iii) incorporating green patents into ESG ratings can be a valuable adjustment.

Acknowledgements

I thank my advisors Motohiro Yogo and Wei Xiong for their patience and guidance. As I recollect exchanges with them, I am again reminded that I will remain but a student. There were times when I was unequivocally convinced that I wanted an academic career; there were more times when I was convinced that I was not fit for one. If not for Moto and Wei, I would have been left teetering. On a separate note, thank you, Wei and Lin for the turkey.

I am much obliged to Markus Brunnermeier for his advice and backing throughout the program; he was always much too generous with his time and advice, even through his injury. Also, I gratefully acknowledge invaluable teaching and feedback from the faculty at Princeton. Special thanks to Yacine Aït-Sahalia, Caio Almeida, Moritz Lenel, Adrien Matray, Atif Mian, Jonathan Payne, Batchimeg Sambalaibat, and David Schoenherr in this regard.

I have also learned so much from my friends Sangmin Oh and Jihong Song; thank you two. Many thanks should also go to Simon Schmickler, Amora Elsaify, Donghwa Shin, Michael Junho Lee, Karsten Müller, Marcelo Medeiros, and other seminar participants at Princeton's Bendheim Center for Finance for crucial suggestions and comments. I owe much to my friends Philippe van der Beck, Goutham Gopalakrishna, Yucheng Yang, Shumiao Ouyang, Nicolas Hommel, Ziang Li, and Filippo Palomba in getting through the job market. I would like to thank Laura Hedden, Kelsey Richardson, and Jessica Roethel for their help in navigating life at Princeton.

While things started in my undergraduate years at University of Pennsylvania, I never got a chance to thank my teachers from back then. I express my gratitude to Christian Opp, Nikolai Roussanov, and Urban Jermann for their encouragement despite my shortcomings.

Life at Princeton would have been too dull without my roommates and buddies Francisco Cabezón, Seonghyeok Cho, Jinseok Lee, Yeongjun Ryu, and Seungwook So; thank you all.

In the end, all there is to acknowledge is that none of this matters but for my beloved. I acknowledge that my wife, my parents, my family and friends have been the reason. What was nebulous then but clear now is that I work and wrestle life to be a better person to them—marginally.

To Yeayeun, my *umma* and *appa*, and my friends.

Contents

Abstract . . . iii

Acknowledgements . . . iv

List of Tables . . . ix

List of Figures . . . xi

1 Do Trading Pauses Protect Retail Investors? 1

1.1 Introduction . . . 1

1.2 The Korean Stock Market . . . 6

1.2.1 Notable Aspects . . . 6

1.2.2 Korea Exchange (KRX) Data . . . 10

1.3 Ex-Post Effect of VI on Retail Trading . . . 14

1.3.1 Stylized Facts . . . 14

1.3.2 Empirical Results . . . 16

1.3.3 Mechanisms . . . 21

1.4 Ex-Ante Effect of VI Rules on Retail Trading . . . 23

1.4.1 Stylized Facts . . . 23

1.4.2 Empirical Results . . . 26

1.4.3 Trading Strategy . . . 28

1.5 Price Reversion and Retail Wealth Loss . . . 32

1.5.1 Price Reversion around Threshold Breaches . . . 32

1.5.2 Retail Wealth Losses . . . 33

1.5.3 Effect on Overall Market Conditions . . . 38

1.6 Conclusion . . . 43

Appendix to Chapter 1 45

A1 Data Description . . . 45

A2 Trading Strategy with a Tree-Based Classifier . . . 46
A3 Additional Figures and Tables . . . 50

2 Retail Trading Intensity and the Overnight-Intraday Return Gap 58
2.1 Introduction . . . 58
2.2 Model . . . 63
2.2.1 Model with One Risky Asset . . . 63
2.2.2 Extension to Multiple Assets . . . 65
2.3 Data and Methodology . . . 67
2.3.1 Data and Stylized Facts . . . 67
2.3.2 Empirical Design . . . 73
2.4 Empirical Results . . . 75
2.4.1 Retail and Institutional Flow . . . 75
2.4.2 Portfolio Sorts on Retail Investor-Related Variables . . . 77
2.4.3 Return Gap and Retail Volume Share . . . 79
2.5 Conclusion . . . 83

Appendix to Chapter 2 85
B1 Model Details . . . 85
B1.1 Proof of Proposition 1 . . . 85
B1.2 Proof of Proposition 2 . . . 86

3 Unpacking the Demand for Sustainable Equity Investing 88
3.1 Introduction . . . 88
3.2 Data and Stylized Facts . . . 94
3.2.1 Data . . . 94
3.2.2 Stylized Facts: Evidence from Valuation Regressions . . . 96
3.3 The Asset Demand System with Sustainability . . . 98
3.3.1 Setup and Notation . . . 98
3.3.2 Implementation . . . 100
3.4 Estimated Demand for Sustainability . . . 101
3.4.1 Summary Statistics of Estimated Demand Coefficients . . . 101
3.4.2 Within-Investor Demand or Across-Investor AUM Shift? . . . 103
3.4.3 Cross-sectional Patterns . . . 104

3.5 Investor Pressure for Sustainability and Real Effects . . . 106
3.5.1 Quantifying Investor Pressure . . . 106
3.5.2 Investor Pressure and Future Environmental Performance . . . 109
3.6 Counterfactuals and Asset Pricing Implications . . . 110
3.6.1 Impact of ESG-Agnostic Mandates . . . 110
3.6.2 Hypothetical ESG Ratings: Incorporating Green Innovation . . . 112
3.7 Conclusion . . . 113

Tables for Chapter 3 **115**

Figures for Chapter 3 **123**

Appendix to Chapter 3 **128**
C1 Data . . . 128
C1.1 MSCI Data . . . 128
C1.2 PatentsView Data . . . 129
C1.3 Portfolio Holdings, Asset Prices, and Characteristics . . . 129
C2 Additional Empirical Results . . . 130
C2.1 Relationship among Measures of Firm Sustainability . . . 130
C2.2 Robustness to Alternative Environmental Scores . . . 131
C2.3 Counterfactual: Within-Investor Demand Shift or Across-Investor AUM Shift 132
C3 Theory . . . 133
C3.1 Including Sustainability as a Characteristic . . . 133
C3.2 Derivation of Investor Pressure in Proposition 3 . . . 135
C4 Computing Counterfactuals . . . 136
C5 Additional Tables and Figures . . . 138

List of Tables

1.1 Summary Statistics of Listed Stocks . . . 12
1.2 Panel Regression Coefficient Estimates . . . 18
1.3 Trading Strategy Returns . . . 29
1.4 VI Prediction . . . 31
1.5 Retail Trading Losses around Static-Up VIs . . . 35
1.6 Retail Trading Volume and Losses on Extreme Return Days . . . 37
1.7 Diff-in-Diff Coefficient Estimates for Volatility . . . 41
1.8 Diff-in-Diff Coefficient Estimates for Liquidity . . . 42
1.9 Diff-in-Diff Coefficient Estimates for Volatility . . . 43
1.10 Diff-in-Diff Coefficient Estimates for Volatility . . . 43
1.A1.1 Available Variables from KRX . . . 45
1.A2.2 Features and Definitions . . . 47
1.A2.3 Grid Search for Hyperparameters . . . 49

2.1 Summary Statistics of Stock-Month Observations . . . 72
2.2 Correlation Matrix of the Main Variables . . . 72
2.3 Persistence of Morning Flows . . . 76
2.4 Persistence of Afternoon Flows . . . 76
2.5 Portfolio Returns for Univariate Sorts . . . 78
2.6 Fama-MacBeth Regressions . . . 79
2.7 Return Gap vs Retail Volume Share (Panel OLS) . . . 80
2.8 Return Gap vs Retail Volume Share (Diff-in-Diff) . . . 82

3.1 Summary Statistics . . . 115
3.2 Valuation Regressions . . . 116
3.3 Demand Coefficients: Summary Statistics . . . 117

3.4 Demand for Emissions Intensity: Time Trend Analysis . . . 118
3.5 Demand for Sustainability and Investor Characteristics . . . 119
3.6 Investor Pressure and Future Environmental Performance . . . 120
3.7 Counterfactual Exercise: Impact of ESG-Agnostic Mandates . . . 121
3.8 Counterfactual Exercise: Hypothetical ESG Ratings with Green Patents . . . 122
3.C5.1 Relationship Between Green Characteristics . . . 138
3.C5.2 Sustainable Demand and UNPRI Signatory Status: Cross-Sectional . . . 139
3.C5.3 Sustainable Demand and UNPRI Signatory Status: Time-Series . . . 140
3.C5.4 Investor Pressure: Summary Statistics . . . 141
3.C5.5 Investor Pressure and Future Environmental Performance: Longer Horizons . . . 142
3.C5.6 Counterfactual Exercise: Impact of ESG-Agnostic Mandates, Alternative Order . . . 143
3.C5.7 Stock Characteristics and Investment Universe Coverage . . . 144
3.C5.8 Valuation Regression with Alternative Environment Scores . . . 145
3.C5.9 Counterfactual Exercise: Demand vs. AUM Shifts . . . 146

List of Figures

1.1 Retail Trading Volume Share . . . 7
1.2 Retail Net Buy Patterns . . . 8
1.3 Description of VI Auction Mechanism . . . 9
1.4 Description of Static VI vs. Dynamic VI . . . 10
1.5 Number of Daily VI Occurrences by VI Type . . . 12
1.6 Retail Volume around Static VIs . . . 15
1.7 Retail Net Buy around Static VIs . . . 16
1.8 Interval Construction . . . 17
1.9 Retail Volume and Net Buy around 10% Breaches . . . 20
1.10 Features of Korean Mobile Trading Systems . . . 21
1.11 Retail Volume around 10% Breaches by Number of Previous Breaches . . . 22
1.12 Retail Activity in Common Stocks around Preferred Stocks' VIs . . . 23
1.13 Return Bin Construction . . . 24
1.14 Average Retail Net Buy by Return Bins . . . 25
1.15 Relative Ask/Bid Order Depth by Return Bin . . . 26
1.16 Difference in Ex-Ante Volume and Net Buy by Return Bins . . . 27
1.17 Prices around Static VIs . . . 33
1.18 Retail Trading Volume and Profits on Normal vs. Extreme Return Days . . . 36
1.19 Measures of Volatility and Rule Adoption . . . 40
1.A2.1 Log-Loss for Different Hyperparameter Values . . . 48
1.A2.2 Mean Decrease Accuracy of Features . . . 49
1.A3.3 Number of VIs by Size . . . 50
1.A3.4 Foreign Volume and Net Buy around Static VIs . . . 51
1.A3.5 Pension Volume and Net Buy around Static VIs . . . 52
1.A3.6 Prop Volume and Net Buy around Static VIs . . . 53

1.A3.7 Retail Volume and Net Buy around Dynamic-Up VIs 54
1.A3.8 Retail Volume and Net Buy around Dynamic-Down VIs 54
1.A3.9 Retail Volume and Net Buy around +5% Breaches 55
1.A3.10 Prices around $\pm$5% Breaches 55
1.A3.11 Retail Net Buy by Return Bin 56
1.A3.12 Market Conditions around Rule Adoption 57

2.1 Cumulative Overnight and Intraday Returns 59
2.2 Monthly Volume by Investor Type 61
2.3 Time Setting of the Model 63
2.4 Overnight and Intraday Returns 68
2.5 Intraday Net Buy Patterns 69
2.6 Intraday Volume Patterns 70
2.7 Relative Retail Trade Prices 71
2.8 Retail Shares Over Time by Initial RS Quintiles 74
2.9 Monthly Retail Volume by RS Quintile 74
2.10 Binscatter of Return Gap against Volume Share. 81

3.1 Valuation Regressions: Time-Series of Coefficients 123
3.2 Binned Scatterplots of Coefficients across Investors 124
3.3 Trends in Average Coefficients 125
3.4 Demand Coefficient for Emission: Role of Within-Investor Preference Shift 126
3.5 Time Series of Average Investor Pressure 127
3.6 Counterfactual: Impact of ESG-Agnostic Mandates 127
3.C5.1 Average Price Inelasticity Over Time 147
3.C5.2 Distribution of Environmental Scores for Oil & Gas and Banking Stocks 148
3.C5.3 Demand Coefficients for Environmental Score: Alternative Environment Score Constructions 149
3.C5.4 Demand Coefficients for Emission Intensity: Alternative Environment Score Constructions 150
3.C5.5 Counterfactual: Demand vs. AUM Shifts 151

Chapter 1

Do Trading Pauses Protect Retail Investors?*

1.1 Introduction

On March 9, 2020, when COVID-19 induced fears overtook global investors, circuit breakers at the US stock exchanges kicked in. There were three more days in that month during which such market-wide circuit breakers were triggered. Above and beyond the news that the S&P 500 Index fell by more than 7% within a single trading day—the threshold for a level 1 circuit breaker—the term *circuit breaker* conjures up an ominous emotion, and certainly makes for a good headline. Given the daily volatility of stocks, the economic news contents in the headlines (1) *S&P 500 Index fell by 6.5%* and (2) *S&P 500 Index triggered a level 1 circuit breaker* must not be substantially different. However, the salience of this regulatory measure makes the latter caption simply harder to ignore.

Both market-wide circuit breakers and stock-level trading pauses are triggered by either extremely positive or negative intraday returns, and are meant to serve as a speed bump to panic selling and manic buying. They aim to achieve this by providing market participants time to reassess the situation and react, thereby protecting investors from extreme price swings.[1] In this sense, salience is a feature rather than a bug. But what types of reactions are these interventions

* This research was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant. I am grateful to my advisors Motohiro Yogo, Wei Xiong, and Markus Brunnermeier for their guidance and support. I also thank Stella Park, Sangmin Oh, Jihong Song, Francisco Cabezón, Jonathan Payne, Simon Schmickler, Michael Junho Lee, Nicolas Hommel, Ziang Li, Philippe van der Beck, Goutham Gopalakrishna, Donghwa Shin, Karsten Müller, Batchimeg Sambalaibat, Caio Almeida, Moritz Lenel, Yacine Aït-Sahalia, Atif Mian, Jiwon Lee, Sohye Yoon, Amora Elsaify, and participants of the Princeton Finance Workshops for feedback.

[1] Deutsche Börse describes their trading pause rules as "protective mechanisms" and Korea Exchange states that trading pause rules are put in place to "prevent damages to investors."

soliciting? What if they stir up more animal spirit than liquidity provision? Should we be worried about any unintended consequences?

In this paper, I exploit high frequency, investor type-level trade flow data from the Korean stock market to show that one such unintended consequence is a consistent and systematic wealth transfer from retail investors to institutional investors. I come to this conclusion by documenting that (1) retail investors as a whole make bets on price continuation around trading pauses—net buying (selling) stocks that trigger a pause by appreciating (depreciating) in price—and that (2) prices revert after these trading pauses. Effectively, retail investors are net buyers at the peak around an upward trading pause and net sellers at the trough around a downward trading pause. On the other hand, foreign institutions absorb more than 80% of these net flows and emerge as the biggest winners through trades around trading pauses. This pattern is strikingly consistent across tens of thousands of trading pause events.

As a benchmark, I compute the transfer of wealth from retail investors to foreign institutions under the assumption that wealth is marked to market at the end of the trading day using the daily closing price of a stock. For instance, if retail investors net bought one share at \$100 around the pause and later this share closes at \$90, this would imply a daily loss of \$10 under this simplifying assumption. Back-of-the-envelope calculation suggests that in the first two years of rule adoption, retail investors as a whole transferred around USD 150 million of wealth to foreign institutions during stock-days on which trading pauses were triggered. Moreover, retail investors' heightened trading activity around a trading pause is also subject to brokerage fees and trade taxes that amount to 0.3% of total trading value.[2] These costs amount to an additional USD 550 million, making the estimated losses exceed USD 700 million. Indeed, these are not causal estimates because it may well be the case that retail losses could be similar on extreme return days—days during which prices move outside the $\pm 10\%$ band relative to opening prices—even without the trading pause rule. However, even after using such pre-rule extreme return days as reference points, the resulting estimates do not change significantly.

To be more specific about the setting, I study a stock-level, 2-minute trading pause rule—called the volatility interruption (VI) mechanism—that was introduced to the Korean stock market on June 15, 2015. Abstracting from a few details, VI is triggered when a security listed on the Korea Exchange (KRX) moves more than 10%, either up or down, relative to its opening price.[3] Once VI is triggered, continuous trading halts for 2 minutes and an auction phase begins. These 2-minute VI

[2] This is a conservative estimate using a relatively low brokerage trading fee. Losses are translated to USD terms using the approximate average KRW/USD exchange rate of 1,100 during the years 2014-2016.

[3] The exact rule is discussed in more detail in Section 1.2.1

auctions work in the same way as do opening and closing auctions in NYSE or NASDAQ: traders can add, modify, or cancel orders, these orders are accumulated to the order book, and they are cleared only at the end of the auction phase. The state of the aggregate order book is public so that anyone can observe the indicative clearing price and volume throughout the auction phase. As with opening and closing auctions, the goal is to concentrate liquidity and to promote price discovery.

The KRX data prove to be particularly well suited for several reasons. The first is that KRX provides exact identifiers for investor types, such as retail investors and foreign institutions, at the *individual trade level*. For every single executed order, I observe the investor types of the buyer and seller, on top of the usual variables like order type, order volume, and executed price.[4] The second is that the investor type-level trade flows are exhaustive. Earlier papers such as Barber and Odean (2008) and Hirshleifer et al. (2008) use retail flows from brokerage firms. More recent innovations by Boehmer et al. (2021) also enable researchers to impute retail flows from the US trade and quote (TAQ) data. However, neither of these methodologies capture the exhaustive flows of different investor types. Because KRX is the sole stock exchange on which all stock trades happen, I can make definitive statements about the aggregate flows between investor types. Lastly, the KRX data exist both for the periods before the rule adoption (pre-rule period) and after the rule adoption (post-rule period). This allows for the sampling of *pseudo-pauses*—a stock's breaching of the pause trigger limit, $\pm$10%, during the pre-rule period—which can serve as control events.[5]

With this detailed dataset, I first document that retail investors' trading volume increases disproportionately during the minutes immediately following the trading pauses. Earlier studies (Christie et al. 2002; Hautsch and Horvath 2019) have also found heightened volume following trading halts or pauses. I refine these findings by showing that more than 95% of this extra activity is coming from retail investors. Similar to an exercise in Seasholes and Wu (2007), which shows that attention-grabbing effects of halts are weaker when there are multiple contemporaneous halts, I demonstrate that the extra retail trading garnered by a trading pause decreases monotonically with the number of preceding pauses on the same day. This adds weight to the narrative that retail investors are reacting to the salience of this regulatory device.

In addition to these ex-post patterns, I also present notable changes in retail investors' ex-ante trading behavior near the 10% VI thresholds. Generally speaking, retail investors act as liquidity providers in the Korean market: net selling stocks whose prices have gone up and vice versa for

[4] While this is true in principle, I have trade-level data from June 2014 and May 2016. This is due to funding constraints. For periods extending to May 2018, I use 1-minute frequency data instead.

[5] The term and concept are adopted from Lee et al. (1994) and Hautsch and Horvath (2019).

stocks that depreciated.[6] This pattern holds regardless of whether I use contemporaneous—with respect to retail flows—returns or lagged returns, and whether I look at shorter horizons (hours) or longer horizons (days and weeks). This is unsurprising given that institutional flows are informed and persistent. During the pre-rule period, retail investors' tendency to net sell becomes monotonically stronger as a stock appreciates relative to its opening price. However, in the post-rule period, this monotone pattern breaks down, and retail investors suddenly become net buyers when a stock's price nears the $+10\%$ threshold—i.e., when a stock appreciates by 7, 8, or 9% relative to its opening price. Why do retail investors suddenly become buyers of such *expensive* stocks from being sellers? I attribute this to the anticipation of the ex-post trading pattern and overconfidence. If a trader expects to correctly foresee an upward VI and the ensuing late-arriving net buy orders, this trader has an incentive to buy. By backward induction, this creates a cascade of incentives to net buy when the probability of upward VI rises. I demonstrate the existence of this incentive by constructing a trading strategy that successfully implements this idea.

Notwithstanding these concerning trading patterns, circuit breaker-like rules are ubiquitous. This leads us to expect that trading pause rules introduce other benefits to overall market conditions that are not yet discussed. I exploit the arbitrariness of the pause-triggering threshold to identify the rule's effects on volatility and liquidity. The arbitrarily chosen 10% threshold is *tight* for a security with a high daily volatility whereas it is practically nonexistent for a security like Samsung Electronics' common stock. In the benchmark specification, stocks are sorted by return volatility in the past 60 trading days. Then, the top quintile is considered the treatment group while the bottom quintile is considered the control group. The outcome variables used are monthly averages of 5-minute realized volatility, high-low price range, bid-ask spread, order book depth, Amihud measure, and the number of $\pm 10\%$ breaches during a given month. I find statistically insignificant effects on measures of volatility and mixed effects on measures of liquidity. For robustness, I also sort by the number of extreme returns and daily high-low price range in the past 60 trading days, and find similar results. One concern is that more volatile stocks may be more sensitive to overall market conditions relative to less volatile stocks. To mitigate this concern, the main specification includes quintile-specific slopes to market volatility and market returns as controls.

Altogether, the results demonstrate how the interplay of behavioral biases (salience and extrapolation) and regulatory factors (trading pause rule) confounds this seemingly simple exchange rule. As epitomized by the Lucas critique (Lucas 1976), economists heed much attention to endogenous responses of market participants. I provide an example where failing to account for *behavioral*

[6] Barrot et al. (2016) find the same pattern in the French stock market.

responses to rule changes attenuates, and even reverses, the intended effects.

My results are closely related to the main narrative of Seasholes and Wu (2007) and Chen et al. (2019b): a group of naive traders, who are often small retail investors, flock to stocks that recently breaches an upward limit, and sophisticated traders who correctly anticipate this actively buy near the upper price limits only to unload positions to the latecomers. I further the understanding about what drives both the ex-post and ex-ante incentives of retail trading around pauses. In addition, the high resolution investor type-level data reveal that a significant amount of wealth transfer happens from retail investors to institutions, and also allow me to quantify this as well.

This paper also contributes to the literature on attention-induced trading by individual investors and its perils (Barber and Odean 2008; Engelberg et al. 2012; Barber et al. 2022b). Most papers in this area use performance measures related to hypothetical returns to measure trading losses. Because exhaustive retail flows are observed at a high frequency, this paper is able to compute the amount of wealth outflow directly. Similar to the findings in Barber et al. (2009), retail investors' performance suffer due to unnecessary trading, while foreign institutions reap most of the direct benefits in my setting as well. One takeaway from Barber et al. (2022b) is that how trading venues are designed and presented to retail investors matter for the trading outcome of retail investors. Results here echo this lesson.

Various other papers also study the intended and unintended effects of trading pauses. Studies on trading pause rules have faced challenges to identifying the effect of adopting these rules because exchanges around the world often introduce them in a non-staggered manner—i.e., to all stocks on the same day. A common workaround has been looking at volatility and liquidity around realized pauses. Studies have found that markets experience heightened activity around pauses, contrary to the exchanges' aim of cooling-off markets (Lee et al. 1994; Christie et al. 2002; Hautsch and Horvath 2019). However, this alternative approach faces the limitation that only *local* statements can be made. Markets may be more turbulent around realized pauses, but they may experience less pause-triggering price movements in the first place. I show that this is unlikely to be case through a new workaround that exploits the arbitrariness of the pause-triggering threshold.

Studies have also considered how the endogenous reaction of traders to circuit breaker-like rules can affect trading during normal times. A classic example is the "magnet effect" which refers to a situation where traders drive the price towards the limit when prices are close to the interruption thresholds (Subrahmanyam 1994; Cho et al. 2003; Yan Du et al. 2009). On downwards swings, for instance, leveraged buyers may unravel positions in expectation of inability to rebalance their portfolios, thereby adding to the selling pressure (Chen et al. 2019a). This paper's result on the

ex-ante trading behavior of retail investors suggests that they indeed behave in a way consistent with the magnet effect hypothesis.

1.2 The Korean Stock Market

1.2.1 Notable Aspects

Korea Exchange (KRX) is the sole exchange for all stocks, ETFs, and many derivatives in Korea. It consists of three divisions: the KOSPI market division, the KOSDAQ market division, and the derivatives market division. KOSPI houses ETFs and larger blue-chip stocks such as those of Samsung Electronics, Hyundai Motor, POSCO and LG Electronics, while KOSDAQ houses smaller technology stocks. There are 793 companies listed on KOSPI and 1,497 companies listed on KOSDAQ.[7]

The total market capitalization of stocks listed on KRX amounts to USD 1.67 trillion. This places Korea's public equity market in between those of Germany and Australia in terms of size. More interestingly, the average monthly trading volume is USD 277 trillion, trailing only the US, China, and Japan in this department. On top of its considerable size and activity, the Korean stock market possesses several institutional details that make it an attractive setting to study retail investor behavior.

Active Retail Participation. First, the Korean stock market is a retail-driven market. As shown in Figure 1.1, the retail share of monthly trading volume in KRW terms has hovered around 65% and even reached highs of 80% during the pandemic-induced retail trading surge. The same statistic in the US moved from around 15% to 20% during the same periods.[8] According to the Korean Securities Depository (KSD), a custodial service provider, retail investors account for 28% of total stock holdings. Consistent with Odean (1999), Korean retail investors trade much more intensively compared to institutional investors.

[7] All market statistics such as number of listed companies, market capitalization, or trading volume are as of September 2022.

[8] Source: Bloomberg Intelligence.

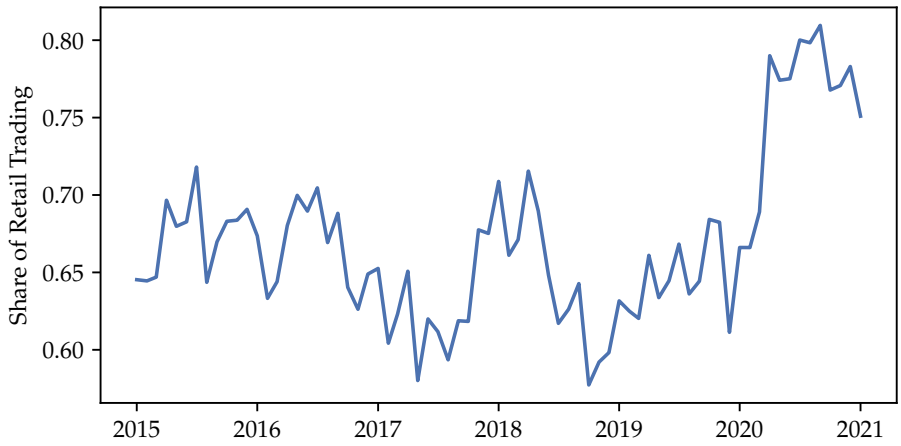

Figure 1.1: Retail Trading Volume Share. This figure plots the share of retail investors' daily trading volume relative to total trading volume in the Korean stock market. Both total trading volume and retail trading volume are measured in KRW terms and aggregated at the monthly frequency. The ratio of the sums are reported in the figure.

Known locally as "ants," Korean retail investors are both an object of media's ridicule and market's scrutiny, respectively due to their poor performance and sizable influence. As Kim and Kim (2022) finds, by analyzing approximately 20,000 individual brokerage accounts, Korean retail investors indeed demonstrate many of the well documented behavioral traits such as overconfidence (Scheinkman and Xiong 2003), disposition effect (Shefrin and Statman 1985), penchant for lottery-like stocks (Bali et al. 2011), and herding. Thus, interesting retail trading patterns come in both variety and quantity in the Korean stock market.

However, not all retail investors should be considered equal. The largest 0.5% of retail accounts, whose balances are larger than KRW 1 billion, constitute around half of total holdings amount. Similar to Chen et al. (2019b), Kim and Kim (2022) also find that larger accounts are less prone to the aforesaid behavioral biases. Given these facts, it is more reasonable to consider the aggregate retail trade flows as coming from a mix of *sophisticated retail investors* and *naive retail investors*.

Investor Type-Level Trade Flows. The most attractive aspect of the KRX data is that KRX provides identifiers for the investor type at the *individual trade-level* through their trades and quotes (TAQ) data, albeit at a significant price. The provided investor types are: (1) retail investors, (2) foreign institutions, (3) proprietary traders, (4) asset managers, (5) pension funds, (6) banks, (7) insurance companies, and (8) private equity. All investor types except retail investors and foreign institutions refer to domestic institution types.

At a more reasonable price, KRX also provides trade flows aggregated at the investor type level

at 1-minute frequencies. This is the data that most of the analyses in this paper will rely on. For example, I can observe how many shares of Samsung Electronics retail investors bought and sold as a group between 9:01 and 9:02 a.m. on March 15th, 2020. This allows me to compute, for instance, retail net buys (shares bought minus shares sold) or retail volume during any 1-minute interval within my sample period. The same applies to all other provided investor types.

This establishes two important advantages over the previously used data on retail flows: exhaustiveness and precision. Various studies using brokerage data (Hirshleifer et al. 2008, Luo et al. 2022) are subject to the constraint that they have to work with a representative sample of retail trades. In a recent study, Boehmer et al. (2021) provides an innovative way to identify retail orders from the US TAQ data using regulatory restrictions on price improvement. However, the method is still a conservative imputation that applies to marketable retail orders, rather than marketable and limit orders together. With the KRX data, I can safely assume that I am working with *the* retail flow.

Using such data, I can immediately document the following interesting facts, for instance: Figure 1.2a shows that retail investors were net buyers throughout the pandemic-induced trading frenzy and Figure 1.2b shows that they are net buyers in the morning.

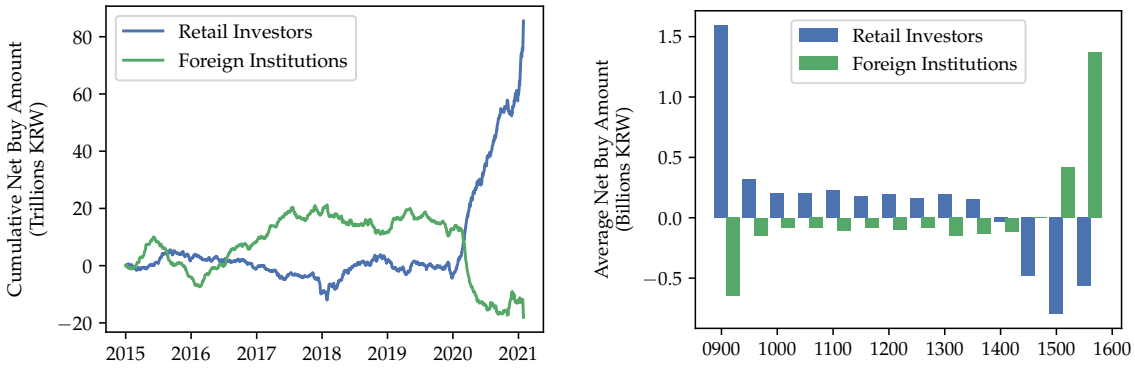

(a) Cumulative Retail Net Buy

(b) Retail Net Buy by Time of Day

Figure 1.2: Retail Net Buy Patterns. Panel (a) plots the cumulative net buy of retail investors (in green) and foreign institutions (in blue) from Jan 2015 to Mar 2021. Retail (foreign) net buy is measured each day by summing up retail buy volume in all stocks in KRW terms and subtracting the sell volume in all stocks. Then, the cumulative sum of these daily net buys are plotted. Panel (b) plots the average net buy volume during a given 30-mintue interval over all stock-day observations.

Volatility Interruption (VI) Mechanism. VI is a trading pause rule that the KRX adopted in June 15th, 2015. It is triggered when the to-be-executed trade price lies outside the "pre-defined price range," which will be discussed in the next paragraph. Once a VI is triggered, continuous trading is paused for 2 minutes and the "auction phase" begins. During the auction phase, orders can

be entered, changed, and cancelled, but will not be immediately executed. Each of these actions will be accumulated on the limit order book and will affect the prevailing supply and demand curves. At any moment during the auction phase, indicative price and volume—the intersection of supply and demand given the state of the order book at that moment—is broadcasted to market participants. At the end of the 2-minute auction phase, the stock emerges with a single equilibrium price and continues to trade normally. VI differs from a trading *halt* in this respect: a trading halt precludes trading for the rest of the trading day.[9]

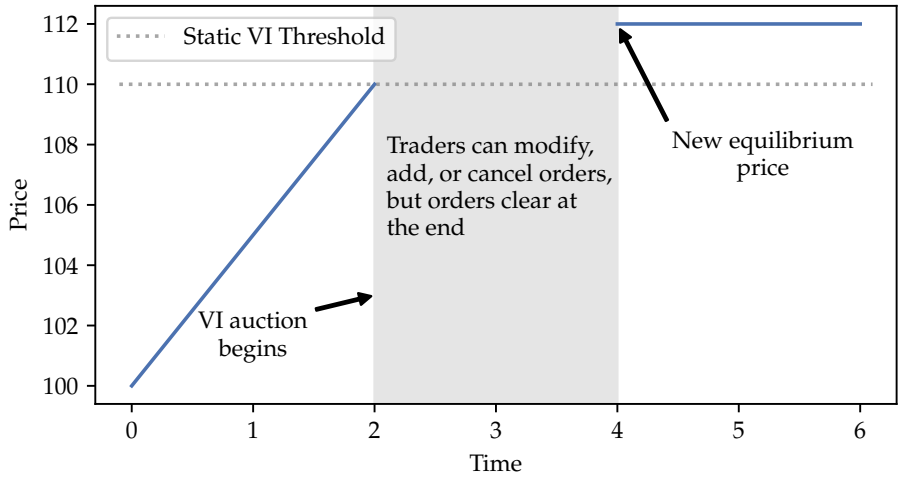

Figure 1.3: Description of VI Auction Mechanism. This figure illustrates what happens when a static VI is triggered. The blue line is an example path of a stock's price while the gray dotted line is the static VI threshold (10% relative to opening price). Once the price breaches this threshold, a 2-minute VI auction phase begins. At the end of the 2 minutes, the uniform price that clears the market becomes the new equilibrium price and continuous trading resumes. This pictorial description is adopted from Deutsche Börse's website and modified for my use.

There are two types of VIs at KRX: static and dynamic. A static VI is triggered when the to-be-executed trade price lies outside $\pm 10\%$ relative to the reference price which is defined as "the most recent auction price." This will be the price determined by the most recent of closing auction, opening auction, or VI auction. For most situations, it is safe to assume that the reference price is the opening price, unless a stock has already triggered a VI on the same day. A dynamic VI is triggered when the to-be-executed trade price lies outside $\pm 3\%$ relative to its previously executed price for constituents of the KOSPI 200 and $\pm 6\%$ for other securities. It is triggered only if a *single order* moves the potential trade price outside the price range. Thus, a static VI is targeted toward large cumulative intraday price moves while a dynamic VI targets situations that resemble fat-finger trades. Going forward, I refer to a static VI triggered by an upward price movement as a

[9] The description of the VI mechanism and Figures 1.3 and 1.4 rely heavily on the Deutsche Börse's descriptions of the same rule.

static-up VI. Other three cases are abbreviated similarly.

KRX adopted the dynamic VI rule on September 1, 2014, and the static VI rule on June 15, 2015. They were introduced as a way to mitigate the impact of widening the daily stock-level price limits, thresholds at which stocks halt trading for the entire day, from ±15% to ±30%. Since June 15, 2015, all stocks and ETFs are subject to daily price limits of ±30%, dynamic VI, and static VI rules. All rule changes applied to all listed stocks from the adoption date onwards, which poses challenges to identification.

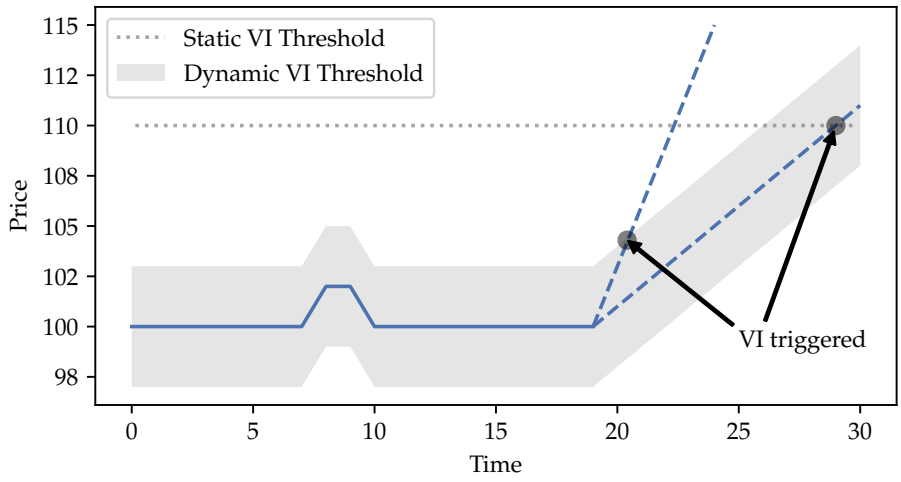

Figure 1.4: Description of Static VI vs. Dynamic VI. This figure illustrates when static and dynamic VIs are triggered. The blue line is an example path of a stock's price, the gray area is the dynamic VI threshold corridor (3% relative to the latest execution price), and the gray dotted line is the static VI threshold (10% relative to the opening price). The two dotted blue lines represent two possibilities: (1) price breaches the dynamic corridor by a *single trade* (left dot) or (2) price breaches the static threshold anytime during the trading day (right dot). In either case, a VI auction as described in Figure 1.3 ensues. This pictorial description is adopted from Deutsche Börse's website and modified for my use.

According to the KRX website, the stated goal of the VI rule is to "prevent damages to investors due to excessive price movements." Similarly, the Deutsche Börse classifies the VI rule as a "protective mechanism." These statements do not explicitly mention retail investors. However, the emphasis on cooling-off of sentiment and providing time to react suggest that non-sophisticated investors were likely targeted at the conception of VI-like rules.

1.2.2 Korea Exchange (KRX) Data

As of September 2022, KRX data covers 2,411 common and preferred stocks, 622 ETFs and various other securities listed on KOSDAQ and KOSPI markets. All data listed below can either be accessed directly on KRX's website or can be purchased from KRX, and span at least one year of the pre-rule

period.

1-Minute Frequency Trade Flow and Price Data. 1-minute frequency data that I work with runs from June 2014 to May 2018. This data can be considered redundant because they can be computed from trade-level data. I rely on the 1-minute frequency data for longer run analysis because trade-level data cost up to ten times the 1-minute frequency data. Trade flow data includes buy/sell volume and buy/sell value (KRW) at 1-minute frequency by investor type. For any 1-minute interval, I can compute the average price retail investors paid for Samsung Electronics common stock by dividing retail investors' buy value by their buy volume. Price related data include open, close, high, low prices, traded volume/value, liquidity measures, and order imbalance measures. All available variables are listed in Table 1.A1.1.

Trade-Level Data. The trade-level data that I work with runs from June 2014 to May 2016. For every trade that is executed on KRX, I observe the investor type of the buyer and seller, volume and price of buy/sell orders, order types (market, limit, stop limit, etc.), order time (microseconds), and snapshots of the order book (up to 10 levels) at trade execution moments. Notice that this is different from the quote-level data which records all changes to the order book. Other variables are listed in Table 1.A1.1.

Other Firm-Level Data. KRX also provides various daily statistics such as short volume, short interest, book-to-market ratio, dividend yield, and shares outstanding on its website. For regulatory reasons, foreign institutional holdings are also reported at a daily frequency, but holdings of other investor types are not directly reported. Another source of firm-level fundamentals is the Data Analysis, Retrieval and Transfer System (DART) which is the EDGAR counterpart in Korea.[10] Other variables are listed in Table 1.A1.1.

Summary Statistics. I begin with some relevant statistics that give us a better perspective. Figure 1.5 plots the number of static and dynamic VI occurrences per day. As one may expect, the number of VIs spike during turbulent days. During the COVID-19 crash the number of VIs exceed the number of listed stocks because a stock may hit multiple VIs in a single day. Because a 10% intraday price movement is large, one may expect that VIs are very rare events. As we see in Table 1.1, the average daily volatility for liquid stocks is around 3%. Using an inflated number of 5% suggests that a 10% intraday move should be two standard deviations under the assumption of normality.

[10] DART provides excellent automation and API access through OpenDART.

The median daily number of static VIs is 104—around 5% relative to the number of listings—which is in the ballpark of our inflated estimate, but it is true that the VI occurrences are more frequent than what a normal return distribution would suggest. This high median suggests that VIs are quite frequent during normal times as well and gives us ample statistical power.

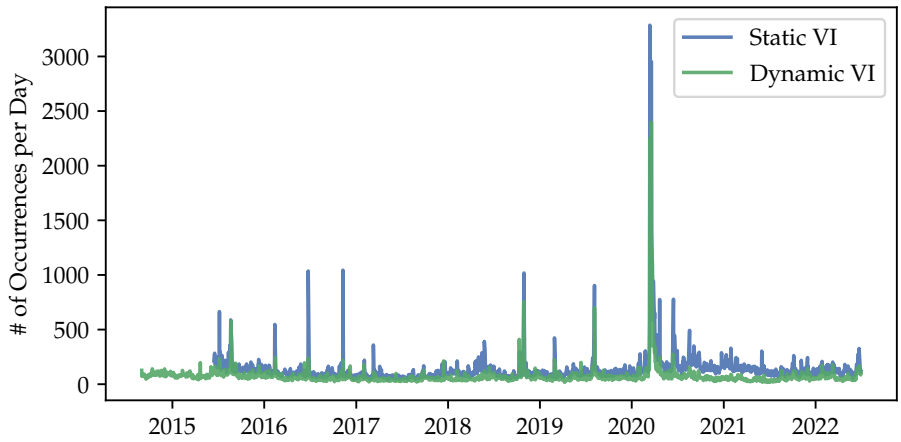

Figure 1.5: Number of Daily VI Occurrences by VI Type. This figure plots the number of daily VI occurrences by VI type. Dynamic VIs were introduced in September 2014 and static VIs in June 2015. Repeated VIs during the same day are counted separately for this figure.

As discussed at the beginning of this section, there are more than 2,000 listings on KRX, which is a large number relative to total market capitalization. This means that most of the stocks will be small and often illiquid. For this reason, I use a subset of 980 stocks (1) that traded every day from 2014 to 2018, (2) whose average market capitalization larger than 50 KRW billion, and (3) whose average daily trading value is larger than 0.5 KRW billion. Table 1.1 provides a summary of the stock-month level observations.

	Count	Mean	SD	1%	25%	50%	75%	99%
Market Cap. (KRW trillion)	84,585	1.55	10.85	0.03	0.10	0.20	0.64	21.70
Avg Daily Return Vol. (%)	84,585	2.99	1.33	1.12	2.04	2.67	3.62	7.26
Market Beta	84,585	0.93	0.68	−0.86	0.52	0.92	1.34	2.63
Price-Earnings Ratio	84,581	48.59	586.19	0.00	0.00	12.06	27.38	527.10
Price-Book Ratio	84,585	2.25	3.85	0.28	0.85	1.40	2.46	14.19
Foreign Holdings (%)	84,585	9.63	12.78	0.00	1.67	4.22	12.56	60.30

Table 1.1: Summary Statistics of Listed Stocks. This table reports the summary statistics for the subset of chosen stocks. Each observation is at the stock-month level. Market capitalization, share of foreign institutional holdings, P/E ratio, and P/B ratios are averaged within the same stock-month. Daily return volatility is the standard deviation of daily returns of a given stock during the month of interest. Market beta is the rolling 60 day market beta of a given stock as of the first day of the month of interest.

Definitions. *Net buy* by investor type is a quantity that will appear repeatedly. It is defined as the number of shares purchased by an investor type during an interval minus the number of shares sold by the same investor type. In order to make the quantity comparable across different stocks, net buy is often normalized by shares outstanding. In few of the analyses, in Figure 1.2a for example, I use KRW value instead of number of shares.

$$\text{NB}_{ijt} = \text{\# shares of stock } i \text{ bought minus sold by investor type } j \text{ during interval } t$$

$$\text{nb}_{ijt} = \frac{\text{NB}_{ijt}}{\text{\# shares outstanding}_{it}}$$

Because markets have to clear during any interval $t \to t+1$, we must have

$$\sum_j \text{NB}_{ijt} = 0 \text{ and } \sum_j \text{nb}_{ijt} = 0 \text{ for any stock } i \text{ and time } t.$$

For this reason, unconditionally we should expect to see zero net buys for any randomly chosen (i, j, t). This means that it is more often the case that negative retail net buy, which mechanically implies positive institutional net buy, over a period is associated with rising prices. Another related flow definition is *aggressive buy*, which is defined analogously as:

$$\text{AB}_{ijt} = \text{\# shares of stock } i \text{ bought by investor type } j \text{ during interval } t \text{ at } \textit{ask price}$$

$$\text{ab}_{ijt} = \frac{\text{AB}_{ijt}}{\text{\# shares outstanding}_{it}}$$

By definition, an aggressive buy cannot move the executed price downward. Of course, if bid and ask quotes adjust downwards as time goes by, even positive aggressive buy can be associated with negative price movements. However, this situation is rarer. It should be noted that aggressive buy includes both limit buy orders that cross the spread and market buy orders.

In the following sections, the majority of the results pertain to static VIs whose reference price is the opening price of the stock in question. For this reason, it is convenient to define *return from open* as the following:

$$\text{return from open}_{it} = \frac{P_{it}}{P^o_{it}} - 1$$

where P^o_{it} is the opening price of stock i on the day in which time t is included. This means that the first static VI occurs when return from open reaches $\pm 10\%$ for the first time.

1.3 Ex-Post Effect of VI on Retail Trading

The first thing I document is that VIs are salient events, and they induce retail investors to trade more intensively. I also document that retail investors bet on price continuation: they net buy on upward VIs and net sell on downward VIs. I start by looking at the retail trading volume and retail net buy around static VIs.

1.3.1 Stylized Facts

A stock that has risen or fallen by 10% relative to its opening price faces a selection problem. Such a big move may be driven by a big news event, liquidity shock, or sudden change in sentiment. Because I have data from periods before the VI rule was put in place, I use the observations from those periods as reference points. Similar to Lee et al. (1994) and Hautsch and Horvath (2019), I call these pre-rule threshold breaches *pseudo-pauses*. For periods before rule adoption, from June 2014 to May 2015, I sample the minutes around the first breach of $\pm 10\%$ for each stock. For periods after rule adoption, from July 2015 to December 2016, I sample the minutes around the first static VI.

Because both volume and net buys at 1-minute intervals are very noisy, I aggregate them at 10-minute intervals instead. One thing to note is that the interval from minute 0 to minute 10 starts with trades that are executed *after* the VI. Recall that once a VI is triggered, a 2-minute auction phase begins. This means that the first 10-minute interval completely includes the auction phase, and none of the trades that happen before the pause is triggered. If abnormal trading and net buying by retail investors happen during the auction period, they will be reflected in the interval from minute 0 to minute 10. I present the results for retail volume in Figure 1.6 and for net buy in Figure 1.7.

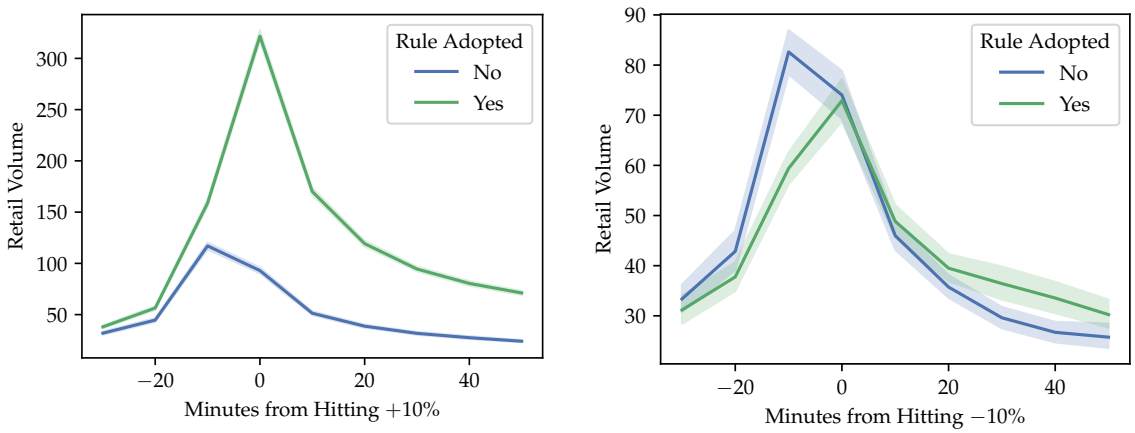

(a) Static-Up VI

(b) Static-Down VI

Figure 1.6: Retail Volume around Static VIs. This figure plots the average retail volume during 10-minute bins around VIs and pseudo-pauses. Retail trading volume in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, volume is expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). Panel (a) plots retail volume around static-up VIs. Panel (b) plots quantities around static-down VIs. 95% confidence intervals are constructed by bootstrapping.

Figure 1.6 shows spikes in volume in the first 10-minute interval that comes immediately after the breaching of 10% thresholds. We are interested in the height of the green line relative to that of the blue line at $x = 0$. Compared to a pre-rule 10% breach, retail volume is around 200 bps of shares outstanding higher. This is a large magnitude considering the brevity of the interval. Furthermore, activity stays heightened for the next hour or so.

The first thing to note is the difference between the upward breach and the downward breach. Retail investors appear to react much more to a static-up VI compared to a static-down VI. This may be due to the fact that there is a large regulatory hurdle for retail investors to short. Retail investors can always make attention-induced buys, but cannot make attention-induced sells without owning the stock.

The second interesting pattern is that volume peaks immediately *before* the breaching of 10% when VI was not yet introduced, and it peaks immediately *after* the breach once VI is in place. A threshold breach happens when an extreme price movement happens. This means a sudden large activity is likely to precede the threshold breaches. Accordingly, we see high volume at $x = -10$ for all cases (up-down and pre-post). The main difference in Figure 1.6a is that volume is pushed up even further following a static-up VI, while it recedes after the 10% breach in a pseudo-pause. Concerns about reverse causality will be discussed further later in the section.

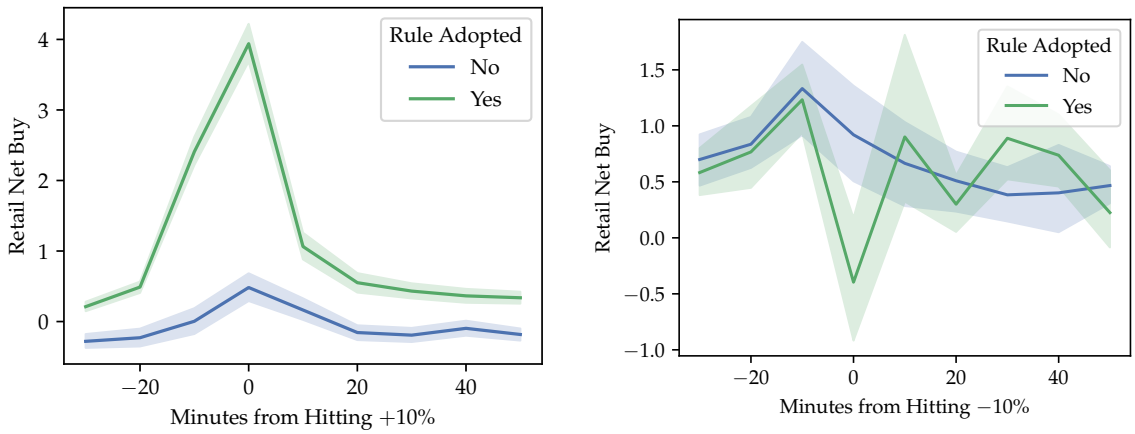

(a) Static-Up VI

(b) Static-Down VI

Figure 1.7: Retail Net Buy around Static VIs. This figure plots the average net buy during 10-minute bins around VIs and pseudo-pauses. Retail net buy is computed as retail buy volume minus retail sell volume. Net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, net buy is expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). Panel (a) plots retail volume around static-up VIs. Panel (b) plots quantities around static-down VIs. 95% confidence intervals are constructed by bootstrapping.

Figure 1.7 repeats the same exercise with retail net buys. Again the pattern is cleaner with upward breaches. Panel (a) shows that that retail investors are the net buyers of a stock around the breaching of 10%. In other words, more retail investors are choosing to buy a stock that has significantly risen in price from the opening of the market. While the absolute magnitude of the spike is small at around 4 bps of total shares outstanding, this is *net* retail buy during a 10-minute interval; unconditionally, this number should be zero.

While clear patterns appear in this unconditional averaging exercise, it is necessary to account for temporal patterns and other stock-level controls. We will see in the following section that adding controls do not affect the pattern.

1.3.2 Empirical Results

Panel OLS. I start by using a simple panel ordinary least squares (OLS) regression to confirm the above patterns. For each stock, I first split the time series into 10-minute bins so that each observation (i, t) is stock i in 10-minute bin t. To clarify, t is not the time of the day, but rather the t-th 10-minute interval during the entire time-series. For days during which a static-up VI occurs, the 10-minute bins are constructed so that the moment of VI occurrence lies at the beginning of a 10-minute bin. Figure 1.8 illustrates an example of this binning. If VI occurs at 9:36 a.m., trade flows following the VI will be included in the interval 9:36 $\rightarrow$ 9:46 a.m. and flows immediately

preceding the VI will be included in the interval 9:26 → 9:36 a.m. The minutes preceding 9:06 a.m. are discarded because 9:00 a.m. is the opening time.

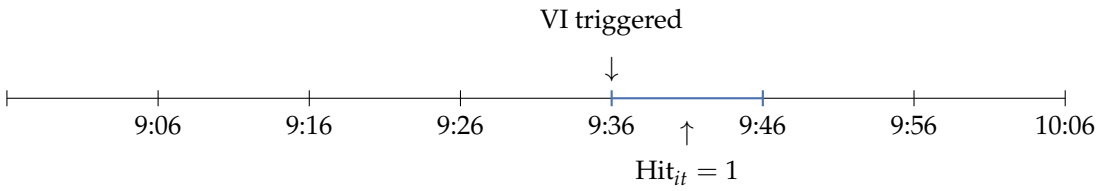

Figure 1.8: Interval Construction. This figure illustrates how 10-minute intervals are constructed from the trade flow data. The figure is an example in which ±10% breach happens at 9:36 a.m. Volume and net buy are aggregated at 10-minute intervals around this breach. Later, the first interval following the breach is assigned $\text{Hit}_{it} = 1$ and other intervals are assigned $\text{Hit}_{it} = 0$.

A 10% threshold breach in the pre-rule period, a pseudo-pause, is treated in the same manner. Once this bin assignment is done, the following dummy variables are assigned:

$$\text{Hit}_{it} = 1 \text{ if stock } i \text{ breached +10\% at } t$$

$$\text{Post}_{it} = 1 \text{ if interval } t \text{ is part of the post-rule period.}$$

In a panel regression, the coefficient on Hit_{it} captures the effect of breaching the 10% threshold, while the coefficient on $\text{Hit}_{it} \times \text{Post}_{it}$ captures the effect of experiencing a VI in addition to the effect of a threshold return breach. The panel OLS specification is as follows:

$$y_{it} = \gamma_1 \cdot \text{Hit}_{it} + \gamma_2 \cdot \text{Post}_{it} + \beta \cdot \text{Hit}_{it} \times \text{Post}_{it} + \Gamma' \mathbf{X}_{it} + \text{Hour FE} + \varepsilon_{it}. \tag{1.1}$$

The outcome variables used are total volume, retail volume, foreign institution volume, retail net buy, and foreign institution net buy of stock i during an interval t. The vector $\mathbf{X}_{it}$ includes control variables such as log market capitalization, average market beta in the past 60 trading days, average Amihud measure in the past 60 trading days, return in the past 20 days, and return in the past 200 days. Lagged flows are also included as controls when the dependent variables are flows because they tend to be persistent.

The coefficient of interest is β. Columns (1), (2), and (3) of Table 1.2 show that trading volume is significantly higher during the first 10-minute after following a static-up VI even relative to a pre-rule pseudo-pause. What is also notable is that almost all the effect is accounted for by rise in retail volume. The magnitude 183.9 bps for retail volume lines up well with the results of the previous subsection.

Columns (4) and (5) show that retail investors are net buyers during this first interval and the foreign institutions are handling most of this net buys. What is also interesting is that foreign institutions' rise in volume and rise in magnitude of net buy are almost the same: 1.14 bps vs. 0.9 bps. This means that foreign institutions are determined in their trade direction, and are mostly selling without buying. The magnitudes for both net buy regressions are small because the flows are aggregated at the investor type level.

	Total Vlm	Retail Vlm	Foreign Vlm	Retail NB	Foreign NB
	(1)	(2)	(3)	(4)	(5)
Hit 10%	16.0***	15.6***	0.22***	0.19***	−0.095***
	(0.21)	(0.21)	(0.01)	(0.03)	(0.02)
Hit 10% × Post	185.0***	183.9***	1.14***	1.87***	−0.90***
	(0.27)	(0.27)	(0.01)	(0.03)	(0.02)
Lagged Vlm	0.62***	0.61***	0.0057***		
	(0.00)	(0.00)	(0.00)		
Lagged Retail NB				0.33***	
				(0.00)	
Lagged Foreign NB					0.30***
					(0.00)
Hour FE	✓	✓	✓	–	–
Industry × Interval FE	–	–	–	✓	✓
Stock Controls	✓	✓	✓	✓	✓
Within R^2	.447	.447	.078	.104	.085
Observations	25,475,118	25,475,118	25,475,118	23,485,541	23,485,541

Table 1.2: Panel Regression Coefficient Estimates. This table reports the coefficient estimates of the panel ordinary least squares (OLS) regression specification in equation (1.1). The different columns correspond to different dependent variables. *, **, and *** represent significance at the 10%, 5%, and 1% level, respectively. Standard errors are clustered at the industry by interval level, and are reported in parentheses.

Challenges to Identification. The fundamental problem is that returns and trade flows are endogenously determined. Although the regression is specified so that the timing of returns, which determines the independent variable (VI occurrence or 10% breach), precedes that of the dependent variable (volumes and flows), there are still lingering concerns. I discuss three possibilities: news about fundamentals, liquidity shocks, and persistence in retail sentiment.

Suppose there is good news about the stock between the day's opening and the moment of static-up VI. Unless prices and flows adjust instantaneously, the prolonged effect of news will affect both returns from open and retail flow throughout the day. Liquidity shocks at large institutions work in the same way: if an institution is forced to make directional trades in a prolonged manner, this meta-order would affect both returns and retail flows. Persistence in retail sentiment is also similar: retail investors become excited about a stock, starts to put in aggressive buy orders, returns breach the threshold due to these orders, and excitement persists after the threshold breach.

The available remedies are (1) conditioning on prices rising by 10% relative to open and (2) controlling for preceding trade flows by different investors. Adding variables that precede the event, such as abnormal volume and liquidity measures, will further help control for news and liquidity shocks. However, we will see in the next subsection that the most concerning problem is retail sentiment *very close to* the VI thresholds.

Difference-in-Differences. In order to tackle the identification problem, I start by adapting the diff-in-diff setup proposed by Hautsch and Horvath (2019). The ideal experiment would be comparing two identical stocks that are under the influence of same news, liquidity shocks, and taste shocks, and only one stock randomly triggers a static VI. A reasonable alternative would be comparing two similar stocks that have both increased by 10%, while letting one stock randomly trigger a static VI. As does Hautsch and Horvath (2019), I face the same problem that the VI rules were introduced to all stocks with the same threshold. Thus, in this section I settle with the viable alternative of comparing a VI event with a pseudo-pause.

I first gather the static-up VI events for each stock i. Then, I keep the VI event if it is the first VI event of the day, then index it by $h = 1, \ldots, H$. There are 14,484 static-up VIs that satisfy this condition, so $H = 14,484$. For each VI event h, 20 minutes before and 60 minutes after the VI are sampled, and 10-minute bins are constructed. These bins are assigned event time τ where $\tau = 0$ includes the 10 minutes immediately following the VI. Analogously, the first breaching of 10% is sampled from the pre-rule period and are given indices $h = H + 1, \ldots, H + K$. There are 9,089 such breaches so that $K = 9,089$. 20 minutes before and 60 minutes after these breaches are sampled in the same way.

Samples from actual VIs are considered treated so that $\text{Treated}_h = 1$ if $h \in \{1, \ldots, H\}$. Time is defined in terms of event time. The first difference is the difference between intervals with $\tau \geq 0$ and the interval $\tau = -2$. The second difference is between the VI and pseudo-pause events. Once we denote by D_τ the dummy variable for event time interval τ, we can write the diff-in-diff

regression specification as follows:

$$y_{i\tau h} = \gamma \cdot \text{Treated}_h + \sum_{\tau=-1}^{5} \gamma_\tau \cdot D_\tau + \sum_{\tau=-1}^{5} \beta_\tau \cdot D_\tau \times \text{Treated}_h + \Gamma' \mathbf{X}_h + \alpha_i + \varepsilon_{i\tau h} \tag{1.2}$$

where $y_{i\tau h}$ is the outcome variable for stock i in event time τ for sample h, and α_i is the stock fixed effect. The outcome variable is either retail volume or net buy in each of the 10-minute interval τ. The controls are the same as in specification (1.1) and the coefficient of interest are β_τ's with $\tau \geq 0$. While γ_0 captures the effect of breaching the 10% threshold on the outcome variable, β_0 captures the additional effect of triggering a static-up VI. The coefficients β_τ's are plotted in Figure 1.9.

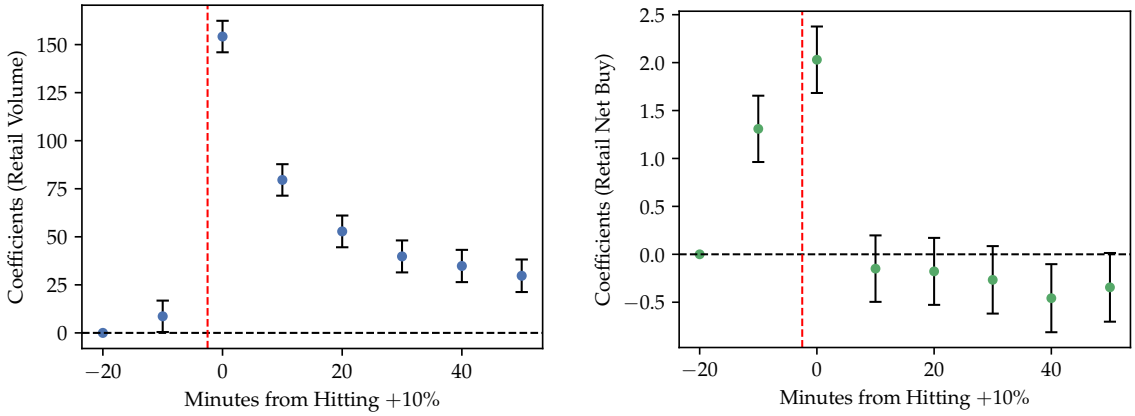

(a) Retail Volume

(b) Retail Net Buy

Figure 1.9: Retail Volume and Net Buy around 10% Breaches. This figure plots the estimates of the coefficients β_τ from the diff-in-diff regression specification in equation (1.2). 95% confidence intervals are displayed.

We see that the coefficient β_0 is positive and significant in both panels of Figure 1.9. This is consistent with previous subsections: retail investors increase trading abnormally, and they net buy stocks that just experienced a static-up VI, effectively betting on price continuation. The magnitudes of the coefficients are also largely unchanged relative to those of the panel OLS specification.

At the same time, we can observe a quite aggressive pre-trend in Figure 1.9b. This raises doubts about the necessary parallel trends assumption. One noticeable fact is that the pre-trend appears very abruptly around 10 minutes prior to the breaching of the threshold. This is a reason to believe that a sudden shift in retail sentiment is the more likely concern rather than news about fundamentals or institutional liquidity shocks. A fundamental news shock that is consistent with a static-up VI is *good news*. At the same time, it is reasonable to think that positive (negative) institutional (retail) net buy will be associated with this shock. This is not consistent with the pre-trend we observe. An explanation based on institutional liquidity shocks face this same issue. It also faces

another issue that institutions are more likely to smooth their meta-orders throughout the day or week, while the pre-trend is consistent with a more sudden move.

Hautsch and Horvath (2019) attribute a similar pre-trend to the magnet effect. However, it is unnatural to consider the coefficient β_{-1} as the causal effect of an imminent trading pause that is yet to realize. The more natural interpretation would be: the fact that *a trading pause rule is in place* causes endogenous ex-ante reaction by traders during normal times and the *realization of trading pause* causes additional reaction by traders ex-post.

1.3.3 Mechanisms

There are reasons to suspect that the heightened retail activity following VIs are due to the salience of the event among retail investors. As Figure 1.10 demonstrates, retail investors can get various alerts and filtered lists of stocks contingent on a stock's VI status through their mobile or desktop trading applications. It is also possible to filter on the proximity of the stock's price relative to the VI thresholds, which is an aspect more relevant to Section 1.4.

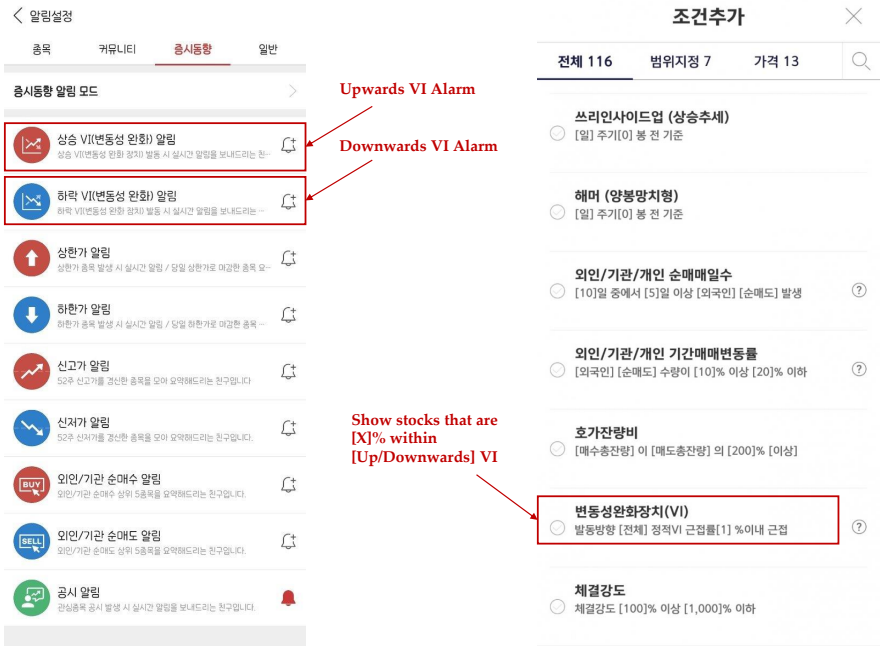

Figure 1.10: Features of Korean Mobile Trading Systems. These figures are screenshots taken from the mobile trading app of a popular Korean brokerage firm Kiwoom Securities. The figure on the left shows that a user can set an alarm when a VI is triggered for a chosen stock. The figure on the right shows that a user can put a filter to show only stocks whose prices are X% within a static-up or static-down VI.

Multiple VIs and Dissipation of Attention. In their study of trading halts in the Chinese market, Seasholes and Wu (2007) found a similar pattern of attention-induced trading for stocks that breach

the upward price limits. Furthermore, they found that the effect becomes attenuated as the number of contemporaneous halts increases. If retail investors are limited in attention and trading capacity, this should also hold in my setting. I run another regression as specified in equation (1.2), but within different subsets of the pauses and pseudo-pauses. For instance, the blue dots in Figure 1.11 present the resulting coefficient estimates by using the VIs and pseudo-pauses that were the n-th event of the day, where $1 \leq n \leq 15$.

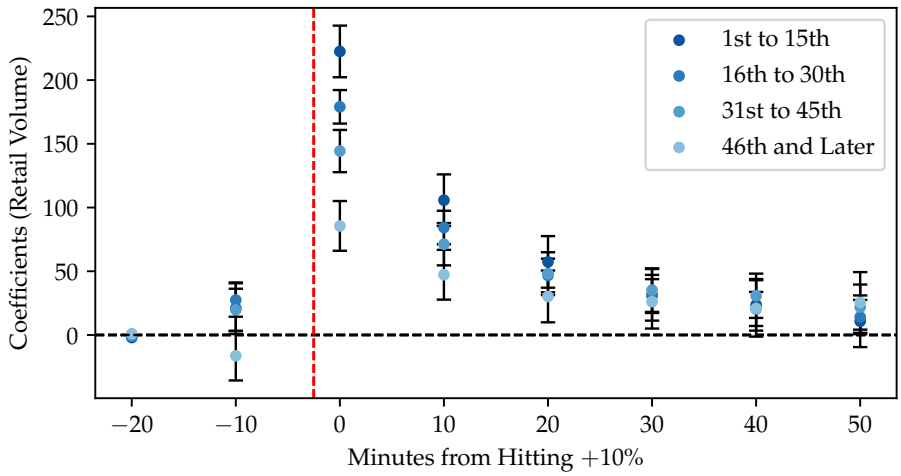

Figure 1.11: Retail Volume around 10% Breaches by Number of Previous Breaches. This figure repeats the plot in Figure 1.9a using four different subsets of the data. Both VIs and pseudo-pauses are ordered within the same day. The first group includes VIs and pseudo-pauses that are between the 1st and 15th to occur during their day of occurrence. Similarly, bins are assigned in increments of 15 occurrences. 95% confidence intervals are displayed.

In Figure 1.11, we see a clear monotone decrease in the coefficient estimates as we move onto subsets with late-coming VIs. This pattern supports the attention-based explanation. To address the concern that later VIs come at later times of the day, during which retail activity subsides, controls include time of the day dummies.

Preferred-Common Stock Pairs. To partially address the concern about reverse causality, I focus on a special case: preferred and common stock pairs. In Korea, around 100 established firms including Samsung, LG, and Hyundai, have floated preferred stocks on top of their common stocks. While the details differ slightly, preferred stocks offer a guarantee to pay extra dividends relative to their common stock counterparts at the cost of giving up voting rights. In theory, they should be subject to the same news about fundamentals and sentiment about the firm compared to their common stock counterparts. One notable aspect is that these preferred stocks have a much lower float, making them much more volatile relative to the common stocks. This means that there are

many cases where a preferred stock breaches the 10% threshold while the common stock does not. I gather the *trades for the common stock* around *breaches of the preferred stocks*, for the pre- and post-rule periods, then run a diff-in-diff with these events. Because the common stock itself did not breach the 10% threshold, this exercise should be free from a taste shock causing the 10% breach and should only capture the attention spillover. Figure 1.12 shows the coefficient estimates, and we see effects for the retail volume, but lack power for retail net buy.

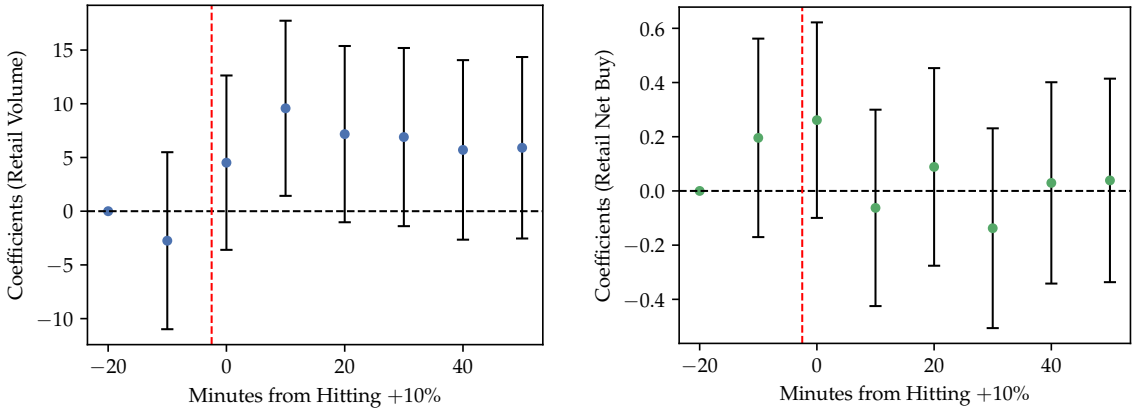

(a) Retail Volume

(b) Retail Net Buy

Figure 1.12: Retail Activity in Common Stocks around Preferred Stocks' VIs. This figure plots the estimates of the coefficients β_τ from the diff-in-diff regression specification in equation (1.2). 95% confidence intervals are displayed.

1.4 Ex-Ante Effect of VI Rules on Retail Trading

While the pre-trend in Section 1.3.2 poses difficulty in terms of identifying the causal effect of VI occurrence on retail trading, it is also what make the phenomenon more interesting. As the probability of triggering a static VI increases, traders appear to anticipate this possibility and alter their behavior (Chen et al. 2019a). Thus, we should expect causal effect of rule adoption on ex-ante trading behavior as well.

1.4.1 Stylized Facts

I begin by presenting a drastic change in retail investor net buy behavior. First, I divide each stock i's time series on day d into 1-minute bins. Before moving on, I drop stock-minute observations that come after breaching of 10% for the pre-rule period and a static VI hit for the post-rule period. Essentially, all observations that are affected by the ex-post effects of threshold breaches are discarded. For each 1-minute interval $t \to t+1$, I use the price at the *beginning* of the interval to

compute the return from that day's opening price:

$$r_{idt} = \frac{P_{idt}}{P^o_{id}} - 1$$

where P_{idt} is the price of stock i at the beginning of the interval $t \to t+1$ and P^o_{id} is the opening price of stock i on day d. Albeit imperfectly, this mitigates the concern that net trades in fact moved the prices to reach r_{idt}. I also compute the retail net buy during this same 1-minute interval, and call it nb_{idt}. If we see a high nb_{idt} associated with a high r_{idt}, this tells us that retail investors are net buying stocks that are expensive, or have appreciated a lot relative to the opening prices.

Mapping this to an illustration in Figure 1.13: if t is 9:05 a.m., r_{idt} will be 1.2% and nb_{it} will be the retail net buy during 9:05 and 9:06 a.m. Each stock-interval is assigned a return bin using the computed r_{idt}. The example interval in Figure 1.13 will be assigned to the [0%, 2%] bin if the bins were spaced out by 2%.

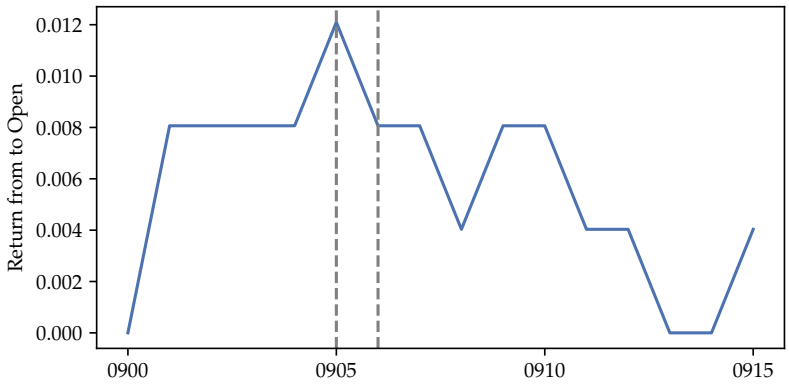

Figure 1.13: Return Bin Construction. This figure illustrates how 1-minute intervals are constructed and are assigned to return bins. The blue line represents a sample price path of a stock in terms of return from open. Return from open is defined as the return of the opening price of a 1-minute interval relative to the day's opening price.

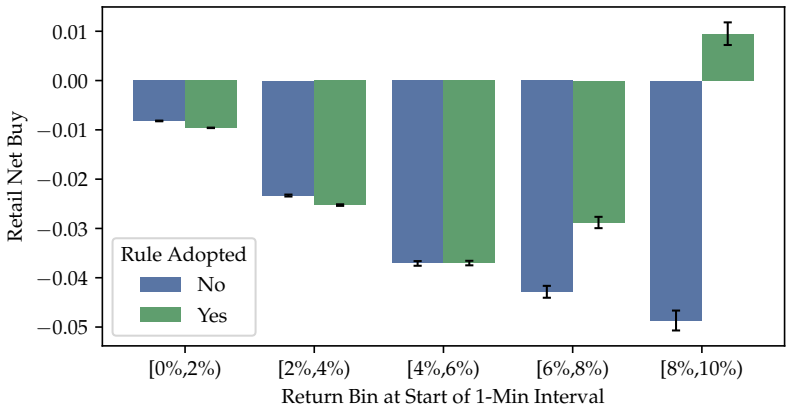

Figure 1.14: Average Retail Net Buy by Return Bins. This figure plots the average 1-minute net buy conditional on the return from open at the beginning of the 1-minute interval. The averages are computed for both pre-rule (blue bars) and post-rule (green bars) periods. 95% confidence intervals are constructed by bootstrapping.

For each return bin, the average retail net buy for the 1-minute intervals that fall in the return bin is computed. The results are shown in Figure 1.14. At medium to long horizons, retail investors provide liquidity to institutions. At the same time, institutions often make informed trades and do so by spreading out their orders. These in conjunction mean that on a day when institutions make informed buys, thereby pushing up prices, retail investors will end up net buying appreciated stocks. This narrative fits the pattern of the blue bars in Figure 1.14. In addition, disposition effect and mental accounting of retail investors are also consistent with the net selling of appreciated stocks.

What is surprising is the pattern of the green bars. After the VI rule adoption, they completely reverse their behavior near the static-up VIs and become net buyers of stocks that appreciated more than 8% intraday. Because all intervals that come after the breaching of $\pm 10\%$ were discarded, these net buy behaviors were affected by anticipation of possible VIs, but not by effects of realized VIs. Figure 1.A3.11 plots the same quantities, but includes other bins.

I also look at the order book depth on the ask and bid sides conditional on the return bin of the current traded price. For these figures, I use a shorter sub-period, Jun 2014–May 2016, for which trade-level data is available. Here I find that Figure 1.15a shows the time-weighted average of order book depth on the ask side. Order book depth on either (bid or ask) side is computed by summing up the outstanding volume at the first three levels of the order book. It is then normalized by the average depth on both the ask and bid sides over the past 20 trading days. Thus, a value of 0.5 should be considered a *normal* relative depth. Time-weighted averages are computed conditional on the return bin of the latest traded price. This procedure is repeated separately for the pre-rule

and post-rule periods.

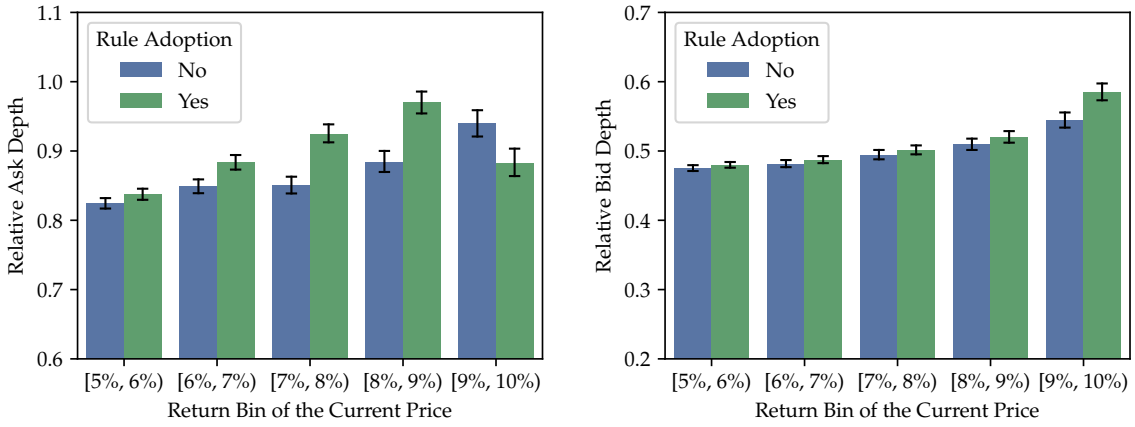

(a) Ask Depth

(b) Bid Depth

Figure 1.15: Relative Ask/Bid Order Depth by Return Bin. This figure plots the time-weighted average depth on the ask and bid sides conditional on the return bin of the current price. Ask (bid) depth is defined as the number of outstanding limit orders at the first three levels of the order book, normalized by the past 20 trading day average depth at both bid and ask combined. The heights of the blue (green) bars correspond to the pre-rule (post-rule) period averages. 95% confidence bands are constructed by bootstrapping.

In Figure 1.15a, we see that the ask depth conditional on the price being very close to the upper threshold is lower. One possibility is that traders postpone selling an appreciated asset when a VI is imminent. On the other hand, we see in Figure 1.15b that buying pressure near the threshold is higher in the post-rule period. These patterns are not consistent with attracting extra liquidity during extreme price upswings.

1.4.2 Empirical Results

I now test whether the patterns survive with various control variables. Also, the binning approach described in Figure 1.13 is an approximation. It works under the presumption that prices will lie in the same bin during the ensuing 1-minute interval. For instance, if prices start at 7.9% at the beginning of the minute and most of the trades in the minute take place around 9%, these trades will be mistakenly assigned to the [6%, 8%] return bin.

To get around this issue, I try a similar procedure with the trade-level data, albeit with a shorter sample period.[11] I follow the same procedure as before at the trade level. First, I discard the observations that come after a $\pm 10\%$ breach, then I assign return bins to each executed trade prices. The return bins are spaced by 1% instead for this analysis. For each stock-day, all retail buys and sells are summed up within the same bin, and the sums are normalized by shares outstanding.

[11] Trade-level data spans June 2014–May 2016, while 1-minute frequency data continues to May 2018.

Consider a scenario where retail investors bought a total of 100 shares of stock i during day d at an executed price (in terms of return from open) of 1.5% and sold a total of 100 shares at 2.5%. Suppose that there are 10,000 shares of stock i outstanding. Then, the retail buy volume for the bin [1%, 2%] is 100 bps, the retail sell volume of the bin [2%, 3%] is 100 bps, and the net buy for the two bins are $+100$ bps and -100 bps respectively. In such a way, stock-day-bin observations are constructed and used for the following regression:

$$y_{idb} = \gamma \cdot \text{Post}_{id} + \sum_{\substack{b=-9 \\ b\neq 0}}^{9} \gamma_b \cdot D_b + \sum_{\substack{b=-9 \\ b\neq 0}}^{9} \beta_b \cdot D_b \times \text{Post}_{id} + \Gamma' \mathbf{X}_{id} + \alpha_i + \varepsilon_{idb} \tag{1.3}$$

where the outcome variables are retail volume and retail net buy. Control variables are the same as the specification in equation (1.1).

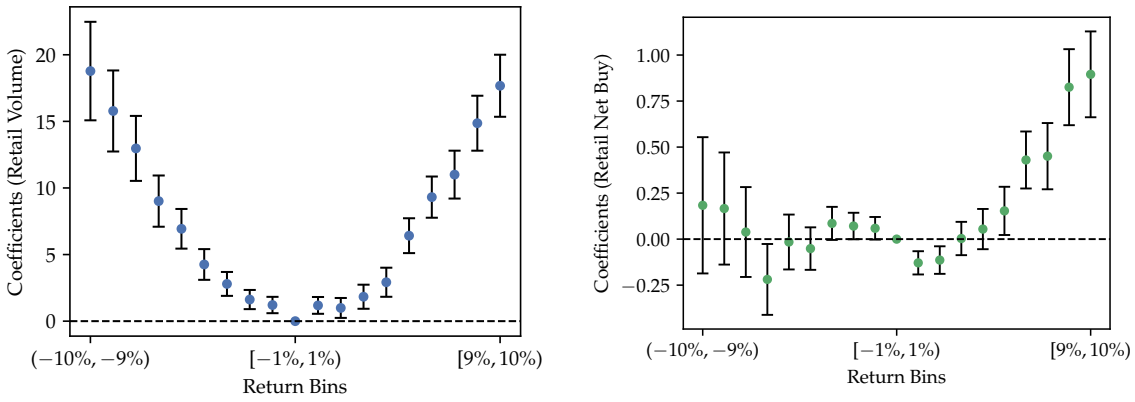

(a) Retail Volume

(b) Retail Net Buy

Figure 1.16: Difference in Ex-Ante Volume and Net Buy by Return Bins. This figure plots the estimates of the coefficients β_b from the diff-in-diff regression specification in equation (1.3). The outcome variable for panel (a) is retail volume and for panel (b) is retail net buy. 95% confidence intervals are displayed.

The resulting coefficient estimates for β_b's are shown in Figure 1.16. The scatter and confidence band at $x = b$ correspond to the bin $[b\%, (b+1)\%]$. Figure 1.16a shows that retail trading becomes more active around extreme prices. The regression controls for the effect of temporal changes through the Post_{id} term. We see that effects do not exist for moderate prices, but only for more extreme prices that are nearer the VI thresholds. Figure 1.16a captures the pattern documented in Figure 1.14: retail investors' tendency to net buy a stock increases when prices are near the static-up VI threshold. Patterns are similar if the trades are further normalized by duration of the prevailing trade prices, so that the outcome variables are in terms of intensity or rate of trades.

This pattern suggests a possible explanation of the pre-trend observed in Figure 1.9b. When we

sample the static-up VIs ex-post, the price path is bound to pass through 8%, 9%, then 10%. The pattern just presented suggests that when prices are around these levels, retail investors are already net buying extensively and this will show up as a pre-trend in Figure 1.9b. This observation by no means extricates the reverse causality concern that aggressive retail buying near the threshold has a (reverse) causal effect on the realization of static-up VIs.

1.4.3 Trading Strategy

Cutoff Strategy. The presented explanation is that retail investors more aggressively buy a stock in anticipation of a static-up VI. Testing this statement directly requires an observable measure of retail investors' subjective probability of a static-up VI and their expectation of late-arriving buying pressure mentioned in Section 1.3. Because it is challenging to convincingly construct these measures, I resort to demonstrating an incentive to purchase that gets stronger as prices approach the upper threshold. I do this by presenting an extremely simple trading strategy: purchasing a stock as soon as its best ask price breaches X% and selling either (1) at the static-up VI auction, or (2) at the best bid once 15 minutes elapses, whichever happens first. This strategy may well be described as *bubble-riding* and similar ideas are also presented in Cho et al. (2003), Seasholes and Wu (2007), and Chen et al. (2019b).

The first thing to notice in Table 1.3 is that the strategy improves as the threshold X% increases from 8% to 9.5%. In other words, this is a strategy that improves as we buy at *more expensive* prices. The average return increases from 0.28% to 0.34% and standard deviation of returns decreases from 3.00% to 2.45%. Given that the sum of round trip brokerage fees and trading taxes amount to 0.3%, this is not profitable strategy in this crude form, especially considering the large standard deviation of returns. The trades are presented to demonstrate that due to the inflow of buy orders that come after the static-up VI, there is an incentive to buy at an already elevated price of 8% or 9%. Furthermore, this incentive becomes even stronger as we reach higher prices.

	Model	9.5%	9.0%	8.5%	8.0%
Count	1,960	6,475	7,210	8,053	9,071
Prob VI	0.70	0.84	0.75	0.64	0.55
Prob Profit	0.67	0.60	0.63	0.59	0.55
Mean	0.70	0.34	0.35	0.30	0.28
SD	3.29	2.45	2.67	2.85	3.00
Min	−12.23	−14.69	−14.59	−14.59	−18.48
25%	1.40	−0.36	−0.73	−1.35	−1.65
50%	1.41	0.38	0.68	0.82	0.64
75%	2.39	1.09	1.45	1.75	2.04
Max	22.40	22.40	22.40	22.40	22.40

Table 1.3: Trading Strategy Returns. This table reports the results from the trading strategy described in this subsection. The first column presents the result from using a classification model discussed in Appendix A2. Other columns refer to results from a different buying cutoff. The top row specifies the number of cutoff breaches across all stocks during the period Jun 2015–May 2016. The second row shows the proportion of those breaches that end up in a static-up VI. The third row shows the proportion of trades that end up making a profit.

VI Prediction. The trade-off in choosing a higher cutoff is that the probability of exiting through a VI auction increases, but the purchase price also increases. If a trader can identify variables that predict a static-up VI, conditional on entering her trade at 8%, this naive strategy may be improved. I start by inspecting variables that may predict short-term retail activity to see if they can explain the occurrence of a VI, at least in-sample. If stock i breaches 8% on day d and time t, a trading position is opened. If this position exits through a VI auction, we let $y_{idt} = 1$, and let $y_{idt} = 0$ otherwise.

Daily covariates $\mathbf{X}_{id}$ include log market capitalization, Amihud measure, foreign investor holdings share, current day's overnight return (return from previous close to current day's open), past 20-day return, price deviation relative to 5-day moving average price, and log 5-day average trade volume. Intraday covariates $\mathbf{Z}_{idt}$ include log 5-minute trading volume, order imbalance (bid order volume over total order volume outstanding), and log order book depth. Using these variables, I run OLS regressions (linear probability model) and logistic regressions according to equations (1.4) and (1.5).

$$y_{idt} = \alpha + \Gamma' \mathbf{X}_{id} + \Lambda' \mathbf{Z}_{idt} + \varepsilon_{idt} \tag{1.4}$$

$$\mathbf{P}\left(y_{idt} = 1 \mid \mathbf{X}_{id}, \mathbf{Z}_{idt}\right) = \frac{\exp(\alpha + \Gamma'\mathbf{X}_{id} + \Lambda'\mathbf{Z}_{idt} + \varepsilon_{idt})}{1 + \exp(\alpha + \Gamma'\mathbf{X}_{id} + \Lambda'\mathbf{Z}_{idt} + \varepsilon_{idt})} \tag{1.5}$$

The results are shown in Table 1.4, where columns (1)–(3) show the coefficient estimates from OLS and columns (4)–(6) show those from logistic regressions. Past returns, past volumes, and intraday variables are omitted from columns (1), (2), (4), and (5) to demonstrate the additional explanatory power of these variables.

The coefficient signs mostly line up well with intuition. We see that characteristics that indicate lower liquidity, such as smaller size, higher Amihud measure, and lower foreign institutional holding, lead to higher chance of static-up VI occurrences. Also, characteristics that indicate heightened sentiment, such as high overnight returns, higher recent trading volume, and higher order imbalance, also lead to higher chance of static-up VI occurrences. Unsurprisingly, the explanatory power of the regression models are still low because the exercise can be seen as attempting to predict future price movements.

While these coefficient estimates provide the intuition that imminent trading activity matters, these generalized linear models cannot capture high-dimensional relationships between VI occurrence and covariates. Because such relationships may provide extra predictive power, I train a tree-based model using the gradient boosting algorithm (Chen and Guestrin 2016) together with various other features. The details are reported in Appendix A2. With this prediction model, I compute the trading results from entering positions at the 8% threshold only if an upward-VI is predicted. This result is reported in the first column of Table 1.3. We see that the average return is 0.70% and the probability of exiting through a VI is 0.70, while the standard deviation of returns is roughly similar as before (4th column). I offer these results to demonstrate that the foreseeable inflow of retail demand in case of a VI presents an incentive for going long—taking liquidity, rather than providing liquidity—near the upper thresholds.

	VI: Linear			VI: Logistic		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	2.58***	2.48***	2.56***	8.85***	8.37***	8.84***
	(0.15)	(0.16)	(0.16)	(0.67)	(0.71)	(0.73)
Log Market Cap	−0.08***	−0.09***	−0.09***	−0.35***	−0.38***	−0.38***
	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)
Amihud	1.30***	1.40***	1.23***	6.00***	6.73***	6.00***
	(0.16)	(0.17)	(0.17)	(0.79)	(0.88)	(0.90)
Foreign Holdings	−0.34***	−0.31***	−0.30***	−1.61***	−1.46***	−1.45***
	(0.06)	(0.06)	(0.06)	(0.29)	(0.29)	(0.29)
Overnight Return		1.87***	1.89***		8.04***	8.26***
		(0.22)	(0.22)		(0.97)	(0.98)
Past 20-Day Return		0.01	0.01		0.05	0.03
		(0.02)	(0.02)		(0.11)	(0.11)
Deviation from 5-Day MA		0.11	0.07		0.46	0.29
		(0.07)	(0.07)		(0.31)	(0.32)
Log 5-Day Trading Volume		0.01**	0.01		0.05**	0.04
		(0.00)	(0.01)		(0.02)	(0.02)
Hours from Open			−0.02***			−0.08***
			(0.00)			(0.01)
Log 5-Minute Trading Volume			0.01***			0.04***
			(0.00)			(0.01)
Order Imbalance			0.10***			0.41***
			(0.02)			(0.10)
Log Book Depth			−0.02***			−0.07***
			(0.00)			(0.01)
Observations	11,375	11,375	11,375	11,375	11,375	11,375
R^2	0.05	0.06	0.07			

Table 1.4: VI Prediction. This table reports the estimates from OLS and probit regressions where the outcome variable is a dummy variable indicating whether the trade exited with a static-up VI. Standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5%, and 1% level, respectively. The sample includes all instances of stocks breaching 8% in terms of return from open during the period Jun 2016–May 2018.

1.5 Price Reversion and Retail Wealth Loss

Previous sections demonstrated retail investors' trading patterns around static VIs. A natural follow-up question is whether such patterns are profitable for them. This section demonstrates that it is not, at least for the retail investors as a whole. Data reveals that prices revert after both static-up and static-down VIs, and that this reversion is persistent in the medium term. This implies that retail investors are buying at the peak for static-up VIs and selling at the trough for static-down VIs. Foreign institutions take the other side of these retail trades. Retail investors effectively end up transferring wealth to foreign institutions through trading activity around VIs.

1.5.1 Price Reversion around Threshold Breaches

I start by plotting the prices, in terms of return from open, around $\pm 10\%$ breaches. For each stock i, prices at the end of each 1-minute interval t are recorded. These prices are converted into return from open as before. For the pre-rule period, the event time is defined in terms of minutes relative to the first t such that return from open breaches $\pm 10\%$ during interval $t \to t+1$. For the post-rule period, it is defined relative to the time of the first static VI of the day. 15 minutes preceding the breaches and 60 minutes following them are sampled. I drop pauses and pseudo-pauses that occur within the first 15 minutes of market open and last 60 minutes of market close, so that all event times can fall in the sampled observations. Because the resolution is at 1-minute frequencies, the closing price for the 0th minute is not always exactly equal to the 10% threshold. For instance, if prices fall from -9.5% to -11% within a 1-minute interval, the first sampled minute at event time 0 will have a price of -11%. This is the reason we see prices that are more extreme than $\pm 10\%$ for the blue lines in Figure 1.17.

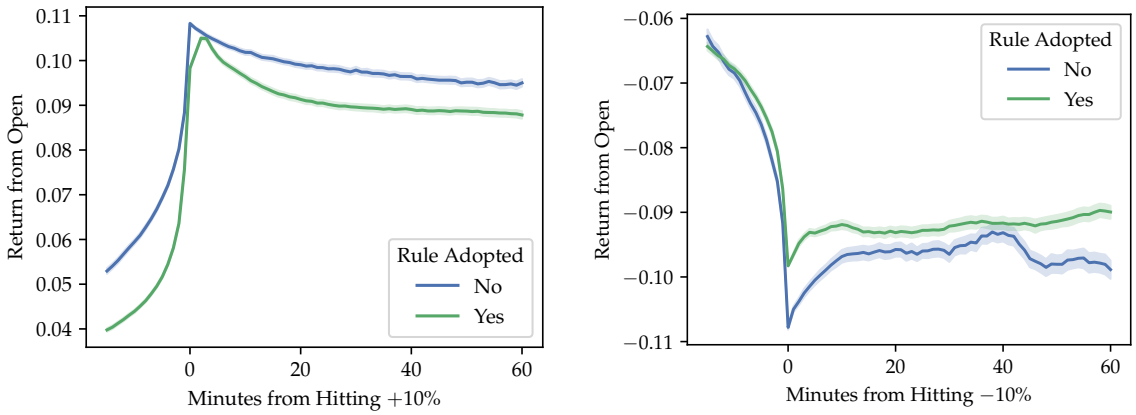

(a) Static-Up VI

(b) Static-Down VI

Figure 1.17: Prices around Static VIs. This figure plots average prices around VIs and pseudo-pauses, in terms of return from open. Return from open is computed using the closing price of each 1-minute interval according to Section 1.2.2. 95% confidence intervals are constructed by bootstrapping.

Figure 1.17 shows the results of the above procedure. Both Figures 1.17a and 1.17b show that prices revert after extreme returns from open. It should be noted that this pattern exists not just around the 10% thresholds, but at other less extreme thresholds such as 6% or 8%. The reversion is weaker at these cutoffs, however. If series of large flows impact prices and liquidity is scarce, we would expect to see such mean reversion if we sample ex-post at these threshold breaches. I do not attempt to explain the source of such reversion, but simply demonstrate that reversion around the VIs are extremely consistent as suggested by the tight confidence bands.

Recall that the first 2 minutes of the green lines are within the VI auction phase. This means that there are no executed prices during this brief period. This is why observations are missing for minute 1 in the figures above. Also, this means that the first observed prices at minute 2 are the clearing prices of the VI auctions. We can see that in both figures, prices tick up slightly for the green lines, indicating that the clearing prices of the VI auctions were slightly higher—about 0.5% on average—than the VI trigger prices. It is natural to ask whether the price reversion in Figure 1.17 is temporary. For both pre-rule and post-rule observations, stocks close about 1.5% lower relative to the price at the moment of the 10% breach.

1.5.2 Retail Wealth Losses

The price and retail trading patterns jointly suggest that retail investors are transferring wealth to foreign institutions around the VIs. I make back-of-the-envelope calculations for the magnitude of these losses by making the assumption that wealth is marked to market using the daily closing

prices. A detailed explanation follows in the next subsections.

The bottom line is that during the two-year period following the rule adoption, retail investors transferred around an extra KRW 170 billion (approx. USD 150 million) to institutional investors on *extreme return stock-days*. A stock-day observation (i, d) during which the stock i breaches $\pm 10\%$ is referred to as an extreme return stock-day. This loss happens over around 13,000 extreme return stock-days during this two-year period. This figure does not include the transfers from less sophisticated retail investors to other retail investors, which is expected to be at least an order of magnitude larger given the relative magnitudes of retail volume and retail net buy. Furthermore, total retail trading volume during these stock-days amount to KRW 210 trillion (approx. USD 190 billion). Due to round trip brokerage fees and trading taxes, which add up to at least 0.3% of trading volume, this costs retail investors an extra KRW 620 billion (approx. USD 550 million) over this period.

Losses around Static VIs. Suppose a static-up VI is triggered at minute t_0. Denote by x_t the number of shares net bought by retail during the interval $t \to t+1$ and y_t the net outflow of dollars due to these net trades. Because the exact holdings data is not available, we ignore the previous holdings and focus on the marginal change in marked-to-market wealth from t_0 onward. If P_{t_0+1} is the price at the end of this interval, retail investors' wealth change after the first interval will be

$$\Delta w_{t_0+1} = x_{t_0} P_{t_0+1} - y_{t_0}.$$

If prices are constant throughout the interval, $x_{t_0} P_{t_0+1} - y_{t_0}$ so that the wealth change is zero. Similarly, the marked-to-market wealth after T periods would be

$$P_{t_0+T+1} \sum_{t=0}^{T} x_{t_0+t} - \sum_{t=0}^{T} y_{t_0+t} \tag{1.6}$$

because over this period retail investors spent $\sum_{t=0}^{T} y_{t_0+t}$ dollars to accumulate $\sum_{t=0}^{T} x_{t_0+t}$ shares, each of which is valued at P_{t_0+T+1}. The implicit assumption is that the terminal price is the fair price, so later on sensitivity analysis needs to be done with respect to T.

All variables are directly observable in the data. Using the above definition, I first compute the loss in wealth for each static-up VI event using $T = 60$ minutes. I find that on average, retail investors lose approximately USD 9,000 per VI event and this amounts to approximately USD 45 million per year. Among these, 80% of the transfer is reaped by foreign institutions and the rest goes to domestic proprietary traders. Results for using different values for the fair price P_{t_0+T+1}

are included in Table 1.5. The results indicate that estimated losses are larger if we use a larger T.

	60 Mins	120 Mins	180 Mins	Day' Close
Count	7,072	7,072	7,072	7,072
Mean	−10.09	−11.36	−12.41	−14.77
SD	80.55	86.57	100.52	114.95
Min	−2119.61	−2119.61	−3692.76	−3934.78
25%	−6.73	−7.44	−7.96	−9.33
50%	−0.48	−0.64	−0.69	−0.89
75%	1.94	1.79	1.83	1.73
Max	1146.77	991.34	991.34	1332.83

Table 1.5: Retail Trading Losses around Static-Up VIs. This table reports the retail trading losses computed according to equation (1.6) for different terminal prices. The last column uses the closing price of the day of the VI occurrence to mark-to-market. The reported values for mean, standard deviation, and percentiles are in terms of KRW millions.

Losses on Days with Extreme Returns. Using the procedure in equation (1.6), I can compute the profit/loss from trading in stock i on any given day d for each investor type. I assume for now that wealth is marked-to-market using stock i's closing price of day d, P_{id}. If retail investors in aggregate paid y_{id} in KRW terms throughout the day in order to accumulate x_{id} shares in net terms, then the profit/loss of the day is

$$\Delta w_{id} = P_{id} x_{id} - y_{id}. \tag{1.7}$$

I want to test if these daily net profits are lower on days that VIs occur. On a day in which a static-up VI occurs, retail investors will be net buying at a high price ($y_{id} \uparrow$) and will end up with more shares ($x_{id} \uparrow$) at the end of the day, but the price of these shares will have reverted ($P_{id} \downarrow$). I test this in a diff-in-diff set up where the first difference is between time relative to the VI rule adoption and the second difference is between normal and extreme return days. More specifically, I assign relative quarters, $q = -3, \ldots, 12$, to dates around the month of the rule adoption, and assign $\text{Treated}_{id} = 1$ if stock i breaches $\pm 10\%$ during day d. The first month after the rule adoption and the last month before the rule adoption are removed to allow for a burn-in period. Then, I run the following regression:

$$y_{idq} = \gamma \cdot \text{Treated}_{id} + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \gamma_q \cdot D_q + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \beta_q \cdot D_q \times \text{Treated}_{id} + \Gamma' \mathbf{X}_{id} + \alpha_i + \varepsilon_{idq} \tag{1.8}$$

where the outcome variables are daily retail volume and daily retail profit in KRW terms. The control variables are as in specification (1.3).

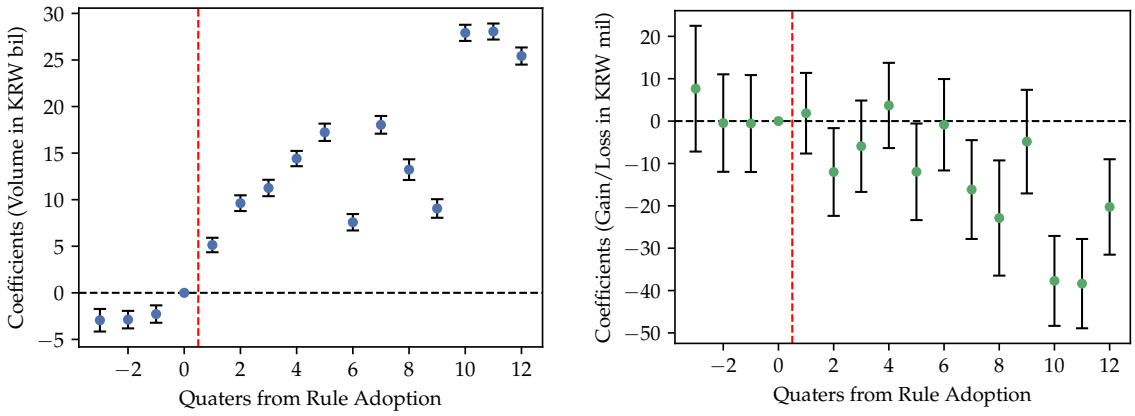

(a) Retail Volume

(b) Retail Profit/Loss

Figure 1.18: Retail Trading Volume and Profits on Normal vs. Extreme Return Days. This figure plots the estimates of the coefficients β_q from the diff-in-diff regression specification in equation (1.8). The outcome variables are retail trading volume in KRW billions and retail trading net profits in KRW millions. 95% confidence intervals are displayed.

The coefficients β_q's for $q > 0$ capture the effect of being in an extreme return day during post-rule quarters, after controlling for quarter fixed effects. Results in Figure 1.18b point towards larger losses, but are too noisy to interpret for most quarters. However, the results in Figure 1.18a show a clear pattern of increased activity on extreme return days. A coefficient estimate of 30 translates to a cost of: KRW 30 billion $\times$ 0.3% = KRW 90 million. This amounts to more than USD 80,000 of trading taxes and brokerage costs per stock-day. The costs add up to substantial amounts because there are more than 10,000 such stock-days in a year.

To gauge the average effect over the quarters adjacent to rule adoption, I run the following diff-in-diff regression:

$$y_{idq} = \gamma_1 \cdot \text{Treated}_{id} + \gamma_2 \cdot \text{Post}_{id} + \beta \cdot \text{Treated}_{id} \times \text{Post}_{id} + \Gamma' \mathbf{X}_{id} + \alpha_i + \gamma_q + \varepsilon_{idq} \tag{1.9}$$

where $\text{Treated}_{id} = 1$ if stock i breached $\pm 10\%$ during day d and $\text{Post}_{id} = 1$ if day d falls in the post-rule period. The stock-day-level control variables are lagged by 20 trading days, except for the past return variables which are lagged by 1 trading day.

	Retail Volume		Retail Profit	
	(1)	(2)	(3)	(4)
Constant	-77.89^{***}	-133.87^{***}	73.53^{***}	114.56^{***}
	(0.34)	(20.78)	(3.75)	(30.37)
Hit $\pm 10\%$	10.92^{***}	10.59^{***}	-20.33^{***}	-19.51^{***}
	(0.21)	(0.95)	(2.26)	(1.71)
Hit $\pm 10\% \times$ Post	16.88^{***}	16.75^{***}	-12.08^{***}	-13.10^{***}
	(0.23)	(2.65)	(2.53)	(4.40)
Log Market Cap	2.91^{***}	5.10^{***}	-2.77^{***}	-4.37^{***}
	(0.01)	(0.78)	(0.13)	(1.14)
Amihud	4.36^{***}	10.84^{***}	-8.79^{**}	-10.36^{***}
	(0.36)	(2.86)	(3.93)	(3.22)
Past Volatility	125.41^{***}	79.88^{***}	-39.10^{**}	-10.75
	(1.60)	(10.52)	(17.63)	(15.33)
Market Beta	0.63^{***}	0.06	-1.11^{***}	-0.74
	(0.03)	(0.31)	(0.32)	(0.54)
Past 20-Day Return	19.51^{***}	19.49^{***}	-4.86^{***}	-4.55^{**}
	(0.11)	(1.92)	(1.26)	(2.27)
Past 200-Day Return	2.22^{***}	1.98^{***}	-1.61^{***}	-1.23^{**}
	(0.03)	(0.69)	(0.33)	(0.57)
Stock FE	–	✓	–	✓
Quarter FE	–	✓	–	✓
Observations	797,293	797,293	797,293	797,293
R^2	0.20	0.18	0.002	0.001

Table 1.6: Retail Trading Volume and Losses on Extreme Return Days. This table reports the estimates from diff-in-diff regressions as specified in equation (1.9) where the outcome variables are daily retail trading volume and daily retail trading profit. Trading volume is in KRW billions and trading profit is in KRW millions. Standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5%, and 1% level, respectively. For columns (2) and (4), standard errors are two-way clustered at the stock and quarter levels.

The coefficient of interest is β, which is reported in the third row of Table 1.6. We see from the second row that retail investors trade more extensively and also make larger losses on extreme return days. The third row shows that both of these effects intensify during the post-rule period. These results are inconsistent with the intention of cooling-off markets in case of extreme price movements.

One concern may be that extreme returns, which determines Treated_{id}, is in fact a result of high

trading volume y_{idq}. This can be addressed to some degree by taking the trading activity that comes *after* the occurrence of $\pm 10\%$ breaches. Because this necessarily drops all preceding trades, average trading activity has to be used instead of daily gross trading activity in order to make a fair comparison. This procedure does not affect the significance of coefficient estimates, but complicates the interpretation.

1.5.3 Effect on Overall Market Conditions

Despite these concerning patterns, circuit breaker-like rules are ubiquitous. Sifat and Mohamad (2019) count 48 stock exchanges with trading halts, 98 with price limits, and 31 with volatility interruption mechanisms out of the 152 exchanges studied, as of 2018. More than 85% of world's stock market capitalization is subject to some form of stock-level trading interruption rules. This leads us to expect that such rules improve price stability and liquidity. I start by testing whether this is true. The difficulty arises from the fact that the rule started to apply to all stocks with the same thresholds on the same day. This precludes the usual empirical designs related to natural experiments since we lack a control group.

Volatility as Treatment Dosage. While there were no variation in the threshold level, I propose a way to turn this nuisance on its head. The threshold for a static VI, namely 10%, was chosen arbitrarily. For a highly volatile stock with daily volatility of 5%, this threshold is *tight*, whereas for a stock with daily volatility of 1%, it is *loose*. In this sense, stocks were treated to a different degree: volatile stocks received a higher dosage of treatment. For instance, Samsung Electronics' common stock never triggers a static VI during the sample period used here.

In the benchmark specification, stocks are sorted by their daily return volatility during the past 60 trading days. Then, the top quintile is considered the treatment group while the bottom quintile is considered the control group. The other three quintiles are discarded. The results remain similar in different variations—e.g., including the other quintiles or using terciles instead. Also, results are similar when sorts are done based on past 20 quarters before the beginning of the pre-rule period.

Effect on Volatility. Outcome variables related to a stock's volatility are: 5-minute realized volatility, daily return volatility, high-low price range, and number of $\pm 10\%$ breaches. Daily return volatility is computed at the monthly frequency using daily close-to-close returns. 5-minute realized volatility for a given day is computed by summing squared 5-minute log returns within each day, then taking the square root of this number. Then, this daily value is averaged at the

monthly frequency. High-low price range is computed by taking the difference between the daily high and daily low, and then dividing this number by the opening price. Again, this is averaged at the monthly frequency. The number of $\pm 10\%$ breaches is computed by counting the number of days during which a stock breaches $\pm 10\%$ relative to its opening price.

Event study plots in Figure 1.19 show the coefficients β_q's from running the regressions under the following specification:

$$y_{ijmq} = \gamma \cdot \text{Treated}_{im} + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \gamma_q \cdot D_q + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \beta_q \cdot \text{Treated}_{im} \times D_q + \Gamma' \mathbf{X}_{ijmq} + \alpha_j + \varepsilon_{ijmq} \tag{1.10}$$

where y_{ijmq} is the outcome variable for stock i in month m, α_j is the industry fixed effect, $\text{Treated}_{im} = 1$ if stock i falls in the fifth quintile in terms of past volatility, and D_q is a dummy variable that is equal to 1 if the observation falls in quarter q. Logarithms are taken for the outcome variables to test for relative changes. Quarters are relative to the rule adoption and the month of the rule adoption, June 2015, is dropped. Vector of controls $\mathbf{X}_{ijmq}$ include log market capitalization, market beta using past 120 trading days, and past 5-day, 20-day, 200-day returns. More importantly it also includes quintile-specific slopes for market volatility z_m and market return x_m in month m:

$$\delta_1 \cdot z_m + \delta_2 \cdot \text{Treated}_{im} \times z_m + \delta_3 \cdot x_m + \delta_4 \cdot \text{Treated}_{im} \times x_m.$$

This is to partially address the inherent limitation of this specification, which is that selection on past volatility happens by design. Since more volatile stocks would respond differently to market conditions, heterogeneous slopes are added.

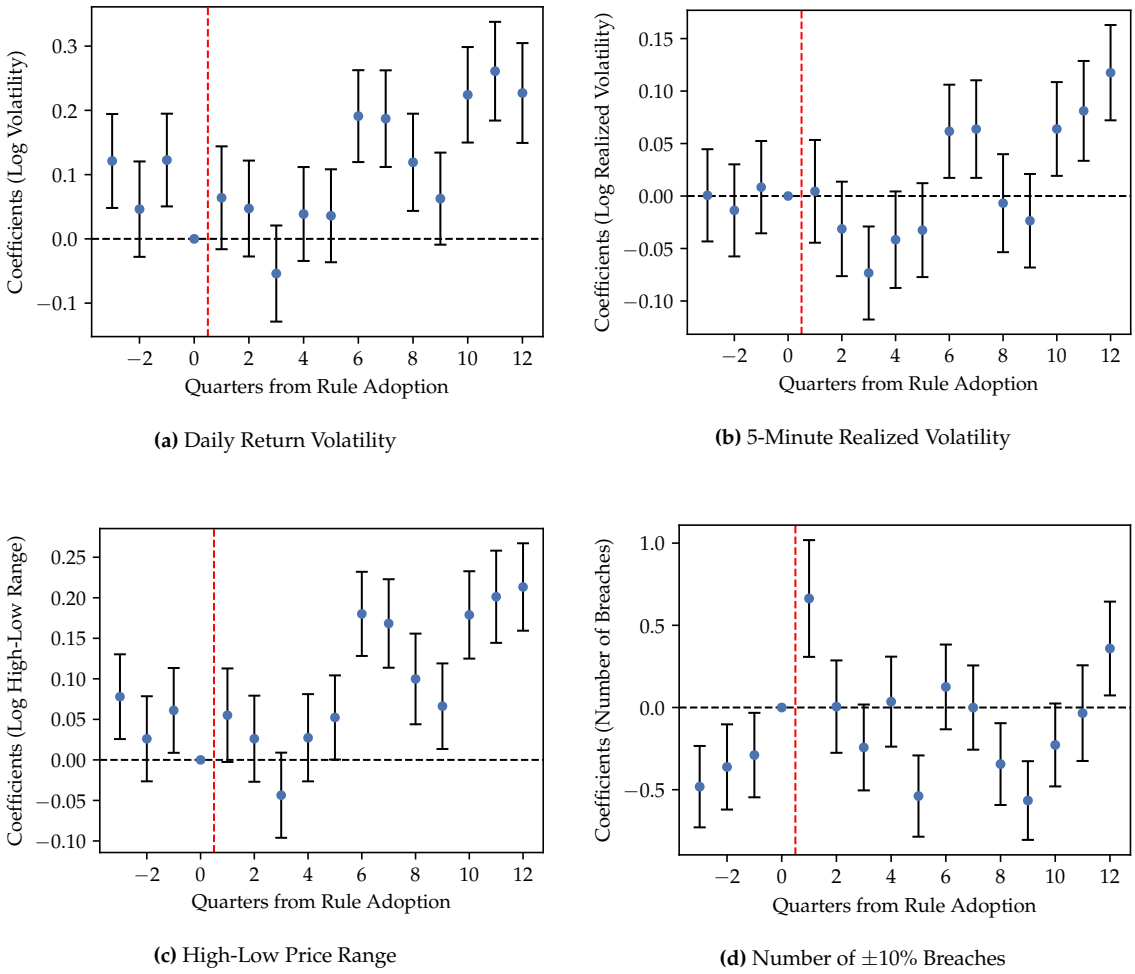

(a) Daily Return Volatility

(b) 5-Minute Realized Volatility

(c) High-Low Price Range

(d) Number of $\pm 10\%$ Breaches

Figure 1.19: Measures of Volatility and Rule Adoption. This figure plots the estimates of the coefficients β_q's from the diff-in-diff regression specification in equation (1.10). Logs are taken for the outcome variables in panels (a), (b), and (c). 95% confidence intervals are displayed.

Coefficients displayed in panels (a), (b), and (c) represent the average *differences in proportional changes* because outcome variables are log-transformed. The expected policy effect would be either decreased volatility or decreased frequency of extreme returns—i.e., negative and significant coefficients. The event study plots are not consistent with this policy effect.

I also run the following regression specification that is closer to the usual diff-in-diff specification:

$$y_{ijmq} = \gamma \cdot \text{Treated}_{im} + \beta \cdot \text{Treated}_{im} \times \text{Post}_m + \Gamma' \mathbf{X}_{ijmq} + \alpha_j + \gamma_q + \varepsilon_{ijmq} \tag{1.11}$$

where $\text{Post}_m = 1$ if month m comes after the rule adoption date, α_j and γ_q are stock and quarter fixed effects respectively, and all other variables are the same as before. Both quarter and industry fixed effects are included. Here I use the 4 quarters around the rule adoption instead.

	Return Vol	Realized Vol	High-Low Range	Num Breaches
	(1)	(2)	(3)	(4)
Treated × Post	−0.02	−0.03*	−0.01	0.30
	(0.03)	(0.02)	(0.02)	(0.19)
Stock Controls	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Observations	7,558	7,558	7,558	7,558
Within R^2	0.38	0.52	0.54	0.22

Table 1.7: Diff-in-Diff Coefficient Estimates for Volatility. This table reports the coefficient estimates of the diff-in-diff regression specification in equation (1.11). Logs are taken for the outcome variables in columns (1), (2), and (3). Standard errors are two-way clustered at the industry and quarter level, and are shown in parentheses. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

Table 1.7 reports the coefficients β's from running the regressions according to equation (1.11). We see that only the coefficient in column (2) is weakly negative, and that the point estimate in column (4) is in fact positive. The introduction of static VI does not appear to depress the price variability in the treatment group.

Effect on Liquidity. I repeat the exercise in the previous subsection with measures of liquidity. Outcome variables for liquidity are: relative bid-ask spread, order book depth, and the Amihud measure. The relative bid-ask spread is computed by normalizing the bid-ask spread by the mid-price, and then averaging this number at the monthly frequency. Order book depth is the monthly average of outstanding total volume at the first three levels of the order book on both the ask and bid sides. The Amihud measure is computed at the monthly frequency as in Amihud (2002), except volume is expressed in KRW billion.

	Bid-Ask Spread	Book Depth	Amihud Measure
	(1)	(2)	(3)
Treated × Post	−0.04	0.05	−0.14*
	(0.03)	(0.10)	(0.07)
Stock Controls	✓	✓	✓
Industry FE	✓	✓	✓
Quarter FE	✓	✓	✓
Observations	7,239	7,239	7,239
Within R^2	0.31	0.08	0.75

Table 1.8: Diff-in-Diff Coefficient Estimates for Liquidity. This table reports the coefficient estimates of the diff-in-diff regression specification in equation (1.11). Logs are taken for the outcome variables. Standard errors are two-way clustered at the industry and quarter level, and are shown in parentheses. *, **, and *** denote 10, 5, and 1% statistical significance respectively.

Again, coefficients displayed in Table 1.8 represent differences in proportional changes because outcome variables are log-transformed. If the policy works as expected, we should see negative coefficients in columns (1) and (3), and a positive coefficient in column (2). Only the coefficient in column (3) is consistent with the policy objectives. However, the coefficient is statistically significant only at the 10% level.

Preferred Stocks as Treatment Group Using the same reasoning as before, I use preferred-common stock pairs as treatment-control pairs: preferred stocks are more volatile and can be considered to face a tighter threshold. In this case, adding firm fixed effects mitigates the selection problem in the previous section. I repeat the regression in equation (1.10) with underlying firm fixed effects instead of industry fixed effects.

	Return Vol	Realized Vol	High-Low Range	Num Breaches
	(1)	(2)	(3)	(4)
Treated × Post	−0.02	0.05	0.07	0.48
	(0.10)	(0.06)	(0.08)	(0.46)
Firm FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Observations	3,526	3,526	3,526	3,526
Within R^2	0.02	0.07	0.04	0.02

Table 1.9: Diff-in-Diff Coefficient Estimates for Volatility. This table reports the coefficient estimates of the diff-in-diff regression specification in equation (1.11). Standard errors are clustered at the quarter level, and are shown in parentheses. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

	Bid-Ask Spread	Book Depth	Amihud Measure
	(1)	(2)	(3)
Treated × Post	-0.19^{***}	-0.25^{***}	-0.68^{***}
	(0.05)	(0.08)	(0.22)
Firm FE	✓	✓	✓
Quarter FE	✓	✓	✓
Observations	3,379	3,379	3,379
Within R^2	0.58	0.69	0.76

Table 1.10: Diff-in-Diff Coefficient Estimates for Volatility. This table reports the coefficient estimates of the diff-in-diff regression specification in equation (1.11). Standard errors are clustered at the quarter level, and are shown in parentheses. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

The results are presented in Table 1.9 and Table 1.10. As before, we do not see improved price stability in the treatment group. In fact, we see positive, but insignificant point estimates in columns (2), (3), and (4) of Table 1.9. In Table 1.10, we see improvements in terms of bid-ask spread and the Amihud measure. However, the coefficient estimate in column (2) is negative, which is contrary to the expected effect of the rule introduction. Again, results are at best mixed, and it is difficult to make definitive statements on the efficacy of VI rule adoption on overall market conditions.

1.6 Conclusion

Using data from the Korean stock market, I present a series of unintended consequences of stock-level trading pauses, whose stated goal is "investor protection." During the years following rule

introduction, retail investors as a whole have transferred hundreds of millions of dollars of wealth to foreign institutions on extreme return days which triggered trading pauses. These losses emerge due to the tendency of retail investors to bet on price continuation around these pauses. They have lost even more wealth through trading taxes and costs related to excessive attention-induced trading caused by the salience of this mechanism. Furthermore, this paper finds scant evidence of improvements in market conditions, such as volatility or liquidity, that benefit these retail investors.

The Korean stock market proves to be particularly well-suited in documenting these empirical facts due to its institutional features related to trade flow data. First, investor type-level flow data exists for a relatively long time series at a high resolution, and spans both the pre- and post-rule periods. Second, these equity flows are exhaustive because KRX is the sole exchange through which all stock transactions happen. Lastly, retail investors actively, and directly, participate in the stock market, so that clear patterns emerge around extreme price movements even at high frequencies. These edges allow me to quantify the retail wealth losses that occur immediately around trading pause events, and on days during which pauses are triggered.

Although such uniqueness is what enables the project, the lessons we can draw are broader. Previous studies on trading pauses and halts, whose settings span all major markets in the US, China, and Europe, hint at the universality of post-pause market instability. The similarities lead us to suspect that retail investors are being left vulnerable in other major markets as well. A more general takeaway is that preponderance of behavioral actors may derail even simple regulatory interventions.

There are a couple of venues that this paper can be pushed further. By relying on the arbitrariness of the pause thresholds, I attempt to circumvent the identification challenge posed by the fact that the rule was adopted to all stocks on the same day. However, this remains an imperfect alternative. Working directly with the KRX to carry out relevant randomized control experiments would address identification concerns. Also, explicitly modeling and estimating the proportion of inattentive arbitrageurs and inattentive behavioral actors will allow us to think about counterfactual policies that can better achieve the original regulatory goals.

Appendix to Chapter 1

A1 Data Description

Data Category	Variables
1-Minute Frequency Price Data	Open, High, Low, Close, Trading volume, Trading value, Accumulated trading volume, Accumulated trading value, Best bid order price, Best ask order price, Residual quantity at best bid order, Residual quantity at best ask order, Bid-ask midpoint, Depth-weighted average limit price, Total cancelled IOC, Total cancelled FOK, Cancelled order ratio, Cancelled order volume ratio, Buy-sell imbalance in number of trades, Buy-sell imbalance in trading volume, Order imbalance in number of orders, Order imbalance in order volume, Average quoted spread, Average percentage quoted spread, Average effective spread, Average realized spread, HS adverse selection costs, Depth in volume, Depth in number of orders
1-Minute Frequency Trade Flow Data	Investor code, Accumulated buy trading volume, Accumulated buy trading value, Accumulated sell trading volume, Accumulated sell trading value
Trade-Level Data	Expected matching price, Expected matching quantity, Total residual quantity at ask order, Ask step 1-10 price, Ask step 1-10 residual quantity, Total residual quantity at bid order, Bid step 1-10 price, Bid step 1-10 residual quantity, Order trade acceptance no., Accumulated trading volume by market, Accumulated trading value by market, Aggregated residual quantity at step1-10 best ask order, Aggregated residual quantity at step1-10 best bid order, Last Step 1 best ask order price, Last Step 1 best bid order price, Total residual quantity at last ask orders, Total residual quantity at last bid orders, etc.
Other Market Data	Short volume, Short interest, Book-to-market ratio, Dividend yield, Shares outstanding, Foreign holdings

Table 1.A1.1: Available Variables from KRX. This table lists the available variables for each dataset. More details are available on the KRX website.

A2 Trading Strategy with a Tree-Based Classifier

In this section, I provide the details of the training procedure of the gradient boosting algorithm that predicts imminent VIs. Gradient boosting is a classification (and possibly a regression) technique that employs weak decision tree classifiers. The individual tree predictions are aggregated to form an ensemble classifier. I use XGBoost (Chen and Guestrin 2016) to implement the algorithm.

The sample consists of stock-time observations when the stock first breaches 8% return relative to its opening price. This is a classification problem whose outcome variable $y = 1$ if this stock triggers a static-up VI within the next 15 minutes and $y = 0$ otherwise.

There are two limitations to the data I work with. The first is that trade-level data is only available for the period June 2014–May 2016. In order to compute realistic trading performance, I assume that purchases happen at the ask and sales happen at the bid. If I use only the trade-level data to train and test the model, I am forced to work with less than 10,000 data points. I choose the alternative of training the data using 1-minute frequency data, which span June 2014–May 2018, and using the trade-level data only when computing the trading performance. I perform cross-validation on the data from June 2016–May 2018, without using any of the trading period, and also train the final model on this same period afterwards. This unfortunately means that the classification model uses data from the future. Still, none of the trading period data is used in the training and cross-validation process.

The second limitation is the relative dearth of data. The above sampling procedure yields around 15,000 to be used in training and cross-validation, and around 9,000 trading instances. Because a tree-based method is chosen to capture high-dimensional relationship between the features, this is a small sample. One limiting factor comes from the pool of chosen stocks. To be consistent with the main analyses of the paper, much of the small-cap stocks are discarded. Because these are the stocks that most often satisfy the sampling criterion, I end up with a much smaller sample.

First, features are constructed according to the definitions in Table 1.A2.2. Then, 5-fold cross-validation is used to compute the average out of sample prediction performance for each hyperparameter combination. As with other machine learning algorithm, the gradient boosting algorithm is prone to overfitting and most of the hyperparameters control the amount of regularization imposed. Baseline values for tree depth, number of estimators, and learning rate were taken from Hastie et al. (2009).

Variable Name	Definition
Log Market Capitalization	log(shares outstanding $\times$ day's closing price), lagged by 5 days
Market Beta	Beta w.r.t. KOSPI, lagged by 5 days
Amihud Measure	$10^9 \times \frac{\lvert\text{day's return}\rvert}{\text{day's volume in KRW}}$, lagged by 5 days
Foreign Holdings	Percentage of foreign investor holdings, lagged by 5 days
Overnight Return	Return from previous day's close to the trading day's open
5-Day Return	Closing price return in the past 5 days, lagged by 1 day
20-Day Return	Closing price return in the past 20 days, lagged by 1 day
200-Day Return	Closing price return in the past 200 days, lagged by 1 day
5-Day Disparity	$\frac{\text{buy price}}{\text{average closing price in the past 5 days}} - 1$
20-Day Disparity	$\frac{\text{buy price}}{\text{average closing price in the past 20 days}} - 1$
40-Day Disparity	$\frac{\text{buy price}}{\text{average closing price in the past 40 days}} - 1$
Disparity from 5-Day High	$\frac{\text{buy price}}{\text{past 5-day high price}} - 1$
Disparity from 20-Day High	$\frac{\text{buy price}}{\text{past 20-day high price}} - 1$
Disparity from 200-Day High	$\frac{\text{buy price}}{\text{past 200-day high price}} - 1$
5-Day Trading Volume	Trading volume in KRW in the past 5 days
5-Day Normalized Trading Volume	$\frac{\text{trading volume in KRW in the past 5 days}}{\text{trading volume in KRW in the past 40 days}}$
10-Day Normalized Trading Volume	$\frac{\text{trading volume in KRW in the past 10 days}}{\text{trading volume in KRW in the past 40 days}}$
Opening Trading Volume	Trading volume during the first minute of the trading day
Normalized Opening Trading Volume	$\frac{\text{opening trading volume}}{\text{40-day average of opening trading volume}}$
5-Day Normalized Opening Volume	$\frac{\text{5-day average of opening trading volume}}{\text{40-day average of opening trading volume}}$
10-Day Normalized Opening Volume	$\frac{\text{10-day average of opening trading volume}}{\text{40-day average of opening trading volume}}$
5-Minute Normalized Trading Volume	$\frac{\text{trading volume in KRW in the past 5 minutes}}{\text{trading volume in KRW in the past 40 days}}$
Order Imbalance	$\frac{\text{bid volume outstanding in order book}}{\text{bid volume outstanding in order book}+\text{ask volume outstanding in order book}}$
Order Book Depth	Sum of volume at the first bid and first ask
Hours from Open	Hours relative to market open

Table 1.A2.2: Features and Definitions. This table lists the variables used for VI prediction. *Trading day* refers to the day that the stock in question breaches 8%. *Buy price* refers to the trading day's opening price $\times$ 1.08.

Preliminary analyses on performance were run as shown in Figure 1.A2.1 to come up with the hyperparameter grid/region. Then, a grid search is performed over the combinations of hyperparameters listed on Table 1.A2.3.

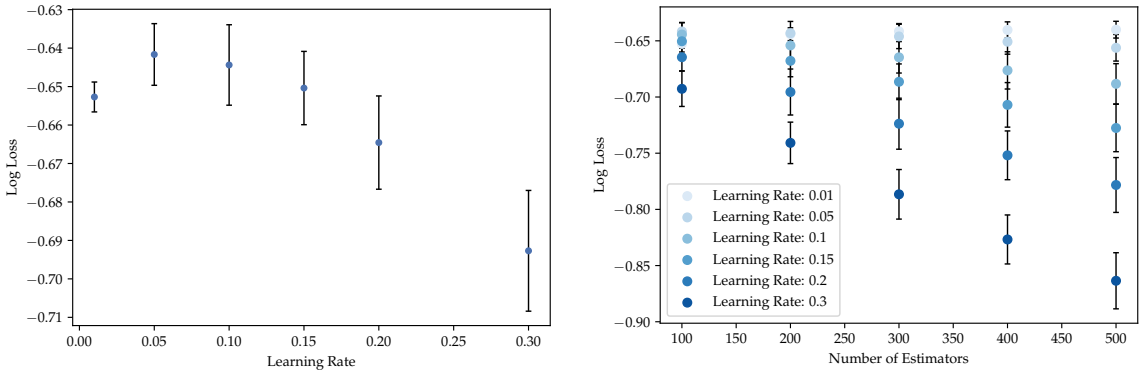

(a) Learning Rate

(b) Number of Estimators and Learning Rate

Figure 1.A2.1: Log-Loss for Different Hyperparameter Values. This figure plots the average log-loss, $\frac{1}{N}\sum_i^N [y_i \log p_i + (1-y_i)\log(1-p_i)]$ where p_i is the predicted probability of $y_i = 1$, among the cross-validation sets. Log-losses are computed for different learning rates and number of estimators while fixing other hyperparameters as default values.

The final hyperparameter values are marked in bold in Table 1.A2.3. A more intensive search and optimization is not performed because it is not the main focus of this paper and because the bottleneck of performance is likely the small sample size. The final model has an average accuracy of 0.63, precision of 0.61, and AUC of 0.67 using a 5-fold cross-validation. I use a more conservative predicted probability cutoff of 0.65 to get the final result reported in the first column of Table 1.3.

As a reference, Figure 1.A2.2 shows the mean decrease accuracy (MDA) for each feature. MDA is defined roughly as the loss in prediction accuracy, relative to the full model, when a given feature is randomly permuted. This measure is accepted to be most relevant for out-of-sample predictions because the computation of MDA naturally involves out-of-sample predictions. Figure 1.A2.2 shows that the most relevant features are size, 5-minute volume, and minutes from open. A value of 0.03 means that permuting a given feature decreases the prediction accuracy by 3 percentage points. The final model excludes features that have a negative MDA value.

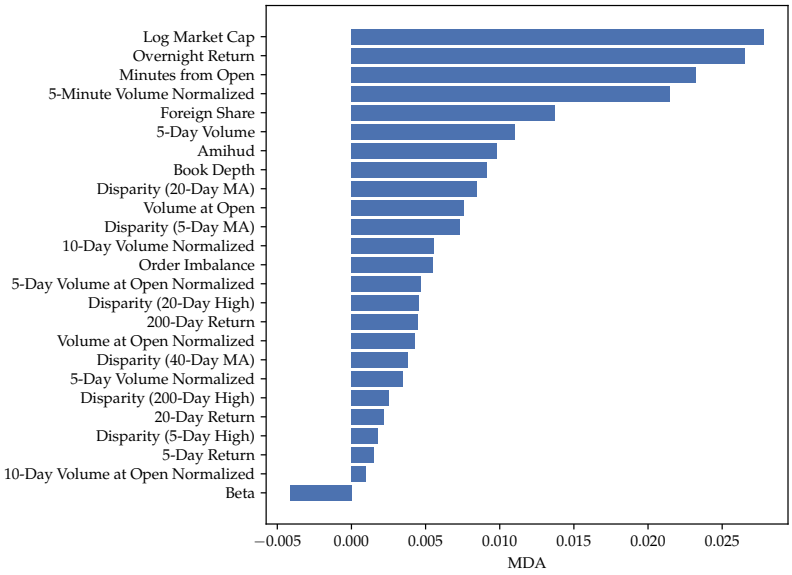

Figure 1.A2.2: Mean Decrease Accuracy of Features. This figure illustrates the MDA of the features used in the prediction exercise. 10 test samples were created using 0.2 of the data for each test sample. For each train-test set, a model was trained with one randomly permuted column and its performance was compared to the full model on the same train-test set. The average differences on the 10 sets are recorded and plotted.

Hyperparameter Name	Definition	Grid
Learning Rate	Step size shrinkage used in update to prevent overfitting.	[**0.01**, 0.05, 0.1, 0.15, 0.2, 0.3]
Number of Estimators (+)	Number of weak learners (trees).	[100, 200, 300, 400, **500**]
Maximum Depth (−)	Maximum depth of a tree.	[3, 4, **5**, 6, 7]
Minimum Loss Reduction (+)	Minimum loss reduction required to make a further partition on a leaf node of the tree.	[0, 0.25, **1**, 10]
Subsample Ratio (−)	Subsample ratio of the training instances at each tree growth.	[0.5, 0.7, **0.9**]
Column Subsample Ratio (−)	The subsample ratio of columns (features) when constructing each tree.	[0.3, **0.5**, 0.8]

Table 1.A2.3: Grid Search for Hyperparameters. This table lists the hyperparameters of XGBoost. Definitions are taken from the documentation page. Sign next to the hyperparameter name indicates the direction of the hyperparameter value that leads to the classifier's variance reduction.

A3 Additional Figures and Tables

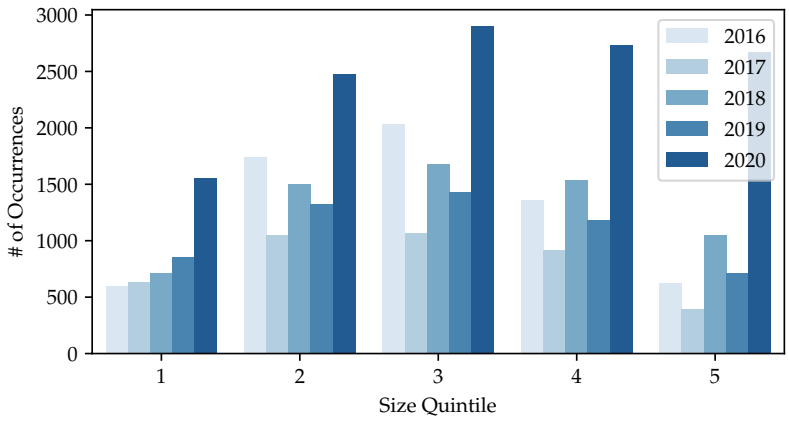

Figure 1.A3.3: Number of VIs by Size. This figure illustrates the number of annual VI occurrences by size quintile. Bins for size were made based on market capitalization at the beginning of each month. VIs include both static and dynamic VIs.

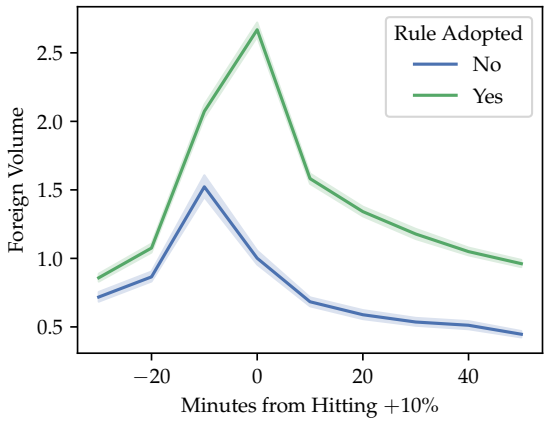

(a) Volume around Static-Up VI

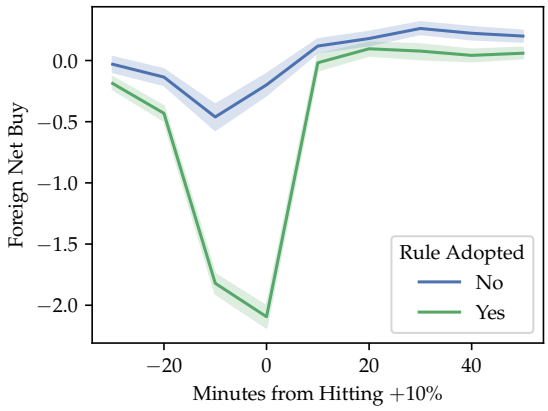

(b) Net Buy around Static-Up VI

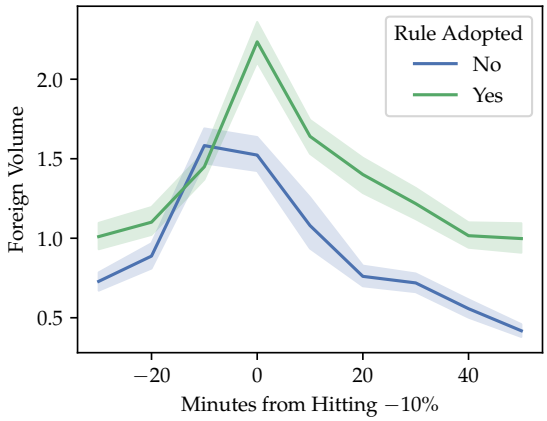

(c) Volume around Static-Down VI

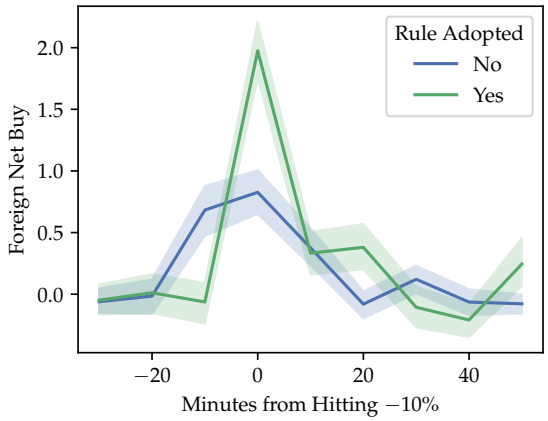

(d) Net Buy around Static-Down VI

Figure 1.A3.4: Foreign Volume and Net Buy around Static VIs. This figure plots the average foreign institutions' volume and net buy during 10-minute bins around VIs and pseudo-pauses. Trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, they are expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). 95% confidence bands are constructed by bootstrapping.

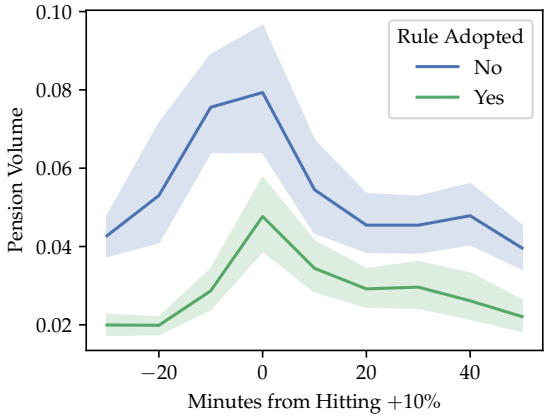

(a) Volume around Static-Up VI

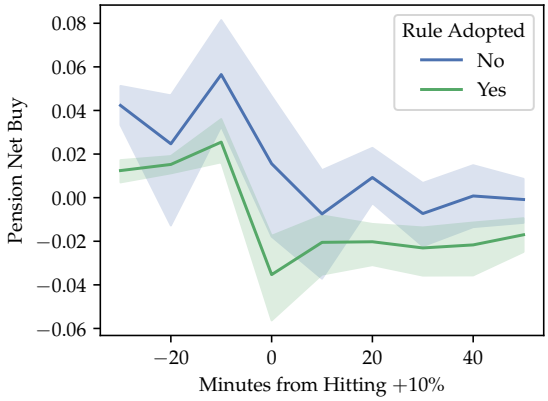

(b) Net Buy around Static-Up VI

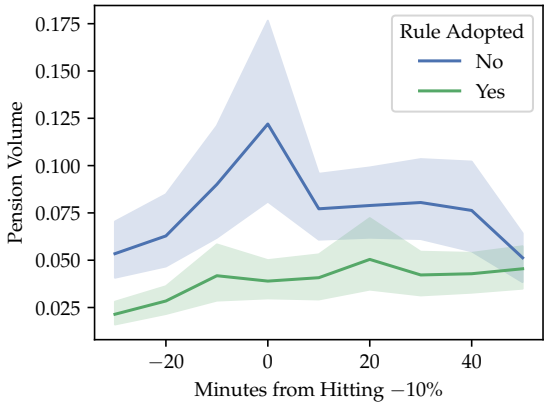

(c) Volume around Static-Down VI

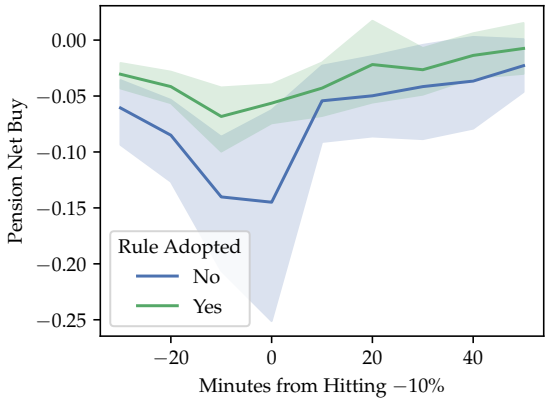

(d) Net Buy around Static-Down VI

Figure 1.A3.5: Pension Volume and Net Buy around Static VIs. This figure plots pension funds' average volume and net buy during 10-minute bins around VIs and pseudo-pauses. Trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, they are expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). 95% confidence bands are constructed by bootstrapping.

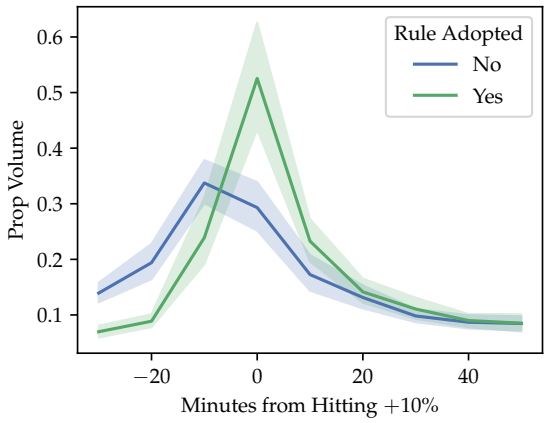

(a) Volume around Static-Up VI

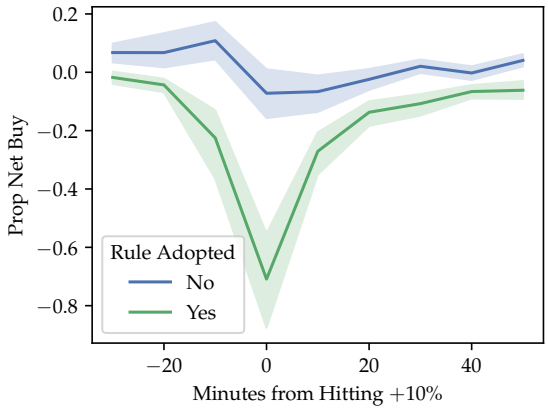

(b) Net Buy around Static-Up VI

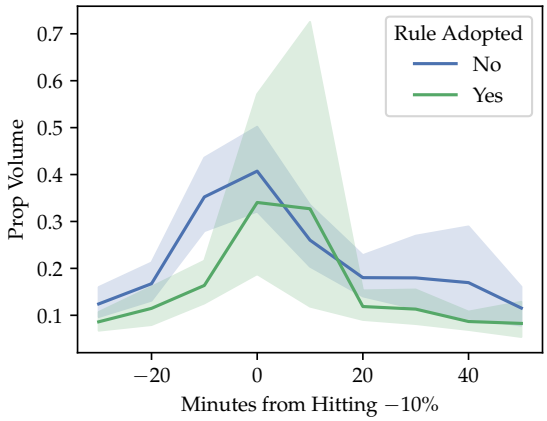

(c) Volume around Static-Down VI

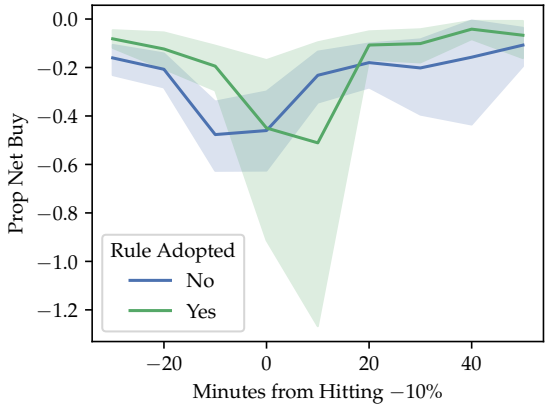

(d) Net Buy around Static-Down VI

Figure 1.A3.6: Prop Volume and Net Buy around Static VIs. This figure plots proprietary traders' average volume and net buy during 10-minute bins around VIs and pseudo-pauses. Trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, they are expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). 95% confidence bands are constructed by bootstrapping.

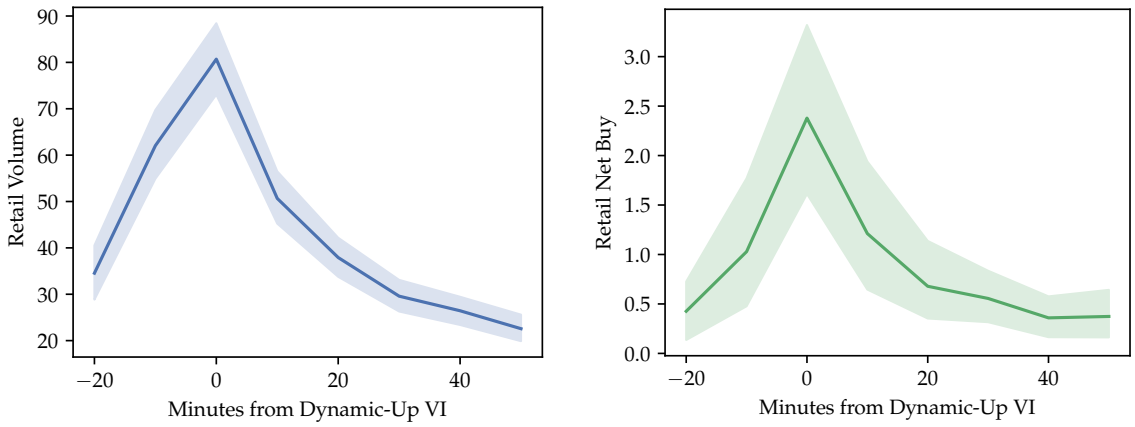

(a) Retail Volume

(b) Retail Net Buy

Figure 1.A3.7: Retail Volume and Net Buy around Dynamic-Up VIs. This figure plots the average retail volume and net buy during 10-minute bins around dynamic-up VIs. Retail trading volume in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. 95% confidence bands are constructed by bootstrapping.

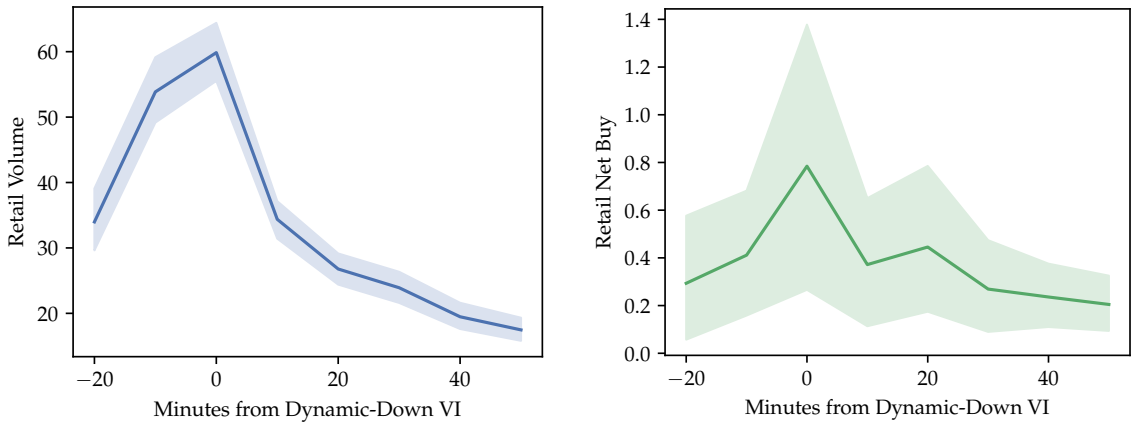

(a) Retail Volume

(b) Retail Net Buy

Figure 1.A3.8: Retail Volume and Net Buy around Dynamic-Down VIs. This figure plots the average retail volume and net buy during 10-minute bins around dynamic-down VIs. Retail trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. 95% confidence bands are constructed by bootstrapping.

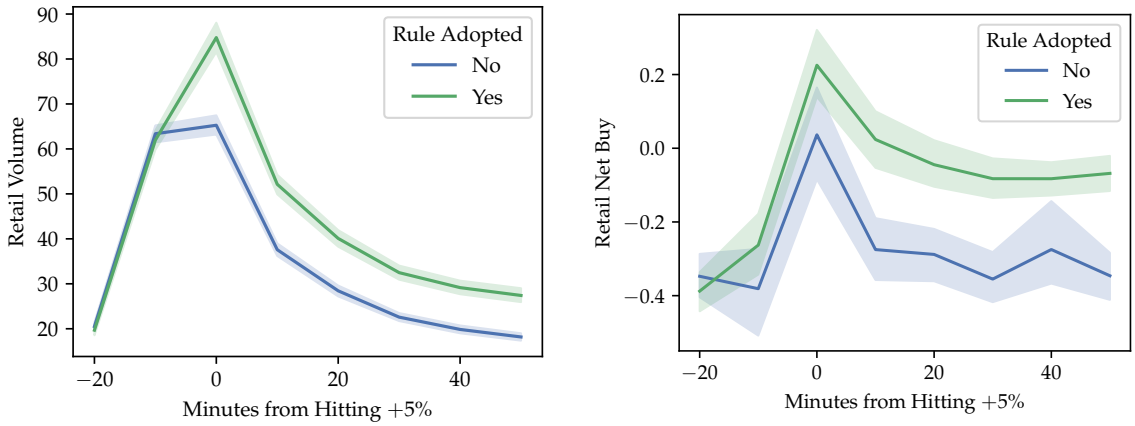

(a) Retail Volume

(b) Retail Net Buy

Figure 1.A3.9: Retail Volume and Net Buy around +5% Breaches. This figure plots the average retail volume and net buy during 10-minute bins around +5% breaches. Retail trading volume in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. 95% confidence bands are constructed by bootstrapping.

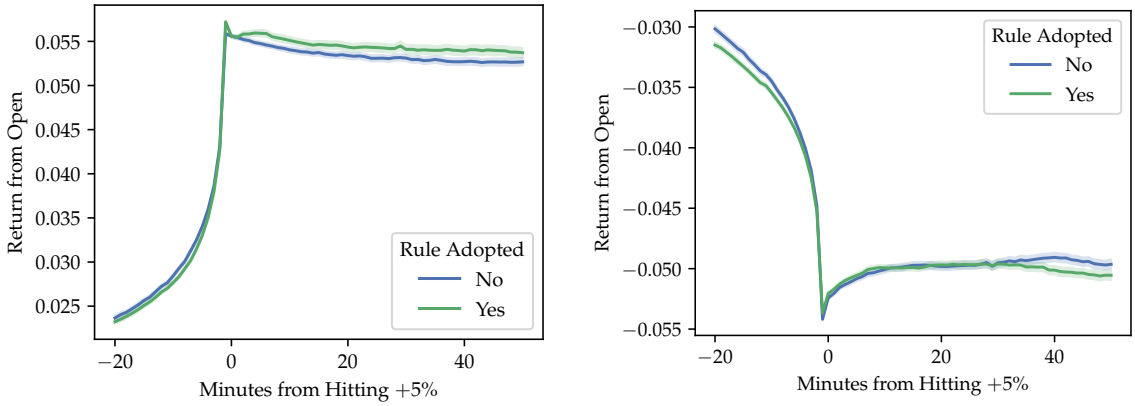

(a) +5% Breach

(b) −5% Breach

Figure 1.A3.10: Prices around ±5% Breaches. This figure plots average prices around ±5% breaches, in terms of return from open. Return from open is computed using the closing price of each 1-minute interval according to Section 1.2.2. 95% confidence intervals are constructed by bootstrapping.

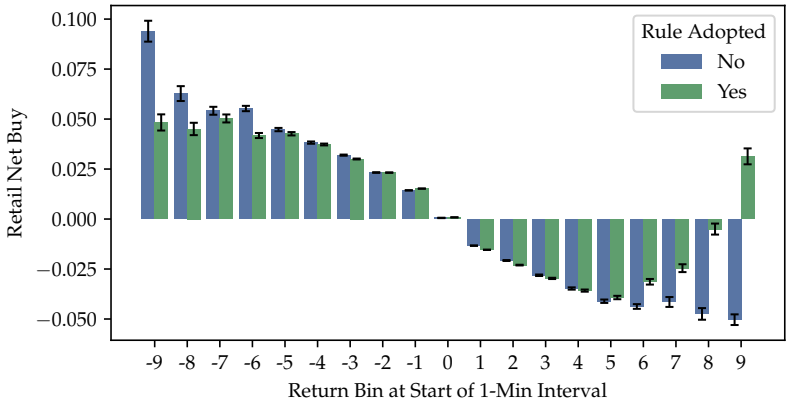

Figure 1.A3.11: Retail Net Buy by Return Bin. This figure plots the average 1-minute net buy conditional on the return from open at the beginning of the 1-minute interval. The averages are computed for both pre-rule (blue bars) and post-rule (green bars) periods. Bars at $x = b$ refers to the group of intervals whose price, in terms of return from open, at the beginning fall between $[b\%, (b+1)\%]$. 95% confidence bands are constructed by bootstrapping.

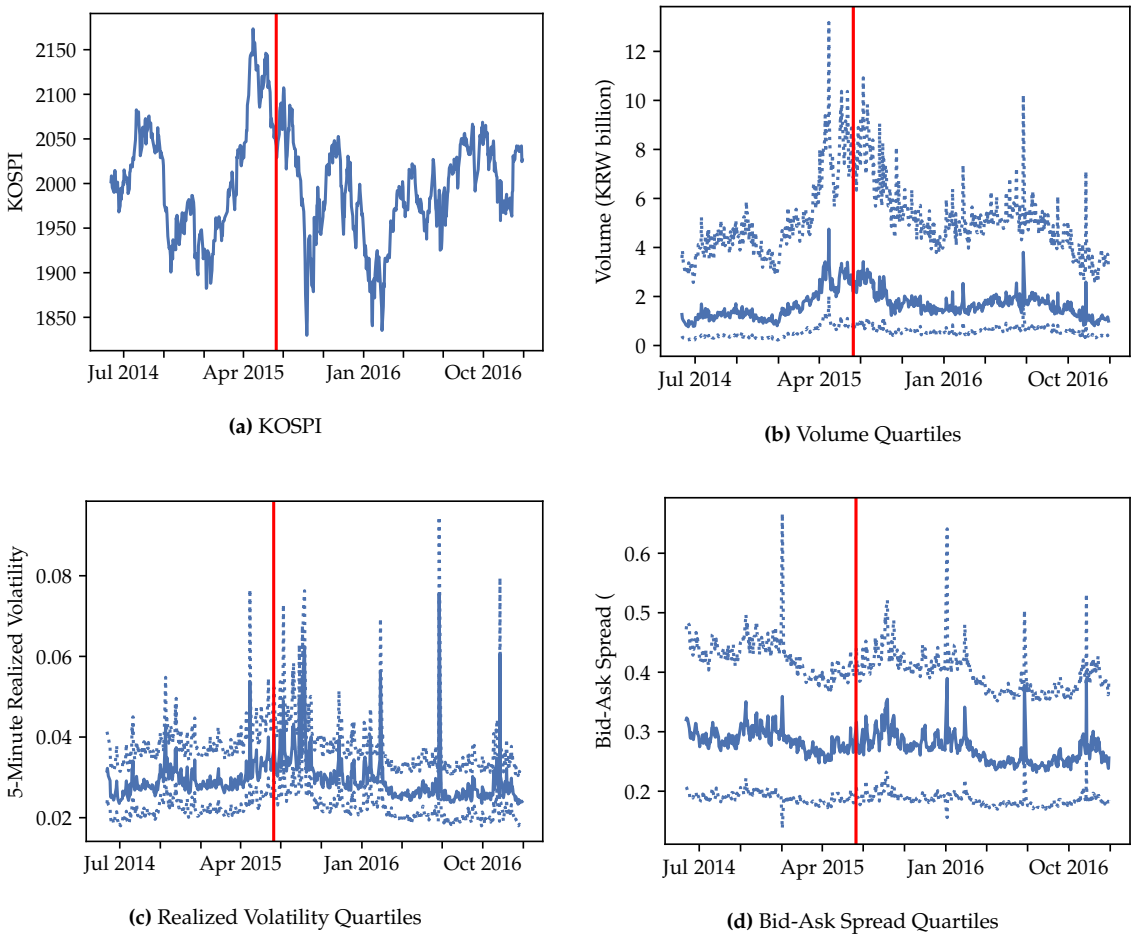

(a) KOSPI

(b) Volume Quartiles

(c) Realized Volatility Quartiles

(d) Bid-Ask Spread Quartiles

Figure 1.A3.12: Market Conditions around Rule Adoption. This figure plots the overall market conditions around the rule adoption. Panel (a) shows the level of KOSPI. Panel (b)–(d) plots the 25th, 50th, and 75th percentiles of volume, 5-minute realized volatility, and daily average bid-ask spreads in the cross-section of stocks.

Chapter 2

Retail Trading Intensity and the Overnight-Intraday Return Gap*

2.1 Introduction

In a continuing stream of work, researchers have documented a large gap between overnight and intraday returns (Kelly and Clark 2011; Berkman et al. 2012). Overnight returns (close-to-open) tend to be much higher than intraday returns (open-to-close) in both developed and emerging stock markets. Figure 2.1 replicates this pattern in the Korean stock market which is the setting of this paper. Daily returns are decomposed into their overnight and intraday components, and are compounded separately. One can easily notice the "overnight-intraday return gap" that has caught the attention of many academics and traders alike. Given that much of the corporate announcements arrive outside of trading hours, this may appear to be a natural pattern. However, papers have shown that the pattern and extent of the divergence are not fully explained by this narrative, and have sought other explanations.

* This research was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant. We thank Motohiro Yogo, Wei Xiong, Sangmin Oh, Jihong Song, Francisco Cabezón, Nicolas Hommel, Ziang Li, Michael Junho Lee, Jiwon Lee, Sohye Yoon, Narek Alexanian, Ziqiao Zhang, and Jiwon Choi for feedback and comments.

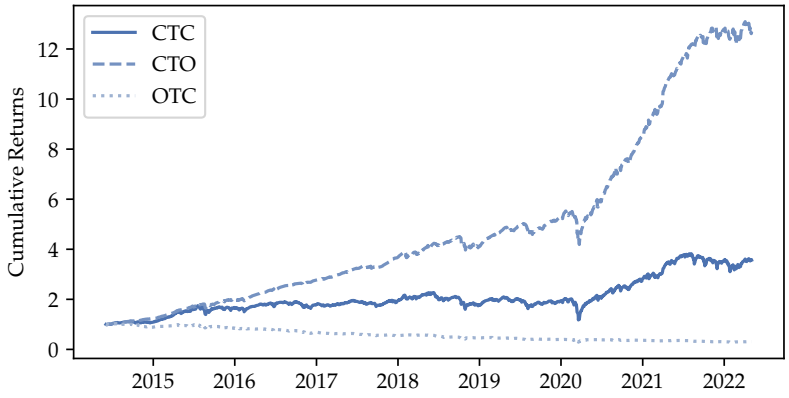

Figure 2.1: Cumulative Overnight and Intraday Returns. This figure shows cumulative returns for overnight (close-to-open), intraday (open-to-close), and daily (close-to-close) returns of the aggregate Korean stock market separately. The exact definition of each is given in Section 2.3.1.

Several studies may be described as proposing the "price impact narrative." Berkman et al. (2012) exemplifies this: the paper points to the fact that retail investors are much more active around the opening hours and attributes the return gap to opening price inflation caused by attention-induced buying by retail investors. A temporary upward price pressure induced by an inflow of retail investors near the market open creates a wave-like pattern that is consistent with both high overnight returns and ensuing reversals in intraday returns. In a similar spirit, Lou et al. (2019) presents a more general "clientele effect." The authors show that high overnight (intraday) returns during a given month predicts high overnight (intraday) returns in the following month. This is to be expected if a certain group of investors trades during a certain hour of the day, and if they do so in the same direction over several days or weeks.

A different view is the "overnight risk narrative." Similar to the view put forward by Savor and Wilson (2014), which shows that days with macroeconomic announcements pose inherently different risks, perhaps overnight periods are also distinct from intraday periods due to extreme illiquidity. Consistent with this idea, Hendershott et al. (2020) shows that variation in market beta explains the cross-section of overnight returns well, notwithstanding the well-documented fact that CAPM performs poorly in explaining the cross-section of close-to-close returns (Fama and French 2004). For a theoretical basis for this argument, we can also look to Hong and Wang (2000) who explicitly model the regular closure of markets.

In this paper, we bring both narratives together by considering the role of noise traders.[1] We

[1] Boehmer et al. (2021) shows evidence that retail flows appear to be informed, in the sense that retail order imbalances predict higher future returns. Barber et al. (2022c) shows that this alone may not necessarily imply that retail investors should be considered informed and Kim and Kim (2022) provide further evidence that the aggregate behavior of Korean retail investors tend to be prone to well known behavioral biases.

make two advances in this direction. The first is that we explicitly model the varying prominence of noise traders in the morning and afternoon to draw several empirical predictions about overnight and intraday returns that are consistent with previous empirical findings.[2] We extend the model developed in De Long et al. (1990) to include N assets that are differentially subject to aggregate noise trader misperception. Our model shows that the more exposed a stock is to noise trader risk, the higher the overnight-intraday return gap would be.

The second is empirical: we verify that the overnight-intraday return gap, at the stock-month level, is indeed closely tied to retail investor prevalence, measured by the retail investor volume share. This is made possible through a unique advantage of the Korean stock market data. Because Korea Exchange (KRX) reports the identity of the buyers and sellers for each stock trade up to the investor type, we can exactly identify the trades coming from retail investors. While recent innovation made in Boehmer et al. (2021) allows researchers to impute retail flows from the US trade and quotes (TAQ) data, these imputed flows remain a proxy for the exhaustive retail activity. Barber et al. (2022a) also raise questions about the accuracy and comprehensiveness of the methodology. Due to these issues, the Korean stock market proves to be especially suitable for this study.

Furthermore, we propose a new identification strategy. Previous studies discussed above explore the cross-sectional correlation between overnight returns and lagged characteristics, such as market beta or proxies of retail attention, but stop short of establishing a causal relationship. A naive specification under our setup is to regress the return gap on retail investor intensity. An obvious problem is that we have a price-related variable on the left while we have a quantity-related variable on the right.

To address this, we exploit the shock to aggregate retail investor trading intensity caused by COVID-19. Perhaps due to the lack of entertainment under stay-at-home orders or generous stimulus packages, a retail trading frenzy ensued in the months following March 2020 (Ozik et al. 2021). The Korean stock market experienced a similar surge in retail trading, as shown in Figure 2.2. Consistent with our prediction, the overnight-intraday return gap also widens significantly in the post-pandemic period.

[2] We do not currently microfound why noise trader intensity is higher in the morning but instead point to the robust empirical patterns shown in Figure 2.6.

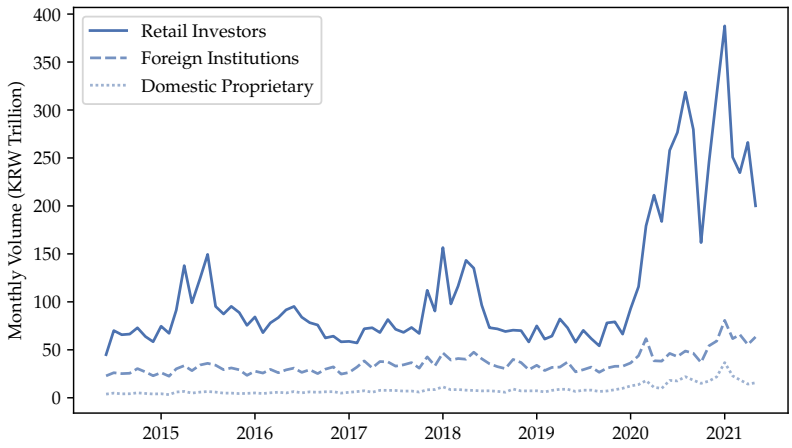

Figure 2.2: Monthly Volume by Investor Type. This figure shows the monthly trading volume of different investor types in KRW terms. The investor types considered are retail investors, foreign institutions, and domestic proprietary institutions.

But more importantly, we go from this increase in aggregate retail intensity to cross-sectional variation in retail volume shares by exploiting the fact that *shares* are capped at 1. Suppose there are only two stocks: a retail-dominant stock with 99% retail volume share and another stock with 50% retail volume share. Furthermore, suppose that an exogenous increase to general retail trading activity arrives and that retail investors allocate comparable volume to these two stocks. The first stock will experience a negligible increase in retail volume share relative to the second stock due to the fact that it was already close to the bound 100%. A more detailed example is discussed in Section 2.3.2 and Figure 2.8 shows this pattern empirically.

We present additional evidence backing the narratives related to heterogeneous groups of investors exerting price pressure in different times of the day. In our empirical analysis, we start by documenting intraday trading patterns of retail and institutional investors consistent with such clientele-based explanations. In particular, we verify that long-term institutions' net buys in the afternoon during a given month predict their afternoon net buys in the next month. We also reaffirm that, unconditionally, retail investors are net buyers in the morning while institutions are net buyers in the afternoon.

We also provide reasons to believe that overnight returns may be a proxy for underlying, unobservable retail activity (Aboody et al. 2018; Akbas et al. 2022). While the negative return predictability results of the above papers are not replicated in Korea, our results show why overnight returns are closely linked to sentiment through the retail (noise) traders. Because the prominence of retail investors manifest in the morning, high morning prices end up being good proxies of stock-level sentiments.

Previous studies have also looked at the divergence between overnight and intraday returns. Kelly and Clark (2011) compare Sharpe ratios of overnight and intraday returns of 5 common stock ETFs. They find that risk-adjusted returns are significantly higher during non-trading hours. Branch and Ma (2015) further show that high overnight returns tend to be followed by low intraday returns. While they speculate on several possible explanations, such as market maker trading patterns or retail investor activity near open, they do not investigate this further.

Studies are somewhat split on why the overnight-intraday return gap exists and persists. As discussed above, Berkman et al. (2012) and Lou et al. (2019) point to price pressure by certain groups of investors during certain times of the day. Heston et al. (2010) also find evidence of investors having a predictable demand for immediacy at certain times of the day. In a similar spirit, Jones et al. (2022) attribute this pattern to price pressure by extrapolative trading by investors. They find that high intraday returns lead to high morning order imbalances and high overnight returns in the following day, which is expected in presence of unsophisticated and short-constrained investors. On the other hand, Hendershott et al. (2020) views the patterns in overnight returns as compensation for risks of different nature while markets are closed. Bogousslavsky (2021) finds that mispricing corrects in the morning and worsens at the end of the day and attributes this to capital-constrained arbitrageurs closing their position at close.

Other more recent papers also investigate overnight returns, but focus on different aspects. Akbas et al. (2022) test the relationship between the number of daytime reversals—positive overnight returns followed by negative intraday returns—and subsequent monthly returns. They find that more reversals lead to higher returns in the following month. Boyarchenko et al. (2023) compute market index futures returns for each hour of the trading day and find large positive returns during opening hours of European markets (2-3 a.m. EST). They argue that this is due to overnight resolution of order imbalances from end of previous U.S. trading day.

This paper is also related to a broader literature on noise traders, investor sentiment, and retail flows. De Long et al. (1990) presents a model with rational arbitrageurs and optimistic noise traders. They show that prices can diverge from fundamentals when the unpredictable noise trader demand cannot be fully diversified away. We extend this model to multiple assets using an approach close to that of Hong and Sraer (2016). The difference is that Hong and Sraer (2016) focus on differential loading on fundamental risk, while we focus on loading on noise trader risk.

2.2 Model

Before moving on to the empirical results, we present a simple model to guide our empirical analyses. We adapt the model of noise traders in De Long et al. (1990). The main distinguishing feature we add is the notion of "morning" and "afternoon." Based on the robust empirical pattern of high retail investor intensity near the market open, we exogenously assume that morning periods have a larger mass of noise traders.[3] Empirical patterns in Figure 2.6 strongly support this assumption.

2.2.1 Model with One Risky Asset

Setup. Time is discrete and indexed by t. Each interval from t to $t+1$ should be considered a "day," and it is further divided into a "morning" and an "afternoon" sub-period. Figure 2.3 shows this graphically. As with De Long et al. (1990), we consider an overlapping generations model in which each agent lives for two sub-periods. There is no first period consumption nor labor supply decision. They only decide their asset allocation when young.

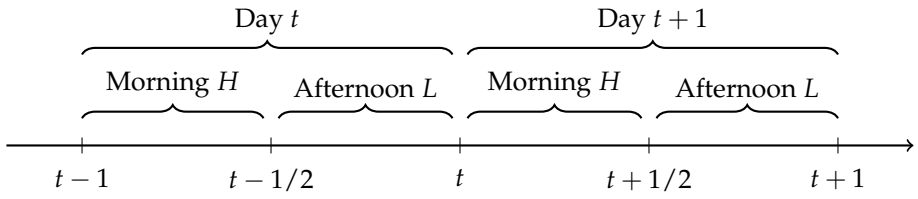

Figure 2.3: Time Setting of the Model. This figure shows how each period is divided into a morning and an afternoon period. Morning periods are marked by a larger mass of noise traders.

Assets. There are two types of assets that the agents can trade: the safe asset (s) and the risky (unsafe) asset (u). Both assets yield a dividend of r each period. The safe asset is in perfectly elastic supply, but the risky asset is in limited supply. Quantity of the risky asset is normalized to be 1. The price of the safe asset is fixed at 1 and serves as the numeraire. The price of the risky asset at time t is P_t.

Traders. There are two types of traders: rational arbitrageurs (a) and noise traders (n). The total mass of traders is 1 every period. During the morning periods, there is a mass μ_H of noise traders, while there is a mass μ_L of noise traders in the afternoon periods. Mornings have a larger mass of noise traders: $\mu_H > \mu_L$.

[3] Industry professionals admit that institutions avoid trading in the morning for this reason. In light of this, it is possible to imagine a model that endogenizes the sorting of noise traders and rational traders into the mornings and afternoons.

Both types of traders have CARA utilities with risk aversion γ, so that they maximize:

$$U = -\exp(-2\gamma w) \quad \text{or equivalently maximize} \quad \mathbf{E}(w) - \gamma \text{Var} w \tag{2.1}$$

where w is their wealth when old. A noise trader misperceives the expected price of the risky asset to be higher by:

$$\rho_t \sim N(\rho^*, \sigma_\rho^2) \tag{2.2}$$

where $\rho^* \geq 0$ measures the average "optimism" of the noise traders. The expectations of the noise traders are distinguished by the superscript n, while the t subscript denotes conditioning on time t information.

Each trader makes an allocation decision when young, which is the amount λ to invest in the risky asset and the amount ν to invest in the safe asset. Given that there are no labor income, λ fixes the amount of safe assets, through the budget constraint: $\lambda \cdot p_t + \nu \cdot 1 = 0$. Each unit of safe asset yields $(1 + r)$ while each risky asset yields $(p_{t+1} + r)$ in the next period. The optimal demand and market clearing conditions are identical to those in De Long et al. (1990), and are not included here for brevity (see Appendix B1).

Prices. Here, we consider a stationary equilibrium that is similar to the original paper. The only difference is that we now have two prices with their respective distributions: P^H and P^L. Suppose we are in a morning period with μ_H mass of noise traders. The market clearing condition for the risky asset gives us:

$$P_t = \frac{1}{1+r}(r + \mathbf{E}_t(P_{t+1/2}) - 2\gamma \text{Var}_t(P_{t+1/2}) + \mu_H \rho_t). \tag{2.3}$$

We will have an analogous equation for the afternoon period. By assuming fixed distributions for the prices at t and $t + 1/2$ respectively for $t \in \mathbb{Z}^+$, we can find the expected prices. The important point to note is that the difference between these expected prices will be positive:

$$\mathbf{E}(P^H) - \mathbf{E}(P^L) = \frac{1}{2+r}(\mu_H - \mu_L)\left(2\gamma(\mu_H + \mu_L)\left(\frac{\sigma_\rho}{1+r}\right)^2 + \rho^*\right) \geq 0. \tag{2.4}$$

This comes from the fact that the larger mass of optimistic noise traders raises prices and from the fact that risky assets in the afternoon faces a one-period larger noise trader risk. Also, as long as the μ_H and μ_L scale linearly, the overnight-intraday return gap increases with the general level of

noise trader mass.[4]

2.2.2 Extension to Multiple Assets

We now extend our model to an economy with N assets instead. The main difference would be that idiosyncratic "noise trader risk" may be diversified away when N is a very large number. However, we later demonstrate that the empirical reality may be far from this limit.

Factor Structure in Misperception. During day $t+1$, noise traders misperceive the expected morning and afternoon prices of stock i by $\rho_{i,t}$ and $\rho_{i,t+1/2}$ respectively:

$$\mathbf{E}_t^n(P_{i,t+1/2}) = \mathbf{E}_t(P_{i,t+1/2}) + \rho_{i,t} \qquad \mathbf{E}_{t+1/2}^n(P_{i,t+1}) = \mathbf{E}_{t+1/2}(P_{i,t+1}) + \rho_{i,t+1/2}. \tag{2.5}$$

The stock-level misperception has a factor structure:

$$\rho_{i,t} = b_i z_t + \varepsilon_{i,t} \tag{2.6}$$

where b_i is the factor loading, z_t is the common factor, and $\varepsilon_{i,t}$ is the idiosyncratic noise with $\mathrm{Cov}(z_t, \varepsilon_{i,t}) = 0$. The common factor z_t can be interpreted as the market-wide misperception.

The joint distribution of z_t and $\varepsilon_{i,t}$ is as follows:

$$z_t \overset{i.i.d.}{\sim} N(z^*, \sigma_z^2) \tag{2.7}$$

$$\varepsilon_{i,t} = \phi^{1/2}\varepsilon_{i,t-1/2} + \eta_{i,t} \text{ where } \eta_{i,t} \overset{i.i.d.}{\sim} N(0, \sigma_{\eta,i}^2) \text{ and } \mathbf{E}(\eta_{i,t}\eta_{j,t}) = 0,\ \forall i \neq j. \tag{2.8}$$

z^* is a measure of average market optimism of noise traders. We allow for persistence in the ε's to better match empirical findings such as those in Lou et al. (2019), but it is not crucial for the main cross-sectional predictions.

Assets. There are two types of tradable assets: stocks and the riskless bond. The supply of the riskless bond is perfect elastic and its daily risk-free rate is exogenously fixed at r. There are N stocks, and they are indexed by $i = 1, \ldots, N$. Each stock i has a fixed quantity Q_i of supply and pays dividends $d_{i,H}$ and $d_{i,L}$ in the morning and afternoon respectively. The prices of stock i in the morning and afternoon of day t are denoted by $P_{i,t-1/2}$ and $P_{i,t}$, and are endogenously determined in equilibrium.

[4] One simple example would be that $\mu_H = (1+\alpha)\mu$ and $\mu_L = \mu$, where μ is the general level of noise trader prevalence.

Traders. We use bold letters to denote vectors: e.g., $\boldsymbol{P}_t = (P_{1,t}, \ldots, P_{N,t})'$ and $\boldsymbol{Q} = (Q_1, \ldots, Q_N)'$. The agents' preferences are same as before, but now each type-τ investor, where $\tau \in \{n, a\}$, must choose her own *portfolio* demands $\boldsymbol{\theta}^\tau$ to maximize the expected utility from wealth W in next half-day. Specifically, $\theta^\tau_{i,t}$ is the type-τ investor's portfolio holding of stock i at time t. Thus, the problems for an investor of type τ in the morning and afternoon can be written as below.

$$\begin{aligned} \text{Morning:} \quad & \max_{\boldsymbol{\theta}^\tau_t} \mathbf{E}^\tau_t(W_{t+1/2}) - \gamma \text{Var}^\tau_t(W_{t+1/2}) \\ & \text{s.t.} \quad W^\tau_{t+1/2} = (\boldsymbol{d}_H + \boldsymbol{P}_{t+1/2} - (1+r)^{1/2}\boldsymbol{P}_t)'\boldsymbol{\theta}^\tau_t \end{aligned} \tag{2.9}$$

$$\begin{aligned} \text{Afternoon:} \quad & \max_{\boldsymbol{\theta}^\tau_{t+1/2}} \mathbf{E}^\tau_{t+1/2}(W_{t+1}) - \gamma \text{Var}^\tau_{t+1/2}(W_{t+1}) \\ & \text{s.t.} \quad W^\tau_{t+1} = (\boldsymbol{d}_L + \boldsymbol{P}_{t+1} - (1+r)^{1/2}\boldsymbol{P}_{t+1/2})'\boldsymbol{\theta}^\tau_{t+1/2} \end{aligned} \tag{2.10}$$

Prices and Returns. We solve for the stationary equilibrium price of stocks in morning and afternoon using the same strategy as before. The proposition below characterizes the equilibrium. As in Hong and Sraer (2016), we impose the condition $\boldsymbol{b}'\boldsymbol{Q} = 1$, which says that share-weighted misperception beta is equal to 1.

Proposition 1. *In the steady-state equilibrium, the price vector in the afternoon and morning are given by:*

$$\textit{Afternoon:} \quad \boldsymbol{P}_t = \frac{\boldsymbol{m}_H}{r} - \frac{2\gamma}{r}\Sigma_H \boldsymbol{Q} + \frac{\mu_H \boldsymbol{b}(z_t - z^*)}{(1+r)^{1/2}} + \frac{m_{H,\varepsilon}}{1+r-\phi}\boldsymbol{\varepsilon}_t \tag{2.11}$$

$$\textit{Morning:} \quad \boldsymbol{P}_{t+1/2} = \underbrace{\frac{\boldsymbol{m}_L}{r} - \frac{2\gamma}{r}\Sigma_L \boldsymbol{Q}}_{\textit{benchmark}} + \underbrace{\frac{\mu_L \boldsymbol{b}(z_{t+1/2} - z^*)}{(1+r)^{1/2}}}_{\textit{market misperception}} + \underbrace{\frac{m_{L,\varepsilon}}{1+r-\phi}\boldsymbol{\varepsilon}_{t+1/2}}_{\textit{idiosyncratic misperception}} \tag{2.12}$$

where $\boldsymbol{m}_H$ and $\boldsymbol{m}_L$ are $N \times 1$ vectors, Σ_H and Σ_L are $N \times N$ covariance matrices, and $m_{H,\varepsilon}$ and $m_{L,\varepsilon}$ are scalars. All of them are functions of $\boldsymbol{d}_H, \boldsymbol{d}_L, \boldsymbol{b}, \mu_H, \mu_L, r, \phi, z^*, \sigma^2_z, (\sigma^2_{\eta,i})_i$.

Proof. See Appendix B1. □

Proposition 2. *Asset i's equilibrium expected intraday return and overnight return are given by:*

$$\begin{aligned} R^H_{i,t} &= d_{i,H} + \mathbf{E}_t(P_{i,t+1/2}) - (1+r)^{1/2} P_{i,t} \\ &= 2\gamma \left(\frac{\mu_L^2}{1+r} b_i \sigma_z^2 + \frac{m_{L,\varepsilon}^2}{(1+r-\phi)^2} \sigma_{\eta,i}^2 Q_i \right) - \mu_H \rho_{i,t} \end{aligned} \tag{2.13}$$

$$\begin{aligned} R^L_{i,t+1/2} &= d_{i,L} + \mathbf{E}_{t+1/2}(P_{i,t+1}) - (1+r)^{1/2} P_{i,t+1/2} \\ &= 2\gamma \left(\frac{\mu_H^2}{1+r} b_i \sigma_z^2 + \frac{m_{H,\varepsilon}^2}{(1+r-\phi)^2} \sigma_{\eta,i}^2 Q_i \right) - \mu_L \rho_{i,t+1/2}. \end{aligned} \tag{2.14}$$

Proof. See Appendix B1. □

Similar to the single asset case, we find that the overnight-intraday return gap increases with the difference $(\mu_H^2 - \mu_L^2)$. In the cross-section, we see that higher misperception factor loading b_i and higher idiosyncratic variance in misperception $\sigma_{\eta,i}^2$ both lead to higher expected return gaps. As the number of assets $N \to \infty$, the latter term disappears so that only the systematic component matters. The main takeaway we want to emphasize is that as long as $(\mu_H^2 - \mu_L^2)$ is increasing in the general level of retail or noise trader activity, stocks exposed to more noise trader risk will see higher overnight-intraday return gaps.

2.3 Data and Methodology

2.3.1 Data and Stylized Facts

Overnight and Intraday Returns. We use 1-minute frequency price data for common and preferred stocks purchased from KRX. Our price data covers the period from June 2014 to May 2021. There are approximately 2,400 stocks listed on KRX as of 2021, but many of them are small and illiquid compared to the average listed US company. For this reason, we use a subset of 901 stocks (1) that traded every day throughout our sample period, (2) whose average market capitalization larger than KRW 50 billion, and (3) whose average daily trading volume is larger than KRW 500 million.

The overnight return, or the close-to-open (CTO) return, and the intraday return, or the open-

to-close (OTC) return, for stock i on day d are computed as follows:

$$\text{CTO}_{id} = \frac{P_{id}^{\text{open}} - P_{id-1}^{\text{close}}}{P_{id-1}^{\text{close}}} \tag{2.15}$$

$$\text{OTC}_{id} = \frac{P_{id}^{\text{close}} - P_{id}^{\text{open}}}{P_{id}^{\text{open}}}. \tag{2.16}$$

The usual daily return, or the close-to-close (CTC) return, can be recovered through the following:

$$1 + \text{CTC}_{id} = (1 + \text{CTO}_{id}) \times (1 + \text{OTC}_{id}) = \frac{P_{id}^{\text{close}}}{P_{id-1}^{\text{close}}}. \tag{2.17}$$

This decomposition is as in Berkman et al. (2012) and Figure 2.4 demonstrates this breakdown graphically.

There are two institutional details worth mentioning. The first is that, as in many other stock exchanges, opening and closing prices are determined through opening and closing auctions in Korea. These auctions aggregate buy and sell orders during the 30 minutes (10 minutes) prior to market open (close), and determine the opening (closing) uniform auction prices. These opening and closing prices are P^{open} and P^{close} above respectively. The second detail is that the closing time of the Korean stock market was extended by 30 minutes in August 2016, namely from 3:00 p.m. to 3:30 p.m.

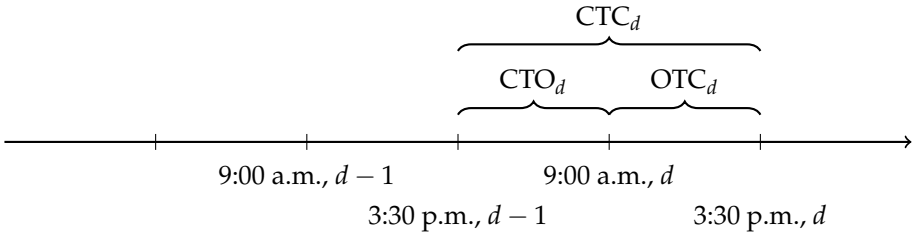

Figure 2.4: Overnight and Intraday Returns. This figure shows how overnight and intraday returns are defined. Each day's return (close-to-close) are decomposed into overnight return (close-to-open) and intraday return (open-to-close). The Korean stock market opens at 9:00 a.m. each day and closes at 3:30 p.m. During periods before August 1st of 2016, the closing time was 3:00 p.m.

Figure 2.1 shows the cumulative returns, ignoring transaction costs, for the CTO and OTC returns separately. For each day d, equal weighted overnight and intraday returns are computed by averaging across the different stocks. Figure 2.1 plots the following against time:

$$\prod_d (1 + \text{CTO}_d) = \prod_d \left(1 + \frac{1}{N}\sum_i \text{CTO}_{id}\right) \quad \text{and} \quad \prod_d (1 + \text{OTC}_d). \tag{2.18}$$

We can see that average overnight returns are much higher than average intraday returns. Also, this divergence becomes even starker after the COVID-19 crash of March 2020.

Later in our analyses, we aggregate observations to the stock-month level, as in Lou et al. (2019). We compute the *overnight-intraday return gap* of stock i during month t as the difference between the cumulative CTO returns and the cumulative OTC returns:

$$\text{“Return Gap”} = \text{RG}_{it} = \frac{\prod_{d \in \text{month } t}(1 + \text{CTO}_{id})}{\prod_{d \in \text{month } t}(1 + \text{OTC}_{id})} - 1. \tag{2.19}$$

Trade Flows by Investor Type. An important advantage of the KRX data is that KRX provides identifiers for the investor type at the *individual trade level* through their trades and quotes (TAQ) data, albeit at a significant price. The provided investor types are: (1) retail investors, (2) foreign institutions, (3) proprietary traders, (4) asset managers, (5) pension funds, (6) banks, (7) insurance companies, and (8) private equity. All investor types except retail investors and foreign institutions refer to domestic institution types. At a more reasonable price, KRX also provides trade flows aggregated at the investor type level at 1-minute frequencies. We use the 1-minute frequency trade flow data over the period from June 2014 to May 2021.

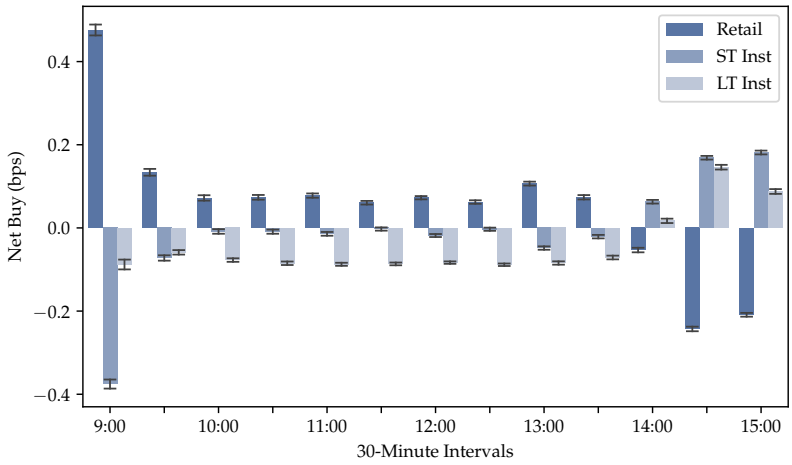

Figure 2.5: Intraday Net Buy Patterns. This figure shows the pooled average values for normalized net buys by different investor types during different times of the day. Net buys are normalized by shares outstanding. Each observation is at the stock-day-interval-investor level and each time interval is 30 minutes long. 95% confidence intervals are computed based on bootstrapping and are displayed.

Figure 2.5 plots the average *net buy* by retail investors, short-term institutions (foreign institutions and domestic proprietary traders), and long-term institutions (other institutions) at different times of the day. In some analyses, institutions are grouped this way because many of the institutional types, such as private equity and insurance companies, play too small a role in sub-daily

frequencies. For each stock i and investor type j, trades are aggregated at the 30-minute interval level, starting from 9:00 a.m. Suppose these intervals are indexed by t. Then, net buy and normalized net buy are computed as follows:

$$\text{NB}_{ijt} = \#\text{ shares of stock } i \text{ bought minus sold by investor type } j \text{ during interval } t \tag{2.20}$$

$$\text{nb}_{ijt} = \frac{\text{NB}_{ijt}}{\#\text{ shares outstanding}_{it}}. \tag{2.21}$$

Because markets have to clear during any interval $t \to t+1$, we must have:

$$\sum_j \text{NB}_{ijt} = 0 \text{ and } \sum_j \text{nb}_{ijt} = 0 \text{ for any stock } i \text{ and time } t. \tag{2.22}$$

Unconditionally, we should expect to see zero net buys for any randomly chosen (i, j, t).

We compute the pooled averages—across stocks and days—of the net buys nb_{ijt} at different 30-minute intervals. Figure 2.5 plots these averages for retail investors and short-term institutions. We see a clear pattern: retail investors are net buyers in the morning and net sellers in the afternoon.

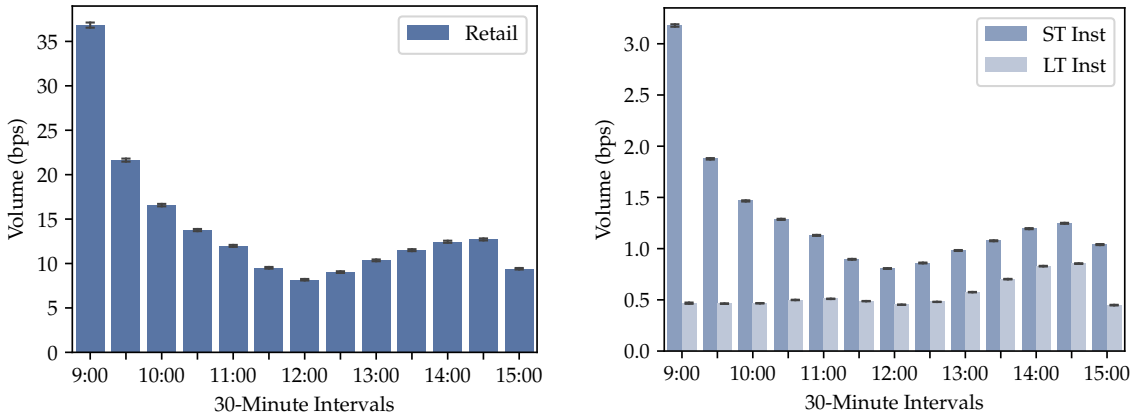

(a) Retail Volume at Different Times

(b) Institutional Volume at Different Times

Figure 2.6: Intraday Volume Patterns. This figure shows the pooled average values for normalized volume by different investor types during different times of the day. Volumes are normalized by shares outstanding. Each observation is at the stock-day-interval-investor level and each time interval is 30 minutes long. 95% confidence intervals are computed based on bootstrapping and are displayed.

We repeat the same procedure with (normalized) volume by investor types. The results are displayed in Figure 2.6. While volume is higher in the morning for both retail and short-term institutions, long-term institutions trade more in the afternoon. This pattern is consistent with short-term institutions providing liquidity to retail investors, while long-term institutions avoiding the turbulent mornings. If we sum up all three parties, we see that the volume share of retail

investors is higher in the morning.

Another interesting exercise we can conduct is to compute average *relative buy-sell prices* by investor types. With these 1-minute frequency trade flow data, we can compute the average buy and sell prices for investor type j type as follows:

$$\text{Avg Buy (Sell) Price}_{ijt} = \frac{\text{Buy (Sell) Amount}_{ijt}}{\text{Buy (Sell) Volume}_{ijt}}. \tag{2.23}$$

Because the volume is in number of shares, above gives the average price per share within the interval t. To make fair comparisons, we keep 30-minute intervals during which all investor types have nonzero buy and sell trades. Then, for each i and t, we take the average retail buy (sell) price, divide it by the average buy (sell) price for the short-term institutions, take the pooled average across stocks and days, and subtract one. This produces average relative price deviation for each time of the day. Figure 2.7 plots the relative price deviation against times of the day.

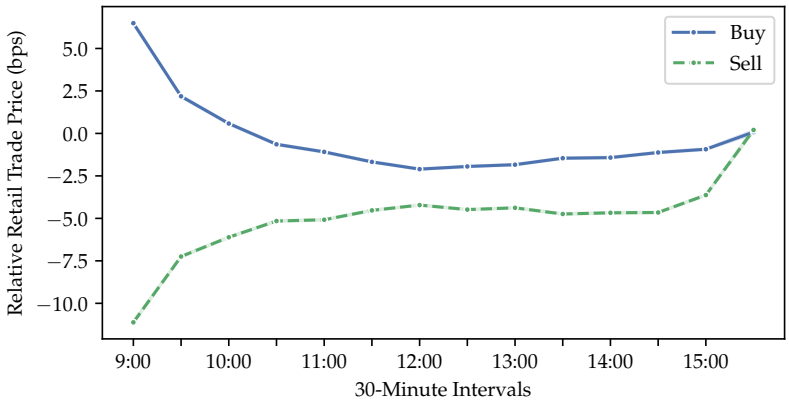

Figure 2.7: Relative Retail Trade Prices. This figure shows the pooled average values for the retail investors' buy and sell price deviations relative to those of the short-term institutions. The deviations are expressed in basis points. Each observation is at the stock-day-interval-investor level and each time interval is 30 minutes long.

We see that retail investors buy (sell) at higher (lower) prices in the morning relative to short-term institutions. In other words, they are demanding more immediacy close to the open.

Another important variable of interest is the *retail volume share*. This is simply the proportion of retail volume relative to the total volume for any stock i and time interval t:

$$\text{"Retail Share"} = \text{RS}_{it} = \frac{\text{Retail Buy Volume}_{it} + \text{Retail Sell Volume}_{it}}{\text{Total Buy Volume}_{it} + \text{Total Sell Volume}_{it}}. \tag{2.24}$$

Firm Characteristics. KRX also provides various daily statistics such as short volume, short interest, book-to-market ratio, dividend yield, and shares outstanding on its website. For regulatory reasons, foreign institutional holdings are also reported at a daily frequency, but holdings of other investor types are not directly reported. These firm characteristics are obtained directly from the KRX website and values from the latest trading day are used later as control variables in stock-month-level regressions. We present the summary statistics for the return gap, firm characteristics, and retail shares in Table 2.1.

	Mean	SD	25th	50th	75th
Retail Vlm Share	75.14	20.65	64.33	83.75	90.29
Cum. Overnight Ret	3.02	9.11	−1.64	1.91	6.09
Cum. Intraday Ret	−1.31	12.49	−8.43	−2.29	4.41
Monthly Ret Gap	6.31	18.59	−4.25	4.29	13.95
Foreign Ownership	10.02	13.73	1.63	4.22	12.94
Log Market Cap	26.19	1.61	25.18	25.91	27.05
Book-to-Market	0.90	0.71	0.42	0.73	1.19
Amihud	4.62	18.05	0.24	0.69	1.77

Table 2.1: Summary Statistics of Stock-Month Observations. This table reports the summary statistics for the subset of chosen stocks. Each observation is at the stock-month level and variables are winsorized at 1%. Share variables are expressed in percentages. Cumulative overnight returns, intraday returns, and return gaps are computed according to equations (2.18) and (2.19). Amihud measure is computed using absolute returns in percentages and trading volume in KRW billions. Market capitalization is expressed in KRW billions and then log-transformed.

	Ret Gap	Retail Share	Log Market Cap	Foreign Own	B/M
Retail Vlm Share	0.17				
Log Market Cap	−0.07	−0.82			
Foreign Ownership	−0.09	−0.68	0.65		
Book-to-Market	−0.10	−0.10	−0.04	0.07	
Amihud	0.01	0.32	−0.44	−0.23	0.15

Table 2.2: Correlation Matrix of the Main Variables. This table reports the time-series average of the cross-sectional correlation matrices of each month. First, the cross-sectional correlations among the variables for a given month are computed. Then, these matrices are averaged over the months.

2.3.2 Empirical Design

Difference-in-Differences. In one of our empirical exercises, we regress the return gap on the retail share with other control variables as a starting point. Threat to identification arises from the fact that this naive specification has a price-related variable on the left and a quantity-related variable on the right. It may also be subject omitted variables bias. Consider, for example, hard-to-value stocks. These stocks may attract a lot of retail trading (Laarits and Sammon 2022) and also experience large price swings that lead to large return gaps. Thus, we need an exogenous source of variation in retail shares.

Suppose there are two stocks (A and B), and two agents (the representative retail investor and the representative institution). Also, suppose that the retail (institutional) investor allocates α ($\hat{\alpha}$) share of volume to asset A. If retail and institutional volume are V^r and V^i respectively, the retail volume share of assets A and B are as the following:

$$\text{RS}_A = \frac{\alpha V^r}{\alpha V^r + \hat{\alpha} V^i} \qquad \text{RS}_B = \frac{(1-\alpha) V^r}{(1-\alpha) V^r + (1-\hat{\alpha}) V^i}. \tag{2.25}$$

Consider a scenario where αV^r is large relative to $\hat{\alpha} V^i$ so that $\text{RS}_A \approx 1$. In other words, A is a retail-dominated stock. When an exogenous shock to V^r arrives, such as the retail frenzy following the COVID-19 stock market crash, we can see that RS_A will not increase significantly due to the fact that *shares* are capped at 1. On the other hand, RS_B increases more significantly, assuming $\text{RS}_B \ll 1$.

Figure 2.8 demonstrates this idea empirically. Here, we compute each stock's average daily retail volume share during the years 2014 to 2018. Then, we sort stocks into quintiles based on this initial retail volume shares. We next compute monthly retail volume shares for each stock in the months around the COVID-19 shock. The average retail volume share for different quintiles are plotted in Figure 2.8. We see that the 5th quintile is always close to 1, while the 1st quintile experiences a significant increase around the COVID-19 shock.

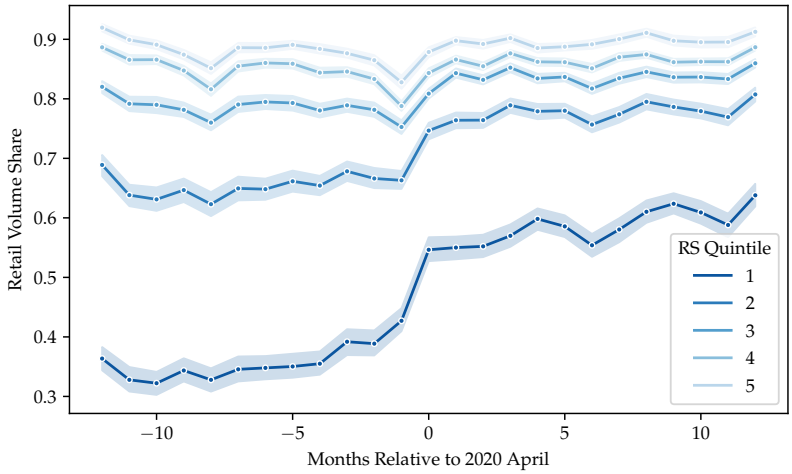

Figure 2.8: Retail Shares Over Time by Initial RS Quintiles. This figure shows the average retail investor volume shares for stocks within each initial retail share quintile. Stocks are first assigned into quintile portfolios based on their average retail volume share during the years 2014 to 2018. Then, average monthly retail volume shares are computed for each quintile group and plotted. 95% confidence intervals are computed based on bootstrapping and are displayed.

The implicit assumption behind the argument related to equation (2.25) is that the allocation of retail volume is somewhat stable over time. Figure 2.9 provides some support for this assumption. We see that even after the large jump in total retail volume shown in Figure 2.9a, the proportion of trading volume flowing to different initial retail share quintiles shown in Figure 2.9b is quite stable.

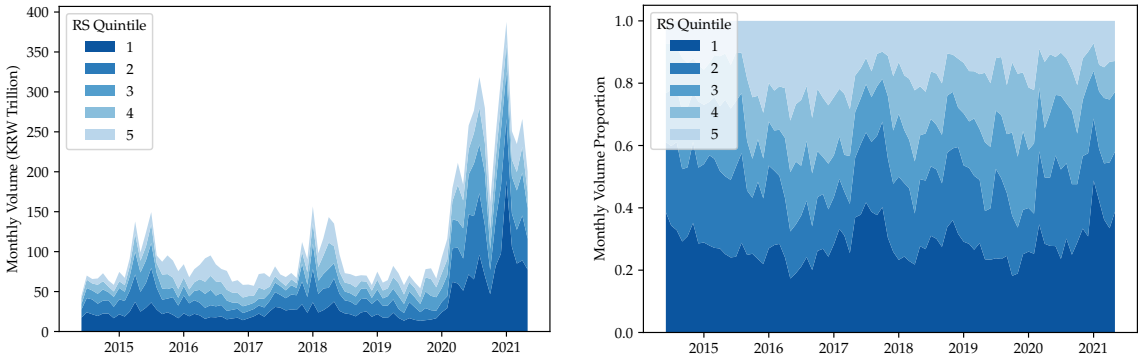

(a) Monthly Retail Volume by RS Quintile

(b) Monthly Retail Volume Proportion by RS Quintile

Figure 2.9: Monthly Retail Volume by RS Quintile. This figure shows the total monthly retail investor volume for each RS quintile in terms of KRW trillions and in terms of proportions. Stocks are assigned into quintile portfolios based on their current month's retail volume share and retail trading volume is summed within each quintile.

Portfolio Sorts. We also test whether variables related to retail investor intensity can predict future return gaps. Conceptually, these variables are intended to capture either the idiosyncratic and systematic components of noise trader risk. We will demonstrate that variables such as lagged re-

tail volume share, variance of daily retail net buy, and "the retail net buy beta" are all positively correlated with future overnight-intraday return gaps. However, one issue we will encounter is that model-based quantities do not exactly map one-to-one to the empirical ones at the current stage. To include other firm-level control variables, we accompany the results with Fama-MacBeth regressions.

2.4 Empirical Results

2.4.1 Retail and Institutional Flow

As shown in Figures 2.5 and 2.6, different groups of "clientele" tend to prefer different times of the day to trade. A more specific claim made in Lou et al. (2019). Because different groups of investors tend to trade during different times of the day, and because the demands tend to be persistent, this should lead to intraday price patterns being persistent as well. The authors analyze the intraday price patterns instead because the trade flows of different clientele are not generally available.

Our setup has the advantage that such claims can be directly tested to some extent. One approach is to examine the morning and afternoon net buys of different investor types. For concreteness, we define mornings to be the first two hours of the opening and the afternoon to be the last three hours of trading. Also, as before, we divide the groups of investors into (1) retail investors, (2) short-term institutions, and (3) long-term institutions. We then aggregate the morning net buys, normalized by shares outstanding, of each group at the stock-month level. One limitation is that heterogeneity within one of the three groups cannot be considered.

The regression specification as follows:

$$\mathrm{NB}_{ijt}^{\mathrm{morning}} = \beta_1 \cdot \mathrm{NB}_{ijt-1}^{\mathrm{morning}} + \beta_2 \cdot \mathrm{NB}_{ijt-2}^{\mathrm{morning}} + \beta_3 \cdot \mathrm{NB}_{ijt-3}^{\mathrm{morning}} + \Gamma' \mathbf{X}_{ijt} + \alpha_i + \gamma_t + \varepsilon_{ijt} \tag{2.26}$$

where $\mathrm{NB}_{ijt}^{\mathrm{morning}}$ is the morning net buys of investor type j in stock i during month t. Controls $\mathbf{X}_{ijt}$ include size, book-to-market, Amihud measure, foreign ownership, and past month's return. Stock fixed effects α_i and month fixed effects γ_t are included. The same specification is used for $\mathrm{NB}_{ijt}^{\mathrm{afternoon}}$ as well.

	Retail NB Morning	ST Inst Morning	LT Inst NB Morning
	(1)	(2)	(3)
NB Lag 1	0.002	0.023***	0.020***
	(0.005)	(0.004)	(0.004)
NB Lag 2	0.004	0.004	−0.008**
	(0.005)	(0.004)	(0.004)
NB Lag 3	0.004	0.000	−0.003
	(0.005)	(0.004)	(0.004)
Stock FE	✓	✓	✓
Month FE	✓	✓	✓
Stock Controls	✓	✓	✓
Observations	41,920	58,930	46,528
R^2	0.002	0.003	0.003

Table 2.3: Persistence of Morning Flows. This table reports the regression results using the specification in equation (2.26). Only lagged morning net buys of the same investor type are added as explanatory variables. Standard errors are two-way clustered at the stock and month level. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

	Retail NB Afternoon	ST Inst Afternoon	LT Inst NB Afternoon
	(1)	(2)	(3)
NB Lag 1	0.069***	0.023***	0.039***
	(0.004)	(0.004)	(0.004)
NB Lag 2	0.025***	0.004	0.014***
	(0.004)	(0.004)	(0.004)
NB Lag 3	0.012***	0.000	−0.006
	(0.004)	(0.004)	(0.004)
Stock FE	✓	✓	✓
Month FE	✓	✓	✓
Stock Controls	✓	✓	✓
Observations	72,657	58,930	57,159
R^2	0.008	0.003	0.003

Table 2.4: Persistence of Afternoon Flows. This table reports the regression results using the specification in equation (2.26), but uses afternoon net buys instead. Only lagged afternoon net buys of the same investor type are added as explanatory variables. Standard errors are two-way clustered at the stock and month level, and are shown in parentheses. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

Results are shown in Table 2.3 and Table 2.4. it should be noted that the R^2 on all the regressions

are very low because net buys are very noisy. Still, we see that the afternoon results are consistent with the Lou et al. (2019) hypothesis. We see that between 2% and 7% of each investor group's afternoon flows persist to the next month on average. The pattern is stronger for long-term investors, which is reasonable given that they would move to gradually accumulate positions. One caveat, however, is that net buys have to sum to zero for any given stock i and time t. Thus, the positive coefficient on retail investors in Table 2.4 may well be merely reflecting the pattern induced by long-term institutions.

2.4.2 Portfolio Sorts on Retail Investor-Related Variables

Here, we examine whether proxies for the model-based quantities have predictive power in terms of future overnight-intraday return gaps. As a proxy for the stock-level optimism $\rho_{i,t}$, we use daily retail net buys normalized by shares outstanding. Analogously, we use the daily aggregate retail net buy, KRW amount of retail net buy normalized by KRW amount of total market capitalization, as the proxy for z_t. Using these quantities, we compute the rolling one year factor loadings for each stock. More specifically, we compute the rolling betas in the following way:

$$\mathrm{NB}_{it}^{\mathrm{retail}} = \alpha_i + \beta_i \cdot \mathrm{NB}_t^{\mathrm{retail}} + \varepsilon_{it} \tag{2.27}$$

where $\mathrm{NB}_t^{\mathrm{retail}} = \frac{\text{Aggregate Retail NB in Day } t}{\text{Aggregate Market Cap in Day } t-1}$. We call the estimated betas *retail net buy (NB) betas*. We repeat the procedure with retail gross buy volume normalized by shares outstanding and call the resulting betas *retail buy volume (BV) betas*. Finally, we also compute the rolling variances of the stock-level retail net buys and use it as a sorting variable as well. The raw variances are meant to proxy for the $\sigma_{\eta,i}^2$ values in the model. While the model predicts that this should not affect the return gap with large N, empirically we find that the variance has more predictive power. At the end of every month, stocks are sorted based on the rolling betas and variances, and are assigned deciles. Equal-weighted average return gaps of the deciles are computed for the following month. As a reference, we repeat the exercise using the lagged retail volume share as the sorting variable.

Sort Variable	Decile	Overnight Ret	Intraday Ret	Ret Gap
Retail NB Var	1	1.61%	−0.08%	3.17%
		(3.4)	(−0.28)	(7.98)
	10	3.81%	−3.02%	9.15%
		(7.6)	(−7.88)	(14.41)
	10–1	2.20%	−2.94%	5.98%
		(5.25)	(−7.48)	(12.62)
Retail NB Beta	1	2.63%	−1.81%	6.30%
		(5.72)	(−6.57)	(10.5)
	10	3.60%	−2.40%	7.95%
		(7.44)	(−7.08)	(13.77)
	10–1	0.97%	−0.59%	1.65%
		(4.98)	(−1.84)	(3.45)
Retail Vlm Share	1	0.21%	0.28%	0.77%
		(0.42)	(1.41)	(1.51)
	10	5.93%	−3.80%	13.85%
		(6.6)	(−8.14)	(9.51)
	10–1	5.72%	−4.08%	13.08%
		(6.62)	(−8.7)	(8.89)

Table 2.5: Portfolio Returns for Univariate Sorts. This table reports overnight/intraday return patterns. At the end of each month, all stocks are sorted into deciles based on the different sort variables shown. We compute *t*-statistics, shown in parentheses, based on standard errors corrected for serial-dependence with 6 lags.

Table 2.5 displayed the results. At first sight, we see that portfolios of high beta and high variance assets have higher return gaps, which is consistent with the model. However, there are a few concerning issues at this stage.

One issue is that at the stock-day level, retail net buys are extremely noisy. This means that the one-factor specification does not have much explanatory power and most of the variation comes from the idiosyncratic noise component.[5] In other words, our proxies for b_i are not good enough. This also causes the problem of a U-shaped decile portfolio returns, rather than a monotone shape.

Another issue is that a simple sort based on retail volume share does a much better job, as we can see in the last three rows of Table 2.5. We will see in the next section that contemporaneous

[5] Dou et al. (2022) offers a different way to compute "flows." The authors residualize the flows first to remove the time fixed effects and return-chasing effects. We try this specification as well, but results are similar due to the extremely noisy nature of daily net buys.

retail volume share is an important explanatory variable for the overnight-intraday return gap. In addition, as discussed in Laarits and Sammon (2022), retail volume share tends to be very stable over time. In this respect, our model-based quantities are currently not capturing more information related to noise trader risks relative to this simpler measure.

For completeness, we repeat the findings with Fama-MacBeth regressions and the results are displayed in Table 2.6.

	Overnight-Intraday Return Gap					
	(1)	(2)	(3)	(4)	(5)	(6)
Retail NB Var	0.01***	0.01***				
	(0.00)	(0.00)				
Retail NB Beta			0.74***	0.74***		
			(0.18)	(0.16)		
Retail Vlm Share					0.14***	0.23***
					(0.01)	(0.02)
Log Market Cap		−0.20		−0.54***		1.60***
		(0.16)		(0.16)		(0.18)
Book-to-Market		−1.88***		−2.32***		−1.64***
		(0.24)		(0.25)		(0.25)
Amihud Measure		0.31***		0.03		0.01
		(0.09)		(0.09)		(0.09)
Foreign Holdings		−0.07***		−0.07***		0.02**
		(0.01)		(0.01)		(0.01)
Observations	70,324	64,613	70,324	64,613	70,324	64,613
R^2	0.01	0.02	0.00	0.01	0.02	0.03

Table 2.6: Fama-MacBeth Regressions. This table reports Fama-MacBeth regressions of monthly excess stock returns on lagged firm characteristics. The dependent variable is the monthly overnight-intraday return gap. Each column uses different explanatory variables that capture exposure to retail investor flows. Standard errors are shown in brackets. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

2.4.3 Return Gap and Retail Volume Share

Considering the results above, we next examine the contemporaneous relationship between monthly returns gaps and retail volume shares. Our single-asset model predicts a positive relationship between the two quantities, although our multi-asset model does not directly link this quantity to the

return gap at this point.[6] We start with a naive panel ordinary least squares (OLS) specification with stock-month-level controls and fixed effects. The regression specification is as follows:

$$\text{RG}_{it} = \beta \cdot \text{RS}_{it} + \Gamma' \mathbf{X}_{it} + \alpha_i + \gamma_t + \varepsilon_{it} \tag{2.28}$$

where RG_{it} is the overnight-intraday return gap for stock i during month t and RS_{it} is the retail volume share for stock i during month t. Controls $\mathbf{X}_{it}$ include size, book-to-market, Amihud measure, and foreign ownership. Stock fixed effects α_i and month fixed effects γ_t are included.

	Ret Gap		Overnight Ret		Intraday Ret	
	(1)	(2)	(3)	(4)	(5)	(6)
Retail Vlm Share	0.28^{***}	0.29^{***}	0.17^{***}	0.19^{***}	-0.07^{***}	-0.05^{***}
	(0.01)	(0.02)	(0.00)	(0.02)	(0.00)	(0.01)
Log Market Cap	2.21^{***}	6.75^{***}	1.16^{***}	0.20	-0.92^{***}	-5.51^{***}
	(0.08)	(0.65)	(0.04)	(0.27)	(0.06)	(0.42)
Book-to-Market	-1.42^{***}	-0.80	-0.12^{***}	-0.76^{***}	0.94^{***}	-0.17
	(0.09)	(0.63)	(0.05)	(0.27)	(0.07)	(0.36)
Amihud Measure	0.01^{**}	0.01	0.01^{***}	0.01	0.00	-0.00
	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)
Foreign Holdings	0.03^{***}	0.03	0.04^{***}	0.04^{***}	0.01^{**}	0.01
	(0.01)	(0.02)	(0.00)	(0.01)	(0.01)	(0.02)
Stock FE	–	✓	–	✓	–	✓
Month FE	–	✓	–	✓	–	✓
Observations	73,635	73,635	73,635	73,635	73,635	73,635
Within R^2	0.03	0.04	0.04	0.04	0.01	0.03

Table 2.7: Return Gap vs Retail Volume Share (Panel OLS). This table reports the regression results using the specification in equation (2.28). In columns (2), (4), and (6), standard errors are two-way clustered at the stock and month level. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

Results are shown in Table 2.7. As expected, coefficient estimates for β are positive and significant for the first two columns. The coefficient estimate in column (2) indicates that a one percentage point increase in retail volume share leads to a 0.29% increase in the monthly return gap. We also see that this comes from both the overnight and intraday legs in the other four columns.

What is interesting is that the coefficient on size is positive. This appears at first counterintuitive

[6] While aggregate values of μ_H and μ_L are discussed, it is unclear how to translate this into cross-sectional variation of the retail volume shares.

given that Berkman et al. (2012) finds smaller stocks to have higher overnight returns. The pattern is due to the fact that retail volume share is the main driver and size is negatively correlated with this variable. Figure 2.10 is revealing in this respect. In Figure 2.10, stock-month observations are sorted on the retail volume share and are assigned into 20 bins of equal sizes. A binscatter plot is done using this binning, and this processed is repeated separately for different size terciles. We see that the return gap increases with the retail volume share within each size tercile, but also that smaller stocks have a high retail volume share in aggregate.

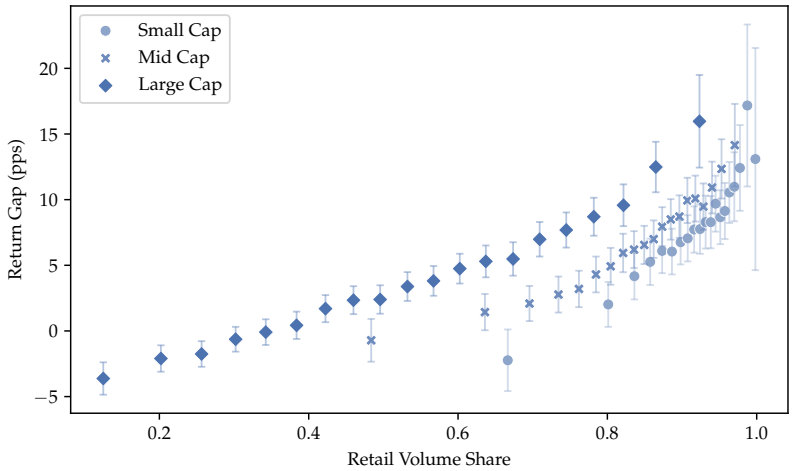

Figure 2.10: Binscatter of Return Gap against Volume Share. This figure shows a binscatter of the monthly return gap against the retail volume shares. Each observation is at the stock-month level. Binscatter is conducted using 20 bins and bins are created so that each bin contains the same number of observations. The procedure is done separately for different size terciles. 95% confidence bands are displayed.

As discussed in Section 2.3.2, this first-pass specification may be subject to omitted variables bias. To address this concern, we focus on the fact that retail volume shares are capped at 1. When the plausibly exogenous shock to aggregate retail trading intensity arrived with the pandemic—either due to the higher number of retail investors working from home or the government stimulus packages—only stocks whose retail volume shares were sufficiently away from 1 experienced significant increases in the retail volume share. Figure 2.8 demonstrates this pattern.

In light of this, we define our artificial *treatment* group to be the stocks in the lowest initial retail share quintile. The initial retail shares are computed using the months between June 2014 and December 2018. Then, the *control* group is defined as the stocks in the highest initial retail share quintile. The period from January 2019 to December 2019 is defined as the *pre* period and the period from April 2020 to March 2021 is defined as the *post* period. The first quarter of 2020 is dropped.

The diff-in-diff regression specification is as follows:

$$\text{RG}_{it} = \gamma_1 \cdot \text{Treated}_{it} + \gamma_2 \cdot \text{Post}_{it} + \beta \cdot \text{Treated}_{it} \times \text{Post}_{it} + \Gamma' \mathbf{X}_{it} + \alpha_i + \gamma_t + \varepsilon_{it} \tag{2.29}$$

where Treated_{it} and Post_{it} are defined as discussed above. Controls $\mathbf{X}_{it}$ include size, book-to-market, Amihud measure, and foreign ownership. Stock fixed effects α_i and month fixed effects γ_t are included. In specifications with stock and month fixed effects, the treated and post dummies are absorbed.

	Q1 vs Q5		Q1/Q2 vs Q4/Q5	
	(1)	(2)	(3)	(4)
Treated	-7.42^{***}		-4.11^{***}	
	(0.83)		(0.49)	
Post	1.87^{***}		2.42^{***}	
	(0.59)		(0.40)	
Treated × Post	2.52^{***}	3.79^{*}	1.79^{***}	2.41
	(0.78)	(2.18)	(0.55)	(1.56)
Log Market Cap	0.89^{***}	13.78^{***}	0.69^{***}	12.66^{***}
	(0.22)	(1.72)	(0.15)	(1.20)
Book-to-Market	-1.31^{***}	0.05	-1.81^{***}	-0.37
	(0.22)	(0.87)	(0.16)	(0.97)
Amihud Measure	0.03	0.21	-0.04	0.18
	(0.10)	(0.33)	(0.06)	(0.15)
Foreign Holdings	-0.10^{***}	-0.35^{***}	-0.11^{***}	-0.30^{***}
	(0.02)	(0.12)	(0.01)	(0.08)
Stock FE	–	✓	–	✓
Quarter FE	–	✓	–	✓
Observations	7,858	7,858	16,648	16,648
Within R^2	0.02	0.05	0.02	0.04

Table 2.8: Return Gap vs Retail Volume Share (Diff-in-Diff). This table reports the regression results using the specification in equation (2.29). In columns (2) and (4), standard errors are two-way clustered at the stock and month level. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

Results are shown in Table 2.8. Columns (3) and (4) are results when we repeat this exercise with two quintiles for the treatment group and two quintiles for the control group. The coefficient estimates of interest are β's which are displayed in the third row. We see that the estimates are

positive and significant in columns (1) and (3). Also, the estimates for γ_1 and γ_2 displayed in first and second rows are also consistent with our priors. As the treatment groups consist of stocks with lower initial retail volume share, we expect the estimates for γ_1 to be negative. Also, because retail trading increased in general following the outbreak of the pandemic, we expect the estimates for γ_2 to be positive.

We would interpret the results cautiously, however. While we believe that the pandemic-induced increase in retail activity is an exogenous shock, we admit that there are still issues with making a strong causal claim. First, we must believe that it is the *retail volume share* that drives the return gap, rather than any general measure of *retail activity*. This is because our strategy also relies on the fact that shares are mechanically capped at 1.

Another issue with this specification is the difficulty of interpretation. As opposed to finding an instrument for retail volume share, this artificial treatment-control setup deters us from drawing a quantitative interpretation at a per unit-basis. We can only comment that the treatment group experienced a larger average differential increase in return gap relative to the change in return gap for the control group.

2.5 Conclusion

The persistent and large gap between overnight and intraday returns in many stock markets has drawn the attention of many researchers. Studies have pointed to both attention-induced trading and extrapolative trading—behavioral patterns often related to retail investors—in the morning as possible drivers of this global phenomenon. Still, retail flow data for the US has been either partial (e.g., brokerage data) or imputed (e.g., Boehmer et al. 2021). Furthermore, studies on the overnight-intraday return decomposition often fall short of establishing causal relationships due to the lack of exogenous variation in the explanatory variables of interest.

Our setup resolves the data problem and also addresses the identification challenge. KRX provides exact investor type-level identifiers at high frequencies. This means that our retail flow data are accurate and exhaustive. In addition, we devise a way to exploit the retail trading frenzy during the pandemic era as an exogenous variation to retail volume shares. We find that retail volume shares are positively related to the overnight-intraday return gap at the monthly frequency. Along the way, we also develop a simple cross-sectional asset pricing model that incorporates the daily ebb and flow of noise trader prominence. We are able to derive empirical predictions that are not only consistent with the overnight-intraday return gap, but also with existing evidence provided

by recent studies such as Aboody et al. (2018) and Lou et al. (2019).

Appendix to Chapter 2

B1 Model Details

B1.1 Proof of Proposition 1

Proof. We start with solving the optimal portfolio choice in the morning for two types of investors. The corresponding first-order condition with respect to maximization problem in equation (2.9) is:

$$
\begin{aligned}
2\gamma \mathrm{Var}_t(\boldsymbol{P}_{t+1/2})\boldsymbol{\theta}_t^q &= \left(\boldsymbol{d}_H + \mathbf{E}_t(\boldsymbol{P}_{t+1/2}) - (1+r)^{1/2}\boldsymbol{P}_t\right) \\
2\gamma \mathrm{Var}_t(\boldsymbol{P}_{t+1/2})\boldsymbol{\theta}_t^n &= \left(\boldsymbol{d}_H + \mathbf{E}_t(\boldsymbol{P}_{t+1/2}) - (1+r)^{1/2}\boldsymbol{P}_t + \boldsymbol{\rho}_t\right).
\end{aligned}
$$

Once we impose the market clearing condition $\mu_H\boldsymbol{\theta}_t^n + (1-\mu_H)\boldsymbol{\theta}_t^q = \boldsymbol{Q}$, we have:

$$
(1+r)^{1/2}\boldsymbol{P}_t = \boldsymbol{d}_H + \mathbf{E}_t(\boldsymbol{P}_{t+1/2}) - 2\gamma \mathrm{Var}_t(\boldsymbol{P}_{t+1/2})\boldsymbol{Q} + \mu_H\boldsymbol{\rho}_t. \tag{2.30}
$$

Similarly, for the afternoon trading session $(t+1/2, t+1)$, we can derive the following:

$$
(1+r)^{1/2}\boldsymbol{P}_{t+1/2} = \boldsymbol{d}_L + \mathbf{E}_{t+1/2}(\boldsymbol{P}_{t+1}) - 2\gamma \mathrm{Var}_{t+1/2}(\boldsymbol{P}_{t+1/2})\boldsymbol{Q} + \mu_L\boldsymbol{\rho}_{t+1/2}. \tag{2.31}
$$

Next, we derive the afternoon prices $P_t, P_{t+1}, \ldots$ in a stationary equilibrium. We adopt the assumption that conditional variance of price is constant as in De Long et al. (1990).

$$
\begin{aligned}
(1+r)\boldsymbol{P}_t = \boldsymbol{m}_H + \mathbf{E}_t(P_{t+1}) - 2\gamma\Sigma_H\boldsymbol{Q} + (1+r)^{1/2}\mu_H\boldsymbol{b}(z_t - z^*) + m_{H,\varepsilon}\varepsilon_t \qquad (2.32) \\
\text{where} \quad \boldsymbol{m}_H = (1+r)^{1/2}\left(\boldsymbol{d}_H + \mu_H\boldsymbol{b}z^*\right) + \boldsymbol{d}_L + \mu_L\boldsymbol{b}z^* \\
\Sigma_H = (1+r)^{1/2}\mathrm{Var}_t(\boldsymbol{P}_{t+1/2}) + \mathrm{Var}_{t+1/2}(\boldsymbol{P}_{t+1}) \\
m_{H,\varepsilon} = (1+r)^{1/2}\mu_H + \mu_L\phi^{1/2}
\end{aligned}
$$

Recursive substitution leads to:

$$\boldsymbol{P}_t = \frac{\boldsymbol{m}_H}{r} - \frac{2\gamma}{r}\Sigma_H \boldsymbol{Q} + \frac{\mu_H \boldsymbol{b}(z_t - z^*)}{(1+r)^{1/2}} + \frac{m_{H,\varepsilon}}{1+r-\phi}\varepsilon_t. \tag{2.33}$$

Similarly, the morning prices can be written as:

$$\begin{aligned} \boldsymbol{P}_{t+1/2} = \frac{\boldsymbol{m}_L}{r} - \frac{2\gamma}{r}\Sigma_L \boldsymbol{Q} + \frac{\mu_L \boldsymbol{b}(z_t - z^*)}{(1+r)^{1/2}} + \frac{m_{L,\varepsilon}}{1+r-\phi}\varepsilon_t \qquad (2.34) \\ \text{where} \quad \boldsymbol{m}_L = (1+r)^{1/2}\left(\boldsymbol{d}_L + \mu_L \boldsymbol{b} z^*\right) + \boldsymbol{d}_H + \mu_H \boldsymbol{b} z^* \\ \Sigma_L = (1+r)^{1/2}\mathrm{Var}_{t+1/2}(\boldsymbol{P}_{t+1}) + \mathrm{Var}_t(\boldsymbol{P}_{t+1/2}) \\ m_{L,\varepsilon} = (1+r)^{1/2}\mu_L + \mu_H \phi^{1/2}. \end{aligned}$$

Let $\Sigma_\eta = \mathrm{diag}(\sigma^2_{\eta,1}, ..., \sigma^2_{\eta,N})$, then we can see that:

$$\begin{aligned} \mathrm{Var}_t(\boldsymbol{P}_{t+1/2}) &= \frac{\mu_L^2}{1+r}\boldsymbol{b}\boldsymbol{b}'\sigma_z^2 + \left(\frac{m_{L,\varepsilon}}{1+r-\phi}\right)^2 \Sigma_\eta \\ \mathrm{Var}_{t+1/2}(\boldsymbol{P}_{t+1}) &= \frac{\mu_H^2}{1+r}\boldsymbol{b}\boldsymbol{b}'\sigma_z^2 + \left(\frac{m_{H,\varepsilon}}{1+r-\phi}\right)^2 \Sigma_\eta. \end{aligned}$$

We can substitute for Σ_H, Σ_L for explicit expressions.

□

B1.2 Proof of Proposition 2

Proof. Using the equilibrium prices derived in Proposition 1, we can write the return in morning:

$$\begin{aligned} &\boldsymbol{d}_H + \mathbf{E}_t(\boldsymbol{P}_{t+1/2}) - (1+r)^{1/2}\boldsymbol{P}_t \\ &\quad = \boldsymbol{d}_H + \frac{1}{r}\left(\boldsymbol{m}_L - (1+r)^{1/2}\boldsymbol{m}_H\right) \\ &\qquad - \frac{2\gamma}{r}(\Sigma_L - (1+r)^{1/2}\Sigma_H)\boldsymbol{Q} - \mu_H \boldsymbol{b}(z_t - z^*) + \frac{m_{L,\varepsilon}\phi^{1/2} - (1+r)^{1/2}m_{H,\varepsilon}}{1+r-\phi}\boldsymbol{\varepsilon}_t. \end{aligned}$$

Some useful relations are presented here for reference:

$$\begin{aligned} \boldsymbol{m}_L - (1+r)^{1/2}\boldsymbol{m}_H &= -r(\boldsymbol{d}_H + \mu_H \boldsymbol{b} z^*) \\ \Sigma_L - (1+r)^{1/2}\Sigma_H &= -r\mathrm{Var}_t(\boldsymbol{P}_{t+1/2}) \\ m_{L,\varepsilon}\phi^{1/2} - (1+r)^{1/2}m_{H,\varepsilon} &= -\mu_H(1+r-\phi) \end{aligned}$$

Rewriting the expressions in terms of the vector of morning returns $\boldsymbol{R}_t^H$, we get:

$$\boldsymbol{R}_t^H = 2\gamma \text{Var}_t(\boldsymbol{P}_{t+1/2})\boldsymbol{Q} - \mu_H \boldsymbol{\rho}_t.$$

Afternoon returns can be derived analogously.

□

Chapter 3

Unpacking the Demand for Sustainable Equity Investing*

3.1 Introduction

Sustainable investing seeks to direct capital towards companies with positive environmental and social impact and away from those with negative impact, thereby affecting their cost of capital. Interest in sustainable investing has grown significantly in the last decade with ESG (Environmental, Social, and Governance) assets predicted to reach over $53 trillion by 2025 (Bloomberg, 2021). This trend has been particularly pronounced in equity markets, resulting in a plethora of academic research attempting to understand equity investors' demand for sustainability (Coqueret, 2021).

While existing research shows that the demand for sustainable assets in financial markets has strengthened enough in aggregate to affect equity valuations (van der Beck, 2021; Pástor et al., 2022), this trend alone masks the underlying heterogeneity in demand for sustainability across investors. Often-cited threats to sustainable investing include concerns about price-elastic investors picking up "brown" stocks divested by sustainable investors and political backlash against ESG such as Florida's banning of ESG considerations from state pension investments (Bloomberg, 2022).[1]

* This work was previously circulated under the title "Measuring Institutional Pressure for Greenness: A Demand System Approach." For comments and discussions, we thank Moritz Lenel, Ernest Liu, Ralph Koijen, Thomas Kroen, Lars Hansen, Stefan Nagel, Christian Opp, Jonathan Payne, Lubos Pastor, Roberto Rigobon, Zacharias Sautner, Simon Schmickler, David Schoenherr, Xuan Tian, Wei Xiong, Motohiro Yogo, and participants at the 2020 MFR Young Scholars Workshop, 2020 CAFM, Chicago Booth Finance Brownbag, Chicago Econ Dynamics Working Group, and Princeton Finance Student Workshop. This research was funded in part by the John and Serena Liew Fellowship Fund at the Fama-Miller Center for Research in Finance, University of Chicago Booth School of Business.

[1] Popular media outlets often report hedge funds and private equity firms purchasing stakes in polluting firms that are being divested by climate-conscious institutions (Fletcher and Brower, 2021; Gilbert, 2021). The usual story is that institutional investors who face fewer constraints or less pressure from their clients happily scoop up divested shares, thereby

Not every investor shares the enthusiasm for sustainability.

Taking this heterogeneity into account is of first-order importance for the outstanding questions related to sustainable investing. Can investors meaningfully affect firms' cost of capital even if the hedge fund sector collectively invests in divested brown firms? If states or countries ban ESG considerations, how will the cross-section of stock valuations be affected? Do firms improve environmental performance in response to price pressure from sustainable shareholders?[2] Answers to these questions cannot be deduced from valuation patterns alone as they depend on the joint distribution of price elasticities, strength of demand for sustainability, and the specific dimensions of sustainability that investors care about.

In this paper, we investigate the heterogeneity in investor demand for sustainability and study its implications for firm decisions and asset prices. Not only do we estimate each investor's demand separately, but we also incorporate various dimensions of sustainability that a particular investor may care about. As a result, we go beyond a simple dichotomy of "green" vs. "brown" investors and provide a more nuanced classification of who green investors are. We then apply an equilibrium asset pricing framework to derive model-based empirical quantities and to consider counterfactual scenarios that allow us to answer policy-relevant questions.[3]

Our findings reveal both the merits and limitations of sustainable equity investing. On one hand, we show that the demand for firm sustainability has been rising for both active and passive investors and is not driven by potential correlation of sustainability with other firm characteristics. In addition, we find that investors not only demand firms with high third-party sustainability ratings but also those with low emissions intensity. On the other hand, they do not demand firms that innovate in sustainable technologies, which are the firms that could benefit the most from cost of capital reductions. In addition, investor pressure generated by these demand patterns seem to translate into only limited improvements in firm sustainability.

As the first step in understanding the investor demand for sustainability, we start by measuring three key dimensions of firm-level sustainability: emissions intensity, environment score, and green patents. To measure emissions intensity, we gather data on Scope 1 greenhouse gas emissions from

attenuating the effectiveness of sustainable investment mandates.

[2] In our framework, the first question will be the concern that price-elastic investors absorbing the excess supply (demand) of brown (green) assets. A perfectly price-elastic market will completely "undo" the valuation effect of sustainable investing, which also implies zero impact on the cost of capital of green and brown firms. The second question will be a counterfactual scenario in which we "turn off" the demand coefficient on sustainability for a subset of investors.

[3] Many existing equilibrium asset-pricing models on sustainable investing are based on one or two sources of investor heterogeneity. Pástor et al. (2021), Pedersen et al. (2021), and Goldstein et al. (2021) all present models that feature investors with different preferences or beliefs about ESG stocks. Zerbib (2022) offers a model that has investors with both different ESG preferences and investment universes. Our paper builds on these works by analyzing multiple sources of investor heterogeneity.

S&P Trucost, which are emissions that come directly from sources controlled or owned by the firm. Next, we create environment scores using data from the MSCI ESG ratings database. Since these scores are correlated with greenhouse gas emissions among firms, we extract the component of the score that is unrelated to emissions. Finally, we use data from PatentsView to create company-level production of green patents. By adopting this comprehensive approach, we obviate the need to assume a specific perspective on how investors perceive sustainability, thus improving upon existing studies that typically examine only one dimension.

Using the firm sustainability measures, we present initial evidence that the three aspects of sustainability are valued differently by investors. Our cross-sectional valuation regressions from 2013 to 2021 reveal that third-party environment scores have been consistently valued by investors over this period, and emissions intensity has been negatively valued only after 2018. In particular, during the post-2018 period, a one standard deviation lower emission intensity is associated with 6.45% lower market-to-book ratio, and one standard deviation higher environmental score is associated with 11.5% higher market-to-book ratio. On the other hand, the production of green patents is not valued by investors. These results suggest that different aspects of sustainability may be valued heterogeneously in the equity market.

To gain a deeper understanding of this heterogeneity, we construct an asset demand system that incorporates our sustainability measures into investor demand curves, along with traditional stock characteristics that are known to influence investor demand. We model and estimate investor demand using the approach in Koijen and Yogo (2019), which provides a tractable model of investor demand that allows for rich heterogeneity. We also illustrate that sustainability can enter investor demand in two ways: when sustainability is informative about expected returns and when the investor faces a minimum sustainability constraint.

We document substantial heterogeneity in demand for sustainability across investors. On average, investors have a positive preference for higher environment score and emissions intensity but not for green patents. Comparing across investors, we find that active investors—those who deviate more from the market benchmark weights—exhibit not only higher price elasticities but also stronger demand for sustainability. This finding suggests that active investors are not counteracting sustainable investing as commonly assumed, but rather playing a crucial role along with passive investors in shaping the valuation patterns in equity markets.

The trend of increasing overall demand for low-emissions firms can come from two sources: (i) a within-investor preference shift towards low-emission stocks, or (ii) a shift of AUM away from "brown" investors who prefer high-emission stocks, towards "green" investors who prefer low-

emission stocks. By examining trends based on estimated coefficients, we provide evidence that this trend is mainly driven by within-investors shifts in preferences rather than shifts in capital across investors. We also confirm this evidence through counterfactuals based on the estimated demand system.

Having documented the heterogeneity across investors, we then examine the implications of these patterns in investor demand. We first consider implications for firm decisions. We use the estimated demand curves to quantify investor pressure for sustainability, which captures the price pressure a firm receives, through investor demand, to become more sustainable. Our closed-form expression shows that the pressure is determined by the average demand for sustainability of the firm's investors, adjusted for their collective price elasticity. Our estimated demand parameters also suggest that, on average, firms have faced increased pressure to improve sustainability. However, we find that higher investor pressure today only weakly predicts future improvements in firm sustainability.

We next consider the asset pricing implications of two counterfactual scenarios that reflect important developments in sustainable investing. In the first, we study the impact of introducing ESG-agnostic mandates by analyzing the valuation patterns when certain investors "shut off" their demand for sustainability. This exercise shows that active and passive investors contribute roughly equally to the valuation patterns, challenging the often held belief that active investors may "undo" sustainable investing.[4]

In the second scenario, we examine the effect of introducing hypothetical ESG ratings. Specifically, we consider how leading sustainability score providers could adjust their methodology to incorporate firm-level green productivity. We find that such a change leads to a meaningful increase in the valuation of top green patent producers while maintaining the valuation gap between low- and high-emission stocks. These results suggest that the proposed change in rating methodology can be implemented without compromising the objective of sustainable investing.

Related Literature and Contribution

Our paper contributes to four main strands of literature in sustainable equity investing and asset pricing. First, our paper contributes to the literature on the asset-pricing implications of sustainable investing, to which Giglio et al. (2021) and Coqueret (2021) provide a comprehensive review. Our focus is on the equity market, and existing papers in this literature study the return gap between

[4] In the language of demand system asset pricing, this would be the concern that price-elastic investors absorbing the excess supply (demand) of brown (green) assets. A perfectly price-elastic market will completely "undo" the valuation effect of sustainable investing, which also implies zero impact on the cost of capital of green and brown firms.

green and brown stocks both through the lens of a theoretical framework (Heinkel et al., 2001; Pástor et al., 2021; Pedersen et al., 2021 Zerbib, 2022) and through empirical analyses based on realized returns (Görgen et al., 2020; Bolton and Kacperczyk, 2021; Derrien et al., 2021; ; Glossner, 2021; Hsu et al., 2022; Pástor et al., 2022).[5] To this literature, we provide two contributions. First, we provide new evidence on how three different measures of environmental performance are priced in the cross-section of stock valuations and how these pricing relationships have evolved over time. Our findings complement Choi et al. (2022) who find that carbon emissions intensity is negatively correlated with stock valuation in 26 countries as well as Cohen et al. (2020) who show that ESG investors lack incentive to invest in firms with high green innovation capacity. The coverage of our results on all three green characteristics also adds to the discussion regarding "ESG confusion" in Berg et al. (2022) that highlights the low correlation between different sustainable metrics. Our second contribution is to illustrate how different changes in sustainable demand could affect these valuation relationships via counterfactuals, which add to the analysis of green-brown expected returns in Berk and van Binsbergen (2021) and the analysis of realized returns in van der Beck (2021).

Second, our paper contributes primarily to the growing literature that directly studies investor demand for sustainable assets. While some papers use the survey instruments (Krueger et al., 2020; Gormsen et al., 2023), most papers analyze their portfolio choice decisions directly. For institutional investors, Gibson et al. (2020) computes a portfolio-level sustainability measure for all 13F investors and shows that institutions with high portfolio sustainability earned higher returns after 2010. Others have focused on subsets of institutional investors by studying inflows into sustainable mutual funds (Hartzmark and Sussman, 2019; van der Beck, 2021; Baker et al., 2022) or the greenwashing behaviors of hedge funds and active mutual funds (Liang et al., 2021; Kim and Yoon, 2023). We contribute to this literature by providing a comprehensive estimate of sustainable demand for institutional investors in the U.S. stock market. The estimates of investor demand thus shed new light on both the cross-sectional differences in sustainable demand across investors as well as the time-series evolution of sustainable demand for each characteristic. While our paper shares the same focus on asset demand as in Koijen et al. (2022), we provide more richness by analyzing a comprehensive set of sustainability metrics, studying more counterfactual scenarios, and providing real impact analysis through our model-based investor pressure measure.

Third, our paper contributes to the literature on the real impact of sustainable equity investing.

[5] There also exists work with respect to demand for other types of assets such as bonds (Flammer, 2021) and options (Ilhan et al., 2021).

Theoretically, Broccardo et al. (2022) shows in a model of firm incentives that divestment ("exit") tends to be less effective than engagement ("voice"), and Edmans et al. (2022) highlights the limits of blank exclusion of such full divestment strategies. Berk and van Binsbergen (2021) also argues based on a CAPM calibration that even a large substitution from brown to green stocks would only marginally increase the cost of capital for brown firms. The empirical evidence is also generally mixed regarding the impact of sustainable investing on real firm decisions (Heath et al., 2021, Gantchev et al., 2022; Hartzmark and Shue, 2023). We contribute to this literature by deriving a new measure of a firm's incentive to improve its environmental performance and showing that while firms on average have positive incentive from investor pressure to reduce carbon emission, the correlation between a firm's investor pressure and future environmental performance is only weakly positive.

Finally, our paper contributes to the burgeoning literature that studies questions in asset pricing based on estimation of asset demand in markets ranging from equity, corporate bonds, and country-level assets (Koijen and Yogo, 2019, 2020; Koijen et al., 2022; Bretscher et al., 2020; Jiang et al., 2022). In particular, our paper relates to studies that apply demand estimation to specific asset-pricing questions, including Gabaix et al. (2022) on the asset demand of U.S. households, Huebner (2022) on the source of equity momentum, Jansen (2021) on long term bond demand, and van der Beck and Jaunin (2021) on retail investor demand.[6] Our paper contributes to the latter part of this literature by providing a structural analysis of sustainable equity investing through our emphasis on asset demand of individual investors.

Roadmap

We first describe the data as well as stylized facts from valuation regressions (Section 3.2). We then set up the asset demand system that includes sustainability characteristics in investor demand curves (Section 3.3), which we estimate and highlight key patterns (Section 3.4). With the estimated demand system, we explore implications for both firm decisions and asset prices. In Section 3.5, we quantify investor pressure for sustainability and examine whether higher pressure today leads to greater future sustainability. In Section 3.6, we consider two counterfactual scenarios related to investor preferences and measures of sustainability to examine how valuation patterns changes.

[6] Another strand of the literature focuses on asset demand elasticity: Gabaix and Koijen (2021) estimates low "macro elasticity" of equity demand and proposes the inelastic market hypothesis, Davis et al. (2022) proposes an explanation of inelastic demand based on a model of information acquisition, and Haddad et al. (2021) estimates moderate strategic substitution in the price elasticity of investors.

3.2 Data and Stylized Facts

We measure firm sustainability in three ways, an approach that improves upon the previous literature that usually focuses on one of the three dimensions (Section 3.2.1). We then show that each sustainability characteristic is priced differently in the cross-section of stocks, indicating that these characteristics may enter investor demand curves in different ways (Section 3.2.2).

3.2.1 Data

We construct three measures that capture different dimensions of firm-level sustainability: (i) emissions intensity using data from S&P Trucost, (ii) environment score using data from MSCI, and (iii) green patents using data from PatentsView. We later combine this data with detailed holdings of institutional investors from FactSet as well as stock characteristics and firm-level variables from CRSP and Compustat.

Measuring Firm Sustainability

Emissions Intensity. We use firm-level Scope 1 greenhouse gas emissions from S&P Trucost. We choose carbon emission as our primary measure because it is one of the most important objectives for sustainable investing and also the most quantifiable. We also focus on Scope 1 emissions, which are the emissions that emanate directly from sources controlled or owned by the firm. For our measure, we use the logarithm of Scope 1 emissions intensity, which is defined as a company's annual Scope 1 emissions divided by annual revenue.[7] We henceforth refer to this measure as *emissions intensity*.

Environment Score. We obtain firm-specific measures of environmental performance from MSCI ESG Ratings database, which succeeds the MSCI KLD database used in previous studies related to ESG investing. We choose MSCI ESG ratings over other ESG rating datasets with a similar motivation as in Pástor et al. (2022): MSCI covers more firms than other raters, exhibits the least noise (Berg et al., 2019), and is based on a comprehensive set of corporate documents, government data, and news media.

Following Pástor et al. (2022), we use a combination of the Environment Pillar Score and Environment Pillar Weight from MSCI to measure sustainability for asset n at time t. Specifically, let $E_t(n)$ be the asset n's most recent Environment Pillar Score before month t, looking back no

[7] Fewer than 0.1% of all firm-year observations in the Trucost data have zero Scope 1 emission. We add the minimum positive level of Scope 1 emissions intensity to the actual values before taking the logarithm.

more than 12 months, and let $w_t^E(n)$ be the most recent Environment Pillar Weight before month t, looking back no more than 12 months.

We define $g_t(n)$ as the following, which we call the *raw environment score*:

$$g_t(n) = \frac{-\left(10 - E_t(n)\right) w_t^E(n)}{100}$$

The product is a combination of how far the asset's rating is from a perfect score $\left(10 - E_t(n)\right)$ and the relative importance of environmental issues for the firm $\left(w_t^E(n)\right)$. Due to the minus sign, $g_t(n)$ is always negative and a value close to zero implies higher level of sustainability. Appendix C1.1 illustrates the importance of adjusting MSCI Environmental Pillar Scores by the Environmental Pillar Weights by comparing the scores between oil & gas and banking stocks. Appendix C2.2 shows that our main results are robust to two alternative definitions of environmental score, based on different transformations of environmental pillar scores and weights.

The quarterly average cross-sectional correlation between $g_t(n)$ and log Scope 1 emissions intensity is -0.58, which suggest that emission is an important part in MSCI's environmental ratings. In order to improve precision and interpretability of our analysis, we regress environmental score on log emissions intensity in each quarter and use this residual in all cross-sectional analyses. This residual thus captures the component of environmental score orthogonal to current level of carbon emission, and we henceforth refer to this residual as the *environment score.*

Green Patents. We construct firm-specific measures of green technology innovation based on data of granted U.S. patents from PatentsView. We follow the method developed by Haščič and Migotto (2015) and used in Cohen et al. (2020) to identify "green patents" related to technologies for reducing emission, mitigating pollution, or improving environmental performance in general. We use the firm identifier links provided by Autor et al. (2020) and the WRDS Patents database to merge PatentsView data with our stock universe.

For cross-sectional analyses, we take the number of green patents that each firm develops in the past five years, and scale the number of patents by the firm's total asset. We henceforth refer to this measure as *green patents* and define *non-green patents* analogously using the number of non-green patents[8]. Appendix C1.2 provides further details on processing the PatentsView data.

[8] The correlation between the two measures is 0.54, which highlights the importance of including non-green patents as a control variable that captures a firm's overall innovation capacity. The correlations between all other stock characteristics are below 0.1.

Portfolio Holdings, Asset Prices, and Characteristics

We use quarterly institutional portfolio holdings of U.S. stocks from the FactSet database, which sources its data primarily from Securities and Exchange Commission Form 13-F filings. All institutional investment managers with more than $100 million of asset under management must file the form every quarter. The data comes at the investment manager level rather than individual fund level (e.g. Vanguard files its aggregated holdings as one institution) and reports only long positions. We use the data from 2013Q1 to 2021Q3 and merge the data with our quarterly stock universe from CRSP and Compustat via CUSIP. Appendix C1.3 provides additional details.

Summary Statistics

Table 3.1 provides summary statistics of investor and stock characteristics. Panel (a) first provides the distribution of stock characteristics. In an average quarter, 10.1% of stocks by market cap does not have environmental score data, and 3.9% of stocks by market cap does not have emissions data. For these observations with missing data, we impute the environmental scores or emission intensities based on industry-quarter averages using Fama-French 12-industry classifications. Among observations with non-missing annual emissions intensity, the median and mean are 13.6 and 189.7 tons per million dollars of revenue. In an average quarter, 64.8% of stocks (constituting 67.3% of total market cap) have at least one non-green patent granted in the previous 5 years, and 33.0% of stocks (constituting 29.6% of total market cap) have at least one green patent granted in the previous 5 years. In panel (b), we summarize the investors by type. Institutional investors constitute 69.3% of total AUM in an average quarter and the top 25 largest investment advisors constitute 30.7% of the total AUM. Hedge funds have the highest average portfolio active share at 0.711. In Appendix C2.1, we also summarize the relationship between the three measures of firm sustainability and show that the environment score has real information content with respect to future emissions and green innovation.

3.2.2 Stylized Facts: Evidence from Valuation Regressions

Next we provide results from cross-sectional valuation regressions during the period from 2013 to 2021. We find that environment score is consistently positively priced; emissions intensity is negatively priced only after 2018; and green patent productivity is not priced. These results imply that different aspects of sustainability may be demanded differently by investors, thereby motivating a more structural analysis via the asset demand system in subsequent sections.

Specifically, we estimate the following valuation regressions:

$$mb_t(n) = a_t + \boldsymbol{\lambda}' \mathbf{x}_t(n) + \epsilon_t(n) \tag{3.1}$$

where the dependent variable $mb_t(n)$ is the market-to-book ratio of firm n at time t and $\mathbf{x}_t(n)$ is a vector of time-varying firm characteristics that includes both sustainability and non-sustainability characteristics that are cross-sectionally standardized in each quarter.

Table 3.2 summarizes the coefficients from the regressions. Column (1) shows the results based on the entire sample. First, among the three green characteristics, only the environmental score is significantly reflected in the cross-section of valuations: a one standard deviation higher environmental score is associated with 10.7% higher market-to-book ratio, with t-statistic of 8.2 based on clustered standard errors. Second, emissions intensity is negatively related to valuation, but the coefficient is not statistically significant over the entire sample. Finally, although a firm's overall innovation capacity is positively reflected in valuations (strongly positive coefficient for the non-green patents), green patents are not reflected in the cross-sectional of valuations.

Columns (2) - (3) of Table 3.2 provides subsample analyses based on data from 2013-2017 and 2018-2021 separately. First, the results for environmental score are similar over the two sample periods, i.e., the environmental score is consistently valued in the cross-section. Second, the coefficient for emissions intensity is not significant for the 2013-2017 subsample but is strongly negative and significant for the 2018-2021 subsample. In the 2018-2021 sample, one standard deviation higher emissions intensity is associated with 6.45% lower market-to-book ratio in the cross-section. This difference across two sample periods suggests a significant shift in investor demand that prefers low-emission over high-emission stocks in the recent years.

Figure 3.1 plots the coefficients from estimating Equation (3.1) cross-sectionally for each quarter. The time series of coefficients confirm both a consistent valuation gap for environmental score and a strengthening valuation gap against emissions intensity. In fact, the valuation regression coefficient for emissions intensity was positive and statistically significant from 2013 to 2014 and then turned negative and statistically significant since 2018.

Overall, we find that the three sustainability characteristics are priced differently in the cross-section of stocks and also for different time periods. These results suggest that characteristics may enter investor demand curves in different ways, which we explore next.

3.3 The Asset Demand System with Sustainability

We set up the asset demand system that includes sustainability characteristics in investor demand curves. We allow heterogeneity across both investors and time, which allows us to examine how each of the sustainability measures are demanded differently across investors.

3.3.1 Setup and Notation

We adapt the setting and notation used in Koijen and Yogo (2019), which we partly introduce here while omitting some details to avoid repeating the entire setup. A key addition is the introduction of sustainability characteristics. Investors may care about sustainability either for pecuniary or non-pecuniary reasons, and evidence can be found for both (e.g. Barber et al., 2021 and Bansal et al., 2018). While we remain agnostic on what the more prominent motivation is, we show that sustainability should enter the characteristics-based demand in at least two cases: sustainability is informative about expected returns or investors are constrained to hold a sustainable portfolio (e.g. due to investment mandates or pressure from clients).

Consider an economy with N assets indexed by $n = 1, \ldots, N$ and I investors indexed by $i = 1, \ldots, I$. We denote the outside asset as the 0th asset. Furthermore, let $P_t(n)$ and $S_t(n)$ denote the price and shares outstanding of asset n at time t respectively. We denote the logarithms of these variables in lowercase letters and the N-dimensional vectors in boldface. Suppose each asset has K characteristics indexed by $k = 1, \ldots, K$ so that the kth characteristics of asset n at time t is denoted $x_{kt}(n)$ and the vector of characteristics is denoted $\mathbf{x}_t(n)$.

Investor Decisions. Investor i optimally chooses at each time t her weights on these assets $\mathbf{w}_{it}$. Denoting the asset under management of investor i at time t by A_{it}, investor i maximizes expected terminal wealth $\mathbb{E}_{it}[\log(A_{iT})]$ under the intertemporal budget constraint.[9] Investors face short-sale constraints, $\mathbf{w}_{it} \geq \mathbf{0}$ and $\mathbf{1}'\mathbf{w}_{it} < 1$. Investors have heterogeneous beliefs about expected returns of assets, which they form by considering the observed characteristics. Investor i's unobserved latent demand for asset n is denoted $\log(\epsilon_{it}(n))$. Then investor i's information set for asset n can be written as

$$\hat{\mathbf{x}}_{it}(n) = \begin{bmatrix} me_t(n) \\ x_t(n) \\ \log(\epsilon_{it}(n)) \end{bmatrix} \tag{3.2}$$

[9] As in Pástor et al. (2021), we can have sustainability enter the utility directly, but we derive our results without doing so for now.

and an Mth-order polynomial of this vector can be written as

$$\mathbf{y}_{it}(n) = \begin{bmatrix} \hat{\mathbf{x}}_{it}(n) \\ vec(\hat{\mathbf{x}}_{it}(n)\hat{\mathbf{x}}_{it}(n)') \\ \vdots \end{bmatrix}, \tag{3.3}$$

which determines the investors' beliefs about expected returns.

Factor Structure. We maintain Assumption 1 of Koijen and Yogo (2019), so that the covariance of log excess returns, relative to the outside asset, is $\Sigma_{it} = \Gamma_{it}\Gamma_{it}' + \gamma_{it}\mathbf{I}$, where Γ_{it} is a vector of factor loadings and $\gamma_{it} > 0$ is idiosyncratic variance, and that expected excess returns and factor loadings are polynomial functions of characteristics:

$$\begin{aligned} \mu_{it}(n) &= \mathbf{y}_{it}(n)'\Phi_{it} + \phi_{it} \\ \Gamma_{it}(n) &= \mathbf{y}_{it}(n)'\Psi_{it} + \psi_{it} \end{aligned} \tag{3.4}$$

where Φ_{it} and Ψ_{it} are vectors and ϕ_{it} and ψ_{it} are scalars that are constant across assets. In other words, returns have a one-factor structure and an asset's own characteristics are sufficient for its factor loadings.

Sustainability as Characteristics. Importantly, we further assume that firm-level sustainability metrics are included in the vector of characteristics $\mathbf{x}_t(n)$. In Appendix C3.1, we use an example of one sustainability metric to show that sustainability can enter an investor's characteristic-based demand function if either it is either informative about the expected returns or the investor faces a "minimum sustainability constraint". Moreover, Appendix A of Koijen and Yogo (2019) shows that a particular coefficient restriction implies that the investors' optimal portfolio weights follow logit functions of prices, characteristics, and latent demand. In other words, optimal portfolio weight for stock n, for investor i, at a given period t satisfies:

$$\frac{w_{it}(n)}{w_{it}(0)} = \exp\left(b_{0,it} + \beta_{0,it} me_t(n) + \beta_{1,it}'\mathbf{x}_t(n)\right)\epsilon_{it}(n) \tag{3.5}$$

with sustainability entering as part of the characteristics $\mathbf{x}_t(n)$.

3.3.2 Implementation

We estimate the demand model for investor i for a given quarter t, which can be written as:

$$\forall i, \forall t : \frac{w_{it}(n)}{w_{it}(0)} = \exp\left(b_{0,it} + \beta_{0,it} mb_t(n) + \beta'_{1,it} \boldsymbol{s}_t(n) + \beta'_{2,it} \mathbf{x}^*_t(n)\right) \epsilon_{it}(n) \tag{3.6}$$

where $mb_t(n)$ is the log market-to-book ratio of asset n at time t. $\boldsymbol{s}_t(n)$ denotes the cross-sectionally standardized sustainability characteristics: emissions intensity, environmental score, and green patents. $\mathbf{x}^*_t(n)$ denotes other cross-sectionally standardized characteristics: log book equity, profitability, investment, dividend to book equity, market beta, and non-green patents. Note that we follow Koijen et al. (2022) to use log market-to-book ratio as the measure for price. The coefficients $\beta'_{1,it}$ measure investor i's demand for the three sustainability characteristics, after controlling for all other stock characteristics.

We use the same identification assumption as Koijen and Yogo (2019) for estimating Equation (3.6): we assume the latent demand $\epsilon_{it}(n)$ is exogenous to all stock characteristics except log market-to-book ratio, all investors' AUM A_{it}, and all investors' investment universes $\mathcal{N}_{it}$. Under these assumptions, $mb_t(n)$ is the only endogenous regressor in (3.6) as $mb_t(n)$ is correlated with latent demand $\epsilon_{it}(n)$ through market clearing.

To instrument for $mb_t(n)$ in the demand estimation for investor i, we construct counterfactual log market capitalization of stock n if all investors other than i or the household sector holds an equal-weighted portfolio of their investment universes:

$$\widetilde{me}_{i,t}(n) = \log\left(\sum_{j \neq i, HH} A_{jt} \frac{1\{n \in \mathcal{N}_{it}\}}{1 + |\mathcal{N}_{it}|}\right) \tag{3.7}$$

Based on the identification assumptions above, the instrument $\widetilde{me}_{i,t}(n)$ for investor i is exogenous to the investor's latent demand $\epsilon_{it}(n)$, and thus the instrument satisfies the exclusion restriction:

$$\mathbb{E}_t\left[\epsilon_{it}(n) \mid \mathbf{x}_t(n), \mathbf{s}_t(n)\right] = 1.$$

The instrument satisfies the relevance condition because all else equal, stocks held by more and larger investors tend to have higher market capitalization and thus higher market-to-book ratio. Koijen and Yogo (2019) documents that the instrument has high first-stage t-statistics that pass the Stock and Yogo (2005) test for weak instruments. We use non-linear GMM to estimate the demand equation (3.6) based on the instrument $\widetilde{me}_{i,t}(n)$ and all non-price characteristics.

3.4 Estimated Demand for Sustainability

In this section, we summarize the estimated demand curves and highlight three key patterns. First, we document strong heterogeneity in demand for sustainability across investors—in particular, more price-elastic investors exhibit higher demand for sustainability in terms of both environmental score and emission intensity. Second, we show that the increasingly negative demand for emissions intensity is primarily driven by within-investor demand shifts rather than across-investor shifts in AUM. Third, we find that higher active share, and lower portfolio turnover are associated with stronger sustainability demand, while indicators for value investors or signatories of the United Nations Principles for Responsible Investment (UNPRI) are not significantly correlated with sustainability demand.

3.4.1 Summary Statistics of Estimated Demand Coefficients

The estimated demand system reveals a strong heterogeneity in demand for sustainability across investors behind the valuation patterns. In particular, active investors are more price-elastic and exhibit higher demand for sustainability, which suggests that they are not "undoing" sustainable investing but rather playing a key role in driving and maintaining the valuation gap.

Table 3.3 provides summary statistics of our estimated demand coefficients. We compute the summary statistics across investors in every quarter, and then take an equal-weighted average across quarters. First, the demand for environmental score is positive on average, with an AUM-weighted average coefficient of 0.031. The demand for emissions intensity is negative on average, with an AUM-weighted average coefficient of -0.023. These coefficients mean that an average investor increases its demand by 3.1% per one standard deviation higher environmental score, and decreases its demand by 2.3% per one standard deviation higher emissions intensity. Therefore, investors have positive demand for sustainability on average, and they have positive preference for two orthogonal measures of environmental performance—emissions intensity and environmental score. Moreover, the demand coefficients for these two green characteristics have comparable magnitudes with coefficients for the five non-green characteristics.

The demand for green patents is near zero on average, with an AUM-weighted average coefficient of -0.003 for green patents. In comparison, the demand for a firm's overall innovation (measured by non-green patents) is positive, with an AUM-weighted average coefficient of 0.02. Therefore, the average investor in our sample does not have specific preference for green patents. In addition, we observe strong heterogeneity of demand across investors: the equal-weighted 10th/90th

percentile of demand coefficients are $-0.287/0.343$ for environmental score, and $-0.394/0.190$ for emissions intensity. In line with Koijen and Yogo (2019) and Koijen et al. (2022), this result highlights the importance of allowing cross-investor heterogeneity for understanding demand for sustainability.

Figure 3.2 summarizes the relationship between pairs of demand coefficients across investors through binscatter plots. Panel (a) shows that price-elastic investors have higher demand for sustainability on average for both environment score and emissions intensity. Because active investors tend to be more price-elastic,[10] these results provide the first evidence that active investors are not "undoing" sustainable demand by aggressively buying brown stocks. In contrast, the stronger sustainable demand for price-elastic investors suggests that they play an important role in creating the valuation gap between green and brown stocks, as the demand of price-elastic investors have higher relative impact on valuation (Koijen et al., 2022). Furthermore, panel (b) of shows that investors with higher demand for environment score also tend to have larger negative demand for emissions intensity. These results suggest that "green investors" consider both the actual emission of a firm and its third-party environmental rating when making sustainable investment decisions.

Finally, in Figure 3.3, we study the time trends in the AUM-weighted average coefficients for all investors as well as by broad investor types, which provides a validation of our exercise with respect to the valuation regression results. First, across all investor types, the average demand coefficient for environmental score is positive and stable over time, while the average demand coefficient for emissions intensity is negative and increasing in magnitude over time. These results are consistent with the time series of valuation regression coefficients in Figure 3.1, where the valuation gap between stocks with high and low environmental scores is positive and stable over time, and the valuation gap between high- and low-emission stocks is negative and opening up over time.

Second, we observe that the increasing demand for low-emission stocks is driven by both active and passive institutional investors. For active institution types, the average demand coefficient for emissions intensity decreased from -0.013 in 2013Q1 to as low as -0.105 in 2020Q2. For passive institution types, this coefficient decreased from -0.008 in 2013Q1 to -0.040 in 2021Q3. These results suggest that the shift in asset demand towards low-emission firms is a broad-based trend across different investor types.

[10] The average quarterly cross-sectional correlation between active share and price inelasticity is -0.30.

3.4.2 Within-Investor Demand or Across-Investor AUM Shift?

The trend of increasing overall demand for low-emission firms could come from two sources: (i) a within-investor preference shift towards low-emission stocks, or (ii) a shift of AUM away from "brown" investors who prefer high-emission stocks, towards "green" investors who prefer low-emission stocks. We next show that the increasingly negative demand for emissions intensity is primarily driven by within-investor demand shifts rather than across-investor shifts in AUM.

To quantify the relative importance of each in driving the shift in overall demand, Figure 3.4 plots the time series of average demand coefficient for emissions intensity against a counterfactual series where there is no within-investor preference shift. To construct this counterfactual, let $T_{0,i}$ be the first quarter when investor i appears in our sample, and let $\beta_{i,GHG,t}$ be investor i's demand coefficient for emission in quarter t. With these notations, $\beta_{i,GHG,T_{0,i}}$ is investor i's demand coefficient for emission in its earliest quarter in our sample. In each quarter, we compare the AUM-weighted average of demand coefficients $\beta_{i,GHG,t}$ against the counterfactual coefficients if there was no preference shift over time, $\beta_{i,GHG,T_{0,i}}$:

$$\bar{\beta}_{GHG,t} \coloneqq \frac{\sum_i A_{i,t}\beta_{i,GHG,t}}{\sum_i A_{i,t}} \text{ (Actual Data)}$$

$$\bar{\beta}^*_{GHG,t} \coloneqq \frac{\sum_i A_{i,t}\beta_{i,GHG,T_{0,i}}}{\sum_i A_{i,t}} \text{ (No Within-Investor Preference Shift)}$$

Figure 3.4 shows that the decrease of $\bar{\beta}_{GHG,t}$ over time—from -0.009 in 2013Q1 to -0.035 in 2021Q3—is almost entirely driven by preference shift within investor. If each investor's demand coefficient had stayed constant over time, the average demand coefficient for emissions intensity would only decrease from -0.009 in 2013Q1 to -0.014 in 2021Q3. Therefore, the overall demand shift towards low-emission stocks is mostly driven by portfolio rebalancing decisions of each investor, rather than a shift of AUM from "brown" to "green" investors.

To formally test the analysis above, we also regress the emission demand coefficient $\beta_{i,GHG,t}$ on a time trend with investor fixed effects. Within each quarter, the investors are weighted by their AUM. If the time trend of decreasing coefficients is purely driven by the shift of AUM across different investors, the time trend should not be significant after controlling for investor fixed effects. On the other hand, if the time trend is driven by within-investor change of demand, the time trend should remain significant after controlling for investor fixed effects.

Table 3.4 displays the regression results. Column (1) shows the time trend without controlling

for investor fixed effects, and column (2) show the time trend after controlling for investor fixed effects. We observe that adding investor fixed effects has little impact on the time trend coefficient: the time trend coefficients in columns (1) and (2) are not statistically significantly different from each other. Therefore, these results further bolster our finding that the time trend of increasing demand towards low-emission stocks is mainly driven by within-investor demand shifts, rather than a shift of AUM across investors. In Appendix C2.3, we further confirm this finding using a counterfactual analysis where we reverse the investor-changes in demand for sustainability as well as the changes in AUM.

We note that our results here are complementary to and independent from the finding of van der Beck (2021) that the outperformance of sustainable stocks from 2016 to 2021 can be entirely attributed to the flow-driven price pressure from end investors' portfolio reallocation from non-sustainable to sustainable mutual funds. The finding of van der Beck (2021) shows that the AUM shift between mutual funds is important, but in our institution-level data, shift of AUM between funds provided by a same institution will be reflected as a within-institution demand shift in our analysis. For example, if many retail investors shifted their investment from a Vanguard value-stock fund into a Vanguard sustainable fund, the total AUM of Vanguard would not change, but Vanguard's green demand coefficients will increase. Therefore, as long as most flows from non-sustainable to sustainable mutual funds occur among funds provided by a same institution, our results are consistent with van der Beck (2021). Furthermore, our sample of all 13F institutional investors is much larger than the equity mutual funds and thus our findings provide new insights on the recent growth of sustainable investing.

3.4.3 Cross-sectional Patterns

In this section, we examine the relationship between green demand and investor characteristics via cross-sectional regressions. In the cross-section of investors, we find that higher price elasticity, higher active share, and lower portfolio turnover are all associated with stronger sustainability demand. On the other hand, value investors and signatories of the UNPRI do not have significantly different demand for sustainability.

Investor Characteristics

For each of our three sustainability measures, we regress the investor-quarter level demand coefficients on four continuous (price inelasticity, log AUM, active share, and quarterly portfolio

turnover) and seven investor style or type indicators. For the indicators, we include one dummy variable for non-US investors, two investment style indicators for value and growth, and four investor type indicators for hedge funds, private banking, long-term, and broker-dealer (the left-out type is investment advisors). We control for time fixed effects to make the comparison across investors within each quarter, and we exclude the household sector from this analysis.

Table 3.5 reports the cross-sectional regression results. First, columns (1) and (2) of the table show that higher price elasticity, higher active share, and lower portfolio turnover are all associated with stronger sustainability demand, for both environmental score and emissions intensity. Also, growth investors have stronger sustainability demand than "generalist" investors who do not classify as either value or growth. We also find that foreign (non-U.S.) investors have stronger sustainability demand than U.S. investors, which complements the findings in Dyck et al. (2019) and may reflect Krueger et al. (2020) which highlights the importance of geographical origins for investors' sustainability preferences.

Second, compared with investment advisors, hedge funds have higher demand coefficients for both environmental score and emissions intensity—i.e., hedge funds have stronger preference for stocks with higher environmental score and *higher* emission. On one hand, to the extent that emission is a cleaner measure of firms' current environmental performance than environmental scores, these result could be interpreted as evidence of "green window dressing" for hedge funds. Hedge funds might buy more high-environmental score stocks to boost their portfolio score, but at the same time buy more high-emission stocks[11]. On the other hand, because higher environmental score predicts lower future emission in the cross-section of stocks (Table 3.C5.1), we can also interpret the results as evidence that hedge funds prefer stocks that have high current emission but better emission reduction potential in the future.

Finally, investors classified as "value" by FactSet do not have significantly different sustainability demand compared with "generalist" investors. Together with our finding that price-elastic investors tend to have stronger green demand, this result alleviates the popular concern of value investors "undoing" sustainable investing by buying up brown stocks.

[11] We only propose "green window dressing" as one potential explanation of the hedge fund coefficients in Table 3.5. While we do not seek to prove the existence of "green window dressing", we note that this proposition is in line with the finding of Liang et al. (2021) that many hedge funds engage in "greenwashing"—i.e., holding a brown portfolio while advertising themselves as green investors.

UNPRI Signatory

We also examine whether signatories of the United Nations Principles for Responsible Investment (UNPRI) have stronger sustainable demand. The UNPRI describes itself as the "leading proponent of responsible investment" and one of its six main principles is that "[investors] will be active owners and incorporate ESG issues into our ownership policies and practices."[12] We use fuzzy string matching to match the UNPRI signatory list with the institutional investors in FactSet data. In an average quarter in our sample, 9.9% of institutional investors are UNPRI signatories, but they constitute 38.9% of total institutional AUM. The fraction of institutional AUM controlled by UNPRI signatories has grown from 18.1% in 2013Q1 to 52.1% in 2021Q3.

However, we show from cross-sectional and time-series regressions that UNPRI signatory status is not associated with higher sustainable demand. Table 3.C5.2 shows that after controlling for other investor characteristics, UNPRI signatories do not exhibit higher sustainable demand than non-signatories.[13] Table 3.C5.3 further shows that signing the UNPRI is not significantly associated with any within-investor demand change, after controlling for within-investor time trends. Our results are consistent with the findings of Kim and Yoon (2023) and Liang et al. (2021) that investors who signed the UNPRI have not necessarily increased their sustainable portfolio holdings.

3.5 Investor Pressure for Sustainability and Real Effects

Using the estimated demand system, we next study the real effects of sustainable investing. To achieve this goal, we first use the estimated demand curves to quantify investor pressure for sustainability, which captures the price pressure a firm receives, through investor demand, to become more sustainable. While the average firm has experienced greater investor pressure to become more sustainable, we find that higher investor pressure today predicts greater future sustainability only to a limited extent, thus highlighting the limited real impact of sustainable equity investing.

3.5.1 Quantifying Investor Pressure

We first derive a closed-form expression for investor pressure and show that the average firm has experienced greater investor pressure to become more sustainable. Importantly, we show that the investor pressure is the average of the coefficients on sustainability of the firm's investors adjusted

[12] See UNPRI's website for these descriptions of the main principles (Link to UNPRI Principles) and the signatory list (Link to UNPRI Signatories), last accessed in December 2022.

[13] Without controls, UNPRI signatories have statistically significant higher demand only for environmental score.

for their price elasticity.

We define the investor pressure of firm n for the kth sustainability characteristic $s_{k,t}(n)$ as the equilibrium price impact of a small change in $s_{k,t}(n)$, holding all other stock characteristics and all investor demand curves constant:

$$\frac{\partial mb_t(n)}{\partial s_{k,t}(n)}. \tag{3.8}$$

This can be computed analytically from the demand system as below, with the proof in Appendix C3:

Proposition 3. *The price impact of a change in the value of sustainability characteristic k for firm n, denoted as* $\mathbf{M}$*, is given as the nth diagonal element of the matrix*

$$\mathbf{M} := \frac{\partial \mathbf{p}}{\partial \mathbf{s}_k} = \left(\mathbf{I} - \sum_i \beta_{0i} A_i \mathbf{H}^{-1} \mathbf{G}_i\right)^{-1} \left(\sum_i \beta_{ki} A_i \mathbf{H}^{-1} \mathbf{G}_i\right) \tag{3.9}$$

where β_{ki} *is investor i's demand coefficient for sustainability characteristic* s_k*, and we omit time subscripts for simplicity. The matrices* $\mathbf{H}$ *and* $\mathbf{G}_i$ *are defined as follows:*

$$\mathbf{H} := diag\left(\sum_i A_i \mathbf{w}_i\right) = \sum_i A_i diag\left(\mathbf{w}_i\right)$$
$$\mathbf{G}_i := diag\left(\mathbf{w}_i\right) - \mathbf{w}_i \mathbf{w}_i'.$$

The quantity $\mathbf{M}_{n,n}$, which is the nth diagonal entry of $\mathbf{M}$, can be interpreted as the price pressure that a firm receives through investor demand. Put differently, it represents the firm's marginal benefit in terms of stock valuations derived from increasing its sustainability characteristic $s_{k,t}(n)$.[14] Public firms have incentive to increase their stock valuations as they are related to both equity cost of capital and value of share-based compensations.

Because we hold latent demand constant in the calculations above, the measure of investor pressure only captures the intensive margin of investor demand and thus is a lower bound on the actual investor pressure that a firm may receive. If substantial variation in holdings operates through the extensive margin, then the current methodology understates how $mb_t(n)$ could change

[14] We recognize that ideally, we need a fully micro-founded model with the supply side, or the firm side, of the demand system to relate this quantity back to the firms' objectives. Only in this way can we also account for the adjustment cost of making the marginal change, but this is outside the scope of this paper. Instead, we control for observed firm characteristics and industry classification in our empirical analysis and argue that doing so we can compare firms with similar adjustment or marginal cost of changing the characteristic in question.

with $s_t(n)$ as new investors would start to hold the stock if the firm improves sufficiently. We note, however, that such response on the extensive margin is not a first-order concern in our setup, as Koijen and Yogo (2019) shows that the set of stocks that institutions invest in is usually small and highly persistent.

The matrix inside the inverse in Equation (3.9) is the aggregate demand elasticity. Therefore, the valuation of a stock reacts more to a change in characteristic if the stock is held by less price-elastic investors.[15] In addition, the nth diagonal entry of the second term is

$$\frac{\sum_i \beta_{ki} A_i w_i(n)(1 - w_i(n))}{\sum_i A_i w_i(n)}. \tag{3.10}$$

This quantity can be viewed as an AUM weighted average of the coefficients on the sustainability characteristic. Therefore, investor pressure for a given firm n is a weighted average of sustainability demand coefficients of its institutional owners, adjusted for their price elasticity. If a firm faces a representative owner who is price inelastic and exhibits a high coefficient on the sustainability characteristic, this firm faces a large investor pressure on that dimension of sustainability.

Table 3.C5.4 reports the summary statistics for investor pressure of each sustainability characteristic. The average investor pressure for environmental score is 0.063: an average firm could expect its valuation to increase by 0.63% if it improves environmental score by 0.1 standard deviation. The average pressure for emissions intensity is -0.084: an average firm could expect its valuation to increase by 0.84% if it improves Scope 1 carbon emissions intensity by 0.1 standard deviation. The average pressure for green innovation is close to zero. In addition, the investor pressure for each characteristic shows strong heterogeneity across stocks. For example, the investor pressure for emission is -0.192 at 10th percentile and 0.004 at 90th percentile.

Figure 3.5 plots the average of investor pressure across stocks for each sustainability characteristic in each quarter. The investor pressure for environmental score is positive and roughly constant over time, and the investor pressure for carbon emissions intensity is negative and increasing over time in magnitude (a near twofold increase from -0.039 in 2013Q1 to -0.110 in 2021Q3). The investor pressure for green innovation is around zero for the entire sample. These results are in line with the time series evolution of sustainable demand coefficients, as discussed in Section 3.4.2.

The increasing investor pressure against carbon emission complements our previous finding in Section 3.4.3 that the trend towards higher sustainable demand has not been "undone" by price-elastic investors. Because investor pressure for a characteristic is increasing in the demand for that

[15] The intuition is that a change of characteristic induces change in characteristic-based demand. A same amount of demand shift creates stronger price pressure in stocks with less price-elastic demand curves.

characteristic but decreasing in investors' price elasticity, one might be concerned *ex ante* that if investors have become more price-elastic over time, the increasing demand for low-emission stocks we documented in Section 3.4.2 may not translate into higher investor pressure. Figure 3.C5.1 shows that this concern is not realized: the investors on average has in fact become more price inelastic during our sample period, as the average price inelasticity coefficient ($\beta_{0,i,t}$) increased from 0.67 in 2013Q1 to 0.72 in 2021Q3. Therefore, both the increasing demand for low-emission stocks and the decreasing price elasticity contribute to higher investor pressure for the emission characteristic.

3.5.2 Investor Pressure and Future Environmental Performance

We next examine whether a firm's investor pressure is associated with its future environmental performance. Table 3.6 shows results from regressing future 1-year environmental performance on a stock's characteristics and investor pressures. Column (1) shows that in the cross-section, 1 standard deviation higher pressure for environmental score predicts 0.0187 standard deviation higher change of environmental score over the next year. Column (2) shows that a 1 standard deviation higher pressure for emission (more negative pressure) predicts 0.0251 standard deviation more reduction of emissions intensity over the next year. Column (3) shows similar results for green innovation. Table 3.C5.5 further shows that the results are robust for 2-year and 3-year future environmental performances.

We draw two conclusions from the regression results. On the one hand, the coefficients highlighted above are positive and statistically significant after controlling for a stock's current non-green and green characteristics. Therefore, higher investor pressure predicts better future environmental outcomes in the cross-section, even after controlling for current green performance. On the other hand, the relationship between investor pressure and future 1-year environmental outcome is very small in magnitude. The small magnitude might be due to both the indirect relationship between stock valuation and firm decision-making and the long time needed to make environmental improvements.[16]

Taken together, our results in this section show that the average firm has experienced greater investor pressure to reduce its carbon emissions intensity, but higher investor pressure today only has weak predictive power for future sustainability performance.

[16] Stock valuation could affect firm decision-making in two ways: first through expected stock return and thus the firm's cost of equity capital, and second through the incentive for firm management through stock-based compensations. Both relationships are indirect in nature, therefore a large change of valuation might be needed to generate a small change of management incentives and real investment decisions.

3.6 Counterfactuals and Asset Pricing Implications

We next study the asset pricing effects of sustainable equity investing through two counterfactual scenarios related to investor preferences and measures of sustainability. Specifically, we calculate counterfactual asset prices under each scenario and examine how the valuation patterns from Section 3.2.2 subsequently change. The steps for computing counterfactuals closely follow the algorithm in Koijen and Yogo (2019), which we detail in Appendix C4. In the first counterfactual, we explore the impact of imposing ESG-agnostic mandates on select investors. In the second, we consider incorporating the firm's green patents, which is not demanded by investors, into the construction of the environment score, a characteristic that is heavily demanded.

3.6.1 Impact of ESG-Agnostic Mandates

In the first counterfactual analysis, we "shut off" the demand for sustainability of select investors to examine the contribution of different investors to the valuation of sustainability characteristics. Our counterfactual exercise corresponds to "ESG-agnostic" policies that force a subset of investors to not consider sustainability in their portfolio choice decisions.

Specifically, we consider three counterfactual scenarios where we set the demand coefficients for all three sustainable characteristics to zero for: (a) active institutions, (b) all institutions, and (c) all investors including the household sector.[17] In each scenario, we regress counterfactual valuations on stock characteristics as in Section 3.2.2, and we focus on the coefficients ("valuation gap") on the three green characteristics. We attribute the difference between actual data and scenario (a) to active institutions, the difference between scenarios (a) and (b) to passive institutions, and the difference between scenarios (b) and (c) to households.[18] Any remaining valuation gap in counterfactual scenario (c) for green characteristics are driven by the extensive margin of asset demand.[19]

Table 3.7 shows the valuation regression coefficients on three green characteristics under each counterfactual scenario: panel (a) for the full sample from 2013 to 2021, and panel (b) for the subsample from 2018 to 2021. First, the coefficients for environmental score in panel (a) show that the valuation gap between high-environmental score and low-environmental score stocks are not

[17] Same as Section 3.4.2, we define "passive" institutional investors as large investment advisors, medium- or small-passive investment advisors, and long-term investors; and we define "active" institutional investors as medium- or small-active investment advisors, hedge funds, private banking, and brokers.

[18] Note that this attribution is subject to the order in which we shut off green demand for different types of investors. Table 3.C5.6 shows that our results are similar when we shut off the demand for passive institutions before active institutions.

[19] The asset demand system we estimate only models the intensive margin of asset demand as a function of stock characteristics. If a stock characteristic is correlated with the extensive margin of portfolio choice (i.e., the fraction of investors that hold the stock in their portfolios, or have the stock in their investment universes), then the stock characteristic will still be positively correlated with stock valuation even without any demand for it on the intensive margin.

driven by the green demand of institutional investors. In the data, stocks with 1 standard deviation higher environmental score is associated with 10.7% higher market-to-book ratio on average. This valuation gap becomes 9.51% if we shut off the green demand of all institutional investors in column (3), and it becomes 5.4% if we shut off green demand of all investors in column (4). Therefore, the positive relationship between environmental score and valuation is driven roughly in half by the households' demand for high environmental score stocks, and in another half by the extensive margin of asset demand. On the extensive margin, Table 3.C5.7 provides suggestive evidence that stocks with higher environmental score appear in the investment universe of a higher fraction of institutional investors, after controlling for other characteristics. Therefore, stocks with higher environmental score would have higher valuation even without any sustainable demand on the intensive margin.

Second, the coefficients for emissions intensity in both panels show that the gradually increasing valuation gap between low- and high-emission stocks is entirely driven by the green demand of institutional investors. Column (3) of panel (a) shows that emissions intensity would become *positively* correlated with valuation in the full sample if we shut off the green demand of all institutions. Column (3) of panel (b) also shows that the negative relationship between emissions intensity and valuation from 2018 to 2021 would disappear if we shut off green demand of all institutions. Together, these results show that the demand of institutional investors against high-emission stocks is the main driver of their lower valuations.

Finally, the coefficients for emissions intensity in panel (b) also show that active and passive institutions contribute equally to the valuation gap between low and high-emission stocks. In the 2018-2021 sample, 1 standard deviation higher emissions intensity is associated with a 6.45% lower market-to-book ratio. This valuation gap is reduced to 3.45% if we shut off the green demand for active institutions in column (2) and becomes a statistically insignificant at 0.31% if we shut off the green demand for all institutions in column (3). Figure 3.6 also provides visual evidence by plotting the time series of quarterly valuation regression coefficients for emissions intensity in the data and the counterfactual scenarios. Together, these results reinforce our previous finding that active institutional investors as a whole are contributing to sustainable investing thanks to their demand for low-emission stocks instead of "undoing" sustainable investing by buying up high-emission stocks.

3.6.2 Hypothetical ESG Ratings: Incorporating Green Innovation

In previous analyses, we show that our measure of green innovation—ratio of green patents to asset—is not correlated with valuation in the cross-section, and does not have positive characteristic demand from investors on average. We also showed in panel (b) of Table 3.C5.1 that our green patents measure is negatively correlated with environmental score. Because investors have positive demand for the environmental score on average, we can possibly increase the valuation of top green innovators—thereby encouraging their green innovation effort—by incorporating green innovation into the construction of environmental score. For example, if MSCI adjusts its environmental score definition to make the score higher for top green innovators, investors with positive demand for environmental score will be attracted to these top innovators, which will in turn push up their valuation.

In the second counterfactual analysis, we show that this change would increase the valuation of top green innovators that produce the most green patents, but at the same time the valuation gap between low- and high-emission stocks will still persist in the counterfactual equilibrium. These results show that the proposed change could increase firms' incentive to produce green patents without jeopardizing the main effect of sustainable investing on the valuation of emission.

Setup and Assumption

To proceed with the counterfactual, we start by increasing the environmental score for top 20% green innovators, ranked by our green patent characteristic. Specifically, we increase the standardized environmental score by 1.0 for top 20% green innovators in each quarter, and re-standardize the modified environmental score within each quarter. In an average year, the top 20% of green innovators develop 8,163 green patents, which make up 95.5% of the annual green patents developed by the firms in our sample.

Importantly, we keep all demand coefficients unchanged in these counterfactual simulations. Thus, our analyses is based on the assumption that modifying the content of environmental score will not change investors' demand curves for the standardized environmental score. While this assumption is strong and cannot be empirically verified, we offer two supporting arguments for making this assumption. First, institutional investors' demand for high-environmental score stocks is at least partially driven by an environmental score target, which can be either required by end investors or self-imposed in order to attract flows (see, e.g., Hartzmark and Sussman 2019 for evidence of mutual fund investors chasing funds with high sustainability rating on Morningstar).

This type of demand for high-environmental score stocks is likely unrelated to the specific content of environmental score and thus unlikely to change in our counterfactual scenario. Second, green innovation is more likely to be perceived by investors as a positive signal for better future environmental performance (or financial performance, if the green technology like electric vehicle can be monetized). Therefore, it is unlikely that investors will reduce their environmental score demand if they know that environmental score becomes more correlated with green innovation. Based on these two reasons, we keep the demand coefficients unchanged in these counterfactual simulations.

Results

In Table 3.8, we examine the difference of log market-to-book ratios between actual data and our counterfactual simulation. Column (1) shows that in the counterfactual simulation, the average log market-to-book ratio for the top 20% green innovators will increase by 0.047, and the average log market-to-book for other stocks will decrease by 0.022. These can be interpreted as a 4.7% increase of market-to-book for the top green innovators, and a 2.2% decrease of market-to-book for other stocks.[20] Columns (2) and (3) show that the relative valuation change between green innovators and non-innovators stays roughly the same after controlling for time fixed effects and non-green characteristics. Therefore, our proposed change of environmental score construction would move stock valuations in the intended direction.

Given the positive correlation between green innovation and current emissions intensity, one might be concerned that our proposed change of environmental score would also increase the valuation for high-emission stocks. Column (4) of Table 3.8 shows that this concern is unnecessary: the valuation regression coefficient on emissions intensity does not change significantly in the counterfactual simulation. This result shows that our proposed change of environmental score can be implemented without attenuating the existing achievement of sustainable investing (i.e., the valuation gap between low- and high-emission stocks).

3.7 Conclusion

In this paper, we investigate the heterogeneity in investor demand for sustainability and study its implications for firm decisions and asset prices. By utilizing a comprehensive set of measures for

[20] These results are based on the average of log market-to-book ratios across innovator or non-innovator stocks. Alternatively, we can treat each type of stocks as one asset and computes its aggregated market-to-book (sum of market equity divided by sum of book equity). Using this method, the valuation for top 20% of green innovators will increase by 7.17% in an average quarter, and the valuation for non-green-innovators will decrease by 3.62% in an average quarter.

firm sustainability as well as a structural asset demand system, we provide a more nuanced understanding of the heterogeneous demand for sustainability among investors. Our findings suggest that while investors do value third-party environment scores and low emissions intensity, they do not value firms that innovate in sustainable technologies. In addition, investor pressure generated by these demand patterns translate into only limited improvements in firm sustainability.

Overall, we contribute to the ongoing conversation about sustainable equity investing by providing a more detailed understanding of investor demand for sustainability and its implications for the effectiveness of sustainable equity investing. Our framework could be used by future research to explore how the demand for sustainability varies across different regions (e.g., U.S. vs Europe) or asset markets (e.g., green bonds), or to study sustainable demand at a more granular level (e.g., for individual mutual funds or households). We leave these directions for future research.

Tables for Chapter 3

Table 3.1: Summary Statistics

(a) Distribution of Characteristics

Characteristic	Count	Mean	Std. Dev.	Min	Q10	Q50	Q90	Max	% Miss
Market Equity	64,184	14.0	51.7	0.3	0.8	3.0	27.4	2324.4	0
Market / Book	64,184	8.8	160.2	0.0	1.1	2.8	10.5	18596.4	0
Profit / Asset	64,184	0.26	0.22	−0.50	0.04	0.23	0.55	0.98	0
Asset Growth	64,184	0.10	0.20	−0.52	−0.07	0.06	0.33	1.20	0
Dividend / Book Equity	64,184	0.04	0.05	0.00	0.00	0.02	0.10	0.29	0
CAPM Beta	64,184	1.24	0.60	−0.34	0.52	1.18	2.02	3.34	0
Non-Green Patents (bps)	64,184	315	687	0	0	11	1065	3525	0
Green Patents (bps)	64,184	15	43	0	0	0	42	293	0
Environment Score	54,142	−1.56	1.33	−8.55	−3.50	−1.29	−0.17	0.00	0.101
Emissions Intensity	47,510	189.7	660.8	0.1	0.7	13.6	343.3	6315.6	0.039

(b) Investors by Type (Average Across Quarters)

	Count	AUM ($bil)	AUM Share	Active Share
Large IA	25	9178.2	0.307	0.301
Medium-Passive IA	42	2156.0	0.073	0.287
Medium-Active IA	43	1899.6	0.066	0.592
Small-Passive IA	854	1523.0	0.051	0.412
Small-Active IA	917	2074.4	0.072	0.659
Long Term	104	1060.6	0.037	0.344
Hedge Funds	351	975.4	0.034	0.711
Brokers	83	910.6	0.031	0.511
Private Banking	816	690.1	0.023	0.530
Household	1	9065.5	0.307	0.212

This table reports the summary statistics for select variables. Panel (a) provides the distribution of stock characteristics. Panel (b) summarizes the investor types in our sample. The provided statistics are computed for each quarter and then averaged across quarters. In panel (a), profitability is defined as revenue minus cost of goods sold; asset growth rate is computed over 1 year; dividend to book equity is computed based on rolling 1-year dividend; and CAPM beta is estimated based on rolling 60 months of data. In panel (b), active share is defined as one half of sum of the difference between the weight of each stock in the portfolio and the market weight.

Table 3.2: Valuation Regressions

	Dependent Variable: Log Market-to-Book		
	(1) Full Sample	(2) 2013-2017	(3) 2018-2021
Log Book Equity	-0.390^{**}	-0.349^{**}	-0.451^{**}
	[−18.78]	[−16.80]	[−17.66]
Profit / Asset	0.198^{**}	0.188^{**}	0.217^{**}
	[12.08]	[10.95]	[10.35]
Asset Growth	0.121^{**}	0.0990^{**}	0.147^{**}
	[9.542]	[6.816]	[7.952]
Dividend / Book Equity	0.274^{**}	0.261^{**}	0.292^{**}
	[16.68]	[13.85]	[15.54]
CAPM Beta	-0.0482^{**}	-0.0442^{**}	-0.0535^{*}
	[−3.483]	[−3.218]	[−2.199]
Non-Green Patents	0.142^{**}	0.114^{**}	0.184^{**}
	[7.961]	[6.517]	[8.175]
Environment Score	0.107^{**}	0.0952^{**}	0.115^{**}
	[8.191]	[6.738]	[6.767]
Emissions Intensity	−0.0156	0.0133	-0.0645^{**}
	[−1.153]	[1.056]	[−4.261]
Green Patents	0.00359	0.0121	−0.0115
	[0.237]	[0.715]	[−0.603]
Year-Quarter FE	✓	✓	✓
Within R^2	.393	.377	.425
Observations	64,184	37,933	26,251

This table summarizes the results from the valuation regression as shown in Equation (3.1). Specifically, we regress the market-to-book ratio on a vector of time-varying firm characteristics and year-quarter fixed effects. As sustainability characteristics, we include environment score, emissions intensity, and green patents as described in Section 3.2.1. Column (1) displays the results based on the entire sample, while columns (2) and (3) display results for sub periods. Standard errors are clustered by year-quarter.

Table 3.3: Demand Coefficients: Summary Statistics

	AUM-Weighted		Equal-Weighted				
	Mean	SD	Mean	SD	Q10	Q50	Q90
Log Market to Book	0.699	0.365	0.349	0.615	−0.485	0.454	0.990
Log Book Equity	1.275	0.406	0.693	0.598	−0.053	0.674	1.498
Profitability	0.010	0.179	0.046	0.371	−0.376	0.034	0.480
Asset Growth	0.031	0.151	0.082	0.327	−0.279	0.057	0.476
Dividend / Book	0.079	0.205	0.027	0.336	−0.370	0.019	0.436
Market Beta	−0.028	0.176	−0.092	0.374	−0.556	−0.067	0.339
Non-Green Patents	0.020	0.210	−0.052	0.472	−0.562	−0.021	0.423
Environment Score	0.031	0.130	0.023	0.272	−0.287	0.013	0.343
Emissions Intensity	−0.023	0.126	−0.085	0.251	−0.394	−0.061	0.190
Green Patents	−0.003	0.168	−0.026	0.376	−0.407	0.004	0.341

This table provides summary statistics for the demand coefficients estimated from Equation (3.6). The three sustainability characteristics—environment score, emissions intensity, and green patents—are as described in Section 3.2.1. For each coefficient, we compute the summary statistics across investors in every quarter and then construct an AUM(equal)-weighted average across quarters.

Table 3.4: Demand for Emissions Intensity: Time Trend Analysis

	Dep Var: Emission Demand Coef	
	(1)	(2)
Time Trend	-0.00114^{**}	-0.00121^{**}
	$[-4.718]$	$[-4.670]$
Investor FE		✓
AUM-Weighted	✓	✓
Observations	113,231	112,942

This table summarizes the within-investor time trend of demand coefficients for emissions intensity. Data is at the investor-quarter level, and each observation is weighted by the investor's AUM share in a quarter.

Table 3.5: Demand for Sustainability and Investor Characteristics

	Dep Var: Demand Coefficient		
	(1) Environment Score	(2) Emissions Intensity	(3) Green Patents
Price Inelasticity	-0.0796^{**}	0.155^{**}	0.00875
	[−6.268]	[12.03]	[0.732]
Log AUM	0.0521^{**}	−0.0109	−0.00745
	[5.236]	[−1.102]	[−0.992]
Active Share	0.0223^{+}	-0.0671^{**}	-0.0814^{**}
	[1.891]	[−5.174]	[−7.298]
Turnover	-0.0184^{+}	0.0389^{**}	0.00395
	[−1.909]	[3.843]	[0.406]
1{Non−USA}	0.122^{**}	-0.0581^{+}	-0.126^{**}
	[4.724]	[−1.764]	[−4.935]
1{Style=Value}	0.0321	0.0509	0.0231
	[1.014]	[1.607]	[0.789]
1{Style=Growth}	0.0885^{**}	-0.165^{**}	0.0112
	[3.201]	[−5.420]	[0.418]
1{Hedge Fund}	0.0680^{+}	0.0981^{*}	-0.158^{**}
	[1.984]	[2.543]	[−4.598]
1{Priv. Banking}	0.00648	−0.0333	−0.0104
	[0.283]	[−1.386]	[−0.465]
1{Long Term}	−0.0570	−0.00978	-0.0583^{+}
	[−1.407]	[−0.210]	[−1.864]
1{Broker/Dealer}	0.0514	0.242^{**}	0.0246
	[1.280]	[4.470]	[0.772]
Time FE	✓	✓	✓
Within R^2	.013	.041	.013
Observations	113,231	113,231	113,231

This table summarizes the relationship between demand for sustainability and investor characteristics via cross-sectional regressions. The three sustainability characteristics—environment score, emissions intensity, and green patents—are as described in Section 3.2.1. Data is at the investor-quarter level., and all variables are cross-sectionally standardized.

Table 3.6: Investor Pressure and Future Environmental Performance

	Forward 1-Year Outcome		
	(1) Environment Score	(2) Emissions Intensity	(3) Green Patents
Pressure: Environment Score	0.0187^{+}	-0.000276	-0.0147^{*}
	[2.006]	[−0.0259]	[−2.248]
Pressure: Emission Intensity	-0.0279^{**}	0.0251^{**}	0.00184
	[−3.097]	[2.880]	[0.319]
Pressure: Green Patents	0.00181	0.00322	0.00983^{+}
	[0.206]	[0.342]	[1.996]
Environment Score	-0.228^{**}	-0.0617^{**}	-0.0103^{+}
	[−19.53]	[−2.963]	[−1.740]
Emission Intensity	-0.0323^{*}	−0.0178	0.00130
	[−2.648]	[−1.225]	[0.268]
Green Patents	−0.00625	0.00685	0.769^{**}
	[−0.706]	[0.703]	[37.07]
Time FE	✓	✓	✓
Non-Green Controls	✓	✓	✓
Within R^2	.053	.007	.668
Observations	51,065	51,065	51,065

This table summarizes the cross-sectional relationship between investor pressure and future environmental performance. The dependent variables are: future one-year change in environment score, future one-year change in emissions intensity, and the future one-year green patents. The main independent variables are the investor pressure for three green characteristics. "Non-Green" control variables include log book equity, investment, profitability, market beta, and dividend to book equity. All outcome variables and regressors are cross-sectionally standardized in each quarter.

Table 3.7: Counterfactual Exercise: Impact of ESG-Agnostic Mandates

(a) Counterfactual Valuation Regressions: Full Sample (2013 - 2021)

	Data	CF: Shut off Green Demand		
	(1)	(2) Active Inst	(3) All Inst	(4) All Inst + HH
Environment Score	0.107^{**}	0.0947^{**}	0.0951^{**}	0.0540^{**}
	[8.191]	[7.463]	[7.194]	[4.103]
Emissions Intensity	-0.0156	0.0101	0.0330^{*}	0.0387^{**}
	[-1.153]	[0.759]	[2.634]	[3.065]
Green Patents	0.00359	0.0309^{+}	0.0249	0.0161
	[0.237]	[1.942]	[1.518]	[0.993]
Time FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	64,184	64,184	64,184	64,184

(b) Counterfactual Valuation Regressions: Subsample (2018 - 2021)

	Data	CF: Shut off Green Demand		
	(1)	(2) Active Inst	(3) All Inst	(4) All Inst + HH
Environment Score	0.115^{**}	0.107^{**}	0.114^{**}	0.0698^{**}
	[6.767]	[6.305]	[6.529]	[4.010]
Emissions Intensity	-0.0645^{**}	-0.0345^{*}	0.00311	0.00803
	[-4.261]	[-2.266]	[0.199]	[0.506]
Green Patents	-0.0115	0.0241	0.0150	0.00618
	[-0.603]	[1.262]	[0.784]	[0.327]
Time FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	26,251	26,251	26,251	26,251

This table presents the results from valuation regressions in counterfactual scenarios where we "shut off" green demand, i.e., set the demand coefficients for all three sustainability characteristics to zero for one or more types of investors. Specifically, we consider three scenarios in which we "shut off" green demand for all active institutions (column (2)), all institutions (column (3)), and all investors including the household sector (column (4)). After obtaining counterfactual valuations, we re-estimate the valuation regression as shown in Equation (3.1). We present results for both the full sample as well as the 2018-2021 subsample.

Table 3.8: Counterfactual Exercise: Hypothetical ESG Ratings with Green Patents

	Dep Var: Log M/B, Counterfactual − Actual			
	(1)	(2)	(3)	(4)
1{Top 20% Green Innovator}	0.0470**	0.0692**	0.0684**	
	[15.15]	[18.79]	[17.80]	
1{Not Top Green Innovator}	−0.0222**			
	[−26.40]			
Emissions Intensity				0.000495
				[0.870]
Time FE		✓	✓	✓
Controls			✓	✓
Observations	64,184	64,184	64,184	64,184

This table presents the results from the counterfactual exercise that explores the consequences of adjusting the environmental score definition to reward companies with high green innovation. Specifically, we increase the environmental score by 1.0 for the top 20% green innovators as ranked by our green patent characteristic in each quarter. We then re-standardize the modified environmental score within each quarter. We keep the demand coefficients unchanged in these counterfactual simulations. After obtaining the counterfactual valuations, we compute the difference of log market-to-book ratios between the actual data and the counterfactual. In columns (3-4), control variables include all non-green stock characteristics in the valuation regressions plus the environmental score.

Figures for Chapter 3

Figure 3.1: Valuation Regressions: Time-Series of Coefficients

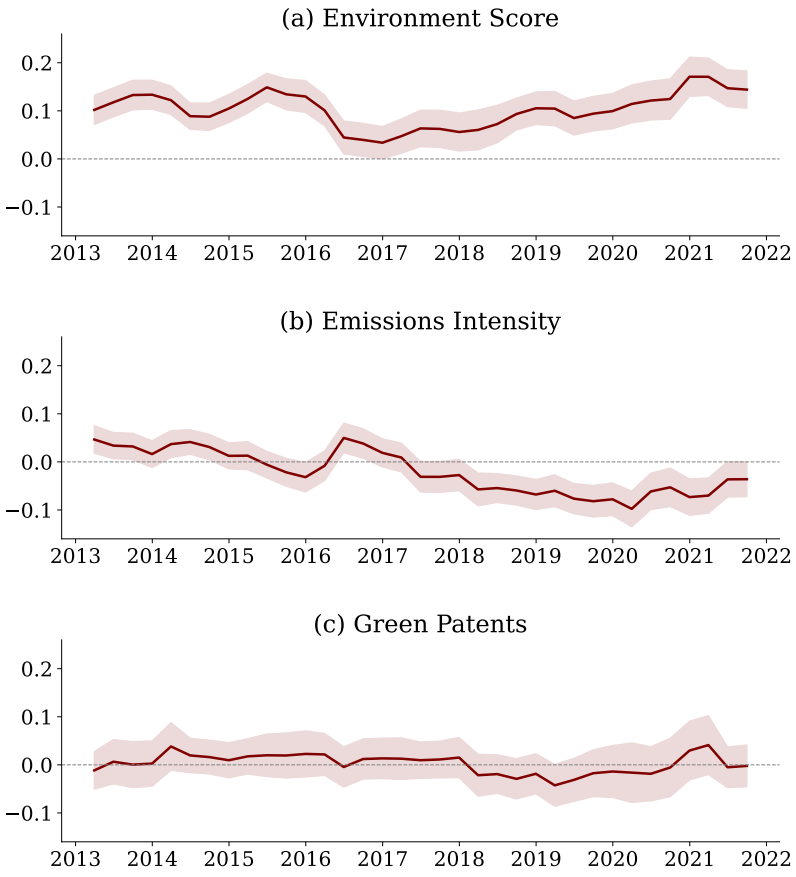

This figure plots the time-series of the coefficients obtained from the valuation regression as shown in Equation (3.1). We first estimate Equation (3.1) cross-sectionally for each quarter. The shaded area represents the 95% confidence interval around the mean.

Figure 3.2: Binned Scatterplots of Coefficients across Investors

(a) Price Elasticity vs. Demand for Sustainability

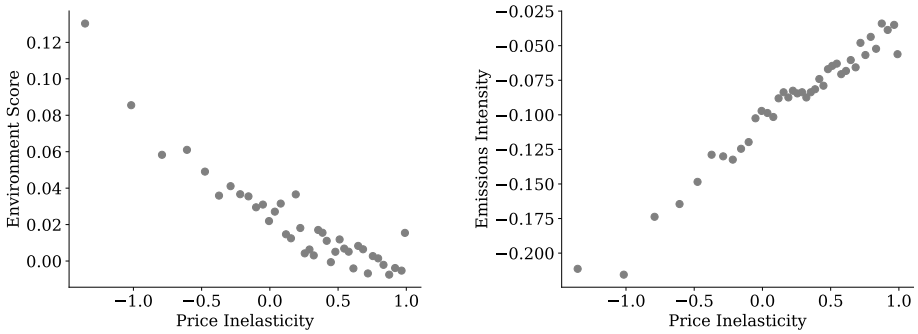

(b) Demand across Sustainability

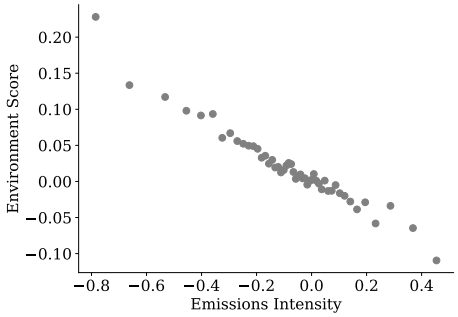

This figure summarizes the relationship between pairs of demand coefficients across investors. In panel (a), we plot the binscatter of coefficient on environment score against price inelasticity (left) as well as the coefficient on emissions intensity on price inelasticity (right). In panel (b), we plot the coefficient on environment score against the coefficient on emissions intensity.

Figure 3.3: Trends in Average Coefficients

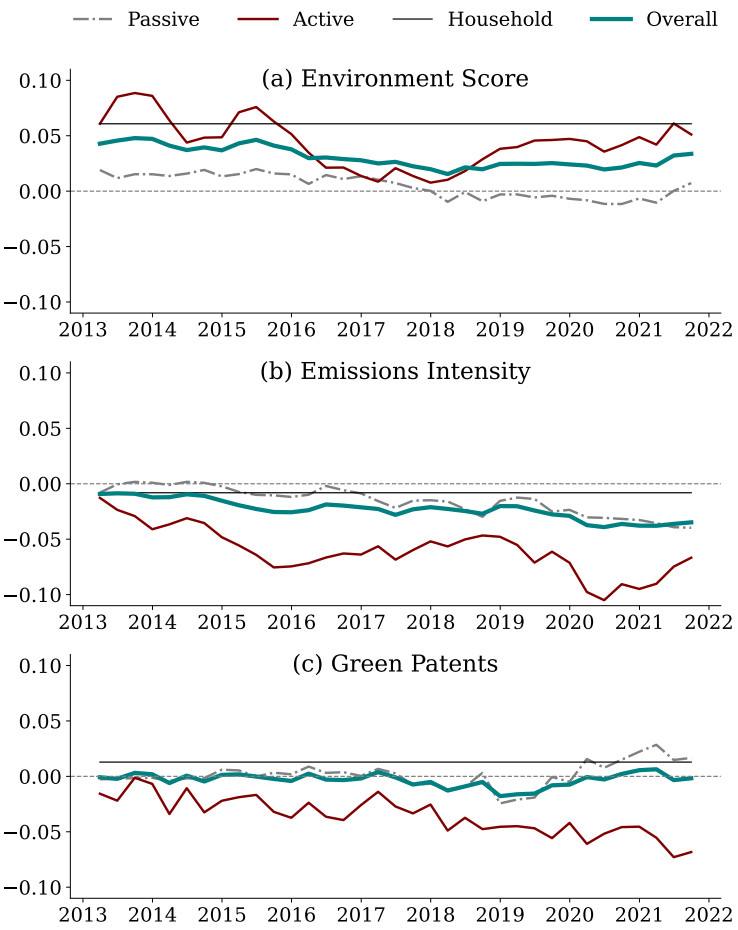

This figure summarizes the time trends in the AUM-weighted average coefficient of investors. We define "passive" institutional investors as large investment advisors, medium- or small-passive investment advisors, and long-term investors; and we define "active" institutional investors as medium- or small-active investment advisors, hedge funds, private banking, and brokers.

Figure 3.4: Demand Coefficient for Emission: Role of Within-Investor Preference Shift

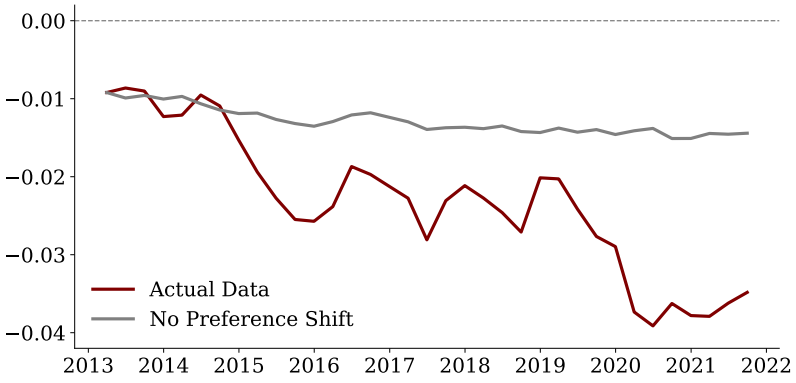

This figure plots the AUM-weighted average demand efficient on emissions intensity across time both in the actual data and in the counterfactual where there is no within-investor shift in preference. Specifically, let $T_{0,i}$ be the first quarter when investor i appears in our sample and let $\beta_{i,GHG,t}$ be investor i's demand coefficient for Log GHG1 intensity in quarter t. For each quarter, we compute the average of demand coefficients in the data $\left(\bar{\beta}_{GHG,t}\right)$:

$$\bar{\beta}_{GHG,t} := \frac{\sum_i A_{i,t}\beta_{i,GHG,t}}{\sum_i A_{i,t}}$$

as well as the average of the counterfactual coefficients if there is no shift in preferences $\left(\bar{\beta}^*_{GHG,t}\right)$:

$$\bar{\beta}^*_{GHG,t} := \frac{\sum_i A_{i,t}\beta_{i,GHG,T_{0,i}}}{\sum_i A_{i,t}}$$

We then plot the time-series of both $\bar{\beta}_{GHG,t}$ and $\bar{\beta}^*_{GHG,t}$.

Figure 3.5: Time Series of Average Investor Pressure across Stocks

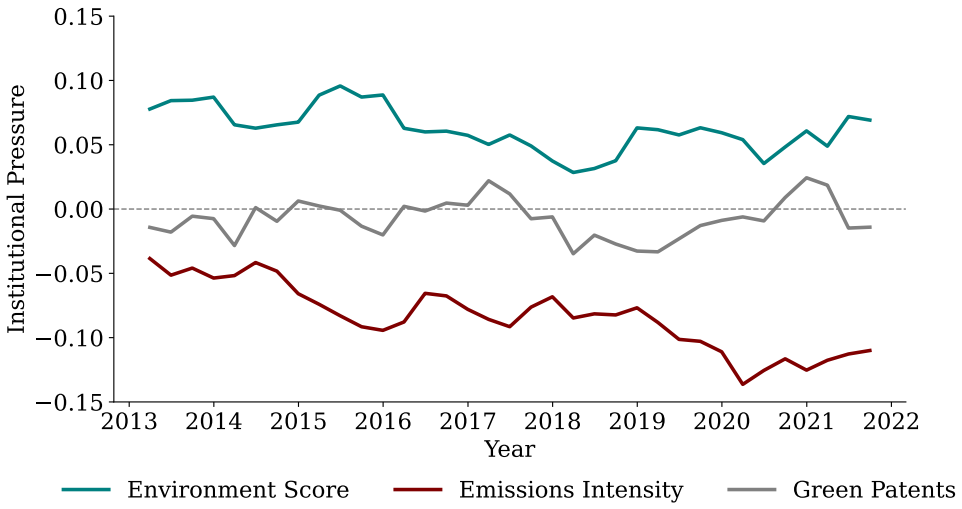

This figure plots the average investor pressure across stocks for each quarter and each green characteristic. We compute the investor pressure for sustainability for each stock using the procedure described in Section 3.5.

Figure 3.6: Counterfactual: Impact of ESG-Agnostic Mandates

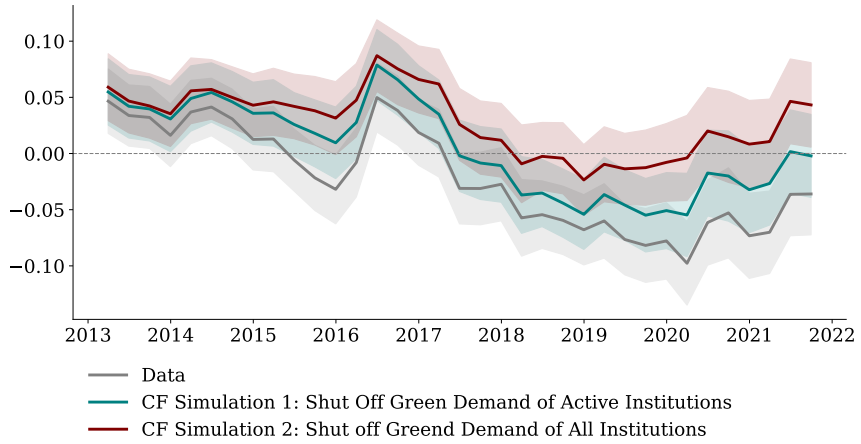

This figure plots the time-series of the quarterly valuation regression coefficients on emissions intensity in the data and in the counterfactual where we shut off green demand of select investors.

Appendix to Chapter 3

C1 Data

C1.1 MSCI Data

As discussed in Section 3.2.1, we define the *raw environment score* as:

$$g_t(n) = \frac{-\left(10 - E_t(n)\right) w_t^E(n)}{100}$$

where $E_t(n)$ is the environmental pillar score provided by MSCI, and $w_t^E(n)$ is the environmental pillar weight provided by MSCI. $E_t(n)$ measures a firm's environmental performance relative to peers, and $w_t^E(n)$ measures the importance of environmental issue for the firm. $E_t(n)$ ranges from 0 (worst) to 10 (best) and represents the weighted average score across various dimensions related to environmental issues.

Figure 3.C5.2 illustrates the distribution of $E_t(n)$, $w_t^E(n)$, and $g_t(n)$ for stocks in the Oil & Gas and Banking industries. The distribution of MSCI environmental pillar score $E_t(n)$ is similar across the two industries, which indicates that $E_t(n)$ is constructed based on peer-to-peer comparisons. However, the MSCI pillar weights $w_t^E(n)$ for oil & gas stocks are significantly higher than those for banking stocks. Therefore, the pillar weights are necessary for measuring the absolute (rather than peer-adjusted) level of companies' environmental performance. Figure 3.C5.2 confirms that our raw environmental score $g_t(n)$ is significantly higher for banking stocks compared to oil & gas stocks.

C1.2 PatentsView Data

We obtain data of granted U.S. patents from the PatentsView database provided by the U.S. Patents and Trademark Office.[21] For each patent granted from 1975 to 2021, we have the name of patent assignee (the company that applied for the patent) and two patent classification codes (Cooperative Patent Classification and International Patent Classification codes—CPC and IPC).

We first match assignee names in the patent data to companies in CRSP. We start from the mappings from individual patents to Compustat firm identifiers (GVKEY) provided by Autor et al. (2020) from 1975 to 2012 and by the WRDS US Patents database from 2011 to 2019. Based on these two mappings, we construct a mapping from the PatentsView assignee names to Compustat GVKEY, and extrapolate these mappings to patents from 2020 to 2021.[22] Based on these mappings, we match 2,319,184 out of 3,642,855 patents from 2008 to 2021 to Compustat firms.

We then classify each patent as "green" or "non-green" based on a modified list of environment-related CPC or IPC codes provided by Haščič and Migotto (2015).[23] Compared with the original list of Haščič and Migotto (2015), we add three more CPC subclasses related to climate change mitigation.[24] Among the 2,319,184 patents matched to Compustat firms from 2008 to 2021, 210,716 (9.1%) are classified as green based on our criteria.

C1.3 Portfolio Holdings, Asset Prices, and Characteristics

The FactSet database provides an "Entity Sub-type" for each reporting institution. We first follow Koijen et al. (2022) to classsify the type codes into five categories:

- Investment Advisors: IA, MF, IC
- Long-term Investors: FO, IN, PF, SV
- Hedge Funds: HF, FF, FH, FS
- Private Banking: PB, FY, CP, VC
- Brokers: BM, BR

Because "investment advisors" constitute 80% of total AUM of all institutional investors, we break down this category further based on AUM and active share. We classify the largest 25 investment

[21] `https://patentsview.org/download/data-download-tables`

[22] Because the WRDS data covers fewer patents than the Autor et al. (2020) data between 2011 and 2012, we also extrapolate the patent assignee to Compustat mapping from Autor et al. (2020) to patents from 2013 to 2019 that are not covered by the WRDS data.

[23] Cohen et al. (2020) uses the exact list provided by Haščič and Migotto (2015) to define "green" patents.

[24] The CPC subclasses added are: Y02D (climate change mitigation technologies in information and communication), Y02P (climate change mitigation technologies in the production or processing of goods), and Y02W (climate change mitigation technologies related to wastewater treatment or waste management).

advisors (who constitute 50% of total AUM of all investment advisors) as "large". We then divide all other investment advisors into 4 groups of equal total AUM (medium-passive, medium-active, small-passive, small-active) based on AUM and active share.

We also aggregate all non-institutional holdings of stocks in our universe into a household sector. In cases where total institutional holdings exceeds shares outstanding due to missing short positions or misreporting, we scale back the holdings of each institution to ensure total institutional holdings equals shares outstanding. The demand estimation procedure also requires the investment universe for each institution-quarter. Following Koijen and Yogo (2019), we define the investment universe for each institution as all stocks the institution has held in the previous three years.

The data on stock prices, dividends, returns, and shares outstanding are from the Center for Research in Security Prices (CRSP) Monthly Stock Database. We restrict our sample to ordinary common shares (i.e., share codes 10, 11, 12) that trade on NYSE, AMEX, and Nasdaq (i.e., exchange codes 1, 2, and 3). We further restrict our sample to stocks with non-missing price and shares outstanding. Accounting data are from the Compustat North America Fundamentals Annual and Quarterly Databases.

To mitigate the impact of missing data and make sure our results are not driven by micro-caps, we only use the largest stocks that collectively constitute 99% of total market cap in each quarter for analysis. There are 64,184 stock-quarter observations from 2013Q1 to 2021Q3 in our sample, averaging to 1,834 stocks per quarter.

C2 Additional Empirical Results

C2.1 Relationship among Measures of Firm Sustainability

Table 3.C5.1 summarizes the relationship between our three measures of firm sustainability. In Panel (a), we examine the relationship between green patents and other dimensions of sustainability based on cross-sectional regressions. Column (1) of the panel shows that in the cross-section, firms with higher emissions intensity and lower environmental score tend to have more green patents, after controlling for non-green patents. In the cross-section, 1 standard deviation higher environmental score (emissions intensity) is associated with 0.08 standard deviation lower (0.04 standard deviation lower) green innovation. Column (2) shows that the negative relationship between environmental score and green innovation is entirely across-industry, while the positive re-

lationship between emission and green innovation still holds within-industry.

Panel (b) examines whether our three green characteristics could predict future change of emissions intensity in the cross-section. In column (1) of the panel, we regress future 1-year change in emissions intensity on all stock characteristics, controlling for year-quarter fixed effects. The regression show that 1 standard deviation higher environmental score is associated with 4.14% lower emissions intensity over the next year, with t-statistic of 2.5 based on clustered standard errors. Columns (2) and (3) of the panel also show that this predictive relationship is robust to adding industry fixed effects or using industry-time fixed effects. These result show that the environmental score has real information content as it has predictive power for future emissions reduction.[25]

C2.2 Robustness to Alternative Environmental Scores

In the paper, we construct environmental score as $g_t(n) = \frac{-\big(10-E_t(n)\big)w_t^E(n)}{100}$, where $E_t(n)$ is the MSCI environmental pillar score and $w_t^E(n)$ is the MSCI environmental pillar weight; we then residualize $g_t(n)$ against emission intensity in each quarter to arrive at the environmental score we use for valuation regressions and demand estimations. $g_t(n)$ will be high for stocks that have superior environmental performance compared to their peers (high $E_t(n)$), or for stocks in industries where environmental issues are not important (low $w_t^E(n)$).

In this appendix, we show that our results are robust to two alternative constructions for $g_t(n)$:

- Alternative Environment Score (1): $g_t^{(1)}(n) = E_t(n)$: use MSCI environmental pillar score directly without adjusting for environmental pillar weights. $g_t^{(1)}(n)$ will only be high for stocks with good peer-adjusted environmental performance.
- Alternative Environment Score (2): $g_t^{(2)}(n) = \frac{\big(E(n)-5\big)w_t^E(n)}{100}$. In this construction, we use the environmental pillar weight $w_t^E(n)$ (which ranges from 0 to 100) to scale the peer-adjusted environmental performance $E_t(n) - 5$ (which ranges from -5 to 5). $g_t^{(2)}(n)$ will be high for stocks in environmentally-important industries that have better environmental performance relative to its peers.

Importantly, compared to $g_t(n)$, neither the two alternative scores will be high for stocks in environmentally-unimportant industries such as banking.

[25] Igan et al. (2021) find that higher ESG scores do not predict larger decrease in carbon emissions. It should be noted that the independent variable in their work is the firm-level ESG score from Refinitiv (Thomson Reuters), while it is the residualized environment score in our case.

Valuation Regression Table 3.C5.8 shows the valuation regression results based on the alternative scores defined above. The alternative scores $g_t^{(1)}(n)$ and $g_t^{(2)}(n)$ are both residualized against log scope-1 emission intensity in each quarter to maintain consistency with our main specification in the paper. Table 3.C5.8 first shows that both alternative scores are consistently positively priced in the cross-section of valuations. Columns (2-3) also show that the MSCI environmental pillar score has a significantly higher valuation coefficient in the 2018-2021 subsample compared to the 2013-2017 subsample. Comparing Table 3.C5.8 with the valuation regression results in the main paper (Table 3.2), we find that the coefficients for emission intensity and green patents are almost the same.

Demand Coefficients Figure 3.C5.3 shows the AUM-weighted average demand coefficient for environmental score by broad investor types under different environmental score definitions. The figure confirms that there is consistently positive demand for stocks with high third-party environmental score throughout our sample period.

Figure 3.C5.4 shows the demand coefficients for emission intensity under different environmental score definitions. The figure shows that the demand coefficients for emission intensity are not affected by how we construct the environmental scores.

C2.3 Counterfactual: Within-Investor Demand Shift or Across-Investor AUM Shift

In this counterfactual analysis, we undo the within-investor changes in demand for sustainability as well as the changes in AUM to quantify the relative importance of within-investor demand shift versus across-investor AUM shift. We find that the increasing valuation difference between low- and high-emission stocks is almost entirely driven by within-investor demand shifts. This result highlights that the growth of sustainable investing is primarily driven by institutional investors rebalancing their portfolios from high-emission to low-emission stocks.

In the first counterfactual simulation, we shut off within-investor change of all three green demand coefficients as in Section 3.4.2 by setting each investor's green demand coefficients to their values in the first quarter when the investor appears in our sample. In the second counterfactual simulation, we first shut off within-investor demand change in the same way as above, and then further shut off AUM shift across investors by reallocating AUM across investors based on their AUM in 2013Q1 (we set counterfactual AUM to zero for investors not in our data in 2013Q1). We at-

tribute the difference between actual data and the first counterfactual simulation to within-investor shift of green demand, and we attribute the difference between the first and second counterfactual simulation to the shift of AUM between green and brown investors.

Figure 3.C5.5 plots the quarterly valuation regression coefficients for the emission characteristic, based on the data or the two counterfactual simulations. In the data, the valuation gap per 1 standard deviation higher emissions intensity is 4.65% in 2013Q1, which flips in sign to -3.60% in 2021Q3 and subsequently reaches as low as -9.77% in 2020Q1 with t-statistic greater than 5. If we shut off within-investor demand shift (the line labelled "Simulation 1" in the figure), this valuation gap is still a positive 2.83% in 2021Q3, and is never statistically significantly negative in any quarter. Columns (1) and (2) of Table 3.C5.9 further supports this results based on a pooled valuation regression on the 2018-2021 data, showing that the valuation gap between low- and high-emission stocks mostly disappear if there had been no within-investor demand change since 2013Q1. Moreover, the comparison between two counterfactual simulations show that the shift of AUM across investors has little impact on the green-brown valuation gap.

In sum, these results reinforce our finding that the increasing valuation difference between low-emission and high-emission stocks is almost entirely driven by a preference shift towards low-emission stocks within institutional investors.

C3 Theory

C3.1 Including Sustainability as a Characteristic

In this section, we show that sustainability enters the investor's characteristic-based demand if either it is informative about expected returns or investors face a minimum sustainability constraint.

If sustainability is informative about the expected returns, it immediately follows from the same line of argument as in Koijen and Yogo (2019) that it should enter the characteristics-based demand. Suppose on the other hand that sustainability is not informative about the expected returns, but investors face a minimum sustainability constraint instead, similar to Pástor et al. (2021). More concretely, suppose for some $c > 0$ investor i faces, on top of short-sale constraints, an extra con-

straint[26]

$$\mathbf{b}_{it}'\mathbf{w}_{it} = (d_i\mathbf{g}_t)'\mathbf{w}_{it} > c \tag{3.11}$$

where $\mathbf{b}_{it}$ is an $N \times 1$ vector of non-pecuniary benefits which is a product of d_i, investor i's ESG sensitivity, and $\mathbf{g}_t$, the vector of firms' sustainability. Let $\nu_{it} \geq 0$ be the Lagrange multiplier associated with this new constraint. Also, let us denote the kth elementary vector by $\mathbf{e}_k$. Then we have the following result:

Proposition 4. *If an investor faces a sustainability constraint, the optimal portfolio weight on asset n for which the short-sale constraint is not binding is*

$$\mathbf{w}_{it}(n) = \mathbf{y}_{it}(n)'\Pi_{it} + \pi_{it},$$

where

$$\Pi_{it} = \frac{1}{\gamma_{it}}(\tilde{\Phi}_{it} - \Psi_{it}\tilde{\kappa}_{it}), \quad \pi_{it} = \frac{1}{\gamma_{it}}\left(\phi_{it} - \lambda_{it} - \psi_{it}\tilde{\kappa}_{it}\right)$$

are constant across assets. The modified factor loading is given by

$$\tilde{\Phi}_{it} = \Phi_{it} + \nu_{it}d_i\mathbf{e}_k,$$

the modified constant is given by

$$\tilde{\kappa}_{it} = \frac{\Gamma_{it}^{(1)'}(\tilde{\mu}_{it}^{(1)} - \lambda_{it}\mathbf{1})}{\Gamma_{it}^{(1)'}\Gamma_{it}^{(1)} + \gamma_{it}},$$

and $\tilde{\mu}_{it}$ is the expected returns adjusted for the shadow benefits of sustainability

$$\tilde{\mu}_{it} = \mu_{it} + \nu_{it}\mathbf{b}_{it}.$$

Proposition 4 is identical to Proposition 1 in Koijen and Yogo (2019) but with a slight modification to the constant terms to account for the shadow benefit of sustainability, $\nu_{it}\mathbf{b}_{it}$. This addition comes from the fact that green assets are valuable beyond their expected returns because they relax the sustainability constraint. Even with the new constraint, the key content remains: variation in characteristics $\mathbf{y}_{it}(n)$ is the only source of variation in the portfolio weights. Furthermore, the

[26] The current formulation implicitly assumes that green stocks counteract the effects of brown ones. This simplifies the argument, and we motivate it by referring to Morningstar's ESG rating methodology which rates each fund using the weighted average of the fund's Sustainalytics scores. In order to incorporate negative screening against a group of stocks, the sensitivity d_i can be changed to a vector $\mathbf{d}_i$ with a very large $\mathbf{d}_i(n)$ value if stock n is screened. You can find Morningstar's description by following this link.

expression for $\tilde{\Phi}_{it}$ reveals that even if investors do not believe sustainability is informative about expected returns (the factor loading on sustainability is zero in Φ_{it}), the optimal portfolio weights will still be positively related to sustainability.

C3.2 Derivation of Investor Pressure in Proposition 3

To compute **M**, recall the following identity that holds by market clearing:

$$\mathbf{p} = \log\left(\sum_i A_i \mathbf{w}_i\right) - \mathbf{s} \tag{3.12}$$

Differentiating both sides by $\mathbf{p}$:

$$\begin{aligned}
\mathrm{I} &= \begin{pmatrix} \left(\frac{1}{\sum_i A_i w_i(1)}\right)\left(\frac{\partial}{\partial \mathbf{p}(1)}\sum_i A_i w_i(1)\right) & \cdots & \left(\frac{1}{\sum_i A_i w_i(1)}\right)\left(\frac{\partial}{\partial \mathbf{p}(n)}\sum_i A_i w_i(1)\right) \\ \\ \left(\frac{1}{\sum_i A_i w_i(n)}\right)\left(\frac{\partial}{\partial \mathbf{p}(1)}\sum_i A_i w_i(n)\right) & \cdots & \left(\frac{1}{\sum_i A_i w_i(n)}\right)\left(\frac{\partial}{\partial \mathbf{p}(n)}\sum_i A_i w_i(n)\right) \end{pmatrix} \\
&= \begin{pmatrix} \frac{1}{\sum_i A_i w_i(1)} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{1}{\sum_i A_i w_i(n)} \end{pmatrix} \begin{pmatrix} \frac{\partial\left(\sum_i A_i w_i(1)\right)}{\partial \mathbf{p}(1)} & \cdots & \frac{\partial\left(\sum_i A_i w_i(1)\right)}{\partial \mathbf{p}(n)} \\ \vdots & & \vdots \\ \frac{\partial\left(\sum_i A_i w_i(n)\right)}{\partial \mathbf{p}(1)} & \cdots & \frac{\partial\left(\sum_i A_i w_i(n)\right)}{\partial \mathbf{p}(n)} \end{pmatrix} \\
&\equiv \mathbf{H}^{-1}\frac{\partial}{\partial \mathbf{p}}\left(\sum_i A_i \mathbf{w}_i\right)
\end{aligned} \tag{3.13}$$

where

$$\mathbf{H} := \mathrm{diag}\left(\sum_i A_i \mathbf{w}_i\right) = \sum_i A_i \mathrm{diag}\left(\mathbf{w}_i\right) \tag{3.14}$$

Furthermore, it can be shown that:

$$\frac{\partial w_i(n)}{\partial p(n)} = \beta_{0i} w_i(n)(1 - w_i(n)), \quad \frac{\partial w_i(n)}{\partial p(m)} = -\beta_{0i} w_i(n) w_i(m)$$

$$w_i(n) \equiv \frac{\delta_i(n)}{1 + \sum_\ell \delta_i(\ell)}$$

which can be rewritten as

$$\frac{\partial \mathbf{w}_i}{\partial \mathbf{p}} = \beta_{0i}\mathbf{G}_i, \quad \mathbf{G}_i = \mathrm{diag}\left(\mathbf{w}_i\right) - \mathbf{w}_i\mathbf{w}_i'$$

Through analogous steps, it can be shown that the derivative with respect to the kth characteristic is

$$\frac{\partial \mathbf{w}_i}{\partial \mathbf{x}_k} = \beta_i \mathbf{G}_i$$

Now going back to the market clearing condition (3.12) and differentiating both sides by $\mathbf{x}_k$:

$$\mathbf{M} := \frac{\partial \mathbf{p}}{\partial \mathbf{x}_k} = \mathbf{H}^{-1} \left(\sum_i \beta_{0i} A_i \mathbf{G}_i \right) \mathbf{M} + \mathbf{H}^{-1} \left(\sum_i \beta_{ki} A_i \mathbf{G}_i \right)$$

Rearranging yields the desired expression:

$$\mathbf{M} = \left(\mathbf{I} - \sum_i \beta_{0i} A_i \mathbf{H}^{-1} \mathbf{G}_i \right)^{-1} \left(\sum_i \beta_{ki} A_i \mathbf{H}^{-1} \mathbf{G}_i \right)$$

C4 Computing Counterfactuals

We summarize the steps for computing counterfactual which closely follows the algorithm in Koijen and Yogo (2019). For each quarter t, given AUM $\{A_{i,t}\}$ and demand coefficients $\{\boldsymbol{\beta}_{i,t}\}$ for all investors i, non-price stock characteristics $\{\boldsymbol{x}_t(n), g_t(n)\}$ for all stocks n, and latent demand $\{\varepsilon_{i,t}(n)\}$, we compute the equilibrium log market capitalization $\widetilde{me}_t(n)$ as follows:

1. Start from an initial guess $\widetilde{me}_t(n;0)$.
2. Plug $\widetilde{me}_t(n;0)$ and $\{\boldsymbol{\beta}_{i,t}, \boldsymbol{x}_t(n), g_t(n), \varepsilon_{i,t}(n)\}$ into Equation 3.5 for quarter t to compute portfolio weights $w_{i,t}(n)$.
3. Use the market-clearing condition to compute the prices corresponding to portfolio weights: $\widetilde{me}^*_t(n) = \log \sum_i A_{i,t} w_{i,t}(n) - \log BE_t(n)$.
4. Update the guess for prices: $\widetilde{me}_t(n;1) = \widetilde{me}_t(n;0) + k_t(n) \cdot \left[\widetilde{me}^*_t(n) - \widetilde{me}_t(n;0) \right]$.
   (a) $k_t(n)$ is the stock- and time-specific speed of update. Koijen and Yogo (2019) uses the optimal $k^*_t(n) = \left[1 - \frac{\sum_i \beta_{0,i,t} A_{i,t} w_{i,t}(n)[1 - w_{i,t}(n)]}{\sum_i A_{i,t} w_{i,t}(n)} \right]^{-1}$ derived based on Newton's Method.
   (b) To improve numerical stability, we use a smaller update speed $k_t(n) = 0.25 \cdot k^*_t(n)$. If the procedure fails to converge, we try $k_t(n) = 0.1 \cdot k^*_t(n)$ and $k_t(n) = 0.05 \cdot k^*_t(n)$.
   (c) If the procedure still fails to converge, we then try fixed update speeds $k_t(n)$= 0.8 / 0.5 / 0.2.
5. Iterate back to step (1) until the price vector converges: $\min |\widetilde{me}_t(n;1) - \widetilde{me}_t(n;0)| < 10^{-4}$.

For each set of primary inputs, we can compute the counterfactual portfolio holdings and investor

pressures after computing the prices.

C5 Additional Tables and Figures

Table 3.C5.1: Relationship Between Green Characteristics

(a) Green Patent and Other Green Characteristics

	Dep Var: Green Patents	
	(1)	(2)
Non-Green Patent / Asset	0.545**	0.476**
	[16.53]	[14.49]
E-Score Residual	−0.0782**	−0.00300
	[−5.653]	[−0.200]
Log GHG1 Intensity	0.0439**	0.0444*
	[4.570]	[2.282]
Time FE	✓	✓
FF-48 Industry FE		✓
Within R^2		
Observations	64,184	64,184

(b) Green Characteristics and Future Change in Emissions Intensity

	Dep Var: Future 4Q Change in Emissions Intensity		
	(1)	(2)	(3)
E-Score Residual	−0.0414*	−0.0386*	−0.0436*
	[−2.511]	[−2.603]	[−2.670]
Log GHG1 Intensity	−0.00867	−0.0686*	−0.0268*
	[−1.667]	[−2.503]	[−2.173]
Green Patent / Asset	0.00239	0.00187	0.00289
	[0.782]	[0.537]	[0.903]
Non-Green Controls	✓	✓	✓
Time FE	✓	✓	
FF-48 Industry FE		✓	
FF-12 Industry × Time FE			✓
Within R^2	.021	.032	.022
Observations	51,065	51,065	51,065

This table summarizes the relationships between the different measures of sustainability. Panel (a) examines the relationship between the measure of green innovation and other characteristics in the cross-section. Panel (b) examines the predictability of each characteristic for future one-year ahead change in emissions intensity in the cross-section.

Table 3.C5.2: Sustainable Demand and UNPRI Signatory Status: Cross-Sectional

	E-Score Resid		Emission		Green Patent	
	(1)	(2)	(3)	(4)	(5)	(6)
1{UNPRI Signatory}	0.0578*	−0.0000795	−0.0252	−0.0225	−0.0275	−0.0108
	[2.268]	[−0.00309]	[−0.801]	[−0.749]	[−1.058]	[−0.414]
Time FE	✓	✓	✓	✓	✓	✓
Investor Controls		✓		✓		✓
Within R^2	0	.013	0	.041	0	.013
Observations	113,231	113,231	113,231	113,231	113,231	113,231

This table summarizes the relationship between green demand and investor characteristics via cross-sectional regressions. Data is at the investor-quarter level, and all continuous variables are cross-sectionally standardized. Controls include all investor characteristics in Table 3.5.

Table 3.C5.3: Sustainable Demand and UNPRI Signatory Status: Time-Series

	Demand Coefficients		
	(1) E-Score Resid	(2) Emission	(3) Green Patent
1{UNPRI Signatory}	−0.00135	−0.00916	−0.00391
	[−0.123]	[−0.945]	[−0.244]
Time Trend	−0.000883**	−0.00144**	−0.00157**
	[−3.052]	[−4.682]	[−4.286]
Investor FE	✓	✓	✓
Within R^2	.001	.004	.002
Observations	112,942	112,942	112,942

This table summarizes the within-investor change in demand for sustainability after an investor becomes a UNPRI signatory via time-series regressions.

Table 3.C5.4: Investor Pressure: Summary Statistics

	Mean	SD	Q10	Median	Q90
E-Score Residual	0.063	0.084	−0.033	0.056	0.162
Emission	−0.084	0.085	−0.192	−0.071	0.004
Green Patent	−0.008	0.092	−0.118	0.008	0.082

This table summarizes the investor pressure. Statistics are computed in each quarter, and then averaged across quarters.

Table 3.C5.5: Investor Pressure and Future Environmental Performance: Longer Horizons

	Forward 2-Year Outcome			Forward 3-Year Outcome		
	(1) E-Score Resid	(2) GHG	(3) Green Patent	(4) E-Score Resid	(5) GHG	(6) Green Patent
Pressure: E-Score Residual	0.0236^{+}	-0.0145	-0.0137^{+}	0.0131	-0.0453^{**}	-0.0155
	[1.923]	[−1.074]	[−1.792]	[0.929]	[−3.245]	[−1.689]
Pressure: GHG	-0.0437^{**}	0.0339^{**}	0.00289	-0.0629^{**}	0.0305^{*}	0.00243
	[−3.566]	[2.788]	[0.420]	[−4.176]	[2.164]	[0.321]
Pressure: Green Patent	−0.000842	0.0113	0.0152^{*}	0.00769	0.0109	0.0200^{**}
	[−0.0653]	[1.064]	[2.680]	[0.490]	[0.855]	[2.983]
E-Score Residual	-0.300^{**}	-0.123^{**}	-0.0138^{+}	-0.349^{**}	-0.178^{**}	-0.0174^{+}
	[−20.03]	[−4.839]	[−1.916]	[−19.25]	[−6.342]	[−2.007]
Log GHG1 Intensity	-0.0414^{*}	-0.0429^{**}	0.00233	-0.0438^{*}	-0.0530^{**}	0.00291
	[−2.620]	[−2.816]	[0.364]	[−2.332]	[−3.067]	[0.380]
Green Patent / Asset	0.00293	0.00564	0.784^{**}	−0.000242	0.00360	0.781^{**}
	[0.216]	[0.469]	[29.65]	[−0.0144]	[0.220]	[24.48]
Time FE	✓	✓	✓	✓	✓	✓
Non-Green Controls	✓	✓	✓	✓	✓	✓
Within R^2	.094	.029	.71	.134	.067	.722
Observations	41,486	41,486	40,754	33,182	33,182	32,099

This table summarizes the cross-sectional relationship between investor pressure and future environmental performance. The dependent variables are: future 2/3-year change in environmental score, future 2/3-year change in emissions intensity, and the future 2/3-year number of green patents divided by total assets. The main independent variables are the investor pressure for three green characteristics. "Non-Green" control variables include log book equity, investment, profitability, market beta, and dividend to book equity. All outcome variables and regressors are cross-sectionally standardized in each quarter.

Table 3.C5.6: Counterfactual Exercise: Impact of ESG-Agnostic Mandates, Alternative Order

(a) Counterfactual Valuation Regressions: Full Sample (2013 - 2021)

	Data	CF: Shut off Green Demand		
	(1)	(2) Passive Inst	(3) All Inst	(4) All Inst + HH
E-Score Residual	0.107**	0.107**	0.0951**	0.0540**
	[8.191]	[7.927]	[7.194]	[4.103]
Log GHG1 Intensity	−0.0156	0.00697	0.0330*	0.0387**
	[−1.153]	[0.552]	[2.634]	[3.065]
Green Patent / Asset	0.00359	−0.00211	0.0249	0.0161
	[0.237]	[−0.136]	[1.518]	[0.993]
Time FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	64,184	64,184	64,184	64,184

(b) Counterfactual Valuation Regressions: Subsample (2018 - 2021)

	Data	CF: Shut off Green Demand		
	(1)	(2) Passive Inst	(3) All Inst	(4) All Inst + HH
E-Score Residual	0.115**	0.123**	0.114**	0.0698**
	[6.767]	[7.005]	[6.529]	[4.010]
Log GHG1 Intensity	−0.0645**	−0.0275+	0.00311	0.00803
	[−4.261]	[−1.789]	[0.199]	[0.506]
Green Patent / Asset	−0.0115	−0.0209	0.0150	0.00618
	[−0.603]	[−1.107]	[0.784]	[0.327]
Time FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	26,251	26,251	26,251	26,251

This table presents the results from valuation regressions in counterfactual scenarios where we "shut off" green demand, i.e., set the demand coefficients for all three sustainability characteristics to zero for one or more types of investors. Importantly, we shut off the demand for passive institutions before active institutions, which is in contrast to the baseline procedure as described in Section 3.6.1. Specifically, we consider three scenarios in which we "shut off" green demand for all passive institutions (column (2)), all institutions (column (3)), and all investors including the household sector (column (4)). After obtaining counterfactual valuations, we re-estimate the valuation regression as shown in Equation (3.1). We present results for both the full sample as well as the 2018-2021 subsample.

Table 3.C5.7: Stock Characteristics and Investment Universe Coverage

	Investment Universe Coverage	
	(1) Equal-Weighted	(2) AUM-Weighted
E-Score Residual	0.00971**	0.00331
	[6.164]	[1.687]
Log GHG1 Intensity	−0.000324	−0.00155
	[−0.187]	[−0.788]
Green Patent / Asset	0.000919	−0.000244
	[0.522]	[−0.136]
Log Book Equity	0.109**	0.0794**
	[38.18]	[33.82]
Profit / Asset	0.0258**	0.0263**
	[16.28]	[17.32]
Asset Growth	0.00556**	0.00499**
	[5.262]	[5.440]
Dividend / Book Equity	0.0319**	0.0146**
	[12.99]	[6.570]
CAPM Beta	0.00131	0.00212
	[1.001]	[1.204]
Non-Green Patent / Asset	0.0150**	0.0108**
	[8.604]	[6.097]
Time FE	✓	✓
Dep. Var. Mean	.196	.679
Within R^2	.659	.412
Observations	64,184	64,184

This table summarizes the results from regressing investment universe coverage on stock characteristics. For column (1), the equal-weighted investment universe coverage for a stock is the fraction of investors by number that have the stock in their investment universes. For column (2), the AUM-weighted investment universe coverage for a stock is the fraction of investors by AUM that have the stock in their investment universes.

Table 3.C5.8: Valuation Regression with Alternative Environment Scores

	Dependent Variable: Log Market-to-Book					
	(1) Full Sample	(2) Full Sample	(3) 2013-2017	(4) 2013-2017	(5) 2018-2021	(6) 2018-2021
Alternative E-Score (1)	0.170** [11.94]	0.141** [10.79]	0.212** [11.43]			
Alternative E-Score (2)				0.124** [9.308]	0.114** [8.236]	0.139** [7.452]
Emission Intensity	−0.0133 [−1.005]	0.0159 [1.279]	−0.0646** [−4.458]	−0.0138 [−1.016]	0.0152 [1.202]	−0.0641** [−4.303]
Green Patent	−0.00608 [−0.420]	0.0000154 [0.000947]	−0.0156 [−0.851]	−0.0113 [−0.782]	−0.00230 [−0.142]	−0.0262 [−1.428]
Year-Quarter FE	✓	✓	✓	✓	✓	✓
Within R^2	.411	.39	.452	.397	.382	.431
Observations	64,184	37,933	26,251	64,184	37,933	26,251

This table presents the valuation regression results based on two alternative measures of environmental scores defined in Appendix C2.2. Both non-sustainable and sustainable characteristics are included in the regression, but the coefficients for sustainable characteristics are presented to save space.

Table 3.C5.9: Counterfactual Exercise: Demand vs. AUM Shifts

	Data	CF Simulations	
	(1)	(2) Simulation 1	(3) Simulation 2
Environment Score	0.115**	0.119**	0.118**
	[6.767]	[6.725]	[7.164]
Emissions Intensity	−0.0645**	−0.0124	0.00715
	[−4.261]	[−0.793]	[0.476]
Green Patents	−0.0115	0.0706**	0.0607**
	[−0.603]	[3.584]	[3.059]
Time FE	✓	✓	✓
Controls	✓	✓	✓
Observations	26,251	26,251	26,251

This table presents the results from valuation regressions in counterfactual described in Appendix C2.3, which is designed to juxtapose the relative importance of demand shifts *within* investors against the AUM shift *across* investors. In the first scenario, we shut off the within-investor change of all green demand coefficients by setting each investor's green demand coefficients to their values in the investor's first quarter. In the second scenario, we first shut off the within-investor demand change in the same way as above and then further shut off the AUM shift across investors by reallocating AUM across investors based on their AUM in 2013Q1. After obtaining counterfactual valuations, we re-estimate the valuation regression as shown in Equation (3.1).

Figure 3.C5.1: Average Price Inelasticity Over Time

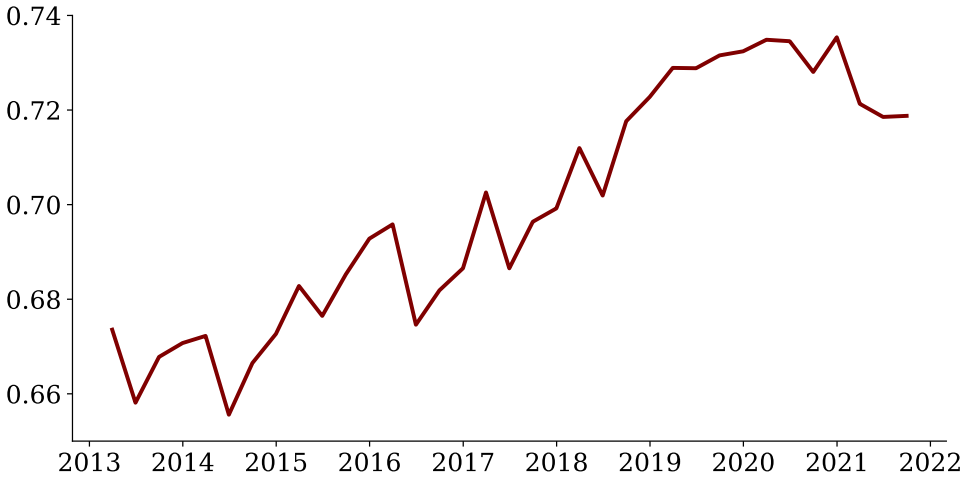

This figure summarizes the time trends in the AUM-weighted average price inelasticity coefficient ($\beta_{0,i,t}$) across all investors.

Figure 3.C5.2: Distribution of Environmental Scores for Oil & Gas and Banking Stocks

(a) MSCI Environmental Pillar Score: $E_t(n)$

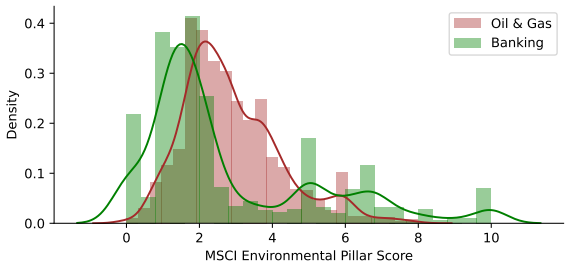

(b) MSCI Environmental Pillar Weight: $w_t^E(n)$

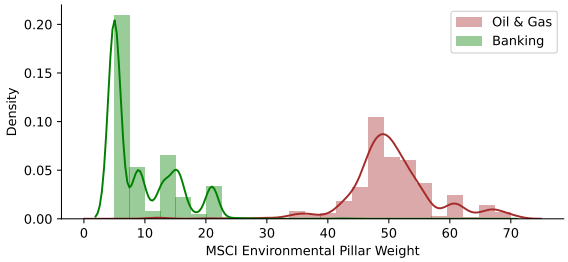

(c) Raw Environmental Score Following Pástor et al. (2022): $g_t(n)$

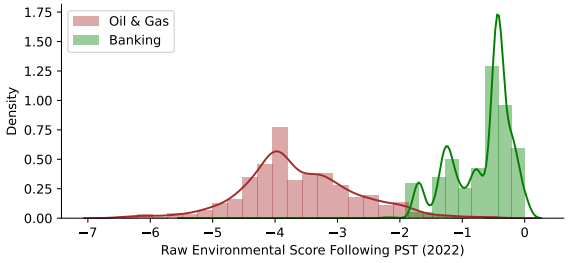

This figure illustrates the distribution of MSCI environmental pillar scores, MSCI environmental pillar weights, and the raw environmental scores across stock-quarters from oil & gas and banking industries. We use the 48-industry classifications from Ken French's website to select observations for each industry.

Figure 3.C5.3: Demand Coefficients for Environmental Score: Alternative Environment Score Constructions

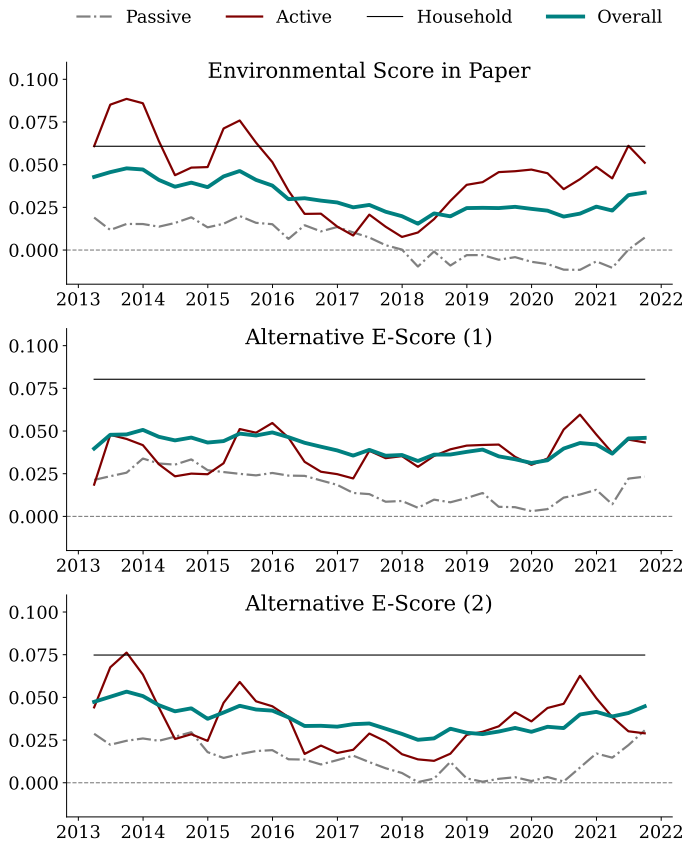

This figure plots the average demand coefficient for environmental score by broad investor types (passive institutions, active institutions, and the household sector) based on different definitions of environmental score. The alternative scores are defined in Appendix C2.2. The investor classifications are the same as in the main paper.

Figure 3.C5.4: Demand Coefficients for Emission Intensity: Alternative Environment Score Constructions

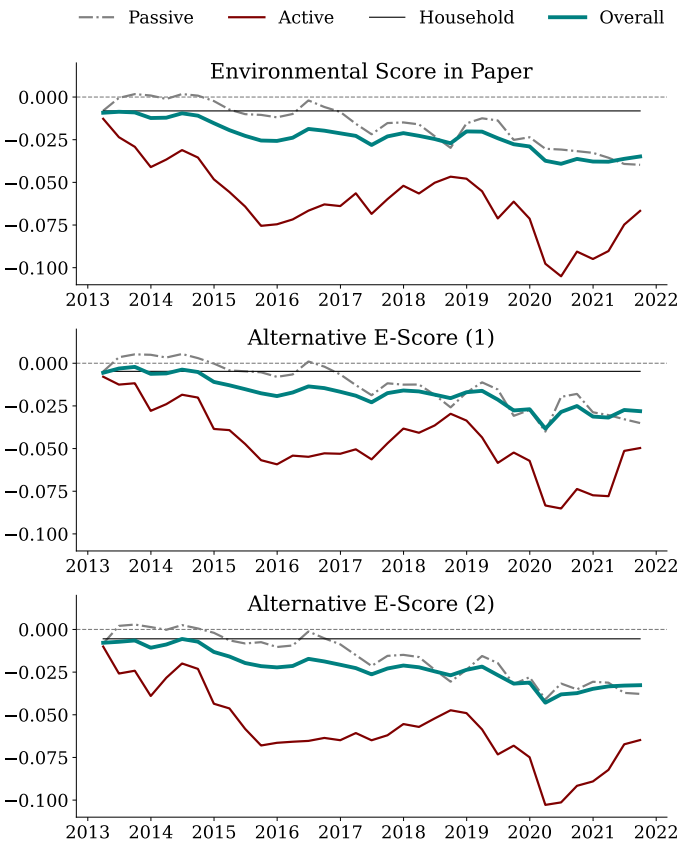

This figure plots the average demand coefficient for emission intensity by broad investor types (passive institutions, active institutions, and the household sector) based on different definitions of environmental score. The alternative scores are defined in Appendix C2.2. The investor classifications are the same as in the main paper.

Figure 3.C5.5: Counterfactual: Demand vs. AUM Shifts

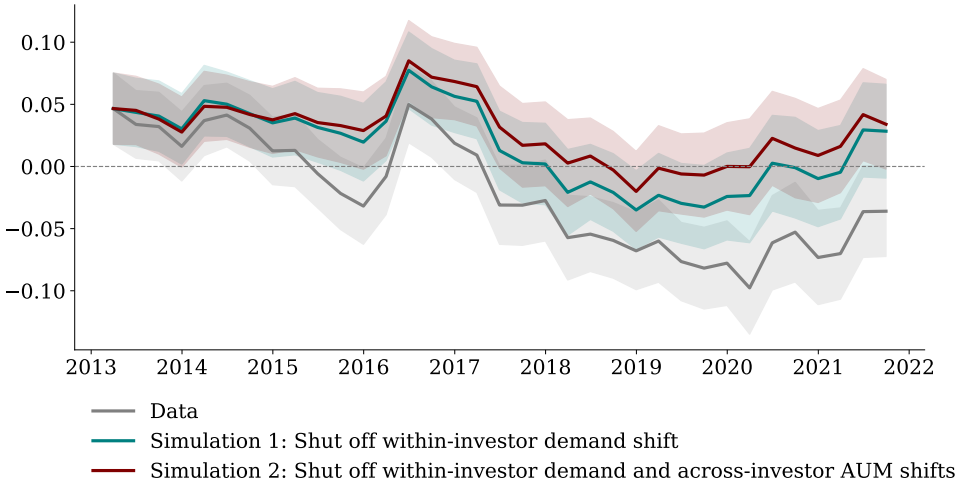

This figure plots the time-series of the quarterly valuation regression coefficients on emissions intensity in the data in the counterfactual described in Appendix C2.3, which is designed to juxtapose the relative importance of demand shifts *within* investors against the AUM shift *across* investors.

Bibliography

Aboody, D., Even-Tov, O., Lehavy, R., and Trueman, B. (2018). Overnight Returns and Firm-Specific Investor Sentiment. *Journal of Financial and Quantitative Analysis*, 53(2):485–505.

Akbas, F., Boehmer, E., Jiang, C., and Koch, P. D. (2022). Overnight Returns, Daytime Reversals, and Future Stock Returns. *Journal of Financial Economics*, 145(3):850–875.

Amihud, Y. (2002). Illiquidity and Stock Returns: Cross-Section and Time-Series Effects. *Journal of Financial Markets*, 5(1):31–56.

Autor, D., Dorn, D., Hanson, G. H., Pisano, G., and Shu, P. (2020). Foreign Competition and Domestic Innovation: Evidence from US Patents. *American Economic Review: Insights*, 2(3):357–74.

Baker, A. C., Larcker, D. F., and Wang, C. C. Y. (2022). How Much Should We Trust Staggered Difference-in-Differences Estimates? *Journal of Financial Economics*, 144(2):370–395.

Bali, T. G., Cakici, N., and Whitelaw, R. F. (2011). Maxing Out: Stocks as Lotteries and the Cross-Section of Expected Returns. *Journal of Financial Economics*, 99(2):427–446.

Bansal, R., Di, W., and Yaron, A. (2018). Is Socially Responsible Investing a Luxury Good? *Unpublished Working Paper*.

Barber, B. M., Huang, X., Jorion, P., Odean, T., and Schwarz, C. (2022a). A (Sub)penny For Your Thoughts: Tracking Retail Investor Activity in TAQ. *Unpublished Working Paper*.

Barber, B. M., Huang, X., Odean, T., and Schwarz, C. (2022b). Attention-Induced Trading and Returns: Evidence from Robinhood Users. *The Journal of Finance*, 77(6):3141–3190.

Barber, B. M., Lee, Y.-T., Liu, Y.-J., and Odean, T. (2009). Just How Much Do Individual Investors Lose by Trading? *The Review of Financial Studies*, 22(2):609–632.

Barber, B. M., Lin, S., and Odean, T. (2022c). Resolving a Paradox: Retail Trades Positively Predict Returns but are Not Profitable. *Unpublished Working Paper*.

Barber, B. M., Morse, A., and Yasuda, A. (2021). Impact Investing. *Journal of Financial Economics*, 139(1):162–185.

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

Barrot, J.-N., Kaniel, R., and Sraer, D. (2016). Are Retail Traders Compensated for Providing Liquidity? *Journal of Financial Economics*, 120(1):146–168.

Berg, F., Koelbel, J. F., Pavlova, A., and Rigobon, R. (2022). ESG Confusion and Stock Returns: Tackling the Problem of Noise. *NBER Working Paper*.

Berg, F., Koelbel, J. F., and Rigobon, R. (2019). Aggregate Confusion: The Divergence of ESG ratings. *MIT Sloan Working Paper*.

Berk, J. and van Binsbergen, J. H. (2021). The Impact of Impact Investing. *Unpublished Working Paper*.

Berkman, H., Koch, P., Tuttle, L., and Zhang, Y. (2012). Paying Attention: Overnight Returns and the Hidden Cost of Buying at the Open. *Journal of Financial and Quantitative Analysis*, 47(4):715–741.

Bloomberg (2021). ESG Assets May Hit $53 Trillion by 2025, a Third of Global AUM. *Bloomberg Intelligence*.

Bloomberg (2022). Gop Fury Over ESG Triggers Backlash with Us Pensions at Risk. *Bloomberg*.

Boehmer, E., Jones, C. M., Zhang, X., and Zhang, X. (2021). Tracking Retail Investor Activity. *The Journal of Finance*, 76(5):2249–2305.

Bogousslavsky, V. (2021). The Cross-Section of Intraday and Overnight Returns. *Journal of Financial Economics*, 141(1):172–194.

Bolton, P. and Kacperczyk, M. (2021). Do Investors Care About Carbon Risk? *Journal of Financial Economics*, 142(2):517–549.

Boyarchenko, N., Larsen, L. C., and Whelan, P. (2023). The Overnight Drift. *The Review of Financial Studies, Forthcoming*.

Branch, B. S. and Ma, A. J. (2015). Overnight Return, the Invisible Hand Behind Intraday Returns? *Unpublished Working Paper*.

Bretscher, L., Schmid, L., Sen, I., and Sharma, V. (2020). Institutional Corporate Bond Pricing. *Swiss Finance Institute Research Paper*.

Broccardo, E., Hart, O., and Zingales, L. (2022). Exit Versus Voice. *Journal of Political Economy*, 130(12):3101–3145.

Chen, H., Petukhov, A., and Wang, J. (2019a). The Dark Side of Circuit Breakers. *Unpublished Working Paper*.

Chen, T., Gao, Z., He, J., Jiang, W., and Xiong, W. (2019b). Daily Price Limits and Destructive Market Behavior. *Journal of Econometrics*, 208(1):249–264.

Chen, T. and Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA. Association for Computing Machinery.

Cho, D. D., Russell, J., Tiao, G. C., and Tsay, R. (2003). The Magnet Effect of Price Limits: Evidence from High-Frequency Data on Taiwan Stock Exchange. *Journal of Empirical Finance*, 10(1):133–168.

Choi, D., Gao, Z., Jiang, W., and Zhang, H. (2022). Carbon Stock Devaluation. *Unpublished Working Paper*.

Christie, W. G., Corwin, S. A., and Harris, J. H. (2002). Nasdaq Trading Halts: The Impact of Market Mechanisms on Prices, Trading Activity, and Execution Costs. *The Journal of Finance*, 57(3):1443–1478.

Cohen, L., Gurun, U. G., and Nguyen, Q. H. (2020). The ESG-Innovation Disconnect: Evidence from Green Patenting. *NBER Working Paper*.

Coqueret, G. (2021). *Perspectives in ESG Equity Investing*.

Davis, C., Kargar, M., and Li, J. (2022). An Information-Based Explanation for Inelastic Demand. *Unpublished Working Paper*.

De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy*, 98(4):703–738.

Derrien, F., Krueger, P., Landier, A., and Yao, T. (2021). ESG News, Future Cash Flows, and Firm Value. *Unpublished Working Paper*.

Dou, W. W., Kogan, L., and Wu, W. (2022). Common Fund Flows: Flow Hedging and Factor Pricing. *Unpublished Working Paper*.

Dyck, A., Lins, K. V., Roth, L., and Wagner, H. F. (2019). Do Institutional Investors Drive Corporate Social Responsibility? International Evidence. *Journal of Financial Economics*, 131(3):693–714.

Edmans, A., Levit, D., and Schneemeier, J. (2022). Socially Responsible Divestment. *European Corporate Governance Institute–Finance Working Paper*.

Engelberg, J., Sasseville, C., and Williams, J. (2012). Market Madness? The Case of "Mad Money". *Management Science*, 58(2):351–364.

Fama, E. and French, K. (2004). The Capital Asset Pricing Model: Theory and evidence. *Journal of Economic Perspectives*, 18(3):25–46.

Flammer, C. (2021). Corporate Green Bonds. *Journal of Financial Economics*, 142(2):499–516.

Fletcher, L. and Brower, D. (2021). Hedge Funds Cash in as Green Investors Dump Energy Stocks. *Financial Times*.

Gabaix, X., Koijen, R. S., Mainardi, F., Oh, S., and Yogo, M. (2022). Asset Demand of u.s. Households. *Unpublished Working Paper*.

Gabaix, X. and Koijen, R. S. J. (2021). In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. *Unpublished Working Paper*.

Gantchev, N., Giannetti, M., and Li, R. (2022). Does Money Talk? Market Discipline Through Selloffs and Boycotts. *Review of Finance, Forthcoming*.

Gibson, R., Glossner, S., Krueger, P., Matos, P., Steffen, T., et al. (2020). Responsible Institutional Investing Around the World. *Swiss Finance Institute Working Paper*.

Giglio, S., Kelly, B., and Stroebel, J. (2021). Climate Finance. *Annual Review of Financial Economics*, 13:15–36.

Gilbert, M. (2021). Hedge Funds Are Feasting on ESG's Profit Leftovers. *Bloomberg*.

Glossner, S. (2021). Repeat Offenders: ESG Incident Recidivism and Investor Underreaction. *Unpublished Working Paper*.

Goldstein, I., Kopytov, A., Shen, L., and Xiang, H. (2021). On ESG Investing: Heterogeneous Preferences, Information, and Asset Prices. *Unpublished Working Paper*.

Görgen, M., Jacob, A., Nerlinger, M., Riordan, R., Rohleder, M., and Wilkens, M. (2020). Carbon Risk. *Unpublished Working Paper*.

Gormsen, N. J., Huber, K., and Oh, S. (2023). Climate Capitalists. *University of Chicago Working Paper*.

Haddad, V., Huebner, P., and Loualiche, E. (2021). How Competitive is the Stock Market? Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing. *Unpublished Working Paper*.

Hartzmark, S. M. and Shue, K. (2023). Counterproductive Impact Investing: The Impact Elasticity of Brown and Green Firms. *Unpublished Working Paper*.

Hartzmark, S. M. and Sussman, A. B. (2019). Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows. *The Journal of Finance*, 74(6):2789–2837.

Haščič, I. and Migotto, M. (2015). Measuring Environmental Innovation Using Patent Data. *OECD Working Paper*.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer, New York, NY.

Hautsch, N. and Horvath, A. (2019). How Effective Are Trading Pauses? *Journal of Financial Economics*, 131(2):378–403.

Heath, D., Macciocchi, D., Michaely, R., and Ringgenberg, M. C. (2021). Does Socially Responsible Investing Change Firm Behavior? *European Corporate Governance Institute–Finance Working Paper*, (762).

Heinkel, R., Kraus, A., and Zechner, J. (2001). The Effect of Green Investment on Corporate Behavior. *Journal of Financial and Quantitative Analysis*, 36(4):431–449.

Hendershott, T., Livdan, D., and Rösch, D. (2020). Asset Pricing: A Tale of Night and Day. *Journal of Financial Economics*, 138(3):635–662.

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

Hirshleifer, D. A., Myers, J. N., Myers, L. A., and Teoh, S. H. (2008). Do Individual Investors Cause Post-Earnings Announcement Drift? Direct Evidence from Personal Trades. *The Accounting Review*, 83(6):1521–1550.

Hong, H. and Sraer, D. A. (2016). Speculative Betas. *The Journal of Finance*, 71(5):2095–2144.

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

Hsu, P.-H., Li, K., and Tsou, C.-Y. (2022). The Pollution Premium. *Journal of Finance, Forthcoming*.

Huebner, P. (2022). The Making of Momentum: A Demand-System Perspective. *Unpublished Working Paper*.

Igan, D., Kirti, D., and Elmalt, D. (2021). Limits to Private Climate Change Mitigation. *IMF Working Paper*.

Ilhan, E., Sautner, Z., and Vilkov, G. (2021). Carbon Tail Risk. *The Review of Financial Studies*, 34(3):1540–1571.

Jansen, K. A. (2021). Long-Term Investors, Demand Shifts, and Yields. *Unpublished Working Paper*.

Jiang, Z., Richmond, R. J., and Zhang, T. (2022). A Portfolio Approach to Global Imbalances. *NBER Working Paper*.

Jones, C. S., Pyun, S., and Wang, T. (2022). Return Extrapolation and Day/Night Effects. *Unpublished Working Paper*.

Kelly, M. A. and Clark, S. P. (2011). Returns in Trading Versus Non-Trading Hours: The Difference Is Day and Night. *Journal of Asset Management*, 12(2):132–145.

Kim, J. and Kim, M. (2022). Behavioral Biases and the Trading of Individual Investors in the Korean Stock Markets. Technical report, Korea Capital Market Institute.

Kim, S. and Yoon, A. (2023). Analyzing Active Fund Managers' Commitment to ESG: Evidence from the United Nations Principles for Responsible Investment. *Management Science*, 69(2):741–758.

Koijen, R. S. J., Richmond, R., and Yogo, M. (2022). Which Investors Matter for Equity Valuations and Expected Returns?

Koijen, R. S. J. and Yogo, M. (2019). A Demand System Approach to Asset Pricing. *Journal of Political Economy*, 127(4):1475–1515.

Koijen, R. S. J. and Yogo, M. (2020). Exchange Rates and Asset Prices in a Global Demand System. *Unpublished Working Paper*.

Krueger, P., Sautner, Z., and Starks, L. T. (2020). The Importance of Climate Risks for Institutional Investors. *The Review of Financial Studies*, 33(3):1067–1111.

Laarits, T. and Sammon, M. (2022). The Retail Habitat. *Unpublished Working Paper*.

Lee, C. M. C., Ready, M. J., and Seguin, P. J. (1994). Volume, Volatility, and New York Stock Exchange Trading Halts. *The Journal of Finance*, 49(1):183–214.

Liang, H., Sun, L., and Teo, M. (2021). Greenwashing: Evidence from Hedge Funds. *Singapore Managment University Working Paper*.

Lou, D., Polk, C., and Skouras, S. (2019). A Tug of War: Overnight Versus Intraday Expected Returns. *Journal of Financial Economics*, 134(1):192–213.

Lucas, R. E. (1976). Econometric Policy Evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46.

Luo, C., Ravina, E., Sammon, M., and Viceira, L. M. (2022). Retail Investors' Contrarian Behavior Around News, Attention, and the Momentum Effect. *Unpublished Working Paper*.

Odean, T. (1999). Do Investors Trade Too Much? *American Economic Review*, 89(5):1279–1298.

Ozik, G., Sadka, R., and Shen, S. (2021). Flattening the Illiquidity Curve: Retail Trading During the Covid-19 Lockdown. *Journal of Financial and Quantitative Analysis*, 56(7):2356–2388.

Pástor, Ľ., Stambaugh, R. F., and Taylor, L. A. (2021). Sustainable Investing in Equilibrium. *Journal of Financial Economics*, 142(2):550–571.

Pástor, Ľ., Stambaugh, R. F., and Taylor, L. A. (2022). Dissecting Green Returns. *Journal of Financial Economics*, 146(2):403–424.

Pedersen, L. H., Fitzgibbons, S., and Pomorski, L. (2021). Responsible Investing: The ESG-Efficient Frontier. *Journal of Financial Economics*, 142(2):572–597.

Savor, P. and Wilson, M. (2014). Asset Pricing: A Tale of Two Days. *Journal of Financial Economics*, 113(2):171–201.

Scheinkman, J. A. and Xiong, W. (2003). Overconfidence and Speculative Bubbles. *Journal of Political Economy*, 111(6):1183–1220.

Seasholes, M. S. and Wu, G. (2007). Predictable Behavior, Profits, and Attention. *Journal of Empirical Finance*, 14(5):590–610.

Shefrin, H. and Statman, M. (1985). The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. *The Journal of Finance*, 40(3):777–790.

Sifat, I. M. and Mohamad, A. (2019). Circuit Breakers as Market Stability Levers: A Survey of Research, Praxis, and Challenges. *International Journal of Finance & Economics*, 24(3):1130–1169.

Stock, J. and Yogo, M. (2005). Asymptotic Distributions of Instrumental Variables Statistics with Many Instruments. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 6:109–120.

Subrahmanyam, A. (1994). Circuit Breakers and Market Volatility: A Theoretical Perspective. *The Journal of Finance*, 49(1):237–254.

van der Beck, P. (2021). Flow-Driven ESG Returns. *Unpublished Working Paper*.

van der Beck, P. and Jaunin, C. (2021). The Equity Market Implications of the Retail Investment Boom. *Swiss Finance Institute Research Paper*, (21-12).

Yan Du, D., Liu, Q., and Rhee, S. G. (2009). An Analysis of the Magnet Effect Under Price Limits. *International Review of Finance*, 9(1-2):83–110.

Zerbib, O. D. (2022). A Sustainable Capital Asset Pricing Model (S-CAPM): Evidence from Environmental Integration and Sin Stock Exclusion. *Review of Finance*, 26(6):1345–1388.