

Three Essays in Financial Economics

A Ph.D. thesis submitted to the University of Bern

at the faculty of Business, Economics and Social Sciences,
supervised by Prof. Dr. Philip Valta,
in partial fulfillment of the requirements for the degree of

Doctor Rerum Oeconomicarum

by

Sascha Jakob

February, 2023

Original document stored on the web server of the Universitätsbibliothek Bern.



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

You are free to:

- Share — copy and redistribute the material in any medium or format. The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:

- Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
- NonCommercial — You may not use the material for commercial purposes.
- NoDerivatives — If you remix, transform, or build upon the material, you may not distribute the modified material.
- No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

- You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.
- No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

A detailed version of the license agreement is available at:

- <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>

*Learn,
Discover,
Heal,
Create – and Make the World Ever Better*

— University of Rochester

Acknowledgements

I would like to acknowledge and give my warmest thanks to my supervisor and co-author Prof. Dr. Philip Valta who made this thesis possible. His guidance and advice carried me through all the stages of writing my thesis. I would also like to thank my secondary advisor Prof. Dr. Julien Cujean for his fruitful comments and support.

My gratitude also goes to my co-author Dr. Tim Glaus for sharing the efforts, and to my independent examiner Prof. Dr. Thorsten Hens for taking the time to evaluate this thesis.

Special thanks go to my working colleagues and fellow Ph.D. students at the Institute for Financial Management at the University of Bern - Christian, Jan, Karin, Marc, Philipp, and Samuel - for their helpful and constructive thoughts, witty banters, and moral support.

I thank my family, friends, and colleagues, who supported me in the process of composing this thesis - and all those that I met along this journey but forgot to mention.

...Finally, and from the bottom of my heart: Thank you, Mom and Dad, for what needs no further elaboration...

Sascha Jakob

Preface

In 1970, Nobel laureate Eugene F. Fama proposed his Efficient Market Hypothesis in "Efficient Capital Markets: a Review of Theory and Empirical Work", which has become one of the most central tenets of modern economics about financial markets:

The primary role of the capital market is allocation of ownership of the economy's capital stock. In general terms, the ideal is a market in which prices provide accurate signals for resource allocation: that is, a market in which firms can make production-investment decisions, and investors can choose among the securities that represent ownership of firms' activities under the assumption that security prices at any time "fully reflect" all available information. A market in which prices always "fully reflect" available information is called "efficient."

At the time of its proposition by [Fama \(1970\)](#), there was sufficient evidence that markets are efficient to render the Efficient Market Hypothesis' position as a fundamental pillar of financial economics irrefutable. To this day, Eugene F. Fama is a staunch defender of the Efficient Market Hypothesis and has continuously contributed compelling evidence that markets fully reflect information and are efficient.

However, a decade after Eugene F. Fama shaped one of the most important concepts for financial economics, Nobel laureate Joseph E. Stiglitz and Sanford J. Grossman shared their formal conclusion in "On the Impossibility of Informationally Efficient Markets" that it is impossible for financial markets to be fully efficient at all times:

We have argued that because information is costly, prices cannot perfectly reflect the information which is available, since if it did, those who spent resources to obtain it would receive no compensation. There is a fundamental conflict between the efficiency with which markets spread information and the incentives to acquire information.

The statement of [Grossman and Stiglitz \(1980\)](#) poses a major challenge for the Efficient Market Hypothesis of [Fama \(1970\)](#) and is a catalyst for the ongoing debate about whether financial markets are efficient and reflect all available information at all times, or not. What reconciles both notions is that information assumes a critical role in financial markets.

My thesis encompasses three essays, each of which examines the role of information in a specific setting arising in financial economics. Thus, each essay contributes to the literature about the role of information in financial markets and to the debate whether financial markets are efficient or not. In what follows, I briefly summarize their contents and contextualize their aggregate findings within the context of this thesis to provide a collective conclusion.

Essay 1: What do Market Participants Learn from Share Repurchases? Evidence from a Return Decomposition

The first essay is a joint work with my supervisor Philip Valta and is the research product of a Swiss National Science Foundation grant. We aim at understanding what market participants learn from corporate repurchase announcements and the objective is to deepen our understanding about the nature of information contained in repurchase announcements of firms. By applying a method from asset pricing we extract information about a firm's cash flows and its risk from stock returns. We present evidence that repurchase announcements contain information about a firm's risk when that firm is underpriced. More specifically, we show that market participants learn that their current assessment of the firm's risk is inaccurate and too high given the information that is available to them. Importantly, no new fundamental information about the firm's risk is contained in the repurchase announcement. This initiates a correction of the perceived risk ultimately leading to an appreciation in the firm's stock price.

Our paper makes at least two contributions to the literature. First, we contribute to the literature on the anomalous behavior of stock returns around share repurchase announcements. Second, our paper adds to the literature on the information content of share repurchases.

Essay 2: Skill in the Game

The second essay is single-authored. The main objective of this study is to provide novel insights on whether mutual fund managers possess skill and are not simply lucky when allocating their assets. To that end, I introduce a novel measure that captures whether fund managers can anticipate how stock prices will react to changes in the aggregate market's expectations about the values of stocks, and adjust their fund holdings accordingly. In my setting, the market changes its expectations about the value of a stock due to firm-specific information and information about the entire financial market. I show that fund managers are able to anticipate changes in market expectations that are driven by firm-specific information but not those driven by information that affects the entire financial market. This suggests that mutual fund managers excel and are more precise at acquiring, processing, and using firm-related information for investment decisions. Furthermore, I show that this ability is only prevalent when a fund management consists of a team but not when

it consists of one individual manager only. Finally, I show that firm-specific information in stock prices is less complete than market-wide information, Thus, I provide one possible mechanism that explains why anticipation of changes in the market's expectations driven by firm-specific information is rendered possible.

The contribution of this paper is at least threefold. First, I contribute to the vast literature on the skill of mutual fund managers. Second, I enrich the literature devoted to examining the skill difference between team-managed funds and single-managed funds. Third, the paper examines whether the informational inefficiency of stocks is related to managerial skill.

Essay 3: Entropic Market Timing

The third paper is co-authored with Tim Glaus. We propose an information theoretic approach to measure the extent to which prices in financial markets reflect all available information. Our measures draw on the idea of return predictability and are directly linked to the Efficient Market Hypothesis. The primary duty of the measures is to identify periods where assets or entire financial markets are inefficient in that they do not reflect all available information, such that active asset allocation might become profitable. Using these measures, we propose market timing strategies and provide timing measures for two of the most important and most established financial market phenomena, value and momentum. We also document that market efficiency is cyclical for the U.S. stock market and varies over time.

The contribution of this study is severalfold. In general, our study contributes to the discussion about efficient markets. We first contribute to the literature on market timing. Second, we add to the literature that examines the performance of active investment. Third, we contribute to the literature of price efficiency measures. Finally, our paper contributes to the research that adopts ideas from information theory and maps it to financial markets.

Collectively, my thesis contributes to the debate whether financial markets are efficient or not. The first essay finds that financial market participants have erroneous expectations about the risk of certain firms, suggesting that their information processing is flawed at times. The second essay shows that skilled fund management structures are able to anticipate how the aggregate market's expectations about a stock will shift, implying that information is not always fully and instantly reflected in financial markets. The third essay finds that assets can be informationally incomplete such that financial markets can be timed.

Thus, I conclude in the spirit of [Grossman and Stiglitz \(1980\)](#): *It is impossible for financial markets to be informationally efficient at all times.*

References

- Fama, E. F., 1970. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance* 25, 383.
- Grossman, S. J., Stiglitz, J. E., 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.

Contents

Essay 1: What do Market Participants Learn from Share Repurchases? Evidence from a Return Decomposition	15
I Introduction	16
II Related literature and hypotheses	19
III Data and method	21
A Return decomposition: Cash flow and cost of capital news	22
B Transition matrix	25
C Pricing errors	26
D Abnormal returns	26
IV Results	27
A Short- and long-run stock performance	27
B Changes in cash flow and discount rate news	30
C Discount rate dynamics	38
D Variance of cash flows and discount rates	41
E Firm size and discount rate precision	43
V Conclusion	46
References Essay 1	48
Appendix Essay 1	51
A Return decomposition	51
B Market-to-book decomposition	53
C Additional tables	55
Essay 2: Skill in the Game	59
I Introduction	60
II Related literature and hypotheses	63
A Fund manager skill	64
B Team managed funds	66
III Measuring skill	67

A	Return decomposition	68
B	Informed trading	69
C	Skill	70
IV	Data	72
V	Skill in the U.S. mutual funds industry	74
A	Anticipation of changes in market expectations	74
B	Skill and fund performance	75
C	Skill persistence	77
D	Robustness	83
E	Team-managed funds	87
VI	Informational inefficiency and skill	89
A	Measuring informational inefficiency	90
B	Results	93
VII	Conclusion	96
	References Essay 2	97
	Appendix Essay 2	101
A	Data cleaning and variable construction	101
A.1	Fund universe and share class aggregation	101
A.2	Mutual fund managers	101
A.3	Fund holdings and returns	102
A.4	Return decomposition	103
B	Name cleaning measures	104
B.1	Levenshtein measure	104
B.2	Cosine similarity measure	105
C	Additional tables	106
	Essay 3: Entropic Market Timing	108
I	Introduction	109
II	Methodology	113
A	Information complexity	113
B	Enhanced Kontoyiannis entropy estimator	114
C	Entropy estimation	115

D	Efficiency metrics	118
E	Data	120
III	Results	120
A	Cross-sectional analysis of stock price efficiency	120
B	Timing value	127
C	Timing short selling	136
IV	Conclusion	142
	References Essay 3	144
	Appendix Essay 3	148
A	Entropy Estimator	148
B	Variable construction	149
C	Additional Figures	150
	Statement of Authorship	152

List of Tables

Essay 1: What do Market Participants Learn from Share Repurchases? Evidence from a Return Decomposition		15
I	Descriptive statistics	24
II	Vector Auto Regression (VAR) transition matrix	25
III	Stock performance after repurchase announcements and pricing errors	28
IV	Change in the level of cash flow and discount rate news	31
V	Pricing errors and changes in the level of cash flow and discount rate news	32
VI	Levels of cash flow and discount rate news around repurchase announcements	33
VII	Multivariate analysis of cash flow and discount rate news levels	35
VIII	Proximity to the repurchase announcement date and changes in discount rates	37
IX	VAR implied levels of discount rates around repurchase announcements	40
X	Variances of cash flow and discount rates	42
XI	Firm size, discount rates, and cash flows	44
C.I	VAR Transition Matrix - Middle tercile	55
C.II	VAR Transition Matrix - Bottom tercile	56
C.III	VAR Transition Matrix - Top tercile	57
C.IV	Mispricing and changes in the level of cash flow and discount rate news	58
 Essay 2: Skill in the Game		 59
I	Descriptive statistics	73
II	Annualized performance of skilled and unskilled funds	75
III	Outperformance of skilled fund managers	77
IV	Skill autocorrelation	78
V	Beta Persistency	79
VI	Multivariate Beta Analysis	82
VII	Reverted skill regression and lagged news	84
VIII	Erikson-Whited errors-in-variables regression	85
IX	Team-managed vs. single-managed funds	88
X	Asset efficiency	94
XI	Portfolio efficiency	95
C.I	Vector Auto Regression (VAR) transition matrix	106
C.II	Beta autocorrelation for partially skilled funds and non-zero holding changes	107
 Essay 3: Entropic Market Timing		 108
I	Cross-sectional efficiency	121
II	Value and buy and hold assets	130
III	Timing value with RIC	132

IV	Cumulative return of timing strategy	133
V	Alphas for the active timing strategy	134
VI	Timing value with Delay	136
VII	Alphas for the active timing strategy using delay measures	137
VIII	Futures contracts	139
IX	Timing short positions	140

List of Figures

Essay 1: What do Market Participants Learn from Share Repurchases? Evidence from a Return Decomposition	15
1 Level of discount rates around repurchase announcements for the bottom quintile	39
Essay 2: Skill in the Game	59
1 Funds and assets under management	74
2 Performance of skilled and unskilled funds	76
Essay 3: Entropic Market Timing	108
1 Parameter space for number of past returns and return quintiles	116
2 Efficiency of the cross-section of U.S. stocks between 1985 and 2020	122
3 Price delay of the cross-section of U.S. stocks between 1985 and 2020	126
4 Inefficiency of value over time	129
5 Timing value using RIC as an indicator of inefficiency	131
6 Timing short positions using information complexity	141
A.I Cumulative Return of Active Timing Strategy	150

Essay I:

What do Market Participants Learn from Share Repurchases? Evidence from a Return Decomposition^{*}

Sascha Jakob and Philip Valta[†]

ABSTRACT

This paper analyzes cash flow and cost of capital dynamics around share repurchase announcements of publicly traded US firms by decomposing stock returns into news related to cash flows and discount rates. After repurchase announcements, the cost of capital decreases significantly, while cash flows do not change. The decrease in the cost of capital is largest for firms that appear underpriced. These firms also experience the highest long-term returns after repurchase announcements. The volatility of the discount rate and cash flows also decreases but is not systematically related to long-term returns. The findings suggest that market participants learn about a temporary overestimation of the cost of capital when firms announce share repurchases.

Keywords: Share repurchases, Buyback Anomaly, Return Decomposition, Cost of Capital

JEL Classification Numbers: G14, G32, G35

^{*}The paper is currently under revision following a revise and resubmit recommendation at the Journal of Corporate Finance.

[†]Jakob and Valta are both at the University of Bern. Corresponding author is Philip Valta. Email: philip.valta@unibe.ch. Address: Institute for Financial Management, Engehaldenstrasse 4, 3012 Bern, Switzerland. Phone: +41 31 684 37 12. We thank Julien Cujean, Pierre Collin-Dufresne, Dusan Isakov, Sonia Konstantinidi, Tiago Loncan, Giorgia Simion, and seminar and conference participants at King's College London, University of Bern, the 2022 French Finance Association meetings, 2022 European FMA meetings, and the 2022 Swiss Finance Institute Research Days for valuable comments and suggestions. Funding from the Swiss National Science Foundation is gratefully acknowledged.

I Introduction

Despite the Covid pandemic, firms are repurchasing their own shares in record volume.¹ Share repurchase announcements are associated with a positive stock market reaction, suggesting that these announcements affect market participants' expectations about cash flows and the cost of capital (see, e.g., [Farre-Mensa et al. \(2014\)](#)). A substantial research effort has been devoted to understand firms' repurchase motives and the information that these decisions convey to financial market participants.² Yet, despite significant progress, continued disagreement exists about the interpretation of the information that repurchase announcements signal to market participants. Specifically, research provides evidence that share repurchase announcements contain information about the underpricing of an underlying stock (see, e.g., [Dann \(1981\)](#); [Vermaelen \(1981\)](#); [Kahle \(2002\)](#); [Jagannathan et al. \(2000\)](#); [Dittmar \(2000\)](#); [D'Mello and Shroff \(2000\)](#); [Brav et al. \(2005\)](#)), and that the positive stock performance after the repurchase announcement is a correction of this underpricing (see, e.g., [Peyer and Vermaelen \(2009\)](#); [Manconi et al. \(2018\)](#)). However, other papers advocate that share repurchases signal lower future risk (see, e.g., [Grullon and Michaely \(2004\)](#); [Michaely et al. \(2020\)](#)). Distinguishing between these alternative interpretations is challenging because we do not observe the signal that share repurchase announcements convey to market participants.

The objective of this paper is to advance this discussion and shed light on the information contained in repurchase announcements by applying a method from asset pricing to extract news related to changes in future cash flows and risk (discount rates) from stock returns. Existing research has used the decomposition productively (see, e.g., [Campbell \(1991\)](#); [Vuolteenaho \(2002\)](#); [Michaely et al. \(2020\)](#)). The method is based on a vector auto regression (VAR) and produces estimates of cash flow and discount rate news before and after repurchase announcements.³ We can thus infer news about cash flows and discount rates from market prices rather than financial statement data. This advantage is instrumental to our analysis, as we aim at understanding what market participants learn from repurchase announcements.

Consider a discounted cash flow model of a firm where the market determines the value of a stock by the expected cash flows and the cost of capital of the firm (e.g., [Vuolteenaho \(2002\)](#)). We assume that the market price of the firm is subject to a potential pricing error at the time of the repurchase announcement. Thus, the market price is the sum of the fundamental value and a pricing error, which captures non-fundamental perceptions of market participants about future cash flows and the cost of capital. A negative pricing error implies that the market price is below the fundamental value and indicates underpricing. When

¹See, e.g., "US companies buy back shares in record volumes", Financial Times, March 27, 2022.

²Section II reviews this literature in detail.

³Conceptually, *news* in our setting is the difference between the VAR's prediction using the latest fundamental information and the market's contemporaneous assessment of cash flows and discount rates. Differences arise because the market obtains and processes information that is not incorporated in the VAR's prediction. We infer the informational nature of *news* by estimating changes in news and changes in the VAR predictions around repurchase announcements.

a firm announces a share repurchase, the signal can contain information about the firm's fundamental value and/or about the pricing error. If the repurchase announcement contains news about the firm's pricing error, we expect the stock returns after the repurchase announcement to be higher for firms with negative pricing errors at the time of the repurchase announcement. Moreover, in such case we also expect the changes in cash flow and discount rate news to be systematically related to the pricing error. Firms with a negative pricing error should see a decrease in the cost of capital or an increase in cash flows.

We analyze these hypotheses using a sample of 2,417 open market share repurchase announcements for 1,844 distinct publicly listed U.S. firms for the years 1995 to 2019. In a first step, we show that the short-term cumulative abnormal returns (CARs) are positive and statistically significant across all repurchase announcements, corroborating existing research (e.g., [Farre-Mensa et al. \(2014\)](#)). We then introduce a measure of pricing error, which is based on the market-to-book decomposition of [Rhodes-Kropf et al. \(2005\)](#). This method allows decomposing the market-to-book ratio into a firm-specific and industry-specific pricing error. The firm-specific pricing error measures deviations from valuations implied by sector valuation multiples. The industry-specific pricing error captures the deviations of short-run multiples from their long-run average values. Using the the firm-specific pricing error from the third model in [Rhodes-Kropf et al. \(2005\)](#), we group firms into five quintiles. The bottom quintile contains firms that are likely to be underpriced at the time of the repurchase announcement (i.e., they have a negative pricing error). The top quintile contains firms that are likely to be overpriced. We show that the short-term CARs are more than twice as large for firms in the bottom quintile compared to firms in the top quintile. Moreover, the long-term abnormal stock performance (post-announcement drift) is strongly positive and significant for stocks with a negative pricing error (bottom quintile). These firms experience significant abnormal returns between 11% and 25% over the 12 to 36 months period after repurchase announcements. For stocks with a positive pricing error, the post-announcement drift is close to zero and not statistically significant. These findings support existing evidence on the share buyback anomaly (e.g., [Peyer and Vermaelen \(2009\)](#)) using an alternative measure of mispricing and suggest that financial market participants gradually learn about and initiate a correction of a potential pricing error at the time of the announcement.

To analyze the information revealed in share repurchase announcements, we investigate the changes in the levels of cash flows and the cost of capital around repurchase announcements. Specifically, we analyze the difference between the level of discount rate (cash flow) news from one quarter before until one quarter, one year, two years, and three years after the announcement. We document a significant average decrease in discount rate news of 40 to 50 basis points around repurchase announcements. The decrease is persistent until three years after the announcement. The level of cash flows news, by contrast, does not significantly change, consistent with [Grullon and Michaely \(2004\)](#), who show that operating performance does not improve after repurchase announcements.

We further show that discount rate news decrease most significantly in the bottom quintile of firms with a negative pricing error, ranging from 20 basis points (for one quarter) to 90 basis points (for 3 years). For these firms, we also observe the highest initial stock market reaction and long-term returns. For firms in the top quintile, the change in the level of discount rate news is quantitatively smaller (only about 30 basis points). These results hold in a univariate setting and a multivariate regression analysis that controls for firm characteristics and year and industry fixed effects. The findings are also robust to estimating separate VARs for each group of pricing error, using different estimation windows, or relaxing data restrictions regarding overlapping repurchase events. The findings suggest that share repurchase announcements reveal information about the cost of capital when firms are potentially underpriced. By contrast, we observe only little changes in cash flow news. In further tests, we show that the level of discount rate news is significantly positive one quarter before the repurchase announcement and reverts to zero after the announcement. The observed dynamics suggest that market participants learn that a stock is underpriced because the cost of capital is temporarily too high. The repurchase announcement thus helps correcting investors' perception about firms' systematic risk, which generates positive short-term and long-term abnormal returns after the repurchase announcement and gradually corrects the mispricing.

Next, we analyze the changes in the volatility of cash flows and discount rates. Our analysis shows that the volatility of cash flows declines significantly after repurchase announcements. This finding is consistent with recent evidence in [Michaely et al. \(2020\)](#), who show that payout decisions signal lower future cash flow volatility. In addition, we find that the volatility of discount rates also significantly decreases by 11.3% after repurchase announcements. This result is novel and suggests that financial market participants learn about the distribution of firms' systematic risk through repurchase announcements and assess the cost of capital with more precision after the announcement.

Furthermore, we analyze whether the changes in the volatility of cash flows and costs of capital depend on the firm-specific pricing error. We show that the decrease in cash flow volatility for stocks in the bottom quintile is much smaller (-15.3%) compared to the decrease for stocks in the top quintile (-24.2%). We observe a similar pattern for the reduction of the volatility of the cost of capital (-19.2% for the bottom quintile and -21.3% for the top quintile). Hence, while the relatively overpriced firms experience the largest decrease in cash flow and cost of capital volatility, these stocks are not associated with positive long-term abnormal returns. By contrast, relatively underpriced firms experience a smaller decrease in cash flow and cost of capital volatility but experience large and significant long-term abnormal returns. We conclude that, while share repurchase announcements are associated with lower future cash flow and cost of capital volatility on average, the decrease in volatility cannot explain the cross-sectional differences in long-run abnormal returns.

Our paper makes several contributions to the literature. First, the paper contributes to the literature on the share buyback anomaly (see, e.g., [Lakonishok and Vermaelen \(1990\)](#); [Ikenberry et al. \(1995\)](#); [Vermaelen \(1981\)](#); [Chan et al. \(2004\)](#); [Peyer and Vermaelen \(2009\)](#);

Manconi et al. (2018); Evgeniou et al. (2018)). We document for a large sample of share repurchase announcements a persistent post-announcement drift for underpriced firms, where we measure underpricing using the firm-specific pricing error of Rhodes-Kropf et al. (2005). In contrast to existing buyback anomaly papers, we present a systematic analysis of the information revealed in share repurchase announcements using the Campbell (1991) decomposition of stock returns. One advantage of this method is that it allows extracting information on both cash flows and costs of capital from stock returns rather than from infrequent accounting data. Our findings suggest that underpricing among repurchasing firms, and thus the buyback anomaly, is driven by a temporary overestimation of the cost of capital. As firms announce share repurchases, investors gradually learn their mistake in assessing the firms' systematic risk, and the stock price increases.

Second, our paper adds to the literature on the information content of share repurchases. The most closely related paper is by Michaely et al. (2020). They decompose returns with the Campbell (1991) decomposition and show that payout decisions are associated with a decrease in the volatility of cash flows. While our analysis confirms their finding of a decrease in average cash flow volatility, our paper differs from and advances their analysis in at least four important ways. First, Michaely et al. (2020) analyze mostly dividends. The focus of our analysis are share repurchases, which are associated with a persistent post-announcement drift. Second, we relate the results from the return decomposition to the cross-sectional variation of long-term abnormal returns and show that news about the decrease in cash flow volatility are not able to explain the long-term abnormal returns after repurchase announcements. Thus, it appears that changes in cash flow volatility are not priced in the context of repurchase announcements. Third, we show that the variance of discount rate news significantly decreases. This result suggests that market participants receive valuable information to assess firms' systematic risk more precisely at the repurchase announcement. Fourth, discount rate changes are an economically important component to understanding the stock return dynamics around repurchase announcements. As such, our results complement those in Grullon and Michaely (2002) and Grullon and Michaely (2004), who argue that dividends and share repurchases signal lower systematic risk because firms will be more mature and have fewer growth options in the future. Our interpretation, however, is different. We focus on mispricing by relating changes in firms' systematic risk (cost of capital) to pricing errors and show that share repurchases help correcting investors' perception about the firm's current cost of capital for firms that are underpriced. Our results thus contribute to the debate about whether repurchase announcements signal maturity (e.g., Grullon and Michaely (2004)) or underpricing (e.g., Peyer and Vermaelen (2009)), suggesting that long-term abnormal returns are the result of temporary mispricing.

II Related literature and hypotheses

A robust finding in the literature on share repurchases is that the initial market response to the announcement of share repurchases is positive, economically large, and statistically

significant. This finding holds across many empirical studies, for the U.S., and for other countries in the world (see, e.g., [Dann \(1981\)](#); [Vermaelen \(1981\)](#); [Peyer and Vermaelen \(2009\)](#); [Manconi et al. \(2018\)](#)). Moreover, several papers document that firms, which announce share repurchases, tend to outperform other firms that do not repurchase shares in the long-term. This "share buyback anomaly" was first reported by [Lakonishok and Vermaelen \(1990\)](#) and [Ikenberry et al. \(1995\)](#). These papers document superior stock performance of repurchasing firms up to four years subsequent to the announcement of the repurchase program. Other papers confirm the existence of this anomaly (e.g., [Chan et al. \(2004\)](#), [Grullon and Michaely \(2004\)](#), or [Peyer and Vermaelen \(2009\)](#)) and show that the post-announcement drift is robust to controlling for the five factors of [Fama and French \(2015\)](#) and for stock liquidity (see, e.g., [Evgeniou et al. \(2018\)](#)). Recent evidence shows that the anomaly persists in markets outside the U.S. (see, e.g., [Manconi et al. \(2018\)](#)).⁴

While the literature broadly agrees on the robustness of the positive initial price reaction, there is disagreement about i) the interpretation of the information that triggers the initial market response and ii) the robustness and interpretation of the post-announcement drift. A dominant view in the literature that analyses the post-announcement drift is that financial markets are not fully efficient and that repurchase announcements contain information about the stock's mispricing. Several papers provide evidence supporting this mispricing hypothesis (see, e.g., [Dann \(1981\)](#); [Vermaelen \(1981\)](#); [Kahle \(2002\)](#); [Jagannathan et al. \(2000\)](#); [Dittmar \(2000\)](#); [D'Mello and Shroff \(2000\)](#)). Survey evidence also suggests that underpricing is an important motive for firms' repurchase decision ([Brav et al. \(2005\)](#)).

However, other papers offer alternative explanations for the initial stock market response and long-term returns. For example, [Grullon and Michaely \(2004\)](#) show that the operating performance of repurchasing firms does not improve after repurchase announcements, and that the positive initial market response is consistent with lower future risk of these firms. [Kumar et al. \(2008\)](#) show that after share repurchase announcements the stock's equity beta, i.e., its systematic risk and the standard error associated with this beta, decrease. More recently, [Michaely et al. \(2020\)](#) argue that both dividends and share repurchases signal lower future cash flow volatility. Alternatively, [Dittmar and Field \(2015\)](#) confirm the existence of the timing ability of repurchasing firms to purchase underpriced stock, however, only in a sample of firms that repurchase shares infrequently. They conclude that mispricing alone cannot explain all the repurchases.⁵ Other papers suggest that the post-announcement drift can be explained by increased takeover risk exposure ([Bargeron et al. \(2017\)](#); [Lin et al. \(2014\)](#)) and that it has strongly declined in recent years due to improved financial market efficiency ([Fu and Huang \(2016\)](#)). Our analysis contributes to this research by analyzing

⁴Actual repurchases are spread over periods of up to three years after the repurchase announcement, and firms often repurchase less than the actual amount announced (see, e.g., [Stephens and Weisbach \(1998\)](#); [Oded \(2005\)](#); [Bonaimé et al. \(2014\)](#)). Moreover, actual repurchases are associated with a positive market reaction, a firm-specific release of information, higher liquidity, and more insider trading (e.g., [Chung et al. \(2007\)](#); [Ben-Rephael et al. \(2014\)](#)). [Bonaimé \(2012\)](#) analyzes the reputation of firms and how completion rates are associated with returns.

⁵[Lee et al. \(2020\)](#) discuss alternative explanations for share repurchases, such as managerial self interest and hubris.

the dynamics of cash flows and costs of capital by using the stock return decomposition.

Consider a discounted cash flow model of a firm, where the value of a stock is determined by the expected cash flows and the cost of capital of the firm (e.g., Vuolteenaho (2002)). We assume that the firm is subject to a potential pricing error at the time of the repurchase announcement, such that its market price deviates from its fundamental value. The market price is thus the sum of the fundamental value, $V(CF, r)$, and a pricing error, ϵ , which captures non-fundamental perceptions of future cash flows and/or the cost of capital: $P = V(CF, r) + \epsilon$. A negative pricing error (i.e., $\epsilon < 0$) implies that the market price is below the fundamental value and indicates underpricing. A possible interpretation of a negative pricing error is that market participants are irrationally pessimistic about the firm's cash flows and/or its systematic risk (cost of capital).

A repurchase announcement can reveal information about the firm's fundamental value and about its pricing error. If the repurchase announcement reveals news about fundamentals, a positive initial market reaction at the repurchase announcement date could signal higher cash flows, lower systematic risk, or both. Alternatively, repurchase announcements could also contain information about the firm's pricing error and thus help correcting investors' perception about the firm's true future cash flow and cost of capital. In this case, we expect the stock returns after the repurchase announcement to be systematically related to the pricing error at the time of the repurchase announcement. Firms with a negative pricing error should have higher post-announcement returns compared to firms with no or with a positive pricing error. Moreover, if the repurchase announcement signals information about the pricing error, we also expect the changes in cash flow or cost of capital news to be related to the pricing error at the time of the repurchase announcement. Firms with a negative pricing error should see a decrease in the cost of capital or an increase in cash flow news. Before turning to the empirical analysis, we summarize the testable hypotheses:

Hypothesis 1: *For firms with a negative pricing error, the short-term and long-term CARs are higher compared to firms with no or with a positive pricing error.*

Hypothesis 2: *The change in cash flow and cost of capital news is systematically related to the pricing error. Firms with a negative pricing error experience a decrease in the cost of capital or increase in cash flows.*

III Data and method

Our sample construction starts with all open share repurchase programs announced by publicly listed U.S. firms between 1995 and 2015 that are available in the SDC Platinum database. The sample period for repurchase announcements stops in 2015 because we use 20 quarters of post-announcement data to estimate long-run returns and the vector auto regression (we require at least 12 quarters of data prior and subsequent to the announcement). We make sure that the estimation window is not overlapping with another repurchase announcement of the same firm. We also drop firms with missing or invalid CUSIP and firms in

financial (SIC codes 6000-6999) and other regulated industries (4900-4999). Monthly stock price data come from the CRSP database, and balance sheet data come from the quarterly COMPUSTAT database covering the years 1995 to 2019. We end up with a sample of 2'417 share repurchases announced by 1'844 distinct firms.

A Return decomposition: Cash flow and cost of capital news

We follow [Michaely et al. \(2020\)](#) and [Vuolteenaho \(2002\)](#) and decompose stock returns into the first (average) and second moments (variance) of expected cash flows and costs of capital using monthly returns and quarterly COMPUSTAT data. Consistent with a model that values a firm's stock as the present value of expected cash flows discounted with the expected costs of capital, the decomposition extracts news related to the average and variance separately for cash flows and the discount rate. We extract this information for both the period before and after the repurchase announcements, and then analyze how news related to cash flows and the discount rate change around the repurchase announcements. The method allows obtaining estimates of changes in cash flows and discount rates (and their volatility) and thus maps directly into our testable hypotheses. An important advantage of the method is that it infers news from market prices rather than balance sheet data. Furthermore, balance sheet data is only available at a low frequency. By using market prices, we can circumvent this problem and obtain more frequent estimates of cash flow and discount rate news around repurchase announcements.

We express unexpected returns as a function of changes in expectations about cash flows and costs of capital and use a Vector Auto Regression (VAR) to implement the return and return variance decomposition as in [Vuolteenaho \(2002\)](#) and [Michaely et al. \(2020\)](#). Let $\mathbf{z}_{i,t}$ be a vector at time t that contains firm-specific state variables of firm i and has returns as its first component. Assuming that a first-order VAR is sufficient to describe the evolution of the state variables in $\mathbf{z}_{i,t}$, we can write the VAR system as follows:

$$\mathbf{z}_{i,t} = \mathbf{\Gamma}\mathbf{z}_{i,t-1} + \mathbf{u}_{i,t}. \quad (1)$$

$\mathbf{\Gamma}$ is the transition matrix of the VAR system, and $\mathbf{u}_{i,t}$ is an error term. We further define the vector $\mathbf{e}\mathbf{1}' = [1 \ 0 \ 0]$ and rewrite unexpected stock returns as:

$$r_t - \mathbb{E}_{t-1}[r_t] = \mathbf{e}\mathbf{1}'\mathbf{u}_{i,t}. \quad (2)$$

Using the transition matrix as well as the residuals of the VAR allows us to decompose equation (2) into cash flow news and discount rate news components such that the level of discount rate news for a given quarter, t , is:

$$\eta_{r,t} = \boldsymbol{\lambda}'\mathbf{u}_{i,t}, \quad (3)$$

where $\boldsymbol{\lambda}$, among other things, contains $\mathbf{\Gamma}$. The level of cash flow news is:

$$\eta_{cf,t} = (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}')\mathbf{u}_{i,t}. \quad (4)$$

The second moments are the variance of the levels before and after the repurchase announcement. Unexpected returns map to the difference in cash flow news and discount rate news:

$$r_t - \mathbb{E}_{t-1}[r_t] = \eta_{cf,t} - \eta_{r,t} \quad (5)$$

Appendix A describes the method in greater detail.

From equations (3) and (4), it is apparent that news, η_t , are a function of the VAR residuals. Specifically, cash flow and discount rate news are the portion of the residuals, or unexpected returns, that are associated with cash flows and the discount rate, respectively. As such, they capture implicit information in stock returns that are associated either with cash flows or costs of capital.⁶ The variances of cash flow and discount rate news measure the precision with which the market assesses cash flows and costs of capital with respect to the equilibrium model.

We follow Michaely et al. (2020) and Vuolteenaho (2002) and construct the variables for the return decomposition. We compute the simple quarterly stock return as the cumulative monthly return within a fiscal quarter, recorded from m to $m + 2$ for $m \in \{\text{February, May, August, November}\}$. As in Michaely et al. (2020), we assume a de-listing return of 30% if a firm is de-listed for a known cause and has a missing de-listing return. The return r_t is the market-adjusted log return defined as log return less the cross-sectional average log return (see, e.g., Vuolteenaho (2002)). Market equity is defined as the total firm market equity as recorded in CRSP at the end of each quarter. If quarter t market equity is missing, we let the previous quarter’s market equity grow with the rate of return in that quarter (without dividends).

Book equity is defined as shareholders’ equity plus balance-sheet deferred taxes and investment tax credit (item TXDITCQ) if available, minus the book value of preferred stock. We use stockholders’ equity (item SEQQ), or common equity (item CEQQ) plus the carrying value of preferred stock (item PSTKQ), or total assets (item ATQ) minus total liabilities (item LTQ) in that order as shareholders’ equity. We use the redemption value (item PSTKRQ) if available, or carrying value for the book value of preferred stock. Whenever book equity is unavailable, we proxy for it by the last period’s book equity plus earnings, less dividends, assuming that the clean-surplus relation holds. If neither earnings nor book equity are available, we assume that the book-to-market ratio has not changed from quarter $t - 1$ to quarter t , and compute the book-equity proxy from the last quarter’s book-to-market ratio and this quarter’s market equity. We exclude firms with a quarter

⁶Note that *news* in this context are not a specific, directly observable piece of information such as news about earnings announcements or a corporate event. Rather, news is undefined latent information conveyed through prices that either relates to cash flows or to costs of capital. Formally, the term *news* relates to the information one obtains by comparing the market levels of cash flows and costs of capital with what the equilibrium model (the VAR model) expects the levels to be given all information that is available before the announcement. It measures the extent to which the market perception of future cash flows and costs of capital differs relative to the expected value after observing a signal. The difference could be a result of new fundamental information about the firm, noise, or a combination of the two.

Table I: Descriptive statistics

This table reports descriptive statistics. Panel A summarizes the variables of the vector auto regression (VAR) system. r_q is the quarterly return, BM-Ratio $_q$ is the book-to-market ratio, and RoE_q is the return on equity. Panel B reports summary statistics of the variables we use in the analysis. The firm-specific and industry-specific pricing errors are the mispricing measures of Rhodes-Kropf et al. (2005) in the repurchase quarter. Log(Market Cap) is the logarithm of the market capitalization at the time of the repurchase announcement. Age is approximated by the difference between the firm year and the first year the firm reports in CRSP. Debt-to-Assets is the ratio between interest bearing long-term and short-term debt relative to total assets. The sample period is from 1995 to 2019.

Panel A: Variables of the VAR system						
	N	Mean	Median	5%	95%	sd
r_q	134'634	0.00	0.00	-0.30	0.26	0.17
BM-Ratio $_q$	134'634	0.73	0.64	0.28	1.46	0.43
RoE_q	134'634	0.03	0.03	-0.02	0.08	0.04
Panel B: Descriptive statistics of variables in the analysis						
	N	Mean	Median	5%	95%	sd
Firm-specific pricing error	2'417	0.19	0.16	-0.68	1.13	0.57
Industry-specific pricing error	2'417	0.06	0.06	-0.35	0.50	0.26
Log(Market Cap)	2'417	6.66	6.58	3.68	10.00	1.92
Debt-to-Assets	2'413	0.19	0.16	0.00	0.51	0.17
Age	2'417	14.43	14.00	5.00	28.00	7.21

$t - 1$ market equity of less than USD 10 million and a book-to-market ratio of more than 100 or less than 0.01. Moreover, we set negative or zero book-equity values to missing.

ROE is defined as earnings over beginning of quarter's book equity. To compute the ROE we use earnings available for common equity. When earnings are missing, we use the clean-surplus formula to approximate earnings. We drop observations with a ROE lower than - 100%. Each quarter, we log transform market equity, stock returns, and return on equity and cross-sectionally demean it. A log transformation can be cumbersome if returns are close to 1 or if book-to-market ratios are close to zero or infinity. We mitigate these concerns by following Michaely et al. (2020) and Vuolteenaho (2002) and redefine each firm as a portfolio that consists of 90% common stock and 10% Treasury bills using market values. Each period, the portfolio is rebalanced to reflect these weights.

Panel A of Table I reports the descriptive statistics of the variables that we use for the VAR estimation. The quarterly return is zero on average. The average (median) book-to-market ratio is 0.73 (0.64), implying that the equity of the average firm trades above its book value. The average and median quarterly returns on equity are 3%.

Table II: Vector Auto Regression (VAR) transition matrix

This table reports the point estimates of a panel VAR for all repurchasing firms using the method outlined in section III. r_t denotes the centered excess log stock return, θ is the centered log book-to-market ratio, and RoE is the centered log return-on-equity. The sample period is from 1995 to 2019.

	r_{t-1}	θ_{t-1}	RoE_{t-1}
r_t	-0.005 (-1.36)	0.039 (12.21)	0.266 (15.18)
θ_t	-0.291 (-69.94)	0.844 (252.38)	-0.266 (-12.76)
RoE_t	0.018 (18.48)	-0.021 (-22.36)	0.219 (30.92)

B Transition matrix

An essential element of the analysis is the estimate of the transition matrix Γ of the VAR system and the discount factor ρ to generate the averages and variances of cash flow and cost of capital news. As in Vuolteenaho (2002) and Michaely et al. (2020), we define a new variable that is equal to the excess log ROE minus the excess log stock return plus the lagged book-to-market ratio. ρ is the coefficient estimate of a regression of this variable on the contemporaneous book-to-market ratio. We obtain a coefficient of 0.95, which is close to Vuolteenaho (2002) (coefficient of 0.97). The VAR system is estimated using data from 20 quarters before to 20 quarters after the repurchase announcement. We require that a repurchase announcement has at least 12 quarters of data prior and subsequent to the announcement and that the estimation window is not overlapping with another repurchase announcement of the same firm. Using these restrictions, we end up with a sample of 2'417 repurchase announcements. A large sample is necessary to get precise estimates of the transition matrix Γ . Therefore, estimating separate VARs for each repurchase announcement can lead to efficient but imprecise estimates of Γ . Hence, we estimate one panel VAR over all repurchase events to obtain the estimates of the transition matrix Γ . Then, we estimate separate VARs before and after each repurchase announcement to obtain the VAR residuals.

Table II shows the point estimates of the constant panel VAR with one lag for all repurchase announcements. The autocorrelation of market adjusted log returns, r_t , is close to zero and statistically insignificant. Centered market adjusted log returns load positively on the lagged log book-to-market ratio and lagged log ROE. The book-to-market ratio is significantly autocorrelated and loads negatively on lagged market adjusted log returns and on lagged ROE. ROE is positively autocorrelated, positively correlated with lagged returns, and negatively correlated with the lagged book-to-market ratio. These dynamics in the state variables are broadly consistent with the findings in Michaely et al. (2020). Two differences are worth mentioning. First, in contrast to Michaely et al. (2020), lagged market adjusted returns are not correlated with contemporaneous returns in our sample. The correlation

is positive but close to zero in [Michaely et al. \(2020\)](#). Second, the log book-to-market ratio loads negatively on lagged returns. [Michaely et al. \(2020\)](#) report a marginally positive value. The differences are likely due to different sample periods and the fact that our sample focuses on repurchase announcements, while their analysis focuses on dividend events.

C Pricing errors

To measure pricing errors, we decompose the market-to-book ratio using the method developed by [Rhodes-Kropf et al. \(2005\)](#).⁷ The method decomposes the market-to-book ratio into a firm-specific and industry-specific pricing error. The firm-specific pricing error measures deviations from valuations implied by sector valuation multiples. The industry-specific pricing error captures deviations of short-run multiples from their long-run average values. In the following, we use the firm-specific component as our main measure of pricing error. This measure of pricing error maps well into our conceptual framework, in which the observed market price is the sum of the fundamental value and a pricing error. [Appendix B](#) describes the construction of the measure in detail.

We estimate the pricing errors using the full merged CRSP and Compustat sample. [Panel B of Table I](#) reports the descriptive statistics for the market-to-book decomposition. The average (median) firm-specific pricing error is marginally positive with an average (median) of 0.19 (0.16). These statistics are similar to those reported in [Rhodes-Kropf et al. \(2005\)](#). The average (median) firm size measured as the logarithm of market capitalization is 6.66 (6.58).

D Abnormal returns

To estimate the initial market response around share repurchase announcements, we adopt standard event study methods using the market model. To analyze long-run abnormal returns at the portfolio level, we use [Ibbotson \(1975\)](#)'s IRATS approach and a calendar time approach (see, e.g., [Peyer and Vermaelen \(2009\)](#)) using a Fama-French five factor model augmented with momentum for both approaches. Specifically, for the IRATS approach, we compute long-run cumulative average abnormal returns (CAR) subsequent to repurchase announcements over time and across security using the following cross-sectional regression each month in event time, j :

$$(r_{i,t} - r_{f,t}) = \alpha_j + b_j(r_{m,t} - r_{f,t}) + c_jSMB_t + d_jHML_t + e_jUMD_t + f_jRMW_t + g_jCMA_t + \varepsilon_{j,t}, \quad (6)$$

⁷This decomposition has been widely used in finance research. For example, [Golubov and Konstantinidi \(2019\)](#) reassess the interpretation of the value premium in the cross-section of stock returns and show that the market-to-value component from this decomposition drives all of the value premium. [Bonaimé et al. \(2014\)](#) use this decomposition to analyze the effects of misvaluation and leverage on the economic gains of share repurchases. [Cziraki et al. \(2021\)](#) use the decomposition to analyze investor sentiment after insider trades before repurchase announcements.

where $r_{i,t} - r_{f,t}$ denotes the excess return of stock i in calendar month t that corresponds to event month j . $j = 1$ is the first month after the repurchase announcement for each security, $j = 2$ is the second month after the repurchase announcement, etc. $r_{m,t}$ denotes the return of the value-weighted CRSP index, SMB , HML , RMW , and CMA are the size, value, profitability, and investment factor of the [Fama and French \(2015\)](#) five factor model, and UMD is the momentum factor of [Carhart \(1997\)](#). Summing over all intercepts from the cross-sectional regressions in every event month j subsequent to the announcements yields the cumulative abnormal return. We consider long-run abnormal returns for 12, 24, and 36 months after the repurchase announcement.

For the calendar time approach, each month we form portfolios that only include stocks of firms that announced a repurchase program during the last 12, 24, or 36 months. We regress the time series of portfolio returns on the factors in equation (6). The intercept (constant), α , is the average monthly abnormal return (AAR) of the portfolio. For both the IRATS and calendar time approach we use excess returns relative to the risk-free rate.

IV Results

In this section, we present the main results. We start by analyzing how relative mispricing relates to stock returns of repurchasing firms in the short- and long-run. Second, we analyze the changes in levels of cash flow and discount rate news as a function of the pricing error to understand how repurchase announcements affect market expectations of cash flows and the cost of capital. Finally, we examine whether changes in the variances of cash flow and discount rate news are associated with long-run abnormal returns.

A Short- and long-run stock performance

We assign each firm that announces a share repurchase program to a quintile of relative mispricing using the firm-specific measure of [Rhodes-Kropf et al. \(2005\)](#) in the quarter when the repurchase program is announced. Panel A of Table III reports the short-term cumulative abnormal returns (CARs) over various event windows for the five mispricing quintiles. The CARs are positive for all quintiles, confirming the robust finding of a positive initial market reaction in the literature (see, e.g., [Farre-Mensa et al. \(2014\)](#)). Moreover, for stocks in the bottom quintile, the CARs are significantly higher compared to the CARs for stocks in the top quintile. These results support the first hypothesis and are consistent with evidence that analyses short-term CARs as a function of relative mispricing (see, e.g., [Bonaimé et al. \(2014\)](#)).

Panel B of Table III reports the annualized long-run abnormal returns using the Fama-French IRATS approach for stocks in the five quintiles over horizons of 12, 24, and 36 months. Panel C reports the average abnormal returns (AAR) using the Fama-French calendar time approach. Relatively underpriced firms are associated with significant positive long-term abnormal returns of 10.94%, 19.41%, and 25.44% over one, two, and three years, respectively.

Table III: Stock performance after repurchase announcements and pricing errors

This table reports the post-announcement stock performance conditional on the firm-specific pricing error. We form quintiles based on the pricing error at the time of the repurchase announcement. Stocks in the bottom quintile are likely underpriced and vice versa. Panel A shows the short-term CARs from the market model. Panel B reports long-term CARs estimated with [Ibbotson \(1975\)](#)'s IRATS approach. The IRATS approach computes long-run cumulative average abnormal returns subsequent to repurchase announcements using the following cross-sectional regression each event month j :

$$(r_{i,t} - r_{f,t}) = \alpha_j + b_j(r_{m,t} - r_{f,t}) + c_jSMB_t + d_jHML_t + e_jUMD_t + f_jRMW_t + g_jCMA_t + \varepsilon_{j,t},$$

where $r_{i,t} - r_{f,t}$ denotes the excess return of stock i in calendar month t which corresponds to event month j , $r_{m,t}$ denotes the return on the market, SMB , HML , RMW , and CMA denote the size, value, profitability, and investment factor, respectively and UMD is the momentum factor. The CAR is the sum of intercepts from the cross-sectional regressions. Panel C reports the monthly average abnormal return (AAR) using the same risk factors for a calendar time approach with portfolios that include stocks that announced a repurchase program in the previous 12, 24, or 36 months. The sample period is from 1995 to 2019. t -statistics are in parentheses. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	Bottom quintile (N=480)	2 nd quintile (N=482)	Middle quintile (N=483)	4 th quintile (N=480)	Top quintile (N=481)
Panel A: Short-term CARs (Event-study approach)					
CAR (± 1 day)	2.36%*** (6.57)	2.12%*** (6.73)	2.24%*** (7.44)	1.59%*** (5.95)	1.50%*** (6.01)
CAR (± 3 days)	1.75%*** (3.52)	2.27%*** (5.28)	2.47%*** (6.41)	1.68%*** (4.26)	1.61%*** (4.54)
CAR (± 5 days)	1.81%*** (3.21)	1.61%*** (3.21)	2.16%*** (5.08)	1.80%*** (3.99)	1.21%*** (3.08)

Table III: continued

	Bottom quintile (N=480)	2 nd quintile (N=482)	Middle quintile (N=483)	4 th quintile (N=480)	Top quintile (N=481)
Panel B: Long-term CARs (IRATS approach)					
CAR (12 months)	10.94%*** (3.12)	3.23% (1.61)	2.89% (1.12)	2.84% (1.09)	-2.68% (-1.08)
CAR (24 months)	19.41%*** (4.10)	0.92% (0.24)	1.20% (0.31)	5.86% (1.44)	-5.42% (-1.61)
CAR (36 months)	25.44%*** (4.09)	5.24% (1.12)	3.19% (0.63)	8.35% (1.63)	-3.13% (-0.67)
Panel C: Long-term AARs (Calendar time approach)					
AAR (12 months)	0.78%*** (2.72)	0.19% (0.76)	0.48%** (1.97)	0.29% (1.10)	-0.27% (-1.18)
AAR (24 months)	0.61%*** (2.61)	-0.12% (-0.65)	0.07% (0.33)	0.22% (1.24)	-0.28% (-1.52)
AAR (36 months)	0.59%*** (2.93)	-0.05% (-0.31)	0.06% (0.40)	0.19% (1.19)	-0.18% (-1.08)

The long-term abnormal returns of relatively overpriced firms (top quintile) are marginally negative and not statistically significant. We confirm this pattern when examining the AARs using the calendar time approach. Firms in the bottom quintile exhibit significant positive abnormal monthly returns of 0.78%, 0.61%, and 0.59% over 12, 24, and 36 months, respectively, while firms that are overpriced have insignificant negative abnormal returns. Using our alternative measure of mispricing, we thus confirm the findings in the literature that the long-term stock performance after repurchase announcements is highest for potentially underpriced stocks (see, e.g., [Peyer and Vermaelen \(2009\)](#)).

B Changes in cash flow and discount rate news

[Michaely et al. \(2020\)](#) analyze averages of cash flow and discount rate news around corporate payout events and find no significant changes. There are, however, two concerns related to the analysis of changes in averages in our setting. First, averages are a function of the VAR residuals, which have zero expectation:

$$\eta_{cf} = \mathbb{E}[\eta_{cf,t}] = \mathbb{E}[(\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}')\mathbf{u}_{i,t}] = (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}')\mathbb{E}[\mathbf{u}_{i,t}] = 0 \quad (7)$$

$$\eta_r = \mathbb{E}[\eta_{r,t}] = \mathbb{E}[\boldsymbol{\lambda}'\mathbf{u}_{i,t}] = \boldsymbol{\lambda}'\mathbb{E}[\mathbf{u}_{i,t}] = 0 \quad (8)$$

Thus, the resulting empirical first moments of cash flow and discount rate news are (close to) zero by construction. Second, averages over many quarters even out any temporary changes in cash flow or discount rate news that only exist around the payout announcement itself. For example, firms could become underpriced in the quarters just prior to repurchase announcements (e.g., [Peyer and Vermaelen \(2009\)](#)). An average over longer periods will fail to capture any temporary dynamics in cash flow and discount rate expectations that are potentially related to *temporary* underpricing.

Therefore, we propose an alternative way to measure changes in cash flows and the discount rate. Specifically, we estimate the change in the levels of cash flow and discount rate news from one quarter before the repurchase announcement to a specific time after the repurchase announcement (i.e., one quarter, one year, two years, or three years). This type of analysis is more likely to capture information about potential transient deviations from fundamentals around repurchase announcements.

Table [IV](#) reports the changes in cash flow and discount rate news from one quarter before to one quarter, one year, two years, and three years after the repurchase announcements for the full sample. In column 1, we observe a small increase in the level of cash flow news in the short term (one quarter). However, over longer horizons, there is no significant change in cash flows news around repurchase announcements. By contrast, the change in discount rate news is negative, statistically significant, and highly persistent until three years after the repurchase announcement (column 2). This result suggests that repurchase announcements are associated with, on average, a decrease in the cost of capital, consistent with [Grullon and Michaely \(2004\)](#).

Table IV: Change in the level of cash flow and discount rate news

This table reports the average change in levels of cash flow and discount rate news for the full sample. The level of news, η_t , is defined as cash flow or discount rate news in equations (3) and (4) for a specific quarter t . We estimate changes from one quarter before the repurchase announcement to one quarter, one year, two years, and three years after the announcement. The VAR is estimated for the entire sample of share repurchases. The sample period is from 1995 to 2019. t -statistics are in parentheses below the coefficient estimates. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

Horizon	Δ Cash flow news	Δ Discount rate news
1 quarter	0.004** (2.36)	-0.004*** (-6.08)
1 year	0.001 (0.55)	-0.004*** (-6.67)
2 years	0.001 (0.67)	-0.004*** (-6.55)
3 years	-0.001 (-0.77)	-0.005*** (-7.85)
Observations	2'417	2'417

In Panel A of Table V, we analyze the changes in cash flow news separately for quintiles based on the firm-specific pricing error. Similar to the full sample, we do not observe any persistent and statistically significant change in cash flow news around repurchase announcements. Only for the bottom quintile, we observe a significant decrease in cash flow news over three years. From Table III we know that firms in the bottom quintile display the highest post-announcement drift. However, a decrease in cash flow news should be associated with lower, and not higher stock returns. A possible explanation of the negative relation between cash flow news and stock returns could be related to agency costs. Specifically, a decrease in cash flows can reduce the free cash flow problem (see Jensen (1986)) and thus lower agency costs, which, in turn, can trigger a positive market response. However, it is not clear why this agency channel should only be present in undervalued firms. If a decrease in cash flow news signals a decrease in agency costs, we should observe a negative association between cash flow news and stock returns in all quintiles. Therefore, although we observe a significant decrease in cash flow news for the bottom quintile, this decrease is hard to reconcile with the large positive post-announcement drift for this set of firms.

Panel B of Table V shows the same analysis for discount rate news. We observe a significant decrease in discount rate news across all quintiles and horizons. The sharpest decrease is in the bottom quintile containing firms with a negative pricing error, ranging between 20 (one quarter) and 90 (3 years) basis points. The decrease in the middle and top quintiles is roughly half the magnitude of the bottom quintile and also statistically significant. The last column of Table V reports the difference in the change of discount rate news between firms in the bottom and top quintile and shows that bottom quintile firms

Table V: Pricing errors and changes in the level of cash flow and discount rate news

This table reports the average change in levels of cash flow and discount rate news for the five pricing error quintiles. The level of news, η_t , is defined as cash flow or discount rate news in equations (3) and (4) for a specific quarter t . We estimate changes from one quarter prior to the repurchase announcement to one quarter, one year, two years, and three years subsequent to the announcement. The VAR is estimated for the entire sample of share repurchases. Panel A shows the changes in the level of cash flow news for quintiles conditional on the firm-specific mispricing measure of Rhodes-Kropf et al. (2005), where quintiles are formed at the time of the repurchase announcement. Panel B shows the changes in the level of discount rate news. The sample period is from 1995 to 2019. t -statistics are in parentheses below the coefficient estimates. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	Bottom quintile (N=484)	2 nd (N=483)	Middle quintile (N=484)	4 th (N=483)	Top quintile (N=483)	Δ 1 - 5
Panel A: Δ Cash flow level and firm-specific pricing error.						
Horizon						
1 quarter	0.008** (2.16)	0.006* (1.93)	0.005* (1.88)	0.000 (0.00)	-0.002 (-0.68)	0.011** (2.05)
1 year	-0.002 (-0.49)	0.000 (0.11)	0.001 (0.23)	0.000 (0.01)	0.005 (1.48)	-0.007 (-1.35)
2 years	-0.006 (-1.52)	0.006* (1.95)	0.005 (1.45)	-0.004 (-1.30)	0.004 (1.14)	-0.010* (-1.89)
3 years	-0.009** (-2.26)	0.000 (0.04)	0.001 (0.40)	-0.001 (-0.45)	0.003 (0.91)	-0.012** (-2.31)
Panel B: Δ Discount rate level and firm-specific pricing error.						
Horizon						
1 quarter	-0.002 (-1.13)	-0.004*** (-2.75)	-0.002 (-1.53)	-0.005*** (-3.48)	-0.006*** (-4.88)	0.005** (2.38)
1 year	-0.006*** (-3.87)	-0.005*** (-3.60)	-0.004*** (-2.99)	-0.003** (-2.16)	-0.003** (-2.13)	-0.003 (-1.60)
2 years	-0.008*** (-4.45)	-0.004** (-2.79)	-0.003** (-2.34)	-0.003** (-2.39)	-0.003** (-2.38)	-0.004* (-1.86)
3 years	-0.009*** (-5.68)	-0.005*** (-3.78)	-0.004*** (-2.86)	-0.004*** (-3.13)	-0.003* (-1.83)	-0.006*** (-3.03)

Table VI: Levels of cash flow and discount rate news around repurchase announcements

This table reports cash flow and discount rate news levels, η_t , before and after the repurchase announcements. The news level before the announcement is defined as the level of cash flow (discount rate) news in the quarter before the repurchase announcement (see equations 3 and 4). The news level after the announcement is defined as the level of cash flow (discount rate) news one year after the announcement. We form quintiles based on the firm-specific pricing error at the time of the repurchase announcement. Stocks belonging to the bottom quintile are likely to be underpriced at the time of the repurchase announcement. Stocks from the top quintile are likely to be overpriced. The sample period is from 1995 to 2019. t -statistics are in parentheses below the coefficient estimates. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

News level (η_t)	Bottom quintile (N=484)		2 nd quintile (N=483)		Middle quintile (N=484)		4 th quintile (N=483)		Top quintile (N=483)	
	η_{pre}	η_{post}	η_{pre}	η_{post}	η_{pre}	η_{post}	η_{pre}	η_{post}	η_{pre}	η_{post}
Cash Flow News	0.008*** (2.76)	0.005* (1.90)	-0.002 (-0.98)	-0.002 (-0.97)	-0.002 (-0.97)	-0.001 (-0.55)	-0.001 (-0.29)	-0.000 (-0.08)	-0.005** (-1.98)	0.000 (0.06)
Discount Rate News	0.008*** (8.01)	0.002 (1.26)	0.005*** (4.94)	-0.000 (-0.06)	0.004*** (4.18)	-0.000 (-0.46)	0.003*** (3.38)	0.000 (0.46)	0.003*** (2.76)	-0.000 (-0.24)

exhibit a significantly stronger decrease in discount rate news compared to top quintile firms.

One concern of this approach could be that firms with negative and positive pricing errors have different dynamics in the exogenous VAR variables. Hence, using the *same* VAR coefficient estimates for all firms across different quintiles may not be appropriate. To address this concern, we estimate separate VARs for different groups based on the firm-specific pricing error.⁸ Table C.IV in Appendix C reports the changes in the level of cash flow and discount rate news around repurchase announcements conditional on the pricing error when estimating separate VARs. The results are not only very similar to those in Table V, but amplify our insight that the reduction in the level of discount rate news is significantly larger for potentially underpriced firms.

Our results so far show that share repurchase announcement strongly affect the cost of capital when a firm is potentially underpriced and much less when it is not. This finding suggests that the market implied cost of capital may be too high for firms with a negative pricing error. The repurchase announcement thus serves as a signal that corrects the cost of capital downwards for these firms.

In a next step, we analyze the levels of cash flow and discount rate news one quarter before and one year after the repurchase announcement for the five quintiles of pricing error. Table VI reports the results. The discount rate news for bottom quintile firms before the announcement is positive and three times larger compared to firms in the top quintile. Furthermore, the level of discount rate news shrinks to zero for all quintiles after the repurchase announcement. These results suggests that before the repurchase announcement, the market uses an abnormally high discount rate for firms with a negative pricing error, possibly after the market overreacts to bad news (see, e.g., Peyer and Vermaelen (2009)). The fact that we observe a much larger decrease in the cost of capital for underpriced compared to overpriced firms suggests that this overreaction is due to an overestimation of the cost of capital.

In Table VII, we extend this analysis to a multivariate setting and estimate the following regression:

$$\eta_{i,t} = \gamma_0 + \gamma_1 \text{Quintile}_t + \mathbf{\Gamma}' \mathbf{X}_t + \varepsilon_{i,t}, \quad (9)$$

where the dependent variable is the discount rate (cash flow) news one quarter before the announcement. *Quintile* denotes the quintile of the firm-specific pricing error. The excluded quintile in the regression is the middle quintile. \mathbf{X}_t is a vector of control variables that includes firm age, market capitalization, the debt-to-asset ratio, and the book-to-market

⁸To estimate different transition matrices and to capture the dynamics that are specific to certain degrees of pricing error, we split our data set into terciles. For this part of the study, i.e., the estimation of the transition matrix $\mathbf{\Gamma}$ separately for groups of different pricing errors, we split the sample into terciles, and not into quintiles. The reason is that we need to make sure that we have sufficient data to precisely estimate the parameters of the transition matrix. The Tables C.I, C.II, and C.III in Appendix C report the VAR coefficient estimates for the three groups of firms.

Table VII: Multivariate analysis of cash flow and discount rate news levels

This table reports the results for the following specification:

$$\eta_{i,t} = \gamma_0 + \gamma_1 \text{Quintile}_t + \mathbf{\Gamma}' \mathbf{X}_t + \varepsilon_{i,t},$$

where $\eta_{i,t}$ is the level of news one quarter prior to the announcement, or the change from one quarter prior to one year after. Quintile_t denotes the quintile based on the firm-specific pricing error at the announcement. The coefficients are estimated relative to the middle quintile. \mathbf{X}_t is a vector of control variables. The level of news, η_t , is defined as in equations (3) and (4) for a specific quarter t . Panel A reports the coefficients for discount rate news while panel B reports the coefficients for cash flow news. The VAR is estimated for the entire sample of share repurchases. The sample period is from 1995 to 2019. Standard error are adjusted for heteroskedasticity and clustered at the year-quarter level. t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

Panel A: Discount rate news			
Bottom quintile	0.0041** (2.45)	0.0035** (1.97)	0.0028* (1.68)
2 nd quintile	0.0016 (1.20)	0.0012 (0.89)	0.0013 (0.97)
4 th quintile	-0.0005 (-0.35)	-0.0003 (-0.25)	-0.0007 (-0.50)
Top quintile	-0.0010 (-0.81)	-0.0013 (-0.92)	-0.0016 (-1.12)
Market Cap		-0.0005 (1.46)	-0.0005 (-1.39)
Age		-0.0012 (-1.17)	-0.0017 (-1.53)
Debt-to-Asset		-0.0032 (-0.96)	-0.0027 (-0.72)
B/M-Ratio		0.0004 (1.26)	0.0002 (0.72)
Constant	0.0039*** (3.45)	0.0066* (1.92)	0.0083** (2.36)
Adj. R-squared	0.01	0.01	0.02
Observations	2417	2413	2409
Industry-fixed Effects	×	×	✓
Year-fixed Effects	×	×	✓

Table VII: continued

Panel B: Cash flow news			
Bottom quintile	0.0105** (2.54)	0.0083* (1.92)	0.0085* (1.98)
2 nd quintile	-0.0001 (-0.03)	-0.0011 (-0.32)	-0.0002 (-0.05)
4 th quintile	0.0016 (0.42)	0.0024 (0.64)	0.0015 (0.41)
Top quintile	-0.0028 (-0.82)	-0.0002 (-0.06)	-0.0008 (-0.20)
Market Cap		-0.0010 (-1.04)	-0.0009 (-0.82)
Age		0.0027 (0.97)	0.0011 (0.38)
Debt-to-Assets		0.0003 (0.05)	0.0085 (0.92)
B/M-Ratio		-0.0006 (-0.70)	-0.0010 (-1.12)
Constant	-0.0023 (-0.71)	-0.0078 (-0.93)	-0.0042 (-0.49)
Adj. R-squared	0.01	0.00	0.01
Observations	2417	2413	2409
Industry-fixed Effects	×	×	✓
Year-fixed Effects	×	×	✓

Table VIII: Proximity to the repurchase announcement date and changes in discount rates

This table reports the change in level of discount rate news from 1,2,3, and 4 quarters prior to the announcement to one year after. The level of news, η_t , is defined as discount rate news in the spirit of equations (3) for a specific quarter t . The VAR is estimated for the entire sample of share repurchases. We form mispricing quintiles at the time of the repurchase announcement using the measure of Rhodes-Kropf et al. (2005). The sample period is from 1995 to 2019. t -statistics are in parentheses below the coefficient estimates. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	Bottom quintile (N=484)	2 nd (N=483)	Middle quintile (N=484)	4 th (N=483)	Top quintile (N=483)	Δ 1 - 5
Proximity						
-1 quarter	-0.006*** (-3.87)	-0.005*** (-3.60)	-0.004*** (-2.99)	-0.003** (-2.16)	-0.003** (-2.13)	-0.003 (-1.60)
-2 quarter	-0.001 (-0.92)	-0.002* (-1.92)	-0.003* (-1.94)	0.005 (0.38)	0.001 (0.63)	-0.002 (-1.11)
-3 quarter	-0.000 (-0.09)	-0.000 (-0.34)	0.000 (0.31)	0.001 (0.79)	0.001 (0.65)	-0.001 (-0.47)
-1 year	0.002 (1.03)	0.000 (0.18)	0.000 (0.38)	0.000 (0.20)	0.002 (1.20)	0.000 (0.04)

ratio. Standard errors are clustered at the year-quarter level. Panel A presents the results for discount rate news. Column 1 shows that discount rate news are significantly higher for firms in the bottom quintile compared to firms in the middle quintile. The coefficient estimate on the bottom quintile dummy is positive and statistically significant. This result is robust to the inclusion of control variables (column 2) and two-digit SIC industry and year fixed effects (column 3). These findings further support the interpretation that in the quarter prior to the repurchase announcement, the market applies a discount rate that is too high for firms with a negative pricing error. For cash flow news (Panel B), the coefficient on the bottom quintile dummy is also positive and significant, suggesting that cash flow news are significantly higher for firms in the bottom quintile. However, as argued before, reductions in cash flow news are an unlikely driver of the post-announcement drift.

An implication of a *temporary* overestimation of the cost of capital due to an overreaction of bad past news is that the reduction in discount rate news should be most pronounced when measured relative to the quarter immediately before the announcement. Hence, if we estimate the change in the level of discount rate news from six months or one year before the announcement, we should observe a smaller change or no change at all as the the market has not yet received and overreacted to the bad news at this point in time. Table VIII reports the changes in the level of discount rate news from 1, 2, 3, and 4 quarters before to one year after the repurchase announcement for the five pricing error quintiles. The change from one quarter before the announcement is statistically significant and highest for the bottom quintile. The further we move away from the repurchase announcement, the smaller is the

change in discount rate news. These results strengthen the interpretation of a temporary overestimation of the cost of capital around repurchase announcements.

Overall, the results in this section suggest that the strong positive post-announcement drift observed among underpriced repurchasing firms is the consequence of a slow correction of the cost of capital. Specifically, the results suggest that the market gradually starts to overestimate the cost of capital in the quarters before the announcement, possibly due to an overreaction to bad news, resulting in an underpriced stock. The decrease in discount rate news triggered by the repurchase announcement implies a correction of the cost of capital. Furthermore, we find no evidence that changes in expectations about cash flows dictate the post-announcement drift, neither for underpriced firms nor for the entire sample.

C Discount rate dynamics

So far, our results establish that bottom quintile firms experience the sharpest decrease in the level of discount rate news, and we conjecture that this decrease maps to a correction of the discount rate itself. In this subsection, we substantiate this interpretation of the results with additional analyses.

In our framework, "news" is the difference between the discount rate predicted by the VAR for some period and the effective discount rate used by the market in that same period. This difference possibly arises due to new information. However, the VAR prediction is a linear combination of past state variables, making it impossible for contemporaneous information to be reflected in the discount rate prediction. If the prediction eventually incorporates the news in subsequent periods, the discount rate news would simply disappear because that information is now reflected in the VAR prediction. In this case, the discount rate does not change but only its components, the VAR prediction and news, do.

To gain a deeper understanding of the discount rate dynamics around repurchase announcements, we compare the discount rate predicted by the VAR model with the discount rate used by the market. Specifically, using the return decomposition, we estimate the expected discount rate by the VAR and the market discount rate, i.e., the effective discount rate used by the market to discount cash flows. We start by rearranging equation (5) and express realized returns as:

$$r_t = \mathbb{E}_{t-1}[r_t] + \eta_{cf,t} - \eta_{r,t}, \quad (10)$$

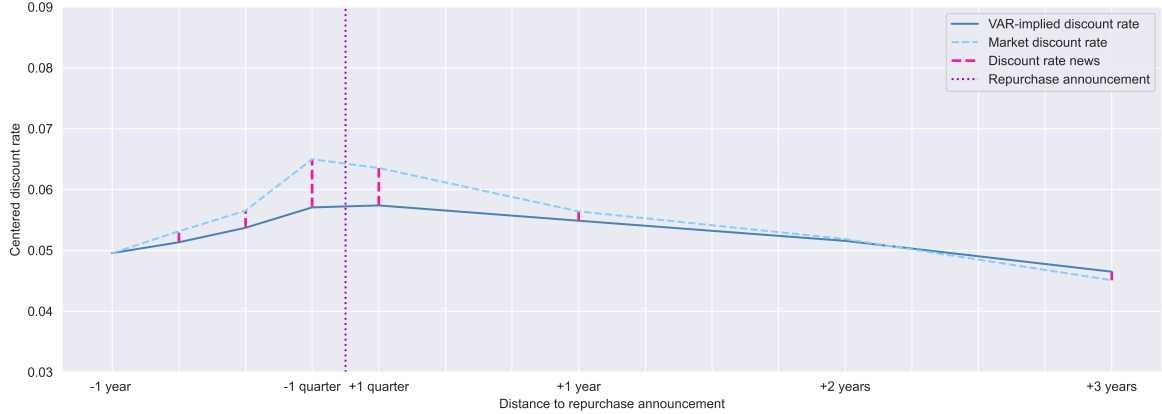
such that it is a function of expected returns and news. To measure discount rate news in equation (10), we decompose the VAR residuals, or unexpected returns, into cash flow and discount rate news. To measure the VAR's prediction of the discount rate, $\hat{\rho}_{i,t}$, we decompose the expected returns of the VAR into a cash flow and discount rate component such that the predicted discount rate maps to:

$$\hat{\rho}_{i,t} = \boldsymbol{\lambda}' \mathbb{E}_{t-1}[\mathbf{r}_{i,t}], \quad (11)$$

where $\mathbb{E}_{t-1}[\mathbf{r}_{i,t}]$ is the vector of expected values under the VAR, with returns in its first

Figure 1: Level of discount rates around repurchase announcements for the bottom quintile

This figure shows the prediction of the centered discount rate by the VAR model, the centered discount rate used by the market, and the level of discount rate news for potentially underpriced firms using the pricing error of Rhodes-Kropf et al. (2005) around repurchase announcements. The discount rate used by the market is the sum of the VAR prediction as defined in equation (11) and discount rate news as defined in equation (3). The VAR is estimated on the full sample of repurchasing firms. The sample period is from 1995 to 2019.



position, for the repurchasing firm i .⁹ The market discount rate then maps to:

$$\rho_{i,t} = \hat{\rho}_{i,t} + \eta_{r_{i,t}}, \quad (12)$$

which is the sum of the predicted discount rate and the level of discount rate news.

If the market learns about an underpricing due to an overestimation of the discount rate, we should observe that the expected discount rate does not decrease significantly after the announcement but the level of discount rate news does. In this case, the overall discount rate used by the market decreases because the discount rate news represent a non-fundamental component that gets corrected as the market learns about the pricing error it makes. Figure 1 shows the evolution of the expected discount rate, $\hat{\rho}_{i,t}$, the market discount rate, $\rho_{i,t}$, and the level of discount rate news, $\eta_{r_{i,t}}$, for firms in the bottom pricing error quintile (underpriced) of our sample.

In the year before the announcement, we observe an increase in the predicted discount rate (solid blue line). It increases up to the quarter before the announcement and then reverts back to its initial level over the next three years. Such dynamics are expected if share repurchases are initiated because of bad news (e.g., Peyer and Vermaelen (2009)). At the same time, the level of discount rate news gradually increases over the year before the announcement, peaking in the quarter just before and then decreasing to zero in the subsequent quarters (dashed red lines). Importantly, the predicted discount rate does not absorb and reflect the level of news from prior quarters. Combining this insight with the results in Table III that show that during the first year after the announcements, firms in

⁹We assume that the proportion of the information regarding cash flows and discount rates is the same in expected returns as it is for unexpected returns.

Table IX: VAR implied levels of discount rates around repurchase announcements

This table reports the VAR-implied levels of the centered discount rate and changes therein for the bottom, middle, and top quintile. To estimate the VAR-implied levels we use equation (3) and substitute the unexpected return (residuals), $u_{i,t}$, with the fitted return from the VAR model, $\hat{r}_{i,t}$. Panel A reports levels while panel B reports changes in levels. The VAR implied discount rate then maps to $\lambda' \hat{r}_{i,t}$. We form quintiles based on the firm-specific pricing error at the time of the repurchase announcement. Stocks belonging to the bottom quintile are likely to be underpriced at the time of the repurchase announcement. Stocks from the top quintile are likely to be overpriced. The sample period is from 1995 to 2019. t -statistics are in parentheses below the coefficient estimates. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

Panel A: VAR implied centered discount rate levels			
	Bottom quintile (N=484)	Middle quintile (N=484)	Top quintile (N=483)
1 quarter prior	0.057	-0.005	-0.069
1 quarter after	0.057	-0.002	-0.059
1 year after	0.055	-0.001	-0.056
2 years after	0.052	-0.003	-0.053
3 years after	0.047	-0.001	-0.051

Panel B: Changes in VAR implied centered discount rate levels relative to one quarter prior to the announcement.			
Horizon	Bottom quintile (N=484)	Middle quintile (N=484)	Top quintile (N=483)
1 quarter	0.000 (0.12)	0.004** (2.38)	0.010*** (5.15)
1 year	-0.002 (-0.74)	0.005** (2.35)	0.012*** (5.33)
2 year	-0.005* (-1.72)	0.002 (1.02)	0.016*** (5.64)
3 year	-0.011*** (-3.04)	0.004 (1.56)	0.018*** (5.94)

the bottom quintile exhibit an abnormal return of 10.94%, our results suggest that this price appreciation likely results from a reversal of the overestimation of the discount rate.

To strengthen this interpretation, Table IX shows the predicted discount rates from the VAR in Panel A and the changes relative to the quarter prior to the announcement in Panel B. For firms in the bottom quintile, there is no significant change observable for the predicted discount rate in the quarter right after the announcement or over a horizon of one year. However, during the same period, there is a significant decrease in the level of discount rate news (see Table V), resulting in a discount rate that decreases. Only after three years, there is a strong significant decrease in the expected discount rate. Put differently, the VAR prediction does not incorporate the discount rate news and remains constant while discount rate news decrease and disappear.

These dynamics of discount rate news and the discount rate itself support our interpretation of the findings. Among underpriced firms, discount rate news in the quarter before the repurchase announcements are of non-fundamental nature, causing an overestimation of the cost of capital. Through the repurchase announcement, market participants learn about their incorrect perception of the cost of capital and correct it. Thus, the positive post-announcement drift for firms in the bottom quintile is likely the consequence of a downward correction of the cost of capital by market participants.

D Variance of cash flows and discount rates

Michaely et al. (2020) show that the variance of cash flow news decreases after repurchase announcements, suggesting a decrease in cash flow risk. This finding is related to the interpretation of the results in Grullon and Michaely (2004), where a decrease in risk drives returns after repurchase announcements. The finding in Michaely et al. (2020), however, raises a question. Do we expect variances of news to be systematically associated with long-term stock returns? Put differently, is the variance of cash flow and discount rate news a priced (systematic) risk? In the return decomposition, the variance captures the *precision*, with which the market assesses expected cash flows and discount rates relative to an equilibrium model. This precision does not necessarily correspond to a source of risk according to asset pricing theory, such as the systematic risk in the Capital Asset Pricing Model. In the return decomposition, systematic risk is captured by the *level* of the discount rate rather than by the volatility of cash flows and discount rates. Thus, we do not expect changes in variances of cash flow and discount rate news to be systematically related to long-term returns.

To support this argument empirically, we analyze changes in the variance of cash flow and discount rate news as a function of the pricing error using the five mispricing quintiles. If a decrease in the variance of news is systematically associated with stock performance, we should observe that underpriced firms experience the largest decrease in the variance of cash flow and discount rate news. We follow Michaely et al. (2020) and estimate the change in variance from before to after the repurchase announcement relative to the cross-sectional average variance prior to repurchase announcements.

Table X: Variances of cash flow and discount rates

This table reports the changes in the variances of cash flow and discount rate news conditional on the firm-specific pricing error using the measure of Rhodes-Kropf et al. (2005). We form quintiles based on the pricing error at the time of the repurchase announcement. Stocks in the bottom quintile are likely underpriced and vice versa. The change in variance is defined as the average change in the variance of news scaled by the cross-sectional average variance prior to the repurchase announcement. Panel A reports changes in second moments for the full sample while panel B reports changes for the mispricing quintiles. The VAR is estimated on the full sample of repurchasing firms. The sample period is from 1995 to 2019. t -statistics are in parentheses. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

Panel A: Full sample (N=2'417)					
Cash flow news			Discount rate news		
$\text{VAR}[\eta_{cf,pre}]$	$\text{VAR}[\eta_{cf,post}]$	$\Delta\text{VAR}[\eta_{cf}]$	$\text{VAR}[\eta_{r,pre}]$	$\text{VAR}[\eta_{r,post}]$	$\Delta\text{VAR}[\eta_r]$
0.0207	0.0170	-0.1776*** (-9.85)	0.0009	0.0008	-0.1133*** (-4.99)

Panel B: Changes in variances and firm-specific pricing errors					
	Bottom quintile (N=484)	2 nd (N=483)	Middle quintile (N=484)	4 th (N=483)	Top quintile (N=483)
Cash flow news	-0.153*** (-3.55)	-0.146*** (-3.59)	-0.119*** (-2.99)	-0.211*** (-5.36)	-0.242*** (-6.25)
Discount rate news	-0.192*** (-2.92)	-0.062 (-1.06)	-0.032 (-0.57)	-0.140** (-2.53)	-0.213*** (-3.45)

Panel A of Table X reports the average change in variances for the entire sample. The variance of cash flows declines by 17.76% compared to the cross-sectional mean of cash flow variance before repurchase announcements. This decline is statistically and economically significant and consistent with the findings in Michaely et al. (2020), who report a decline of 14.79%. Moreover, the variance of the discount rate also declines significantly by 11.33%. This result is new and not reported in Michaely et al. (2020).

Panel B of Table X reports the changes of variances around repurchase announcements for the five mispricing quintiles. Cash flow and discount rate variances decrease for all quintiles. However, there are differences between the top and bottom quintiles, especially for the variances of cash flow news. The variance of cash flow news of top quintile (positive pricing error) firms decreases by 24.2%, while firms in the bottom quintile (negative pricing error) experience a reduction of only 15.3%. For discount rate news, the variance decreases by 21.3% for firms in the top quintile and by 19.2% for firms in the bottom quintile. For firms with little pricing error (middle quintile), the variances of cash flow and discount rate news decrease by 11.9% and 3.2% (not statistically significant), respectively.

These dynamics are inconsistent with the hypothesis that a decrease in variances is responsible for the post-announcement drift. Firms with a positive pricing error experience the highest reductions in cash flow and discount rate variances, yet Table III shows that

their short-term and long-term CARs are indistinguishable from zero and significantly lower compared to the CARs of firms with a negative pricing error. These underpriced firms experience smaller reductions in cash flow news variance and similar reductions in discount rate news variance. Put differently, both set of firms experience similar decreases in the variance of discount rate news, yet display significantly different long-term abnormal returns. Furthermore, a smaller decrease in the variance of cash flow news is associated with a higher post announcement drift, contradicting the idea that a decrease in cash flow volatility drives long-run stock return performance. The implication of these results is that the decrease in cash flow and discount rate variance is unlikely to explain the cross-sectional differences in long-term abnormal returns.

E Firm size and discount rate precision

The previous section suggests that changes in the volatility of news are not systematically related to the post-announcement drift. However, the variance of cash flow and discount rate news may have second-order effects on stock performance through firm size. Existing research shows that the buyback anomaly is mostly pronounced for small firms (see, e.g., [Peyer and Vermaelen \(2009\)](#)). There is generally less information available for smaller firms, and the lack of information potentially leads to more imprecise estimates of discount rate news and exacerbates the overestimation of the cost of capital. Hence, smaller firms could exhibit higher variances than larger firms, and the correction of the cost of capital that we document could be larger for small firms as there is more to correct in the first place.

To understand the influence of firm size on our results, in each year, we divide our sample at the median market capitalization and assign firms to five firm-specific pricing error quintiles separately for small and large firms. Next, we estimate our main tests for small and large firms separately. Table [XI](#) reports the results. Panel A shows that the patterns in long-run performance are comparable to those of the overall sample. Firms in the bottom quintile experience the highest long-term abnormal returns over 1,2, and 3 years, although the returns are only statistically significant for the bottom quintile of small firms. This lack of statistical power likely arises due to the small sample size. The fact that the post-announcement drift is strongest for small undervalued firms is consistent with existing research (e.g., [Peyer and Vermaelen \(2009\)](#)).

Next, we analyze the change in the level of discount rate news separately for small and large firms. Panels B (small firms) and C (large firms) in Table [XI](#) report these changes from one quarter before to one quarter, one year, two years, and three years after the announcement. The results reinforce our main finding for the full sample that, in the long-term, firms with a negative pricing error (bottom quintile) experience a significant and much sharper decrease in discount rate news around repurchase announcements than firms in the top quintile. This finding holds for both small and large firms. The decrease among smaller firms is, on average, more pronounced than for larger firms. This suggest that the market overestimates the discount rate for both small and large firms, but the error is amplified

Table XI: Firm size, discount rates, and cash flows

This table reports the changes in cash flow and discount rate news around repurchase announcements when dividing our sample into sub samples containing firms that are below or above the median market value in given year. We then form our quintiles of firm-specific mispricing at the time of the repurchase announcement using the Rhodes-Kropf et al. (2005) measure within the two size groups separately. Firms belonging to the bottom quintile are likely to be underpriced. Firms from the top quintile are likely to be overpriced. Panel A reports the long-term cumulative abnormal returns from the IRATS approach. Panel B shows the decrease in the level of discount rate news for firms that are below the median market value, i.e., smaller firms. The level of news is defined as discount rate news in the spirit of equation (3) for a specific quarter t . We estimate changes from one quarter prior to the repurchase announcement to one quarter, one year, two years, and three years subsequent to the announcement. Panel C reports the same analysis for larger firms, i.e., firms that are above the median market value in a given year. Panel D reports the change in variances. The change in variance is defined as the average change in the variance of cash flow (discount rate) news scaled by the cross-sectional average variance of cash flow (discount rate) news prior to the repurchase announcement. The VAR is estimated on the full sample of repurchasing firms. The sample period is from 1995 to 2019. t -statistics are in parentheses below the coefficient estimates. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

Panel A: Long-term CARs				
	Small Firms (N=1'219)		Large Firms (N=1'198)	
	Bottom quintile (N=244)	Top quintile (N=243)	Bottom quintile (N=240)	Top quintile (N=239)
CAR (12 months)	11.49%* (1.75)	-0.00% (-0.01)	0.86% (0.20)	0.00% (0.00)
CAR (24 months)	13.20% (1.49)	-2.33% (-0.23)	5.91% (0.97)	-1.10% (-0.16)
CAR (36 months)	21.62%** (2.09)	10.18% (0.77)	5.64% (0.67)	-7.83% (-0.88)

Table XI: continued

Panel B: Δ level of discount rate news for small firms				
Horizon	Bottom quintile (N=244)	Middle quintile (N=244)	Top quintile (N=243)	Δ 1 - 5
1 quarter	-0.003 (-1.07)	-0.002 (-0.93)	-0.008*** (-3.64)	0.006* (1.74)
1 year	-0.006** (-2.27)	-0.004 (-1.56)	-0.006*** (-2.62)	-0.000 (-0.01)
2 years	-0.012*** (-4.70)	-0.005** (-2.41)	-0.008*** (-3.44)	-0.004 (-1.16)
3 years	-0.011*** (-4.32)	-0.005** (-3.44)	-0.006*** (-2.67)	-0.005 (-1.18)

Panel C: Δ level of discount rate news for large firms				
Horizon	Bottom quintile (N=240)	Middle quintile (N=239)	Top quintile (N=239)	Δ 1 - 5
1 quarter	-0.003 (-1.40)	-0.005*** (-2.63)	-0.006*** (-3.14)	0.003 (1.13)
1 year	-0.007*** (-3.22)	-0.003 (-1.56)	0.000 (0.01)	-0.007** (-2.37)
2 years	-0.006*** (-2.72)	-0.000 (-0.23)	-0.002 (-1.11)	-0.004 (-1.22)
3 years	-0.006*** (-2.96)	-0.002 (-1.33)	-0.002 (-0.86)	-0.004 (-1.56)

Panel D: Change in discount rate variance for small and large firms				
	Small Firms (N=1'219)		Large Firms (N=1'198)	
	Bottom quintile (N=244)	Top quintile (N=243)	Bottom quintile (N=240)	Top quintile (N=239)
Δ Cash flow news variance	-0.168*** (-2.61)	-0.335*** (-4.89)	-0.102* (-1.89)	-0.306*** (-5.60)
Δ Discount rate news variance	-0.324** (-3.14)	-0.294*** (-3.02)	-0.020 (-0.23)	-0.253*** (-2.86)

among smaller firms. This result is consistent with our conjecture that cost of capital estimates are less precise for smaller firms. In an unreported test, we find that the variance (precision) of discount rate news of smaller firms is significantly higher (lower) than that of larger firms. Hence, the deviations from the true value are larger among smaller firms, and the overestimation of the cost of capital is amplified and the correction larger.

Panel D of Table XI presents the results for changes in the variance of cash flow and discount rate news. We observe that for both small and large firms, the decrease in the variance of cash flow and discount rate news is stronger for the top quintile of pricing errors. Moreover, the decrease is strongest for the top quintile of small firms. The decrease in cash flow variance is not statistically significant for small firms in the bottom quintile, and the change in discount rate variance is positive yet statistically insignificant for large firms in the bottom quintile. Thus, and further corroborating our main findings, we conclude that the positive post-announcement drift is hard to reconcile with the decrease in cash flow and discount rate variance.

In sum, controlling for size, we find the sharpest decrease in the level of discount rate news among underpriced firms and no systematic relationship between a decrease in variances and long-run returns. Furthermore, the results in this section suggest that smaller underpriced firms exhibit a stronger downward correction of the discount rate than larger underpriced firms, while also displaying higher post-announcement drifts and higher variances of second moments. These findings have at least two implications. First, they reinforce our argument that the post-announcement drift is systematically associated with a decrease in the cost of capital, but not with a decrease in the variance of cash flows or costs of capital. Second, small firms display stronger corrections in the cost of capital and higher post-announcement drifts because their discount rate estimates are less precise before the announcements. As a consequence, the overestimation of the cost of capital documented in the previous sections is amplified when a firm is small, and the downward correction is more pronounced. These results suggest that changes in the level of discount rate news have a direct systematic effect on long-run returns while variances, but not their changes, affect the long-run returns indirectly through firm size in that they aggravate the overestimation due to a lower precision among small firms.

V Conclusion

We analyze changes in cash flows and discount rates around repurchase announcements by applying a return decomposition method from asset pricing. This method provides us with estimates of the average and volatility of cash flow and discount rate news before and after share repurchase announcements. We then analyze the changes in cash flows and discount rates as a function of pricing error at the time of the repurchase announcement.

We find that share repurchase announcements are associated with a significant reduction in the level of discount rates. The sharpest decrease occurs among firms with a negative pricing error, i.e., firms that are potentially underpriced. These are also the firms that

experience the highest long-term returns. These findings suggest that the market learns about the cost of capital for this particular set of firms and is consistent with the idea that repurchase announcements convey information about a temporary overestimation of firms' systematic risk.

We also document that cash flow and discount rate volatility decrease after repurchase announcements. Hence, after a share repurchase announcements, discount rates and cash flows are estimated with more precision. However, the decrease in cash flow and discount rate volatility is smallest for firms that are likely to be underpriced, i.e., among firms that experience the highest initial market response and long-term returns. These results suggest that the decrease in volatility cannot explain the cross-sectional variation in long-term returns. We finally show that discount rate volatility, but not its change, affects the post-announcement drift indirectly through firm size in that it aggravates the overestimation due to a lower precision among small firms.

References

- Bargeron, L., Bonaime, A., Thomas, S., 2017. The timing and source of long-run returns following repurchases. *Journal of Financial and Quantitative Analysis* 52, 491–517.
- Ben-Rephael, A., Oded, J., Wohl, A., 2014. Do firms buy their stock at bargain prices? evidence from actual stock repurchase disclosures. *Review of Finance* 18, 1299–1340.
- Bonaimé, A. A., 2012. Repurchases, reputation, and returns. *Journal of Financial and Quantitative Analysis* 47, 469–491.
- Bonaimé, A. A., Öztekin, O., Warr, R. S., 2014. Capital structure, equity mispricing, and stock repurchases. *Journal of Corporate Finance* 26, 182–200.
- Brav, A., Graham, J. R., Harvey, C. R., Michaely, R., 2005. Payout policy in the 21st century. *Journal of Financial Economics* 77, 483–527.
- Campbell, J. Y., 1991. A variance decomposition for stock returns. *The Economic Journal* 101, 157.
- Campbell, J. Y., Shiller, R. J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1, 195–228.
- Carhart, M. M., 1997. On persistence in mutual fund performance. *The Journal of Finance* 52, 57–82.
- Chan, K., Ikenberry, D., Lee, I., 2004. Economic sources of gain in stock repurchases. *Journal of Financial and Quantitative Analysis* 39, 461–479.
- Chung, D. Y., Isakov, D., Pérignon, C., 2007. Repurchasing shares on a second trading line. *Review of Finance* 11, 253–285.
- Cziraki, P., Lyandres, E., Michaely, R., 2021. What do insiders know? Evidence from insider trading around share repurchases and SEOs. *Journal of Corporate Finance* 66.
- Dann, L. Y., 1981. Common stock repurchases. An analysis of returns to bondholders and stockholders. *Journal of Financial Economics* 9, 113–138.
- Dittmar, A., Field, L. C., 2015. Can managers time the market? Evidence using repurchase price data. *Journal of Financial Economics* 115, 261–282.
- Dittmar, A. K., 2000. Why do firms repurchase stock? *Journal of Business* 73, 331–355.
- D’Mello, R., Shroff, P. K., 2000. Equity undervaluation and decisions related to repurchase tender offers: An empirical investigation. *Journal of Finance* 55, 2399–2424.
- Evgeniou, T., de Fortuny, E. J., Nassuphis, N., Vermaelen, T., 2018. Volatility and the buyback anomaly. *Journal of Corporate Finance* 49, 32–53.
- Fama, E. F., French, K. R., 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116, 1–22.
- Farre-Mensa, J., Michaely, R., Schmalz, M. C., 2014. Payout Policy. *Annual Review of Financial Economics* 6, 75–134.
- Fu, F., Huang, S., 2016. The persistence of long-run abnormal returns following stock repurchases and offerings. *Management Science* 62, 964–984.

- Golubov, A., Konstantinidi, T., 2019. Where is the risk in value? evidence from a market-to-book decomposition. *Journal of Finance* 74, 3135–3186.
- Grullon, G., Michaely, R., 2002. Dividends, share repurchases, and the substitution hypothesis. *Journal of Finance* 57, 1649–1684.
- Grullon, G., Michaely, R., 2004. The information content of share repurchase programs. *Journal of Finance* 59, 651–680.
- Ibbotson, R. G., 1975. Price performance of common stock new issues. *Journal of Financial Economics* 2, 235–272.
- Ikenberry, D., Lakonishok, J., Vermaelen, T., 1995. Market underreaction to open market share repurchases. *Journal of Financial Economics* 39, 181–208.
- Jagannathan, M., Stephens, C. P., Weisbach, M. S., 2000. Financial flexibility and the choice between dividends and stock repurchases. *Journal of Financial Economics* 57, 355–384.
- Jensen, M. C., 1986. Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review* 76, 323–329.
- Kahle, K. M., 2002. When a buyback isn't a buyback: Open market repurchases and employee options. *Journal of Financial Economics* 63, 235–261.
- Kumar, P., Sorescu, S. M., Boehme, R. D., Danielsen, B. R., 2008. Estimation risk, information, and the conditional CAPM: Theory and evidence. *Review of Financial Studies* 21, 1037–1075.
- Lakonishok, J., Vermaelen, T., 1990. Anomalous price behavior around repurchase tender offers. *The Journal of Finance* 45, 455–477.
- Lee, I., Park, Y. J., Pearson, N. D., 2020. Repurchases after being well known as good news. *Journal of Corporate Finance* 62.
- Lin, J. C., Stephens, C. P., Wu, Y. L., 2014. Limited attention, share repurchases, and takeover risk. *Journal of Banking and Finance* 42, 283–301.
- Manconi, A., Peyer, U., Vermaelen, T., 2018. Are buybacks good for long-term shareholder value? Evidence from buybacks around the world. *Journal of Financial and Quantitative Analysis* .
- Michaely, R., Rossi, S., Weber, M., 2020. Signaling safety. Working Paper .
- Oded, J., 2005. Why do firms announce open-market repurchase programs? *Review of Financial Studies* 18, 271–300.
- Peyer, U., Vermaelen, T., 2009. The nature and persistence of buyback anomalies. *Review of Financial Studies* 22, 1693–1745.
- Rhodes-Kropf, M., Robinson, D. T., Viswanathan, S., 2005. Valuation waves and merger activity: The empirical evidence. *Journal of Financial Economics* 77, 561–603.
- Stephens, C. P., Weisbach, M. S., 1998. Actual share reacquisitions in open-market repurchase programs. *Journal of Finance* 53, 313–333.
- Vermaelen, T., 1981. Common stock repurchases and market signalling. *Journal of Financial Economics* 9, 139–183.

Vuolteenaho, T., 2002. What drives firm-level stock returns? *Journal of Finance* 57, 233–264.

Appendix Essay 1

The appendix describes the methods in section III. Part A provides a formal treatment of the return decomposition that we use to estimate the first and second moment of cash flow and discount rate news. Part B discusses the market-to-book decomposition of Rhodes-Kropf et al. (2005). Part C contains additional tables and analyses.

A Return decomposition

The decomposition of Vuolteenaho (2002) draws on the decomposition of the dividend discount model of Campbell and Shiller (1988) but instead of dividend growth the decomposition takes ROE (earnings over book equity) as cash flow fundamental. In order to derive the ROE-based present value model three assumptions must hold. First, book equity, BE , dividends, D , and market equity, ME , are assumed to be strictly positive. Second, earnings, X , dividends, and book equity satisfy the clean surplus identity, that is, earnings equal the change in in book value of equity, ΔB_t , minus dividends. Third, both the difference between the logarithm of book equity, be , and the logarithm of market equity, me , as well as the difference between the logarithm of dividends, d , and the the logarithm of book equity are assumed to be stationary. With these assumptions we can write the logarithm of the book-to-market ratio, θ , as:

$$\theta_{t-1} = k_{t-1} + \sum_{j=0}^{\infty} \rho^j r_{t+j} - \sum_{j=0}^{\infty} \rho^j (e_{t+j} - f_{t+j}), \quad (13)$$

where e denotes the logarithm of ROE defined as $\log(1 + X_t/BE_{t-1})$ and r denotes the logarithm of the excess stock return defined as $\log(1 + R_t + F_t) - f_t$. The simple excess return and discrete interest rate are denoted by R and F , respectively, such that f maps to the logarithm of one plus the discrete interest rate. The discount factor is ρ , and k summarizes linearization constants.

We follow Campbell (1991) to obtain stock return news from changes in expectations from $t - 1$ to t . To that end, we can rewrite equation (13) as an identity for unexpected returns:

$$r_t - \mathbb{E}_{t-1}[r_t] = \Delta \mathbb{E}_t \left[\sum_{j=0}^{\infty} \rho^j (e_{t+j} - f_{t+j}) \right] - \Delta \mathbb{E}_t \left[\sum_{j=1}^{\infty} \rho^j r_{t+j} \right], \quad (14)$$

where $\Delta \mathbb{E}_t$ denotes the change in expectations from $t - 1$ to t . i.e., $\mathbb{E}_t[\cdot] - \mathbb{E}_{t-1}[\cdot]$. We can now use equation (14) and write unexpected returns as the difference between cash flow news and expected return news:

$$r_t - \mathbb{E}_{t-1}[r_t] = \eta_{cf,t} - \eta_{r,t} \quad (15)$$

We follow Vuolteenaho (2002) and Michaely et al. (2020) and use a vector autoregression (VAR) to implement the return and return variance decomposition. Let $\mathbf{z}_{i,t}$ be a vector at time t that contains firm-specific state variables of firm i and contains returns as its first component. Assuming that a first-order VAR is sufficient to describe the evolution of the state variables in $\mathbf{z}_{i,t}$ we can write the VAR system as:

$$\mathbf{z}_{i,t} = \mathbf{\Gamma}\mathbf{z}_{i,t-1} + \mathbf{u}_{i,t}, \quad (16)$$

where $\mathbf{\Gamma}$ is the transition matrix of the VAR system. We further define the vector $\mathbf{e}\mathbf{1}' = [1 \ 0 \ 0]$ and rewrite unexpected stock returns as:

$$r_t - \mathbb{E}_{t-1}[r_t] = \mathbf{e}\mathbf{1}'\mathbf{u}_{i,t} \quad (17)$$

Discount rate news can then be written as:

$$\begin{aligned} \eta_{r,t} &= \mathbf{e}\mathbf{1}' \sum_{j=1}^{\infty} \rho^j \mathbf{\Gamma}^j \mathbf{u}_{i,t+j} \\ &= \mathbf{e}\mathbf{1}' \rho \mathbf{\Gamma} (\mathbb{1} - \rho \mathbf{\Gamma})^{-1} \mathbf{u}_{i,t} \end{aligned} \quad (18)$$

which we, for simplicity, define as:

$$\eta_{r,t} = \boldsymbol{\lambda}' \mathbf{u}_{i,t} \quad (19)$$

where $\mathbb{1}$ is an identity matrix with matching dimensions. We can then write cash flow news as:

$$\eta_{cf,t} = (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}') \mathbf{u}_{i,t} \quad (20)$$

and the variance of cash flows as:

$$\delta(\eta_{cf,t}) = (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}') \boldsymbol{\Sigma} (\mathbf{e}\mathbf{1} + \boldsymbol{\lambda}) \quad (21)$$

where $\boldsymbol{\Sigma}$ denotes the covariance matrix of $\mathbf{u}_{i,t+1}$ which is assumed to be independent of the information set at $t - 1$. Similarly, the variance of discount rate news is:

$$\delta(\eta_{r,t}) = \boldsymbol{\lambda}' \boldsymbol{\Sigma} \boldsymbol{\lambda} \quad (22)$$

and the covariance between the two news components is given by:

$$\Omega(\eta_{cf}, \eta_r, t) = \boldsymbol{\lambda}' \boldsymbol{\Sigma} (\mathbf{e}\mathbf{1} + \boldsymbol{\lambda}) \quad (23)$$

In order to estimate the transition matrix $\mathbf{\Gamma}$ we need an estimate for the discount factor ρ . To that end, we follow Michaely et al. (2020) and estimate ρ as the regression coefficient of the excess log ROE minus the excess log stock return, plus the lagged book-to-market ratio on the book-to-market ratio.

B Market-to-book decomposition

Conceptually, the market-to-book ratio can be decomposed as follows:

$$\text{Market-to-Book} = \text{Market-to-Value} \times \text{Value-to-Book} \quad (24)$$

where *Value* is an estimate of a firm's fundamental value. Using lower case letters to denote logs, we can rewrite the above equation as follows:

$$m - b = (m - v) + (v - b) \quad (25)$$

The first term in this identity, $m - v$, denotes the stock price deviation from fundamental value, while the second term, $v - b$, denotes the difference between fundamental value and book value. If a stock is correctly priced, the term $m - v$ is zero and $m - b = v - b$. However, if it is mispriced, the term $m - v$ is different from zero. Specifically, if the term is positive the stock is above its fundamental value and vice versa. [Rhodes-Kropf et al. \(2005\)](#) estimate v using annual cross-sectional industry-specific regressions of equity values on firm fundamentals. The obtained coefficients map to valuation multiples that account for variation in investors' expectation of stock returns and growth over time and across different industries. The obtained coefficients are averaged over time and used to estimate v based on current firm-specific fundamentals. Due to the time-varying nature of the industry-specific coefficients, stock price deviations from fundamental value, $m - v$, can be dissected into firm-specific deviation from current industry-implied valuation (firm-specific error) and the deviation of current industry-implied valuation from the long-run industry valuation (industry-error):

$$m_{i,t} - b_{i,t} = \underbrace{m_{i,t} - v(\theta_{i,t}; \alpha_{j,t})}_{\text{firm-specific error}} + \underbrace{v(\theta_{i,t}; \alpha_{j,t}) - v(\theta_{i,t}; \alpha_j)}_{\text{industry error}} + \underbrace{v(\theta_{i,t}; \alpha_j) - b_{i,t}}_{\text{Value-to-Book}} \quad (26)$$

where i is an indicator at the firm-level, t indicates the time period, j represent the industry, and where $v(\theta_{i,t}; \alpha_{j,t})$ are the fitted values from the cross-sectional regressions of equity values on firm-specific fundamentals and $v(\theta_{i,t}; \alpha_j)$ map to the estimated long-run industry valuation multiples, i.e., the averaged coefficients over time. Firm-specific and industry error add up to the total error, $m - v$.

We follow the third model in [Rhodes-Kropf et al. \(2005\)](#) to estimate the terms $v(\theta_{i,t}; \alpha_{j,t})$ and $v(\theta_{i,t}; \alpha_j)$ in equation (26) and proceed as follows. We group firms according to the 12 Fama and French industries and run the following cross-sectional regression for each industry and year:

$$m_{i,t} = \alpha_{0,j,t} + \alpha_{1,j,t} b_{i,t} + \alpha_{2,j,t} \ln(NI)_{i,t}^+ + \alpha_{3,j,t} I_{<0} \ln(NI)_{i,t}^+ + \alpha_{4,j,t} LEV_{i,t} + \varepsilon_{i,t}, \quad (27)$$

where $m_{i,t}$ is the logarithm of market equity, $b_{i,t}$ is the logarithm of book equity, $NI_{i,t}^+$ is the absolute value of net income and $I_{<0}$ is an indicator function for negative net income. i , j , and t index firm, industry, and year, respectively. Using the estimated coefficients from equation (27), we construct the fundamental value of the firm as $v(\theta_{i,t}; \alpha_{j,t}) = \hat{\alpha}_{0,j,t} + \hat{\alpha}_{1,j,t} b_{i,t} +$

$\hat{\alpha}_{2j,t} \ln(NI)_{i,t}^+ + \hat{\alpha}_{3j,t} I_{<0} \ln(NI)_{i,t}^+ + \hat{\alpha}_{4j,t} LEV_{i,t}$. This value represents the fundamental value of the firm obtained by applying annual, sector-average regression multiples to firm-level accounting variables. We then also average the alphas over time to get $\bar{\alpha}_j = N^{-1} \sum \hat{\alpha}_{j,t}$ and obtain $v(\theta_{i,t}; \alpha_j) = \bar{\alpha}_{0j} + \bar{\alpha}_{1j} b_{i,t} + \bar{\alpha}_{2j} \ln(NI)_{i,t}^+ + \bar{\alpha}_{3j} I_{<0} \ln(NI)_{i,t}^+ + \bar{\alpha}_{4j} LEV_{i,t}$. This is the fundamental value of the firm obtained by applying long-run industry average multiples to firm-level accounting variables.

Using these two constructed variables, the firm specific error is defined as $m_{i,t} - v(\theta_{i,t}; \alpha_{j,t})$. This error results from firm-specific deviations from valuations implied by sector valuation multiples. We obtain the sector error as $v(\theta_{i,t}; \alpha_{j,t}) - v(\theta_{i,t}; \alpha_j)$. This error captures the deviations of short-run multiples from their long-run average values. Finally, the long-run value to book ratio is defined as $v(\theta_{i,t}; \alpha_j) - b_{i,t}$. It shows the difference between valuations implied by long-run multiples and current book values.

C Additional tables

Table C.I: VAR Transition Matrix - Middle tercile

This table reports the point estimates of a panel VAR using the method outlined in section III. The VAR is estimated for the sample of firms in the middle tercile based on the firm-specific pricing error following Rhodes-Kropf et al. (2005). r_t denotes the centered excess log stock return, θ is the centered log book-to-market ratio, and RoE is the centered log return-on-equity. The sample period is from 1995 to 2019.

	r_{t-1}	θ_{t-1}	RoE_{t-1}
r_t	-0.076 (-10.64)	0.050 (5.13)	0.262 (5.98)
θ_t	-0.204 (-31.87)	0.716 (79.95)	-0.209 (-4.93)
RoE_t	0.015 (7.91)	-0.044 (-15.20)	0.190 (11.34)

Table C.II: VAR Transition Matrix - Bottom tercile

This table reports the point estimates of a panel VAR using the method outlined in section III. The VAR is estimated for the sample of firms in the bottom tercile based on the firm-specific pricing error following Rhodes-Kropf et al. (2005). r_t denotes the centered excess log stock return, θ is the centered log book-to-market ratio, and RoE is the centered log return-on-equity. The sample period is from 1995 to 2019.

	r_{t-1}	θ_{t-1}	RoE_{t-1}
r_t	-0.034 (-4.70)	0.070 (10.06)	0.419 (11.30)
θ_t	-0.242 (-28.31)	0.722 (61.12)	-0.356 (-7.44)
RoE_t	0.013 (7.15)	-0.021 (-11.85)	0.123 (10.68)

Table C.III: VAR Transition Matrix - Top tercile

This table reports the point estimates of a panel VAR using the method outlined in section III. The VAR is estimated for the sample of firms in the top tercile based on the firm-specific pricing error following Rhodes-Kropf et al. (2005). r_t denotes the centered excess log stock return, θ is the centered log book-to-market ratio, and RoE is the centered log return-on-equity. The sample period is from 1995 to 2019.

	r_{t-1}	θ_{t-1}	RoE_{t-1}
r_t	-0.012 (-1.74)	0.009 (1.18)	0.10 (3.60)
θ_t	-0.209 (-30.41)	0.825 (110.18)	-0.108 (-3.89)
RoE_t	0.013 (4.14)	-0.038 (-10.09)	0.247 (12.50)

Table C.IV: Mispricing and changes in the level of cash flow and discount rate news

This table reports the average change in levels of cash flow and discount rate news. The level of news, η_t , is defined as cash flow or discount rate news in the spirit of equations (3) and (4) for a specific quarter t . We estimate changes from one quarter prior to the repurchase announcement to one quarter, one year, two years, and three years subsequent to the announcement. The VAR is estimated for the entire sample of share repurchases. Panel A shows the changes in the level of cash flow news for quintiles conditional on the firm-specific mispricing measure of Rhodes-Kropf et al. (2005), where quintiles are formed at the time of the repurchase announcement. Panel B shows the changes in the level of discount rate news. The sample period is from 1995 to 2019. t -statistics are in parentheses below the coefficient estimates. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	Bottom quintile	2 nd	Middle quintile	4 th	Top quintile	Δ 1 - 5
Panel A: Δ Cash flow level and firm-specific pricing error.						
Horizon						
1 quarter	0.006 (1.14)	0.000 (0.11)	0.006** (2.04)	0.002 (0.70)	-0.000 (-0.13)	0.006 (1.12)
1 year	-0.007 (-1.44)	0.000 (0.09)	-0.003 (-0.86)	0.005* (1.67)	0.004 (1.38)	-0.011** (-2.05)
2 years	-0.012** (-2.28)	0.006 (1.42)	-0.003 (-0.96)	0.004 (1.29)	0.006** (2.02)	-0.017*** (-3.16)
3 years	-0.009* (-1.69)	-0.005 (-1.41)	-0.001 (-0.34)	0.003 (1.08)	0.003 (1.13)	-0.012** (-2.16)
Panel B: Δ Discount rate level and firm-specific pricing error.						
Horizon						
1 quarter	-0.005** (-2.26)	-0.007*** (-3.77)	-0.004*** (-2.87)	-0.001** (-2.01)	-0.003*** (-5.27)	-0.002 (-1.14)
1 year	-0.011*** (-4.55)	-0.008*** (-3.60)	-0.006*** (-4.06)	-0.001 (-1.31)	-0.001** (-2.04)	-0.010*** (-4.42)
2 years	-0.014*** (-5.81)	-0.005*** (-2.63)	-0.006*** (-4.46)	-0.001 (-0.86)	-0.000 (-0.88)	-0.013*** (-6.18)
3 years	-0.012*** (-5.08)	-0.010*** (-4.95)	-0.006*** (-4.35)	-0.002*** (-02.51)	-0.001* (-1.75)	-0.011*** (-5.12)

Essay II:
Skill in the Game*

Sascha Jakob[†]

ABSTRACT

Using the return decomposition of Vuolteenaho (2002) on holdings of 5'357 U.S. equity mutual funds, I examine whether mutual fund managers have the ability to anticipate future changes in aggregate market expectations. I find that the average mutual fund management is able to do so. However, this is only true for expectation changes that are driven by firm-specific information but not for changes driven by systematic information. By showing that asset prices are significantly less complete with respect to firm-specific information than to systematic information, I provide a potential economic mechanism that explains my findings. Further, I show that this skill only exists among fund management structures that include multiple individuals. In effect, team-managed funds may displace their individually managed counterparts as a consequence of competition in the U.S. mutual funds industry. This is consistent with the rapidly shrinking fraction of single-managed funds documented in the literature.

Keywords: Mutual Funds, Skill, Fund performance, Informational efficiency, Behavioral decision making theory

JEL classification Numbers: C40, C58, G14, G17

*I am grateful for the fruitful comments of Lorenzo Bretscher, Julien Cujean, Samuel Jäger, and Philip Valta.

[†]University of Bern, Faculty of Business, Economics and Social Sciences. Engehaldenstrasse 4, 3012 Bern, Switzerland. Email: sascha.jakob@unibe.ch

I Introduction

More than \$35 trillion is managed by U.S. mutual funds by the end of 2021, making them a fundamental part of the average U.S. investor's overall portfolio. In the same year, investors paid a combined \$23 billion in management fees to the U.S. mutual funds industry. A large fraction of this amount is supposed to be a compensation for the fund management's superior skill in allocating capital and generating marginal value to investors. However, under the Efficient Market Hypothesis, it should be impossible for the entire U.S. mutual funds industry to persistently generate marginal value (see e.g., [Fama and French \(2010\)](#)), raising the question whether a fund management simply extracts an economic rent through fees or indeed possess superior skill allowing it to provide marginal value in return (see e.g., [Barras et al. \(2022\)](#); [Berk and van Binsbergen \(2015\)](#); [Berk and Green \(2004\)](#)) .

A considerable effort of research has been devoted to whether mutual fund managers are equipped with superior skills that allow them to consistently outperform the market and provide value to investors. [Treyner and Mazuy \(1966\)](#) and [Jensen \(1968\)](#) are among the first to examine whether actively managed funds outperform the market and arrive at the conclusion that they are not. More specifically, [Gruber \(1996\)](#) shows that the average mutual fund underperforms passive market indices by about 65 basis points per year. Similarly, [Carhart \(1997\)](#) concludes that investors are, on average, better off by investing in passively managed funds as there is no persistence in mutual fund returns. On the other side of the spectrum, there is also compelling evidence that some fund managers are able to consistently outperform the market. For instance, [Brown and Goetzmann \(1995\)](#) conclude that there exists some persistence in fund returns after accounting for survivorship. More recently, [Pastor et al. \(2015\)](#) and [Pastor et al. \(2017\)](#) provide novel evidence that funds can outperform the market.

What these studies have in common is that they confine their understanding of skill to fund outperformance, either before or after fees, i.e., gross or net alpha. However, [Berk and Green \(2004\)](#) argue that the net alpha of a fund is dictated by competition between investors, and not by the skill of managers. Extending this thought, [Berk and van Binsbergen \(2015\)](#) argue that analyzing the gross alpha instead, appears to be flawed as well. This follows from the observation of [Berk and Green \(2004\)](#) where competition drives net alpha to zero. This implies that the gross alpha is equal to the fee the fund charges. Thus, neither net nor gross alpha need be correlated with managerial skill in the cross section. Furthermore, [Chen et al. \(2004\)](#) show that fund performance decreases with fund size. In a rational and efficient market, one could assume that money follows skill such that skilled funds grow in size. As a fund obtains more capital to allocate, at a certain point it may be forced into making conscious suboptimal choices, rendering fund performance lower. In this case, fund performance is a poor indicator of skill.

The main objective of this paper is to provide novel insights on whether fund managers possess skill by using an approach that abandons the notion that skill is equivalent or related to fund performance. To that end, I introduce a novel measure that analyses the

dynamics of a fund management's stock holdings as a function of changes in future market expectations and assesses whether a fund management persistently and correctly anticipates future changes in the market's expectations about the stocks that the fund management holds in its portfolio.

To do so, I employ the return decomposition of [Vuolteenaho \(2002\)](#) to decompose stock returns into changes in aggregate market expectations about a firm's cash flows and its discount rate. I then analyze whether fund managers adjust their holdings in a given period according to changes in aggregate market expectations in the next period. To that end, I estimate the sensitivity between changes in fund holdings and subsequent changes in aggregate market expectations about cash flows and the discount rate for a fund manager in a given quarter. A positive sensitivity indicates that a fund manager adjust holdings in accordance with future expectation changes while a negative sensitivity suggests that the opposite is true. The larger the sensitivity the stronger the adjustment. A stronger adjustment can be because a fund manager either adjusts many of her holdings correctly, the correct adjustment of individual positions is large, or both.

It is apparent that a similar and simpler measure would be to assess whether holdings change in accordance with future abnormal returns. However, [Vuolteenaho \(2002\)](#) finds that changes in expectations about cash flows are driven by firm-specific (idiosyncratic) information while changes in expectations about the discount rate are dictated by market-wide (systematic) information. Using the return decomposition effectively allows me to analyze whether fund managers possess the skill to anticipate changes in expectations driven by firm-specific and/or systematic information. Put differently, it disentangles a fund management's ability to acquire and process firm-specific information from its ability to process systematic information.

By examining fund managers' persistence to anticipate changes in market expectations, I find that fund managers are skilled to anticipate changes that are driven by firm-specific information, suggesting that fund managers possess skill in acquiring and processing idiosyncratic information. However, fund managers are not able to anticipate changes in expectations driven by systematic information. My results are robust with respect to several variations of the estimation approach underlying my skill measure, as well as with respect to errors-in-variables.

Under the Efficient Market Hypothesis, it should not be possible for the entire U.S. equity mutual funds industry to have skill, implying that my results may be driven by a subset of fund managers and pertain to a particular fund management structure only. This naturally raises question which types of fund managements are endowed with the skill to anticipate changes in market expectations. In a competitive market, one would assume that funds with no skill will cease to exist in favor of skilled funds. [Adams et al. \(2018\)](#) documents that the fraction of team-managed funds relative to solo-managed funds increased sharply over time, while I show that the total number of mutual funds remained fairly constant over the last decade. This growing displacement of individually-managed funds by team-managed funds

suggests that latter may have superior skill and should be preferable to investors.

In this spirit, the second objective of this paper is to analyse whether management structures encompassing multiple individuals are more skilled than a single fund manager. I provide novel evidence on the skill differential between individuals and teams and find that fund management structures involving multiple individuals have superior skill. Specifically, team-managed funds are able to persistently anticipate future changes in aggregate market expectations driven by firm-specific information while a fund manager on her own is unable to do so. This implies that team-managed funds are better at processing (firm-specific) information. This finding is well supported by behavioral decision making theory, according to which, groups are better at processing information under uncertainty (see e.g., [Hinsz et al. \(1997\)](#)) and are able to process larger amounts of information (see e.g., [Vollrath et al. \(1989\)](#)).

The finding that skill in the U.S. equity mutual funds industry exists with respect to acquiring and processing idiosyncratic information but not systematic information raises the question why this may be the case. The ability to consistently anticipate changes in expectations implies that there exists some degree of predictability in returns. In efficient markets, returns should, however, be unpredictable (see e.g., [Fama \(1970\)](#), and more recently [Fama and French \(2010\)](#)). On the other side, [Grossman and Stiglitz \(1980\)](#) establish that a competitive economy cannot be in equilibrium and be informationally efficient at all times. Hence, when markets are informationally inefficient, the aggregate market's change of expectations about an asset upon arrival of new information can be flawed. This error in information processing results in prices that do not represent all available information such that returns, and hence changes in aggregate market expectations, may be anticipated by skilled fund managements. Put differently, persistent and correct anticipation of changes in market expectations mandates that a stock be informationally inefficient, else all relevant information would already be incorporated in the price.

In this spirit, the third goal of this paper is to examine whether differences in informational efficiency with respect to firm-specific and systematic information at the asset level are related to a fund management's ability to anticipate changes in aggregate market expectations driven by firm-specific information but not those driven by systematic information. More specifically, I examine whether the informational efficiency with respect to systematic and idiosyncratic information differs in the mutual funds asset universe.

To do so, I split stock returns of fund holdings into a systematic component and an idiosyncratic component using standard asset pricing factor models. I then compute informational efficiency using the entropy measure of [Gao et al. \(2008\)](#), which can be used to measure the completeness of information in and the predictability of a time series, as proposed by [Prado \(2018\)](#). I estimate entropy for both the systematic and idiosyncratic component of stock returns.

I find that the asset universe of U.S. equity mutual funds is informationally less efficient with respect to firm-specific information compared to systematic information, and

this difference is statistically significant. This is consistent with my finding that some fund managements are able to anticipate changes in aggregate market expectations driven by firm-specific information but not those driven by market-wide information. Stocks are informationally less complete with respect to idiosyncratic information, allowing skilled fund managements to anticipate future changes in market expectations.

The contribution of this paper is at least threefold. First, I contribute to the vast literature on skill of mutual fund managers. I show that the average mutual fund management is able to anticipate changes in aggregate market expectations driven by firm-specific information, suggesting they possess skill in acquiring and processing idiosyncratic information. I do so by proposing a holdings based measure. I therefore join the strand of literature that abandons fund performance as a measure of skill and looks at fund holdings instead (see e.g., [Daniel and Titman \(1997\)](#); [Kacperczyk and Seru \(2007\)](#); [Kacperczyk et al. \(2008\)](#); [Cremers and Petajisto \(2009\)](#); [Baker et al. \(2010\)](#); [Kacperczyk et al. \(2014\)](#)).

Second, I enrich the literature devoted to examining the skill difference between team-managed funds and single-managed funds. I show evidence that team-managed funds have superior skill compared to individually managed funds using a holdings based measure of skill. This complements the superior fund performance of team-managed funds documented by [Han et al. \(2017\)](#). On a larger scale, my paper also contributes to the literature devoted to behavioral decision making theory, using an economic setting. I reaffirm the findings of [Hinsz et al. \(1997\)](#) and [Vollrath et al. \(1989\)](#) that teams excel at processing information and make sounder decisions under uncertainty.

Third, to the best of my knowledge, my paper is the first paper to examine whether informational efficiency of holdings is related to managerial skill. By showing that holdings are informationally less efficient with respect to firm-specific information than to systematic information, I provide a potential economic mechanism that explains my findings. I therefore contribute an alternative opinion to [Fama and French \(2010\)](#), who argue that fund outperformance, and hence skill, is not persistent and due to luck, because markets are efficient.

The remainder of this paper is structured as follows. Section [II](#) reviews relevant research and establishes testable hypotheses, section [III](#) outlines the construction of the skill measure, section [IV](#) provides an overview of the data and the variable construction, and sections [V](#) and [VI](#) concern the main analysis. Section [VII](#) concludes the paper.

II Related literature and hypotheses

In this section I review the relevant literature and establish the testable hypotheses. Specifically, I start by reviewing the literature that uses holdings based measures of skill and discuss the two measures that are most closely related to mine. I then proceed with summarising the literature that analyses the performance of team- and single-managed funds.

A Fund manager skill

The empirical evidence on the presence of managerial skill is ample yet inconclusive (see section I). One recurring concern about early studies is the use of fund performance, or fund alpha, as a measure of skill. Some scholars argue that fund performance alone is an inadequate measure of skill because they are driven by competition and size, and not skill (see e.g., [Berk and Green \(2004\)](#), [Berk and van Binsbergen \(2015\)](#) or [Chen et al. \(2004\)](#)).

In order to infer skill from fund performance, [Berk and van Binsbergen \(2015\)](#) suggest that a fund's gross alpha should be multiplied by the fund size to obtain the fund's value-added. As such, value-added is similar to the economic rent of a firm, which is defined as the markup price of its product times the quantity sold. In a recent paper, [Barras et al. \(2022\)](#) extend this idea to additionally include scalability to address the documented diseconomies of scale by [Chen et al. \(2004\)](#) and more recently [Zhu \(2018\)](#). Using their measure, they report that skill exists in the mutual funds industry.

An entirely different approach to measure skill of a fund is to analyse its holdings rather than measures related to a fund's performance. Afterall, [Daniel and Titman \(1997\)](#) and [Grinblatt et al. \(1995\)](#) attribute much of the fund performance to fundamental characteristics of the stocks held by the fund, such as size, profitability, growth expectations, or past performance. Early studies to examine holdings rather than fund performance include [Grinblatt and Titman \(1989\)](#) and [Grinblatt and Titman \(1993\)](#) who conclude that mutual fund managers indeed have the ability to pick stocks that outperform their benchmarks, before any costs are deducted. More specifically, [Wermers \(2000\)](#) finds that funds, on average, hold stocks that outperform the market by 1.3% per year but their net returns underperform by one percent. Of this 2.3% difference 0.7% is due the underperformance of non-stock holdings whereas 1.6% is due to expenses and transaction costs such that managers pick stocks well enough.

One caveat of these measures is that they do not account for the fact that fund holdings can be the result of passive decisions. As a consequence, looking at changes in holdings, i.e., trading, may be a better way to measure managerial skill. [Chen et al. \(2000\)](#) analyze the trades of funds and find that firms for which mutual funds increased their holdings have higher returns compared to those where holdings decreased. By taking the viewpoint of the firm, these findings, however, only provide evidence that the mutual funds industry in aggregate possesses some skill.

A measure of managerial skill which is closely related to mine is proposed by [Kacperczyk and Seru \(2007\)](#). They define skill as the sensitivity of changes in holdings relative to changes in public and private information. The lower the sensitivity with respect to public changes, the more skilled a fund management structure is. Put differently, a manager that uses public information to trade is considered unskilled while one that uses private information is deemed skilled. They find that managers with a low sensitivity have higher fund performance and show that skill is present in the cross-section.

However, the measure of [Kacperczyk and Seru \(2007\)](#) is subject to flaws. As they use the R^2 from the regression of changes in holdings on changes in public information, their measure does not distinguish between a fund manager who trades in accordance with the effect public information has on the return or exactly opposite. Put differently, it does not capture whether fund managers make correct use of the information. My measure is different to that of [Kacperczyk and Seru \(2007\)](#) in at least two ways. First, it captures whether fund managers change holdings in accordance with the effect information has on returns. As such, it measures whether fund managers process and interpret information correctly. Second, I look at the sensitivity with respect to firm-specific and systematic information rather than public and private information. My measure therefore complements that of [Kacperczyk and Seru \(2007\)](#).

A second paper that is closely related is that of [Kacperczyk et al. \(2014\)](#). They propose a new definition of skill as the ability to pick stocks and time the market, and find evidence for stock picking skill in booms and market timing skill in recessions. Specifically, market timing captures whether funds overweight, relative to the market portfolio, high beta stocks when markets are about to rise and underweight them when markets are about to decline (and vice versa for low beta stocks). Stock picking measures whether funds overweight stocks who have high future idiosyncratic returns and underweight stocks with low future idiosyncratic returns. Their measures are variants of [Grinblatt and Titman \(1993\)](#) and [Daniel et al. \(1997\)](#) who distinguish fund performance based on aggregate market returns from that based on the idiosyncratic component of returns. This is similar to my skill measure, which disentangles skill related to processing systematic information from processing idiosyncratic information. However, my measure has at least three notable advantages. First, it differs in that it measures *changes* in holdings and as a consequence, can account for the fact that the current portfolio weights are also the result of passive decisions or exogenous constraints. Second, and closely related, by taking first differences, my measure of skill is *contemporaneous* in that it measures how fund managers react to *new* information. Looking at current asset weights rather than changes therein is *backward-looking*, as weights summarize all *past* decisions from *past* information. Third, I do not use a market model to measure systematic and idiosyncratic returns as in [Kacperczyk et al. \(2014\)](#). Instead, I use a Vector Autoregression to decompose abnormal returns into changes in market expectations using past returns and firm fundamentals as state variables. My measure of the idiosyncratic component is therefore a function of the dynamics at the firm level directly rather than the residual from a market model. Finally, my measure differs along the conceptual dimension as well. It does not measure market timing or stock picking by definition. Instead, it measures whether fund management structures are better than the aggregate market at acquiring and processing firm-specific and systematic information, and better at understanding how this information affects stock returns.

Ultimately, I combine the ideas of [Kacperczyk and Seru \(2007\)](#) and [Kacperczyk et al. \(2014\)](#) while addressing the respective inherent concerns and provide a more robust measure of holdings based skill and novel evidence on the presence of skill among mutual fund

management structures. A first test of the presence of skill maps to testing the following hypothesis:

*$H_{1,null}$: Management structures of U.S. mutual equity funds are **unable** to persistently anticipate changes in aggregate market expectations and adjust their holdings accordingly.*

The corresponding alternative hypothesis is:

*$H_{1,alternative}$: Management structures of U.S. mutual equity funds are **able** to persistently anticipate changes in aggregate market expectations and adjust their holdings accordingly.*

In a second step, I add a layer of granularity by examining whether a potential presence of skill concerns the ability to process firm-specific and/or market-wide. Therefore, I test the following hypotheses:

*$H_{2,null}$: Management structures of U.S. mutual equity funds are **unable** to consistently anticipate changes in aggregate market expectations driven by **firm-specific information**.*

*$H_{3,null}$: Management structures of U.S. mutual equity funds are **unable** to consistently anticipate changes in aggregate market expectations driven by **market-wide information**.*

The corresponding alternative hypotheses are:

*$H_{2,alternative}$: Management structures of U.S. mutual equity funds are **able** to consistently anticipate changes in aggregate market expectations driven by **firm-specific information**.*

*$H_{3,alternative}$: Management structures of U.S. mutual equity funds are **able** to consistently anticipate changes in aggregate market expectations driven by **market-wide information**.*

B Team managed funds

According to behavioural decision making theory, accumulation and interpretation of information is better when made by teams. For instance, [Snizek and Henry \(1989\)](#) compare individual and group decisions and find that groups recall and recognise relevant information better than individuals. [Vollrath et al. \(1989\)](#) finds that groups can recall a larger volume of information than a single individual ever could. When analyzing decisions of mutual fund managements and their asset allocation, one particular dimension of interest is decision making under uncertainty, as capital markets are characterized by risk. [Hinsz et al. \(1997\)](#) find that when the task is complex and completed under high levels of uncertainty, group members tend to pool and integrate their resources and correct each other's errors. Drawing on this notion, a team of fund managers should be able to make better asset allocation decisions in an environment shaped by uncertainty and hence, display higher skill.

In contrast, the classical decision making theory is founded on a rational choice model. Specifically, [Kahneman and Tversky \(1984\)](#) define choice making as a maximization process

incorporating optimal decisions and [Arrow \(1986\)](#) argues that from a classical utility theory perspective, a rational decision should lead to the same maximizing choice and optimal performance outcome, regardless of whether that decision is made by an individual or a team. Thus, one could expect that individual decision makers and group decision makers would not vary in their degree of skill.

As the classical and behavioral decision making theory are at odds concerning the skill of team-managed and solo-managed funds, I test:

***$H_{A,null}$** : Fund management structures consisting of multiple individuals and a single individual can anticipate changes in aggregate market expectations equally well.*

The alternative hypothesis is:

***$H_{A,alternative}$** : There is a difference in the ability to anticipate changes in aggregate market expectations between Fund management structures consisting of multiple individuals and a single individual.*

Empirical evidence on the performance of team managed funds relative to single managed funds is at odds as well. [Chen et al. \(2004\)](#) find that team managed funds underperform funds that are managed by a single fund manager. [Bär et al. \(2011\)](#) provide a potential rationale therefore. They find that teams follow less extreme investment styles, their portfolios are less industry concentrated, and are, as a consequence, less likely to achieve extreme performance outcomes. On the other hand, [Prather and Middleton \(2002\)](#) and [Bliss et al. \(2008\)](#) find that there is no difference between the two, and [Han et al. \(2017\)](#) show evidence that team managed funds perform better than their individually managed counterparts. My paper differs in that it takes the idea of holdings based skill measures and casts it on the question whether team-managed funds and solo-managed funds differ in skill when it comes to processing information and anticipating changes in market expectations. It also adds an additional layer of granularity in that it disentangles decisions made with firm-specific information from decisions made with systematic information, in order to assess whether teams and individuals differ among those dimensions.

III Measuring skill

In this section, I describe my main measure of skill. I start by outlining the underlying return decomposition of [Vuolteenaho \(2002\)](#) and proceed by estimating informed trading for the changes in fund holdings. Using the decomposed returns of holdings and informed trading at the fund manager level, I then construct a skill measure that captures to what extent a fund management structure is able to anticipate changes in aggregate market expectations and engages in informed trading as a consequence.

A Return decomposition

Skill allows fund managers to identify stocks that will perform well in the future. Assuming a dividend discount model, a stock price changes because the market's expectations about the stock's cash flows and/or the discount rate change. If fund managers are better at acquiring and processing information, they should be able to anticipate the aggregate market's changes in expectations (news) about cash flows and the discount rate and rebalance their portfolio accordingly. In this spirit, I measure skill as the sensitivity of changes of fund holdings in the manager's portfolio with respect to the ex-post cash flow and discount rate news of the stocks that were traded. In order to estimate news I use the return decomposition of [Vuolteenaho \(2002\)](#). The decomposition of [Vuolteenaho \(2002\)](#) draws on the decomposition of the dividend discount model of [Campbell and Shiller \(1988\)](#) but instead of dividend growth the decomposition takes ROE (earnings over book equity) as cash flow fundamental. In order to derive the ROE-based present value model three assumptions must hold. First, book equity, BE , dividends, and market equity are assumed to be strictly positive. Second, earnings, X , dividends, and book equity satisfy the clean surplus identity, that is, earnings equal the change in in book value of equity, ΔB_t , minus dividends. Third, both the difference between the logarithm of book equity, and the logarithm of market equity, as well as the difference between the logarithm of dividends, and the the logarithm of book equity are assumed to be stationary. With these assumptions we can write the logarithm of the book-to-market ratio, θ , as:

$$\theta_{t-1} = k_{t-1} + \sum_{j=0}^{\infty} \rho^j r_{t+j} - \sum_{j=0}^{\infty} \rho^j (e_{t+j} - f_{t+j}), \quad (1)$$

where e denotes the logarithm of ROE defined as $\log(1 + X_t/BE_{t-1})$ and r denotes the logarithm of the excess stock return defined as $\log(1 + R_t + F_t) - f_t$. The simple excess return and discrete interest rate are denoted by R and F , respectively, such that f maps to the logarithm of one plus the discrete interest rate. The discount factor is ρ , and k summarizes linearization constants.

I follow [Campbell \(1991\)](#) to obtain stock return news from changes in expectations from $t - 1$ to t . To that end, I rewrite equation (1) as an identity for unexpected returns:

$$r_t - \mathbb{E}_{t-1}[r_t] = \Delta \mathbb{E}_t \left[\sum_{j=0}^{\infty} \rho^j (e_{t+j} - f_{t+j}) \right] - \Delta \mathbb{E}_t \left[\sum_{j=1}^{\infty} \rho^j r_{t+j} \right], \quad (2)$$

where $\Delta \mathbb{E}_t$ denotes the change in expectations from $t - 1$ to t . i.e., $\mathbb{E}_t[\cdot] - \mathbb{E}_{t-1}[\cdot]$. I can now use equation (2) and write unexpected returns as the difference between cash flow news and expected return news:

$$r_t - \mathbb{E}_{t-1}[r_t] = \eta_{cf,t} - \eta_{r,t} \quad (3)$$

I follow [Vuolteenaho \(2002\)](#) and use a vector autoregression (VAR) to implement the return and return variance decomposition. Let $\mathbf{z}_{i,t}$ be a vector at time t that contains firm-specific

state variables of firm i and contains returns as its first component. Assuming that a first-order VAR is sufficient to describe the evolution of the state variables in $\mathbf{z}_{i,t}$, the VAR system can be written as:

$$\mathbf{z}_{i,t} = \mathbf{\Gamma}\mathbf{z}_{i,t-1} + \mathbf{u}_{i,t}, \quad (4)$$

where $\mathbf{\Gamma}$ is the transition matrix of the VAR system. I further define the vector $\mathbf{e}\mathbf{1}' = [1 \ 0 \ 0]$ and rewrite unexpected stock returns as:

$$r_t - \mathbb{E}_{t-1}[r_t] = \mathbf{e}\mathbf{1}'\mathbf{u}_{i,t} \quad (5)$$

Then, discount rate news can be written as:

$$\begin{aligned} \eta_{r,t} &= \mathbf{e}\mathbf{1}' \sum_{j=1}^{\infty} \rho^j \mathbf{\Gamma}^j \mathbf{u}_{i,t+j} \\ &= \mathbf{e}\mathbf{1}' \rho \mathbf{\Gamma} (\mathbf{1} - \rho \mathbf{\Gamma})^{-1} \mathbf{u}_{i,t} \end{aligned} \quad (6)$$

which, for simplicity, is defined as:

$$\eta_{r,t} = \boldsymbol{\lambda}' \mathbf{u}_{i,t} \quad (7)$$

where $\mathbf{1}$ is an identity matrix with matching dimensions. And so, cash flow news becomes:

$$\eta_{cf,t} = (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}') \mathbf{u}_{i,t} \quad (8)$$

In order to estimate the transition matrix $\mathbf{\Gamma}$, I need an estimate for the discount factor ρ . To that end, I estimate ρ as the regression coefficient of the excess log ROE minus the excess log stock return, plus the lagged book-to-market ratio on the book-to-market ratio. Finally, the variances of discount rate news and cash flow news are given by:

$$\text{VAR}[\eta_r] = \boldsymbol{\lambda}' \boldsymbol{\Sigma} \boldsymbol{\lambda} \quad (9)$$

$$\text{VAR}[\eta_{cf}] = (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}') \boldsymbol{\Sigma} (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}), \quad (10)$$

and the covariance between the two news components is given by:

$$\text{COV}[\eta_{cf}, \eta_r] = \boldsymbol{\lambda}' \boldsymbol{\Sigma} (\mathbf{e}\mathbf{1}' + \boldsymbol{\lambda}) \quad (11)$$

where $\boldsymbol{\Sigma}$ denotes the covariance matrix of $\mathbf{u}_{i,t}$. The transition matrix $\mathbf{\Gamma}$ and the covariance matrix of the news components are reported in table C.I in appendix C.

B Informed trading

Next, I need a variable that captures the changes in fund holdings of manager j at time t for each holding i that are due to skill and not explained by other sources. Specifically, the main motivation to adjust portfolio holdings is to generate profits based on skilled informed trading. However, the structure of mutual funds also leads them to trade for other

reasons. Following [Alexander et al. \(2007\)](#), the reasons for uninformed trading are at least threefold. First, inflows and outflows of investor money force mutual fund managers to rebalance their portfolios to control liquidity, so-called liquidity-motivated trading (see e.g., [Chordia \(1996\)](#); [Edelen \(1999\)](#); and [Nanda et al. \(2000\)](#)). Second, fund managers may buy recent winners and sell recent losers to make their portfolio look more attractive to investors, so-called window dressing (see e.g., [Lakonishok et al. \(1991\)](#); and [Musto \(1999\)](#)). Finally, changes in fund holdings can arise from tax optimization strategies where fund managers sell unprofitable assets to realize losses at the cutoff-date at the end of the calendar year and buy them back after the cutoff-date (see e.g., [Gibson et al. \(2000\)](#); and [Huddart and Narayanan \(2002\)](#)).¹

In all of the above cases, a change in fund holdings is not a result of skilled informed trading. I therefore estimate a regression at the manager level of the form:

$$\Delta Holdings_{j,i,t} = \pi_{0,j} + \tau_{t,j} + \pi_{1,j} WindowDressing_{i,t} + \varepsilon_{j,i,t}; \quad \forall j \quad (12)$$

where $WindowDressing_{i,t}$ captures whether a stock is a past winner or loser to account for window-dressing, and $\tau_{t,j}$ is a quarter fixed effect that captures characteristics that are constant across all holdings in a given quarter, namely tax optimization and liquidity-motivated trading. The intercept $\pi_{0,j}$ captures time invariant fund manager characteristics that do not change over time.²

C Skill

The residuals in equation (12), $\varepsilon_{j,i,t}$, capture informed trading of holdings due to manager j 's skill. If fund managers have skill to anticipate changes in expectations about cash flows and discount rates, the unexpected change of fund holdings, $\varepsilon_{j,i,t}$, should be related to future cash flow and discount rate news. Positive cash flow news indicate higher cash flows, which should make a stock more profitable. An increase in discount rate news indicates higher discount rates making a stock less profitable as future cash flows are discounted stronger and so the share price declines. A skilled fund manager changes her holdings accordingly

To infer whether fund managers can potentially anticipate cash flow and discount rate news, I need a measure that captures to what extent changes in holdings are associated with next period's cash flow and discount rate news. A simple measure would be the conditional correlation. This, however, does not reveal how strongly a fund manager's adjustment of holdings is in anticipation of future cash flow and discount rate news. To also capture this dimension of anticipation, I estimate a sensitivity. Specifically, I estimate cross-sectional regressions of unexpected changes in fund holdings on next period's cash flow and discount

¹Generally, the cutoff-date is at the end of the calendar year suggesting that fund managers sell losers in December to reduce taxes and buy them back in January.

²This implicitly also requires that fund manager skill changes over time

rate news for each fund manager over all his fund holdings at a given point in time:

$$\varepsilon_{j,i,t} = \omega_{j,t} + \beta_{j,t,cf}\eta_{cf,i,t+1} + \beta_{j,t,r}(-\eta_{r,i,t+1}); \quad \forall t, j \quad (13)$$

One advantage of this specification is that it disentangles the adjustment due to cash flow news from that due to discount rate news. It is important to note that this regression cannot measure any causal relationship as both exogenous variables are only observed in the next period. Its sole purpose is to determine whether there is a linear relationship between changes in holdings and future news and to capture the extent of that relationship. Nevertheless, to address potential econometric concerns related to the form of the specification in equation (13), I also run the following linear least squares specifications:

$$\eta_{cf,i,t+1} = \omega_{j,t} + \beta_{j,t,cf}\varepsilon_{j,i,t} + \gamma_{j,t,r}(-\eta_{r,i,t+1}) + \nu_{j,i,t}; \quad \forall t, j \quad (14)$$

$$-\eta_{r,i,t+1} = \omega_{j,t} + \beta_{j,t,r}\varepsilon_{j,i,t} + \gamma_{j,t,cf}\eta_{cf,i,t+1} + \nu_{j,i,t}; \quad \forall t, j \quad (15)$$

The beta regression coefficients of the manager-specific cross-sectional regressions in equations (13) to (15) then serve as a proxy for skill of a fund manager in a given quarter. Specifically, a positive coefficient indicates that the fund manager adjusts fund holdings in line with future changes in aggregate market expectations related to cash flows or the discount rate, implying that her anticipation was correct. If a fund manager adjusts her holdings opposite to future expectation changes, her anticipation is flawed. Therefore, . If a fund manager fails to anticipate cash flow and/or discount rate news, the respective beta is negative. Therefore, a straight forward measure of skill is an indicator whether at least one beta is positive:

$$\mathcal{S}_{j,t}^{\beta+} = \mathbf{1}(\mathbf{1}(\beta_{j,t,cf} > 0) + \mathbf{1}(\beta_{j,t,r} > 0) \geq 1), \quad (16)$$

where $\mathbf{1}$ is the indicator function, and $\beta_{j,t,cf}$ and $\beta_{j,t,r}$ denote the regression coefficients for cash flow and discount rate news from the cross-sectional regressions in equations (13) to (15) of fund manager j in quarter t . For each cross-sectional regression, I require at least 10 observations of unexpected changes in fund holdings. Note that the conditional correlation between changes in holdings and cash flow or discount rate news is the same, regardless of whether I estimate the coefficients using equation (13) or equations (14) and (15). Therefore, the signs of the regression coefficients in equation (13) and those in equations (14) and (15) are identical, and so must be the respective outcome of the indicator function.

It is not required that fund managers anticipate cash flow news and discount rate news equally well. In fact, some managers may only be able to anticipate cash flow news while others demonstrate skill at anticipating discount rate news, yet others are capable of both or neither. Insight on which component fund managers are able to anticipate is a novel avenue. [Vuolteenaho \(2002\)](#) finds that cash flow news are driven by firm specific, fundamental information while discount rate news are driven by systematic, market-wide information. Put differently, fund managers that display skill at anticipating cash flow news are likely skilled at acquiring and processing firm-specific, fundamental information, while fund man-

agers that excel at anticipating discount rate news are skilled at acquiring and processing systematic information. On an aggregate level, this fosters understanding of whether potential skill in the mutual funds industry is manifested by superior processing of firm-specific information, systematic information, or both.

In this spirit, I introduce a second layer of skill measures by recording which betas of the regression in equation (13) are positive in case of skill. If both are positive, a fund manager correctly anticipates the market’s changes in expectations of both cash flows and the discount rate. In this case, the fund manager is skilled at processing firm-specific information and systematic signals. This complete skill is defined as:

$$\mathcal{S}_{j,t}^* = \mathbb{1}(\mathbb{1}(\beta_{j,t,cf} > 0) \times \mathbb{1}(\beta_{j,t,r} > 0) = 1) \quad (17)$$

In order to determine whether a fund manager is able to anticipate changes in expectations driven by firm-specific, idiosyncratic information in particular, it is required that the beta related to cash flow news be positive:

$$\mathcal{S}_{j,t}^{cf} = \mathbb{1}(\beta_{j,t,cf} > 0) \quad (18)$$

Similarly, to identify whether a fund manager is specifically able to anticipate changes driven by market-wide, systematic signals, the beta associated with discount rate news is required to be positive:

$$\mathcal{S}_{j,t}^r = \mathbb{1}(\beta_{j,t,r} > 0) \quad (19)$$

One possibility when analyzing whether skill manifests itself in processing firm-specific or systematic information is that latter two might be correlated. Thus, correct anticipation of firm-specific information might partially be due to correctly anticipating systematic information, and vice versa. Indeed, table C.I in appendix C shows that the covariance matrix of changes in market expectations is non-diagonal for the sample, indicating that cash flow and discount rate news are correlated. The regression approach in equations (13) to (15) alleviates this concern by partialling out the effect one has through the other. As such, it yields a measure that disentangles the ability to anticipate firm-specific information from the ability to anticipate systematic information.

IV Data

I consider all U.S. domestic equity mutual funds, i.e., U.S. mutual funds that exclusively allocate their assets into U.S. equity. I obtain all mutual funds data from the CRSP Mutual Funds Database. Security prices are obtained from the Center for Research on Security Prices (CRSP) and fundamental data of securities at the quarterly frequency are obtained from Compustat. The risk-free rate and the risk premia of factor models are obtained from Kenneth French’s website. Reported holdings in the CRSP Mutual Funds database date back to 2003. Thus, the sample period for the analysis is 2003 to 2021. I lay out the data

Table I: Descriptive statistics

This table reports descriptive statistics. Panel A shows the sample characteristics of the mutual funds while panel B reports summary statistics of the variables used in the analysis. Informed trading is defined as in equation (12) and additionally reported for positive and negative changes. Betas pertaining to the anticipation of cash flow and discount rate news are estimated according to equation (13). The sample period is from 2003 to 2021.

Panel A: Sample Characteristics						
	N					
Mutual funds	5'347					
Single managed funds	2'755					
Team managed funds	3'669					
Holdings	16.49 Mio.					
Fund quarters	121'844					
Panel B: Variables	N	Mean	Median	5%	95%	sd.
Informed trading	11.32 Mio.	0.01	-0.02	-1.60	2.05	2.39
Informed trading ⁺	5.28 Mio.	1.01	0.19	0.01	5.97	2.40
Informed trading ⁻	6.05 Mio.	-0.86	-0.26	-3.76	-0.01	2.02
Quarterly fund return	121'844	0.03	0.04	-0.15	0.16	0.09
Total Net Assets	121'844	1'628.29	306.85	15.20	6'330.06	6'311.04
Beta cash flow news	121'844	0.34	0.00	-4.23	4.65	17.13
Beta discount rate news	121'844	-1.85	-0.06	-30.13	26.78	100.73

cleaning steps and variable construction extensively in appendix A and proceed with the sample characteristics.

Panel A of table I reports the characteristics of the final sample. Between 2003 and 2021, there are a total of 5'347 domestic equity funds in the U.S. 3'669 funds are managed by more than one person at some point in the sample period while only 2'755 mutual funds are managed by a single manager at some point in time. The sample covers 16.49 million reported holdings spread across 121'844 fund quarters.


Panel B of table I reports the summary statistics of the most important variables. For 11.32 Mio. reported holdings I observe trading activity during the fund quarter. Put differently, almost 70% of holdings are adjusted during the fund quarter, on average. In 5.28 Mio. cases the trading maps to an increase in the number of shares while in 6.05 Mio. observations the number of shares held decreases. On average, an U.S. Mutual equity fund has Total Net Assets (TNA) of 1.63 Bn. USD and a quarterly return of 3%.

Supplementary, figure 1 shows the evolution of the number of active U.S. domestic equity funds for the sample together with their aggregate assets under management (AUM)³. After the Global Financial Crisis at the end of 2008, the number of active funds is lowest with less than 500 complete fund quarters. This number then quickly grows to a stable level of 2'500 funds as of 2011. A similar pattern can be observed for assets under management. AUM is lowest at a level of less than half a trillion USD after the financial crisis. This number

³Due to only a very small number of funds having complete observations between 2003 and 2005, I show the trend for the years 2006 to 2021

Figure 1: Funds and assets under management

This figure shows the number of domestic U.S. equity mutual funds and their assets under management for the period 2006 to 2021.



grows onwards to an amount of roughly 5 trillion USD in recent years.

V Skill in the U.S. mutual funds industry

In this section I present my main results. I start by analyzing the ability of the U.S. equity mutual funds industry to anticipate changes in aggregate market expectations and how this is related to fund performance. I then analyze whether the ability to anticipate changes is persistent over fund quarters. Finally, I examine whether teams of fund managers are better at anticipating changes in market expectations and make better investment decisions than individual fund managers.

A Anticipation of changes in market expectations

The last two rows of panel B in table I report the summary statistics for the betas obtained from the regression in equation (13), which ultimately proxy for a fund management’s skill. A positive beta indicates that the fund management correctly adjusts the holdings and the magnitude of beta measures the scale of the adjustment. In order to analyze whether fund managers have skill in that they possess the ability to anticipate changes in aggregate market expectations about cash flows and the discount rate, I start by looking at the averages of betas. The sample average of the beta associated with changes in expectations related to cash flows is significantly positive with an average beta of 0.34 (t-statistic: 6.83). This conjectures that, on average, fund managers adjust their fund holdings in accordance with future changes in aggregate market expectations about cash flows. The opposite is true for

Table II: Annualized performance of skilled and unskilled funds

This table reports annualized returns of skilled and unskilled funds. A fund is treated as skilled if the indicator function in equation (17) is one (column 1). A fund is unskilled if the indicator function in equation (16) is zero (column 2). The first row reports annualized quarterly fund returns. The second and third row report annualized quarterly excess returns over the S&P 500 and the CRSP value-weighted index, respectively. The last column reports the performance difference between skilled and unskilled funds. The sample period is 2003 to 2021. t -statistics are in parentheses. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	$\mathcal{S}_{j,t}^* = 1$	$\mathcal{S}_{j,t}^\beta = 0$	Δ
Absolute fund return	14.06%*** (58.20)	11.93%*** (50.29)	2.12%*** (6.26)
Relative to S&P 500	0.34%*** (2.70)	-0.67%*** (-5.56)	1.01%*** (5.78)
Relative to CRSP value-weighted index	-0.90%*** (-7.42)	-1.73%*** (-15.11)	0.83%*** (4.98)

expectations related to the discount rate. The average beta is significantly negative and amounts to -1.85 (t-statistic: -6.42). Hence, fund managers are, on average, not able to anticipate changes in market expectations related to the discount rate.

Both betas display a large standard deviation, especially discount rate news. The standard deviation of the beta associated with cash flow news is 17.13 while that for discount rate news amounts to 100.73. This suggests that the precision with which fund managers anticipate changes in cash flow news is higher than the precision about discount rate news. This is consistent with the conjecture that they can anticipate cash flow news but not discount rate news.

In the spirit of Vuolteenaho (2002), the descriptive statistics about the betas of cash flow and discount rate news suggest that fund managers appear to excel at acquiring and processing firm-specific information, but are inferior with respect to market-wide, systematic information.

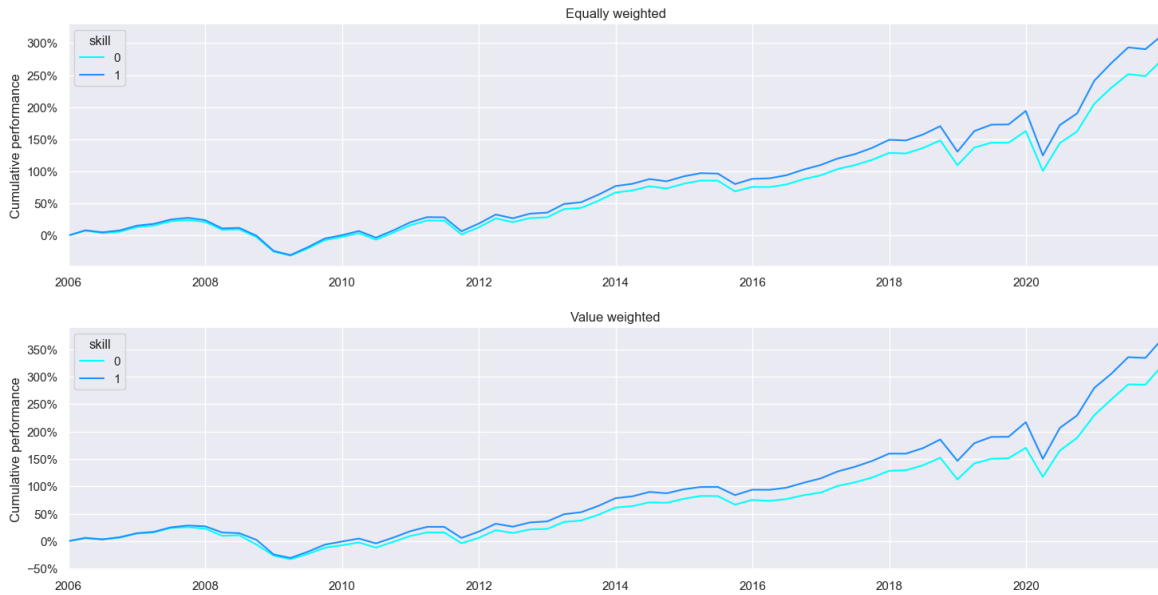
B Skill and fund performance

My skill measure is not based on fund performance. However, it is still desirable for a skill measure to tell apart funds that perform better from those that perform worse. Therefore, I analyze the performance of funds that have no skill and funds that have complete skill as defined in equation (16) and (17).

Table II reports the average annualized fund return for skilled and unskilled fund managements. The average annualized return of skilled fund managers amounts to 14.06% and the average annualized fund return for unskilled fund managers maps to 11.93%. The difference maps to 2.12% and is statistically significant, suggesting that skilled fund managers outperform unskilled fund managers on average. The second row reports the performance

Figure 2: Performance of skilled and unskilled funds

This figure shows the performance of portfolios of domestic U.S. equity mutual funds that have skill and those which do not for the period 2006 to 2021.



relative to the S&P 500. Funds that have complete skill outperform the S&P 500 slightly by 34 basis points. Funds with no skill underperform the index by 67 basis points on average, which would be consistent with the 65 basis points underperformance of funds relative to market indices documented by Gruber (1996). In both cases, the relative performance is statistically significant. Panel C reports the fund performance relative to the value-weighted CRSP index. Both type of funds underperform the index on average. However, unskilled fund managements do significantly worse than skilled fund managements as they perform 83 basis points worse. In conclusion, irrespective of the performance metric, skilled funds statistically outperform unskilled funds. The skill measures that I construct are able to discriminate worse performing funds from better performing funds, based on the fund management’s ability to anticipate changes in aggregate market expectations.

To further demonstrate my skill measure’s ability to identify funds that perform better, I construct two portfolios that contain funds with complete skill and funds with no skill, respectively. Specifically, in each calendar quarter I assign all reporting mutual funds to a portfolio conditional on whether a fund has complete skill as defined in equation (17) or no skill such that the indicator function in (16) equals zero. I compute the portfolio return as the equal or value-weighted average quarterly return of all funds that are in the portfolio in a given quarter. Figure 2 displays the cumulative performance of the portfolios over time. The figure suggests that over the period from 2006 to 2021, skilled fund managements perform better overall. This finding holds impartial of the weighting scheme. In case of an equal-weighted approach, the outperformance of skilled funds relative to unskilled funds appears to be larger.

Table III: Outperformance of skilled fund managers

This table reports the results for the following performance regression:

$$r_{\mathcal{S}_{j,t}^*=1} = \alpha_{skill} + \beta_{unskilled} r_{\mathcal{S}_{j,t}^{\beta^+=0}} + \varepsilon_{\mathcal{S}_{j,t}^*=1}$$

where $r_{\mathcal{S}_{j,t}^*=1}$ denotes the return on a portfolio containing all skilled funds in a given quarter t , and $r_{\mathcal{S}_{j,t}^{\beta^+=0}}$ denotes the return on a portfolio containing unskilled funds in the same month. The intercept of the regression, α_{skill} , maps to the quarterly outperformance of the skilled portfolio relative to the unskilled portfolio. The first column reports the results when funds in the portfolio are equally weighted and the second column weights funds according to their total net assets (TNA). t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	Equally-weighted	Value-weighted
$\beta_{unskilled}$	0.9977*** (115.40)	0.9814*** (48.37)
α_{skill}	0.0022*** (2.91)	0.0030* (1.71)
Observations (quarters)	64	64
R^2	0.99	0.97

In order to quantify the outperformance of the skilled portfolio relative to the unskilled portfolio, I regress quarterly returns of the portfolio containing skilled funds on the quarterly returns of the portfolio containing unskilled funds:

$$r_{s,t} = \alpha_s + \beta_u r_{u,t} + \varepsilon_{s,t}, \quad (20)$$

where $r_{s,t}$ denotes the return on the portfolio containing skilled funds at time t and $r_{u,t}$ represents the return on the portfolio with unskilled funds. The intercept of this performance regression, α_s , maps to a Jensen's alpha and captures the quarterly outperformance of the portfolio containing skilled funds over unskilled funds. Table III reports the results for both equal and value-weighted portfolios. For both weighting schemes, the alpha is positive and significant. In economic terms, the equal-weighted skilled portfolio outperforms the unskilled portfolio by 88 basis points on an annual basis. Using a value-weighted approach, the outperformance amounts to 120 basis points.

The results in this section show that the proposed measure of skill also captures (relative) skill in terms of economic performance. Specifically, mutual fund managers that are able to anticipate both cash flow news and discount rate news outperform their peers who are neither able to anticipate news related to cash flows nor news related to the discount rate.

C Skill persistence

The previous section established that the U.S. equity mutual funds industry in aggregate is able to anticipate changes in market expectations driven by firm-specific information. However, this finding might be the result of luck. Specifically, in every fund quarter a subset of fund managers might be able to anticipate next quarter's changes in market expectations.

Table IV: Skill autocorrelation

This table reports the results for the following OLS and logistic panel autoregression:

$$\mathcal{S}_{j,t}^i = c + \phi \mathcal{S}_{j,t-1}^i + \varepsilon_{j,t}, \quad i \in \{\beta, *, cf, r\}$$

where $\mathcal{S}_{j,t-1}^i$ is the skill indicator in the spirit of equations (16) to (19). Column one reports the results when a fund manager has any form of skill as defined in equation (16). Column 2 reports the result when a manager has full skill as defined in equation (17). Column 3 and 4 report the results for partial skill with respect to cash flow news and discount rate news, respectively, as defined in equations (18) and (19). The sample period is 2003 to 2021. t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	$\mathcal{S}_{j,t}^\beta$	$\mathcal{S}_{j,t-1}^*$	$\mathcal{S}_{j,t}^{cf}$	$\mathcal{S}_{j,t}^r$
Skill autocorrelation (OLS)	0.0050* (1.66)	-0.0017 (-0.57)	-0.0018 (-0.59)	0.0047 (1.55)
Skill autocorrelation (Logit)	0.0315* (1.66)	-0.0116 (-0.57)	-0.0072 (-0.59)	0.0188 (1.55)
Observations	121'844	121'844	121'844	121'844

This subset of managers may change, however, every quarter. For skill to exist at the fund manager level rather than at the industry level, it is required that skill be persistent from one quarter to the next.

To assess whether the quarterly skill indicators in equations (16) to (19) display autocorrelation in the cross-section, I start by running the following regression:

$$\mathcal{S}_{j,t}^i = c + \phi_q \mathcal{S}_{j,t-1}^i + \varepsilon_{j,t}; \quad i \in \{\beta, *, cf, r\}, \quad (21)$$

If managerial skill, or lack thereof, is persistent, the regression coefficients of lagged quarterly skill, ϕ_q , should be positively correlated. Put differently, it measures whether skillful managers tend to have been skillful in previous quarters and whether unskillful managers also lacked skill in previous quarters.

Table IV reports the persistence in skill indicators for fund managers in the entire cross-section. The first column reports the autocorrelation of the overall skill measure as defined in equation (16) using a simple OLS regression as well as a logit regression to accommodate the fact that all variables in the regression are binary variables. The coefficient is significant and positive, indicating that skill and/or lack thereof is persistent from one quarter to the next. When looking at the other indicators in columns 2 to 4, there is no statistically significant autocorrelation observable. One possible issue with this approach is that it fails to disentangle the dynamics of skill from the dynamics of no skill. Specifically, it assumes that both skill and lack of skill are persistent or not persistent at the same time, which clearly need not be the case.

In order to analyze the dynamics of skill and no skill separately, I look at the betas obtained in equation (13) for skilled and unskilled funds separately. I use the betas obtained in equations (14) and (15) to show that my subsequent main findings are robust in section

Table V: Beta Persistency

This table reports the results for the following panel autoregression:

$$\beta_{j,t}^s = c + \phi\beta_{j,t-1}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\}$$

where $\beta_{j,t}^s$ is the regression coefficient obtained in equation (13) for cash flow or discount rate news. The first two columns report the partial autocorrelations in beta associated with cash flow news when a fund management structure is unable (column 1) or able (column 2) to anticipate cash flow news in a given quarter. The last two columns report the partial autocorrelations in beta associated with discount rate news when a fund management structure is unable (column 3) or able (column 4) to anticipate discount rate news in a given quarter. The sample period is 2003 to 2021. t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	$\beta_{j,t}^{cf}$		$\beta_{j,t}^r$	
	$\mathcal{S}_{j,t}^{cf} = 0$	$\mathcal{S}_{j,t}^{cf} = 1$	$\mathcal{S}_{j,t}^r = 0$	$\mathcal{S}_{j,t}^r = 1$
$\beta_{j,t-1}$	-0.0147*** (-3.90)	0.0607*** (13.64)	0.0209*** (4.78)	-0.0314*** (-8.13)
Intercept	-3.8401*** (-59.51)	4.4053*** (57.97)	-25.9842*** (-60.21)	23.8512*** (60.67)
Observations	53'653	54'545	55'936	52'262
Adj. R^2	0.01	0.01	0.01	0.01

V.D. Specifically, I examine whether there is autocorrelation between past and current betas when a fund manager is skilled in the current quarter. I do the same for fund managers that are unskilled in a given quarter. If skill is persistent, both a skilled fund manager's current beta and past betas should be positive and therefore be positively correlated. If skill is not persistent, a skilled fund manager's current beta is positive but not positively correlated with past betas, as they are likely negative. An analogous logic applies when testing the persistence of betas when a fund manager is unskilled in a given quarter.

To that end, I estimate the following regression:

$$\beta_{j,t}^s = c + \phi_q\beta_{j,t-1}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\} \quad (22)$$

where $\beta_{j,t}^s$ denotes the beta obtained in equation (13) for cash flow or discount rate news and $\beta_{j,t-1}^s$ is the lag thereof. I run the analysis separately for anticipation of changes in expectations about cash flows and expectations about the discount rate, and do so separately for skilled and unskilled funds in the spirit of equations (18) and (19). Table V reports the findings.

Column one and two report the persistence of betas for cash flow news while column 3 and 4 report coefficients for discount rate news. When a fund manager has no skill with respect do discount rate news in the current quarter (column 3), her beta is negative and positively correlated with the previous period's beta, suggesting that latter was negative as well. This indicates that fund managers who are not able to anticipate changes in market expectations about the discount rate in a given quarter were also not able to do so in previous

quarters. The beta of managers that are skilled (column 4) shows a negative association with the previous quarter’s beta, suggesting that the ability to anticipate changes in market expectations about the discount rate is not persistent.

Turning to skill related to anticipating changes in expectations about cash flows reveals that skilled fund managers’ betas are positively associated with past betas (column 2), which suggests that positive betas are preceded by betas that are positive as well. This establishes that, at least some, fund managers are able to persistently anticipate changes in market expectations about cash flows. For unskilled managers (columns 1), the relationship is negative, suggesting that lack of skill is not persistent. Interestingly, unskilled managers tend to have a positive cash flow beta in the previous quarter, indicating that in every quarter a subset of previously skilled fund managers becomes unable to persistently anticipate changes in market expectations about cash flows. This is consistent with the notion that managerial skill varies over time, as documented by [Kacperczyk et al. \(2014\)](#).

The persistence of skill might be dictated by characteristics of the fund or the time period. For instance, [Chen et al. \(2004\)](#) find that fund performance decreases with fund size. Following the notion that fund performance is a measure of skill, this would imply that fund managers of smaller funds are more likely to exhibit persistence in skill. Analogously, the results so far might be driven by smaller funds too. Similarly, a fund manager that manages more assets may be less likely to show persistent skill, simply because she needs to anticipate more changes in aggregate market expectations eventually. As a consequence, persistence in skill might be driven by fund managers with few assets to cover. On a different notion, [Kacperczyk et al. \(2014\)](#) and [Leippold and Rueegg \(2020\)](#) find that skill and profitability, respectively, of fund managers are cyclical, indicating that skill might be specific to certain types of market regimes, such that persistence in skill is dictated by certain economic periods. Finally, a fund’s style might also be related to the manager’s ability to consistently anticipate changes in aggregate market expectations. To address these concerns, I examine the persistence in beta in a multivariate setting using the following panel regression:

$$\beta_{j,t}^s = c + \delta_i + \tau_t + \mathbf{\Gamma}' \mathbf{X}_{j,t} + \sum_{q=1}^p \phi_q \beta_{j,t-q}^s + \sum_{q=1}^p \gamma_q \beta_{j,t-q}^s \mathcal{S}_{j,t}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\} \quad (23)$$

where $\beta_{j,t}^s$ denotes fund manager j ’s beta in quarter t with respect to cash flows or the discount rate, $\mathcal{S}_{j,t}^s$ captures skill as defined in equations (18) or (19), δ_i are fund fixed effects, τ_t captures time fixed effects, and \mathbf{X}_t is the control matrix, including the natural log of the total net assets and the natural log of the number of holdings manager j is managing in quarter t . Standard errors are clustered at the fund level.⁴

In this setting, the association between the current beta and the beta q quarters ago is given by $\phi_q + \gamma_q \mathcal{S}_{j,t}^s$. Since $\mathcal{S}_{j,t}^s$ is an indicator variable, the persistence in betas for unskilled fund managers is simply given by ϕ_q while it is $\phi_q + \gamma_q$ for skilled fund managements. Table

⁴I account for a possible errors-in-variables problem in equation (23) in section V.D

VI reports the results. For discount rate news (columns 3 and 4), there are no persistent dynamics present, neither among skilled nor unskilled fund managers. In fact, when a fund manager is skilled in a given quarter, the coefficients associated with past betas are negative, suggesting that past and current betas are negatively correlated. Put differently, a fund manager that is able to anticipate changes in market expectations driven by market-wide information was not able to do so in past quarters. For fund management structures that are unskilled in a given quarter, the current and past betas are uncorrelated.

Columns 1 and 2 in table VI report the relation between the current beta and past betas associated with cash flow news. The results reinforce the findings so far. The betas of fund management structures that are able to anticipate changes in market expectations driven by firm-specific information in a given quarter are positively correlated to past betas. Specifically, for skilled fund managers, past betas tend to be positive, indicating that they were also skilled with respect to idiosyncratic information in previous quarters. These results suggest that fund managers that are able to anticipate how idiosyncratic information affects changes in market expectations can do so with some persistence. As the magnitude of beta indicates how well fund managers adjust their holdings, the results also indicate that those fund managers that adjusted their holdings well in previous periods also adjust well in the current period. In economic terms, for a one standard deviation increase of all past betas associated with cash flow news, the current beta associated with cash flow news increases by 1.90.

For fund management structures that are unable to anticipate changes in expectations driven by firm-specific information, i.e. fund managers with a negative contemporaneous beta, past betas are negatively correlated with the current betas, indicating that they tended to be positive in previous periods. Hence, past betas of both skilled and unskilled fund managers tend to be positive, which seems contradictory at first. However, this is expected for two reasons. First, persistence is only present among skilled fund managers but not unskilled fund managers. Put differently, skilled fund managers tend to stay skilled but unskilled fund managers do not stay unskilled. Rather, a set of fund managers alternates between being able to anticipate changes in expectations in one quarter and being unable to do so in the next. Hence, for some fund management structures, skill is random and is simply luck in the spirit of Fama and French (2010). Also, and closely related, skill is time varying as documented by Kacperczyk et al. (2014). Therefore, among fund management structures that are persistently able to anticipate changes in expectations driven by firm-specific information, a subset becomes unskilled, and vice versa.

Overall, the results in this section suggest that there is some skill in the U.S. equity mutual funds industry. This finding is consistent with previous literature that uses holdings based measures of skill, such as Daniel et al. (1997), Kacperczyk and Seru (2007), Kacperczyk et al. (2014), and Chen et al. (2000). The results also complement the persistence in fund outperformance documented by Grinblatt and Titman (1989), Grinblatt and Titman (1993), Wermers (2000), and Berk and van Binsbergen (2015). Furthermore, skill among

Table VI: Multivariate Beta Analysis

This table reports the results for the following panel regression:

$$\beta_{j,t}^s = c + \delta_i + \tau_t + \mathbf{\Gamma}' \mathbf{X}_{j,t} + \sum_{q=1}^p \phi_q \beta_{j,t-q}^s + \sum_{q=1}^p \gamma_q \beta_{j,t-q}^s \mathcal{S}_{j,t}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\}$$

where δ_i and τ_t are fund and quarter fixed effects, $\mathbf{X}_{j,t}$ are controls, $\beta_{j,t}^s$ denotes the beta associated with cash flow or discount rate news from equation (13), and $\mathcal{S}_{j,t}^s$ represents the skill dummy with respect to cash flow or discount rate news in the sense of equation (18) or (19), respectively. The first two columns report the results for betas related to cash flow news while the last two columns report the results for betas associated with discount rate news. The sample period is 2003 to 2021. t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively. Standard errors are clustered at the fund level.

	$\beta_{j,t}^{cf}$		$\beta_{j,t}^r$	
	1 lag	4 lags	1 lag	4 lags
$\beta_{j,t-1}$	-0.0489** (-2.42)	-0.0438* (-1.89)	-0.0086 (-0.45)	-0.0128 (-0.48)
$\beta_{j,t-2}$		-0.0797*** (-3.77)		-0.0102 (-0.51)
$\beta_{j,t-3}$		-0.0499** (-2.46)		0.0018 (0.07)
$\beta_{j,t-4}$		-0.0317 (-1.48)		0.0018 (0.07)
$\beta_{j,t-1} \times \mathcal{S}_{j,t}^s$	0.0714*** (2.63)	0.0790*** (2.60)	-0.0649** (-2.29)	-0.0600 (-1.63)
$\beta_{j,t-2} \times \mathcal{S}_{j,t}^s$		0.0760** (2.47)		-0.0740** (-2.11)
$\beta_{j,t-3} \times \mathcal{S}_{j,t}^s$		0.0899*** (3.08)		-0.0719** (-2.06)
$\beta_{j,t-4} \times \mathcal{S}_{j,t}^s$		0.0712*** (2.59)		-0.0657* (-1.77)
$\mathcal{S}_{j,t}^s$	8.2577*** (19.77)	8.1812*** (17.45)	50.0672*** (20.70)	49.1597*** (18.21)
ln(Holdings)	0.3328 (1.39)	0.4853 (1.57)	-0.0589 (-0.05)	0.3778 (0.22)
ln(Fund size)	0.2817*** (2.74)	0.3078** (2.23)	-0.5014 (-0.80)	-0.8510 (-0.99)
Intercept	-6.8715*** (-5.73)	-7.6891*** (-4.90)	-23.0190*** (-3.96)	-22.4488*** (-2.60)
Quarter-fixed Effects	✓	✓	✓	✓
Fund-fixed Effects	✓	✓	✓	✓
Observations	108'198	77'388	108'198	77'388
Adj. R^2	0.06	0.07	0.07	0.07

fund management structures is characterised by the ability to persistently anticipate changes in market expectations that are driven by firm-specific information, which is a novel finding. Put differently, fund management structures are skilled at processing and capitalizing idiosyncratic information. The results also suggest that this type of skill is time-varying.

D Robustness

In this section I address potential concerns inherent to my analysis of skill among U.S. equity mutual fund managers presented in the previous section.

Main concerns

So far, I have used the regression coefficients of equation (13) throughout my analysis. Two potential concerns related to this regression are that the exogenous variables are in the future, as noted earlier, and possible omitted variables. Regarding the first concern, the sole purpose of this regression is to measure a potential linear relationship between changes in holdings and future news. It does not aim at establishing any causal effect on its own. To show that my results are not driven by this specification, I reformulate the regression in a more conventional way in equations (14) and (15) and redo my main analysis pertaining to equation (23).

Regarding possible omitted variables, changes in holdings might partly also be the result of past and current changes in aggregate market expectations. Therefore, I extend the skill regression in equation (13) and augment it with past and contemporaneous changes in aggregate market expectations about cash flows and the discount rate:

$$\begin{aligned} \varepsilon_{j,i,t} = & \omega_{j,t} + \beta_{j,t,cf}\eta_{cf,i,t+1} + \beta_{j,t,r}(-\eta_{r,i,t+1}) \\ & + \gamma_{j,t,cf}\eta_{cf,i,t} + \gamma_{j,t,r}(-\eta_{r,i,t}) + \phi_{j,t,cf}\eta_{cf,i,t-1} + \phi_{j,t,r}(-\eta_{r,i,t-1}) + \nu_{j,i,t}; \quad \forall t, j \end{aligned} \quad (24)$$

Furthermore, the estimate for changes in holdings, $\varepsilon_{j,i,t}$, accounts for other variables through equation (12). Table VII reports the results. The first two columns report the findings for the reversed regression setting and the last two columns report the coefficients when current and lagged cash flow and discount rate news are added. The previous findings are robust and remain unchanged in all cases. There is evidence that fund management structures possess the ability to anticipate expectation changes related to firm-specific information but not those related to information that is systematic.

Another potential problem is that of errors-in-variables in the regression in equation (23) induced through the estimation of betas in equation (13). In linear regression models, errors-in-variables can lead to attenuation bias in regression coefficients, i.e., bias towards zero, when independent variables are measured with error, and it can increase standard errors of coefficients if there is measurement error in the dependent variable. Clearly, the betas and their lags in equation (23) might be subject to an errors-in-variables problem.

Table VII: Reverted skill regression and lagged news

This table reports the results for the following panel regression:

$$\beta_{j,t}^s = c + \delta_i + \tau_t + \mathbf{\Gamma}' \mathbf{X}_{j,t} + \sum_{q=1}^4 \phi_q \beta_{j,t-q}^s + \sum_{q=1}^4 \gamma_q \beta_{j,t-q}^s \mathcal{S}_{j,t}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\}$$

where δ_i and τ_t are fund and quarter fixed effects, $\mathbf{X}_{j,t}$ are controls, $\beta_{j,t}^s$ denotes the beta associated with cash flow or discount rate news from equations (13) to (15), and $\mathcal{S}_{j,t}^s$ represents the skill dummy with respect to cash flow or discount rate news in the sense of equation (18) or (19), respectively. The first two columns report the results when betas are estimated using the reverted regressions in equations (14) and (15). Columns 3 and 4 report the results when skill in equation (13) is estimated with current and past cash flow and discount rate news as additional independent variables. The sample period is 2003 to 2021. t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively. Standard errors are clustered at the fund level.

	Reverted skill regression		Lagged news	
	$\beta_{j,t}^{cf}$	$\beta_{j,t}^r$	$\beta_{j,t}^{cf}$	$\beta_{j,t}^r$
$\beta_{j,t-1}$	-0.0112*** (-2.00)	-0.0036 (-0.49)	-0.0328 *** (-3.76)	-0.0130 (-1.09)
$\beta_{j,t-2}$	-0.0103* (-1.77)	-0.0023 (-0.30)	-0.0133 (-1.56)	-0.0054 (-0.43)
$\beta_{j,t-3}$	-0.0130** (-2.22)	-0.0007 (-0.09)	-0.0267*** (-3.29)	-0.0103 (-0.99)
$\beta_{j,t-4}$	-0.0052 (-0.87)	-0.0069 (-0.88)	-0.0042 (-0.53)	-0.0186* (-1.77)
$\beta_{j,t-1} \times \mathcal{S}_{j,t}^s$	0.0339*** (3.40)	-0.0624*** (-6.11)	0.0639*** (4.62)	-0.0634*** (-3.68)
$\beta_{j,t-2} \times \mathcal{S}_{j,t}^s$	0.0299*** (3.20)	-0.0490*** (-4.80)	0.0417*** (3.01)	-0.0549*** (-3.28)
$\beta_{j,t-3} \times \mathcal{S}_{j,t}^s$	0.0236** (2.44)	-0.0541*** (-5.49)	0.0399*** (3.00)	-0.0493*** (-3.16)
$\beta_{j,t-4} \times \mathcal{S}_{j,t}^s$	0.0266*** (2.83)	-0.0420*** (-4.33)	0.0319** (2.26)	-0.0496*** (-3.18)
$\mathcal{S}_{j,t}^s$	0.1198*** (69.88)	0.0172*** (66.79)	1.3664*** (54.06)	9.3144*** (51.84)
ln(Holdings)	-0.0029*** (-7.76)	0.0002 (1.17)	-0.0231*** (-5.02)	-0.0737 (-0.56)
ln(Fund size)	-0.0002 (-1.13)	0.0000 (0.06)	-0.0007 (-0.30)	-0.0163 (-0.22)
Intercept	-0.0433*** (-21.08)	-0.0099*** (-10.89)	-0.5531*** (-23.09)	-4.4123*** (-6.29)
Quarter-fixed Effects	✓	✓	✓	✓
Fund-fixed Effects	✓	✓	✓	✓
Observations	71'658	71'658	70'902	70'902
Adj. R^2	0.34	0.34	0.09	0.09

Table VIII: Erikson-Whited errors-in-variables regression

This table reports the Erikson-Whited estimators for the following panel regression:

$$\beta_{j,t}^s = c + \delta_i + \tau_t + \mathbf{\Gamma}' \mathbf{X}_{j,t} + \sum_{q=1}^4 \phi_q \beta_{j,t-q}^s + \sum_{q=1}^4 \gamma_q \beta_{j,t-q}^s \mathcal{S}_{j,t}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\}$$

using the approach of [Erickson and Whited \(2002\)](#) and more recently [Erickson et al. \(2014\)](#) to address a potential errors-in-variables problem due to the foregoing estimation of $\beta_{j,t}$. I estimate the EIV-regression of [Erickson and Whited \(2002\)](#) specifying that all four lags of $\beta_{j,t}$ are potentially mismeasured and set the highest order of moments to 6. δ_i and τ_t are fund and quarter fixed effects, $\mathbf{X}_{j,t}$ are controls, $\beta_{j,t}^s$ denotes the beta associated with cash flow or discount rate news from equation (13), and $\mathcal{S}_{j,t}^s$ represents the skill dummy with respect to cash flow or discount rate news in the sense of equation (18) or (19), respectively. The sample period is 2003 to 2021. Standard errors are estimated using the bootstrap approach of [Erickson and Whited \(2002\)](#) and the corresponding z -scores are in parentheses. Estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively. The goodness-of-fit of the model is denoted by ρ^2 .

	Cash flow news	Discount rate news
$\beta_{j,t-1}$	-0.0112*** (-269.40)	0.0347*** (253.18)
$\beta_{j,t-2}$	-0.0384** (-2.29)	0.0363** (2.05)
$\beta_{j,t-3}$	-0.0233 (-1.46)	0.0387 (1.64)
$\beta_{j,t-4}$	-0.0092 (-0.53)	0.0242 (1.18)
$\beta_{j,t-1} \times \mathcal{S}_{j,t}^s$	0.0806*** (3.84)	-0.0640*** (-3.22)
$\beta_{j,t-2} \times \mathcal{S}_{j,t}^s$	0.0718*** (2.83)	-0.0593** (-2.01)
$\beta_{j,t-3} \times \mathcal{S}_{j,t}^s$	0.0883*** (3.44)	-0.0706** (-2.25)
$\beta_{j,t-4} \times \mathcal{S}_{j,t}^s$	0.0517** (2.10)	-0.0497* (-1.73)
$\mathcal{S}_{j,t}^s$	5.9959*** (1147.80)	37.7388*** (1078.28)
ln(Holdings)	-0.0369*** (-11.42)	-0.6012*** (-33.17)
ln(Fund size)	0.1402*** (94.78)	0.8293*** (122.15)
Intercept	-3.4016*** (-67.95)	-22.4962*** (-60.29)
Quarter-fixed Effects	✓	✓
Fund-fixed Effects	✓	✓
Observations	77'388	77'388
ρ^2	0.05	0.05

I address this concern by using the errors-in-variables regression of [Erickson and Whited \(2002\)](#) and [Erickson et al. \(2014\)](#) and redo the main analysis pertaining to equation (23).

The results are reported in table VIII. A potential errors-in-variables problem does not drive my findings. The results reinforce my main finding that fund managers can anticipate changes in aggregate market expectations driven by firm-specific information but not those driven by systematic information.

Minor concerns

Another concern could be that my results so far are driven by management structures that are able to anticipate both cash flow and discount rate news at the same time in the sense of equation (17) and are therefore exceptionally skillful. In order to rule out that this subset of fund managers drives the findings, I rerun the analysis and constrain my sample to funds that only have partial skill but not complete skill. Formally, I redefine the partial skill measures in equations (18) and (19) as:

$$\mathcal{S}_{j,t}^{cf} = \mathbb{1}(\beta_{j,t,cf} > 0 \mid \beta_{j,t,r} < 0) \quad (25)$$

$$\mathcal{S}_{j,t}^r = \mathbb{1}(\beta_{j,t,r} > 0 \mid \beta_{j,t,cf} < 0) \quad (26)$$

A further concern of the measure established by [Kacperczyk and Seru \(2007\)](#) is that it considers a fund with few or no changes in holdings as skilled because it is not sensitive to changes in public information as a result. While I do not look at the sensitivity related to public and private information, I measure the sensitivity of the changes in fund holdings with respect to future changes in market expectations. Hence, a similar argument could be made for my case. Specifically, I treat the decision not to change holdings as a conscious choice of the fund management, i.e., the decision not to trade is assumed to be informed trading driven by skill. By construction, a fund that engages in no trading at all has a beta of zero and is unskilled. A fund that only trades one holding, and in accordance with next quarter's changes in expectations, has a positive beta and is considered skilled. This fund's beta, however, is very close to zero due to the funds inaction. Hence, funds that are persistently inactive may have persistently low betas by making only one or a few correct trades in a quarter. Put differently, the persistence in beta documented so far may be driven by funds that remain mostly inactive. To examine whether persistence in betas is driven by inactive equity mutual funds, I estimate fund manager skill in equation (13) using non-zero informed trading only and require at least 10 changes in holdings for a fund manager in a given quarter and redo my analysis.

Table C.II in appendix C reports the results for the persistence in beta when only considering partially skilled funds (columns 1 and 2) and when estimating skill using non-zero informed trading only (columns 3 and 4). The findings remain robust. Skilled fund management structures have the ability to persistently anticipate changes in expectations driven by firm-specific information. This result holds when discarding funds that are universally skilled and looking at funds that are skilled with respect to firm-specific information only.

The results further hold when skill is measured based on active trading only, such that the decision not to trade does not affect skill.

E Team-managed funds

During the sample period, 70% of all U.S. mutual equity funds are managed by a team at some point. [Adams et al. \(2018\)](#) document a large increase in the fraction of team-managed funds in recent years, yet [figure 1](#) reminds us that the total number of U.S. mutual equity funds remained fairly constant during the last decade. Consequently, an increase in the fraction of team-managed funds seems to imply that the number of funds managed by individuals must decrease. Drawing on the notion of [Berk and Green \(2004\)](#) where fund performance is driven by competition, this decrease in single-managed funds may be due to the fact that they are outcompeted by their team-managed counterparts. Specifically, a group of individuals managing a fund may have higher skill than an individual managing a fund. If competition is sufficiently strong, funds managed by individuals will cease to exist as a consequence. This raises the question whether teams of fund managers have superior skill in processing information and anticipating changes in aggregate market expectations compared to a single fund manager.

In behavioral decision making theory, various studies suggest that teams are better at processing information, yet empirical evidence on the performance of team-managed funds relative to single-managed funds is mixed, as discussed in [section II.B](#). In my sample, team-managed funds have an average quarterly return of 3.08% while single-managed funds achieve a quarterly return of 3.37%. This difference is statistically significant and hence, when looking at fund performance, team-managed funds seem to underperform.

As documented earlier, there are several reasons as to why measuring the skill of a fund management using its fund's performance may be flawed. First, skill is dictated by competition and fund size, as postulated by [Berk and Green \(2004\)](#), and [Berk and van Binsbergen \(2015\)](#). More specifically, [Chen et al. \(2004\)](#) find that fund performance decreases with fund size. Therefore, skill simply decreases as funds turn larger. In my sample, the average TNA of single-managed funds amounts to 1.4 billion USD. Meanwhile, team-managed funds are significantly larger and have total net assets of 1.8 billion USD, on average. Put differently, team-managed funds are almost 30% larger, on average. Combining this with the insight of [Chen et al. \(2004\)](#) suggests that team-managed funds may perform worse because they are significantly larger, on average, but are not necessarily less skilled. Second, in efficient markets one would assume that capital flows into skilled funds, making them ultimately larger than unskilled funds. The fact that team-managed funds are significantly larger implies that team-managed funds should be more skilled, or else there would be a severe misallocation of capital by market participants. Third, the fact that 70% of funds in my sample are team-managed, combined with the finding of [Adams et al. \(2018\)](#) that the fraction of team-managed funds increases, also seems to imply that fund companies find this structure of fund management preferable.

Table IX: Team-managed vs. single-managed funds

This table reports the results for the following panel regression:

$$\beta_{j,t}^s = c + \delta_i + \tau_t + \mathbf{\Gamma}' \mathbf{X}_{j,t} + \sum_{q=1}^4 \phi_q \beta_{j,t-q}^s + \sum_{q=1}^4 \gamma_q \beta_{j,t-q}^s \mathcal{S}_{j,t}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\}$$

where δ_i and τ_t are fund and quarter fixed effects, $\mathbf{X}_{j,t}$ are controls, $\beta_{j,t}^s$ denotes the beta associated with cash flow or discount rate news from equation (13), and $\mathcal{S}_{j,t}^s$ represents the skill dummy with respect to cash flow or discount rate news in the sense of equation (18) or (19), respectively. The first two columns report the results for a sample that only includes fund management structures consisting of a single individual. Columns 3 and 4 report the results for a sample that only includes observations for fund management structures consisting of multiple individuals. The sample period is 2003 to 2021. t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively. Standard errors are clustered at the fund level.

	Single-managed		Team-managed	
	$\beta_{j,t}^{cf}$	$\beta_{j,t}^r$	$\beta_{j,t}^{cf}$	$\beta_{j,t}^r$
$\beta_{j,t-1}$	-0.0538 (-1.42)	-0.0357 (-0.82)	-0.0546* (-1.79)	-0.0062 (-0.20)
$\beta_{j,t-2}$	-0.0883** (-2.48)	0.0059 (0.19)	-0.0869*** (-3.38)	-0.0303 (-1.21)
$\beta_{j,t-3}$	-0.0168 (-0.59)	0.0212 (0.50)	-0.0967*** (-3.19)	-0.0207 (-0.82)
$\beta_{j,t-4}$	-0.0302 (-0.97)	-0.0143 (-0.36)	-0.0472 (-1.56)	0.0140 (0.43)
$\beta_{j,t-1} \times \mathcal{S}_{j,t}^s$	0.0700 (1.45)	-0.0540 (-0.84)	0.0850** (2.01)	-0.0597 (-1.61)
$\beta_{j,t-2} \times \mathcal{S}_{j,t}^s$	0.0897* (1.82)	-0.0962 (-1.62)	0.0604 (1.55)	-0.0596 (-1.49)
$\beta_{j,t-3} \times \mathcal{S}_{j,t}^s$	0.0712 (1.57)	-0.0601 (-1.01)	0.1088*** (2.82)	-0.0852** (-2.22)
$\beta_{j,t-4} \times \mathcal{S}_{j,t}^s$	0.0355 (0.83)	-0.0649 (-1.05)	0.1037*** (2.82)	-0.0663 (-1.54)
$\mathcal{S}_{j,t}^s$	8.4783*** (10.71)	51.0600*** (11.27)	7.8453*** (14.81)	47.4140*** (15.16)
ln(Holdings)	0.0555 (0.12)	1.2066 (0.37)	0.6726 (1.61)	-1.2289 (-0.73)
ln(Fund size)	0.3516 (1.26)	0.4709 (0.23)	0.2266 (1.44)	-1.1848 (-1.35)
Intercept	-6.2873** (-2.44)	-34.2801* (-1.86)	-7.8065*** (-3.65)	-13.0564 (-1.52)
Quarter-fixed Effects	✓	✓	✓	✓
Fund-fixed Effects	✓	✓	✓	✓
Observations	35'345	35'345	42'043	42'043
Adj. R^2	0.07	0.07	0.07	0.08

To provide new evidence on whether team- and single-managed funds differ in skill, I approach this question with the holdings based skill measures used in the previous sections. To that end, I split my sample into a subsample containing team-managed funds and a subsample containing single-managed funds, and rerun the analysis in equation (23). Table IX reports the results.

For both type of fund management structures, there is no skill evident that would suggest that fund managements are able to anticipate changes in market expectations driven by systematic information (columns 2 and 4). This is consistent with the finding for the entire sample. The persistence in skill to anticipate changes in expectations related to cash flows (columns 1) is no longer present for the subsample containing only single-managed funds. This suggests that a fund manager on her own is not able to acquire and process firm-specific information better than the aggregate market in order to anticipate changes in market expectations. When turning to team managed funds (column 3 and 4), however, the results suggest that this fund management structure tends to have the ability to persistently anticipate changes in market expectations driven by firm-specific information. Put differently, a group of individuals is superior at acquiring and processing firm-specific information, which allows them to anticipate changes in aggregate market expectations and adjust their holdings accordingly. Hence, if multiple fund managers combine their resources, they have superior skill compared to a single fund manager. This finding is strongly supported by the literature of behavioral decision making theory, where teams are better at acquiring and processing information under uncertainty, (Hinsz et al. (1997)) and are able to process larger amounts of information (Vollrath et al. (1989)). It can also be reconciled with the notion that team-managed funds are larger because rational capital allocation would dictate that capital follows skill, and the fact that the fraction of team-managed U.S. equity mutual funds increases (Adams et al. (2018)). Finally, the results complement the superior fund performance of team-managed funds documented by Han et al. (2017). In the context of this paper, I conclude that skill in the U.S. equity mutual funds industry is related to processing firm-specific but not market-wide information, and that this skill is only present among management structures consisting of multiple individuals.

VI Informational inefficiency and skill

The previous sections establish that the U.S. equity mutual funds industry exhibits skill in that team-managed fund structures are able to consistently anticipate changes in aggregate market expectations driven by firm-specific information but not those driven by market-wide information. Following Fama and French (2010), in an efficient market, fund managers should not be able to persistently anticipate changes in aggregate market expectations. Hence, my findings are inconsistent with the Efficient Markets Hypothesis (EMH). In this section, I provide a potential economic rationale in the form of informational inefficiency at the asset level to justify and support my findings. I also provide a potential explanation why skill is only persistent with respect to idiosyncratic but not systematic information.

A Measuring informational inefficiency

To skillfully anticipate changes in the aggregate market's expectations about future cash flows and the discount rate of a stock, it is required that this particular stock be informationally inefficient in that its price does not (yet) reflect all available information, unless correct anticipation is a consequence of private information. The informational inefficiency can be related to firm-specific information (cash flow news) and/or systematic information (discount rate news). A fund manager that correctly anticipates changes in aggregate market expectations about cash flows may be able to do so because she allocates capital into assets that are informationally inefficient with respect to cash flow news. Likewise, a fund manager that excels at anticipating changes in discount rate news may allocate assets to stocks that are informationally inefficient with respect to discount rate news.

In order to measure the informational inefficiency of a particular asset, I estimate the completeness of information in stock returns. Incomplete information in prices suggests that the market has not incorporated all available information yet. At some point, however, the market likely will incorporate the remaining amount of information which triggers changes in its expectations, and ultimately returns. Fund managements that are able to identify stocks that reflect incomplete information may be able to anticipate changes in market expectations as long as information in prices will be complete eventually.

Information completeness

In order to measure the completeness of information of a security, I use entropy. Specifically, incomplete amounts of information in a price series gives raise to predictable patterns (Prado (2018)). Patterns occur when a price series contains redundant information, which enables its compression. Entropy measures the degree of compression, or redundancy of information. When prices are decompressed and unpredictable, the amount of information in prices is complete. If prices are compressed and can be predicted, information is incomplete.

Undoubtedly, other informational efficiency measures exists. However, for this exercise, it is of towering importance to use an efficiency measure that allows measuring the informational efficiency with respect to both systematic and idiosyncratic information, and allows to contrast the outcomes. For instance, the delay measure of Hou and Moskowitz (2005) measures informational efficiency with respect to systematic information only. Similarly, the cross-sectional return predictability measure of Heston et al. (2010) deducts an overall market effect, making it only suitable to estimate informational efficiency with respect to idiosyncratic information. The Hasbrouck pricing errors use a VAR with state variables that cannot be decomposed into a systematic and an idiosyncratic part (see Hasbrouck (1993); Boehmer and Kelley (2009)). The same is true for intraday return predictability measures that use order imbalance or efficiency measures derived from options. Hence, suitable measures are those that draw on the idea of predictability of and autocorrelation in returns, and returns only.

Many estimators for entropy have been proposed in the literature. In order to estimate

the informational completeness of stock returns, I use compression algorithms to compute the length of repeating patterns in the return sequence. The most commonly used algorithms in this field date back to [Ziv and Lempel \(1978\)](#) and decompose a data sequence into a set of non-redundant substrings. The larger the set of unique substrings relative to the total length of the data, the more informationally efficient a message appears, and the higher its entropy. Intuitively, as the size of the set containing the Lempel-Ziv substrings increases, the patterns in the return sequence are more unique and shorter.

The estimators of [Kontoyiannis et al. \(1998\)](#) and [Gao et al. \(2008\)](#) draw on the idea of Lempel-Ziv algorithms and assess a data sequences' informational efficiency by searching at each position in the data sequence for the longest matching pattern subsequent to that position in the data sequence, with respect to the segment of the data sequence of a certain length just prior to that position in the data sequence. Specifically, using a return sequence, $\{R^n\}$, with length n , for every position in $\{R^n\}$ and window length $w \geq 1$, find the length ℓ of the longest return pattern $r_i^{i+\ell}$ in the return sequence $\{R^n\}$ starting at position i that also appears in the window r_{i-w}^{i-1} preceding position i . Formally, define:

$$\begin{aligned} L_i^w &= L_i^w \{R_1^n\} = L_i^w (x_{i-w}^{i+w-1}) \\ &= 1 + \max\{0 \leq \ell \leq w : x_i^{i+\ell-1} = x_j^{j+\ell-1} \text{ for some } i-w \leq j \leq i-1\}, \end{aligned}$$

which corresponds to 1 plus the longest match ℓ . [Ornstein and Weiss \(1993\)](#) establish that:

$$\lim_{w \rightarrow \infty} \frac{L_i^w}{\log_2(w)} = \frac{1}{H}, \quad (27)$$

where H corresponds to the entropy of the data sequence $\{R_1^n\}$. [Kontoyiannis et al. \(1998\)](#) use the reciprocal of equation (27), $\log_2(w)/L_i^w$, to estimate the entropy of a data sequence. In order to reduce variance and make more efficient use of the data, [Gao et al. \(2008\)](#) suggest that the average of various match-lengths, L_i^w , at different positions be taken.

To that end, I adopt the entropy measure proposed by [Gao et al. \(2008\)](#) and refined by [Prado \(2018\)](#) with an expanding window. Given some return sequence, $\{R^n\}$, of length n , for every position $i > 1$ in $\{R^n\}$ up to position $i = \lfloor n/2 \rfloor + 1$, find the length $\max\{\ell^i\}$ of the longest return pattern $r_i^{i+\ell}$ starting at position i that also appears in the window r_1^{i-1} , i.e., the entire return sequence preceding position i . Rather than using a shifting window, this measure uses an expanding window that expands as the position i increases and has length $i - 1$ rather than a constant length. In this case for a return sequence of length n the entropy estimator H is defined as:

$$\hat{H} = \left[\frac{1}{\lfloor n/2 \rfloor} \sum_{i=2}^{\lfloor n/2 \rfloor + 1} \frac{L^i}{\log_2(i)} \right]^{-1} \quad (28)$$

where L^i denotes $1 + \max\{\ell^i\}$ for position i . The intuition is that if the average length of the longest matches is large the data sequence consist of only a few longer return patterns and thus displays low entropy. On the other hand, if the average length is short the message

consist of many short and unique return patterns casting the return sequence to display high entropy, ultimately rendering it unpredictable and thus informationally efficient.

Identifying redundant return patterns requires redundant return values. As returns are continuous by nature, It is required that returns be assigned to return bins. Furthermore, the length of the return sequence must be decided. In order to set the number of past return observations and the number of discretized bins returns can be assigned to, I resort to the fact that the entropy of a Gaussian IID process maps to 1.42. This allows me to identify the correct parameter values by benchmarking a Gaussian IID process for various tuples of parameter values. Specifically, I maximize the negative absolute difference between the average of rolling entropy estimates from equation (28) over a sufficiently long Gaussian IID process and the true entropy of a Gaussian IID process over the number of past returns n and return bins q in the parameter space Θ :

$$\arg \max_{n \in \Theta, q \in \Theta} - \left| (T - n + 1)^{-1} \sum_{t=n}^T \left[\left(\frac{1}{\lfloor n/2 \rfloor} \sum_{i=2}^{\lfloor n/2 \rfloor + 1} \frac{L^i}{\log_2(i)} \right)^{-1} \right]_{t=t} - 1.42 \right|, \quad (29)$$

which yields an optimal sequence length of 50 returns that can be assigned to 4 discrete bins. I then use these parameters to estimate the informational efficiency of mutual fund holdings.

Systematic and idiosyncratic inefficiency

To assess the systematic and idiosyncratic informational efficiency, I use the estimator in equation (28) to estimate the entropy of the systematic and idiosyncratic return portion of fund holdings. Ideally, I use cash flow and discount rate news as defined in equations (8) and (7) as the return's idiosyncratic and systematic return portion. However, the return decomposition yields one single observation for changes in expectations (unexpected returns) in a quarter, which renders estimating a contemporaneous informational efficiency in a given quarter impossible, regardless of what efficiency measure is applied. In order to estimate a quarter-specific estimate of informational efficiency, I need the systematic and idiosyncratic return components at the daily frequency.⁵

To estimate the systematic and idiosyncratic return portion at the daily frequency, I draw on the approaches presented in Kacperczyk et al. (2014) and Daniel et al. (1997) to disentangle skill related to idiosyncratic returns from skill related to systematic returns, and regress excess returns of a security on the market's excess returns. Specifically, I estimate a Conditional Capital Asset Pricing Model to decompose returns into a systematic and an idiosyncratic component:

$$r_{i,t} - rf_t = \alpha_{i,t} + \beta_{i,t}(r_{m,t} - rf_t) + \varepsilon_{i,t} \quad (30)$$

Here, $(r_{m,t} - rf_t)$ is the excess return on the market and $\beta_{i,t}$ measures the time-varying

⁵Many efficiency measures that are based on return predictability even use intraday frequency (see e.g., Röscher et al. (2017))

sensitivity of stock i with respect to this systematic factor. Thus, the residuals of the regression, $\varepsilon_{i,t}$, capture any return variation that is not explained by the security’s systematic exposure. I estimate rolling betas over one year of daily returns. Systematic (excess) returns, $r_{i,t}^s$, and idiosyncratic (excess) returns, $r_{i,t}^i$, are then defined as:

$$r_{i,t}^s = \hat{\alpha}_{i,t} + \hat{\beta}_{i,t}(r_{m,t} - rf_t) \quad (31)$$

$$r_{i,t}^i = \hat{\varepsilon}_{i,t} \quad (32)$$

Alternatively, I use the [Fama and French \(1992\)](#) three-factor model, and the systematic and idiosyncratic return components are defined analogously. I then proceed with estimating informational inefficiency using the entropy estimator in equation (28) on systematic and idiosyncratic returns. In order to interpret the entropy estimate in a meaningful way, I express it relative to the entropy of a Gaussian IID process. A Gaussian IID process displays maximal entropy by nature and its unconditional entropy maps to 1.42. I scale the entropy estimate of security i ’s systematic or idiosyncratic return at time t by the unconditional entropy of a Gaussian IID process and denote it by $\hat{H}_{i,t}^s$ or $\hat{H}_{i,t}^i$, respectively. Thus, this measure is bounded on the unit interval and is easy to interpret. Most importantly, the entropy estimates of the systematic and the idiosyncratic return portions can be directly compared. A value of 1 implies that returns are essentially a Gaussian IID process and fully efficient while a lower value implies a higher degree of inefficiency and predictability.

B Results

To examine whether a difference in informational efficiency with respect to idiosyncratic and systematic information underlies my finding that fund managers can only anticipate changes in expectations driven by firm-specific information, I test the following additional hypothesis:

$H_{5,null}$: *The informational efficiency of assets with respect to firm-specific and market-wide information is not significantly different.*

The corresponding alternative hypothesis maps to:

$H_{5,alternative}$: *The informational efficiency of assets with respect to firm-specific and market-wide information is significantly different.*

I start by examining the informational efficiency with respect to firm-specific and market-wide information for all stocks in the CRSP universe. Table X reports the results. I estimate rolling entropy estimates of the return components using the past 50 daily returns and 4 return bins. I then take the average daily entropy estimate within a quarter to estimate a security’s informational efficiency in that quarter. Panel A reports the sample characteristics of informational efficiency for all stocks in the CRSP universe. At the average, at the median, and across several percentiles, efficiency with respect to idiosyncratic information is systematically lower, regardless of what asset pricing model is used to estimate the systematic and idiosyncratic return portions. Panel B reports t-tests and Mood’s median tests

Table X: Asset efficiency

This table reports the efficiency of systematic and idiosyncratic returns as defined in equations (31) and (32), respectively. The systematic and idiosyncratic return components are estimated using a Capital Asset Pricing Model (CAPM) and the Fama and French (1992) three factor model (FF-3F). Efficiency is estimated using the entropy estimator of Gao et al. (2008) as defined in equation (2) using 50 daily returns and 4 return bins. Panel A reports summary statistics on the efficiency of all assets held by the U.S. equity mutual funds in the sample. Panel B reports differences in means using a standard t-test, and differences in medians using a Mood’s test, between the efficiency of idiosyncratic returns and systematic returns of those assets. The sample period is 2003 to 2021. p -values are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

Panel A: Sample characteristics				
	CAPM		FF-3F	
	$\hat{H}_{i,t}^s$	$\hat{H}_{i,t}^i$	$\hat{H}_{i,t}^s$	$\hat{H}_{i,t}^i$
Observations	2'079'941	2'079'941	2'079'941	2'079'941
Average Efficiency	0.91	0.88	0.93	0.88
Standard deviation	0.10	0.12	0.08	0.12
5%	0.72	0.65	0.77	0.65
25%	0.86	0.82	0.89	0.82
Median Entropy	0.94	0.90	0.95	0.91
75%	0.99	0.97	1.00	0.97
95%	1.00	1.00	1.00	1.00
Panel B: Sample tests				
	CAPM		FF-3F	
	$\hat{H}_{i,t}^i - \hat{H}_{i,t}^s$		$\hat{H}_{i,t}^i - \hat{H}_{i,t}^s$	
Mean	-0.0351*** (0.00)		-0.0525*** (0.00)	
Median	-0.0358*** (0.00)		-0.0497*** (0.00)	

for the differences between average and median efficiencies with respect to idiosyncratic and systematic returns. Across the board, informational efficiency of the idiosyncratic component is lower. This suggests that stock returns are significantly less complete with respect to firm-specific information than they are with respect to market-wide information. This provides one potential rationale that explains fund managers’ ability to anticipate changes in market expectations driven by firm-specific information only.

U.S. equity mutual fund holdings, however, do generally not represent the broad market due to their active asset allocation approach. Furthermore, table X reports material variation in informational efficiency at the asset level. Therefore, the aggregate informational inefficiency of the fund portfolio may significantly differ from the market-wide average. The argument that a low market-wide average efficiency allows the U.S. equity mutual fund industry to anticipate changes in expectations driven by firm-specific information thus requires that funds ultimately also hold assets that are informationally inefficient. Specifically, the U.S. equity mutual fund industry’s ability to anticipate changes in expectations driven

Table XI: Portfolio efficiency

This table reports the difference in efficiency between a portfolio containing all stocks held by U.S. equity mutual funds and the market portfolio for the systematic and idiosyncratic return components. The systematic and idiosyncratic return components are estimated using a Capital Asset Pricing Model (CAPM) and the Fama and French (1992) three factor model (FF-3F). Efficiency is estimated using the entropy estimator of Gao et al. (2008) as defined in equation (2) using 50 daily returns and 4 return bins. The sample period is 2003 to 2021. p -values are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively.

	CAPM		FF-3F	
	$\bar{H}_{i,t}^s$	$\bar{H}_{i,t}^i$	$\bar{H}_{i,t}^s$	$\bar{H}_{i,t}^i$
$\bar{H}_{j,t} - \bar{H}_{m,t}$	-0.01 (p=0.41)	-0.02*** (p=0.00)	0.00 (p=0.71)	-0.02*** (p=0.00)
Obs. (quarters)	71	71	71	71

by firm-specific information requires that the informational efficiency of their portfolio is similar to or lower than the market-wide average documented earlier. Similarly, the inability to anticipate changes driven by systematic information requires that the respective informational efficiency of their portfolio is similar to or higher than the reported market-wide average.

To that end, I assess the informational efficiency of the market portfolio and a portfolio containing all assets held by the U.S. equity mutual funds industry in my sample. For each quarter, I compute an average informational efficiency of the securities in that portfolio. For the market portfolio, this amounts to the average informational inefficiency of all stocks in the CRSP universe in a given quarter. For the portfolio encompassing fund holdings, this maps to the average informational inefficiency across all stocks that are held by any U.S. equity mutual fund in a given quarter. I then test whether the informational efficiency of fund holdings is significantly different from that of the entire market.

Table XI reports the results. Informational efficiency regarding systematic information is the same for fund holdings and the entire market. This is consistent with my previous finding that fund managers are not able to anticipate changes in expectations that are driven by systematic information as information in prices is rather complete. For idiosyncratic information, the average informational efficiency of securities held by U.S. equity mutual funds is significantly lower than the market-wide average. This implies that funds do not only hold assets that have low informational efficiency but are, on top, able to pick stocks where information is less complete than the market-wide average regarding firm-specific information.

In conclusion, idiosyncratic information in assets held by U.S. equity mutual funds is significantly less complete compared to systematic information, making the return component driven by firm-specific information more predictable. In consequence, skilled fund management structures may be able to anticipate changes in expectations driven by idiosyncratic information, as long as the market incorporates all available idiosyncratic information even-

tually. By contrast, systematic information in asset prices is much more complete, making it much harder, or impossible, to consistently anticipate changes in aggregate market expectations driven by systematic information. This provides a potential economic rationale as to why some fund management structures are able to anticipate changes in expectations driven by firm-specific information but not those driven by systematic information.

VII Conclusion

In this paper, I use the return decomposition of [Vuolteenaho \(2002\)](#) to decompose returns of U.S. equity mutual funds holdings into changes in expectations about cash flows and the discount rate and analyze whether mutual fund managers are able to anticipate such changes and adjust their fund portfolio accordingly. I find that mutual fund managers are able to persistently anticipate changes in expectations about cash flows but not changes related to the discount rate. [Vuolteenaho \(2002\)](#) finds that expectation changes about cash flow news are driven by firm-specific, idiosyncratic information while changes in expectations about the discount rate are driven by market-wide, systematic information. Thus, U.S. equity mutual fund managements display skill at acquiring and processing firm-specific information which allows them to make better choices than the aggregate market. Furthermore, I find that this ability is only present among fund management structures that consist of multiple individuals, suggesting that groups are better at acquiring and processing firm-specific information and making investment decisions. Finally, I show that firm-specific information in returns of fund holdings is less complete compared to systematic information contained in the returns of fund holdings. This significantly larger informational inefficiency concerning idiosyncratic information in returns potentially allows skilled fund managers to correctly anticipate changes in market expectations driven by firm-specific information in the first place.

References

- Adams, J. C., Nishikawa, T., Rao, R. P., 2018. Mutual fund performance, management teams, and boards. *Journal of Banking & Finance* 92, 358–368.
- Alexander, G. J., Cici, G., Gibson, S., 2007. Does motivation matter when assessing trade performance? an analysis of mutual funds.
- Arrow, K. J., 1986. Rationality of self and others in an economic system. *The Journal of Business* 59, 358–399.
- Baker, M., Litov, L., Wachter, J. A., Wurgler, J., 2010. Can mutual fund managers pick stocks? evidence from their trades prior to earnings announcements. *Journal of Financial and Quantitative Analysis* 45, 1111–1131.
- Barras, L., Gagliardini, P., Scaillet, O., 2022. Skill, scale, and value creation in the mutual fund industry. *The Journal of Finance* 77, 601–638.
- Berk, J. B., Green, R. C., 2004. Mutual fund flows and performance in rational markets. *Journal of Political Economy* 112, 1269–1295.
- Berk, J. B., van Binsbergen, J. H., 2015. Measuring skill in the mutual fund industry. *Journal of Financial Economics* 118, 1–20.
- Bliss, R. T., Potter, M. E., Schwarz, C., 2008. Performance characteristics of individually-managed versus team-managed mutual funds. *The Journal of Portfolio Management* 34, 110–119.
- Boehmer, E., Kelley, E. K., 2009. Institutional investors and the informational efficiency of prices. *The Review of Financial Studies* 22, 3563–3594.
- Brown, S. J., Goetzmann, W. N., 1995. Performance persistence. *The Journal of Finance* 50, 679–698.
- Bär, M., Kempf, A., Ruenzi, S., 2011. Is a team different from the sum of its parts? evidence from mutual fund managers. *Review of Finance* 15, 359–396.
- Campbell, J. Y., 1991. A variance decomposition for stock returns. *The Economic Journal* 101, 157.
- Campbell, J. Y., Shiller, R. J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1, 195–228.
- Carhart, M. M., 1997. On persistence in mutual fund performance. *Journal of Finance* 52, 57–82.
- Chen, H.-L., Jegadeesh, N., Wermers, R., 2000. The value of active mutual fund management: An examination of the stockholdings and trades of fund managers. *The Journal of Financial and Quantitative Analysis* 35, 343.
- Chen, J., Hong, H., Huang, M., Kubik, J. D., 2004. Does fund size erode mutual fund performance? the role of liquidity and organization. *American Economic Review* 94, 1276–1302.
- Chordia, T., 1996. The structure of mutual fund charges. *Journal of Financial Economics* 41, 3–39.

- Cremers, K. J., Petajisto, A., 2009. How active is your fund manager a new measure that predicts performance. *Review of Financial Studies* 22, 3329–3365.
- Daniel, K., Grinblatt, M., Titman, S., Wermers, R., 1997. Measuring mutual fund performance with characteristic-based benchmarks. *The Journal of Finance* 52, 1035–1058.
- Daniel, K., Titman, S., 1997. Evidence on the characteristics of cross sectional variation in stock returns. *The Journal of Finance* 52, 1–33.
- Edelen, R. M., 1999. Investor flows and the assessed performance of open-end mutual funds. *Journal of Financial Economics* 53, 439–466.
- Erickson, T., Jiang, C. H., Whited, T. M., 2014. Minimum distance estimation of the errors-in-variables model using linear cumulant equations. *Journal of Econometrics* 183, 211–221.
- Erickson, T., Whited, T. M., 2002. Two-step gmm estimation of the errors-in-variables model using high-order moments. *Econometric Theory* 18, 776–799.
- Fama, E. F., 1970. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance* 25, 383.
- Fama, E. F., French, K. R., 1992. The cross-section of expected stock returns. *Journal of Finance* 47, 427–465.
- Fama, E. F., French, K. R., 2010. Luck versus skill in the cross-section of mutual fund returns. *Journal of Finance* 65, 1915–1947.
- Gao, Y., Kontoyiannis, I., Bienenstock, E., 2008. Estimating the entropy of binary time series: Methodology, some theory and a simulation study. *Entropy* 10, 71–99.
- Gibson, S., Safieddine, A., Titman, S., 2000. Tax-motivated trading and price pressure: An analysis of mutual fund holdings. *The Journal of Financial and Quantitative Analysis* 35, 369.
- Grinblatt, B. M., Titman, S., Wermers, R., 1995. Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior. *American Economic Review* 85, 1088–1105.
- Grinblatt, M., Titman, S., 1989. Mutual fund performance: An analysis of quarterly portfolio holdings. *The Journal of Business* 62, 393.
- Grinblatt, M., Titman, S., 1993. Performance measurement without benchmarks: An examination of mutual fund returns. *The Journal of Business* 66, 47.
- Grossman, S. J., Stiglitz, J. E., 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.
- Gruber, M. J., 1996. Another puzzle: The growth in actively managed mutual funds. *Journal of Finance* 51, 783–810.
- Han, Y., Noe, T., Rebello, M., 2017. Horses for courses: Fund managers and organizational structures. *Journal of Financial and Quantitative Analysis* 52, 2779–2807.
- Hasbrouck, J., 1993. Assessing the quality of a security market: A new approach to transaction-cost measurement. *Review of Financial Studies* 6, 191–212.

- Heston, S. L., Korajczyk, R. A., Sadka, R., 2010. Intraday patterns in the cross-section of stock returns. *Journal of Finance* 65, 1369–1407.
- Hinsz, V. B., Vollrath, D. A., Tindale, R. S., 1997. The emerging conceptualization of groups as information processors. *Psychological Bulletin* 121, 43–64.
- Hou, K., Moskowitz, T. J., 2005. Market frictions, price delay, and the cross-section of expected returns. *Review of Financial Studies* 18, 981–1020.
- Huddart, S., Narayanan, V. G., 2002. An empirical examination of tax factors and mutual funds' stock sales decisions. *Review of Accounting Studies* 7, 319–341.
- Jensen, M. C., 1968. The performance of mutual funds in the period 1945-1964. *The Journal of Finance* 23, 389.
- Kacperczyk, M., Nieuwerburgh, S. V., Veldkamp, L., 2014. Time-varying fund manager skill. *Journal of Finance* 69, 1455–1484.
- Kacperczyk, M., Seru, A., 2007. Fund manager use of public information: New evidence on managerial skills. *Journal of Finance* 62, 485–528.
- Kacperczyk, M., Sialm, C., Zheng, L., 2008. Unobserved actions of mutual funds. *Review of Financial Studies* 21, 2379–2416.
- Kahneman, D., Tversky, A., 1984. Choices, values, and frames. *American Psychologist* 39, 341–350.
- Kontoyiannis, I., Algoet, P. H., Suhov, Y. M., Wyner, A. J., 1998. Nonparametric entropy estimation for stationary processes and random fields, with applications to english text. *IEEE Transactions on Information Theory* 44, 1319–1327.
- Lakonishok, J., Shleifer, A., Thaler, R., Vishny, R., 1991. Window dressing by pension fund managers. *American Economic Review* 81, 227–231.
- Leippold, M., Rueegg, R., 2020. How rational and competitive is the market for mutual funds? *Review of Finance* 24, 579–613.
- Levenshtein, V., 1965. Binary codes capable of correcting deletions, insertions, and reversals. *Doklady Akademii Nauk SSSR* 163, 845–848.
- Michaely, R., Rossi, S., Weber, M., 2020. Signaling safety. Working Paper .
- Musto, D. K., 1999. Investment decisions depend on portfolio disclosures. *Journal of Finance* 54, 935–952.
- Nanda, V., Narayanan, M. P., Warther, V. A., 2000. Liquidity, investment ability, and mutual fund structure. *Journal of Financial Economics* 57, 417–443.
- Ornstein, D. S., Weiss, B., 1993. Entropy and data compression schemes. *IEEE Transactions on Information Theory* 39, 78–83.
- Pastor, L., Stambaugh, R. F., Taylor, L. A., 2015. Scale and skill in active management. *Journal of Financial Economics* 116, 23–45.
- Pastor, L., Stambaugh, R. F., Taylor, L. A., 2017. Do funds make more when they trade more? *The Journal of Finance* 72, 1483–1528.

- Prado, M. L. D., 2018. *Advances in financial machine learning*. Wiley & Sons, Inc., first ed.
- Prather, L. J., Middleton, K. L., 2002. Are $n + 1$ heads better than one? the case of mutual fund managers. *Journal of Economic Behavior and Organization* 47, 103–120.
- Rösch, D. M., Subrahmanyam, A., Dijk, M. A. V., 2017. The dynamics of market efficiency. *Review of Financial Studies* 30, 1151–1187.
- Snizek, J. A., Henry, R. A., 1989. Accuracy and confidence in group judgment. *Organizational Behavior and Human Decision Processes* 43, 1–28.
- Treynor, J., Mazuy, K., 1966. Can mutual funds outguess the market? *Harvard Business Review* 44, 131–136.
- Vollrath, D. A., Sheppard, B. H., Hinsz, V. B., Davis, J. H., 1989. Memory performance by decision-making groups and individuals. *Organizational Behavior and Human Decision Processes* 43, 289–300.
- Vuolteenaho, T., 2002. What drives firm-level stock returns? *Journal of Finance* 57, 233–264.
- Wermers, R., 2000. Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses. *Journal of Finance* 55, 1655–1703.
- Zhu, M., 2018. Informative fund size, managerial skill, and investor rationality. *Journal of Financial Economics* 130, 114–134.
- Ziv, J., Lempel, A., 1978. Compression of individual sequences via variable-rate coding. *IEEE Transactions on Information Theory* 24, 530–536.

Appendix Essay 2

A Data cleaning and variable construction

A.1 Fund universe and share class aggregation

I start by filtering for all U.S. mutual funds that are categorized as domestic equity (ED) using the CRSP style code. Following [Pastor et al. \(2015\)](#), I exclude all mutual funds that are flagged as index funds and I exclude funds that have less than 5 million assets under management (AUM) as done by [Berk and van Binsbergen \(2015\)](#).⁶ In some cases, multiple fund numbers report identical holdings. Specifically, when a mutual fund offers multiple share classes they are listed as separate funds in the CRSP Mutual Funds database. While the share classes' respective fee structures and hence, returns may slightly differ, their asset allocation is the same. As a consequence, I aggregate share classes at the fund level based on the Total Net Asset (TNA) of the share class relative to the TNA of all share classes in the fund. Where TNA are missing, I interpolate the missing values linearly. I end up with 5'347 domestic U.S. equity funds, which is in line with other literature.⁷

A.2 Mutual fund managers

For each mutual fund in the universe, I construct the history of fund management at the share class level to identify which manager is managing the share class at a given point in time. Manager names often have various different recordings despite pertaining to the same individual. For instance, an observation includes the middle name in one quarter but does no longer in the subsequent quarters. In other instances, first and last name are swapped from one reporting to the next, yet in other cases names simply contain typos.

To that end, I employ techniques from information theory and data analysis to identify whether two names ultimately correspond to the same individual. I start by removing middle names or initials. Second, I use the measure of [Levenshtein \(1965\)](#) to identify typos. In information theory, the measure of [Levenshtein \(1965\)](#) is a string metric for measuring the difference between two sequences by assessing how many characters need to be changed in order for one word to be identical to the other. A Levenshtein measure that is close to one suggests that two name strings are almost identical and only one or two characters are different. I then replace the flawed entry with the correct one. Third, I use a cosine similarity measure to identify whether the first and last name are swapped. In data analysis, cosine similarity is a measure of similarity between two sequences. While the Levenshtein

⁶Some authors are more restrictive and exclude funds with less than 15 million AUM (see e.g., [Pastor et al. \(2015\)](#)). Although unreported, I warrant that the results are not driven by the selection of this threshold.

⁷For instance, [Berk and van Binsbergen \(2015\)](#) end up with more than 6000 U.S. mutual fund up to the year 2015 but also include non-equity funds, namely bond funds. [Pastor et al. \(2015\)](#) only have around 3'100 domestic U.S. equity but their sample ends in 2011. Furthermore, they are more restrictive in their definition of domestic equity funds.

measure would consider a string where names are swapped as different strings, the cosine similarity measure treats them as identical. As a result, I replace observations where first and last names are swapped such that all names start with the first name followed by the last name. Both the Levenshtein measure and the cosine similarity measure are explained in detail in appendix [B.1](#) and appendix [B.2](#)

Many mutual funds are managed by a team of individuals. For these observations the manager name variable is reported as *Team Managed*. In some cases, however, the names are explicitly reported. For these entries, only last names are provided, which are often recorded in a different order over the reporting dates. I clean the reported names using the measure of [Levenshtein \(1965\)](#) and I use the cosine similarity measure to assimilate all present orders of the same team of individuals recorded in the database. In some cases, different share classes of a fund are reported to be managed by different individuals despite being subject to the same asset allocation. In order to be able to aggregate share classes at the fund level, I assume that the fund is jointly managed by all the managers reported at the share class level. Hence, I replace the fund manager names with *Team Managed*. Where all share classes in a fund are managed by the same individual, no further actions are required in the aggregation.

In a final step, I merge the history of holdings with the history of managers at the share class level to obtain for each fund manager the share class she is managing and the holdings that she has at a given point in time. Since all share classes are subject to the same holdings and names have been cleaned, aggregation at the fund level is straight forward.

A.3 Fund holdings and returns

For each holding of a fund in a given quarter, I compute the change in numbers of shares relative to the last reporting date. Where the last reporting date is not available, I compute the relative change to the closest reporting date available if it does not date back more than 4 quarters. As the reported holdings do not account for stock splits, I adjust the reported number of shares held in a fund portfolio using the CRSP adjustment factor to eliminate changes in holdings that are simply due to stock splits.

I then proceed with estimating changes in holdings due to informed trading by filtering out trading due to uninformed trading as outlined in the methodology section. Where there is no change in shares, the residuals from the regressions are set to zero. Where no residuals are available, the informed trading is proxied by the change in the number of shares, i.e., the unfiltered change. For small changes, in cases where the sign of the residual is different from the sign of the change in shares, informed trading is also proxied by the overall change in shares.

Some funds report their holdings for every month while others report them on a quarterly basis. As the return decomposition estimates changes in aggregate market expectations at a quarterly frequency, it is required that changes in holdings be quarterly as well. Where holdings are reported at the monthly frequency, I aggregate them to a quarterly frequency.

Similarly, fund returns are reported at the monthly frequency. Where fund returns are missing, I compute the return as the relative change in TNA. Finally, I compute the cumulative quarterly return of a fund. To aggregate returns at the fund level, I weight the return of each share class with the share class' TNA relative to the total TNA of all share classes in the fund, which is a standard convention in the mutual funds literature.

A.4 Return decomposition

In what follows, I follow [Michaely et al. \(2020\)](#) and [Vuolteenaho \(2002\)](#) and lay out the variables of interest for the return decomposition. I compute the simple quarterly stock return as the cumulative monthly return within a fiscal quarter, recorded from m to $m + 2$ for $m \in \{February, May, August, November\}$. As in [Michaely, Rossi, and Weber \(2020\)](#), I assume a de-listing return of 30% if a firm is de-listed for a known cause and has a missing de-listing return. The return r_t is the market-adjusted log return defined as log return less the cross-sectional average log return (see, e.g., [Vuolteenaho \(2002\)](#)). Market equity is defined as the total market equity at the firm level as recorded in CRSP at the end of each quarter. If quarter t market equity is missing, I let the previous quarter's market equity grow with the rate of return in that quarter (without dividends).

Book equity is defined as shareholders' equity plus balance-sheet deferred taxes and investment tax credit (item TXDITCQ) if available, minus the book value of preferred stock. I use stockholders' equity (item SEQQ), or common equity (item CEQQ) plus the carrying value of preferred stock (item PSTKQ), or total assets (item ATQ) minus total liabilities (item LTQ) in that order as shareholders' equity. I use the redemption value (item PSTKRQ) if available, or carrying value for the book value of preferred stock. Whenever book equity is unavailable, I proxy for it by the last period's book equity plus earnings, less dividends, assuming that the clean-surplus relation holds. If neither earnings nor book equity are available, I assume that the book-to-market ratio has not changed from quarter $t - 1$ to quarter t , and compute the book-equity proxy from the last quarter's book-to-market ratio and this quarter's market equity. I exclude firms with a quarter $t - 1$ market equity of less than USD 10 million and a book-to-market ratio of more than 100 or less than 0.01. Moreover, I set negative or zero book-equity values to missing.

ROE is defined as earnings over beginning of quarter's book equity. To compute the ROE I use earnings available for common equity. When earnings are missing, I use the clean-surplus formula to approximate earnings. I drop observations with a ROE lower than - 100%. Each quarter, I log transform market equity, stock returns, and return on equity and cross-sectionally demean it. A log transformation can be cumbersome if returns are close to 1 or if book-to-market ratios are close to zero or infinity. I mitigate these concerns by following [Michaely et al. \(2020\)](#) and [Vuolteenaho \(2002\)](#) and redefine each firm as a portfolio that consists of 90% common stock and 10% Treasury bills using market values. Every period, the portfolio is rebalanced to reflect these weights.

B Name cleaning measures

B.1 Levenshtein measure

In information theory, the measure of [Levenshtein \(1965\)](#) is a string metric for measuring the difference between two sequences. Simply put, the Levenshtein measure between two words is the minimum number of single-character edits required to change one word into the other. Formally, for two sequences a and b with lengths i and j , the Levenshtein measure is defined as:

$$lev_{a,b}(i, j) = \begin{cases} \max(i, j) & \text{if } \min(i, j) = 0, \\ \min \begin{cases} lev_{a,b}(i-1, j) + 1 \\ lev_{a,b}(i, j-1) + 1 \\ lev_{a,b}(i-1, j-1) + \mathbb{1}(a_i \neq b_i) \end{cases} & \text{otherwise.} \end{cases}$$

where $\mathbb{1}(a_i \neq b_i)$ is an indicator function that takes the value 1 if $a_i \neq b_i$. I then transform the Levenshtein measure into a ratio bounded on the interval $[0, 1]$ using:

$$LEV_{a,b}(i, j) = 1 - \frac{lev_{a,b}(i, j)}{\max(i, j)}, \quad (33)$$

such that a value of one indicates that two sequences are identical.

B.2 Cosine similarity measure

In data analysis, cosine similarity is a measure of similarity between two sequences of numbers. For defining it, the sequences are viewed as vectors in an inner product space, and the cosine similarity is defined as the cosine of the angle between them. The cosine similarity always belongs to the interval $[-1, 1]$. For example, two proportional vectors have a cosine similarity of 1, two orthogonal vectors have a similarity of 0, and two opposite vectors have a similarity of -1. The cosine of two non-zero vectors can be derived by using the Euclidean dot product formula:

$$\mathbf{A} \cdot \mathbf{B} = \|\mathbf{A}\| \|\mathbf{B}\| \cos\theta \quad (34)$$

Given two vectors of attributes, \mathbf{A} and \mathbf{B} , the cosine similarity, $\cos(\theta)$, is represented using a dot product and magnitude as:

$$\cos(\theta) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}}, \quad (35)$$

where A_i and B_i are components of vector \mathbf{A} and \mathbf{B} , respectively.

C Additional tables

Table C.I: Vector Auto Regression (VAR) transition matrix

This table reports the transition matrix for a panel VAR for all firms in the CRSP/Compustat database using the method outlined in section III, as well as the covariance matrix of the news components. r_t denotes the centered excess log stock return, θ is the centered log book-to-market ratio, and RoE is the centered log return-on-equity. The sample period is from 2003 to 2021.

	r_{t-1}	θ_{t-1}	RoE_{t-1}
r_t	0.001 (0.45)	0.024 (12.35)	0.115 (13.75)
θ_t	-0.276 (-112.47)	0.849 (423.00)	-0.115 (-11.64)
RoE_t	0.014 (14.06)	-0.020 (-20.29)	0.227 (37.81)
COV		η_{cf}	$-\eta_r$
	η_{cf}	0.0306	0.0020
	$-\eta_r$	0.0020	0.0007

Table C.II: Beta autocorrelation for partially skilled funds and non-zero holding changes

This table reports the results for the following panel regression:

$$\beta_{j,t}^s = c + \delta_i + \tau_t + \mathbf{\Gamma}' \mathbf{X}_{j,t} + \sum_{q=1}^4 \phi_q \beta_{j,t-q}^s + \sum_{q=1}^4 \gamma_q \beta_{j,t-q}^s \mathcal{S}_{j,t}^s + \varepsilon_{j,t}, \quad s \in \{cf, r\}$$

where δ_i and τ_t are fund and quarter fixed effects, $\mathbf{X}_{j,t}$ are controls, $\beta_{j,t}^s$ denotes the beta associated with cash flow or discount rate news from equation (13), and $\mathcal{S}_{j,t}^s$ represents the skill dummy with respect to cash flow or discount rate news in the sense of equation (18) or (19), respectively. The first two columns report the results for a sample that only includes partially skilled funds in the spirit of equations (25) and (26). Columns 3 and 4 report the results when skill in equation (13) is estimated on changes in holdings that are not zero only. The sample period is 2003 to 2021. t -statistics are in parentheses and estimates followed by ***, **, or * are statistically significant at the 1%, 5%, or 10% levels, respectively. Standard errors are clustered at the fund level.

	Partially skilled funds		Non-zero changes	
	$\beta_{j,t}^{cf}$	$\beta_{j,t}^r$	$\beta_{j,t}^{cf}$	$\beta_{j,t}^r$
$\beta_{j,t-1}$	-0.0371* (-1.90)	-0.0162 (-0.69)	-0.0414* (-1.71)	0.0058 (0.22)
$\beta_{j,t-2}$	-0.0681*** (-3.79)	-0.0209 (-1.18)	-0.0730*** (-3.34)	-0.0242 (-1.28)
$\beta_{j,t-3}$	-0.0474*** (-2.80)	-0.0153 (-0.73)	-0.0435** (-1.97)	0.0156 (0.58)
$\beta_{j,t-4}$	-0.0239 (-1.35)	-0.0022 (-0.10)	-0.0337 (-1.56)	0.0134 (0.52)
$\beta_{j,t-1} \times \mathcal{S}_{j,t}^s$	0.0725* (1.79)	-0.0416 (-1.02)	0.0694* (1.94)	-0.0822** (-2.30)
$\beta_{j,t-2} \times \mathcal{S}_{j,t}^s$	0.0820** (1.96)	-0.0556 (-1.34)	0.0499 (1.38)	-0.0642* (-1.80)
$\beta_{j,t-3} \times \mathcal{S}_{j,t}^s$	0.1354*** (3.28)	-0.0411 (-1.05)	0.0737** (2.37)	-0.0990*** (-2.77)
$\beta_{j,t-4} \times \mathcal{S}_{j,t}^s$	0.0256 (0.62)	-0.0516 (-1.33)	0.0643** (2.06)	-0.0878** (-2.47)
$\mathcal{S}_{j,t}^s$	8.3176*** (18.09)	50.7958*** (18.12)	13.9393*** (16.05)	84.0967*** (16.48)
ln(Holdings)	0.4059 (1.24)	-0.7220 (-0.35)	0.9178* (1.91)	2.2327 (0.79)
ln(Fund size)	0.1081 (0.79)	-2.2663** (-2.42)	0.6274** (2.48)	-3.2201** (-2.02)
Intercept	-6.1828*** (-3.85)	-9.4755 (-0.94)	-14.3953*** (-5.60)	-35.2493** (-2.25)
Quarter-fixed Effects	✓	✓	✓	✓
Fund-fixed Effects	✓	✓	✓	✓
Observations	63'130	63'130	62'649	62'649
Adj. R^2	0.07	0.07	0.07	0.08

Essay III:

Entropic Market Timing^{*}

Tim Glaus[†] and Sascha Jakob[‡]

ABSTRACT

We propose an information theoretic approach to measure price efficiency of financial assets and aggregate markets. Our measures draw on the idea of return predictability and are directly linked to the weak-form efficiency of the Efficient Market Hypothesis. [Asness et al. \(2013\)](#) document strong persistence of value and momentum anomalies in various financial markets and across different asset classes. Our efficiency measures are able to time both value and momentum in equity returns by identifying inefficiencies in the subsets of assets that are driving value and momentum. The primary duty of the measures is to signal periods where active asset allocation into value or momentum based strategies can be profitable and when passive investment is preferable. We therefore provide timing measures for two of the most important and most established financial market phenomena. We also document that market efficiency is cyclical for the U.S. stock market and varies over time.

Keywords: Market efficiency, Market timing, Return predictability, Entropy, Value, Momentum

JEL classification Numbers: C40, C58, G14, G17

^{*}We are grateful for the fruitful comments of Julien Cujean, Thorsten Hens, and Philip Valta.

[†]University of Zürich, Department of Banking and Finance, Plattenstrasse 14, 8032 Zürich, Switzerland. Email: tim.glaus@uzh.ch

[‡]University of Bern, Faculty of Business, Economics and Social Sciences. Engehaldenstrasse 4, 3012 Bern, Switzerland. Email: sascha.jakob@unibe.ch, corresponding author

I Introduction

The notion that financial time series evolve according to a random walk at least dates back to French mathematician Louis Bachelier, the forefather of mathematical finance and a pioneer in the study of stochastic processes. According to the random walk theory a financial security's price movement is random and should not follow any discernible pattern. Consistent with the notion of the random walk theory is the efficient market hypothesis of Fama (1970) with the premise that investors react instantaneously to any informational advantage they have thereby eliminating arbitrage opportunities such that prices reflect all available information at any time. One direct implication thereof is that the more efficient a market is the more prices follow a random walk. When originally proposed, the efficient market hypothesis was supported by such strong evidence that its position as a central proposition in financial economics seemed untouchable. However, over the past 50 years a large body of academic research questioned the random walk property of security returns and the efficient market hypothesis no longer holds the impervious position it once did. A vast pool of research documents the existence of various anomalies in the market, such as value (Stattman (1980)) or momentum (Carhart (1997)), which are direct violations of the idea of efficient markets.

To this day, the question to what extent markets are efficient is still fiercely debated. If markets are efficient, active investing should not be profitable over passive investing. We contemplate the market is an aggregate representation of its constituents and therefore, market efficiency is a function of the constituents' efficiencies which vary both in the cross-section and over time. The existence of varying efficiency at the security level fosters active investment. By identifying security-specific inefficiencies, active investing related to a security or set of securities may be temporarily profitable.

Research on market efficiency strongly focuses on the role of information in financial markets. Therefore, a natural approach to measure a financial asset's efficiency should be to resort to concepts originating from information theory. Information theory is the statistical and mathematical study of the quantification, storage, and communication of digital information. This definition encompasses many aspects of efficient markets. Specifically, large bodies of academic research on efficient markets examine what information (storage) or how much information (quantification) prices contain, or how fast (communication) new information is reflected in prices. Drawing on this notion, we propose an information theoretic approach to measure efficiency at the security level using information redundancy in time series in order to determine the degree of *information complexity* the financial asset's return series exhibits.

Formally, we understand information complexity as the pattern complexity in the filtration, \mathcal{F}_t , of a stochastic process, $\{X_t\}$. Put differently, we assess how much redundant information there is in \mathcal{F}_t and how often this redundant information appears. A time series that is informationally complex is efficient because it has minimal redundancy, does not follow any pattern and thus qualifies as a random walk. A time series with low information

complexity is governed by redundant information patterns that render the time series inefficient and predictable. Information complexity, and hence efficiency, can vary over time. This notion of time-varying efficiency emerges from the Adaptive Markets Hypothesis of [Lo \(2004\)](#), which allows for varying degrees of market efficiency.

By measuring efficiency using information complexity, we provide a new approach to identify sets of periods where a security is likely inefficient and some form of active investment may be profitable. Identifying and separating inefficient periods from efficient periods is the primary duty of information complexity. Due to the measure's connection to predictability, one natural and intuitive domain of application is the identification of periods where returns are predictable. By correctly identifying such periods, momentum strategies for equities can be timed with information complexity using past returns as signals in the spirit of [Moskowitz et al. \(2012\)](#).

In a broader context, information complexity also measures inefficiencies beyond time series predictability. One manifestation of inefficiency are anomalies. One of the most important anomalies in literature is that value firms outperform growth firms as documented by [Stattman \(1980\)](#). This outperformance, however, is only observed during some time periods. Information complexity helps identifying periods where value stocks are inefficient and value investing appears profitable. [Asness et al. \(2013\)](#) find consistent value and momentum return premia across eight diverse markets and asset classes. By providing a timing indicator for value and momentum, we cover two of the most studied capital market phenomena.

Our measures of information complexity are based on entropy in the spirit of [Shannon \(1948\)](#). We use the enhanced entropy estimator of [Gao et al. \(2008\)](#) and [Lopez De Prado \(2018\)](#) to approximate entropy by analyzing redundant patterns in time series of returns. We lay out some benefits of entropy based efficiency measures in chapter [A](#) of section [II](#). Using entropy estimates we construct measures of absolute efficiency and relative efficiency which we call *absolute information complexity* (AIC) and *relative information complexity* (RIC), respectively. Absolute information complexity measures the degree of efficiency of a security relative to that of a random Gaussian IID process. This process is unpredictable and fully efficient. Relative information complexity measures a security's efficiency relative to the efficiency of a benchmark or any other financial asset.

We start by identifying the optimal parameters to estimate the entropy of asset returns knowing that entropy of a Gaussian IID process has a closed form solution. We then use these optimal parameters for all entropy estimations of financial assets. This approach also warrants comparability of efficiency across securities. We estimate rolling information complexity measures for the entire CRSP universe and for various futures and ETFs representing different asset classes. To examine the overall efficiency in the cross-section of U.S. stocks, we compute a value weighted average of individual securities' efficiencies over time for the CRSP universe. We show that aggregate market efficiency varies over time substantially. This cyclical nature of efficiency is consistent with the Adaptive Markets Hypothesis of [Lo](#)

(2004) and the findings of [Ito et al. \(2014\)](#), [Ito et al. \(2016\)](#), and [Alvarez-Ramirez et al. \(2012\)](#). We also document that the speed with which these cyclical fluctuations in efficiency happen has increased substantially in recent years. To warrant that our measures are valid, we also show that our efficiency measures capture periods of high and low aggregate market efficiency very similarly to efficiency measures of [Hou and Moskowitz \(2005\)](#), which are not based on entropy.

We then provide market timing applications of our efficiency measures for value and momentum, two of the most important and prominent anomalies in then literature, as highlighted by [Asness et al. \(2013\)](#). First, using relative information complexity we demonstrate that our measures are able to time periods where value firms outperform. We analyze various ETFs that invest in value stocks exclusively and do so for different firm sizes. By actively investing in value stocks when value firms are inefficient and invest in the S&P 500 else, we show that in all cases the Sharpe ratio increases compared to a buy and hold strategy in the underlying value asset only. The Sharpe ratio increases in 90% of cases when the value timing strategy is benchmarked against a buy and hold strategy in the S&P 500. Further, we show that our strategies yield significant and positive alphas. This suggests that relative information complexity is able to time value stocks. We also show that measures of relative information complexity fare substantially better at identifying periods where value stocks are inefficient compared to other efficiency measures.

Second, we show that measures of absolute information complexity can be helpful in timing short positions in equities by drawing on time series momentum of [Moskowitz et al. \(2012\)](#). Time series momentum requires return predictability and absolute information complexity measures pattern redundancy in returns. As such, we use absolute information complexity to time periods where return patterns are predictable. If a time series is considered predictable at a certain point in time and this predictability is mainly driven by redundant patterns in negative returns, opening a short position is recommended. We use 42 futures contracts on equity indices, currency pairs, government bonds and commodities to examine the ability of absolute information complexity to time short selling. We find that information complexity helps improving Sharpe ratios by timing short positions for 90% of futures contracts on equity indices and for almost 80% of futures on currency pairs, compared to a buy and hold strategy in the underlying futures contract only. Overall, we show that our efficiency measures can be used to identify profitable investment opportunities in value and momentum.

In general, our paper contributes to the discussion about efficient markets invoked by [Fama \(1970\)](#). We revisit the efficiency of capital markets over time using information complexity and our measure draws on the idea of predictability (or lack thereof) and randomness of prices, an early extension of the efficient market hypothesis (see eg. [LeRoy \(1973\)](#); [Lucas \(1978\)](#); [Lo and MacKinlay \(1988\)](#); [Lo \(1991\)](#)). In detail, our paper contributes to several domains of the Efficient Market Hypothesis.

First, our paper contributes to the literature on market timing. One strand of literature

proposes market timing strategies due to predictability in assets. [Copeland and Copeland \(1999\)](#) show that changes in the Market Volatility Index (VIX) of the Chicago Board Options Exchange are statistically significant leading indicators of daily market returns. [Tang and Whitelaw \(2011\)](#) document predictable time variation in stock market Sharpe ratios. More recently, [Mascio et al. \(2020\)](#) use machine learning based forecasts to time the market. We add to this literature by establishing entropy based efficiency metrics that draw on predictability to time the market. The most prominent strand of literature on market timing is devoted to market timing ability of mutual funds. [Teynor and Mazuy \(1966\)](#), [Merton \(1981\)](#), and [Henriksson and Merton \(1981\)](#) lay the groundwork and develop models to infer the timing ability of mutual funds. Our measures can potentially also be used to assess the timing ability of mutual funds by examining whether mutual funds tend to perform better in periods of low market efficiency.

Second, we add to the literature that examines the performance consistency of active investment. [Leippold and Rüegg \(2019\)](#) find that active investing outperforms passive investment in some periods and underperforms during other periods. We provide a potential novel indicator for the timing of active and passive investing and thereby add to this strand of literature which generally resorts to return dispersion. For instance, [Stivers and Sun \(2010\)](#) use return dispersion to identify periods that foster active investments by engaging into value and momentum strategies. We propose a measure that is more directly related to market efficiency and thus a more natural candidate to identify periods where active value and momentum timing excels passive investments.

Third, we contribute to the literature of price efficiency measures. Conventional measures include intraday return predictability (see eg. [Hasbrouck and Ho \(1987\)](#); [Chan and Fong \(2000\)](#); [Chordia et al. \(2005\)](#), [Heston et al. \(2010\)](#)), variance ratios (see eg. [Lo and MacKinlay \(1989\)](#); [Andersen et al. \(2001\)](#); [Charles and Darné \(2009\)](#)), Hasbrouck pricing errors (see eg. [Hasbrouck \(1993\)](#); [Boehmer and Kelley \(2009\)](#)), and delay measures (see eg. [Hou and Moskowitz \(2005\)](#); [Busch and Obernberger \(2017\)](#)), among others. Our paper introduces new measures to the existing repository of price efficiency measures, which draw on the notion of entropy.

Finally, our paper contributes to the research that adopts the idea of entropy to financial time series and efficient markets. [Shannon \(1948\)](#) establishes entropy in information theory as a measure of uncertainty. The first proposition to use concepts of entropy to study the efficiency of financial markets dates back to [Gulko \(1999\)](#), who proposes the *Entropic Market Hypothesis*, according to which the entropy of consensus beliefs about the future price changes is a suitable measure of uncertainty. Consequently, an informationally efficient market displays the highest entropy and maximum-entropy consensus beliefs prevail. Drawing on the notion that entropy measures disorder of a system, and hence randomness, a series of papers use different entropy estimators to measure market efficiency. [Oh et al. \(2007\)](#) use entropy to measure the randomness in time series in foreign exchange markets and from that, infer how efficient markets are. In a similar spirit, [Risso \(2008\)](#) uses con-

cepts of entropy to measure informational efficiency of stock market indices while [Pascoal and Monteiro \(2014\)](#) use entropy to measure the degree of unpredictability of financial time series. Likewise, [Kristoufek and Vosvrda \(2014\)](#) understand entropy as a measure of complexity of a system and use it to determine market efficiency of stock market indices around the globe. Motivated by evolutionary finance theories and the adaptive market hypothesis, [Alvarez-Ramirez et al. \(2012\)](#) use approximate entropy to measure market efficiency over time by looking at patterns contained in price changes. Similarly, [Zunino et al. \(2009\)](#) show that concepts of entropy and pattern detection is correlated with market efficiency. Other recent literature uses entropy to measure market risk ([Pele et al. \(2017\)](#)), irregularities in financial time series ([Pincus and Kalman \(2004\)](#)), and pattern detection in financial machine learning ([Lopez De Prado \(2018\)](#)). Our paper establishes entropy based measures to identify inefficient periods at the security level which helps time active asset allocation and passive investment approaches.

The remainder of this paper is structured as follows. Section II introduces our efficiency measures and lays out its construction and estimation. Section III presents the empirical results. Specifically, we analyze aggregate market efficiency over time and compare it to aggregate market efficiency using conventional efficiency measures. We then demonstrate the market timing ability of our measures by timing inefficiencies at the asset level. Section IV concludes the paper.

II Methodology

Research on market efficiency strongly focuses on the role of information in financial markets. Therefore, a natural approach to measure market efficiency is to resort to concepts originating in information theory. Information theory is the statistical and mathematical study of the quantification, storage, and communication of digital information. This definition is well reconcilable with the aspects of efficient markets. Specifically, large bodies of academic research examine what information (storage) or how much information (quantification) prices contain, or how fast (communication) new information is reflected in prices. Since the literature in financial economics is usually not well accustomed to the concepts of information theory, what follows is a faithful exposition of the ideas and measures we use to assess the inefficiency of a financial asset.

A Information complexity

In order to identify whether a financial asset is inefficient at a given point in time we estimate the information complexity of securities' return streams as a measure of price efficiency. Formally and in its most general form, we understand information complexity as the pattern complexity in the filtration, \mathcal{F}_t , of a stochastic process, $\{X_t\}$. Put differently, we assess how much redundant information there is in \mathcal{F}_t and how often this redundant information appears. It reveals how "complex and efficient" the *current* information set is.

Importantly, the information contained in \mathcal{F}_t encompasses all public historic information that has been incorporated into the price stream up to time t , i.e., past returns, fundamental values, and trends. Therefore, our measure of information complexity is directly linked to the weak-form efficiency of the Efficient Market Hypothesis. In order to estimate the information complexity of stocks we measure the compression rate of return patterns in a time series. In information theory, the compression rate is well defined and is equivalent to the notion of entropy.

The adoption of entropy to information theory from classical thermodynamics dates back to [Shannon \(1948\)](#). The higher the entropy, the lower the redundancy and the greater the informational content in a message. A return series that has high entropy thus exhibits low redundancy with no discernible patterns rendering any attempt to predict the return stream futile. This stock is decompressed and efficient. At the other end of the spectrum, a return series that is characterized by low entropy contains redundant information giving rise to predictability of its return pattern making the stock compressed and inefficient.

A large body of academic research has looked into various measures of market efficiency. We argue that entropy based measures have some advantages compared to classical measures to assess the efficiency of a financial security. First, the normality assumption of stock returns is often challenged by the data. The concept of entropy can circumvent this issue as it measures disorder of the time series without imposing any constraints on the theoretical probability distribution (see eg. [Bentes et al. \(2008\)](#); [Darbellay and Wuertz \(2000\)](#)). Second, entropy is capable of detecting nonlinear dependence within the return process (see eg. [Maasoumi and Racine \(2002\)](#)) and may therefore be capable of measuring more complex patterns and forms of predictability and in effect, efficiency. Finally, high values of entropy are related to randomness in the evolution of stock prices and are thus naturally and intuitively coherent with the Efficient Market Hypothesis (see eg. [Zunino et al. \(2009\)](#)).

B Enhanced Kontoyiannis entropy estimator

Many estimators for entropy have been proposed in the literature. In order to estimate the information complexity of security returns we use compression algorithms to compute the length of repeating patterns in the data sequence. The most commonly used algorithms in this field date back to [Ziv and Lempel \(1978\)](#) and decompose a data sequence into a set of non-redundant substrings. The larger the set of unique substrings relative to the total length of the message the more complex a message appears and the higher its entropy. Intuitively, as the size of the set containing the Lempel-Ziv substrings increases, the patterns in the data sequence are more unique and shorter.

The estimators of [Kontoyiannis et al. \(1998\)](#) and [Gao et al. \(2008\)](#) draw on the idea of Lempel-Ziv algorithms and assess a data sequences' complexity by searching at each position in the data sequence for the longest matching pattern subsequent to that position in the data sequence with respect to the segment of the data sequence of a certain length just prior to that position in the data sequence. Specifically, using a return sequence, $\{R^n\}$,

with length n , for every position in $\{R^n\}$ and window length $w \geq 1$, find the length ℓ of the longest return pattern $r_i^{i+\ell}$ in the return sequence $\{R^n\}$ starting at position i that also appears in the window r_{i-w}^{i-1} preceding position i . Formally, define:

$$\begin{aligned} L_i^w &= L_i^w\{R_1^n\} = L_i^w(x_{i-w}^{i+w-1}) \\ &= 1 + \max\{0 \leq \ell \leq w : x_i^{i+\ell-1} = x_j^{j+\ell-1} \text{ for some } i-w \leq j \leq i-1\}, \end{aligned}$$

which corresponds to 1 plus the longest match ℓ . [Ornstein and Weiss \(1993\)](#) establish that:

$$\lim_{w \rightarrow \infty} \frac{L_i^w}{\log_2(w)} = \frac{1}{H}, \quad (1)$$

where H corresponds to the entropy of the data sequence $\{R_1^n\}$. [Kontoyiannis et al. \(1998\)](#) use the reciprocal of equation (1), $\log_2(w)/L_i^w$, to estimate the entropy of a data sequence. In order to reduce variance and make more efficient use of the data, [Gao et al. \(2008\)](#) suggest that the average of various match-lengths, L_i^w , at different positions be taken.

To that end, we adopt the entropy measure proposed by [Gao et al. \(2008\)](#) and refined by [Lopez De Prado \(2018\)](#) with an expanding window. Given some return sequence, $\{R^n\}$, of length n , for every position $i > 1$ in $\{R^n\}$ up to position $i = \lfloor n/2 \rfloor + 1$, find the length $\max\{\ell^i\}$ of the longest return pattern $r_i^{i+\ell}$ starting at position i that also appears in the window r_1^{i-1} , i.e., the entire return sequence preceding position i . Rather than using a shifting window, this measure uses an expanding window that expands as the position i increases and has length $i - 1$ rather than a constant length. In this case for a return sequence of length n the entropy estimator H is defined as:

$$\hat{H} = \left[\frac{1}{\lfloor n/2 \rfloor} \sum_{i=2}^{\lfloor n/2 \rfloor + 1} \frac{L^i}{\log_2(i)} \right]^{-1} \quad (2)$$


where L^i denotes $1 + \max\{\ell^i\}$ for position i . The intuition is that if the average length of the longest matches is large the data sequence consist of only a few longer return patterns and thus displays low entropy. On the other hand, if the average length is short the message consist of many short and unique return patterns casting the return sequence to appear complex and to display high entropy ultimately rendering it unpredictable and thus efficient. The estimator is illustrated in more detailed in appendix [A](#).

C Entropy estimation

Estimating entropy of stock returns using equation (2) involves the selection of hyper parameters. The first hyper-parameter is the number of past returns to be considered when estimating entropy, i.e., the length of our data sequence. This is called *message length*. As the message length increases we add more historic information. However, this may happen at the cost of a lagged outcome. Older returns overweight recent returns and the measure might not reflect current levels of entropy in financial assets and is not myopic enough to capture current levels of inefficiency. Choosing a small message length to obtain a more

Figure 1: Parameter space for number of past returns and return quintiles

This figure shows the function value for equation (4) for all possible pairs of message length and alphabet size over a finite parameter space. Parameter values for message length are between 20 and 120 return observations and the size of the alphabet ranges between 2 and 10 return quintiles.



myopic measure, on the other hand, might result in too little information such that it underestimates entropy.

The second hyper-parameter is the number of values returns are assigned to. Security returns are continuous by nature such that each return represents a unique value. The continuous nature maps every return to a unique return pattern in itself which completely defeats the purpose of our measure which is to find redundant patterns. Therefore, return series must be discretized such that each return observation can be assigned to a specific value of a finite set of values. The set of possible values (bins) is called the *alphabet*. We follow [Lopez De Prado \(2018\)](#) and use quantile encoding. Quantile encoding allows for specifying any number of quantiles and assigns each return to the quantile it belongs to. The size of the alphabet then maps to the number of quantiles.

In order to set the number of past return observations and the size of the set of values returns can be assigned to we resort to the fact that the entropy of a Gaussian IID process,

such as White Noise, maps to :

$$H_{WN} = \frac{1}{2} \log_2(2\pi e\mathbb{V}) = 1.42, \quad (3)$$

where H_{WN} denotes the entropy of White Noise. This allows us to identify the correct parameter values by bench marking a White Noise process for a specific tuple of parameter values. Specifically, we maximize the negative absolute difference between the average of rolling entropy estimates from equation (2) over a sufficiently long white noise process and the true entropy of White Noise over the number of past returns n and return quintiles q for a finite parameter space Θ :

$$\arg \max_{n \in \Theta, q \in \Theta} - \left| (T - n + 1)^{-1} \sum_{t=n}^T \left[\left(\frac{1}{\lfloor n/2 \rfloor} \sum_{i=2}^{\lfloor n/2 \rfloor + 1} \frac{L^i}{\log_2(i)} \right)^{-1} \right]_{t=t} - H_{WN} \right| \quad (4)$$

Put differently, if we estimate the entropy of the white noise process, $\hat{H}_{WN,t}(n, q)$, at each time t , the average of these estimates should correspond to the true entropy of a white noise process. Figure 1 depicts the negative absolute values of $(T - n + 1)^{-1} \sum_{t=n}^T \hat{H}_{WN,t}(n, q) - 1.42$ for 20 to 120 return observations (1 month to 6 months of past return data) and alphabets containing 2 to 10 return values (binary to deciles). For the number of return quintiles the strongest support exists when the set contains 4 values. This segmentation of returns appears to be economically sensible: strongly negative (1), weakly negative (2), weakly positive (3), and strongly positive (4). Figure 1 shows that the function value decreases sharply as the number of quintiles change. To the contrary, the sensitivity of the function value with respect to the number of past returns is weak and the function value is roughly constant between 20 and 120 past return observations. Put differently, we can alter the number of past return observation without losing substantial accuracy of the entropy estimate. This has several advantages. First, using only 1 month of past return observations provides a timely and contemporaneous estimate of efficiency. This can be helpful when estimating efficiency for a short and well defined period in time or when it is assumed that efficiency changes quickly. To assess the efficiency during long-lasting large scale economic events such as the global financial crisis, using the past 6 months of returns may provide a more suitable estimate of entropy to measure efficiency during that entire time period. Secondly, some financial and economic time series are only available at a lower frequency, i.e. monthly. In order to obtain a timely measure of efficiency it is required that that number of past observations be short. The ability to reduce the number of observations is a remedy to this problem of low data frequency. Overall, the function value is maximized when the number of past returns, n , is 50 and returns are assigned to 4 possible values, q . In this case, the entropy estimate of a White Noise process, \hat{H}_{WN} is 1.40 which we consider sufficiently close to 1.42.

D Efficiency metrics

The entropy estimator in equation (2) is a deterministic number. Without further notion the entropy estimate of a financial asset is uncoupled to efficiency. In order to assess the information complexity and thus efficiency of an asset's return stream we propose to express entropy relative to that of a benchmark or compare it to some well defined entropy of a fully efficient process.

a) Absolute Information Complexity

Similar to [Boehmer and Kelley \(2009\)](#), we draw on the idea that efficient time series follow a random walk and benchmark a time series against a random walk. Specifically, we express a stock's information complexity relative to that of a random walk with Gaussian white noise. As its name suggests, a random walk should exhibit high complexity as the process is by definition random and hence, does not follow any discernible pattern and is considered efficient. Using the same logic, an information complexity that is inferior to that of a random walk suggests a certain degree of inefficiency. We therefore measure how far away the complexity of a security is at a given point in time, t , from that of a random walk. This is a measure of absolute efficiency. The lower the information complexity relative to that of a random walk the less efficient a time series is. We make use of the fact that the entropy of a Gaussian IID White Noise process is known but adjust for the error we introduce by estimating the entropy of a discretized White Noise process and for allowing only a finite set of values for the parameters message length and alphabet. In effect, the threshold for full efficiency is conditional on the number of past returns, n , and the set of values, q , returns are assigned to. The threshold then maps to the average of entropy estimates of a sufficiently large White Noise process using n past returns and an alphabet containing q possible values:

$$\hat{H}_{WN}(n, q) = \frac{1}{T} \sum_{t=1}^T \hat{H}_{WN,t}(n, q) = \mathbb{E} \left[\hat{H}_{WN,t}(n, q) \right], \quad (5)$$

where $H_{WN}(n, q)$ denotes the entropy estimate of a White Noise process at time t using n past returns and q return quintiles for discretization. Our measure for Absolute Information Complexity for security i at time t maps to:

$$AIC_{i,t} = 1 + \frac{\min\{\hat{H}_{i,t}(n, q) - \hat{H}_{WN}(n, q), 0\}}{\hat{H}_{WN}(n, q)} \quad (6)$$

where $\hat{H}_{i,t}$ denotes the entropy estimate of asset i at time t using the enhanced Kontoyiannis Entropy estimator from equation (2). Note that the entropy estimates $\hat{H}_t(n, q)$ are time dependent and change over time. This allows the entropy of a financial asset to temporarily be higher than the efficiency threshold. However, the mean of an asset's entropy process over time cannot be higher than the mean of the mean entropy of a purely random White Noise process for the same pair of n and q , which acts as our threshold.

We also need a threshold for inefficiency, i.e., a level below which a financial asset is

no longer considered efficient. The simplest indicator is the maximum efficiency threshold itself. Whenever a time series is below the expected entropy of the White Noise process, a stock is considered inefficient:

$$I_{i,t} = \mathbb{1} (AIC_{i,t} < 1) \quad (7)$$

where $\mathbb{1}$ denotes the indicator function.

Another approach is to use the distribution of the entropy estimates of the White Noise process. Consider the White Noise entropy process, $\{\hat{H}_{WN,t}(n, q)\}$, with mean $\hat{H}_{WN}(n, q)$. This entropy process has variance $\mathbb{V}[\{\hat{H}_{WN,t}(n, q)\}]$. We use the corresponding scaled (to be in line with $AIC_{i,t}$) standard deviation as our unit of deviation from efficiency:

$$\hat{\sigma}_{WN}(n, q) = \frac{\sqrt{\mathbb{V}[\{\hat{H}_{WN,t}(n, q)\}]}}{\hat{H}_{WN}(n, q)} \quad (8)$$

This metric maps a one standard deviation change in the White Noise entropy process, $\{\hat{H}_{WN,t}\}$ to a decrease in $AIC_{i,t}$. We can then define inefficiency as:

$$I_{i,t} = \mathbb{1} (AIC_{i,t} < 1 - k\hat{\sigma}_{WN}(n, q)), \quad (9)$$

where k denotes the level of confidence. A one $\hat{\sigma}_{WN}(n, q)$ deviation below the maximum efficiency threshold then corresponds to an absolute efficiency ratio below which the entropy of a fully efficient process only deviates roughly 15% of the time. For two $\hat{\sigma}_{WN}$, this is the case for roughly 2.5% of observations.

b) Relative Information Complexity

As a second measure of efficiency we propose a relative measure of efficiency which we call relative information complexity (RIC). The efficiency is measured against that of some suitable benchmark. This benchmark can be an equity index, a portfolio of assets, or an individual security. Specifically, we run rolling regressions of the difference between the asset's entropy and the benchmark's entropy on the entropy of the benchmark:

$$\hat{H}_{i,t}(n, q) - \hat{H}_{BM,t}(n, q) = \gamma_{i,t}\hat{H}_{BM,t}(n, q) + \varepsilon_{i,t} \quad (10)$$

and use the regression coefficient $\gamma_{i,t}$ as our measure of Relative Information Complexity. The intuition of the rolling regressions in equation (10) is as follows. If security i and the benchmark BM are equivalent in terms of efficiency, $\hat{H}_{i,t}(n, q) - \hat{H}_{BM,t}(n, q)$ is equal to zero. In this case $\gamma_{i,t}$ should not be different from zero. If security i is more efficient than the benchmark BM , $\hat{H}_{i,t}(n, q) - \hat{H}_{BM,t}(n, q)$ is positive. In this case $\gamma_{i,t}$ should be larger than zero. If security i is less efficient than the benchmark BM , $\hat{H}_{i,t}(n, q) - \hat{H}_{BM,t}(n, q)$ is negative. In this case $\gamma_{i,t}$ should be smaller than zero. We can then analyse whether $\gamma_{i,t}$ is significantly different from zero in period t and use it as an indicator of higher or lower efficiency compared to the benchmark. An additional benefit of this approach is that it yields a smoothed measure of efficiency that is less prone to sporadic extreme entropy

values. Furthermore, the measure indicates how strongly efficiency changes relative to that of the benchmark and as such also captures another dimension of efficiency. Finally, the measure allows to compare the degree of efficiency of two distinct financial assets in an intuitive way. The simplest comparison is that of γ directly such that we have:

$$I_{(i,j),t}^+ = \mathbf{1}(\gamma_{i,t} > \gamma_{j,t}) \quad (11)$$

$$I_{(i,j),t}^- = \mathbf{1}(\gamma_{i,t} < \gamma_{j,t}) \quad (12)$$

where $I_{(i,j),t}^+$ denotes that asset i is more efficient than asset j at time t and $I_{(i,j),t}^-$ indicates the opposite.

To accommodate potential estimation error in γ , one can incorporate standard errors of the γ -estimate as a remedy. In effect, if the confidence interval of $\gamma_{i,t}$ does not overlap with the confidence interval of $\gamma_{j,t}$, asset i and j have different efficiency. Formally:

$$I_{(i,j),t}^+ = \mathbf{1}(\gamma_{i,t} - k\hat{se}(\gamma_{i,t}) > \gamma_{j,t} + k\hat{se}(\gamma_{j,t})) \quad (13)$$

$$I_{(i,j),t}^- = \mathbf{1}(\gamma_{i,t} + k\hat{se}(\gamma_{i,t}) < \gamma_{j,t} - k\hat{se}(\gamma_{j,t})) \quad (14)$$

where k denotes the level of confidence.

E Data

All price data on stocks is obtained from the Center for Research in Security Prices. Balance sheet data is from Compustat and data on futures contracts and ETFs is acquired through Refinitiv Eikon and Datastream. The construction of variables is documented in the section where they are used.

III Results

This chapter presents the results of our empirical exercises. Section A examines the efficiency of the cross-section of U.S. stocks over time using our efficiency measures and compares the results to traditional efficiency measures. The remainder of the chapter provides applications for both relative and absolute information complexity. Specifically, section B applies our measure of relative efficiency to identify value cycles and demonstrates how one can time value stocks using an active-passive investment strategy. Section C uses the absolute efficiency measure to time short positions drawing on a time series momentum approach.

A Cross-sectional analysis of stock price efficiency

This section analyses the efficiency in the cross-section of US stocks. Specifically, we consider all stocks encompassed in the CRSP database between 1985 and 2020 and estimate our measures of efficiency. In a second step, we compare our estimates to the price delay measure of [Hou and Moskowitz \(2005\)](#).

a) Information Complexity

Table I: Cross-sectional efficiency

This table reports the mean, standard deviation, minimum and maximum of entropy and various efficiency measures for the CRSP stock universe between 1985 and 2020. Entropy is estimated on a rolling daily basis with the estimator in equation (2) using 50 returns and 4 return bins. AIC denotes the absolute information complexity, as defined in section II. Our measure of relative information complexity, Gamma, is the regression coefficient of equation (10) estimated with rolling regressions using 3 months of past entropy observations. We estimate relative information complexity with the S&P 500 and the value-weighted entropy index as benchmarks.

Efficiency Measure	Mean	Std. Dev.	Min	Median	Max
Entropy	1.24	0.22	0.31	1.29	1.77
AIC	-0.16	0.22	-1.09	-0.11	0.37
$\hat{H}_{i,t} - \hat{H}_{S\&P500,t}$	-0.07	0.26	-1.41	-0.05	1.21
$\hat{H}_{i,t} - \hat{H}_{EntropyIndex,t}$	-0.04	0.22	-1.15	0	0.83
$RIC_{i,S\&P500,t}$	-0.06	0.12	-0.79	-0.05	0.58
$RIC_{i,EntropyIndex,t}$	-0.03	0.11	-0.77	-0.01	0.36

We start by estimating the entropy estimator in equation (2) using daily returns for each security encompassed by the CRSP sample at any given day. We use the optimal parameters identified in section II.C. Specifically, we estimate entropy using the past 50 daily returns and discretize returns into 4 bins. We therefore obtain a daily rolling entropy estimate for each stock in our cross-section. In order to obtain a more timely entropy estimate, we reverse the time series of discretized returns.

Table I reports the most important cross-sectional statistics for entropy, absolute information complexity, and relative information complexity for the CRSP stock sample between 1985 and 2020. The mean entropy estimate over the whole sample amounts to 1.24 while the median is 1.29. This suggests that the median firm is one standard deviation below the expected entropy of a white noise process while the average is roughly 1.5 standard deviations below the mean entropy of a white noise process. This shows that firms are on average and across time less efficient than a fully random process.

In order to measure the efficiency of the cross-section of stocks as a whole and especially over time, we construct an entropy index. To that end, at any given day we compute the value-weighted average of all entropy estimates across all assets for which an entropy estimate using equation (2) is available. The cross-sectional entropy average provides an economically more sensible estimate of aggregate cross-sectional efficiency compared to measuring the entropy of a representative equity index. The entropy index is an average of the entropy estimates of individual securities while the entropy of an index is as single entropy estimate for a time series that is a function of its constituents' stock returns. Naturally, an equity index therefore contains information of many securities and is by definition more complex than a time series of a single security. We then form our measures of absolute and relative information complexity (AIC and RIC) for the aggregate entropy index.

Figure 2 displays the efficiency of the entropy index over time for the period 1985 to 2020. The top panel reports the absolute information complexity (AIC) for the value-

Figure 2: Efficiency of the cross-section of U.S. stocks between 1985 and 2020

This figure shows the efficiency of the value-weighted entropy index of the CRSP sample between 1985 and 2020. To estimate the index, at any given day we compute the value-weighted average of all entropy estimates across all assets for which an entropy estimate using equation (2) is available. The top panel shows the rolling daily Absolute Information Complexity (AIC) of the index over time where the dotted black line indicates the expected efficiency of a white noise process while the red line represents a scaled one standard deviation decrease of the white noise entropy process. The middle panel exhibits the rolling volatility of the daily entropy process. The bottom panel shows the estimate for relative information complexity (RIC) using equation (10) with the entropy of a white noise process as benchmark. Each shade of color represents one standard deviation of the white noise entropy process. Entropy is estimated using a reversed time series using the past 50 returns and 4 return bins.



weighted entropy index. The purple line represents the yearly average of the daily AIC while the light blue reports the monthly average. Most of the time the annual AIC average is above the one standard deviation margin of the white noise entropy process. Ever since the global financial crisis however, the aggregate efficiency in the cross-section appears to decrease, especially so in recent times. This is reconcilable with the stance that markets may no longer reflect fundamental valuations and are becoming inefficient. We arrive at the same conclusion when analyzing the relative information complexity (RIC) over time from equation (10) in the bottom panel of figure 2. The yearly average RIC of the entropy index with respect to a white noise process is mostly within the one standard deviation confidence interval of a white noise entropy process. However, ever since 2012 the yearly mean is outside of this confidence interval which suggests that the cross-section is becoming less and less efficient on average.

Turning to higher frequencies, it emerges that the average efficiency of the cross-section varies strongly and frequently. The variation of efficiency is captured in the middle panel of figure 2. This panel shows the rolling volatilities of the daily entropy estimates of the value-weighted entropy index for a window of 20, 60, and 120 trading days (1 month, 1 quarter, and 6 months). Low volatility indicates that the level of efficiency remains relatively unchanged and prices seem to incorporate information at a constant rate. High volatility suggests that the efficiency alters in short intervals and information processing in the market is flawed. Coupling the volatility with the AIC estimate allows approximating the overall efficiency at a given point in time. If AIC is high and volatility low the time series is in its most efficient state. If AIC is low and persistent the time series is in its most inefficient state. Mixed states suggest semi-efficiency where markets may be inefficient in that prices contain noise but are efficient in that markets still work well enough to render sophisticated investments profitable.

Our sample starts around the oil crisis in the late 80's. During that period the aggregate yearly absolute information complexity resides below the one standard deviation threshold of a white noise entropy process. Put differently, the yearly average of daily cross-sectional AIC was lower than the level below which the daily entropy of a white noise process resides only 15% of the days. At the same time, the volatility of the entropy process is very high indicating that aggregate market efficiency was low during the oil crisis. Specifically, during the oil crisis, the aggregate U.S. market exhibits a degree of efficiency as low as 60%. After the oil crisis during the 90's the efficiency is fairly constant and the aggregate AIC is fairly high, hovering above 90% suggesting that during the 90's the cross-section of stocks were, on average, efficient. This period of efficiency faces an end with the burst of the dotcom bubble when the yearly AIC average falls below the one standard deviation threshold in 2001 and volatility increases. During that period, AIC drops close to 80% efficiency for brief intervals, which maps to more than 2 standard deviations of a white noise entropy process. While the market appears less efficient in wake of the dotcom bubble, aggregate AIC increase and volatility decrease prior to the financial crisis suggesting that markets were efficient. During this period, the aggregate AIC is between 90% and 95%. The inefficiency around the oil

crisis and subsequent to the dotcom bubble, as well as the high efficiency prior to the global financial crisis are consistent with the findings of [Alvarez-Ramirez et al. \(2012\)](#) who use approximate entropy measures to quantify the efficiency of the Dow Jones Industrial index (DJI) over time as an approximation for market efficiency. Specifically, they find that in the period around the oil crisis efficiency was as low as 55% relative to full efficiency and close to 75% after the dotcom bubble. Subsequent to the dotcom bubble and prior to the global financial crisis they show that the efficiency of the DJI is as high as 95%. Our results are also in line with those of [Risso \(2008\)](#) who measures the evolution of informational efficiency for the US stock market by using local entropy and symbolic time series analysis. He shows that lower efficiency is present around market crashes such as the oil crisis and the dotcom bubble.

Starting in 2006 and 2007 however, the yearly AIC falls below the threshold anew but recovers during the global financial crisis. With the arrival of the crisis, then patterns shift completely. Average aggregate absolute information complexity decreases steadily and volatility grows continuously. In 2012 the level of AIC falls short of the one standard deviation threshold and has not climbed this threshold ever since. In perspective, during the period between the oil crisis up to the year before the financial crisis the average of yearly aggregate absolute information complexity is above 93% while for the time after the financial crisis the same statistic amounts to less than 91% which is below the one standard deviation threshold of white noise efficiency which maps to an AIC of 92.8%. This average decreases further to slightly above 89% when only considering the last 5 years which is clearly below the one standard deviation threshold. The rolling one month volatility of aggregate entropy is 2.02% before the financial crisis and 3.34% thereafter. In relative terms, the post financial crisis entropy volatility is 65% higher, or 79% when considering the last 5 years only. This leads to the overall conclusion that the efficiency in the cross-section of US stocks has likely decreased ever since the recovery of the financial crisis and currently is at its lowest level since the oil crisis in the late 80s.

Overall, our results suggests that the degree of efficiency in the stock market varies strongly both across assets and across time. A market with a time-varying degree of efficiency was first introduced in [Lo \(2004\)](#)'s Adaptive Markets Hypothesis, according to which the degree of market efficiency varies over time due to individual adaptation to a changing environment. Our results suggests that at the aggregate level, the degree of efficiency is characterized by cyclical patterns of efficiency such that efficient and inefficient periods alternate. This is consistent with [Ito et al. \(2014\)](#), [Ito et al. \(2016\)](#), and [Alvarez-Ramirez et al. \(2012\)](#) who find that the U.S. stock market has evolved over time and the degree of market efficiency has cyclical fluctuations. We further document that the speed at which the degree of efficiency changes has increased ever since the financial crisis. Prior, the cycles of efficient and inefficient periods had longer periodicity which is consistent with findings of [Ito et al. \(2016\)](#). Finally, we document that the overall level of efficiency has decreases ever since the global financial crisis which is in line with the ongoing decoupling of market capitalization from fundamental values.

b) Comparison to price delay

To warrant that our efficiency measure is comparable to other efficiency measures in the literature we compare it to the price delay measure of [Hou and Moskowitz \(2005\)](#). This metric measures the speed with which stocks incorporate market wide information and as such, it measures an important dimension of efficiency. If prices incorporate new information only slowly it may be possible to achieve abnormal returns using publicly available information. This is a direct violation of the weak-form efficiency. As our measure of information complexity is directly linked to the weak-form efficiency as well, we deem the price delay measure of [Hou and Moskowitz \(2005\)](#) to be a economically sensible alternative.

To measure price delay, we closely follow [Hou and Moskowitz \(2005\)](#). Using 52 weeks of weekly returns, we regress the return of stock i at time t on the contemporaneous market return, $r_{m,t}$, and 4 weekly lags of it:

$$r_{i,t} = \alpha_i + \beta_i r_{m,t} + \sum_{n=1}^4 \delta_i^{(-n)} r_{m,t-n} + \varepsilon_{i,t} \quad (15)$$

If any of the lagged coefficients $\delta_i^{(-n)}$ is significant the stock adjusts to market news with a lag and is considered less efficient. [Hou and Moskowitz \(2005\)](#) propose the following three measures:

$$D1 = 1 - \frac{R^2(\delta_i^{(-n)} = 0, \forall n \in [1, 4])}{R^2}, \quad (16)$$

where the numerator is the R^2 from the regression in equation (15) when the coefficients of the lags are zero and the denominator is the R^2 from the unconstrained regression. Since $D1$ does neither distinguish between longer and shorter lags nor account for precision, the other two measures, $D2$ and $D3$ are:

$$D2 = \frac{\sum_{n=1}^4 n \delta_i^{(-n)}}{\beta_i + \sum_{n=1}^4 \delta_i^{(-n)}} \quad (17)$$

$$D3 = \frac{\sum_{n=1}^4 \left(n \delta_i^{(-n)} / se(\delta_i^{(-n)}) \right)}{(\beta_i / se(\beta_i)) + \sum_{n=1}^4 \left(\delta_i^{(-n)} / se(\delta_i^{(-n)}) \right)}, \quad (18)$$

where $se(\cdot)$ denotes the standard error of the coefficient estimate. While [Hou and Moskowitz \(2005\)](#) compute the measures by the end of June in a given year, we estimate the measures on a rolling basis. Specifically, we compute Wednesday-to-Wednesday returns for all securities and then run rolling regressions using the past 52 weekly return observations and obtain the measures for each week. To obtain monthly delay measures we simply take the average in a given month. A higher value of $D1$, $D2$, and $D3$ implies higher delay and in effect, less efficiency. We follow [Hou and Moskowitz \(2005\)](#) and use the CRSP value-weighted index as market returns.

Figure 3: Price delay of the cross-section of U.S. stocks between 1985 and 2020

This figure shows the average efficiency of the CRSP universe using the price delay measures of [Hou and Moskowitz \(2005\)](#) between 1985 and 2020. To estimate average efficiency, at any given month we compute the value-weighted average of all $D1$, $D2$, and $D3$ estimates across all assets for which an estimate using equations (16), (17), and (18) is available. We aggregate to a monthly level by taking the average of weekly estimates within a month. Panel A reports the value-weighted averages for all three measures. Panel B compares the delay measure $D1$ with the value weighted entropy index of the same CRSP universe used in figure 2. Panel C does the same exercise but compares the delay measures $D2$ and $D3$ with the value weighted entropy index. Entropy is estimated using a reversed time series using the past 50 returns and 4 return bins.



Figure 3 reports the three delay measures as a value-weighted average for the entire CRSP universe, similar to the value-weighted entropy index in figure 2. Panel A reports the three measures over time while Panel B and C compare D1, D2 and D3 with the value-weighted entropy index for the same securities. Overall, panel A shows that, while the three measures have different scales, efficiency measured as price delay changes over time. In panel B and C we directly compare the delay measures to the entropy index for the CRSP sample. It turns out that the entropy index is fairly consistent with the delay measures. A decrease in delay is characterized by an increase in entropy while an increase in delay is matched with a decrease in entropy. For instance, subsequent to the financial crisis the delay measures decrease while the entropy index increases. But ever since 2012 there appears to be a downward trend in efficiency. Overall, the entropy index decreases in the period from 2012 to 2020 and the delay measures increase.

While both measures tend to capture periods of more and less efficient markets in a similar fashion, advantages of entropy based measures crystallize. First, entropy measures can be applied to aggregate market indices while delay measures cannot because they use market returns as relevant news to which securities respond. Second and more importantly, inferring efficiency from entropy based measures provides an indicator of absolute efficiency. Delay measures have no reference point for absolute efficiency and are always relative to a market index. The lack of a reference point also implies no obvious threshold above which a delay measure indicates inefficiency and it impedes finding optimal lags or frequency when estimating regression (15). Inferring efficiency from delay measures is limited to relative efficiency and requires that a security be compared to the cross-section.

This section shows that our entropy based measures are capable of detecting different degrees of efficiency and are consistent with other measures used in the literature to determine price efficiency. The existence of inefficient and semi-efficient markets gives rise to the exploitation of these inefficiencies using active investment strategies. In the following chapter we show how our entropy based efficiency measures can help identify periods of inefficiency in order to engage in profitable active investment strategies and to identify periods where passive investment is recommended. We also use the price delay measures of [Hou and Moskowitz \(2005\)](#) for comparison.

B Timing value

In this section we use our measure of relative information complexity (RIC) to identify periods where active investment strategies following a value approach are profitable. We assume that in periods where predictability of value returns is high, value stocks outperform growth stocks and pursuing an active investment strategy is advantageous. Hence, investors actively time a long position in value assets when value stocks have a lower relative efficiency and are predictable or passively hold a long position in the market. A large body of work finds evidence for the superior performance of value investing strategies compared to growth strategies. [Fama and French \(1992\)](#), [Fama and French \(2012\)](#), and [Asness et al.](#)

(2013), among others, document a positive return differential in US and international markets. More recently, [Yara et al. \(2020\)](#) show that returns to value strategies are predictable in the time series and [Asness et al. \(2021\)](#) document that timed value strategies outperform untimed ones. These studies use valuation metrics, specifically value spreads, defined as the difference between the valuation ratios of highest and lowest book-to-market sorted portfolios, to evaluate the relative price efficiency of value and growth stocks and to determine the timing of value investing strategies. [Yara et al. \(2020\)](#) construct a strategy that times value in equities using a signal that captures deviations of the past 12 month value spread from its historical average and is observable at time t . They find that the linear timing strategy reliably outperforms an unconditional value strategy and that value returns are thus predictable in real time. These results also persist when looking at alphas relative to the market portfolio and suggest that conditional value strategies are attractive in addition to an indexed market strategy. [Asness et al. \(2021\)](#) use deep value episodes, in which valuation differences of cheap stocks relative to expensive ones are unusually large historically, as a timing signal for their strategies. Specifically, when they observe the value spread exceeding the 80th percentile historically, the strategies takes a long position in the value portfolio and remain invested until spreads decline to the 50th percentile. They show significant alphas for a variety of value assets relative to a long value only strategies. A timed value strategy would therefore entail taking long positions in value assets whenever value assets appear cheaper than growth assets by more than their historical average and taking long positions in the market portfolio otherwise. While recent evidence suggest that value strategies continue to outperform in certain periods, a considerable debate remains, however, about the drivers of the time-variation in the value premium. [Fama and French \(1992\)](#), in the context of the efficient market hypothesis, attribute higher returns to their increased risk, [Lakonishok et al. \(1994\)](#) suggest cognitive biases underlying investor behavior as the drivers of value investing, and [Asness et al. \(2000\)](#) show expected return premiums can vary over time as a consequence of mispricing between value and growth stocks and time variation in the relative degree of mispricing. In recent work, [Gerakos and Linnainmaa \(2018\)](#) and [Golubov and Konstantinidi \(2019\)](#) find mixed evidence on whether value is driven by risk or mispricing. The consistent factor, however, seems to be that the value premium is based on a form of inefficient processing of information. Our measure contributes to this ongoing debate on the source of the value premium by considering an efficiency perspective based on information complexity and helping to identify periods where value stocks have a lower relative efficiency than growth stocks and return predictability is high. We argue that the observed outperformance of value is driven by this variation in relative efficiency over time. Specifically, we suggest a strategy that takes advantage of the predictability of returns to time active long value positions and at the same time also identifies periods where returns are not predictable and passive long positions are preferable.

We first study the cross-section of US stocks, using CRSP data between 1985 and 2020, to assess the efficiency of value and growth stocks. We do this by examining their efficiency relative to the stocks we consider neutral and efficient. In specific, we split stocks into

Figure 4: Inefficiency of value over time

This figure shows relative information complexity (RIC) for the High Minus Neutral (HMN) and Low Minus Neutral (LMN) portfolios, RIC_{HMN} and RIC_{LMN} , over time. To identify long value opportunities, we use RIC as an indicator of inefficiency. US stocks in the cross-section are sorted into decile portfolios according to book-to-market equity (BME) and averaged across their monthly RIC at the end of each month. We then compute RIC_{HMN} and RIC_{LMN} as the difference of the monthly RIC of the value (high BME) and neutral (neutral BME) portfolios, and as the difference of the monthly RIC of the growth (low BME) and neutral (neutral BME) portfolios. Whenever RIC_{HMN} is lower than RIC_{LMN} , value stocks are considered inefficient relative to growth stocks.



decile portfolios based on the book-to-market equity (BME), where high BME, i.e. decile 10, stocks represent value stocks, low BME, i.e. decile 1, stocks represent growth stocks, and middle BME, i.e. decile 5¹, stocks represent neutral stocks, and average across their monthly RIC at the end of each month (detailed overview of variable construction given in appendix B). Given periods of high return predictability for value stocks are associated with informational inefficiency as shown in the literature, we assume value stocks to be relatively inefficient compared to neutral stocks. In an unreported t-test we find that value (and growth) stocks are indeed less efficient compared to neutral stocks, as measured by their mean monthly RIC.

We then compute the difference of the monthly RIC of the value (high BME) and neutral (middle BME) portfolios, i.e. the High Minus Neutral (HMN) portfolio, and the difference of the monthly RIC of the growth (low BME) and neutral (middle BME) portfolios, i.e. the Low Minus Neutral (LMN) portfolio. Figure 4 shows RIC for the HMN and LMN portfolios, RIC_{HMN} and RIC_{LMN} , over time. Whenever RIC_{HMN} is lower than RIC_{LMN} , value stocks are considered inefficient relative to growth stocks.

This applies to the periods in the late 1980s, early 1990s and again from around 1997 to 2004. Thereafter, value appears inefficient again after the global financial crisis and around the beginning of the COVID19 pandemic. Relative inefficiencies of value over time coincide well with known growth-value cycles (see Owyong (2012); Asness et al. (2021); Arnott et al. (2021)). Further, given drawdown periods are associated with high risk and uncertainty, and thus informational inefficiency, we would expect relative inefficiency of value stocks to be highest in those periods. With value being consistently inefficient around recession periods

¹We additionally use both decile 6 and the average of decile 5 and decile 6 portfolios to approximate neutral stocks in our analysis. However, the results do not differ fundamentally.

Table II: Value and buy and hold assets

This table reports month of data availability, size, the annualized expected return, annualized volatility, and Sharpe ratio of value and passive buy and hold assets. The top panel reports ETFs tracking major US value indices, and the bottom panel contains the buy and hold asset.

Value asset	Size	Start	Ann. return	Ann. vola.	SR
iShares Russell 2000 Value ETF	Small-Cap	2000-07	8.55%	19.53%	0.36
Vanguard Small-Cap Value ETF	Small-Cap	2004-01	8.22%	19.41%	0.36
iShares Russell Mid-Cap Value ETF	Mid-Cap	2001-07	8.95%	17.20%	0.45
Vanguard Mid-Cap Value ETF	Mid-Cap	2006-08	7.55%	18.26%	0.36
iShares Russell 1000 Value ETF	Large-Cap	2000-05	5.48%	15.22%	0.26
Vanguard Value ETF	Large-Cap	2004-01	6.34%	14.80%	0.35
iShares Russell Top 200 Value ETF	Mega-Cap	2009-09	11.02%	14.16%	0.74
Vanguard Mega Cap Value ETF	Mega-Cap	2007-12	5.54%	16.04%	0.31
iShares Russell 3000 Value ETF	All-Cap	2000-05	6.35%	15.64%	0.31
iShares MSCI USA Value Factor ETF	All-Cap	2013-04	8.62%	16.56%	0.48
SPDR S&P 500 ETF Trust	Large-Cap	1993-01	8.78%	14.74%	0.44

such as the oil crisis, the dotcom bubble and most recently the global financial crisis, this indeed appears to be the case and is in line with Yara et al. (2020), arguing for high expected value returns in bad times for a number of subsequent years. In the spirit of equation (12), our value indicator, $\mathcal{V}_{i,t}$, therefore takes the form:

$$\mathcal{V}_{i,t} = \mathbf{1}(RIC_{HMN,t} < RIC_{LMN,t}) \quad (19)$$

If at time t value stocks are considered inefficient relative to growth stocks, we take an active long position at $t + 1$ into value assets. Otherwise we take a passive long position into the buy and hold asset tracking the S&P 500 Index.

In order to examine the timing ability of long value positions of our measure, we use the most liquid ETFs tracking a set of established value equity indices, covering small-cap, mid-cap, large-cap, mega-cap and all-cap US value stocks, as measured by their relative book-to-market ratio. The benefit of including assets across different market capitalization ranges is that it allows for the testing of potential size effects. Table II reports the summary statistics for the 10 value assets and the passive buy and hold asset we use. The top panel contains ETFs for major value indices, and the bottom panel contains the passive buy and hold asset tracking the S&P 500 Index.

Table III reports aggregate summary statistics for the overall sample and across size. The fourth and fifth columns report the change in Sharpe ratio for the active timing strategy using relative information complexity (RIC) compared to passive strategies that are long the value asset or long the S&P 500 only. The second last and third last columns report the success rate of the timing strategy, i.e. the fraction of assets with Sharpe ratio improvements, and the number of assets per size bucket, respectively. Figure 5 shows the improvement of the Sharpe ratio when timing long positions into value assets relative to a passive strategy

Figure 5: Timing value using RIC as an indicator of inefficiency

This figure shows the improvement of the Sharpe ratio of a strategy that actively times long value positions using RIC as an indicator of inefficiency compared to a passive long only strategy in the same asset or in the S&P 500 index. The strategy takes a long position whenever we observe that:

$$\mathcal{V}_{i,t} = \mathbf{1}(RIC_{HMN,t} < RIC_{LMN,t}),$$

where $\mathbf{1}$ is the indicator function, $RIC_{HMN,t}$ is the RIC for the High Minus Neutral (HMN) portfolio at time t and $RIC_{LMN,t}$ is the RIC for the Low Minus Neutral (LMN) portfolio at time t . Whenever $RIC_{HMN,t}$ is smaller than $RIC_{LMN,t}$, value stocks are considered inefficient relative to growth stocks. Hence, if at time t value stocks are considered relatively inefficient, we take a long position at $t + 1$ into the value asset. Otherwise we take a long position in the buy and hold asset, the SPDR S&P 500 ETF Trust, tracking the S&P 500 index.

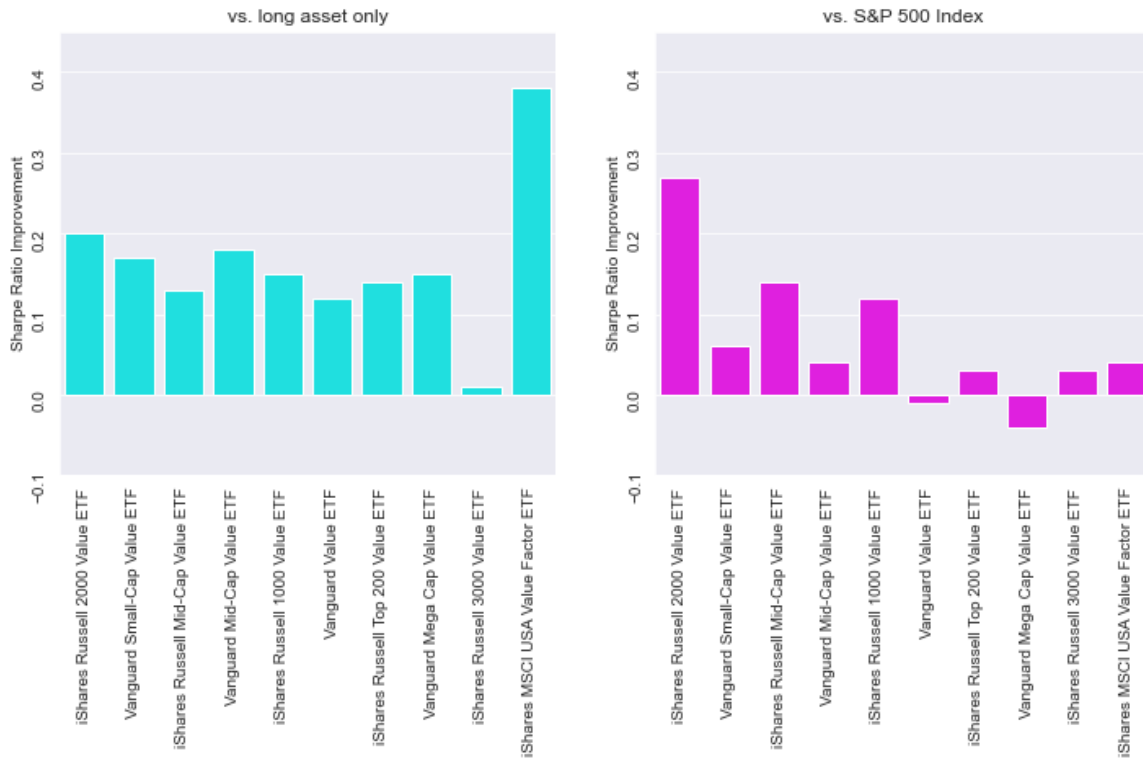


Table III: Timing value with RIC

This table reports aggregate average return and average volatility for the overall sample and across size. The fourth and fifth columns report the improvement of the Sharpe ratio in absolute terms for the active timing strategy using relative information complexity (RIC) compared to passive strategies that are long the asset or long S&P 500 only (i.e. SPDR S&P 500 ETF Trust). The second last and last columns report the success rate of the timing strategy, i.e. the fraction of assets where the Sharpe ratio improves, and the number of assets per size bucket.

Size	Return	Volatility	Δ SR vs. asset	Δ SR vs. S&P 500	Success Rate	Assets
Overall	0.08	0.17	0.16	0.07	80%	10
Small-Cap	0.08	0.19	0.19	0.17	100%	2
Mid-Cap	0.08	0.18	0.16	0.09	100%	2
Large-Cap	0.06	0.15	0.14	0.06	50%	2
Mega-Cap	0.08	0.15	0.15	-0.01	50%	2
All-Cap	0.07	0.16	0.20	0.04	100%	2

that is either long only into the value asset or long into the S&P 500.

We find that the success rate is 80% across the whole sample. The value timing strategy is able to improve the Sharpe ratio compared to both passive strategies for 8 out of 10 assets. For the remaining 2 assets, the timing strategy is at least able to improve the Sharpe ratio compared to a long the value asset only strategy, but fails to outperform the long S&P 500 strategy. The average improvement in the Sharpe ratio is 0.17 and 0.07, respectively, compared to the two passive strategies, which is a substantial improvement. Further, looking at assets across different market capitalization ranges allows us to test whether results are driven by a size effect. The strategy is successful for all small-cap value assets and improves the Sharpe ratio by 0.19 and 0.17 on average relative to passive strategies. Similarly, for mid-cap value assets the success rate is high and the average Sharpe ratio improvement is considerable with 0.16 and 0.09 respectively. Large-cap value assets fair worse in comparison, with a success rate of 50% and with an average improvement of the Sharpe ratio of 0.14 and 0.06. For mega-cap value assets the success rate is also 50%, however, the average Sharpe ratio improvement is negative. Lastly, all-cap value assets, including a diverse mix of stocks across small, mid and large market capitalization, have a success rate at 100%, with the Sharpe ratio improving by 0.20 and 0.04, respectively, relative to the long value asset and long S&P 500 only strategies.

Overall, we find that the conditional, timed value strategy improves Sharpe ratios compared to both an unconditional long only value strategy and a passive buy and hold strategy across market capitalization ranges and for most of the used value assets. The magnitude of the improvement appears to decrease with increasing market capitalization. In particular, the success rates for both large and mega-cap value assets are only at 50%, with mega-caps showing negative Sharpe ratio improvement. With higher market capitalization the investment universe is likely becoming more liquid and more widely covered and thus informationally more efficient, making it harder to find undervalued stocks. Nonetheless, taking into account all-cap assets that look at a broader value universe, results are not

Table IV: Cumulative return of timing strategy

This table reports the cumulative returns for the active timing strategy using RIC, and for both passive strategies that are long the asset or long the S&P 500 only (i.e. long SPDR S&P 500 ETF Trust) for the indicated time period. The first column reports different assets used for the actively timed long position in the strategy.

Active Timing Position	Period	Cumulative return		
		Strategy	Asset	S&P 500
iShares Russell 2000 Value ETF	2000-07 - 2020-12	628.29%	283.50%	161.45%
Vanguard Small-Cap Value ETF	2004-01 - 2020-12	351.94%	189.04%	229.46%
iShares Russell Mid-Cap Value ETF	2001-07 - 2020-12	497.33%	316.01%	257.98%
Vanguard Mid-Cap Value ETF	2006-08 - 2020-12	242.85%	130.54%	186.19%
iShares Russell 1000 Value ETF	2000-05 - 2020-12	276.80%	141.79%	161.79%
Vanguard Value ETF	2004-01 - 2020-12	227.40%	141.54%	229.46%
iShares Russell Top 200 Value ETF	2009-09 - 2020-12	278.78%	204.79%	261.02%
Vanguard Mega Cap Value ETF	2007-12 - 2020-12	136.36%	73.18%	155.71%
iShares Russell 3000 Value ETF	2005-05 - 2020-12	187.90%	185.62%	161.79%
iShares MSCI USA Value Factor ETF	2013-04 - 2020-12	153.08%	73.92%	134.14%

solely driven by a small firm effect. Generally, our timing measure appears to be working well when applied to a variety of value assets. Therefore, relative information complexity appears to be a sound indicator of inefficiency and compellingly times value assets.

To test the robustness of our efficiency measure we examine cumulative returns of our proposed strategies. Table IV reports cumulative returns for the active timing strategy using RIC and for both passive strategies that are either long the value asset or long the S&P 500 only. The first column reports different value assets used for the actively timed long position in the strategy. Overall results support our earlier findings, showing that cumulative returns are higher for the active timing strategy relative to the long asset only strategy for all value assets and higher relative to the long S&P 500 only strategy for 80% of value assets. Across size, the active strategy consistently outperforms passive strategies for small-cap, mid-cap and all-cap value assets. In large- and mega-cap value universes active timing underperforms a passive long S&P 500 approach for 50% of assets. This is consistent with our Sharpe ratio analysis across size, showing that performance is not driven by a specific market capitalization range. In addition, figure A.I in the appendix shows the growth of one dollar invested at the beginning of the respective time period for the active timing strategy and for both passive strategies for each of the included value asset. In line with Leippold and Rüegg (2019), we find that the active timing strategy performs better than passive strategies in some periods and worse during other periods. However, over time and across different assets, active investment appears to generate superior performance.

To further solidify our results, we examine whether our actively timed strategy using relative information complexity produces significant alphas over passive strategies. Table V reports alphas for the timing strategy across different assets. Alphas and t -statistics are obtained by regressing risk-adjusted returns of the timing strategy on the risk-adjusted

Table V: Alphas for the active timing strategy

This table reports the estimated coefficients and t -statistics (in parentheses) obtained by regressing the monthly risk-adjusted returns of the active timing strategy on the risk-adjusted returns of both passive strategies that are long the S&P 500 index (columns 2 to 3) or long the value asset only (columns 4 to 5). The regression model is as follows: $R_{a,t} - R_f = \alpha + \beta(R_{p,t} - R_f) + \epsilon_t$, where $R_{a,t}$ is the monthly return on the actively timed strategy, $R_{p,t}$ is the monthly return on the passive strategy, $R_{f,t}$ is the risk free return, and ϵ_t is the error term. ***, **, and * denote significance at the 1%, 5%, and 10% level. ETFs are sorted according to their date of issuance from oldest to newest.

Actively Timed Asset in Strategy	Period	Alpha	
		vs. S&P 500	vs. asset
iShares Russell 1000 Value ETF	2000-05 - 2020-12	0.0016** (2.389)	0.0019*** (3.590)
iShares Russell 2000 Value ETF	2000-07 - 2020-12	0.0045*** (2.995)	0.0033*** (2.893)
iShares Russell Mid-Cap Value ETF	2001-07 - 2020-12	0.0021** (2.498)	0.0020*** (2.607)
Vanguard Small-Cap Value ETF	2004-01 - 2020-12	0.0012 (1.050)	0.0027** (2.547)
Vanguard Value ETF	2004-01 - 2020-12	-0.0004 (-0.102)	0.0016*** (2.868)
iShares Russell 3000 Value ETF	2005-05 - 2020-12	0.0004 (1.362)	0.0001 (0.321)
Vanguard Mid-Cap Value ETF	2006-08 - 2020-12	0.0008 (0.322)	0.0026*** (2.995)
Vanguard Mega Cap Value ETF	2007-12 - 2020-12	-0.0005 (-0.838)	0.0021*** (2.925)
iShares Russell Top 200 Value ETF	2009-09 - 2020-12	0.0004 (0.922)	0.0019** (2.078)
iShares MSCI USA Value Factor ETF	2013-04 - 2020-12	0.0006 (0.981)	0.0050*** (3.427)

returns of the passive strategies taking a long position into the S&P 500 index or into the value asset, respectively.

Results for the regression on the passive S&P 500 strategy show that alphas of the timed strategy are positive and significant only for the small, mid and large cap Russell value universes, with outperformance decreasing across size. Notably, these time series have the longest tracking record starting in the early 2000s. One issue that arises with all other Value ETFs is their shorter history. Specifically, figure 4 reminds us that during the past ten years the efficiency of value surpasses that of growth stocks whereas before the global financial crisis value stocks appeared to be less efficient. In effect, strategies pertaining to value ETFs originating in the last 10 to 15 years would allocate funds passively to the S&P 500 always or most of the time. As a consequence, these strategy returns essentially map to the return of the asset that tracks the S&P 500 and a positive and statistically significant outperformance is unlikely by construction. On the other hand, strategies that have longer track records seem to be able to exploit the inefficiency of value stocks in earlier periods. Overall, this implies that value timing strategies do not fare worse while offering profitable upside potential at the same time when value stocks become less efficient than growth stocks

over a longer period.

A second performance evaluation considers alphas of value timing strategies relative to simple passive buy and hold strategies of the value assets. For the regression on the passive long asset strategy, results show that alphas are positive and significant in almost every case. Thus, the active timing strategy is able to reliably outperform buy and hold value assets. Overall, an active strategy using RIC as a timing indicator appears to generate alpha over passive strategies for a large number of assets, while its performance is at least equal for other assets, suggesting that active investments are preferable.

Given evidence on the consistency of relative information complexity in timing active long positions into value assets across periods, size, differing assets and relative to other established measures, our proposed efficiency measure appears to be relatively proficient at indicating periods where pursuing active investments into value assets is preferable over passive investments and thus, is a compelling indicator for timing value assets.

An important aspect of an efficiency measure is its validity. To assess the validity of our efficiency measures, we compare the results of active strategies that use relative information complexity and price delay measures as timing indicators. Given higher delay indicates less efficiency, we adapt our value indicator in (19) to take the form:

$$\mathcal{V}_{i,t} = \mathbb{1}(D_{HMN,t} > D_{LMN,t}), \quad (20)$$

Where D is either $D1$, $D2$, or $D3$. If at time t value stocks are considered inefficient relative to growth stocks, i.e. if value stocks have higher delay than growth stocks, we take an active long position at $t + 1$ into value assets. Otherwise we take a passive long position into the buy and hold asset tracking the S&P 500 index.

Table VI reports the improvement of the Sharpe ratio of the strategy that actively times long value positions using price delay measures $D1$, $D2$, and $D3$ proposed by Hou and Moskowitz (2005) compared to the active timing strategy using relative information complexity for the overall sample and across capitalization ranges. The second last and last columns report the success rate of the timing strategy using delay, i.e. the fraction of assets with Sharpe ratio improvements, and the number of assets per size cohort.

Overall, the success rate is 10%. The timing strategy using delay is only able to improve the Sharpe ratio over the timing strategy using relative information complexity for 1 out of 10 assets. Across all three delay measures, the average Sharpe ratio decreases by -0.11. Within the different size cohorts, the Sharpe ratio is consistently lower with the exception of all-cap value assets. For the latter, the timing strategy using delay is on par with the strategy using relative information complexity. However, this only applies to one delay measure for one asset, thus effectively reducing the success rate from 50% to 16%. These results suggest that relative information complexity is able to time value assets while delay measures are not.

Table VI: Timing value with Delay

This table reports the improvement of the Sharpe ratio in absolute terms for the active timing strategy using the price delay measures $D1$, $D2$ and $D3$ proposed by Hou and Moskowitz (2005) relative to the timing strategy using RIC for the overall sample and across size. The second-last and last columns report the fraction of the assets that display a better Sharpe Ratio for the timing strategy using delay and the number of assets used per size cohort.

Size	Sharpe ratio RIC	Δ Sharpe ratio delay vs RIC			Success rate	Assets
		$D1$	$D2$	$D3$		
Overall	0.56	-0.10	-0.08	-0.14	10%	10
Small-Cap	0.55	-0.12	-0.10	-0.24	0%	2
Mid-Cap	0.56	-0.08	-0.07	-0.18	0%	2
Large-Cap	0.44	-0.08	-0.08	-0.09	0%	2
Mega-Cap	0.67	-0.08	-0.08	-0.03	0%	2
All-Cap	0.59	-0.14	-0.09	-0.15	50%	2

In an additional performance evaluation, we examine the alphas of value timing strategies using price delay measures relative to passive buy-and-hold strategies for both the S&P 500 and value assets. Table VII reports alphas for the active timing strategy across different value assets. Alphas and t -statistics are obtained by regressing risk-adjusted returns of the timing strategy using price delay measures $D1$, $D2$, $D3$ on risk-adjusted returns of passive strategies.

Results for the regressions on both passive strategies across all three delay measures show that alphas of the timed strategy are positive and significant only in a few cases. In particular, not only is the active timing strategy not able to significantly outperform the asset tracking the S&P 500 at any point, but, given negative and significant alphas across all measures in a few instances, appears to underperform the latter at times. Further, alphas of the timing strategy relative to the passive buy-and-hold strategy of the value asset are positive and significant in less than 50% of cases and decrease across the different delay measures. Thus, the timing strategy cannot reliably outperform the passive long value asset strategy. Overall, an active strategy that uses delay as a timing indicator does not appear to generate consistent alpha over passive strategies for most assets, while its performance is worse for a few assets. These results provide further evidence that our proposed measure based on relative information complexity is indeed a superior approach to time value assets compared to efficiency measures based on delay.

C Timing short selling

In this chapter we use our measure of absolute information complexity to identify short-selling opportunities through time-series momentum. Specifically, we assume that during certain periods where predictability of asset returns is high, past returns are associated with future returns such that negative returns are entailed with future negative returns. Therefore, investors either passively hold a long position in the asset or actively time a short position when the asset is considered inefficient and predictable with respect to negative returns. Several studies document the predictability of future asset returns with past obser-

Table VII: Alphas for the active timing strategy using delay measures

This table reports the estimated coefficients and t -statistics (in parentheses) obtained by regressing the monthly risk-adjusted returns of the active timing strategy using price delay measures $D1$, $D2$, $D3$ on the risk-adjusted returns of both passive strategies that are long the S&P 500 index or long the value asset only. The regression model is as follows: $R_{a,t} - R_f = \alpha + \beta(R_{p,t} - R_f) + \epsilon_t$, where $R_{a,t}$ is the monthly return on the actively timed strategy, $R_{p,t}$ is the monthly return on the passive strategy, $R_{f,t}$ is the risk free return, and ϵ_t is the error term. ***, **, and * denote significance at the 1%, 5%, and 10% level. ETFs are sorted according to their date of issuance from oldest to newest (see Table V for period reference).

Actively Timed Asset in Strategy	Alpha $D1$		Alpha $D2$		Alpha $D3$	
	S&P 500	asset	S&P 500	asset	S&P 500	asset
iShares Russell 1000 Value ETF	0.0006 (0.745)	0.0008** (2.441)	0.0003 (0.605)	0.0007 (1.019)	0.0003 (0.395)	0.0005 (0.977)
iShares Russell 2000 Value ETF	0.0028 (1.497)	0.0010 (1.642)	0.0021 (1.538)	0.0011 (0.834)	0.0007 (0.466)	-0.0006 (-0.506)
iShares Russell Mid-Cap Value ETF	0.0014 (1.354)	0.0011** (2.106)	0.0011 (1.572)	0.0011 (1.243)	-0.0005 (-0.462)	-0.0008 (-1.234)
Vanguard Small-Cap Value ETF	-0.0002 (-0.155)	0.0011 (1.609)	0.0008 (0.711)	0.0023** (2.221)	-0.0020 (-1.637)	-0.0004 (-0.439)
Vanguard Value ETF	-0.0010 (-1.478)	0.0005** (2.162)	-0.0007 (-1.531)	0.0009* (1.700)	-0.0008 (-1.550)	0.0008* (1.692)
iShares Russell 3000 Value ETF	0.0004 (1.284)	0.0001 (1.289)	0.0003 (1.159)	-0.0000 (-0.145)	-0.0001 (-0.455)	-0.0004** (-2.369)
Vanguard Mid-Cap Value ETF	-0.0004 (-0.394)	0.0013** (2.170)	-0.0001 (-0.145)	0.0017* (1.917)	-0.0016 (-1.600)	0.0003 (0.385)
Vanguard Mega Cap Value ETF	-0.0018** (-2.074)	0.0005** (2.059)	-0.0016*** (-2.696)	0.0009 (1.324)	-0.0010 (-1.451)	0.0015** (2.474)
iShares Russell Top 200 Value ETF	0.0001 (0.053)	0.0010*** (2.683)	-0.0003 (-0.392)	0.0010 (1.354)	0.0004 (0.521)	0.0016** (2.516)
iShares MSCI USA Value Factor ETF	-0.0025 (-1.602)	0.0014 (1.557)	-0.0011 (-0.900)	0.0031** (2.469)	-0.0024* (-1.697)	0.0018 (1.639)

vations. [Jegadeesh and Titman \(2002\)](#) document the existence of cross-sectional momentum and [Moskowitz et al. \(2012\)](#) demonstrate the presence of aggregate momentum in the time-series dimension. This strand of academic literature uses past returns as a predictor for future returns. In the cross-sectional momentum setting of [Jegadeesh and Titman \(2002\)](#), assets that perform poorly relative to the cross-section continue to perform poorly while those assets that record the highest returns will continue to yield high returns. A cross-sectional momentum strategy therefore implies long positions in past winners and short positions in past losers. In the time-series momentum case, [Moskowitz et al. \(2012\)](#) use a simple sign indicator for future returns. Specifically, if the cumulative return over a certain period is positive, future returns are expected to be positive as well and the strategy takes a long position. If the sign is negative the strategy implies a short position in the asset. Both cross-sectional and time-series momentum do not consider the fact that predictability varies over time and across assets. Our measure expands into this dimension and helps identifying periods where a time series is predictable such that past returns are indeed predictive of future returns. Specifically, we provide a strategy that exploits the predictability of negative returns in order to time short positions. At the same time this also identifies periods where returns are not predictable and passive long positions are recommended, either in the asset

itself or in some other suitable asset.

In order to identify short-selling opportunities we use absolute information complexity, $AIC_{i,t}$, as an indicator of inefficiency. Whenever the indicator function in equation (7) takes the value one, the asset is considered inefficient. One peculiarity of this exercise is that we are exclusively interested in timing short positions based on past negative returns. To that end, we refine our measure of inefficiency in that we identify whether the inefficiency originates from redundant sequences of high returns or low returns. If an asset is categorized as inefficient but its inefficiency primarily arises from high returns, then the portion of past low returns need not be predictable of future returns. Only if past low returns cause the asset to appear inefficient investors should open short positions. We construct this additional signal by examining each longest match ℓ^i in our entropy estimate in equation (2). For each longest match corresponding to position i we compute the average return quintile in that longest match. If the longest match mostly contains high returns the average will be high as well and vice versa. In the spirit of equation (2) we then weight each average quintile with the length of the corresponding longest match to obtain an overall measure to pin down where an asset's inefficiency is originating from. From section II.C we know that the optimal number of return quintiles is 4. Therefore, our measure resides between 1 and 4. A magnitude close to one suggests that the redundancy is in low returns while a number close to 4 implies that redundant patterns primarily contain high returns. We consider a magnitude between 1 and 2 as indicative of redundancy in low returns. Therefore, our short-selling indicator, $\mathcal{S}_{i,t}$, takes the form:

$$\mathcal{S}_{i,t} = \mathbb{1} (\mathbb{1} (r_{i,t-1} < 0) + \mathbb{1} (AIC_{i,t-1} < 1) + \mathbb{1} (\bar{Q}_{i,t-1} \leq 2) = 3), \quad (21)$$

where \bar{Q} denotes the weighted average return quintile to identify whether low or high returns foster redundant patterns. Put differently, if at time t the asset return r_i is negative and the asset is considered inefficient and this inefficiency is likely attributable to low returns, we take a short position at time $t + 1$. In all other scenarios we simply buy into the asset and hold.

As our approach is closely related to Moskowitz et al. (2012)'s time-series momentum, we use the same asset classes to examine timing ability of short positions of our measure. Specifically, we use the closest-to-maturity futures available for various equity indices, government bonds, commodities and currency pairs. One benefit of including several asset classes is that it facilitates testing our measure under several market mechanisms. Each asset class's market is governed by different mechanisms, frictions, and conventions. Table VIII reports the summary statistics for the 42 future contracts we use. The top panel and second panel contain futures for major equity indices and government bonds from major financial markets. The third panel embodies commodities and the bottom panel reports major currency pairs for the U.S. dollar.

Table IX reports aggregate summary statistics for the overall sample as well as for equity indices, government bonds, commodities, and currency pairs individually. The third and

Table VIII: Futures contracts

This table reports month of data availability, the annualized expected return and annualized volatility, as well as the average entropy of the future contract's time series. Entropy is estimated using a reversed time series using the past 50 returns and 4 return bins. The top panel contains future contracts for major equity indices, the second panel contains government bonds, the third panel embodies commodities, and the bottom panel reports currency pairs. For every asset, at any given time we use the time series of future contract with the closest maturity.

Future Contract	Start	Annualized return	Annualized Volatility	Sharpe Ratio	Mean entropy
AEX	1988-10	7.03%	18.33%	0.23	1.32
CAC 40	1995-04	5.67%	18.61%	0.19	1.33
DAX	1990-11	9.48%	20.25%	0.35	1.33
FTSE 100	1985-01	5.83%	15.58%	0.17	1.33
IBEX 35	1992-04	6.03%	21.25%	0.18	1.34
MIB	2004-03	1.34%	21.02%	0.01	1.33
Nikkei 225	1988-09	1.33%	19.06%	-0.07	1.32
S&P 500	1997-09	7.21%	15.63%	0.34	1.3
SMI	1992-09	7.57%	15.38%	0.34	1.33
SPI 200	2000-05	4.54%	13.60%	0.22	1.34
AUS 3Y	1992-01	0.28%	1.24%	0.22	1.31
AUS 10Y	1992-01	0.30%	1.06%	0.29	1.34
CAN 10Y	1989-09	1.61%	6.82%	0.24	1.3
EURO 10Y	1990-11	2.67%	5.32%	0.50	1.3
JPN 10Y	1987-09	1.40%	4.65%	0.30	1.26
US 2Y	1990-06	0.37%	1.72%	0.21	1.21
US 5Y	1988-05	0.89%	4.20%	0.21	1.29
US 10Y	1985-01	1.75%	6.67%	0.26	1.31
US 30Y	1985-01	3.04%	10.60%	0.29	1.33
Cattle	1985-01	2.94%	14.70%	-0.01	1.23
Cocoa	1985-01	5.11%	30.11%	0.07	1.33
Coffee	1985-01	6.26%	37.16%	0.08	1.30
Corn	1985-01	5.54%	27.98%	0.08	1.27
Cotton	1985-01	5.22%	29.86%	0.07	1.29
Crude Oil	1988-06	9.74%	34.32%	0.20	1.32
Gas Oil	1985-01	7.70%	34.25%	0.13	1.30
Gold	1985-01	6.21%	15.27%	0.20	1.28
Lean Hogs	1985-01	7.58%	37.03%	0.12	1.04
Platinum	1985-01	6.68%	23.29%	0.16	1.26
Silver	1985-01	7.96%	28.12%	0.17	1.27
Soybeans	1985-01	5.21%	23.81%	0.09	1.30
Sugar	1985-01	12.23%	41.11%	0.22	1.27
Wheat	1985-01	5.75%	28.73%	0.09	1.31
AUD/USD	1987-01	0.11%	11.45%	0.01	1.33
CAD/USD	1985-01	0.36%	7.69%	0.05	1.32
CHF/USD	1985-01	-2.33%	11.34%	-0.21	1.33
EUR/USD	1998-05	0.07%	9.64%	0.01	1.36
GBP/USD	1985-01	0.04%	10.12%	0.00	1.32
JPY/USD	1985-01	-1.87%	10.89%	-0.17	1.29
NOK/USD	2002-05	0.71%	12.36%	0.06	1.32
NZD/USD	1997-05	0.61%	12.88%	0.05	1.33
SEK/USD	2002-05	-0.42%	11.88%	-0.04	1.33

Table IX: Timing short positions

This table reports aggregate average return, average volatility, and average entropy for the overall sample as well as for the individual asset classes. The third and second last columns report the average improvement of the Sharpe ratio in absolute terms and the fraction of assets that display a better Sharpe Ratio using the timing strategy.

Asset Class	Return	Volatility	Entropy	Δ Sharpe Ratio	Success rate	Assets
Overall	0.04	0.21	1.30	0.01	62%	42
Equity Indices	0.06	0.18	1.33	0.08	90%	10
Government Bonds	0.01	0.06	1.29	0.03	44%	9
Commodities	0.07	0.30	1.27	-0.06	43%	14
Currency Pairs	0.00	0.11	1.32	0.04	78%	9

second last columns report results associated with our timing strategy. Specifically, the third last column reports the change in Sharpe ratio for the active timing strategy relative to the passive buy and hold approach. The second last column reports the success rate of the timing strategy, i.e. the fraction of assets where the Sharpe ratio improves. Figure 6 shows the improvement of the Sharpe ratio when timing short positions compared to a passive strategy that is long the asset only.

The overall success rate is 62%. Our strategy is able to improve the Sharpe ratio for 27 out of 42 assets. However, the average improvement of the Sharpe ratio is 0.01. Hence, while across all assets there is a success rate of more than 60% the average improvement is virtually zero. However, when looking at individual asset classes the results are striking. For 9 out of 10 major equity indices the Sharpe ratio improves with our timing indicator. On average, the absolute improvement of the Sharpe ratio amounts to 0.08. When relating this number to the Sharpe ratio of indices this number is considerable. The average Sharpe ratio of the 10 futures contracts on the major equity indices is 0.2. Therefore, the timing strategy improves the Sharpe ratio by almost 40% on average.

The timing measure also does fairly well with respect to currency pairs. The success rate amounts to 78% and the Sharpe ratio improves by 0.04 on average. Due to the depreciation of the U.S. dollar expected returns of passive long strategies yield returns close to zero or negative returns. The average Sharpe ratio is -0.02. The timing strategy renders the average Sharpe ratio slightly positive and therefore prevents losses when passively holding a depreciating currency. The average assessment is strongly driven by the time series of the New Zealand dollar against the U.S. dollar. When excluding this time series, the average improvement of the Sharpe ratio more than doubles to 0.09 For Government bonds and commodities the success rates are below 50% and the average improvement of the Sharpe ratio is marginal or negative.


Overall, the timing measure does a compelling job when applied to futures contracts of equity indices and has some benefit when trading currency pairs. It fails to deliver economic benefit for commodities and Government bonds. Hence, it appears that the measure works well for some asset classes while it does not for others. Two observations stand out when

Figure 6: Timing short positions using information complexity

This figure shows the improvement of the Sharpe ratio of a strategy that actively times short positions using information complexity compared to a passive long only strategy in the same asset. The strategy takes a short position whenever we observe that :

$$\mathcal{S}_{i,t} = \mathbb{1}(\mathbb{1}(r_{i,t-1} < 0) + \mathbb{1}(AIC_{i,t-1} < 1) + \mathbb{1}(\bar{Q}_{i,t-1} \leq 2) = 3),$$

where $\mathbb{1}$ is the indicator function, r_i denotes the return of asset i , AIC_i is the absolute information complexity, and \bar{Q}_i measures whether redundancy, if any, is in high or low returns. Hence, if at time t the asset return r_i is negative and the asset is considered inefficient and this inefficiency is likely attributable to low returns, the strategy takes a short position at time $t + 1$. In all other scenarios the strategy simply buys into the asset and holds. Entropy is estimated using a reversed time series using the past 50 returns and 4 return bins with the estimator in (2). An inefficiency in low returns is identified whenever \bar{Q}_i is equal to or lower than 2.



analyzing the different asset classes in light of the timing strategy. First, both equity indices and currency pairs exhibit a higher entropy in contrast to government bonds and commodities. The average entropy estimate for equity indices and currency pairs is 1.33 for both asset classes and therefore both are above the one standard deviation border of an entropy process of White Noise in the spirit of equation (9). This means that the futures pertaining to these asset classes are likely efficient on average. Turning to government bonds and commodities, the average entropy estimates are 1.29 and 1.27, respectively. Both are therefore outside of the one standard deviation interval of a white noise entropy process and hence, are inefficient on average. Hence it appears that timing and exploiting return predictability works better for those asset classes with higher efficiency. One explanation may be that an asset must be inefficient enough for profitable opportunities to arise but must be efficient enough such that active investment can exploit the inefficiency. Put differently, only if assets are efficient on average, transient inefficiencies can be exploited thorough active investment. Under this notion, especially commodities might be too inefficient, such that past return patterns cannot be exploited. Specifically, figure 6 shows the improvement of the Sharpe ratio for all 42 futures contracts. For the commodity future for lean hogs the timing strategy deteriorates the Sharpe ratio by roughly 0.5. At the same time, table VIII reminds us that lean hogs have the lowest entropy with an average estimate of 1.04. They are a long way clear from the efficiency of any other asset. The time series of lean hogs future contracts is therefore the least efficient. In fact, it may be too inefficient for exploiting past return patterns as market mechanisms through which active trading strategies exploit opportunities fail.

Overall, this section shows that absolute information complexity can be used to time short positions using momentum in time series in the spirit of Moskowitz et al. (2012) and improve Sharpe ratios compare to passive buy and hold strategies. Specifically, absolute information complexity times the persistence of past return patterns into the future and signals when short positions are likely to be profitable. The timing indicator works well for asset classes that are efficient on average such that market mechanisms still work properly most of the time and facilitate active investment to be profitable. The Sharpe ratios improve in almost all cases. For asset classes that are inefficient on average, timing positions is to no avail as market mechanisms may no longer work properly. This suggests that for market timing to be profitable, a minimal degree of efficiency might be required.

IV Conclusion

We construct efficiency measures at the asset level using redundancy in return patterns. Our measures are directly linked to the weak-form efficiency of the Efficient Markets Hypothesis as they measure the degree of predictability. We first use these measures to examine the aggregate efficiency in the U.S. stock market between 1985 and 2020. We find that aggregate market efficiency is cyclical and varies over time. Ever since the Global Financial Crisis, aggregate market efficiency tends do decline.

At the security level, we document substantial variation in efficiency both across assets and time. The main duty of our measures is to separate efficient from inefficient periods at the asset level and indicate periods where a financial asset or a set of financial assets is inefficient enough such that active investment becomes profitable over passive investments. We provide market timing indicators for strategies related to value and momentum, two of the most established phenomena in financial market research. We demonstrate that our measures identify periods where value stocks are inefficient and active asset allocation between value stocks and the S&P 500 is more profitable compared to a passive investment in the S&P 500 or the underlying value stocks only.

Similarly, we provide timing signals for short positions. Specifically, we identify periods of high return predictability where time series momentum in negative returns is persistent such that active asset allocation between long and short positions in equity and currency futures increases the Sharpe ratio relative to a passive long only strategy in the equity and currency futures.

Overall, our efficiency measures appear to be relatively well suited at indicating periods when pursuing active investments is preferable to passive investments for value and momentum. A potential avenue for further research would be the examination of the measures in the context of mutual fund performance and market timing of mutual funds.

References

- Alvarez-Ramirez, J., Rodriguez, E., Alvarez, J., 2012. A multiscale entropy approach for market efficiency. *International Review of Financial Analysis* 21, 64–69.
- Andersen, T. G., Bollerslev, T., Das, A., 2001. Variance-ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns. *Journal of Finance* 56, 305–327.
- Arnott, R. D., Harvey, C. R., Kalesnik, V., Linnainmaa, J. T., 2021. Reports of value’s death may be greatly exaggerated. *Financial Analysts Journal* 77, 44–67.
- Asness, C., Liew, J., Pedersen, L. H., Thapar, A., 2021. Deep value. *Journal of Portfolio Management* 47, 11–40.
- Asness, C. S., Friedman, J. A., Krail, R. J., Liew, J. M., 2000. Style timing: Value versus growth. *Journal of Portfolio Management* 26, 50–60.
- Asness, C. S., Moskowitz, T. J., Pedersen, L. H., 2013. Value and momentum everywhere. *Journal of Finance* 68, 929–985.
- Bentes, S. R., Menezes, R., Mendes, D. A., 2008. Long memory and volatility clustering: Is the empirical evidence consistent across stock markets? *Physica A: Statistical Mechanics and its Applications* 387, 3826–3830.
- Boehmer, E., Kelley, E. K., 2009. Institutional investors and the informational efficiency of prices. *Review of Financial Studies* 22, 3563–3594.
- Busch, P., Obernberger, S., 2017. Actual share repurchases, price efficiency, and the information content of stock prices. *Review of Financial Studies* 30, 324–362.
- Carhart, M. M., 1997. On persistence in mutual fund performance. *The Journal of Finance* 52, 57–82.
- Chan, K., Fong, W. M., 2000. Trade size, order imbalance, and the volatility-volume relation. *Journal of Financial Economics* 57, 247–273.
- Charles, A., Darné, O., 2009. Variance-ratio tests of random walk: An overview.
- Chordia, T., Roll, R., Subrahmanyam, A., 2005. Evidence on the speed of convergence to market efficiency. *Journal of Financial Economics* 76, 271–292.
- Copeland, M. M., Copeland, T. E., 1999. Market timing: Style and size rotation using the VIX. *Financial Analysts Journal* 55, 73–81.
- Darbellay, G. A., Wuertz, D., 2000. Entropy as a tool for analyzing statistical dependences in financial time series. *Physica A: Statistical Mechanics and its Applications* 287, 429–439.
- Davis, J. L., Fama, E. F., French, K. R., 2000. Characteristics, covariances, and average returns: 1929 to 1997. *The Journal of Finance* 55, 389–406.
- Fama, E. F., 1970. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance* 25, 383.
- Fama, E. F., French, K. R., 1992. The cross-section of expected stock returns. *Journal of Finance* 47, 427–465.

- Fama, E. F., French, K. R., 2012. Size, value, and momentum in international stock returns. *Journal of Financial Economics* 105, 457–472.
- Gao, Y., Kontoyiannis, I., Bienenstock, E., 2008. Estimating the entropy of binary time series: Methodology, some theory and a simulation study. *Entropy* 10, 71–99.
- Gerakos, J., Linnainmaa, J. T., 2018. Decomposing value. *The Review of Financial Studies* 31, 1825–1854.
- Golubov, A., Konstantinidi, T., 2019. Where is the risk in value? Evidence from a market-to-book decomposition. *Journal of Finance* 74, 3135–3186.
- Gulko, L., 1999. The entropic market hypothesis. *International Journal of Theoretical and Applied Finance* 02, 293–329.
- Hasbrouck, J., 1993. Assessing the quality of a security market: A new approach to transaction-cost measurement. *Review of Financial Studies* 6, 191–212.
- Hasbrouck, J., Ho, T. S., 1987. Order arrival, quote behavior, and the return-generating process. *The Journal of Finance* 42, 1035–1048.
- Henriksson, R. D., Merton, R. C., 1981. On market timing and investment performance. II. Statistical procedures for evaluating forecasting skills. *The Journal of Business* 54, 513.
- Heston, S. L., Korajczyk, R. A., Sadka, R., 2010. Intraday patterns in the cross-section of stock returns. *Journal of Finance* 65, 1369–1407.
- Hou, K., Moskowitz, T. J., 2005. Market frictions, price delay, and the cross-section of expected returns. *Review of Financial Studies* 18, 981–1020.
- Hou, K., Xue, C., Zhang, L., 2020. Replicating anomalies. *Review of Financial Studies* 33, 2019–2133.
- Ito, M., Noda, A., Wada, T., 2014. International stock market efficiency: a non-Bayesian time-varying model approach. *Applied Economics* 46, 2744–2754.
- Ito, M., Noda, A., Wada, T., 2016. The evolution of stock market efficiency in the US: a non-Bayesian time-varying model approach. *Applied Economics* 48, 621–635.
- Jegadeesh, N., Titman, S., 2002. Cross-sectional and time-Series determinants of momentum returns. *Review of Financial Studies* 15, 143–157.
- Kontoyiannis, I., Algoet, P. H., Suhov, Y. M., Wyner, A. J., 1998. Nonparametric entropy estimation for stationary processes and random fields, with applications to english text. *IEEE Transactions on Information Theory* 44, 1319–1327.
- Kristoufek, L., Vosvrda, M., 2014. Measuring capital market efficiency: Long-term memory, fractal dimension and approximate entropy. *European Physical Journal B* 87.
- Lakonishok, J., Shleifer, A., Vishny, R. W., 1994. Contrarian investment, extrapolation, and risk. *The Journal of Finance* 49, 1541–1578.
- Leippold, M., Rüegg, R., 2019. Fifty shades of active and index alpha. *SSRN Electronic Journal* .
- LeRoy, S. F., 1973. Risk aversion and the martingale property of stock prices. *International Economic Review* 14, 436.

- Lo, A. W., 1991. Long-term memory in stock market prices. *Econometrica* 59, 1279.
- Lo, A. W., 2004. The adaptive markets hypothesis. *Journal of Portfolio Management* 30, 15–29.
- Lo, A. W., MacKinlay, A. C., 1988. Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies* 1, 41–66.
- Lo, A. W., MacKinlay, A. C., 1989. The size and power of the variance ratio test in finite samples. A Monte Carlo investigation. *Journal of Econometrics* 40, 203–238.
- Lopez De Prado, M., 2018. *Advances in financial machine learning*. Wiley & Sons, Inc., Hoboken, New Jersey, first ed.
- Lucas, R. E., 1978. Asset prices in an exchange economy. *Econometrica* 46, 1429.
- Maasoumi, E., Racine, J., 2002. Entropy and predictability of stock market returns. *Journal of Econometrics* 107, 291–312.
- Mascio, D. A., Fabozzi, F. J., Zumwalt, J. K., 2020. Market timing using combined forecasts and machine learning. *Journal of Forecasting* .
- Merton, R. C., 1981. On market timing and investment performance. I. An equilibrium theory of value for market forecasts. *The Journal of Business* 54, 363.
- Moskowitz, T. J., Ooi, Y. H., Pedersen, L. H., 2012. Time series momentum. *Journal of Financial Economics* 104, 228–250.
- Oh, G., Kim, S., Eom, C., 2007. Market efficiency in foreign exchange markets. *Physica A: Statistical Mechanics and its Applications* 382, 209–212.
- Ornstein, D. S., Weiss, B., 1993. Entropy and data compression schemes. *IEEE Transactions on Information Theory* 39, 78–83.
- Owyong, D. T., 2012. Value stocks and the macro cycle. *Journal of Investment Consulting* 13, 25–29.
- Pascoal, R., Monteiro, A. M., 2014. Market efficiency, roughness and long memory in PSI20 index returns: Wavelet and entropy analysis. *Entropy* 16, 2768–2788.
- Pele, D. T., Lazar, E., Dufour, A., 2017. Information entropy and measures of market risk. *Entropy* 19.
- Pincus, S., Kalman, R. E., 2004. Irregularity, volatility, risk, and financial market time series. Tech. Rep. 38.
- Risso, W. A., 2008. The informational efficiency and the financial crashes. *Research in International Business and Finance* 22, 396–408.
- Shannon, C. E., 1948. A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423.
- Stattman, D., 1980. Book values and expected stock returns, Unpublished M.B.A. honors paper, University of Chicago.
- Stivers, C., Sun, L., 2010. Cross-sectional return dispersion and time variation in value and momentum premiums. *Journal of Financial and Quantitative Analysis* 45, 987–1014.

- Tang, Y., Whitelaw, R. F., 2011. Timevarying sharpe ratios and market timing. *Quarterly Journal of Finance* 1, 465–493.
- Teynor, J., Mazuy, K., 1966. Can mutual funds outguess the market?
- Yara, F. B., Boons, M., Tamoni, A., 2020. Value return predictability across asset classes and commonalities in risk premia. Tech. rep.
- Ziv, J., Lempel, A., 1978. Compression of individual sequences via variable-rate coding. *IEEE Transactions on Information Theory* 24, 530–536.
- Zunino, L., Zanin, M., Tabak, B. M., Pérez, D. G., Rosso, O. A., 2009. Forbidden patterns, permutation entropy and stock market inefficiency. *Physica A: Statistical Mechanics and its Applications* 388, 2854–2864.

Appendix Essay 3

A Entropy Estimator

To illustrate the enhanced Kontoyiannis Entropy estimator of [Lopez De Prado \(2018\)](#) with an expanding window in a more tangible fashion, consider a time series of returns with length n , denoted by $\{R^n\}$. For simplicity, a return takes the value 1 if it is positive and 0 elsewhere. Hence, consider the following return sequence of length $n = 10$:

$$\{R^{10}\} = \{R_1, R_2, R_3, \dots, R_{10}\} = \{1, 1, 0, 0, 1, 1, 1, 0, 0, 1\}$$

As we have an expanding window and we compute the length ℓ of the longest return pattern $r_i^{i+\ell}$ in the return sequence $\{R^n\}$ starting at position i that also appears in the window r_1^{i-1} preceding position i , our maximum window length is bounded by $\lfloor n/2 \rfloor$. Furthermore, to guarantee that both matching return sequences are of the same length we have that $\max\{\ell^i\} \leq w$. Hence, in the above case we have that the lengths of our expanding window are $w \in \{1, 2, \dots, 5\}$ and the corresponding positions i that we are interested at are $i \in \{2, 3, \dots, 6\}$. This yields the following sequences over which we search for the the longest match ℓ and average:

$$w = 1, i = 2 : \{R^{n=2}\} = \{R_1 \mid R_2\} = \{1 \mid 1\}$$

$$w = 2, i = 3 : \{R^{n=4}\} = \{R_1, R_2 \mid R_3, R_4\} = \{1, 1 \mid 0, 0\}$$

$$w = 3, i = 4 : \{R^{n=6}\} = \{R_1, R_2, R_3 \mid R_4, R_5, R_6\} = \{1, 1, 0 \mid 0, 1, 1\}$$

$$w = 4, i = 5 : \{R^{n=8}\} = \{R_1, R_2, R_3, R_4 \mid R_5, R_6, R_7, R_8\} = \{1, 1, 0, 0 \mid 1, 1, 1, 0\}$$

$$w = 5, i = 6 : \{R^{n=10}\} = \{R_1, R_2, R_3, R_4, R_5 \mid R_6, R_7, R_8, R_9, R_{10}\} = \{1, 1, 0, 0, 1 \mid 1, 1, 0, 0, 1\}$$

For $i = 2$, finding $\max\{0 \leq \ell \leq w : x_i^{i+\ell} = x_j^{j+\ell} \text{ for some } 1 \leq j \leq i - 1\}$ is trivial as we can only check whether the return string at $\{R_2\} = \{R_1\}$. As this is the case, $\ell^{i=2} = 1$. For $i = 3$ we search the longest sequence in $\{R_3, R_4\}$ which also appears in $\{R_1, R_2\}$. Specifically, we check whether we can find $\{0\}$ or $\{0, 0\}$ in $\{1, 1\}$. This is not the case such that for $i = 3$ then longest match, $\ell^{i=3}$, is zero. For $i = 4$ we search whether we can find $\{0\}$, $\{0, 1\}$, $\{0, 1, 1\}$ in $\{1, 1, 0\}$. Clearly, $\{0\}$ is then longest match such that $\ell^{i=4} = 1$. For $i = 5$ we search whether we can find $\{1\}$, $\{1, 1\}$, $\{1, 1, 1\}$ or $\{1, 1, 1, 0\}$ in $\{1, 1, 0, 0\}$. Here, the longest match is $\ell^{i=5} = 2$ as $\{1, 1\}$ is the longest match in $\{1, 1, 0, 0\}$. Finally, for $i = 6$ we search whether we can find $\{1\}$, $\{1, 1\}$, $\{1, 1, 0\}$, $\{1, 1, 0, 0\}$ or $\{1, 1, 0, 0, 1\}$ in $\{1, 1, 0, 0, 1\}$. For this window, the longest match is $\ell^{i=6} = 5$. The entropy estimate is then an average over a nonlinear function of $L_i = 1 + \max\{0 \leq \ell \leq w : x_i^{i+\ell} = x_j^{j+\ell}, \forall i \in \{2, 3, \dots, 6\}\}$.

We reverse the return series to obtain a more transient estimate of efficiency. Note that the expanding window approach starts at the first position of the return string and then expands through the observations such that the next substring always contains the current one. This implicitly gives more weight to observations at the beginning of the string. By

reversing the order of returns we can assign more weight to the most recent returns and less weight to more distant ones ultimately providing a better estimate of the current efficiency. This is possible because the entropy of $\{1, 1, 1, 0, 1, 1\}$ is the same as for $\{1, 1, 0, 1, 1, 1\}$.

B Variable construction


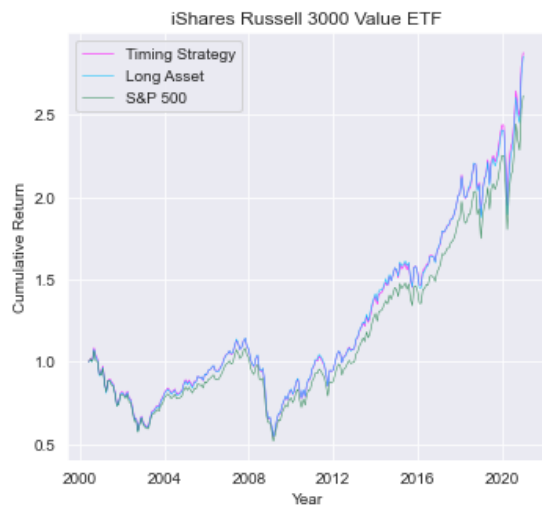


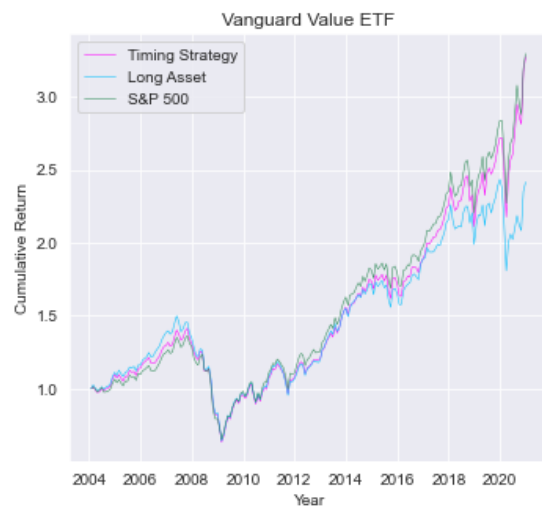

Following [Hou et al. \(2020\)](#), at the end of June every year t , we split stocks into decile portfolios based on the book-to-market equity, which corresponds to the book equity for the fiscal year ending in calendar year $t - 1$ divided by the market equity at the end of December of $t - 1$. Following [Davis et al. \(2000\)](#), book equity is measured as stockholders' book equity, plus balance sheet deferred taxes and investment tax credit (Compustat annual item TXDITC) if available, minus the book value of preferred stock. Stockholder's equity corresponds to the value reported by Compustat (item SEQ), if available. If not, stockholders' equity is measured as the book value of common equity (item CEQ) plus the par value of preferred stock (item PSTK, or depending on availability the redemption value, PSTKRV, or liquidating value, PSTKL), or as the book value of assets (item AT) minus total liabilities (item LT).

C Additional Figures

Figure A.I: Cumulative Return of Active Timing Strategy

This figure shows the cumulative returns for the active timing strategy using RIC, and for both passive strategies that are long the asset or long the S&P 500 only (i.e. long SPDR S&P 500ETF Trust).





Statement of Authorship / Selbständigkeitserklärung

I hereby declare that I have written this thesis independently and have not used any sources other than those indicated. I have marked all co-authorships as well as all passages taken verbatim or in spirit from sources as such. I am aware that otherwise the Senate is entitled to withdraw the title awarded on the basis of this thesis in accordance with Article 36 paragraph 1 letter o of the Law of September 5, 1996 on the University.

Ich erkläre hiermit, dass ich diese Dissertation selbständig verfasst und keine anderen als die angegebenen Quellen benutzt habe. Alle Koautorenschaften sowie alle Stellen, die wörtlich oder sinngemäss aus Quellen entnommen wurden, habe ich als solche gekennzeichnet. Mir ist bekannt, dass andernfalls der Senat gemäss Artikel 36 Absatz 1 Buchstabe o Gesetzes vom 5. September 1996 über die Universität zum Entzug des aufgrund dieser Dissertation verliehenen Titels berechtigt ist.

Sascha Jakob

Bern, 07.02.2023