



Sentiment Analytics and Financial Markets

Dissertation submitted in fulfillment of the requirements for the degree of
Ph.D. in Economics and Management.

Nicolas Moreno

Supervised by Marie Lambert

HEC Liège, Management School of the University of Liège

June 2022

Jury Committee:

Prof. Dr. Marie Lambert, HEC Liège, Management School of the University of Liège,
(Supervisor)

Prof. Dr. Wouter Torsin, HEC Liège, Management School of the University of Liège,
(Chair of the jury)

Prof. Dr. Gabriela Contreras, Radboud University

Prof. Dr. Theodoros Evgeniou, INSEAD

Prof. Dr. Ashwin Ittoo, HEC Liège, Management School of the University of Liège,

Prof. Dr. Thomas Renault, University Paris 1 Panthéon-Sorbonne

Acknowledgements

“The best way out is always through.”

— Robert Frost

This quote easily sounds true in hindsight, when one starts seeing the light at the end of the tunnel. However, I would never have gotten this far, if it was not for you people, to whom I owe it all. This thesis has been a humbling experience for which I feel deeply grateful, and I learned that nothing worthwhile can be achieved alone. My first thoughts go naturally to Marie Lambert, my supervisor. You gave me the chance to start this journey as a researcher, and for this, I will forever be extremely grateful. I love this work, and research in general. I felt lucky, and knew since the beginning that I was a risky bet; this has certainly been confirmed by the process over the past years. Nonetheless, to take a corporate finance analogy, you must have seen somewhere in me a positive NPV. My wish now, is to make your bet pay off, hopefully today, and even more so in the future. I am deeply grateful for your mentoring and for your teachings, both within the academic sphere and beyond. I hope to carry those learnings into my future life. Your attitude and the success that comes from it is inspiring, and an example I look up to. I feel lucky for the patience that you showed, and for being able to forgive and move on. Really, if I am slowly squeezing my way through this difficult challenge, it is thanks to you. I am proud to have been your PhD student and to call you my thesis promoter.

I want to thank the members of my committee for the effort and thought they put into this project. To Gabriela Contreras, Theodoros Eugeniou, and Ashwin Ittoo, I feel lucky to have benefitted from your mentoring and your encouragements over the years. Each of you brought me unique and valuable insights, which contributed greatly to shape what this thesis has become. I hope it lives up to the high standards you continually set each time we met. To Thomas Renault and Wouter Torsin, I know your expertise and knowledge, and feel blessed that you accepted to be part of the jury.

Ludo, I thank you for accepting me on that project, it has been quite the journey. I enjoyed working together, even late at night, it was thrilling and enriching. I am aware this is a chance very few PhD students have. I am grateful as well for the time you took to listen, you have been a cornerstone in me pushing through hardships and go on. Most of all, I want to thank you for that snowy Liege-Bastogne-Liege. You may not know it, but if I went back to racing after more than ten years, it was entirely because of you and that awesome, humbling experience. Your mindset is an example I look up to. I am also

looking forward to the next challenge at a Gran Fondo, or an upcoming Tour of Flanders.

To Denada Ibrushi, I feel lucky to have crossed paths, it is nice to encounter minds with whom it “clicks” rapidly, you have become a friend who is dear to my heart. It is a joy to work with you. For all the efforts we put into it, I very much hope our work pays off.

To Wale Dare, a special thank you, you have been a brilliant colleague, co-author, but also a guide, despite doing everything remotely. I owe you a lot, and your counsel significantly improved the quality of my work.

To my first colleagues at HEC, Boris Fays, Cédric Gillain, and Maxime Ledent, I fondly remember our shared experiences during those years, both the ups and downs. You guys have been really important all along. To my other colleagues at HEC, Alexandre Scivoletto, Prabesh Luitel, and Jérôme Ruth, I enjoyed working and hanging out with you, a shame that the pandemic hit so early. A big thank you to all the staff at HEC Liege, in particular the academic body, who made it all possible and provided me with the education needed to write this dissertation.

I want to thank my family. To my parents, who have provided the foundation that allowed me to build my life upon. The past years have shown that even in the hardest events, I can peacefully rely on you, as those foundations remain rock solid. To my sister and my brother, Elise and Benjamin, on whom I can always count, especially to cheer me up (even if you guys don't realize it). To my “mamy”, Josette, who has always been there supporting her grandkids with love. To my family in Colombia, who has blessed me with a second home. Gracias Gaby, you are a second mom to me.

Keeping balance outside the endeavor that is a PhD is crucial. I thank my friends on and off the bike, who have brought much needed respites and fun, and helped maybe even without realizing it. Thank you Raphaëlle Mattart, more than a fellow PhD student, you have been a caring friend. Esteban, special kudos to you, for being such a driving force behind those great adventures, those that make life the experience it deserves to be.
Thanks, mate.

Finally, Alix, you deserve a special place in this list. You are a rare person, for whom I have an admiration that you probably do not realize. You saw me at my lowest, and provided a support I could never have expected. You saw me at my best, and made things even better. Beyond the big adventures, in which you play a fantastic role, there are “les petits bonheurs”, and this is just my clumsy attempt of expressing how much I cherish them. There is so much happiness in life. Thank you.

Contents

I	Introduction	7
I	Literature Review	11
A	The causal impact of media on financial markets	11
B	Sources of relevant sentiment for economic applications	12
C	Methods to extract sentiment from text	14
II	Research questions and contribution	18
A	Do value stocks earn a premium due to the way they react to earnings announcement news?	18
B	What do commonalities in anomaly portfolio returns tell us about theories seeking to explain market anomalies?	19
C	What drives a firm to become a value and momentum stock?	20
D	Do private equity firms operate at the expense of target firm employees?	20
II	The Earnings Announcement Day News Value Puzzle	23
I	Introduction	24
II	Data	30
A	News <i>tone</i>	30
B	Discussion on news <i>tone</i>	33
C	Sample Statistics	35
III	Empirical Analysis	37
A	The EAD premium of value stocks is a Bad-News premium	37
B	Systematic Risk	42
C	Biased expectations hypothesis	46
D	Investor attention hypothesis	47
E	Biased Expectations vs. Investor Attention	49
IV	Discussion and Conclusion	51

III	Conditional Drivers of Anomaly Comovement	53
I	Introduction	54
II	Detailed decomposition of long-minus-short portfolio return variance	60
III	Estimation of CF and DR news	63
	A The model	63
	B The data	66
	C Empirical estimation of the VAR model	68
IV	Results	74
	A Simulations on random LMS portfolios	74
	B Detailed variance decomposition	76
	C Conditional variance decomposition	79
V	Discussion	84
VI	Conclusion	86
IV	Dynamics of Cash Flow and Discount Rate News for Value and Momentum Stocks	89
I	Introduction	90
II	Data: Quarterly CF and DR news series	96
III	Results	98
	A Drivers of value and momentum price changes	98
	B Relative contribution of CF and DR to past unexpected return differences	102
	C Long term impacts of CF and DR shocks	104
	D Portfolios constructed on past CF and DR shocks	106
	E Robustness concerns	108
IV	Tone correlates with the observed changes in CF and DR	113
V	Discussion	117
VI	Conclusion	119
V	Private Equity and Employee Welfare	121
I	Introduction	123
II	Data and Descriptive Statistics	128
	A Glassdoor Website	128
	B Capital IQ	129
	C Working Sample	130
	D Type of LBOs	132
	E Glassdoor Ratings	132

F	Glassdoor Job Titles & Salaries	133
G	Industries	136
H	Other variables	138
III	Regression Analysis	138
A	Empirical Strategy and Pre-trend Evaluation	138
B	The Cross-Section of Ratings and PE ownership	140
C	Differences across PE sponsors	142
D	Company characteristics	144
E	Robustness Checks	146
IV	Textual Analysis	149
A	The Latent Dirichlet Allocation Approach	150
B	Extracting topics	151
C	Topic Analysis	154
V	Discussion and Conclusion	162
VI Conclusion		165
VII Appendix		191
A	Detailed TRNA sentiment computation process	191
B	News Tone Surprise	194
C	Principal Ratio	196
D	Alternative VAR model estimations	197
E	Scraping Glassdoor	202
F	Glassdoor anecdotal evidence	206
G	Additional Literature Review	211
H	Case Study: Vista	214
I	Details on position classification	216
J	Number of Reviews around transaction date	217
K	Sub-scores	218
L	Cross-effects	220
M	Fitting LDA models: Topic Coherence	223
N	Detailed LDA topics	225
O	Details on code and implementation of LDA procedure	229

Chapter I

Introduction

“[...] the relation between narratives and economic outcomes is likely to be complex and time varying. The impact of narratives on the economy is regularly mentioned in journalistic circles, but without the demands of academic rigor. [...] But, the advent of big data and of better algorithms of semantic search might bring more credibility to the field.” (Shiller, 2017, p. 48)

WE ARE WITNESSING the advent of the digital age. The amount of information available to decision-makers, academics, and any other economic agent is growing at a dizzyingly fast and exponential pace: the World Economic Forum estimates that up until 2020 about 59 zettabytes (= 59 billion terabytes) of data was created worldwide¹. They project this number to reach 175 zettabytes by 2025, a more than threefold increase in just five years.

To harness all the possibilities of this newly available information, statisticians and computer scientists are developing a growing number of tools to handle and make sense of all this “Big Data”. In the context of finance research, Goldstein et al. (2021) identify three key properties that define what constitutes “Big” data: 1) it is large in size; 2) it presents a high dimensionality; 3) it is encoded in a complex structure.

The current growth in data is driven by the rapidly increasing digital footprint individuals are leaving all over the world. A large portion of it comes in the form of various and rapidly evolving social media platforms that capture a growing amount of information and opinions. News outlets and blogs cover events and companies with ever-increasing depth, and their analyses are constantly growing in number. Customers and users leave reviews and opinions about firms and experiences on a growing number of platforms, which allows economic agents to form increasingly

¹www.weforum.org/agenda/2021/05/world-data-produced-stored-global-gb-tb-zb

more informed decisions.

Consequently, an important portion of this new data comes in the form of text, which by nature fits under the “Big Data” definition: textual datasets quickly reach large sizes of multiple gigabytes, generally much bigger than numerical row-column tables traditionally used in finance and economics. Text is also highly dimensional and innately unstructured, requiring dedicated machine learning techniques from the field of Natural Language Processing (NLP) to uncover humanly understandable relationships between written documents and economic quantities.

Numerous economic outcomes are shaped by subjective perceptions and expectations of economic agents. For example, how employees perceive their workplace will have a direct impact on the performance of the firms they work for (Harter et al., 2002; Huang et al., 2015b; Green et al., 2019). Understanding the motivations and perceptions of key stakeholders has the potential to guide decision-makers in taking appropriate action to maximize corporate and societal value.

Another example is financial markets. By definition, they reflect the expectations of investors by medium of stock prices, which correspond to the equilibrium between supply and demand forces. As a result, any change in price, in excess of priced risk factors, is either a consequence of a revision in investor expectations about future cash flows, or a change in perceived risk.

What information do investors rely on to update their expectations? Word-of-mouth and insider information arguably plays a large part², but a growing body of evidence is showing that media and other sources of publicly available information have a causal impact on financial markets. In his presidential address, Shiller (2017) reviews how narratives spread and their importance for investing decisions. Following this argument, if media play a role in shaping narratives, then, by interpreting news using textual analysis, finance researchers have an opportunity to better understand the underlying drivers of market movements. In turn, this can lead to making better investment recommendations, improving capital allocation and guiding investors to satisfy their individual needs and risk aversion.

An intuitive metric that can be extracted from textual data is *sentiment*, for example by relying on classifiers that rank documents on a scale going from negative to positive tone. Both textual analysis and sentiment are notions that have been

²Roll (1988) provocatively states that, even in hindsight, finance researchers are only able to explain a small fraction of the variation in stock returns. This finding has remained remarkably robust over the ensuing thirty years. Recently, Boudoukh et al. (2019) suggested that identifying relevant information from news, might help challenge this assertion. Therefore, with the progress made in NLP techniques, we might be at the dawn of a change in paradigm regarding our ability to explain what drives asset prices.

gaining traction in the economics and finance literature in recent years, indicating their growing potential to address economic problems and bring new valuable insights for finance researchers.

Consider figure I.1 below: it shows the proportion of working papers hosted in the National Bureau for Economic Research (NBER) database³ which include the keywords “*textual analysis*” and “*sentiment*”. It shows that both subjects have received a quickly rising attention from the academic literature. While textual analysis was virtually non-existent before 2012, its interest has risen sharply among finance academics, especially over the past five years. Similarly, sentiment is a topic that is playing a growing role in modern research, especially in the area of financial markets, since the occurrence of this keyword has roughly doubled over the past fifteen years.

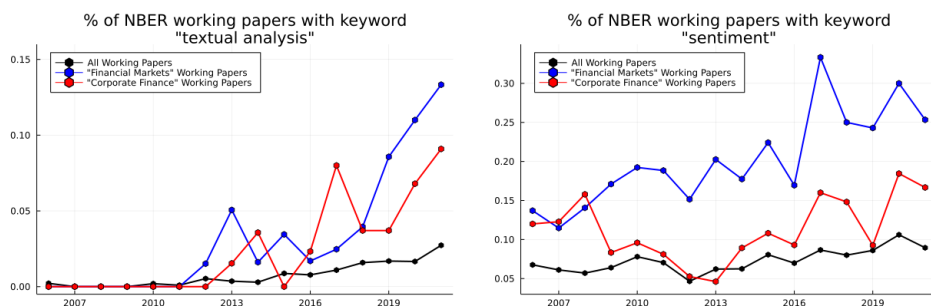


Figure I.1. Proportion of papers hosted on the NBER website including the keywords “*textual analysis*” (left) and “*sentiment*” (right). 15-year trends from January 2006 to December 2021.

Figure I.1 highlights which areas of economics and finance have predominantly used the new opportunities brought forth by text analytics: while overall, less than 3% of all papers in the NBER database make use of textual analysis, this number reaches almost 10% of all papers in the corporate finance section, and concerns just about 15% of all papers related to financial markets. This observation fits naturally with the objectives of this thesis: we use different measures of sentiment extracted from text to address questions related to asset pricing on one hand, and to inform decision-makers in corporate finance on the other.

It is important to define what we mean by sentiment throughout the different chapters of this thesis. Since this notion has become a very popular quantity of interest, different definitions are in use in different contexts. For example, a common definition when linking investor sentiment to uncertainty or investor irrationality is

³www.nber.org/papers?page=1&perPage=50&sortBy=public_date

the following: “*Investor sentiment, defined broadly, is a belief about future cash flows and investment risks that is not justified by the facts at hand.*” (Baker and Wurgler, 2007, p. 129). However, this definition is specific to a particular research question, and is generally not the one we rely on when talking about sentiment.

Instead, we rely on the following definition, which appropriately reflects the multi-faceted nature of sentiment and its ability to provide relevant insights into a wide variety of research questions in economics and finance:

“*Sentiment is the disposition of an entity toward an entity, expressed via a certain medium.*” (Algaba et al., 2020, p. 514).

In particular, in their review of the literature of text analytics in finance, (Algaba et al., 2020, p. 513) coined the term “sentometrics” as being “*the computation of sentiment from any type of qualitative data, the evolution of sentiment, and the application of sentiment in an economic analysis using econometric methods.*”. This encompasses our usage of sentiment in the different chapters of this thesis.

In particular, we make use of two specific sentiment measures. First, we seek to proxy for the content of news stories published in the financial press. We want to understand how news can influence investors in their revision of expectations about a firm. To achieve this, we rely on a sentiment measure computed on the text of those news stories that captures the tone, i.e., the polarity of the information. In this context, sentiment is the disposition of a journalist (the first entity), toward a firm (the second entity), expressed via an online textual news feed (the medium). In particular, we seek to link this sentiment and the associated changes in investor expectations to questions in asset pricing related to the origin of anomaly returns, and help understand why certain firm characteristics predict cross-sectional differences in expected returns.

Second, we seek to proxy for the perception of a key stakeholder in any firm: employees. We achieve this by identifying key complaints that employees (the first entity) make about their employers (the second entity), in written reviews that they post on a social media platform called Glassdoor (the medium). With this form of sentiment, our objective is to inform decision-makers in corporate finance: we seek to assess how changes in firm ownership affect this sentiment measure, and thus the welfare of employees.

The rest of this introduction is structured as follows. In section I, we provide broader context to the thesis and give the reader an overview of the literature in finance dealing with sentiment and textual analysis. First, we motivate the search for sentiment in news by documenting the numerous causal links that the literature

has established between the media and financial markets in section I.A. Section I.B reviews the data sources of “Big Data” this literature has typically resorted to. In section I.C, we document which textual processing techniques the literature has usually relied on to extract sentiment from text, and investigate some state-of-the-art techniques in NLP that could provide new insights from text in future research. Finally, section II highlights the four main research questions that this thesis seeks to address, and section II.D provides an overview of the structure of the dissertation.

I. Literature Review

A. *The causal impact of media on financial markets*

Why study the content of financial media? As shown in figure I.1, in finance and economics, it is the area related to financial markets that has received the most attention. This feature might be explained by the growing literature showing a robust causal link between media and financial markets.

The existence of those links invites researchers to further understand the content of media to guide academics and practitioners alike in their quest to enhance their understanding of drivers of asset prices and updates in investor expectations. In this section, we briefly go over some of these documented causal impacts.

News from the financial press can alleviate market frictions by improving the dissemination of information. Consistent with this hypothesis, Fang and Peress (2009) find that low media coverage firms earn higher returns, even when controlling for usual risk factors. Similarly, Peress (2014) finds evidence that media help to disseminate news and to incorporate them into stock prices. Engelberg and Parsons (2011) propose an experiment that captures the specific effect of media reporting, by analyzing differentiated reactions to the same news in different states receiving different media coverage of the same event. They document that local media coverage is strongly linked to local trading patterns. Dougal et al. (2012) find evidence suggesting that the bias of different journalist cause different market reactions, captured through significant fixed effects on days when those journalists write news in the Wall Street Journal.

Foucault et al. (2016) find that trading rapidly on information is crucial for high-frequency traders, suggesting that public information gets gradually assimilated. Kogan et al. (2021) highlight the role played by news in social media. When fraudulent news are exposed in a company, not only do investors discount all fraudulent information, but the effect spills over to all news published on the social media

platform.

Importantly for chapters II, III and IV, news have also a causal impact on market anomalies. Da et al. (2014) and Hillert et al. (2014) find evidence linking media coverage to the momentum anomaly, as the latter increases with media coverage. Finally, Engelberg et al. (2018) show that anomaly returns are elevated on news-days compared to non-news-days.

B. Sources of relevant sentiment for economic applications

As shown above, we are witnessing a rapid boom in the number of papers in finance that make use of media and text to tackle economic questions. To help others navigate this quickly evolving research corpus, several systematic reviews of this emerging literature have started to flourish⁴. Recently, Loughran and McDonald (2020) provided an update to their initial influential review. Particularly relevant to this thesis, both because of its scope, and because of its extensive focus on “sentometrics” Algaba et al. (2020) provide an extensive review of both the usage of sentiment in economic applications and the methods employed to handle “Big Data” in the context of text.

The field of machine-read text and natural language processing dates back at least to the 1950’s, with an emphasis put on language translation and mapping of syntactic structures (Chomsky, 2009). Unfortunately, those first efforts had little success at their initial objectives. It was only in the late 1980’s that natural language processing experienced a second coming, with a different mindset: the modern idea relies on representing semantic structures using statistical and probabilistic models. Hochreiter and Schmidhuber (1997) were the first to propose recurrent neural networks, marking an important milestone in our ability to capture complex structures inherent to textual data.

From there on, and with the rise of the internet and the need for search engines, NLP has progressed rapidly over the past two decades, being able to perform a wide array of tasks with increasing success. Those include Document Summarization, Machine Translation, Text-to-Speech and Speech-to-Text Conversion, Indexing, and, of particular relevance for this thesis, Topic Discovery and Modeling, and Sentiment

⁴Early attempts include Li et al. (2010), Kearney and Liu (2014), and Loughran and McDonald (2016). Given the continuous development of this research area recently, more updated reviews have followed. Gentzkow et al. (2019a) and Lewis and Young (2019) mainly provide an overview of methods that can be used to extract relevant information from text and applications in the context of finance. Das (2019) and Allen et al. (2020) offer a more global view about the future of financial technologies (FinTech) and the role that big data and natural language processing will have in it.

Analysis.

Those methods have gradually found their way into social sciences, and the fields of economics and finance, in particular, were among the first to leverage their new possibilities. Nonetheless, this early work leveraging NLP only appears at the beginning of the 21st century, often used in a context where the aim is to measure the impact of public information on market movements. Some of the earliest seminal studies include Antweiler and Frank (2004); Das and Chen (2007) and Tetlock (2007). From then on, NLP has been used in a wide variety of social sciences, from politics to history. The rise of social media has further accelerated this trend.

Relevant textual information for social scientists can come in a variety of forms. As mentioned above, in this thesis we focus mainly on two types of sources. First, we use news articles, specifically dedicated about company news. News outlets are an important source of new information, which have been a natural point of interest of the literature. The most important and used data providers in that domain include Dow Jones' (in particular the Wall Street Journal), Factiva, and Reuters, the latter being our chosen source of news.

Making use of those raw data, commercial providers have started providing sentiment metrics databases for news. Those usually assign sentiment scores along a set of predefined dimensions (polarity, uncertainty, emotional indicators, etc.). Examples of such databases used in the financial academic literature include Thomson Reuters MarketPsych (TRMI), Thomson Reuters News Analytics (TRNA), and Ravenpack. In particular, we use TRNA in chapters II and IV of this thesis, and describe its sentiment metrics and construction in detail in section II.

The second data source that we use in this thesis consists of reviews left on websites that can fall under the umbrella term of social media. In particular, we collect reviews of workers about their employers on Glassdoor.com. Glassdoor has the advantage of providing a straightforward sentiment metric which is presented as a score, in which the reviewer rates the company along different dimensions on a scale from one to five stars, similar to what happens on other websites, such as Amazon or TripAdvisor. Several papers have leveraged those ratings. For example, Green et al. (2019) document that firms where employee ratings increase subsequently earn higher returns, or Huang et al. (2020), who find that outlook ratings forecast future operating performance. However, those papers just use the ratings left by employees. Our chapter V provides additional insights into employee sentiment, by refining the approach and grouping the content of the written reviews along a set of topics. This allows us to extract the main dimensions of employee complaints from

the documents.

Other textual sources have also been shown to provide useful information for various finance and economics-related applications. Examples include annual reports (Lopez-Lira, 2020), CEO letters (Boudt and Thewissen, 2019), earnings conference calls (Price et al., 2012; Jiang et al., 2019) and earnings press releases (Huang, Teoh, and Zhang, 2014b; Boudt, Thewissen, and Torsin, 2018; Arslan-Ayaydin, Thewissen, and Torsin, 2021), announcements of central banks (Picault and Renault, 2017), and 10-K filings (Jegadeesh and Wu, 2013; Campbell et al., 2014; Hoberg and Lewis, 2017).

Another popular source of text is Twitter, especially in the present context in which people increasingly gather their news from social media. Azar and Lo (2016) for example, analyze tweets mentioning the Federal Open Market Committee (FOMC) and find that tweets can help predict future returns. Other studies have also used tweets to forecast earnings and returns (Ranco et al., 2015; Bartov et al., 2017). To link social media sentiment to market movements, Renault (2017) and Agrawal et al. (2018) use StockTwits, which is a social media platform specifically dedicated to comment on financial assets.

C. Methods to extract sentiment from text

To classify texts into sentiment metrics, the literature has come up with a plethora of approaches. Here, we group them in three broad categories, as is often done in the literature: lexicon-based classifiers, supervised learning classifiers, and (semi-) unsupervised learning classifiers.

The first approach relies on classifying text using human-constructed lists of words, i.e., dictionaries dedicated to a specific task. Tetlock (2007) and Tetlock et al. (2008) provided seminal pioneering work using a general-purpose lexicon of words (called the “Harvard IV-4 psychosocial dictionary”) to classify news articles based on the fraction of negative terms. Their findings are also economically important, by showing that news have the ability to predict future earnings and market movements.

To improve this approach, Loughran and McDonald (2011) proposed a lexicon specifically dedicated to classify text in financial contexts. This dictionary has been used by others seeking to measure positive vs. negative sentiment in their specific use cases⁵. Other examples include Larcker and Zakolyukina (2012), who build their own

⁵See for example Calomiris and Mamaysky (2019) and Glasserman and Mamaysky (2019) who compute topic-specific news sentiment by leveraging this lexicon, or Jiang et al. (2019) who use it to evaluate manager sentiment in corporate financial disclosures, and find that manager sentiment inversely predicts stock returns. Ahmad et al. (2016) use this lexicon to measure the impact of

model to detect deceptive speech in earnings conference calls, or Baker et al. (2016), who use human-annotated text to quantify economic uncertainty, which serves as a powerful predictor of aggregate investor investment, output, and employment.

Much of the literature has spent significant efforts to build sentiment classifiers based on supervised machine learning approaches. Most often, those methods have shown promising performance, either to classify text with precision out-of-sample, or to nowcast or forecast economic indicators. These types of approaches can take many forms, and several of those have permitted to make discoveries of significant economic importance.

The hallmark of supervised classifiers is that they rely on outcome data that predefines the sentiment categories along which the classification needs to occur. This annotation can either be done manually, or be embedded in the data.

For example, Jegadeesh and Wu (2013) classify words using a regression approach, as either positive or negative, based on how the market reacts to words appearing in 10-Ks. Manela and Moreira (2017) construct a news implied volatility index (NVIX), where news are classified to capture uncertainty using a support vector machine (SVM) depending on the level of the VIX. Their measure appears economically meaningful, as high NVIX appears to be followed by periods of high expected returns. Gentzkow et al. (2019b) and Engelberg et al. (2021) use a LASSO-based estimator to measure how closely speech of republicans and democrats in congress are. This allows them to draw inferences on political partisanship.

Earlier papers often used hand-annotated text to train their machine learning classifiers. Antweiler and Frank (2004) find that bullishness of messages posted on Yahoo! Finance can predict market volatility, and weakly stock returns. Li (2010) documents which firm characteristics best predict tone in firm 10-Ks. Huang et al. (2014a) show that investors react more strongly to negative tone in analyst reports, and that the latter are capable of forecasting future earnings growth. All those studies have in common that their classifier algorithm is a Naive Bayes classifier, which relies on human-annotated datasets.

More sophisticated techniques of machine learning have also been employed to help classify text. One crucial challenge with textual data is its extremely high dimensionality: the number of combinations of words that are possible far exceed anything that a modern computer can handle. Therefore, it is necessary to represent textual data in a more condensed fashion which captures the essential features without losing the essential information.

negative news on stock returns. Others, like Barbaglia et al. (2021) propose their own dictionary by enriching the original Loughran and McDonald (2011) lexicon.

Motivated by this, the literature has made important progress to propose ever more performant vector space models (VSM). The idea of a VSM is to map text into a vector space of reduced dimension, along which all words have a representation which can be used to measure the similarities and relationships between terms. By leveraging advances in deep learning, and the aforementioned enormous amounts of terabytes of data available on the internet as learning corpus, recent models of this sort have achieved impressive results for different tasks in natural language processing, such as question answering, machine translation, or summarization (Radford et al., 2019). Progressively, those models are outperforming any carefully hand-annotated datasets⁶.

Finally, the third group, unsupervised text classifiers, in particular with the aim of deducing topics from text, have found significant traction in the finance literature. One of the most popular models to achieve this task is Latent Dirichlet Allocation (Blei et al., 2003, LDA). The idea of LDA is to represent observations made of atomic elements (in this case, documents made of individual words) as a mixture of an unobserved group of those atomic elements (in this case, topics made from logically connected words). The advantage of LDA, is that it does not require human annotation and achieves highly coherent topics. This is also the method that we employ in chapter V to identify the complaints of employees.

In the literature, multiple papers have used it to answer economic-related questions. Bybee et al. (2021) use it to summarize business news into humanly interpretable topical themes, and find that the representation of certain themes in the media is particularly useful to explain a wide range of economic indicators. Calomiris and Mamaysky (2019) use it to find themes, whose sentiment is particularly useful to predict future market movements. Hoberg and Lewis (2017) use it to find abnormal text in corporate disclosures to detect fraud. Finally, Larsen et al. (2021) use it to extract themes from the media which are useful to predict changes in inflation expectations.

The first of the two sentiment measures used in this thesis is our news *tone* measure, which seeks to serve as a proxy for the content of news stories published in the financial press. We detail the construction and all considerations relating to this

⁶One of the most influential Vector Space Models has been Word2Vec (Mikolov et al., 2013), which is capable of predicting words based on surrounding terms in the text, and can thus leverage context, which is a fundamental feature of language. Pennington et al. (2014) developed GloVe, which is another VSM, which aimed at outperforming Word2Vec. However, today the state-of-the-art method is probably BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2018) and its numerous offspring models that leverage deep learning techniques to develop language models. As of today, this seems like one of the most prominent directions in which the field of natural language processing seems to be heading.

tone measure in section II. Tone is one form of sentiment of an entity about another entity, which captures a measure of polarity ranging on a scale going from “good” to “bad”. We refer to tone as being a semantic measure that aims at capturing how positive (negative) a given text is about an entity. In contrast, we refer to *tone* written in cursive as our specific variable of interest capturing the tonality of news stories about a given firm over a certain period of time. Notice that we may use tone and polarity interchangeably, as in our context both refer to a measure on a continuum between two poles: the likelihood of text about an entity being positive against the likelihood of it being negative.

Tone relies on this second group of classification algorithms, namely on a supervised learning approach. The training set relies on hand-annotated news stories, where human analysts classify them as either positive, negative, or neutral. Furthermore, prior to performing the algorithmic classification task *per se*, the procedure relies on representing the text in a VSM, as described above. The vector space representation is not as advanced as the most recent state-of-the-art models such as BERT described above, but generally does a good job to perform a classification that closely matches the assessment of a human analyst⁷.

Our second measure of sentiment used in this thesis relies on an unsupervised classification method, in particular on LDA⁸. Sentiment of the first entity (the employees) about the second (the employer) is captured by a more subtle measure made up of several dimensions. For instance, we represent each text along a set of 25 topics, which are automatically generated by the algorithm, and that we label to capture the main ideas. Thanks to this representation we are able to find which employees are more susceptible to complain about low salaries, which ones worry to be laid off, or which ones are preoccupied by red tape or inefficient procedures in their company.

Dictionary-based, supervised and unsupervised classification algorithms have their own strengths and weaknesses. For example, when using dictionaries, researchers have the advantage of having a clear mapping of how inputs cause texts to be classified in a specific output. Their main drawbacks however, is that they are limited to the vocabulary defined *ex-ante*, and that they are unlikely to capture complex semantic relationships.

Supervised classifiers can overcome these limitations, especially with recent machine learning advances and the above-described VSM. Generally, they are able

⁷Details on how the VSM is constructed, as well as the performance (precision and recall) of the classification algorithm are provided in appendix A.

⁸Technical details are provided in section II.

to leverage much greater amounts of data and to model complex non-linear relationships. As a consequence, their main drawback is that their classification can ultimately appear too “black box”. It lacks the simplicity and ease of interpretation of a model predicting a linear increase in output “y” if variable “x” increases.

Finally, unsupervised classifiers are able to model complex relationships along naturally relevant dimensions that do not appear evident *ex-ante* to the researcher. However, their main drawback is that their output is often highly sensitive to the input parameters chosen by the researcher and that subsequently, the output might be highly dependent on subjective interpretation. Some procedures to select appropriate input parameters exist, but it is hard to benchmark such models against “optimal outcomes”.

II. Research questions and contribution

The goal of this thesis is not to provide new methods to capture sentiment from media, nor to improve on the methodology of existing ones, seeking to make actionable low-dimensional representations of textual data. Rather, our objective is to tackle economically relevant questions in finance, for which a relevant proxy of sentiment was previously not available. Our goal is to show that leveraging the new sources of information can shed light on phenomena that until now were still poorly understood.

We seek to provide answers to the following four questions.

A. *Do value stocks earn a premium due to the way they react to earnings announcement news?*

Market anomalies earn significantly elevated premia on news- and earnings announcement-days (EAD) (Engelberg et al., 2018). We investigate the origin of this news-premium for value stocks by looking at the content of firm-specific news stories on EAD. We proxy news content with a sentiment measure that we call news *tone*, and find that i) the EAD premium of value stocks is fully concentrated on bad news-days, and ii) growth stocks are much more sensitive to bad news on EAD. An asymmetric response comes about, where value firms appear more resilient to negative media coverage during announcements. Our results fail in attributing the premium to dynamic risk exposures. Nonetheless, we discuss two potential frameworks that do not go against our observed results: the biased investor expectations and the investor at-

tention hypotheses. We review how both explanations have the potential to explain why the EAD value premium is a bad-news premium.

B. What do commonalities in anomaly portfolio returns tell us about theories seeking to explain market anomalies?

What drives anomaly returns? This question has been a primary driver of the asset pricing literature over the past thirty years. As a result, many studies have proposed theories to explain their persistence over time, making predictions about the drivers of anomaly portfolio's unexpected returns. Some theories predict cross-sectional commonalities in revisions to cash flow expectations (i.e., CF news), while others rely on discount rate shock (i.e., DR news) comovement. Still other theories rely on exposures to systematic CF news risk, while others argue that firms with certain characteristics will share greater sensitivities to aggregate DR shocks. To guide research in understanding the source of anomaly returns, we extend the work of Lochstoer and Tetlock (2020), and propose a detailed anomaly portfolio decomposition that allows us to characterize cross-sectional CF and DR news comovement which is specific to each portfolio's long and short leg.

Overall, firms sharing similar anomaly characteristics also share common CF news. Yet, we find significant heterogeneity in unexpected return drivers across anomaly portfolios: Value, small, and loser stocks are those among which we document the most significant comovement in CF news. Contrary to DR shocks, CF news commonalities also vary over time, depending on aggregate CF and DR news indicators. Certain anomalies, such as for value and loser stocks, exhibit stronger comovement during recessions and negative aggregate CF shocks, suggesting a heightened sensitivity to systematic CF risk. In other portfolios, such as among growth and winner firms, we observe that CF commonalities increase in periods of aggregate DR stress, such as high investor sentiment and risk aversion. Overall, our results can help researchers navigate the existing literature on market anomalies, by rejecting certain theories (i.e., theories relying on commonalities in DR shocks) and by reconciling others by separating long and short leg news comovement (e.g., theories explaining the value anomaly based on exposures to aggregate CF risk, and theories relying on aggregate DR risk).

C. What drives a firm to become a value and momentum stock?

The value and momentum anomalies appear as two opposing phenomena. The literature documents several contrasts between the two, suggesting negatively correlated forces driving the returns of the two anomalies. While the literature shows that both anomalies exhibit inverse exposures to systematic risk following market crashes, we shed new light on the opposite drivers behind the two anomalies. We find that value firms are akin to long-term loser socks, which experienced a sustained price decline driven by negative CF shocks and negative news *tone*; yet they are compensated by positive DR shocks, leading to high expected returns. Momentum stocks contrast, in that they experience a price increase driven both by positive CF news and positive media coverage; yet also experience negative DR shocks. The negative DR shocks only materialize years later, as we document a short-term overreaction to CF shocks: we find that cumulated CF shocks over the past four years positively predict the next quarter’s return. Overall, our results contrast the value anomaly as reversal and the momentum anomaly as continuation patterns, echoing the calls of Cochrane (2011) for a joint explanation of both anomalies.

D. Do private equity firms operate at the expense of target firm employees?

There is a profound contrast between the press that private equity (PE) firms receive from politicians in the media⁹, and the often positive aspects brought by PE firms documented by academics¹⁰. Attacks center around PE practices that are supposed to affect employees at target firms, who some argue are at the short end of the stick so that the controlling PE funds can turn in profits. We address this question by reviewing the change in perception of employees at target firms of leveraged buy-out (LBO) deals. We compute this sentiment based on the textual analysis of reviews written on the platform Glassdoor.com.

We find that employee satisfaction decreases following changes in corporate own-

⁹There are many examples, often related to bankruptcies of PE-owned firms that lead to numerous job losses. An example of direct attacks aimed at PE firms includes this quote from senator Elizabeth Warren in a letter to Sun Capital: “years of financial tricks and outright lies to workers are shameful”, and this “financial engineering [...] common in the private equity industry [...] is an egregious betrayal of the people who dedicated their careers to Shopko.” (<https://twitter.com/forrespect/status/1152265448550993920/photo/1>).

¹⁰A few examples that the literature documents include that following LBOs productivity improves (Davis et al., 2014), innovation thrives (Lerner et al., 2011), resilience to economic downturns improves (Bernstein et al., 2019), or that operational processes get streamlined (Bernstein and Sheen, 2016a).

ership, especially following LBOs where the target was previously publicly listed. Moreover, our textual analysis highlights cross-sectional differences across employees induced by the arrival of PE firms: non-managers experience a sharper drop in welfare than managers, which is driven by worries about layoffs, cost-cutting, and lack of management care. The latter issue is particularly worrisome, as it is the complaint affecting employee welfare the most. Nonetheless, we document that certain issues become less frequent following LBOs than after other merger and acquisition deals: issues related to operational processes such as internal politics, lack of communication, slow processes, and overtime, appear less frequently. This allows us to draw a complete and nuanced picture of the problems that employees face when a PE sponsor takes over.

Structure of the dissertation

The rest of the dissertation is structured in such a way that each chapter tackles one of the questions above. Chapter II investigates why value firms earn a premium on earnings announcement, by looking at the content of news stories. Chapter III aims at guiding theories for anomaly premia by uncovering commonalities in CF and DR news shared among stocks with similar anomaly characteristics. Chapter IV investigates why value and momentum portfolios are driven by opposite unexpected return news and highlights the correlated differences in media coverage. Chapter V investigates how employees perceive corporate ownership changes, by extracting sentiment from written employee reviews on Glassdoor.com. Finally, chapter VI concludes by reviewing the main findings of the thesis and discussing implications for future research.

Chapter II

The Earnings Announcement Day News Value Puzzle

DARE WALE[‡] LAMBERT MARIE[‡] MORENO NICOLAS[‡]

APRIL 2022

ABSTRACT

This paper investigates the role played by news on the premium earned by value stocks on earnings announcement days (EAD). We find that the entirety of the dynamic premium is concentrated on *bad-news* EAD, suggesting that accounting for news content matters. Our results fail to support sufficient conditions to attribute the premium to dynamic risk exposures, despite controlling for hundreds of bad-news-specific candidate risk factors. However, our empirical results support necessary conditions for either biased investor expectations or cross-sectional differences in investor attention to hold as valid explanations. The relative resilience of value firms to bad news on EAD, and the asymmetric patterns in news sensitivity when conditioned on trading volume, are pieces of evidence supportive of necessary conditions predicted by those models.

[‡]Wale Dare, Marie Lambert and Nicolas Moreno are affiliated with HEC Liège, Management School of the University of Liège.

I. Introduction

Market anomalies earn elevated returns on earnings announcement days (EAD). Engelberg, Mclean, and Pontiff (2018) show that this phenomenon is robust across a large set of anomalies. The phenomenon even extends beyond EAD: similar return premia for long leg stocks of anomaly portfolios exist on any news-day in general. The source of this premium remains puzzling, as it appears to be difficult to reconcile within a classical systematic risk framework. However, while the premium exists on information-days, the content of the news themselves has until now remained overlooked. This paper aims at bringing a more in-depth understanding of this premium by exploring cross-sectional differences in news content sensitivity.

We focus on the EAD premium concerning value stocks in particular. To capture the media content of a firm, we use news *tone* as a proxy for the news published by the financial press. News *tone* represents an intuitive dimension of any news article: “How likely is the article’s text to be positive or negative for the concerned firm?”. Based on its method of computation, this measure captures the probability that a human analyst would classify a news article from Thomson Reuters as either positive ($tone=+1$) or negative ($tone=-1$), for a given firm of interest. This exogenous and daily proxy for news content allows us to make two novel contributions to understand the premium of value stocks on EAD.

Our first contribution is to show that earnings announcements are not the only important factor in the generation of elevated anomaly returns: the content of the news during EAD also matters. Indeed, compared to other news-days, differences between value and growth stocks are particularly strong on EAD. On those days, growth stocks are more sensitive to *tone*, but this elevated sensitivity only occurs on bad news releases. Importantly, we are able to fully link the EAD anomaly premium to the subset of bad EAD news. Controlling for this “bad news effect” entirely captures the dynamic premium of value stocks on EAD, suggesting that the EAD-value premium is, in fact, a bad news premium.

Beyond cross-sectional differences, the fact that on average stock returns are robustly higher on earnings announcement days has been documented for decades¹. For example, Ball and Kothari (1991) find a strong EAD premium that is robust to increases in systematic risk exposure. Even years later, researchers found that the premium continues to exist (Cohen et al., 2007) and that it is a global phenomenon, occurring worldwide across different markets (Barber et al., 2013). Yet, to this

¹For early papers that document an earnings announcement premium, see for example Penman (1984), Chambers and Penman (1984), Kross and Schroeder (1984), and Chari et al. (1988).

day, the origin of this premium remains subject to debate. Broadly speaking, two theoretical strands currently co-exist that seek to explain its occurrence.

The first theory relies on the idea that announcing firms face a dynamic exposure to a systematic risk factor. It has the advantage of explaining the persistence of the premium, which appears not to be arbitrated away over the years, despite being well-known to investors. Savor and Wilson (2016) propose the following explanation for the nature of this dynamic risk: Announcing firms provide valuable information to investors about the aggregate market that they can use to revise their expectations about systematic factors. Therefore, the more a firm’s announcement informs investors about the aggregate economy, the more the market beta of that firm should spike on EAD, resulting in a positive premium. Their empirical results support this hypothesis and Patton and Verardo (2012) also find that firms with news have elevated betas when looking at intraday data.

Another class of theories seeking to explain the EAD premium revolves around the investor attention hypothesis of Barber and Odean (2008). The idea is that individual investors tend to buy stocks that are getting the most media coverage or significant attention-grabbing events. Because of their limited attention capabilities and the considerable pool of candidate firms to invest in, their focus is overwhelmingly biased towards companies sending the most visible cues. Furthermore, when a negative signal about a company comes up, those same investors will likely sell stocks they already own rather than researching other firms to sell short². Since earnings announcements are a cue driving focus on a firm, the resulting prediction is that attention-constrained investors will push prices up on EAD by creating a net buying pressure on announcing firms. Chapman (2018) provides evidence in favor of this hypothesis by showing that announcements drive up EDGAR searches (i.e., lead to higher investor attention), that first-time investors are biased to buy announcing firms, and that pre-announcing firms see their EAD premium weakened.

Our second contribution aims at investigating whether the “alpha” of value firms on EAD can be explained away in light of the newly identified bad-news nature of the premium. Since our central finding is that cross-sectional differences in the EAD premium depend on the content of the news, we review the theoretical frameworks within which cross-sectional variability can arise on EAD. We identify three hypothe-

²Insofar as arbitraging the behavior of individual investors and taking short positions is required, this theory also predicts a positive link between the EAD premium and trading frictions. For example, Cohen et al. (2007) and Barber et al. (2013) find that the magnitude of the EAD premium increases with idiosyncratic risk. Furthermore, they argue that the premium might be difficult to arbitrage away, given that it would require high levels of portfolio turnover to buy and sell all the announcing firms in a quarter.

ses that could lead to the observed “bad-news” EAD-value premium: cross-sectional differences in either i) dynamic risk exposures ii) mispricing and biased expectations, or iii) investor attention on EAD. Each hypothesis makes different predictions about necessary (and sometimes sufficient) conditions for it to hold, which we can explore empirically.

The first hypothesis relies on the extension of the above notion that announcing firms face dynamic increases in risk sensitivity. If cross-sectional differences exist in this dynamic risk exposure, in particular on bad-news EAD, then it might be possible to link the EAD value premium to this phenomenon. A distinct advantage of this framework is that testing a sufficient condition to infer its causal relationship with the dynamic EAD premium can empirically be set up: If bad-news announcing firms load positively on a systematic risk factor, *and* this exposure fully captures the EAD premium, then it would be sufficient evidence to validate this model.

Engelberg et al. (2018) leverage this notion and explicitly test if the premium for stocks in anomaly long portfolios are caused by a dynamic risk exposure to the market factor or the long-minus-short anomaly portfolio return. Like them, we find a positive exposure to these dynamic risk factors. However, despite controlling for the dynamic risk exposure, the EAD premium of anomalies remains unchanged and strong, suggesting that they do not fit the requirement to be the driver of the EAD value premium. We further extend this approach and investigate if cross-sectional differences in systematic factor sensitivity specifically occur on bad-news EAD; but again the premium remains robust and unchanged.

Furthermore, since traditional factors fail to explain the anomaly EAD premium within the dynamic risk framework, we test if exposure to unconventional systematic factors, including aggregate volatility and news *tone*, could be driving the elevated returns. Again, controlling for those factors leaves the EAD value premium unchanged. We also seek to uncover if any other form of co-movement in stock returns that can be statistically extracted might be an appropriate dynamic risk factor. When decomposing stock returns using a principal component analysis (PCA), we find that no common factor in returns among the first seven-hundred principal components (making up about 99% of all the existing variation in returns) captures the EAD value premium in a dynamic setting. Therefore, over multiple extensions, we strengthen the findings of Engelberg et al. (2018) that cross-sectional differences in the EAD premium are likely not driven by systematic risk exposures.

The second hypothesis that can explain cross-sectional differences in reactions to news on EAD relies on mispricing and biased expectations. The idea is that investors

revise their overly optimistic or pessimistic priors to accurately reflect future cash-flow expectations on days when news are released to the public. This framework predicts an asymmetric response to news of value and growth stocks on EAD. The necessary condition for this hypothesis to hold true is that either value (growth) stocks react less (more) to bad news than growth (value) stocks on EAD, or that value (growth) firms react more strongly (weakly) to good EAD news.

Our results support this necessary condition, as we observe an asymmetric reaction, where value stocks have a relatively more muted response to bad news on EAD relative to growth stocks. Conversely, we find no overreaction to good news of growth stocks, suggesting that if investors form biased expectations, it must be that mispricing gets corrected on bad-news EAD. Specifically, if investors are overly optimistic (pessimistic) about growth (value) stocks, then a piece of bad news for the latter will trigger an abnormally large (small) negative price adjustment.

Basu (1977) already theorized that erroneous expectations will impact the formation of stock prices. A large corpus of research argues in favor of the existence of such expectational errors. In asset pricing, several authors leverage this notion to explain market anomalies³. Greenwood and Shleifer (2014) provide survey evidence in favor of investor over-extrapolation of past performance. Even managers themselves might be subject to such biases: Bradshaw et al. (2006) find that firms over-optimistic in their financing decisions subsequently earn low stock returns.

La Porta (1996) and La Porta et al. (1997) provide evidence in favor of systematic biases differing across book-to-market characteristics. In this framework, the biased expectations hypothesis relies on the idea that growth (value) stocks are more (less) sensitive to bad news because investors over-extrapolate past performance. By definition, growth stocks experienced a more positive past performance compared to value firms given their relatively higher price levels. In turn, those firms get often labeled as “glamour” firms. Our novel news dataset allows us to shed new light and further justify this notion: growth firms tend to get both significantly more media coverage, as well as a significantly more positive *tone* in their media coverage. Following this hypothesis, this “glamorous” bias would lead investors to form overly optimistic expectations about growth firms relative to value firms; which finally get adjusted when reality hits with bad news on EAD.

Engelberg et al. (2018) also argue that the biased expectations explanation likely

³Several popular models based on expectational biases seek to explain asset price behaviors that are otherwise challenging to reconcile in the classical framework of rational expectations (see for example De Bondt and Thaler (1985), Barberis et al. (1998), Daniel et al. (2001) or Barberis et al. (2015)).

fits best with their results. They show that analyst forecasts are overly optimistic about stocks in long legs of anomaly characteristics. Moreover, consistent with biased expectations, Skinner and Sloan (2002) document a similar asymmetry, where growth stock prices have stronger adverse reactions to disappointing earnings surprises. To explain the value anomaly, they rely on quarterly stock returns and earnings surprises. Here lies another key distinction of our study and added value of our dataset: we can leverage daily news *tone* to highlight the unique patterns of earnings announcement days compared to any other news day.

The third hypothesis that can explain the EAD-value premium relies on cross-sectional differences in investor attention (Barber and Odean, 2008). As detailed above, attention-constrained investors will predominantly be driven to buy stocks that attract their focus in the first place, for example through increased media attention. Yet again through the attention mechanism or because of short-sale constraints, those same investors will be more likely to sell stocks they already own when negative news hit those firms. This hypothesis, therefore, predicts as a necessary condition that high-attention stocks will experience a sharper drop in price when bad news are released.

Since growth stocks have particularly strong negative reactions to bad EAD news, if those firms are also on average better “attention-grabbers” than value stocks, then this hypothesis falls in line with our results. Previous studies show that individual investors tend to be more prone to buy stocks that get important or extreme news events (Lee, 1992; Hirshleifer et al., 2008; Lawrence et al., 2018), no matter the direction of the news. Moreover, past performance might also be an important driver of investor attention, as Aboody et al. (2010) show that firms with the largest past price increases tend to earn abnormally high returns on EAD and to attract individual investors. Hence we look at media properties of value and growth stocks. We find that past price increases and low book-to-market correlate with subsequent media coverage and extreme news *tone*, which further suggests that growth firms are more likely to attract investor attention both prior to and during EAD.

Lamont and Frazzini (2007) also connect the premium around EAD to increases in trading volume. Consistent with the investor attention hypothesis, they argue that small, unsophisticated investors might drive these patterns. If it is bad and not good news that matter on EAD for attention-constrained investors, this idea predicts that on bad EAD news with high volume, growth stocks will experience relatively sharper price declines than value stocks. Conversely, differences between value and growth stocks should be much smaller on good EAD news, since the bias

of selling stocks already owned does not apply.

This is exactly the pattern that we observe empirically: the difference in returns between value and growth stocks returns increases exponentially as abnormal trading volume levels increase on bad news EAD. However, there is virtually no return differential between value and growth stocks on good EAD news, even for the highest levels of abnormal trading volume.

Finally, we propose an additional empirical test to compare predictions made by the biased expectations and the investor attention frameworks. We rely on the well-studied post-earnings announcement drift (PEAD) (Ball and Brown, 1968), which is the robustly documented⁴ tendency of stock prices to continue drifting in the direction of the earnings surprise in the days and weeks following the announcement.

Interestingly, the biased expectations hypothesis and the investor attention hypothesis make opposite predictions regarding price drifts following announcements in the cross-section. If bad news on EAD allow for the correction of a mispricing error, we would anticipate no significant price drift for growth stocks following bad EAD news if the price effectively went back to a just level on EAD. However, the investor attention hypothesis relies on strong negative price movements driven by sales of attention-constrained investors. As a result, this theory would predict that cross-sectionally the PEAD should be weaker for bad-news growth stocks, as the initial sharp price drop would in that case not reflect an initial adjustment of mispricing.

Based on this framework, our empirical results provide further support for the investor attention hypothesis. Following good EAD news, we observe no cross-sectional differences in the PEAD behaviors⁵. However, following bad EAD news, growth stocks experience a relative reversal compared to value firms: this suggests that the initial overreaction to bad news on EAD warrants a subsequent re-adjustment and is possibly not fully a reflection of mispricing adjustment.

The rest of the paper is structured as follows. Section II describes the data used in the paper and discusses key differences of our *tone* measure relative to previously used measures of earnings surprises. We document our main results related to negative *tone* asymmetry in section III.A. In section III.B we test a large array of dynamic risk factor candidates. Sections III.C and III.D discuss the biased expect-

⁴Several studies keep documenting the existence of the PEAD in modern times. See for example Bernard and Thomas (1989); Mendenhall (2004); Sadka (2006) and Chordia et al. (2009).

⁵Consistent with the well-documented PEAD, we find that unconditionally, returns keep being significantly higher (lower) in the days following positive (negative) *tone* announcements. This suggests that our *tone* measure is able to capture similar surprises as the standardized unexpected earnings (SUE) metrics used in the PEAD literature.

tations hypothesis and the investor attention hypothesis, respectively, as alternative theoretical frameworks, and section IV concludes.

II. Data

A. News tone

We label the measure of sentiment that we use in this paper news *tone*⁶. We obtain this proxy for news content from the Thomson Reuters News Analytics database. This database recollects all the US firm-specific news stories that were published in the Reuters news feed⁷ starting in January 2003. For each of these news stories, Reuters computes a score⁸ which aims at capturing the tonality employed by the author of the text about a specific company. Importantly, the computation of the semantic score is independent of the interpretation that a potential reader could make based on her prior knowledge. Therefore, it is the way in which the author spins the story and his choice of vocabulary that will be determinant. Consider the following two examples provided by TRNA:

1. “An explosion occurred in Iraq today, killing 20 people.”
2. “A horrific explosion occurred in Iraq today, murdering 20 people.”

The second version of the story contains loaded words (horrific, murdering) that unambiguously point towards a negative event. The tone of the first story is more neutral. Therefore, TRNA will assign a more negative score to the second story.

In the TRNA database, sentiment is computed at the entity level, i.e., at the firm level. Indeed, instead of providing the general tone of a news article, TRNA parses the text to identify entities, such as companies, by their names or tickers for example. This is an important feature for us, as it allows for the computation of firm-specific news tone.

In appendix A, we detail the procedure that Thomson Reuters employs to compute firm-specific news sentiment measures. In a nutshell, the approach consists of

⁶We refer to tone as being a semantic measure that aims at capturing how positive (negative) a given text is about an entity. In contrast, we refer to *tone* written cursive as our specific variable of interest capturing the tonality of news stories about a given firm over a certain period of time. Notice that we may use tone and polarity interchangeably, as in our context both refer to a measure on a continuum between two poles: the likelihood of text about an entity being positive against the likelihood of it being negative.

⁷In a sub-sample of 50.000 news stories, the original source of the news items were traced back to Thomson Reuters (48%), PR Newswire (26%), BusinessWire (23%), Regulatory news services (1%) and other miscellaneous news sources (2%).

⁸The description of the present procedure is based on the Thomson Reuters White Paper for version 1.1 of TRNA.

first representing the text of each news article in a reduced dimensionality space. Their approach is similar to the VSM described in chapter I. They call this step of reducing the vector space “feature extraction”, which they achieve thanks to a model trained by human-annotated data. Finally, they input this low-dimensionality representation of text into a three-layered back-propagation neural network algorithm to compute the probability that a news falls in either of the following three groups: “Negative”, “Neutral”, or “Positive”, such as the output of each news-firm pair leads to the following:

$$\mathbb{P}(pos) + \mathbb{P}(neg) + \mathbb{P}(neut) = 1 \quad (\text{II.1})$$

Appendix A provides statistics on the classification accuracy and precision of the method.

Let’s consider an example to better illustrate the resulting output that we get to work with:

“JP Morgan Chase, Barclays, Goldman Sachs to load CDW IPO.

NEW YORK, March 7 (Reuters) - Technology products retailer CDW, which was taken private by Madison Dearborn Partners LLC and Providence Equity Partners for \$7.3 billion in 2007, has hired banks for an initial public offering later this year, people familiar with the matter said. CDW, which sells products from companies including Apple, Hewlett-Packard and International Business Machines online and through its catalog, has hired JPMorgan Chase, Barclays PLC and Goldman Sachs Group to lead the offering, the people said on Thursday. The proposed IPO could raise about \$750 million.”

In this example, multiple entities are brought up. The article mentions banks, i.e., JP Morgan Chase, Barclays, and Goldman Sachs; private equity sponsors, i.e., Madison Dearborn Partners and Providence Equity Partners; clients, i.e., Apple, Hewlett-Packard, and IBM; and finally there is the target company, CDW. Thomson Reuters computes a separate score for each of those companies, which it recognizes thanks to its hand-annotated dictionary of entities and the part-of-speech tagging described above.

For example, let’s assume that we are interested in gathering information about Goldman Sachs. TRNA will first provide us with a *Relevance* score. This score, which ranges between 0 (the entity is irrelevant in the article) and 1 (the story is fully dedicated to the entity), indicates how prevalent a company is in a given

story. In the news article above, Goldman Sachs has a moderate relevance score of 0.47. This score is again based on the part-of-speech tagging: firms appearing in the headline, being mentioned multiple times and appearing as the main subject will have higher relevance scores. For example, CDW has a relevance of 0.91 in the above article⁹.

Next, TRNA provides us with the classification probabilities. In the above article, Goldman Sachs gets the following score assignment: $\mathbb{P}(\text{pos})=0.62$, $\mathbb{P}(\text{neut})=0.15$, and $\mathbb{P}(\text{neg})=0.23$. Overall, the tone is thus positive, which seems reasonable since Goldman Sachs here is involved in an IPO transaction which probably will earn them fees for their services. Along the same line of reasoning, TRNA assigns a high probability of the news being positive for CDW, as the latter is showing the promise of raising additional money for its operations.

Several studies in economics and finance have utilized TRNA to capture the sentiment content of news stories. Some of the most impactful papers that have used this data include Hendershott et al. (2015) and Heston and Sinha (2017) to handle firm-level problems, Calomiris and Mamaysky (2019) to forecast macroeconomic movements, or even Smales (2014) in the context of commodities. Other papers that have used TRNA include Uhl et al. (2015); Gotthelf and Uhl (2019); Allen et al. (2019); Griffith et al. (2020) and Uhl and Novacek (2021).

Of particular interest for the validity of the data is Calomiris and Mamaysky (2019), who have replicated their results with TRNA and their sentiment classification measure, leading to highly correlated indicators and similar results in their analyses.

Our sample of news stories covers the period going from January 1st 2003 to December 31st 2017. We match the Reuters entities with the CRSP/Compustat database by using CUSIP identifiers.

We apply a set of filters before including a news story to our final sample: First, we impose that a company needs to have a relevance score of at least 0.8 or more¹⁰. This is done to reduce noise and discard news stories that likely will have only a moderate effect on the perception of a company's future expected cash-flows and

⁹As a side note, our relevance score indicates how likely a news is going to have a given firm as a primary protagonist. We acknowledge that an additional limitation of our metric is that all news get the same, and thus "matter the same". In reality, certain news will reflect more strongly a firm's changes in expected cash-flows and discount rates, but our relevance measure is unfortunately not able to distinguish those particularly important news.

¹⁰This threshold is motivated by the work of Boudoukh et al. (2019), who show that to capture unexpected news from the media it is important to filter out poorly relevant news. It also follows previous studies using TRNA who have used relevance thresholds ranging anywhere between 0.35 (Heston and Sinha, 2017) and 1 (Smales, 2015).

discount rates. Second, we seek to eliminate repetition by discarding all news stories with a low novelty. In particular, only novelty scores of 0 over the past 24-hours are considered¹¹, i.e., no similar news about this particular firm has to be published over the past 24-hour cycle, to avoid news repeating the same information over two succeeding days and thus drown the signal in noise. This decision is especially relevant in this chapter, where we work with daily data.

We end up with 2.25 million individual news stories. Since our analysis relies on matching the polarity of said news stories to stock returns, we compute news *tone* for each firm over our chosen trading period as follows:

$$tone_{i,t} = \frac{1}{N_{i,t}} \sum_{j=1}^{N_{i,t}} \mathbb{P}(pos)_j - \mathbb{P}(neg)_j \quad (\text{II.2})$$

where $N_{i,t}$ is the number of news stories for company i over period t , and $\mathbb{P}(pos)_j$ and $\mathbb{P}(neg)_j$ are the individual positive and negative probabilities for each individual story concerning firm i , as defined in equation (II.1). Thus, news *tone* is simply the average difference between positive and negative probabilities. Scaling by the number of stories over the period allows us to compute a variable that is systematically comprised between -1 (i.e., all news over period t about firm i are negative with a probability of 100%) and +1 (all probabilities are 100% positive).

Since we use daily stock returns, the chosen time interval t is equal to one trading day, going from 15 minutes prior to market close (i.e., 3:45 p.m. ET) up to 15 minutes to following market close time¹².

B. Discussion on news tone

Notice that news *tone* is a measure that we obtain from probability outputs of a classification model. We could have used simple categorical polarities (-1, 0, and +1), but chose to capture news *tone* by averaging the individual probabilities over a day for a firm for two reasons. First, appendix A shows the existence of a linear increase in concordance between the TRNA classification probability and human analysts. We interpret this as a sign that the greater the positive (negative) probability, the more distinctively positive (negative) a news is, hence giving us a greater chance

¹¹This decision is directly motivated by Tetlock (2011), who shows that stale news are more likely to cause noisy price movements, usually unrelated to firm cash-flows and that are only transitory in nature.

¹²In chapter IV, we focus on quarterly returns and thus aggregate *tone* over a longer period of one quarter, which matches with the data frequency.

to distinguish moderately bad information from very bad news. Second, since our *tone* measure potentially encompasses multiple news on a given day, working with a continuous variable allows us to give more weight to news with large probabilities, thus diminishing the risk of missing out on important news signals.

We acknowledge that we proxy for the entire content of text with a single tone measure, whereas news text might also contain other types of semantic information, relevant to explaining changes in investor expectations. For example, independently from vocabulary allowing to distinguish “good” from “bad” news, some features of the text might inform investors about the (un)certainty of a given news. Depending on her priors, an investor might interpret certain news as being likely to happen (and thus give them a heavy weight), whereas another might find this information irrelevant. However, capturing all the subtleties of text is very difficult, and humans themselves might not agree on how to interpret the same piece of evidence. This is why we limit ourselves to an easily interpretable metric, i.e., the likelihood that a human analyst would classify the news as good or bad for a firm.

In the literature, measures of sentiment take different forms and can have different objectives. For instance, based on the idea that noise traders can influence asset prices¹³, Baker and Wurgler (2006) seek to define “sentiment” as anything that influences arbitrage forces, either through impacts on investor optimism or pessimism or through their propensity to speculate. When thinking about “investor sentiment”, the literature often makes reference to this popular measure, for which Zhou (2018) provides a review, and to which Huang et al. (2015a) has developed a noteworthy extension to predict aggregate returns.

Much closer to our *tone* measure, Tetlock et al. (2008) quantify firm-specific sentiment by counting the proportion of negative terms in individual news articles. They find that more negative news forecast negative earnings and that news have a stronger impact when focusing on fundamentals. Further focusing on fundamentals, many studies investigating market phenomena around earnings announcements rely on standardized earnings surprises¹⁴, which can for example be computed as the difference between analyst forecasts and realized earnings. Other papers proxy for firm-specific information by extracting manager tone from financial disclosures (Jiang et al., 2019) or by quantifying the tone of press releases present in 8-K filings published on the SEC website (Boudt, Thewissen, and Torsin, 2018).

Our *tone* measure is agnostic about the underlying causes of why a piece of news

¹³See for example the model of De Long et al. (1990).

¹⁴Many papers rely on this kind of measure. Some examples relevant in the context of this study include Ball and Kothari (1991), La Porta et al. (1997), or Skinner and Sloan (2002).

might be “good” or “bad”. News in Thomson Reuters might reflect fundamental information, for example by reporting that a company exceeds (= positive *tone*) or misses (= negative *tone*) expectations. They might also reflect investor sentiment which might appear independent from fundamentals going forward at first sight. For example, if news stories relate past price movements or events related to recent social media perceptions (think about the coverage received by GME following the social media hype in January 2021).

Ultimately, *tone* depends on the spin of the story. Therefore, our underlying assumption is that our measure captures relevant differences in expectations. For example, a 10% growth in revenues might appear like a positive change in a vacuum, but if the expectations were of 20%, the market might be heavily disappointed. This is the strength of our measure: if realizations miss the mark of expectations, it gets translated into the vocabulary employed in the news. Finally, compared to other measures that could be obtained from earnings calls or accounting information, our measure also consistently computes *tone* on every single day with stable assumptions, over a wide range of events and news-worthy information, always differentiating between the news that are presented as being either good or bad.

C. Sample Statistics

For our sample of daily data, the 2.25 million stories translate into 1.50 million news-return day pairs. This corresponds to $\sim 14\%$ of all daily return observations which are matched to a news *tone* measure. On average, news *tone* is positive, as the raw *tone* measure equals 0.21. Therefore, the average news is more positive than a no-news day, to which we assign a *tone* of 0. We observe a lot of variability, as overall the standard deviation of raw *tone* is equal to 0.40.

The rest of our data comes from the merge of the CRSP and Compustat databases. Daily stock returns, prices, number of shares outstanding, and trading volume come from CRSP, while we collect book-equity data from Compustat. We adjust stock returns by subtracting delisting returns if available. We compute market equity at the end of December of year $t-1$ by multiplying the number of shares outstanding by the stock price. We compute book-equity as in Fama and French (1993) at the end of June to ensure valid book-to-market data for the year is public. We classify firms as “Growth” for a year if their book-to-market is in the lowest quintile at the end of June, and classify them as “Value” if they appear in the highest quintile. We adjust daily stock returns by subtracting delisting returns when available. Due to data availability of the *tone* measure, our sample goes from January 1st, 2003 up to

December 31st, 2017.

Finally, notice that we furthermore standardize the tone measure by subtracting the average of all observations and dividing by the standard deviation. This scaling results in *tone* values with a mean centered in 0 and with a standard deviation equal to 1, which eases the interpretation of regression coefficients¹⁵.

Table I - Summary Statistics. This table reports summary statistics of the main variables of interest throughout this study. The left half reports the statistics estimates for the EAD subsample only, while the right half reports the statistics for the sample of observations outside EAD. “Value” and “Growth” refer to all observations classified as value and growth stocks, respectively. “All” includes all firms. “T-stat Diff_{All}” reports the estimate for the Welch t-test difference between the mean estimates between value and all other stocks, and growth and all other firms. *AVol* is the abnormal trading measure as in Lamont and Frazzini (2007), which is equal to the ratio of a day’s trading volume divided by the trailing 365 average. All other variables are described in the main text.

		EAD							Non-EAD				
		Mean	Std	Skew	Kurt	Min	Max	T-stat Diff _{All}	Mean	Std	Skew	Kurt	T-stat Diff _{All}
<i>tone</i>	All	-0.23	0.74	-0.20	0.00	-2.40	1.72		0.00	1.01	-0.23	-0.46	
	VAL	-0.31	0.77	-0.13	-0.10	-2.39	1.72	-12.07	-0.06	0.99	-0.16	-0.38	-33.56
	GRO	-0.21	0.73	-0.20	-0.06	-2.38	1.72	1.07	0.03	1.03	-0.25	-0.53	2.76
<i>surp</i>	All	-0.01	0.19	-0.38	3.49	-1.32	1.10		0.00	0.13	-0.96	21.19	
	VAL	-0.01	0.19	-0.47	4.50	-1.03	0.98	-7.85	0.00	0.12	-0.86	27.23	-4.45
	GRO	0.00	0.20	-0.36	2.63	-1.09	1.10	-0.76	0.00	0.15	-0.95	16.00	-0.68
<i>NDay</i>	All	0.53	0.50	-0.10	-1.99	0.00	1.00		0.14	0.34	2.14	2.56	
	VAL	0.44	0.50	0.24	-1.94	0.00	1.00	-26.23	0.11	0.31	2.53	4.39	-81.09
	GRO	0.59	0.49	-0.38	-1.85	0.00	1.00	24.28	0.17	0.38	1.73	1.01	160.88
<i>Bad</i>	All	0.18	0.38	1.70	0.88	0.00	1.00		0.04	0.20	4.67	19.84	
	VAL	0.17	0.37	1.77	1.14	0.00	1.00	-1.28	0.04	0.18	5.03	23.32	-16.09
	GRO	0.20	0.40	1.53	0.34	0.00	1.00	10.80	0.05	0.22	4.10	14.78	82.01
<i>Good</i>	All	0.08	0.28	3.02	7.11	0.00	1.00		0.05	0.21	4.38	17.15	
	VAL	0.06	0.24	3.58	10.78	0.00	1.00	-12.45	0.03	0.18	5.27	25.81	-74.08
	GRO	0.10	0.30	2.72	5.40	0.00	1.00	7.80	0.06	0.24	3.73	11.88	93.68
<i>AVol</i>	All	3.44	4.16	4.19	22.58	0.00	33.35		1.26	1.98	8.86	110.54	
	VAL	3.33	4.66	4.07	19.92	0.00	33.34	-0.66	1.28	2.26	8.07	89.62	16.44
	GRO	3.77	4.16	4.01	21.12	0.00	33.35	16.97	1.26	1.90	9.13	117.24	9.54

Table I presents summary statistics about the variables of interest in this study, in particular the standardized *tone* measure described above. Overall, we find that value stocks get much more negative media coverage than average. Their average news *tone* is -0.06, which is significantly worse than average, whereas the average news article about growth firms has a *tone* of 0.03. News *tone* is also significantly

¹⁵e.g., a regression coefficient of 5% related to *tone* could be read as follows: “a one standard deviation increase in *tone* implies a 5% increase in return”. We privilege this scale for its intuitive interpretation.

less positive on EAD, and value stocks keep having worse news than the others. Standard deviation of *tone* also decrease on EAD, suggesting that news are more centered around the mean. The distribution of *tone* is negatively skewed, suggesting ample negative tails and a greater proportion of very bad news than very good news. This asymmetry is also reflected in the minimum and maximum values of *tone*.

As a robustness concern, we compute an alternative measure to *tone*, which we label *surp*, for “surprise”. It is defined as the *tone* in excess of the expected *tone*. Indeed, we show in appendix B that *tone* is a highly persistent variable, and past *tone* predicts subsequent *tone*. Since investors might be more susceptible to react to unexpected news than to stale information, (Tetlock, 2011) we model expected *tone* with a linear regression model. It takes past firm *tone*, EAD dummies, and value/growth dummies as independent variables¹⁶. The standard deviation of *surp* is significantly lower than for *tone*, but otherwise, we make similar observations as before, with value firms getting significantly worse news coverage on average.

Not only do growth firms get better news coverage than value firms, but they also get a significantly more extensive amount of coverage. Both on EAD (59% vs. 44%) and on non-EAD (17% vs. 11%), growth stocks have more news-day observations, consistent with their image as glamour firms that attract much attention.

Bad and *good* news-days are defined as observations where a valid *tone* measure falls outside the range of one standard deviation around the mean, i.e., below -0.5 for *Bad* news, and above 0.5 for *Good* news. A greater proportion of observations pertain to the *Bad* than the *Good* news dummy¹⁷ during EAD. However, we observe the opposite on non-EAD. Consistent with their greater overall media coverage, growth stocks get both more *Good* and *Bad* news than value firms.

III. Empirical Analysis

A. The EAD premium of value stocks is a Bad-News premium

Engelberg et al. (2018) show that long legs of anomalies, such as value stocks, earn a premium on news days and on EAD. Our starting point is to understand how value and growth stocks react to public information using our measure of news *tone*. We follow their approach and evaluate how value and growth stock’s returns differ on EAD:

¹⁶The choice of those variables to model expected *tone* is justified in appendix B.

¹⁷The usage of italic versions of *bad* and *good* in this text specifically refer to those two dummies.

$$\begin{aligned}
r_{i,t} = & \alpha_t + \beta_1 EAD_{i,t} + \beta_2 Growth_{i,t} + \beta_3 Value_{i,t} \\
& + \beta_4 EAD_{i,t} * Value_{i,t} + \beta_5 EAD_{i,t} * Growth_{i,t} + \epsilon_{i,t}
\end{aligned} \tag{II.3}$$

where $r_{i,t}$ is the stock return of firm i at day t , EAD, Value and Growth are dummies for matching observations and α_t are intercepts that capture time fixed-effects. β_4 and β_5 capture how much higher value and growth stock returns are on EAD than the average firm (non-value and non-growth). Following Engelberg et al. (2018), if value stocks earn a premium on EAD, we expect a significant positive β_4 and negative β_5 , where the spread between the two captures the magnitude of the EAD-value premium. This specification can be extended to include a variable on which to condition returns of value and growth stocks on EAD as follows:

$$\begin{aligned}
r_{i,t} = & \alpha_t + \beta_1 EAD_{i,t} + \beta_2 Growth_{i,t} + \beta_3 Value_{i,t} \\
& + \beta_4 EAD_{i,t} * Value_{i,t} + \beta_5 EAD_{i,t} * Growth_{i,t} \\
& + \beta_6 Xvar_{i,t} + \beta_7 Xvar_{i,t} * Value_{i,t} + \beta_8 Xvar_{i,t} * Growth_{i,t} \\
& + \beta_9 EAD_{i,t} * Value_{i,t} * Xvar_{i,t} \\
& + \beta_{10} EAD_{i,t} * Growth_{i,t} * Xvar_{i,t} + \epsilon_{i,t}
\end{aligned} \tag{II.4}$$

In such a specification, β_9 and β_{10} capture by how much value and growth stock returns (the dependent variable) increase for a one-point increase in $Xvar$ on EAD. If $Xvar$ is a continuous variable, then those coefficients measure the dynamic sensitivity of value and growth stocks on EAD to this variable. This is useful to capture dynamic increases in exposure to risk factors, as in Engelberg et al. (2018). If β_4 and β_5 drop to zero when including β_9 and β_{10} , we can say that changes in sensitivity to $Xvar$ cause the EAD-value premium.

If $Xvar$ is a dummy, β_9 and β_{10} inform us how much higher returns are on EAD for value and growth stocks, compared to other stocks with earnings announcement, all on days when $Xvar$ equals one. In this case, if β_4 and β_5 drop to zero when including β_9 and β_{10} , it implies that the EAD-value premium occurs on observations where the $Xvar$ dummy is equal to one.

We implement those specifications in Table II, with *tone* replacing $Xvar$. We find that stock prices move in the same direction as news polarity. A one standard deviation increase in firm news *tone* is associated with a 22 basis points (bps) increase in stock returns outside earnings announcement days. On EAD, this relationship is about ten times stronger, showing that media content correlates more strongly with changes in investor expectations when earnings are released.

Specification (1) in Table II documents that the return spread between value and growth stocks increases on EAD by about 33.8 basis points, which is similar in

magnitude to the premium documented by Engelberg et al. (2018). Interestingly, we find a significant difference in sensitivity to *tone* between value and growth firms on EAD. By summing all relevant coefficients in specification (4)¹⁸, we find that a one standard deviation change in *tone* on EAD implies a 296bps difference in return for growth stocks, but only a 146bps change for value stocks.

Table II - Baseline EAD tone sensitivity. In this table, stock returns are multiplied by 100 and are the dependent variable. We capture book-to-market rankings with the Value (highest BM quintile) and Growth (lowest BM quintile) dummies. *tone* is the news content variable as defined above. EAD is a dummy for observations falling on earnings announcement days. We use day fixed-effects and cluster errors by date and firm.

	ret			
	(1)	(2)	(3)	(4)
EAD	0.377*** (15.01)	0.413*** (16.44)	0.629*** (23.22)	0.619*** (22.91)
Growth	0.004 (0.69)	0.003 (0.56)	0.003 (0.56)	0.003 (0.58)
Value	0.042*** (7.85)	0.045*** (8.36)	0.044*** (8.23)	0.044*** (8.26)
EAD & Growth	-0.142*** (-2.77)	-0.138*** (-2.70)	-0.110** (-2.16)	-0.010 (-0.19)
EAD & Value	0.196*** (3.79)	0.200*** (3.87)	0.241*** (4.56)	0.152*** (2.84)
tone		0.301*** (46.70)	0.224*** (36.39)	0.228*** (37.96)
tone & Value			-0.022* (-1.82)	0.006 (0.51)
tone & Growth			0.046*** (4.71)	0.022** (2.52)
tone & EAD			1.977*** (36.90)	1.884*** (30.33)
tone & EAD & Growth				0.827*** (7.43)
tone & EAD & Value				-0.661*** (-5.82)
date-FE	Yes	Yes	Yes	Yes
N	11,086,120	11,086,120	11,086,120	11,086,120
Adjusted R^2	0.16	0.16	0.16	0.16

Growth stocks appear therefore to be twice as sensitive to news on EAD. On this single day, going from a piece of average news (*tone*=0) to a news story that is one standard deviation more positive than the mean (*tone*=1) implies a 1.5% difference

¹⁸We sum all coefficient that concern a one-point increase in *tone* on EAD for value stocks, i.e., $tone, tone\&Value, tone\&EAD$ and $tone\&EAD\&Value = 0.228+0.006+1.884+(-0.661) = 0.146$.

in return between a value and a growth stock, on average. Furthermore, compared to specification (1), we observe in (4) that the EAD premium of value stocks decreases. The spread goes down from 33.8 to 15 basis points but remains significant.


Figure II.1a illustrates how value and growth stocks react to different levels of *tone* polarity on EAD. It is based on a modified version of equation (II.4). Instead of a single *Xvar*, we use multiple dummies, that equal one if *tone* is between -2 and -1.75, between -1.75 and -1.5, and so on, by increments of 0.25 in *tone* for each interval. To ensure sufficient observations for each dummy, the extremes (below -2 on the negative side and +1.5 on the positive side¹⁹) are grouped together. All coefficients are estimated in a single regression; therefore we exclude observations with “neutral” *tone*, between -0.25 and 0, which serve as the reference point. Since we focus on the EAD sample exclusively, we can drop the EAD dummies, such that the tested specification looks as follows:

$$\begin{aligned}
r_{i,t} = & \alpha_t + \beta_1 Growth_{i,t} + \beta_2 Value_{i,t} \\
& + \beta_3 tone_{i,t}^{[min,-1.75]} + \beta_4 tone_{i,t} * Value_{i,t} \\
& + \beta_5 tone_{i,t}^{[min,-1.75]} * Growth_{i,t} \\
& + \beta_6 tone_{i,t}^{[-1.75,-1.5]} + \beta_7 tone_{i,t} * Value_{i,t} \\
& + \beta_8 tone_{i,t}^{[-1.75,-1.5]} * Growth_{i,t} \\
& + \dots \\
& + \beta_n tone_{i,t}^{[1.5,max]} + \beta_{n+1} tone_{i,t} * Value_{i,t} \\
& + \beta_{n+2} tone_{i,t}^{[1.5,max]} * Growth_{i,t} + \epsilon_{i,t}
\end{aligned} \tag{II.5}$$

For the different *tone* intervals for value, we report the sum of coefficients $\beta_2 + \beta_3 + \beta_4$, of $\beta_2 + \beta_6 + \beta_7$, and so on up to $\beta_2 + \beta_n + \beta_{n+1}$. The sum of those coefficients provides the difference in expected return of value stocks for a given level of *tone* compared to the reference observations with “neutral” news. Since we sum three coefficients, we report confidence intervals computed based on the sum of the standard errors, although this is likely a conservative take. It may lead to overly big confidence intervals since it assumes independence between the coefficients. We apply the same approach for growth stocks.

Figure II.1a shows how returns evolve as a function of *tone* on EAD, and Figure II.1b does it for non-EA days. The most considerable difference between the two plots is the distinct asymmetry between value and growth stocks to bad news.

¹⁹This is because of the asymmetry in *tone* on EAD as highlighted in Table I. *tone* has a strong negative skewness. As a result, the 1st percentile of observations with *tone* is at -2.35, while the 99th percentile is at 1.68. This leads us to choose the asymmetric thresholds in Figure II.1.



(a) Return on EAD as a function of *tone* (b) Return on non-EAD as a function of *tone*

Figure II.1. Return as a function of *tone* on EAD and non-EAD. Plots the difference in expected return between growth and value stocks as a function of *tone*. Point estimates are computed as in equation (II.5). Error bars represent 95% confidence intervals. Standard errors are clustered by date and firm and are summed across coefficients to ensure a conservative approach. The green bars report the difference in returns between value and growth firms. Each interval is 0.25 in length. The central observations between -0.25 and 0 are omitted and serve as the reference point for the other coefficients. The extremes (below -2 and above 1.5) are grouped together to ensure bin sizes include sufficient observations. An additional space marks them on the plot. Figure II.1a estimates the coefficients for the EAD sub-sample, and Figure II.1b for non-EAD observations.

As *tone* becomes more negative, returns of growth stocks get even lower, which is highlighted by the bars capturing the gap between value and growth stocks. While growth stocks also appear to react slightly more to good news, the difference with value stocks is much lower. Therefore, the increased sensitivity of growth firms to *tone* on EAD appears to mainly be driven by their strong price reaction to bad news.

Conversely, in Figure II.1b we observe little to no difference between value and growth outside EAD. The price movements for equivalent levels of *tone* are also much weaker than in Figure II.1a, as expected by the large interaction coefficient “*tone&EAD*” in Table II. Based on those results, we conclude that the asymmetrically large sensitivity to bad news of growth stocks is concentrated on EAD. Value stocks appear to be more resilient to bad media coverage during the earnings period.

From Table II, we had seen that controlling for the specific sensitivity to *tone* on EAD reduces the magnitude of the EAD-value premium. Motivated by the bad news sensitivity documented above, Table III investigates if the premium is concentrated on bad news days. We create an interval of one standard deviation around the mean and define *Bad* news as observations where *tone* falls below this interval and *Good* news where they fall above. We employ the same specification as in equation (II.4), except that we replace *Xvar* with a dummy capturing *Good* and *Bad* news

as follows:

$$\begin{aligned}
r_{i,t} = & \alpha_t + \beta_1 EAD_{i,t} + \beta_2 Growth_{i,t} + \beta_3 Value_{i,t} \\
& \beta_4 EAD_{i,t} * Value_{i,t} + \beta_5 EAD_{i,t} * Growth_{i,t} \\
& \beta_6 Bad_{i,t} + \beta_7 Bad_{i,t} * Value_{i,t} + \beta_8 Bad_{i,t} * Growth_{i,t} \\
& \beta_9 EAD_{i,t} * Value_{i,t} * Bad_{i,t} \\
& + \beta_{10} EAD_{i,t} * Growth_{i,t} * Bad_{i,t} + \epsilon_{i,t}
\end{aligned} \tag{II.6}$$

where we would expect β_{10} to be negative (growth stock returns drop more on bad earnings news) and β_9 to be positive (the price of value firms is more resilient to bad EAD news). If the premium is concentrated on *Bad* days, then β_4 and β_5 should drop towards zero. The third column of Table III is consistent with the idea that the EAD-value premium vanishes when controlling for the sensitivity to bad news on EAD. It even slightly reverses, as it goes from the initial 33.8bps to -7.8bps. Moreover, on *Bad* EAD, growth stock returns are comparatively worse than for value. On such days, the former have returns that are 350bps lower. At the same time, value stocks experience a price decrease of only 163bps²⁰, which is consistent with the greater sensitivity to bad news of growth firms documented in Figure II.1a. Since this is coupled with the disappearance of the EAD value premium, we conclude that the relative return increase of value compared to growth stocks occurs on *Bad* news days.

This is further confirmed by the last column of Table III, where we control for the effect of *Good* EAD news. Here again, we see that growth stocks are more sensitive to *tone*, as good news on EAD imply a more significant price increase for those stocks. However, the spread between value and growth is lower than for *Bad* news in column (3), reinforcing the notion of asymmetry documented above. More importantly, controlling for the return difference on *Good* EAD does not capture the EAD value premium, which becomes even more substantial (53bps here vs. 33.8bps in the baseline). This shows that the EAD value premium does not occur on good news days and appears because of the greater resilience to bad news of value stocks on EAD.

B. Systematic Risk

The second part of this paper investigates whether it is possible to trace back to the cause of the premium of value stocks on bad EAD news. First, we investigate if our results can be reconciled within a framework for dynamic systematic risk. As

²⁰This is equal to the sum of all the coefficients including the *Xvar* dummy for *Bad* news.

Table III - The EAD-value premium is captured on *Bad* EAD news. This table reports estimates for equation (II.4) where *XVar* is either the news dummy *Bad* or *Good*. Returns are multiplied by 100, errors are clustered by date and firm and we include date fixed-effects.

Xvar =	ret				
	(0)	(<i>bad</i>)	(<i>bad</i>)	(<i>good</i>)	(<i>good</i>)
	0.377*** (15.26)	0.809*** (27.79)	0.792*** (26.71)	0.179*** (7.10)	0.188*** (7.32)
Growth	0.004 (0.68)	0.012** (2.37)	0.010* (1.89)	-0.003 (-0.63)	-0.002 (-0.41)
Value	0.042*** (7.91)	0.038*** (7.04)	0.040*** (7.41)	0.046*** (8.57)	0.045*** (8.45)
EAD & Growth	-0.142*** (-2.77)	-0.066 (-1.29)	0.122** (2.10)	-0.173*** (-3.40)	-0.246*** (-4.63)
EAD & Value	0.196*** (3.86)	0.181*** (3.41)	0.044 (0.77)	0.245*** (4.73)	0.286*** (5.35)
Xvar		-0.274*** (-21.19)	-0.280*** (-22.26)	0.280*** (33.32)	0.283*** (34.16)
EAD & Xvar		-2.312*** (-33.11)	-2.205*** (-26.39)	2.243*** (29.03)	2.136*** (22.89)
Value & Xvar		0.062*** (2.92)	0.009 (0.47)	-0.018 (-0.89)	0.003 (0.17)
Growth & Xvar		-0.103*** (-4.97)	-0.051*** (-2.78)	0.041*** (2.90)	0.021* (1.65)
EAD & Growth & Xvar			-1.013*** (-6.56)		0.773*** (4.71)
EAD & Value & Xvar			0.856*** (5.54)		-0.699*** (-3.56)
date-FE	Yes	Yes	Yes	Yes	Yes
<i>N</i>	11,086,120	11,086,120	11,086,120	11,086,120	11,086,120
Adjusted <i>R</i> ²	0.16	0.16	0.16	0.16	0.16

discussed in the introduction, one possible explanation for the elevated premium of value stocks on EAD, is that those firm's beta with some priced risk factor increases dynamically during earnings announcements. We can follow the framework of equation (II.4) as in Engelberg et al. (2018) to capture the EAD value premium dynamically. In Table IV we test different candidates as underlying risk factors responsible for the EAD-value premium. First, we control for dynamic exposures to the market return and the high-minus-low (HML) book-to-market portfolio return factor as constructed by Fama and French (1993). While we find that value (growth) stocks become more (less) sensitive to the HML factor on EAD, the EAD-value premium remains virtually unchanged (the spread decreases from 33.8bps to 32.2bps) and stays highly significant. For the market return, the impact is even weaker, as

we observe no dynamic increase in sensitivity specific to value or growth stocks at all. Those results confirm the findings of Engelberg et al. (2018).

Table IV - Systematic Risk Factors. This table reports estimates for equation (II.4) where $XVar$ represents different candidates for systematic risk factors discussed in section III.B. “Mkt” is the excess market return and “HML” the high-minus-low book-to-market portfolio return, both obtained from Kenneth French’s website. “Ivol Mkt” and “Ivol HML” are the aggregate idiosyncratic volatilities of market and of high-minus-low book-to-market stocks. “*tone* Mkt” and “*tone* HML” are the aggregate news *tone* of the market and of the HML portfolio. “EADret” is the value-weighted return of all announcing firms on a given day. Otherwise, regression specifications are identical to Table II (4). $XVar$ only appears in interaction, because all factors only vary in the time dimension and thus are collinear with the day-FE dummies.

Xvar =	ret							
	(0)	(Mkt)	(HML)	(Ivol Mkt)	(Ivol HML)	(<i>tone</i> Mkt)	(<i>tone</i> HML)	(EADret)
EAD	0.377*** (15.26)	0.375*** (15.02)	0.376*** (15.01)	0.230*** (4.46)	0.353*** (14.15)	0.368*** (8.30)	0.353*** (14.15)	0.164*** (6.95)
Growth	0.004 (0.68)	0.001 (0.10)	0.005 (1.23)	0.018 (1.22)	0.004 (0.69)	0.031*** (3.39)	0.004 (0.69)	-0.003 (-0.59)
Value	0.042*** (7.91)	0.049*** (10.93)	0.041*** (8.03)	0.045*** (2.66)	0.035*** (6.67)	0.014 (1.51)	0.035*** (6.67)	0.053*** (9.54)
EAD & Growth	-0.142*** (-2.77)	-0.141*** (-2.76)	-0.133*** (-2.65)	-0.325*** (-3.44)	-0.154*** (-2.97)	-0.158 (-1.64)	-0.154*** (-2.97)	-0.176*** (-3.45)
EAD & Value	0.196*** (3.86)	0.196*** (3.77)	0.189*** (3.70)	0.294*** (2.96)	0.182*** (3.60)	0.271*** (2.98)	0.182*** (3.60)	0.127*** (2.52)
EAD & Xvar		6.006** (2.38)	6.509 (1.08)	0.053*** (2.96)	0.040*** (3.99)	-0.044 (-0.23)	0.040*** (3.99)	46.635*** (22.59)
Growth & Xvar		7.332*** (5.56)	-28.897*** (-12.61)	-0.007 (-0.94)	-0.000 (-0.09)	0.131*** (3.13)	-0.000 (-0.09)	1.196*** (3.65)
Value & Xvar		-16.726*** (-9.67)	15.149*** (5.09)	-0.002 (-0.19)	0.011** (2.52)	-0.130*** (-3.02)	0.011** (2.52)	-2.595*** (-6.24)
EAD & Growth & Xvar		-1.909 (-0.38)	-42.256*** (-4.17)	0.068** (2.13)	0.025 (1.26)	-0.075 (-0.17)	0.025 (1.26)	9.045** (2.50)
EAD & Value & Xvar		2.252 (0.50)	37.429*** (3.17)	-0.034 (-0.99)	0.033 (1.41)	0.359 (0.96)	0.033 (1.41)	16.474** (2.27)
date-FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	11,086,120	11,086,120	11,086,120	11,086,120	11,086,120	11,086,120	11,086,120	10,925,108
Adjusted R^2	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16

Next, we test for additional potential risk factors. First, we construct an aggregate idiosyncratic volatility factor based on Ang et al. (2009), which we call $Ivol^{21}$. Since we found that the EAD value-growth spread varies as a function of firm-specific (i.e., idiosyncratic) news content, we investigate if this difference can be linked to different sensitivities to aggregate idiosyncratic risk shocks. While growth stocks do indeed exhibit a slightly greater sensitivity to aggregate $Ivol$ on EAD, the dynamic premium remains, as the value-growth spread on EAD is equal to 62 basis points.

Another potential source of systematic risk that we consider is the aggregate media *tone*. We compute a value-weighted aggregate *tone* measure for the market,

²¹We compute it as the value-weighted squared residual from a Fama-French three-factor model that includes the excess market return, the value-minus-growth factor and the small-minus-big factor.

as well as the aggregate value-minus-growth *tone*. Neither value nor growth stocks are more sensitive to either aggregate *tone* measure on EAD. Furthermore, the EAD-value premium is virtually unchanged. Thus, we reject aggregate *tone* as a potential candidate for a priced dynamic risk factor on EAD.

Finally, the last candidate for a systematic risk factor is the aggregate return of announcing firms, *EADret*. For this factor, we compute the average returns of all firms that announce their earnings each day. The hypothesis goes similar to the argument of Savor and Wilson (2016). Announcing firms earn a premium because their beta with the market spikes on EAD. Thus, announcing firms that tend to co-move more together on EAD could be more highly exposed to a shared source of risk. Therefore, the goal of the *EADret* factor is to capture co-movement between announcing firms. We find that both value and growth firms co-move more with other announcing firms on EAD. However, the EAD-value premium again remains unchanged. Therefore, this test also fails to support the existence of an underlying dynamic risk premium.

Table V - Principal Component Analysis. This table investigates whether any of the first 700 first common dimensions in stock returns dynamically capture the EAD-value premium. We run a single regression as in Table IV above for each dimension individually. To summarize the results, we group dimensions by fifty and report the minimum V-G EAD spread (i.e., $\beta_4 - \beta_5$ from equation (II.4) among the fifty dimensions in the second column. The bottom row reports the minimum across all 700 dimensions. We also report the maximum spread, the maximum and minimum t-stats of coefficients β_5 and β_4 respectively, and the sum of all t-stats for β_9 and β_{10} (the triple interaction coefficients) that are significant, i.e. whose absolute value is greater than two.

Dimensions	Min V-G EAD spread	Max V-G EAD spread	Max t-stat EAD&Growth	Min t-stat EAD&Value	# $ t >2$ EAD&Value&XVar	# $ t >2$ EAD&Growth&XVar
#1-50	0.33	0.35	-2.62	3.76	2	4
#51-100	0.33	0.35	-2.69	3.77	4	2
#101-150	0.33	0.35	-2.68	3.81	1	4
#151-200	0.33	0.34	-2.71	3.78	3	4
#201-250	0.34	0.34	-2.73	3.79	3	0
#251-300	0.33	0.34	-2.73	3.81	3	1
#301-350	0.33	0.34	-2.72	3.82	0	4
#351-400	0.33	0.34	-2.64	3.81	4	4
#401-450	0.33	0.34	-2.62	3.82	2	5
#451-500	0.34	0.34	-2.71	3.83	1	1
#501-550	0.34	0.34	-2.74	3.79	4	5
#551-600	0.33	0.34	-2.69	3.79	4	3
#601-650	0.33	0.34	-2.70	3.81	2	1
#651-700	0.33	0.34	-2.72	3.8	7	4
Total	0.33	0.35	-2.62	3.76	33	38

Since the set of systematic risk factors used above fails to capture the dynamic EAD premium of value stocks, we test if there is any other common variation in stock returns that can dynamically capture the premium. We compute the 700 first

principal components of the cross-section of stock returns using a PCA to achieve this. We choose this number of dimensions as this captures 98.65% of all the variation in stock returns²². For each dimension, we run equation (II.4) as before in Table IV and report the magnitude of the remaining EAD-value premium captured by coefficients β_4 and β_5 .

Table V summarizes the results. We find that none of the first 700 principal components of returns effectively capture the EAD-value premium. The lowest value-growth spread that we observe is still 33.1 bps and remains highly significant. Therefore, we fail again to find a shared source of systematic risk that captures the dynamic premium of value firms, even from a large set of statistically generated common components in stock returns.

Given the failure of systematic risk factors to capture the premium earned by value firms on EAD, we explore alternative frameworks which can coincide with our observations and that could give rise to the premium.

C. Biased expectations hypothesis

The biased expectations hypothesis predicts that investors form expectations prior to earnings announcements that are systematically biased as a function of a measure varying in the cross-section (in our case, book-to-market). Therefore, a necessary condition for this hypothesis to hold is that relative to value stocks, on EAD, growth stocks underreact to good news, that they overreact to bad news, or both.

This framework might be traced back to the work of Basu (1977), but was popularized with the large literature in behavioral finance. Barberis et al. (1998) and Daniel et al. (1998) are prominent examples of theories that attribute momentum and reversal patterns to systematic biases in expectations. La Porta et al. (1997) already found that value stocks earn higher returns during earnings announcements. They argue that investors form overly optimistic expectations about glamour (i.e., growth) firms and vice versa for value stocks. Those expectational errors can be the result of the extrapolation of past performance, as suggested by Lakonishok et al. (1994). As a result, value (growth) stocks get abnormally positive (negative) earnings surprises on EAD once information about future cash flows becomes publicly available and investors are forced to revise their biased expectations.

Engelberg et al. (2018) show that analyst expectations about long legs of anomalies have a positive bias. Given their failure to identify a dynamic systematic risk

²²See appendix C for the detailed contribution to the variance of each dimension in a scree plot.

premium, they conclude that biased expectations are the most likely candidate to explain the EAD-anomaly premium. The empirical observations we document above section can also fall in line with this explanation.

In particular, we may be able to provide additional nuance to the biased expectations argument: loosely speaking, investors do not appear to be overly pessimistic about value stocks (otherwise we would expect a strong positive reaction on *Good* EAD news for value firms), but rather, they seem overly optimistic about growth firms. Indeed, the large asymmetric response to *Bad* news on EAD suggests that negative news trigger more significant revisions in expectations for growth stocks, consistent with the idea that investors form overly optimistic expectations about the latter.

Those results also closely relate to Skinner and Sloan (2002) who also document an asymmetric reaction of growth stocks on EAD to negative earnings surprises. However, our findings are able to bring additional insights by contrasting the asymmetric reaction on EAD, with the symmetrical reaction on non-EAD. This might suggest that expectations might tend to be systematically biased about earnings news, but not so much for other types of information that might affect the stock price of a company.

Moreover, our news *tone* measure might provide additional insights on why those biased expectations might form in the first place. In table IX in section II we find that growth stocks get significantly higher media coverage. More importantly, this media coverage has a significant negative (positive) bias for value (growth) stocks. This positive tilt for growth stocks is strongest outside EAD. This consistently higher and more favorable media coverage might end up shifting expectations that investors form about a firm's earnings.

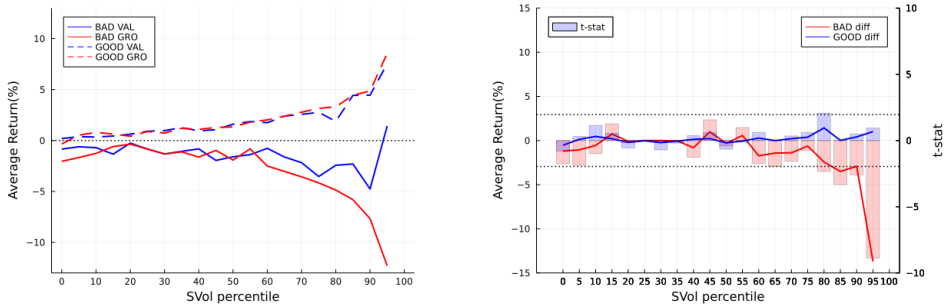
D. Investor attention hypothesis

Alternatively, the results we observe might be consistent with the investor attention hypothesis of Barber and Odean (2008). The argument relies on the idea that individual investors are “net buyers of attention-grabbing stocks”. When searching for a stock to buy, individual investors face the difficulty of choosing among a large pool of potential candidates. In a complex world, firms that grab their attention are more likely to be chosen. Intuitively, growth stocks fall into this category, given their extensive media coverage documented in the data section II. However, when selling following negative signals, investors do not face the same limitations and sell the stocks they already own. Therefore, this framework predicts a strong reaction to

bad news for attention-grabbing stocks and a weak PEAD (Hou et al., 2009), both features that we observe in our data.

Lamont and Frazzini (2007) also argue that the EAD premium is linked to investor attention. They contend that earnings announcements can attract investor awareness to stocks that usually get fewer media coverage and attention. Their results suggest the existence of a link between individual investors' buying behavior and the EAD premium; the latter getting stronger as trading volume increases.

Therefore, we investigate how value and growth stock returns react to *Good* and *Bad* news on EAD as a function of abnormal trading volume (*AVol*). *AVol* of a firm at time t is defined as the ratio of day t trading volume divided by the trailing average of 365 days, as in Lamont and Frazzini (2007). One advantage of this measure is that it is *de facto* scaled around 1 for all stocks (as long as firms do not have run-away increases in trading volume). In Table I we document that growth stocks have more significant increases in *AVol* around earnings announcements, already suggesting that those firms are subject to more substantial increases in trading activity.



(a) Mean returns as function of AVol percentile. (b) Difference as a function of AVol percentile.

Figure II.2. Return as a function of AVol for Bad and Good news. This figure plots the difference in average return between growth and value stocks as a function of abnormal volume (*AVol*), on EAD. *AVol* intervals are divided in 20 buckets. In plot II.2a, the dashed lines capture the mean returns for *Good*, and the lower solid lines average returns for *Bad* news. Value stocks appear in blue and growth stocks in red. Figure II.2b reports the difference in mean returns for a given level of *AVol* between value and growth stocks for *Bad* (red) and *Good* (blue) news. The red bars refer to the t-stat of the difference between value and growth firms for BAD news, and the blue bars to the difference for GOOD news. The black dashed line indicates a t-stat level of ± 1.96 , as indicated on the right axis.

In Figure II.2 we plot the average returns of value and growth stocks as a function of *AVol*. We find that as trading volume rises, the impact of *tone* increases, i.e., good news imply higher returns and bad news lower returns. However, we observe once

again an asymmetry to *Good* and *Bad* news. As abnormal trading volume gets very high on EAD, growth stocks get very negative returns on average. Value stocks also tend to get more negative returns as trading activity increases, but bad news have much less of a negative impact on prices for very high levels of volume. For *Good* news, however, value and growth firms follow very similar patterns: returns increase with the magnitude of trading activity and are not significantly different. Loosely speaking, such a result would be predicted by the hypothesis of Barber and Odean (2008), where investors sell high-attention stocks they already own (i.e., growth stocks). It would, however, be difficult to reconcile with an increased propensity to buy value stocks with good news on EAD, as proposed by Lamont and Frazzini (2007).

As a robustness consideration, throughout the paper, we employ news *tone* as an implicit measure for revisions in investor expectations. Therefore, *tone* should capture the surprise, i.e., the unexpected news. We find that firm *tone* is a persistent variable that can be forecasted by past *tone*. We propose an alternative measure of news *tone*, which we label *surp* and aims to capture the surprise, i.e., the unexpected component of news *tone*. We replicate the above results and report them in appendix B. Those robustness tests conclude that the observations made for *tone* in the analyses above remain qualitatively the same and that *tone* itself is already a relevant measure for news content.

E. Biased Expectations vs. Investor Attention

As a final consideration, we investigate how the biased expectations hypothesis and the investor attention hypothesis stack up in an empirical test to the theoretical predictions of post-earnings announcement drift (PEAD) behavior, first documented by Ball and Brown (1968). The PEAD²³ captures the tendency of stock prices to continue drifting in the direction of the earnings surprise in the days and weeks following the announcement. This framework allows us to study how prices evolve following *Good* and *Bad* news and provide insights about potential over-or under-reaction issues.

Table VI measures the PEAD of value and growth stocks at different horizons. Instead of using analyst forecast errors (the standard approach in the literature) to proxy for earnings surprises, we rely on our measure of *tone*. We find a significant

²³Several potential explanations have been put forward to explain this phenomenon. These include information uncertainty and gradual learning (Brav and Heaton, 2002; Francis et al., 2007), transaction costs (Ng et al., 2008), idiosyncratic risk, (Mendenhall, 2004) and investor attention (Hou et al., 2009).

positive PEAD following good news at the 10, 30, and 90-day horizon. This is in line with previous literature documenting a continuation of returns in the direction of the earnings surprise. We also find a negative continuation of bad news, but only at the 10-day horizon. We find no specific PEAD for either value or growth stocks following good news. However, growth stocks with bad *tone* on EAD tend to get higher returns in subsequent periods.

Table VI - PEAD following *tone* of value and growth. This table reports the magnitude of the post-earnings announcement drift to *Bad* and *Good* news tone. $PEADret_{[1;x]}$ is the cumulative return of an announcing firm starting in day $t + 1$ after the announcement up until day $t + x$. We focus on the EAD sub-sample exclusively and thus drop the EAD dummy. Otherwise, the regression specification is identical to Table II.

	$PEADret_{[1:10]}$		$PEADret_{[1:30]}$		$PEADret_{[1:90]}$	
	(1)	(2)	(3)	(4)	(5)	(6)
Growth	-0.131** (-2.09)	-0.059 (-0.92)	0.050 (0.48)	0.158 (1.51)	0.037 (0.17)	0.239 (1.07)
bad	-0.184** (-2.03)		-0.039 (-0.25)		-0.074 (-0.24)	
Value	0.271*** (3.29)	0.243*** (3.03)	0.464*** (3.47)	0.429*** (3.25)	0.955*** (3.63)	1.179*** (4.28)
Growth & bad	0.459** (2.58)		0.738** (2.50)		1.516** (2.52)	
Value & bad	-0.144 (-0.66)		-0.051 (-0.14)		0.900 (1.31)	
good		0.329** (2.13)		0.475** (2.04)		1.179*** (2.84)
Growth & good		-0.107 (-0.38)		0.031 (0.07)		0.517 (0.55)
Value & good		-0.150 (-0.45)		0.542 (1.08)		-1.593 (-1.56)
date-FE	Yes	Yes	Yes	Yes	Yes	Yes
N	90,776	90,776	90,202	90,202	86,802	86,802
Adjusted R^2	0.25	0.25	0.24	0.24	0.24	0.24

Those results again hint toward an asymmetry to bad news. Growth firms experience a relatively stronger reversal following bad news *tone* on EAD. However, the PEAD of value stocks is not stronger following either *Good* or *Bad* news. This result appears to be running contrary to the biased investor expectations hypothesis. Indeed, if the strong reaction to *bad* news of growth stocks was to adjust for expectational errors, then we would not expect to observe a relatively increased reversal post-EAD specifically following those same *bad* news.

On the other hand, those patterns are consistent with predictions of the investor attention model. Indeed, if the larger price decline to bad EAD news of growth

stocks reflects downward sale pressure from attention-constrained investors, then we would expect an abnormally large negative return on EAD, followed by a reversal and adjustment compensating for the initial downward pressure.

IV. Discussion and Conclusion

Value stocks earn a return premium over growth stocks on news days when public information is released. This effect is particularly strong on earnings announcement days, as documented by Engelberg et al. (2018). By using an exogenous measure of media *tone* to proxy for the content of firm-specific news, we find that the EAD-value premium is concentrated on *Bad* news days. Indeed, a dummy controlling for bad EAD news captures the entirety of the dynamic premium. This occurs because growth stocks are more sensitive than value firms to news *tone* on EAD, and appear significantly more so when news sentiment is negative. There is a distinct asymmetry between value and growth firms: On average, value and growth firms experience similar price increases for an equivalent *Good* news on EAD; however, growth stock prices will decline significantly more for an equivalent *Bad* news.

Despite being concentrated on the subset of bad news days, the EAD-value premium could still result from dynamic increases in risk exposure, if this risk increases on bad EAD specifically. This risk exposure hypothesis posits that value firms will load more on a systematic risk factor during EAD. Controlling for this increase in risk exposure should capture the EAD-value premium. Therefore, the predictions of this model can directly be tested with candidate risk factors. Our results reject a large set of systematic risk factors as candidates for the dynamic risk premium of value stocks, including the aggregate market return and *tone*, high-minus-low book-to-market return and *tone*, aggregate idiosyncratic risk, and aggregate announcing firm return. We also reject the first seven-hundred dimensions of common variations in returns, computed using a PCA, as relevant dynamic risk factors. Therefore, we do not find supporting evidence for the systematic risk hypothesis.

Rejecting systematic risk as the cause for the premium earned by value firms on EAD has the following implication: other candidate explanations need to step in to account for the phenomenon. We discuss multiple potential alternative frameworks, motivated by prior research.

A credible candidate explanation relies on biased investor expectations. Like previous studies, including Engelberg et al. (2018) and Skinner and Sloan (2002), the asymmetry to *Bad* news that we document fits naturally within this framework.

It suggests that investors are overly optimistic about growth stocks, such that bad news on EAD cause a big surprise and a large negative revision in expectations. Looking at it from a different perspective, it can also be interpreted as investors overly anticipating bad news for value stocks, such that the latter are more resilient to negative announcements.

Directly testing this hypothesis is difficult, as the underlying quantity of interest (i.e., aggregate investor expectations about firms, and changes thereof over time) is complex to measure and represents a limitation to our study. The closest we can get is to mimic the approach Engelberg et al. (2018), who found that analysts tend to be overly optimistic (pessimistic) about forecasts of firms in long (short) legs of anomaly portfolios. Investigating whether those surprises to analyst expectations are systematically more biased during bad news, EAD, might be a future research avenue that could further strengthen the case of this hypothesis.

Our results may also concur with explanations relying on investor attention, such as in Barber and Odean (2008), where individual investors tend to overly sell high-attention glamour (i.e., growth) stocks they already own. Again, we do not dispose of the entire detailed order book of the market, and hence cannot determine whether the relative out-(under-)performance of value (growth) firms on bad EAD is driven by holdings and selling activity of individual retail investors. One possible future research avenue is to leverage data containing individual investor trading activity. Kaniel et al. (2012) for example, use trading data from the NYSE's Consolidated Equity Audit Trail Data (CAUD) files to proxy for individual investors' trading activity.

Finally, our results also call for future research extending beyond the value anomaly. While a story relying on investor attention fits nicely with the differences in news coverage and the "glamour" factor of growth and value firms, it remains to be seen if the EAD premium of all anomalies, as in Engelberg et al. (2018) occurs on bad news days. Moreover, despite the lack of evidence in favor of a dynamic risk explanation, it may be possible that to capture the changes in exposures to aggregate risk, high-frequency data is required. An intraday signal, similar to the one used by Patton and Verardo (2012), might serve as a better proxy

Chapter III

Conditional Drivers of Anomaly Comovement

MORENO NICOLAS[‡]

APRIL 2022

ABSTRACT

This paper provides evidence that anomaly-specific commonalities in cash-flow (CF) news can vary as a function of time-varying proxies for market CF and discount-rate (DR) news. Our framework allows us to capture cross-sectional differences highlighting separate comovement sensitivities of stocks in long and short legs of anomaly portfolios. Reconciling multiple theories explaining anomaly returns, we find that value, small, and loser stocks exhibit stronger CF comovement in periods of market CF stress and recessions. Conversely, CF news covariance between growth and winner stocks increases in periods of market DR stress, for example, when investor sentiment is high. Our framework can generate those novel insights by separating anomaly-specific CF and DR variance elements from other CF and DR variance components common to any well-diversified long-minus-short portfolio.

[‡]This paper appears as a single-authored chapter for my PhD dissertation at HEC Liège. It would however never have existed without the tremendous help of Marie Lambert and Denada Ibrushi. Both their guidance and contribution have been essential, and therefore I am hopeful that they may rejoin as co-authors, should this paper proceed in the publication process. At this stage, and as a chapter of my thesis, this paper is not intended for circulation.

I. Introduction

The asset pricing literature faced a dramatic change in paradigm when it gradually came to the realization that variations in expected returns are responsible for the majority of comovement in stock returns. The literature now traces most of the variation in returns of aggregate portfolios of stocks back to common discount-rate (DR) news (Campbell and Ammer, 1993; Cohen et al., 2003; Savor and Wilson, 2016). At the firm level, updates in cash-flow (CF) expectations are the primary driver of unexpected returns (Vuolteenaho, 2002). However, the latter diversify away at the market level due to weaker CF shock correlations in well-diversified aggregate portfolios (Chen et al., 2013).

In parallel, a significant portion of the research in asset pricing has dedicated its focus to understanding the mechanisms driving return predictability in the cross-section (Cochrane, 2011). Researchers have documented numerous so-called market anomalies over the past decades: along specific characteristics, stocks ranking high in certain dimensions consistently outperform those with a low ranking. Therefore, it has become common practice to proxy the payoffs of anomaly exposures by building well-diversified portfolios that go long stocks with a high ranking for a given anomaly characteristic and short those with the lowest ranking. Such long-minus-short (LMS) portfolios can capture the premium of anomalies with a near market-neutral exposure.

However, to understand why persistent anomalous deviations from classical asset models subsist, it is necessary to understand the source of comovement induced by anomaly characteristics. Some theories predict that anomaly returns are mainly driven by shared DR shocks, others by common CF shocks. Still, other theories rely on mechanisms predicting conditional increases in anomaly characteristics' exposure to either CF or DR risk factors. To help researchers navigate the extensive literature seeking to provide explanations for anomaly premia, Lochstoer and Tetlock (2020) propose to empirically estimate the yearly CF and DR variations of well-diversified LMS portfolios. Their main conclusion is that CF news are the main driver of LMS anomaly portfolios and that neither anomaly CF nor DR news correlate with either the business cycle market CF or market DR news.

In this paper, we can shed new light on the drivers of anomaly returns, by making two principal contributions. First, we show that the dominance of DR shock variation at the market level plays a crucial role in investigating LMS anomaly portfolios' drivers. Actually, we show through simulations that any well-diversified and randomly generated LMS portfolio has CF shocks that account for a significantly

greater portion of its variance than DR shocks. By proposing a detailed decomposition of the variance of any LMS portfolio, we are able to highlight the origin of this empirical observation: It is the strong correlation in DR news between any well-diversified long- and short-leg portfolio that ends up pushing the DR variance of all LMS portfolios down.

Indeed, we show that the sum of three terms makes up the CF (DR) return variance of any LMS portfolio:

- (i) Minus two times the covariance between the aggregate long and aggregated short leg portfolio CF (DR).
- (ii) The sum of the individual firm CF (DR) variance terms.
- (iii) The sum of all CF (DR) cross-covariance terms between firms in the long and short leg portfolios.

Notably, the first term (i), capturing the covariance between the long and short portfolio's aggregate CF and DR news, is of consistently more negative magnitude for DR news than for CF news. This holds true for any of the two well-diversified long and short portfolios, because of the dominance of DR news at the market level. As a result, the total DR variance of any LMS portfolio is dramatically reduced so that CF news always dominate.

This observation begs a crucial question: if CF news are the main drivers of any well-diversified LMS portfolio, how can we identify stocks' CF and DR drivers specifically linked to their anomaly characteristics? How can we measure variance drivers of anomaly portfolios that are not influenced by the dominance of DR news at the aggregate level?

Here lies the second and central contribution of the paper. Our detailed LMS portfolio variance decomposition allows us to distinguish comovement patterns within stocks sharing similar anomaly characteristics that significantly differ from comovement among other random groups of stocks. We can measure the magnitude of this comovement thanks to the third term of the decomposition (iii), which captures covariance among stocks with similar anomaly characteristics. This novel approach allows us to bring a fine-grained understanding of the drivers of anomaly returns. We highlight that CF comovement dominates, and we underscore cross-sectional differences. Importantly for theories explaining anomaly returns, we show that CF-comovement across stocks with certain shared anomaly characteristics varies as a function of aggregate CF news and the business cycle, whereas for others, it is sensitive to changes in aggregate DR shocks.

Over the years, the literature has proposed several frameworks predicting the existence of anomaly premia. We can regroup those theories into two broad types of predictions. The first type makes predictions about the primary driver of anomaly return variance.

In theories where DR news are the main drivers of anomalies, cross-sectional differences in returns arise due to commonalities in discount rate revisions linked to anomaly characteristics. For example, noise trading models where costly arbitrage causes positive feedback effects¹ rely on constant CF expectations and prices deviating from fundamentals through revisions in expected returns. Other models posit that irrational investors can cause stocks to comove more with other stocks sharing similar anomaly characteristics than what their underlying cash flows would predict². In such behavioral models of style investing, cross-sectional differences in price deviations from fundamentals arise due to implied later reversals towards fundamentals and thus variations in discount rates.

Our results cast doubt on the idea that this type of theory is the main driver behind anomaly returns. Generally, for all five anomaly portfolios that we study, DR comovement for stocks with similar anomaly characteristics does not exceed DR comovement among other random groups of stocks.

On the other hand, our empirical observations suggest that commonalities in CF news play a crucial role in generating anomaly returns. We observe that almost all stocks sharing similar anomaly characteristics have CF comovement levels that are orders of magnitude higher than randomly sampled firms. Importantly, we can disentangle comovement across firms within each long and short leg separately for all five anomalies that we study. Our results suggest that compared to randomly sampled portfolios, firms with high book-to-market (i.e., value, +204%), low market cap (i.e., small, +282%), and strong negative momentum (i.e., loser +298%), are those where commonalities in CF news are the strongest.

Several theories explaining anomaly returns rely on CF news as the primary source of variation. Usually, those models suppose a negative correlation between anomaly CF and DR news. Consider, for example, frameworks suggesting that anomaly prices deviate because of investor misvaluations. An entire class of models relies on this idea of mistaken expectations and proposes that investors overreact to

¹Examples of such models include Black (1986), De Long et al. (1990) and Shleifer and Vishny (1997).

²Barberis and Shleifer (2003) derive a theoretical model for an economy where investors follow investment styles (e.g., value, growth, momentum) as a function of their popularity. Barberis et al. (2005) and Boyer (2011) are examples of papers that provide evidence for excess comovement among stocks sharing certain characteristics.

cash flow signals³, which in turn lead to later reversals pushed by the DR channel. As Lochstoer and Tetlock (2020) put it, if anomaly characteristics are predictors of exposure to CF shocks, then those characteristics will imply stronger price reactions to CF news, which are further amplified by an opposite DR shock. Therefore, such overreaction models predict strong commonalities in CF movements within anomaly portfolios and a negative correlation between anomaly CF and DR news⁴.

Rational models can also predict persistent anomaly returns while implying negative correlations between CF and DR news. Following, Kogan and Papanikolaou (2013) for example, a positive technology shock (i.e., a CF shock) can simultaneously trigger a decrease in expected returns. If anomaly characteristics are linked to exposures to such shocks, then long-minus-short anomaly portfolios will also exhibit a negative correlation between CF and DR shocks.

A second class of models makes predictions extending beyond the relative importance of CF and DR news in anomaly variance. Indeed, some theories explaining anomaly returns assume that cross-sectional differences in risk exposures to either systematic CF or systematic DR factors can give rise to rational and sustained deviations from classical asset pricing models.

Some theories assume that anomaly premia arise due to exposures to sources of systematic risk linked to aggregate CF shocks. For example, in the case of the value anomaly, Fama and French (1996) provide a compelling argument why high book-to-market stocks (i.e., distressed companies, closer to bankruptcy) should earn a premium due to investor hedging preferences during periods of economic downturn, i.e., a stressful aggregate CF event. This kind of explanation implies a greater sensitivity of value stocks to negative market-wide CF events, such as recessions. It also relies on cross-sectional differences in CF comovement along anomaly characteristics⁵.

Berk et al. (1999) and Zhang (2005) also propose models that explain anomalies by relying on differing sensitivities to market CF movements. In Zhang's model,

³Models for anomalies based on investor overconfidence and extrapolation leading to biased expectations include Daniel et al. (1998, 2001), Barberis et al. (1998, 2015, 2018), Odean (1998), and Hong and Stein (1999). Greenwood and Shleifer (2014) provide survey evidence of investor extrapolation, and Cassella and Gulen (2018) show that levels of investor extrapolation are time-varying and affect the levels of expected returns.

⁴Empirically, this negative CF/DR correlation seems to be sensitive to the choice of empirical specification, but since our baseline parameters are based on Lochstoer and Tetlock (2020), it is a feature that we also observe across anomaly portfolios. This is a result that we do not derive from our detailed portfolio decomposition.

⁵Early papers showed mixed results on the relative importance of the long-horizon CF vs. DR variance between value and growth stocks (Lakonishok, Shleifer, and Vishny, 1994). However, subsequent empirical studies point to the existence of cross-sectional comovement in earnings and profitability (see for example Fama and French (1995) and Cohen, Polk, and Vuolteenaho (2003)).

value firms have a high level of assets in place, making it difficult to smooth the impact of negative aggregate productivity shocks (such as recessions) when they occur. Such a model predicts that value stocks will exhibit strong comovement during stressful aggregate CF events. Consistent with this idea, Campbell and Vuolteenaho (2004) show, also using an empirical decomposition of market CF and DR news, that value firms covary more with market CF shocks. Cohen, Polk, and Vuolteenaho (2009) obtain similarly higher CF betas for value stocks by proxying for CF shocks using accounting measures.

Other theories posit that anomalies might also arise because of differing sensitivities to systematic DR shocks. Lettau and Wachter (2007) argue that growth stocks are long-duration assets, which results in their cash flows being more sensitive to changes in aggregate discount rates. Consistent with this idea, Campbell, Polk, and Vuolteenaho (2010) find that CF variations mainly drive growth firms' price movements and that their common source of risk is linked to market DR shocks. According to such theories linking anomaly returns to time-varying risk aversion (Santos and Veronesi, 2010), one would expect growth stocks to experience spikes in comovement during high sentiment periods (Baker and Wurgler, 2006) or in periods of high risk aversion, for example, when the default spread is large.

Our empirical findings allow us to reconcile those different theories. We document time-varying changes in comovement among stocks with similar anomaly characteristics. Importantly, depending on the anomaly characteristic, those changes are either conditional on variations in aggregate CF news or variations in market DR shocks. To proxy for updates in aggregate CF and DR expectations, we use one-year changes in estimated market CF and DR news, as well as macro-economic indicators related to aggregate CF stress (i.e., recession dummies and 1-year GDP change) and aggregate DR stress (i.e., credit spread and investor sentiment).

A key advantage of our research design is that it enables us to measure dynamic changes in comovement separately for long and short-leg portfolios. This allows us to shed new light on the drivers specific to each characteristic of anomaly portfolios, compared to measures where aggregation in LMS portfolios might end up mixing overlapping signals and exposures to CF and DR risk. Moreover, our specification employing quarterly estimates of CF and DR news further allows us to uncover dynamic risk exposure patterns that could previously not be identified with yearly data.

For example, consider the case of the value anomaly. In hard times, such as recessions, comovement of CF news between stocks increases. This is true within

any random portfolio of stocks, but CF comovement becomes distinctively stronger for stocks with high book-to-market ratios (+337%, t-stat=7.35). Consistent with the theoretical explanations for the value anomaly proposed by Fama and French (1996) and Zhang (2005), we also observe stronger CF comovement within the long leg of the value anomaly in periods where other proxies for market CF shocks are also negative: when 1-year GDP growth and 1-year change in market CF news are in the lower third, we observe comovement among value firms increasing by 300% and 269% respectively. This increase in CF comovement is significantly greater than for periods of positive market CF news. It is also consistent with models suggesting shared risk exposures of value stocks to permanent market contractions. In contrast, we find no patterns distinct from random portfolios when conditioning value stocks' DR news to market CF news. This further cements CF news as the primary driver of anomaly returns.

Growth stocks exhibit opposite patterns. Comovement among stocks with low book-to-market ratios does not significantly spike during negative market CF periods and, on the contrary, tends to increase when market conditions are favorable. However, patterns are noticeably more apparent when conditioning on aggregate market DR news. While CF news of growth stocks do not comove more strongly in positive market DR news periods, we observe a sharp comovement increase in periods of aggregate DR stress. When investor sentiment and 1-year market DR news are the lowest or when credit spread is the highest, growth stock CF comovement significantly increases over baseline (between +75% and +192%). This fits on justifications for the low growth stock premium based on their cash-flow duration (Lettau and Wachter, 2007), and overall reconciles multiple explanations for the value anomaly that might all play together to explain its persistent premium.

Our results can also guide researchers in understanding the drivers of other anomalies. For example, small and loser firms also exhibit strong CF comovement, which seems to be specific to their anomaly characteristics. Both portfolios also show comovement patterns similar to value stocks: the CF covariance of small and loser firms is strongest in periods of recession and negative market CF news. Previous research has robustly documented that the momentum anomaly underperforms in the periods around recessions. Our observations might further strengthen the case of theories suggesting that loser stocks might earn a risk premium specifically following recessions and market-CF panic states. For example, Daniel and Moskowitz (2016) explicitly suggest that "up-market betas" are asymmetrically high for loser stocks (compared to winner stocks) following recessions, which coincides with our

documented elevated CF comovement patterns.

Our study contrasts with Lochstoer and Tetlock (2020) in that our simulations show that aggregate portfolio dynamics have a large weight in long-minus-short portfolio variance, making it necessary to take them apart to understand anomaly-driven variance in the cross-section. Therefore, our detailed decomposition allows us to shed new light on anomaly-specific comovement, thus providing new guidance for researchers seeking to navigate different theoretical explanations for anomaly returns. In particular, we are able to make distinguishing contributions having implications to understand the source of anomaly premia: First, our results are consistent with theories putting forward underlying systematic CF and DR risk factors to explain anomaly returns, thus supporting the extent literature attributing anomalies to different sources of risk. Second, we find that CF news are indeed the main drivers of anomaly comovement, but that differences exist across anomalies, suggesting that a one-size-fits-all explanation for all anomalies is unlikely to make the cut, marking another distinction with Lochstoer and Tetlock (2020).

The rest of the paper is structured as follows: section II lays out our detailed portfolio variance decomposition. Section III presents our empirical approach to estimating quarterly CF and DR news. Section IV.A shows how CF shocks mainly drive any randomly simulated LMS portfolio. Section IV.B illustrates how anomaly CF and DR comovement differs from random portfolios. Section IV.C reports the time-varying conditional CF and DR comovement decompositions. Finally, section V provides a discussion about theoretical frameworks for anomalies and how they relate to our observations, and section VI concludes.

II. Detailed decomposition of long-minus-short portfolio return variance

Based on the decomposition of Campbell and Shiller (1988a), unexpected returns can by definition be split into two components: revisions to cash-flow expectations (i.e., CF news or CF shocks), and revisions to expected returns (i.e., DR news or DR shocks). It follows, that we can also decompose the variance of long-minus-short anomaly portfolio returns into separate CF and DR shocks components. Lochstoer and Tetlock (2020) find that CF variance is always significantly bigger than DR variance for the five anomaly portfolios that they study⁶. We seek to go deeper, by

⁶Those five anomaly portfolios are sorted on Book-to-Market, Size, Profitability, Investment, and Momentum. In this paper, we study the same five anomalies.

further decomposing the portfolio CF and DR shocks. Indeed, LMS portfolios can be expressed as a weighted sum of individual firm's stock returns. Therefore, we can just split it in its different components, to understand why CF shocks dominate the way they do for anomaly portfolios. We can re-write the variance of a portfolio's long-minus-short return (i.e., be it the CF return, DR return or total return) as:

$$\begin{aligned} Var(r^{LMS}) &= Var(r^L - r^S) \\ &= Var(r^L) + Var(r^S) - 2Cov(r^L, r^S) \end{aligned} \quad (III.1)$$

where r denotes the vector of returns, and LMS , L and S the long-minus-short, long-leg, and short-leg portfolios, respectively. Long and short portfolio variances are equal to the value-weighted sum of their constituents, as follows:

$$\begin{aligned} Var(r^L) &= Var\left(\sum_{i=1}^{N_L} r_i \cdot w_i\right) \\ Var(r^S) &= Var\left(\sum_{i=1}^{N_S} r_i \cdot w_i\right) \end{aligned} \quad (III.2)$$

where N_L and N_S are the number of stocks in the high and low portfolios and w_i is the weighting vector based on market equity for cap-weighted portfolios. Since the variance of a sum of random variables can be expressed as the sum of all the individual variances plus the sum of two times all cross-covariances, we can re-write equation (III.2) as follows:

$$\begin{aligned} Var\left(\sum_{i=1}^{N_L} r_i \cdot w_i\right) &= \sum_{i=1}^{N_L} Var(r_i \cdot w_i) \\ &\quad + \sum_{j=1, j \neq k}^{N_L} \sum_{k=1, k \neq j}^{N_L} 2Cov(r_j \cdot w_j, r_k \cdot w_k) \\ Var\left(\sum_{i=1}^{N_S} r_i \cdot w_i\right) &= \sum_{i=1}^{N_S} Var(r_i \cdot w_i) \\ &\quad + \sum_{j=1, j \neq k}^{N_S} \sum_{k=1, k \neq j}^{N_S} 2Cov(r_j \cdot w_j, r_k \cdot w_k) \end{aligned} \quad (III.3)$$

Finally, by replacing equation (III.3) into equation (III.1) and re-arranging the terms we can re-write:

$$\begin{aligned}
Var(r^{LMS}) &= -2Cov(r^L, r^S) \\
&+ \sum_{i=1}^{N_L} Var(r_i \cdot w_i) + \sum_{i=1}^{N_S} Var(r_i \cdot w_i) \\
&+ \sum_{j=1, j \neq k}^{N_L} \sum_{k=1, k \neq j}^{N_L} 2Cov(r_j \cdot w_j, r_k \cdot w_k) \\
&+ \sum_{j=1, j \neq k}^{N_S} \sum_{k=1, k \neq j}^{N_S} 2Cov(r_j \cdot w_j, r_k \cdot w_k)
\end{aligned} \tag{III.4}$$

The first term in this sum corresponds to the covariance between the aggregate returns of the two long and short portfolios. The terms in the second row relate to the variance of individual firm returns, which we label $\sum Var(r_i^L)$ and $\sum Var(r_i^S)$. The terms in the third row capture the cross-covariance components across firms within a same portfolio. We label them $\sum 2Cov(r_{j,k}^L)$ for cross-covariances in the long and $\sum 2Cov(r_{j,k}^S)$ for cross-covariances in the short portfolio. With this notation, we can re-write the variance of any LMS portfolio's CF and DR shocks as follows:

$$\begin{aligned}
Var(CF^{LMS}) &= -2Cov(CF^L, CF^S) \\
&+ \sum Var(CF_i^L) + \sum Var(CF_i^S) \\
&+ \sum 2Cov(CF_{j,k}^L) + \sum 2Cov(CF_{j,k}^S)
\end{aligned} \tag{III.5}$$

$$\begin{aligned}
Var(DR^{LMS}) &= -2Cov(DR^L, DR^S) \\
&+ \sum Var(DR_i^L) + \sum Var(DR_i^S) \\
&+ \sum 2Cov(DR_{j,k}^L) + \sum 2Cov(DR_{j,k}^S)
\end{aligned}$$

We can finally replace these expressions for $Var(CF^{LMS})$ and $Var(DR^{LMS})$ in the following equation that decomposes total portfolio return variance into CF and DR shock components:

$$\begin{aligned}
Var(r^{Ptf}) &= Var(CF^{LMS}) + Var(DR^{LMS}) \\
&- 2Cov(CF^{Ptf}, DR^{Ptf})
\end{aligned} \tag{III.6}$$

By empirically estimating the different components in equation (III.5) and additionally leveraging equation (III.6) we obtain the tools to assess the underlying drivers of anomaly returns. In particular, we want to make use of the decomposition

in (III.5) to answer the following questions:

First, how big is the covariance term between long and short portfolios for aggregate cash flows, i.e., “ $Cov(CF^L, CF^S)$ ”, and for aggregate discount rates, i.e., “ $Cov(DR^L, DR^S)$ ”? More importantly, how specific to anomalies is this portfolio covariance? Empirically, do we observe similar differences between portfolio CF and DR covariance for any well-diversified long and short portfolio, or is it specific to anomalies? For instance, if $Cov(DR^L, DR^S) > Cov(CF^L, CF^S)$ holds true for any well-diversified long and short portfolio, then we would expect any LMS portfolio to be biased towards a greater CF variance component. Indeed, if empirically this relationship always holds true, then $Var(CF^{LMS}) > Var(DR^{LMS})$, as per equation (III.6).

Second, how specific to anomalies are the sum of cross-covariance terms? These terms, displayed in the third line of equations (III.5) capture the comovement across stocks within a same portfolio. Therefore, it estimates how much firms sharing a given anomaly characteristic (e.g., the quintile of firms classified as value or momentum stocks) tend to move together. The question we seek to answer is: how specific are CF and DR comovements for anomalies, compared to comovements occurring across stocks unrelated to those characteristics? Do commonalities in anomalies occur in CF shocks, DR shocks, or both?

Third, what happens when we condition CF and DR comovement on macro variables? Do CF shocks of stocks comove more during negative aggregate CF events, such as recessions? Are variations in the sum of cross-covariances specific to anomaly characteristics? If certain anomalies predict a shared exposure to a systematic risk factor, then we would expect the sum of cross-covariances (i.e., firm comovement) to spike during periods in which firms have to face this risk because it effectively materializes. Based on models such as Zhang (2005), and on the empirical work of Campbell et al. (2010), we would for example expect that CF cross-covariances among value stocks spike during recessions. Meanwhile, based on the model of Lettau and Wachter (2007), we would expect that CF comovement in growth firms spikes when aggregate risk aversion is elevated, like in periods of high sentiment.

III. Estimation of CF and DR news

A. The model

It is possible to estimate updates in investor expectations about future stock returns into CF and DR components following the approximation of Campbell and

Shiller (1988a):

$$\begin{aligned}
r_{t+1} - E_t r_{t+1} &\approx (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} \\
&\quad - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j} \\
&\approx CF_{t+1} - DR_{t+1}
\end{aligned} \tag{III.7}$$

where E_t represents expectation at time t , r_t and d_t denote log-returns and dividends at time t respectively, and ρ is a parameter of linearization that depends on the long-term mean of the log dividend-price ratio. This equation suggests that higher unexpected returns are either due to higher than expected future cash flows, lower than expected future discount rates, or a combination of the two.

A common and easily applicable approach to empirically estimate this return decomposition is that of Campbell (1991), who relies on a Vector Autoregressive (VAR) model of order 1 to construct a proxy for discount rate news based on predictions of future expected returns. He then retrieves cash flow news residually from Equation (III.7). Following this idea, the usual approach in the literature consists of modeling the dynamics of stock returns first, using some variation of the VAR model. Indeed, building the discount rate news proxy first is usually easier than estimating cash-flow expectations directly. It can alleviate certain concerns, like dealing with seasonality in dividends used to estimate firm cash-flows.

In this thesis, the empirical implementation of the decomposition is based on the approach of Lochstoer and Tetlock (2020), which itself is a model that builds on Campbell (1991) and Campbell et al. (2010), and which aims at modelling CF and DR dynamics of firm-level stock returns. It assumes that expected log stock returns can be expressed as a linear function of market-adjusted (*ma*) and market-aggregated (*agg*) characteristics:

$$E_t[r_{i,t+1}] = \alpha + \beta'_1 X_{i,t}^{ma} + \beta'_2 X_t^{agg} \tag{III.8}$$

The market-adjusted variables $X_{i,t}^{ma}$ include widely used characteristics predicting stock returns, which are Book-to-Market (*BM*), Market Cap (Size), Profitability (*Prof*), Investment (*Inv*) and Momentum (*Mom*). These firm-specific components of returns are adjusted for the market by demeaning each characteristic based on its average at time t . The aggregated market variables X_t^{agg} are the value-weighted

averages of the same five characteristics at time t . β_1 captures expected cross-sectional variation while β_2 captures the expected time-series variation in returns.

We implement the return decomposition by running two separate VAR(1) estimations, one for the market-adjusted and one for the aggregate component of returns:

$$Z_{i,t+1} = \mu^{ma} + \Pi^{ma} Z_{i,t}^{ma} + \epsilon_{i,t+1}^{ma} \quad (\text{III.9})$$

$$Z_{t+1} = \mu^{agg} + \Pi^{agg} Z_t^{agg} + \epsilon_{t+1}^{agg} \quad (\text{III.10})$$

where the first element of Z is the stock return and the other elements are the state variables X defined above. Z^{ma} is of size K^{ma} and Z^{agg} is of size K^{agg} . μ is the intercept and ϵ are conditionally mean-zero shocks. DR shocks are by definition updates in expected returns from one period to the next:

$$DR_{t+1} = E_{t+1} \sum_{j=2}^{\infty} \rho^{j-1} r_{t+j} - E_t \sum_{j=2}^{\infty} \rho^{j-1} r_{t+j} \quad (\text{III.11})$$

Therefore it is possible to back out the DR terms from the two VAR systems based on the model of Campbell (1991):

$$DR_{i,t+1}^{ma} = e'_1 \rho \Pi^{ma} (I - \rho \Pi^{ma})^{-1} \epsilon_{i,t+1} \quad (\text{III.12})$$

$$DR_{t+1}^{agg} = e'_1 \rho \Pi^{agg} (I - \rho \Pi^{agg})^{-1} \epsilon_{t+1} \quad (\text{III.13})$$

where e'_1 is a vector of zeros except for its first element which is equal to one. I is the identity matrix and the parameter ρ is a constant relates to the average dividend yield or consumption-wealth-ratio. Given the empirical measures of those variables over the past decades, this constant is usually set to 0.95 in the literature when working with yearly data (Campbell and Vuolteenhao, 2004).

Using the estimated DR shock, the CF shock component of unexpected returns can simply be backed out as follows:

$$CF_{t+1} = r_{t+1} - E_t[r_{t+1}] + DR_{t+1} \quad (\text{III.14})$$

thus, for the two separate estimations:

$$CF_{i,t+1}^{ma} = e'_1 (I + \rho \Pi^{ma} (I - \rho \Pi^{ma})^{-1}) \epsilon_{i,t+1} \quad (\text{III.15})$$

$$CF_{t+1}^{agg} = e'_1 (I + \rho \Pi^{agg} (I - \rho \Pi^{agg})^{-1}) \epsilon_{t+1} \quad (\text{III.16})$$

Finally, total shocks to CF and DR expectations can be obtained by adding the market-adjusted and the aggregate components.

$$DR_{i,t} = DR_t^{agg} + DR_{i,t}^{ma} \quad (\text{III.17})$$

$$CF_{i,t} = CF_t^{agg} + CF_{i,t}^{ma} \quad (\text{III.18})$$

Choosing this approach to estimate updates in CF and DR expectations allows for the estimation of separate coefficients in the VAR for both changes over time and across firms. Both dimensions are important drivers in revisions of expected returns as shown by Vuolteenaho (2002), Chen et al. (2013) and Campbell et al. (2010). This two-step approach allows for the estimation of each component independently, without overlapping them together, which may lead to overweight one of the two dimensions.

The VAR decomposition also comes with its limitations, most notably that CF shocks are backed out as residuals. Hence, when using poor predictors of future expected returns, the CF measure will by construction capture the model’s residual noise and everything unexplained by it. Chen and Zhao (2009) and Chen et al. (2013) provide a more in-depth discussion of the possible drawbacks of estimating DR using the VAR approach. Nonetheless, VAR models remain the standard procedure in the literature to estimate CF and DR news, which is we stick to this approach.

B. *The data*

We rely on Compustat for information about quarterly firm fundamentals. To assign a quarterly observation to the appropriate quarter, we follow the empirical approach of Lyle and Wang (2015) and assume that an observation becomes public the day of the earnings announcement. Hence, we set the date of an observation to the end of the calendar quarter it falls in. Earnings announcements are rarely spaced at perfectly regular three-month intervals. Therefore, if two earnings dates fall in the same quarter, the earlier one is removed. Conversely, if no earning was made public over the 90 days of a calendar quarter, the firm will have no observation reported.

Except for decisions linked to the quarterly frequency of the analysis, we follow Lochstoer and Tetlock (2020) for all the choices related to data selection and treatment. Stock return information for all stocks on the NYSE, AMEX, and NASDAQ comes from CRSP. Due to the data selection process described below, our dataset spans over the period from 1973 up to 2020. We adjust for delisting returns whenever

possible. If the delisting return is not provided, we use the average delisting return of the corresponding delisting type of event. We also remove micro-caps from the sample in the baseline specification⁷. This choice implies that each quarter, we remove all observations that fall below the 20th percentile of market equity. Quintiles are computed only based on NYSE stocks. Additionally, to avoid back-filling biases and ensure data quality, we require observations from Compustat to be preceded by a valid book-equity value in quarter t-1, to have a December fiscal-year close in t-1 and at least two valid earnings and total assets observations in the past five periods. Moreover, book-equity must be strictly positive. Additionally, for CRSP data, a valid market-equity in t-1 and a valid return in the month preceding the period end are required to avoid spurious return predictability due to stale prices.

Furthermore, we drop all duplicate data entries in Compustat caused by restatements or changes in a firm’s fiscal year-end. To merge CRSP data with quarterly fundamentals, we transform all monthly observations to quarterly data, which leads us to compute cumulated returns, the sum of trading volume, and take the last available price and number of shares outstanding during the quarter.

The choice of state variables for the VAR model is motivated by the rationale of selecting a parsimonious set of predictors that have been found to forecast firm-level stock returns. Those variables includes stock returns, book-to-market (BM), size, profitability, investment, momentum (Mom), and return-on-equity (ROE). The choice of these variables is motivated by the asset pricing literature that documents their ability to predict returns out of sample⁸. Those characteristics are also among the most studied and robust factors in asset pricing (McLean and Pontiff, 2016; Harvey et al., 2016).

The return variable, $\ln RealRet$, is computed by taking the logarithm of one plus the quarterly stock return adjusted for inflation. To ensure stationarity of the size variable, $d5_lnME$, we compute it by taking the logarithm of one plus the five-year change in market equity (ME). ME is equal to the number of shares outstanding multiplied by the share price. We carefully sum ME across all classes of common stock for the same firm. Since this variable requires a five-year lag, numerous observations are lost in the earliest years of Compustat. For the momentum variable,

⁷Apart from Lochstoer and Tetlock (2020), several authors have pointed out that the correlation between firm-level CF and DR shocks is sensitive to the inclusion of micro-caps (e.g., Vuolteenaho (2002), Khimich (2012)). We acknowledge that this sample choice might influence the negative correlation between CF and DR that we observe throughout these studies. Nonetheless, the robustness checks reported in table II show that the negative CF-DR correlation at the firm level tends to persist in our sample even when including micro-caps.

⁸See Jegadeesh and Titman (1993) and Carhart (1997) for momentum, Fama and French (1993) for size and book-to-market, and Fama and French (2015) for profitability and investment.

$\ln Mom12$, we compute the cumulated return over the past twelve months and add one to it before taking the natural logarithm. Since the implied holding period is of three months due to the quarterly data, a twelve-month formation period yields the most substantial momentum effect (Jegadeesh and Titman, 1993).

The investment characteristic, $\ln Inv$, is measured by taking the logarithm of the five-year increase in assets (see Cooper et al. (2008)). More specifically, to mitigate seasonality issues, we compute the percentage increase between the average total assets over the most recent year and the average total assets five years in arrears. ROE is defined as earnings divided by the previous' quarter book equity. To smooth out seasonality patterns, quarterly earnings are defined as the rolling average earnings over the past four quarters. If it is missing, we replace it using the clean-surplus accounting identity⁹ and $\ln ROE$ is defined as the logarithm of one plus ROE. Following Fama and French (2015), profitability is equal to revenue minus costs of goods sold, interest expense, and selling and administrative expenses, all divided by book-equity. All variables are computed quarterly. In case of missing observations, we assume it is equal to $0.2 + 0.5ROE$ as a proxy. To mitigate seasonality concerns, profitability is the rolling average over the past four quarters. The final profitability variable, $\ln Prof$, is defined as the logarithm of one plus profitability.

Following Vuolteenaho (2002) before computing the above variables, all assets are transformed to pseudo-firms, i.e., a portfolio which comprises 90% of the firm's stock and 10% of the risk-free asset. This is done because of the log transformation. If a firm goes bankrupt, its return is -1, leading to a problematic logarithm of zero, which is prevented by the pseudo-firm weighting.

C. Empirical estimation of the VAR model

Table I presents the estimation of the VAR model of order 1. The estimated coefficient matrix is the one used for Π in equations (III.9) through (III.16). Panel A presents the model for the market-adjusted (i.e., firm-specific) component. We use one lag and no quarter fixed-effect in the regressions¹⁰. The goal of the procedure is to estimate the expected value of returns, $\ln RealRet$, using past data. All anomaly characteristics significantly forecast returns in the anticipated direction: High book-

⁹In clean-surplus accounting, earnings are equal to change in book equity plus the net payout to shareholders.

¹⁰The assumption is that the smoothing procedure of the variables described in section III.B alleviates seasonality concerns. As a robustness check, chapter III of the thesis estimates the VAR using calendar-quarter fixed-effects, firm fixed-effects, and a VAR of order 4 to capture a full year of past information. CF and DR estimates remain comparable to the results presented here, leading to similar variance contribution and correlation patterns.

Table I - VAR coefficient estimates: This table reports the estimated coefficient matrix Π for equations (III.9) in panel A and (III.10) in panel B. Calendar-quarter fixed are included in the aggregate specification. Errors are clustered by date and firm in the panel regression.)

Panel A: VAR estimates for the firm-specific component							
	$\ln\text{RealRet}_{t0}^{ma}$	$\ln\text{BM}_{t0}^{ma}$	$\ln\text{Prof}_{t0}^{ma}$	$\ln\text{Inv}_{t0}^{ma}$	$\ln\text{ME}_{t0}^{ma}$	$\ln\text{Mom}_{t0}^{ma}$	$\ln\text{ROE}_{t0}^{ma}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\ln\text{RealRet}_{t-1}^{ma}$	-0.02 (-1.00)	0.01 (0.35)	-0.00 (-0.17)	0.00 (0.06)	-0.00 (-0.09)	0.31*** (13.86)	-0.01*** (-8.40)
$\ln\text{BM}_{t-1}^{ma}$	0.01** (2.47)	0.98*** (292.38)	0.00 (1.01)	-0.00*** (-5.40)	-0.00 (-0.10)	0.00 (1.32)	-0.01*** (-16.12)
$\ln\text{Prof}_{t-1}^{ma}$	0.03*** (2.73)	-0.00 (-0.33)	0.32*** (4.20)	0.00*** (2.84)	0.05*** (3.44)	0.04*** (2.94)	0.03*** (5.22)
$\ln\text{Inv}_{t-1}^{ma}$	-0.17*** (-3.70)	0.17*** (3.37)	-0.06*** (-3.08)	0.92*** (260.23)	0.12 (1.52)	-0.26*** (-4.37)	-0.02*** (-2.80)
$\ln\text{ME}_{t-1}^{ma}$	-0.00* (-1.83)	0.01*** (4.59)	0.00 (1.48)	0.00*** (21.85)	0.93*** (197.41)	-0.00* (-1.65)	0.00*** (6.33)
$\ln\text{Mom}_{t-1}^{ma}$	0.03*** (3.00)	-0.01 (-0.86)	0.02*** (4.45)	-0.00 (-0.79)	0.10*** (8.86)	0.69*** (66.83)	0.01*** (15.94)
N	207,096	207,096	207,096	207,096	207,096	207,096	207,096
Adjusted R^2	0.01	0.94	0.11	0.94	0.89	0.58	0.14

Panel B: VAR estimates for the Market-aggregated component							
	$\ln\text{RealRet}_{t0}^{agg}$	$\ln\text{BM}_{t0}^{agg}$	$\ln\text{Prof}_{t0}^{agg}$	$\ln\text{Inv}_{t0}^{agg}$	$\ln\text{ME}_{t0}^{agg}$	$\ln\text{Mom}_{t0}^{agg}$	$\ln\text{ROE}_{t0}^{agg}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\ln\text{RealRet}_{t-1}^{agg}$	0.11 (1.35)	-0.07 (-0.88)	-0.06* (-1.95)	-0.00 (-0.23)	0.31** (2.08)	0.47*** (4.29)	-0.00 (-0.81)
$\ln\text{BM}_{t-1}^{agg}$	0.02 (0.92)	0.96*** (42.74)	-0.00 (-0.10)	0.00 (0.84)	0.02 (0.46)	0.04 (1.24)	0.00** (2.21)
$\ln\text{Prof}_{t-1}^{agg}$	-0.17 (-0.87)	0.29 (1.55)	0.08 (1.12)	0.01*** (3.36)	-0.52 (-1.51)	-0.08 (-0.32)	0.03*** (3.07)
$\ln\text{Inv}_{t-1}^{agg}$	-1.11 (-0.95)	1.83* (1.67)	0.16 (0.36)	0.98*** (58.79)	-1.71 (-0.84)	-1.01 (-0.66)	0.24*** (3.57)
$\ln\text{ME}_{t-1}^{agg}$	-0.00 (-0.12)	-0.01 (-0.58)	-0.00 (-0.53)	0.00 (1.59)	0.92*** (21.37)	0.00 (0.14)	0.00** (2.02)
$\ln\text{Mom}_{t-1}^{agg}$	-0.06 (-1.39)	0.07 (1.57)	0.00 (0.02)	0.00 (0.05)	-0.04 (-0.55)	0.63*** (10.67)	0.00 (0.77)
Quarter-FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	180	180	180	180	180	180	180
Adjusted R^2	0.02	0.95	0.12	0.96	0.84	0.57	0.16

to-market, profitability, and momentum in the previous quarter, as well as small size and low investment, imply higher returns during the holding period.

The R^2 of the regression is low, a bit above 1%, which is in line with prior studies and highlights the difficulty in forecasting future returns. The VAR setting also estimates the dynamics of the variables used to forecast returns. If those variables have strong forecastability, a change in one of those variables will have a

persistent effect on the estimated expected return multiple quarters into the future. This follows from the VAR approach of Campbell (1991) and can be inferred from equations (III.12) and (III.13). Book-to-market, investment, and size are strongly auto-correlated and move slowly. Therefore, a change in one of those characteristics will have a lasting influence on a stock’s expected return. Profitability and momentum are less persistent, and changes in those variables will therefore have a less lasting effect on expected returns. ROE is included in the system of equations as a dependent variable but omitted from the independent variables.

Panel B presents the estimates for the aggregate (i.e., market) component. At the aggregate level, we include calendar-quarter fixed effects to account for market-wide seasonality patterns. The quarterly data yields two principal differences with Lochstoer and Tetlock (2020) which serves as a reference point. First, Profitability is not significantly persistent, showing that there is more quarterly earnings variability and that periodic effects might still cause some noise. Second, the R^2 of the regression forecasting aggregate returns is lower for quarterly than for yearly data. This might be consistent with the predictability of market returns over long periods, depending on aggregate valuation ratios (Campbell and Shiller, 1988b, 2001).

As mentioned above, the literature shows that the VAR approach for the return decomposition is highly sensitive to the chosen specification (e.g., Chen and Zhao (2009)). To make sure that the results we document in this paper are not driven by a particular implementation of the model, we use six different specifications to test the robustness of our results and further contribute to the literature by highlighting to which conditions the decomposition is the most sensitive. We settle for the following specifications:

1. (Baseline, S1) : Quarterly data. $\ln Ret$, $\ln BM$, $d5_lnInv$, $\ln Prof$, $\ln Mom12$, $d5_lnME$ and $\ln ROE$ as dependent variables. The same variables lagged minus $\ln ROE$ as independent variables. Quarter-fixed-effects in the *agg* market decomposition to account for calendar fixed-effects. No micro-caps. Valid sample from 1974 to 2020. Based on Lochstoer and Tetlock (2020). This is the specification reported in table I.
2. (S2) : Same as S1, except for the use of yearly data. $\ln Ret$, $\ln BM$, $d5_lnInv$, $\ln Prof$, $\ln Mom6^{11}$, $d5_lnME$ and $\ln ROE$ are the dependent variables. The same variables lagged minus $\ln ROE$ are the independent variables. No micro-caps. Valid sample goes from 1974 to 2020.

¹¹For a twelve months implied holding period, momentum is strongest when the formation period is based on the prior 6 months

3. (S3) : This approach aims at replicating the approach of Vuolteenaho (2002). We use yearly data and $\ln Ret$, $\ln BM$, and $\ln ROE$ are the dependent variables. Those same variables lagged are the independent variables. Micro-caps are included. For consistency with the reference paper, we use the sample period from 1974 through 2002.
4. (S4) : Same as S3 above, except we use the modern sample period from 2001 through 2020. Again we use yearly data. $\ln Ret$, $\ln BM$, and $\ln ROE$ are the dependent variables and their lagged versions serve as independent variables. Micro-caps are included.
5. (S5) : Same as S1, except that micro-caps are included. $\ln Ret$, $\ln BM$, $d5_lnInv$, $\ln Prof$, $\ln Mom12$, $d5_lnME$ and $\ln ROE$ are the dependent variables. The same variables lagged minus $\ln ROE$ are used as independent variables. Valid sample goes from 1974 to 2020.
6. (S6) : Same as S1, except that that we use only the modern sample period from 2001 through 2020. $\ln Ret$, $\ln BM$, $d5_lnInv$, $\ln Prof$, $\ln Mom12$, $d5_lnME$ and $\ln ROE$ are again the dependent variables. The same variables lagged minus $\ln ROE$ are used as independent variables. Micro-caps are excluded.

We report the estimation of the other five different VAR model specifications in appendix D. Panels A report the firm-specific-component VAR estimation, and panel B the market-aggregate VAR. Since they are identical, the same comments about the decomposition made in table I for the baseline specification apply; including the fact that the coefficients on the lagged variables in the firm-level decomposition are all significant and of the expected sign.

Even if the decomposition estimations remain very similar, we can note a few differences for the robustness specifications. We observe that when going from quarterly to yearly, most variables become less persistent, in particular momentum. Momentum is also the variable losing the most predictive power in the firm specific VAR when going from quarterly to yearly data. The coefficient on lagged book-to-market remains significantly positive throughout most specifications, in particular in the approach following Vuolteenaho (2002), S3. Using the shorter modern samples that coincide with our period of data availability for news *tone* generally leads to lower R^2 and weaker coefficients. The results for specification S5, when including micro-caps, are also very similar. However, the total number of observations only increases slightly, by a few thousand, suggesting that our stringent requirements for high data quality already weed out a significant portion of smaller firms.

At the aggregate level, we observe that the regression R^2 s tend to be higher. In

the baseline quarterly specification, it is however lower than for our reference point from Lochstoer and Tetlock (2020). The difference seems to be driven by the data frequency, since switching to yearly data increases the R^2 , to 0.14 (S2). Most of the coefficients remain non-significant. Including micro-caps (S5) also seems to lead to slightly better fitting models.

In this paper, we are interested in the relative contribution of CF and DR news to unexpected return variance, in particular for anomaly portfolios. Therefore, in table II we report the relative contribution of CF and DR shocks to return variance. Consistent with the literature (Vuolteenaho, 2002), we find in panel A that the CF component dominates at the firm-specific level (making up between 60% and 97% of unexpected return variation across the 6 specification). Conversely, DR shocks account for most of the variation in unexpected returns of the aggregate market portfolio. In our baseline specification, DR accounts for 62% of unexpected return variance and CF shocks only make up 24% of the variance. Previous studies (Vuolteenaho, 2002; Chen et al., 2013) already suggested that CF shocks are relatively more diversified away, such that DR is the main driving force of returns at the aggregate level.

When combining the market-adjusted and aggregate components to compute the total firm CF and DR shocks, we observe that CF updates remain the primary driver of unexpected returns. This stays true across all specifications, making up between 53% and 91% of the variance. DR shocks remain a more modest component of individual stock returns, but in many specifications, the covariance between CF and DR shocks also accounts for a good portion of return variance. This covariance component usually translates into a negative correlation between CF and DR shocks, as in Lochstoer and Tetlock (2020). Like them, however, we notice that this negative correlation is sensitive to the chosen specification. For example, when following the approach of Vuolteenaho (2002) (S3), the correlation becomes slightly positive.

In panel B, we replicate the results for anomaly portfolios. Those portfolios are constructed by cap-weighting individual firm returns based on previous' December's market cap, and then going long the top quintile of the anomaly characteristic and shorting the bottom quintile. Our results are almost perfectly in line with the reference paper. Of particular interest is that the conclusions we can draw from this table seem to be robust to the chosen specification. CF dominates DR shocks for all anomalies and across all specifications. In the baseline specification (S1), CF shocks account for anywhere between 75% and 87% of the variation in returns across all anomaly portfolios. DR shock contributions vary only between 3% and 6%. More

Table II - CF and DR contribution to variance. “Total decomposition” refers to the total unexpected return, as in equation III.7. The contribution to unexpected return variance (in %) is reported for the CF, DR and the covariance components. We report the correlation between the CF and DR return series under “Corr CF/DR”. “Firm-specific component” refers to the unexpected return series based on the market-adjusted VAR, which is reported in panel a of table I. “Aggregate component” reports the variance decomposition for the market-wide component. In panel B, we report the decomposition for the LMS portfolios based on the five anomalies.

Panel A: All firms		S1	S2	S3	S4	S5	S6
Total firm decomposition	%CF	0.70	0.69	0.88	0.91	0.53	0.88
	%DR	0.13	0.12	0.14	0.03	0.18	0.10
	%-2cov	0.17	0.19	-0.02	0.06	0.29	0.02
	Corr CF/DR	-0.29	-0.34	0.03	-0.17	-0.46	-0.03
Firm-specific component	%CF	0.79	0.73	0.68	0.97	0.60	0.98
	%DR	0.05	0.07	0.08	0.01	0.13	0.01
Aggregate component	%CF	0.24	0.34	0.35	0.29	0.18	0.26
	%DR	0.62	0.39	2.29	0.36	0.59	0.73

Panel B: Anomaly Portfolios		S1	S2	S3	S4	S5	S6
Book-to-Market	%CF	0.86	0.70	0.74	0.95	0.57	1.05
	%DR	0.05	0.05	0.04	0.00	0.13	0.01
	%-2cov	0.08	0.25	0.22	0.05	0.30	-0.06
	Corr CF/DR	-0.19	-0.69	-0.61	-0.65	-0.56	0.29
Size	%CF	0.76	0.73	0.64	0.96	0.51	1.00
	%DR	0.03	0.04	0.07	0.00	0.11	0.01
	%-2cov	0.19	0.23	0.28	0.03	0.39	-0.01
	Corr CF/DR	-0.64	-0.70	-0.65	-0.54	-0.83	0.07
Profitability	%CF	0.87	0.67	0.48	0.97	0.60	0.94
	%DR	0.06	0.12	0.15	0.00	0.14	0.02
	%-2cov	0.08	0.21	0.37	0.03	0.26	0.03
	Corr CF/DR	-0.18	-0.38	-0.70	-0.29	-0.46	-0.11
Investment	%CF	0.80	0.71	0.71	0.99	0.60	1.00
	%DR	0.04	0.05	0.06	0.00	0.12	0.01
	%-2cov	0.16	0.23	0.23	0.01	0.28	-0.01
	Corr CF/DR	-0.42	-0.60	-0.55	-0.19	-0.53	0.05
Momentum	%CF	0.75	0.73	0.71	0.99	0.52	0.98
	%DR	0.03	0.04	0.05	0.00	0.10	0.00
	%-2cov	0.22	0.23	0.24	0.01	0.38	0.02
	Corr CF/DR	-0.77	-0.66	-0.67	-0.29	-0.85	-0.16

generally, this is true for all specifications and all anomaly portfolios: CF shocks always contribute more to the portfolio variance (at least 50%) than DR shocks (at most 15%).

Like our reference paper, we observe that the correlation between portfolio CF and DR shocks are consistently negative (except in specification S6). This finding suggests that at the portfolio level, the negative CF/DR correlation is more consistent across specifications than for firm-level returns. This finding tends to be in line with explanations for anomaly premia that suggest either overreaction to CF shocks (Barberis et al., 1998) or changes in risk following CF revisions (Kogan and Papanikolaou, 2013).

IV. Results

A. Simulations on random LMS portfolios

Based on the intuition provided in section II, we explore how portfolios constructed based on randomly selected stocks relate to the decomposition put forward in equation (III.5). The evidence we find in table II, panel A leads us to be cautious about the observed CF dominance in anomaly return variance. Indeed, since CF dominates at the firm level, but DR is the dominant force at the aggregate level, this might suggest that $Cov(DR^L, DR^S) > Cov(CF^L, CF^S)$ (as per equation (III.5)), holds always true.

We simulate 35,000 random portfolios. Every period, we randomly select one fifth of all available firms and assign them to the long leg portfolio (L). Simultaneously, we randomly select another fifth of all firms and assign them to the short leg portfolio (S). We rebalance every period using the same random procedure. We compute the resulting portfolio's CF and DR series as previously, using a market-cap weighted approach. Having established patterns that are roughly consistent across specifications, we focus on the baseline specification, S1, henceforth.

In the average long-minus-short portfolio we simulate, CF shocks are the dominant drivers of returns. 84% of unexpected return variance is due to CF revisions, 6% is due to DR shocks, and 10% can be imputed to the covariance term. We also observe little variation around those mean values. The standard deviation is of 4% for CF and 1% for DR shocks. CF shocks always dominate: their lowest contribution to long-minus-short portfolio variance is 68%, whereas DR variance never exceeds 12%. We also observe that the correlation between CF and DR shocks of long-minus-short portfolios is negative on average (-0.24) with a standard deviation

Table III - Simulations: For 35,000 randomly long-minus-short simulated portfolios, this table reports summary statistics about the estimated relative CF and DR contributions to portfolio variance. The random portfolio assignment procedure is explained in section IV.A. Panel B and C report the summary statistics of the 35,000 long and 35,000 short simulated portfolios (which are as expected very close since they are built following the same procedure). Panel A reports the statistics for the resulting long-minus-short portfolios.

Panel A: Long-minus-Short portfolios							
	Average	Standard deviation	Minimum	25 th percentile	Median	75 th percentile	Maximum
%CF	0.84	0.04	0.68	0.81	0.84	0.86	0.97
%DR	0.06	0.01	0.02	0.05	0.06	0.07	0.12
%-2cov	0.10	0.04	-0.07	0.08	0.11	0.13	0.24
correl CF/DR	-0.24	0.09	-0.55	-0.30	-0.24	-0.18	0.12

Panel B: Long-leg portfolios							
	Average	Standard deviation	Minimum	25 th percentile	Median	75 th percentile	Maximum
%CF	0.25	0.01	0.20	0.24	0.25	0.26	0.32
%DR	0.54	0.02	0.46	0.52	0.54	0.56	0.65
%-2cov	0.21	0.02	0.08	0.19	0.21	0.22	0.29
correl CF/DR	-0.28	0.04	-0.42	-0.31	-0.28	-0.25	-0.10

Panel C: Short-leg portfolios							
	Average	Standard deviation	Minimum	25 th percentile	Median	75 th percentile	Maximum
%CF	0.25	0.01	0.20	0.24	0.25	0.26	0.32
%DR	0.54	0.02	0.45	0.52	0.54	0.56	0.66
%-2cov	0.21	0.03	0.09	0.19	0.21	0.22	0.29
correl CF/DR	-0.28	0.04	-0.43	-0.31	-0.28	-0.25	-0.10

of 0.09.

In panels B and C, we report the same statistics about return recompositions for the simulated long and short portfolios, respectively. As we would have expected, based on the results in table II panel A for the aggregate component of returns, DR shocks systematically dominate in those well-diversified portfolios covering one-fifth of the stock equity universe. DR shocks account for at least 45% of return variation in those portfolios, while CF shocks account at most for 32% of return variation across the 35,000 random portfolios.

Those simulations may cast doubt and incite us to be careful before concluding on the specificity of anomaly characteristics to generate the observed CF dominance in return variance. High CF and low DR variance appears to be a feature shared by any well-diversified long-minus-short portfolio. Even the low DR contribution to

return variance and the negative correlation between CF and DR do not appear to be a unique feature of anomaly portfolios.

B. Detailed variance decomposition

In section II, we laid out the hypothesis that any well-diversified LMS portfolio is mainly driven by CF shocks because the correlation in DR shocks of any two-well diversified portfolios is higher than the corresponding CF correlation. We empirically test this hypothesis by computing each of the elements of portfolio variance as per equation (III.5) and report our results in table IV. In panel A we report the estimated subcomponents for the randomly simulated portfolios. We label the estimated mean and standard deviation of the variance components $\hat{\mu}$ and $\hat{\sigma}$, respectively.

As expected from table III, the CF variance of LMS portfolios is multiple times greater ($\times 14$ on average) than their DR variance. Empirically, this decomposition also supports our first hypothesis from section II: $Cov(DR^L, DR^S) > Cov(CF^L, CF^S)$ always holds, independently of anomaly characteristics. In particular, the covariance term, $-2cov(L, S)$, between long and short portfolios for DR implies a decrease in long-minus-short return variance which is about threefold compared to CF (-44.3 vs. -15.3)¹².

Our results further support the notion that stocks generally exhibit stronger DR than CF comovements on average. The average sum of DR cross-covariances is almost three times larger than the average for CF cross-covariances (21.3 vs. 7.5). This would be expected, given that DR variance dominates in aggregate market portfolios (e.g., Vuolteenaho (2002); Chen et al. (2013)). Panel A of table II also documents the larger variance at the aggregate level of DR shocks.

In panel B of table II, we investigate how these subcomponents of variance relate to anomaly portfolio CF and DR. First, we find that the relative CF dominance of anomaly portfolios can likewise be traced back to high DR correlations between the long and short leg portfolios, just like observed in the simulations. For example, the CF return variance of the HML book-to-market portfolio equals 30.2, whereas its DR variance is of 1.8¹³. This difference (28.4) is almost equivalent to the difference induced by the covariance terms, $-2cov(L, S)$, which is equal to 33.3.

Generally speaking, this observation holds true for all anomalies: The difference

¹²Please note that all variances and covariances are multiplied by 10,000 for easy readability. Those estimators being based on squared returns tend to include multiple zeros in the first decimal places.

¹³Hence why we find that CF shocks account 86% of the HML portfolio variance, and DR shocks for 5%.

Table IV - Detailed portfolio variance decomposition: Shows the detailed variance composition of long-minus-short portfolios for their CF and DR returns, as per equation (III.5). The first column is the total portfolio variance and is equal to the sum of three terms in the following columns as per equation (III.1). The individual portfolio variances can further be decomposed, such that the variance of the long leg, $\text{Var}(L)$, is equal to the sum $\sum \text{Var}(r_i^L)$ and $\sum 2\text{Cov}(r_{j,k}^L)$. The first sum of terms corresponds to total individual firm variances, and the second is the sum of cross-covariances. Similarly, $\text{Var}(S) = \sum \text{Var}(r_i^S) + \sum 2\text{Cov}(r_{j,k}^S)$. In panel A, we report this decomposition for 1,000 randomly simulated portfolios. We denote the estimated means and standard deviations for each component $\hat{\mu}$ and $\hat{\sigma}$, respectively. In panel B, we report the calculated components for the five anomaly portfolios. We report the difference with the above simulations using a standard t-stat formula and report those in brackets. Significant differences are marked by “***”, “**” and “*” for the 1%, 5% and 10% confidence levels, respectively.

Panel A: Random Long-minus-Short portfolio simulations									
		Var(LMS)	-2cov(L,S)	Var(L)	Var(S)	$\sum \text{Var}(r_i^L)$	$\sum \text{Var}(r_i^S)$	$\sum 2\text{Cov}(r_{j,k}^L)$	$\sum 2\text{Cov}(r_{j,k}^S)$
Simulations	$\hat{\mu}$ CF	5.8	-15.3	10.5	10.5	3.0	3.0	7.5	7.6
	$\hat{\sigma}$ CF	(1.8)	(2.1)	(1.9)	(1.9)	(1.2)	(1.2)	(1.8)	(1.9)
	$\hat{\mu}$ DR	0.4	-44.3	22.3	22.3	1.0	1.0	21.3	21.3
	$\hat{\sigma}$ DR	(0.6)	(0.9)	(0.8)	(0.8)	(0.6)	(0.6)	(0.8)	(0.8)
Panel B: Anomaly portfolios									
		Var(LMS)	-2cov(L,S)	Var(L)	Var(S)	$\sum \text{Var}(r_i^L)$	$\sum \text{Var}(r_i^S)$	$\sum 2\text{Cov}(r_{j,k}^L)$	$\sum 2\text{Cov}(r_{j,k}^S)$
Book-to-Market	μ_0 CF	30.2***	-9.2***	26.1***	13.3	3.3	2.4	22.8***	10.9*
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[13.81]	[2.81]	[8.36]	[1.45]	[0.29]	[-0.47]	[8.33]	[1.78]
Size	μ_0 DR	1.8**	-42.5*	22.4	21.9	1.3	0.7	21.1	21.2
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[2.53]	[1.95]	[0.08]	[-0.48]	[0.51]	[-0.41]	[-0.27]	[-0.21]
Profitability	μ_0 CF	20.7***	-17.2	30.2***	7.7	1.5	1.0	28.7***	6.7
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[8.48]	[-0.89]	[10.54]	[-1.49]	[-1.25]	[-1.65]	[11.57]	[-0.45]
Investment	μ_0 DR	0.8	-44.6	23.3	22.1	0.2	0.3	23.1**	21.7
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[0.70]	[-0.41]	[1.21]	[-0.28]	[-1.37]	[-1.12]	[2.19]	[0.49]
Momentum	μ_0 CF	17.2***	-24.2***	22.9***	18.5***	3.4	3.3	19.5***	15.2***
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[6.45]	[-4.19]	[6.63]	[4.17]	[0.33]	[0.25]	[6.54]	[4.10]
Profitability	μ_0 DR	1.1	-45.4	23.6	22.8	1.2	0.8	22.4	22.0
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[1.24]	[-1.23]	[1.56]	[0.63]	[0.34]	[-0.32]	[1.36]	[0.87]
Investment	μ_0 CF	12.6***	-15.8	13.1	15.3**	2.7	4.0	10.4	11.3**
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[3.84]	[-0.27]	[1.40]	[2.49]	[-0.21]	[0.82]	[1.57]	[2.02]
Momentum	μ_0 DR	0.7	-46.3**	22.9	24.1**	0.9	1.1	21.9	23.0**
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[0.48]	[-2.27]	[0.64]	[2.15]	[-0.09]	[0.17]	[0.72]	[2.08]
Profitability	μ_0 CF	48.3***	-3.5***	16.8***	35.0***	3.5	5.0	13.3***	29.9***
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[24.09]	[5.51]	[3.38]	[12.83]	[0.43]	[1.70]	[3.17]	[12.02]
Investment	μ_0 DR	1.8**	-44.8	23.0	23.6	0.9	1.0	22.1	22.6
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[2.48]	[-0.61]	[0.77]	[1.59]	[-0.21]	[0.06]	[0.93]	[1.57]

between portfolio CF and DR variance, $Var(LMS)$, is either similar in magnitude or even smaller than the covariance term $-2cov(L, S)$. Moreover, this difference in covariance between portfolios for CF and DR is generally not significantly different from the simulated portfolios. Even in the cases where this difference with random portfolios is somewhat statistically significant, it does not change our observation that LMS portfolio variance, even for anomalies, is driven by CF shocks because of a DR covariance that is consistently orders of magnitude larger than CF covariance at the portfolio level.

This result again casts doubt on the idea that it is sufficient to measure total portfolio CF and DR variance to uncover CF and DR forces which are specific to anomaly returns. Fortunately, based on equation (III.5), we can re-visit this question by measuring comovement between stocks in different anomaly portfolios directly. In hypothesis 2 of section II, we proposed that it is possible to measure commonalities in CF and DR news which are specific to anomaly characteristics. We achieve this by estimating the sum of cross-covariances terms, i.e., $\sum 2Cov(r_{j,k}^L)$ and $\sum 2Cov(r_{j,k}^S)$ from equation (III.5). We also call this sum of cross-covariances “comovement” between stocks. We report those estimates in the last two columns of table IV.

Our key finding, is that stocks sharing certain anomaly characteristics are subject to CF comovements that significantly exceed CF comovements between the average stocks. On the other hand, stocks sharing similar anomaly characteristics do not exhibit DR commonalities that are statistically different from random. This aligns with the idea that anomaly returns are mainly driven by updates to CF expectations.

We also find significant differences between long and short portfolios of the same anomaly. Value, small, and loser stocks all exhibit CF comovement that significantly exceeds the one observed in random portfolios. Their comovement estimate is between 8.33 and 12.02 standard deviations higher than the average of the simulations. This corresponds to a sum of cross-covariances that are between three and four times larger than for average stocks. However, in the short leg counterparts, the sum of cross-covariances is either much less (i.e., see growth and loser stocks) or not significantly different (i.e., see large-cap stocks) from random.

For the profitability anomaly, CF comovement is also significantly more important than random, but a bit less so than for the three portfolios just mentioned above. This observation is true for both robust (i.e., long-leg) and weak (i.e., short-leg) profitability firms. Investment is the anomaly where CF comovement is the weakest. CF comovement is not significantly higher for conservative (long-leg) stocks, and barely

so for aggressive (short-leg) stocks.

Overall, all those results support the notion that anomaly returns arise from commonalities in CF movements across stocks with similar characteristics. Indeed, we observe that DR comovement in anomaly portfolios is not significantly greater than random. The only exceptions are small-cap and aggressive firms, but the sum of cross-covariance terms is only a few percentage points higher than in the simulations, and the significance level still fails to reach the 1% confidence level.

C. Conditional variance decomposition

Several theories seeking to explain anomaly returns predict heterogeneous risk exposures of CF and DR shocks to aggregate CF and DR news. As mentioned in the introduction, theories such as those proposed by Zhang (2005) or Fama and French (1996) suppose that value stocks face a risk that materializes during adverse CF periods, such as recessions. Others, such as Lettau and Wachter (2007) make the hypothesis that cash flows of growth stocks will be particularly sensitive to aggregate discount rate changes.

The third hypothesis of section II aims at directly testing this kind of hypotheses by asking the following question: when does CF and DR comovement spike? The underlying intuition is that firms exposed to a same risk factor should experience increases in comovement at times when the shared risk materializes. If a set of firms have cash flows that are particularly sensitive to specific market conditions, we would predict to see the sum of CF shock cross-covariances, i.e., $\sum 2Cov(r_{j,k})$, spike during those times.

Empirically, we seek to test those hypotheses in tables V and VI. Campbell et al. (2010) estimate CF and DR betas of value and growth firms with respect to market CF and DR shocks. In the same spirit, we seek to measure comovement as a function of macro variables that relate to the business cycle and aggregate CF events on the one hand (such as NBER recession dummies), and as a function of state variables that reflect aggregate risk aversion on the other (such as investor sentiment).

In table V we report comovement among anomaly portfolios conditioned on different aggregate cash flow events. We choose one-year market CF shocks, recession and expansion dummies, and 1-year GDP growth as proxies for changes in aggregate CF states. Except for the recession dummies, we split each variable into three states, which are separated by the 33rd and 66th percentiles of the variable's time series. Hence, we regroup all observations in either the low, medium, or high state of each conditioning macro variable.

Table V — Conditional portfolio CF and DR comovement – Aggregate CF updates: This table reports the sum of cross-covariances (the last two terms of equation (III.5)) for randomly simulated portfolios and for the long and short legs of the three main anomaly portfolio characteristics identified above as having the strongest CF/DR drivers (i.e., book-to-market, size, and momentum). The figure capturing the sum of cross-covariances is equivalent to the last two columns reported in table IV, except that the computation is reported for the subset of dates where the conditioning macro variable is either low (below its 33th percentile), medium (between its 33th and 66th percentile) or high (above the 66th percentile). Recession and Expansion dates correspond to the ones defined by the National Bureau for Economic Research (NBER). Standard errors for the simulations are reported in parentheses. T-stat for differences with simulations are reported in brackets and computed as in table IV.

Panel A: Portfolio CF conditioned on CF macro variables									
		CF^{agg} Low	CF^{agg} Med	CF^{agg} High	Recession	Expansion	1Y- Δ GDP Low	1Y- Δ GDP Med	1Y- Δ GDP High
Simulations	$\hat{\mu}$ CF	10.62	4.31	4.35	14.64	6.07	8.82	5.73	7.63
	$\hat{\sigma}$ CF	(4.13)	(1.32)	(2.53)	(6.72)	(1.25)	(2.93)	(2.49)	(2.74)
Value	μ_0 CF	39.2***	13.0***	9.5**	64.0***	16.4***	35.2***	13.2***	15.7***
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[6.92]	[6.58]	[2.02]	[7.35]	[8.27]	[9.01]	[3.00]	[2.96]
Growth	μ_0 CF	8.5	9.3***	12.7***	18.4	9.6***	8.5	13.2***	11.7
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[-0.50]	[3.75]	[3.30]	[0.56]	[2.85]	[-0.10]	[3.00]	[1.50]
Small	μ_0 CF	43.2***	23.4***	16.8***	74.7***	21.6***	39.6***	20.6***	24.1***
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[7.90]	[14.44]	[4.92]	[8.95]	[12.39]	[10.49]	[5.97]	[5.99]
Big	μ_0 CF	8.4	4.2	4.3	10.6	5.7	7.0	6.4	6.3
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[-0.53]	[-0.07]	[-0.03]	[-0.61]	[-0.29]	[-0.63]	[0.26]	[-0.48]
Winner	μ_0 CF	17.5*	10.3***	10.0**	23.2	11.1***	14.8**	10.3*	14.6**
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[1.67]	[4.56]	[2.21]	[1.27]	[4.03]	[2.05]	[1.82]	[2.54]
Loser	μ_0 CF	53.1***	20.3***	12.6***	96.4***	19.5***	46.8***	14.5***	28.2***
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[10.28]	[12.07]	[3.26]	[12.17]	[10.70]	[12.94]	[3.53]	[7.48]
Panel B: Portfolio DR conditioned on CF macro variables									
		CF^{agg} Low	CF^{agg} Med	CF^{agg} High	Recession	Expansion	1Y- Δ GDP Low	1Y- Δ GDP Med	1Y- Δ GDP High
Simulations	$\hat{\mu}$ DR	36.38	14.62	12.19	38.22	18.47	21.83	20.15	22.69
	$\hat{\sigma}$ DR	(1.17)	(0.84)	(0.65)	(2.30)	(0.52)	(0.94)	(1.00)	(0.72)
Value	μ_0 DR	36.7	15.0	11.8	39.8	17.9	23.5*	19.3	21.8
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[0.26]	[0.46]	[-0.66]	[0.67]	[-1.01]	[1.82]	[-0.84]	[-1.17]
Growth	μ_0 DR	34.9	15.0	12.6	37.2	18.4	21.3	20.1	22.1
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[-1.25]	[0.45]	[0.64]	[-0.44]	[-0.06]	[-0.51]	[-0.01]	[-0.81]
Small	μ_0 DR	39.6***	15.7	12.9	45.5***	19.4*	23.6*	21.1	24.9***
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[2.75]	[1.31]	[1.07]	[3.18]	[1.73]	[1.91]	[0.93]	[3.12]
Big	μ_0 DR	37.4	14.9	12.4	38.1	18.9	22.1	20.8	23.0
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[0.84]	[0.35]	[0.39]	[-0.06]	[0.88]	[0.30]	[0.68]	[0.37]
Winner	μ_0 DR	38.3	15.2	11.9	36.1	19.5**	21.9	21.1	23.6
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[1.63]	[0.71]	[-0.43]	[-0.94]	[1.96]	[0.08]	[0.95]	[1.31]
Loser	μ_0 DR	38.6*	16.0	13.4*	45.6***	18.8	24.6***	21.4	22.8
	$\frac{\mu_0 - \hat{\mu}}{\sigma}$	[1.86]	[1.61]	[1.86]	[3.22]	[0.65]	[2.97]	[1.28]	[0.19]

In panel A we first report the sum of cross-covariances, $\sum 2Cov(r_{j,k})$, for 1,000 simulated portfolios. We find that on average, CF comovement across firms is more than twice as large during recession (10.62) and negative aggregate CF shock (14.64) times than otherwise (4.31 and 6.07). Overall, DR comovement also increase in hard times (panel B), going from 18.47 in expansion to 38.22 in recession periods. This sort of pattern would be expected: it is hard and costly to hedge against aggregate market declines, such that all stock prices tend to fall together in hard times.

All anomaly portfolios that tended to exhibit significant comovement in table IV, keep showing significant comovement across all macroeconomic states. Nonetheless, we observe differences across conditioning variables. Value stocks comove more during the worst periods of aggregate CF shocks (39.2) than during positive market CF movements (9.5). The difference with the corresponding simulated portfolios is highly significant during negative aggregate CF periods (t-stat = 6.92), but much weaker when aggregate CF shocks are positive (t-stat = 2.02). We observe similar patterns for recessions: the sum of CF cross-covariances is four times larger (64.0 vs. 16.4). Cross-covariances in the average random portfolio only increase by 2.4 times when going from expansion to recession. Similarly, we also find that in the periods of lowest GDP growth, value stocks tend to comove more together than in strong economic growth periods. Such patterns are loosely consistent with the idea that value stocks are exposed to a common risk linked to aggregate CF shocks, which in turn triggers greater comovement during those periods.

In other anomaly portfolios, CF comovement seems to be less influenced by aggregate CF movements. Small firms seem to be similar to value stocks, as they tend to comove more during periods of negative aggregate CF updates. Nonetheless, their comovement patterns are so strong across all macroeconomic states, that differences are less evident than for value firms. Robust firms also present consistently higher CF comovement, which remains proportionally similar to the random portfolios (about three times higher). We find only weak, or no patterns, for big, winner, conservative and aggressive stocks, on which it seems preferable not to draw conclusions. However, loser stocks present patterns similar to value firms: proportionally, their comovement increases more than for average stocks during hard times (i.e., during periods of recession, negative aggregate CF and low GDP growth).

When looking at cross-covariances in portfolio DR shocks (panel B), we find that the differences across anomaly portfolios are much smaller. At most, we observe that the sum of cross-covariances is a few percentage points higher than in the simulations. Notice that the reported variability from the simulations is also much lower than

for CF comovement. Only small and aggressive (and to some extent robust) stocks seem to have higher DR comovement than average. Conditional patterns are not clearly apparent, although DR comovement appears to increase a bit further in hard times.

Table VI — Conditional portfolio CF and DR comovement – Aggregate DR updates: This table reports the sum of cross-covariances as in table V above. Conditioning variables are again split by the 33rd percentile.

Panel A: Portfolio CF conditioned on DR macro variables										
		DR^{agg} Low	DR^{agg} Med	DR^{agg} High	Sentiment Low	Sentiment Med	Sentiment High	Spread Low	Spread Med	Spread High
Simulations	$\hat{\mu}$ CF $\hat{\sigma}$ CF	6.25 (2.44)	5.40 (1.90)	10.22 (3.13)	9.52 (2.74)	6.39 (2.65)	6.12 (2.96)	5.60 (2.99)	6.76 (1.92)	9.72 (2.52)
Value	μ_0 CF $\frac{\mu_0 - \hat{\mu}}{\sigma}$	22.6*** [6.69]	13.2*** [4.09]	28.8*** [5.92]	28.4*** [6.88]	21.0*** [5.50]	15.8*** [3.27]	16.3*** [3.60]	15.2*** [4.40]	32.7*** [9.15]
Growth	μ_0 CF $\frac{\mu_0 - \hat{\mu}}{\sigma}$	15.2*** [3.65]	9.6** [2.23]	8.5 [-0.56]	7.6 [-0.71]	8.0 [0.62]	17.9*** [4.00]	7.0 [0.45]	8.8 [1.05]	17.0*** [2.90]
Small	μ_0 CF $\frac{\mu_0 - \hat{\mu}}{\sigma}$	31.3*** [10.25]	18.2*** [6.74]	35.7*** [8.12]	21.7*** [4.43]	31.1*** [9.31]	32.4*** [8.89]	17.1*** [3.85]	28.9*** [11.51]	39.3*** [11.77]
Big	μ_0 CF $\frac{\mu_0 - \hat{\mu}}{\sigma}$	5.4 [-0.36]	5.6 [0.09]	9.0 [-0.40]	8.9 [-0.23]	5.8 [-0.23]	5.4 [-0.23]	6.0 [0.12]	5.7 [-0.56]	7.8 [-0.76]
Winner	μ_0 CF $\frac{\mu_0 - \hat{\mu}}{\sigma}$	16.4*** [4.16]	8.1 [1.43]	14.8 [1.45]	14.9* [1.95]	11.3* [1.84]	15.6*** [3.21]	8.3 [0.90]	10.1* [1.72]	21.1*** [4.51]
Loser	μ_0 CF $\frac{\mu_0 - \hat{\mu}}{\sigma}$	29.8*** [9.66]	16.1*** [5.61]	43.3*** [10.55]	27.6*** [6.61]	37.3*** [11.64]	25.4*** [6.51]	18.8*** [4.41]	25.4*** [9.69]	45.0*** [14.02]
Panel B: Portfolio DR conditioned on DR macro variables										
		DR^{agg} Low	DR^{agg} Med	DR^{agg} High	Sentiment Low	Sentiment Med	Sentiment High	Spread Low	Spread Med	Spread High
Simulations	$\hat{\mu}$ DR $\hat{\sigma}$ DR	12.12 (1.04)	18.23 (0.81)	21.93 (1.01)	20.72 (0.90)	19.10 (1.03)	23.93 (1.23)	13.25 (1.06)	17.14 (0.73)	32.03 (1.26)
Value	μ_0 DR $\frac{\mu_0 - \hat{\mu}}{\sigma}$	11.4 [-0.72]	16.7* [-1.86]	24.8*** [2.88]	23.7*** [3.27]	17.8 [-1.26]	22.7 [-1.02]	12.4 [-0.80]	18.4* [1.79]	32.5 [0.33]
Growth	μ_0 DR $\frac{\mu_0 - \hat{\mu}}{\sigma}$	12.8 [0.64]	18.5 [0.35]	20.8 [-1.16]	19.7 [-1.11]	18.9 [-0.17]	24.3 [0.30]	13.4 [0.09]	16.4 [-1.00]	32.1 [0.05]
Small	μ_0 DR $\frac{\mu_0 - \hat{\mu}}{\sigma}$	13.3 [1.18]	18.6 [0.46]	25.4*** [3.48]	23.0** [2.49]	18.9 [-0.22]	27.3*** [2.73]	13.2 [-0.05]	19.1*** [2.72]	35.6*** [2.85]
Growth	μ_0 DR $\frac{\mu_0 - \hat{\mu}}{\sigma}$	12.7 [0.53]	18.9 [0.82]	22.4 [0.46]	21.0 [0.31]	20.0 [0.88]	24.2 [0.26]	13.8 [0.51]	17.3 [0.17]	33.0 [0.76]
Winner	μ_0 DR $\frac{\mu_0 - \hat{\mu}}{\sigma}$	14.2** [2.02]	18.1 [-0.11]	21.5 [-0.41]	21.4 [0.72]	20.8* [1.68]	23.9 [0.01]	14.1 [0.84]	17.5 [0.45]	32.9 [0.70]
Loser	μ_0 DR $\frac{\mu_0 - \hat{\mu}}{\sigma}$	11.1 [-0.96]	18.7 [0.56]	26.3*** [4.38]	22.5* [1.93]	19.6 [0.45]	26.1* [1.79]	13.4 [0.14]	18.1 [1.27]	35.7*** [2.95]

In table VI, we condition comovement on variables that proxy for aggregate DR movements. Patterns are notably different. In panel A, we observe that growth stocks exhibit patterns that go in the opposite direction of their long-leg portfolio counterpart. We find stronger comovement when proxies for time-varying risk aversion are high. For example, when investor sentiment is high, comovement is about three times higher than average within those stocks. In periods of high credit spread, comovement among those firms is observed to be double above baseline and

significant. We also condition on aggregate DR shocks and find that growth stocks covary more in periods when market discount rates decrease.

Overall, when considering the results from table V, our results seem consistent with the findings of Campbell et al. (2010): CF shocks of value stocks are particularly exposed to sources of systematic CF risk (leading to stronger comovement during market-wide CF stress), whereas growth stock CF covariance is more heavily influenced by sources of systematic DR risk (as proxied by our variables motivated by aggregate risk aversion).

Indeed, when looking at value stocks we find their CF shocks to comove more than average in all conditional DR states, without consistent patterns. The only exception would be for investor sentiment, which decreases with value stock's CF comovement. We also find in panel A that growth stocks covary more during expansion and positive CF periods. Such observations tend to be consistent with the negative market CF and DR correlation that we document in table II.

Winner stocks exhibit similar patterns as growth stocks. Comovement across those firms increases with investor sentiment and the credit spread, and is high when market DR shocks are low. Furthermore, in table V, winner and growth stocks already followed similar patterns. It is worth noting that we find similar CF comovement patterns in the opposite legs of the book-to-market and momentum anomalies. Therefore, comovement in value and momentum portfolios seem to follow opposite patterns. This type of pattern is reminiscent of Asness et al. (2013) and Daniel and Moskowitz (2016), who find that momentum (WML) and book-to-market (HML) portfolios are negatively correlated. In particular, the betas of HML and WML to the market factor are of opposite sign during times of market decline. It also aligns with the overall theme of chapter IV of this thesis: value and momentum anomalies appear to be opposite sides of a same coin. The evidence provided here is roughly consistent with the idea that the two anomalies will earn premia in different states of the economy, due to their opposite comovement patterns; which might be due opposite systematic CF and DR risk exposures.

For the other anomalies we study, patterns are less clear when conditioning on proxies for aggregate risk aversion. Confirming previous results, small firm CF comovement is consistently higher, but we find no conditional difference. Similarly, big stocks do not exhibit stronger comovement than the random simulated portfolios, no matter the conditioning variable. Robust profitability firm CF's consistently comove more than average. Their sum of CF cross-covariances increases slightly more during periods of high sentiment and credit spread. Patterns are less intuitive

and somewhat surprising for weak profitability firms, as their sum of CF cross-covariances increases more during periods of low sentiment and high credit spread. We find less comovement for investment portfolios than for the other anomalies. Nonetheless, Conservative stocks seem to exhibit a comovement that increases with high investor sentiment and large credit spreads.

Finally, in panel B of table VI we report how DR comovement evolves as a function of proxies for aggregate risk aversion. The sum of DR cross-covariances is again much closer to the average simulated portfolio, as in panel B of table V, with few significant differences. Nonetheless, loser, weak, small and value stocks tend to exhibit DR comovement that increases a bit in periods of high investor sentiment and high credit spread.

V. Discussion

In this section, we discuss which theoretical frameworks best fit with our results, which models are unlikely to be reconciled with our findings, suggest avenues for future research, and link our observations to other parts of this thesis. Overall, many of our results align with the findings of Lochstoer and Tetlock (2020). However, our detailed decomposition allows for significant new findings providing the following insights.

First, our results suggest that firms with similar anomaly characteristics share common CF shocks that significantly exceed those of random pairs of stocks. This would imply that theories relying on cross-sectional heterogeneity to explain anomalies would fit with our observations. Rational models such as Berk et al. (1999), Kogan and Papanikolaou (2013), or Hou et al. (2015), which predict cross-sectional differences in expected returns depending on firms' investment opportunities (which represent CF shocks) all align with this line of reasoning. Behavioral models, relying on cross-sectional differences in investor misvaluations leading to CF shock overreaction, also predict that CF shocks are the dominant driver of anomalies. Some of those models include Daniel et al. (1998), Barberis et al. (1998), and Hong and Stein (1999).

Second, our empirical results support stories relying on cross-sectional heterogeneity in firms' systematic cash-flow risk. For example, stories where value firms are more sensitive to market contraction due to being closer to bankruptcy (Fama and French, 1996) or because of their high level of assets in place (Zhang, 2005) are consistent with the high conditional comovement patterns we document. They also

fit naturally with models such as Campbell and Vuolteenhao (2004), or evidence in Cohen et al. (2009), where value firms have a particularly high sensitivity to aggregate CF shocks, which commands a higher risk premium than aggregate DR exposure due to intertemporal concerns.

However, this is not true for all anomaly portfolios. For example, the model of Lettau and Wachter (2007) predicts that growth stocks comove more with systematic DR shocks. Our framework separating long and short legs of the anomalies allow us to reconcile such theories with empirical observations. Growth and momentum both seem to have a higher exposure to factors related to time-varying risk-aversion. The simultaneous exposures of long and short legs of the anomaly to different sources of risk (i.e., systematic CF and DR) closely relate to the work of Santos and Veronesi (2010). They also coincide with the empirical findings of Campbell et al. (2010), who show that value firms load more strongly on systematic CF shocks, while CF shocks of growth firms are more exposed to sources of DR risk.

In table II we find that across all anomalies, CF and DR shocks are negatively correlated. This aligns with theories of overreaction, such as Barberis et al. (1998) or Hong and Stein (1999), but also with models where positive CF shocks cause a decrease in risk, and thus a negative DR update, as in Kogan and Papanikolaou (2013). Those results are consistent with Lochstoer and Tetlock (2020) and are robust across all except one of our specifications.

Our results suggest that anomaly returns are not driven by heterogeneous DR shocks across anomalies. This implies that theories relying on positive CF-DR feedback effects, such as noise trader models (De Long et al., 1990), do not match with our empirical results.

chapter IV of this thesis focuses on the joint dynamics of value and momentum firms. Those anomalies exhibit consistently opposite patterns in their CF and DR news. The opposition between the anomalies is also documented in the literature. Asness et al. (2013) and Daniel and Moskowitz (2016), find that momentum (WML) and book-to-market (HML) portfolios are negatively correlated. In particular, the betas of HML and WML to the market factor are of opposite sign during times of market decline.

Our findings might provide a useful framework for an explanation that would seek to tackle this appearing puzzling link between value and momentum stocks. The long legs of the two anomalies face a shared source of risk that occurs at different times: value CF shocks comove more during negative aggregate CF events, while winner CF comovement appears to be dependent on aggregate risk aversion. Simi-

larly, the short legs (i.e., growth and loser stocks) also exhibit spikes in comovement driven by opposite systematic events. If the two types of systematic events driving stock comovement are negatively correlated (as the negative correlation between aggregate CF and DR suggests), then the negative comovement between the value and momentum anomalies could be driven by an underlying systematic risk factor. This observation calls for future research to reconcile both anomalies through a rational risk-based channel.

Finally, we suggest some avenues of future research. First, our empirical evidence is strongest for anomalies with well-established theories that rely on cross-sectional heterogeneity in comovement of CF shocks. It would be interesting to test how other anomalies fare in this framework, and if those lacking theoretical justifications to back up their premia struggle to distinguish themselves from random portfolios. It could prove useful to anchor the existing “factor zoo” (Cochrane, 2011) to economic risk factors.

Second, alternative measure for systematic CF and DR risk might be used. One interesting alternative might come from the work of Bybee et al. (2021): They propose a set of news attention factors related to a wide array of topics based on news. Certain anomalies might be particularly sensitive to certain news, and conditioning on such variables might provide valuable insights into the underlying risk factors of anomalies.

VI. Conclusion

In the asset pricing literature, multiple theories propose explanations for the origin of predictable anomaly returns. They often result in competing predictions about the cross-sectional dynamics of CF and DR shocks. With the aim of guiding theories seeking to explain cross-sectional patterns linked to anomaly characteristics, Lochstoer and Tetlock (2020) show that long-minus-short (LMS) portfolios sorted on such characteristics are mainly driven by CF shocks. This paper invites to be cautious before drawing this conclusion: the variance of returns of any well diversified LMS portfolio appears to systematically be driven by revisions to cash flow expectations. In a simulation exercise where we use 35,000 randomly assigned long and short portfolios, we find that the variance of the resulting LMS portfolios is consistently driven by CF shocks (84%) and that DR shocks account for a much smaller portion of unexpected return variance (6%). This result appears to hold true across multiple VAR specifications motivated by previous literature.

We decompose LMS portfolios' variance into its subcomponents and find that the large CF contribution to their variance is due to the strong positive correlation in DR shocks between any two well diversified portfolios. This high DR correlation between the long and short portfolios reflects the DR shock dominance as the main driver of unexpected market returns. Motivated by this result, we further decompose portfolio variance into its individual terms, which allows us to highlight specific comovement across firms with similar anomaly characteristics. We find that most anomalies share CF commonalities that far exceed those found among random groups of stocks. Value, small, and loser stocks are the anomaly portfolios among which we find the strongest levels of CF comovement. On the other hand, we find that comovement in DR news does not significantly increase above random levels.

Finally, we test for conditional effects: Does comovement across stocks with similar anomaly characteristics change across different states of the market, and thus with varying aggregate CF and DR news? We find that CF shocks of value and loser stocks tend to comove more strongly together during periods of negative aggregate CF shocks, such as recessions. On the other hand, growth and winner stocks exhibit increases in CF comovement at times when proxies for aggregate risk aversion are elevated, for example when investor sentiment or credit spread levels are high.

Those findings can help in selecting theories explaining anomaly returns. Overall, they seem consistent with the idea that anomaly returns are mainly driven by cross-sectional differences in CF shocks, and that for certain anomalies, those CF comovements are dependent on macroeconomic conditions, either linked to aggregate CF or DR shocks. Those findings are in the same spirit as Campbell et al. (2010), and highlight that drivers of anomalies vary from one portfolio to the next, suggesting that separate economic theories for each specific anomaly might be needed. Our results show that it is possible for a same anomaly to reconcile theories relying on exposures to aggregate CF shocks, such as Zhang (2005) (for value stocks), as well as theories relying on exposures to systematic DR updates, such as Lettau and Wachter (2007) (for growth stocks).

Chapter IV

Dynamics of Cash Flow and Discount Rate News for Value and Momentum Stocks

MORENO NICOLAS[‡]

APRIL 2022

ABSTRACT

We find that a firm's past cash flow (CF) and discount rate (DR) news predict future expected returns: sustained increases in DR lead to higher returns, as do (surprisingly) positive CF shocks over the past quarters, consistent with a continuation pattern. Those patterns also appear in the value and momentum anomalies. The former is driven by past DR increases, whereas the latter is subject to positive CF news and a decrease in DR. Contrasting patterns further arise when comparing reactions of value and momentum stocks to their media tone coverage. Bad news (as for value firms) imply positive DR revisions, whereas good media coverage (from which momentum stocks benefit) trigger positive CF shocks and negative DR revisions. Our results shed new light on the previously documented negative correlation between value and momentum, and call for a joint explanation of both anomalies.

Keywords: Value, Momentum, Cash-flow and Discount-rate Revisions, News, Big Data

[‡]This paper appears as a single-authored chapter for my PhD dissertation at HEC Liège. It would however never have existed without the tremendous help of Marie Lambert and Denada Ibrushi. Both their guidance and contribution have been essential, and therefore I am hopeful that they may rejoin as co-authors, should this paper proceed in the publication process. At this stage, and as a chapter of my thesis, this paper is not intended for circulation.

I. Introduction

In the years prior to the portfolio holding period, value firms tend to be “loser” stocks, i.e., firms that were subject to a decrease in price¹. Gerakos and Linnainmaa (2018) show that the premium earned by value firms is specific to this change in market equity over the past five years. This study aims at understanding the underlying drivers of those price changes leading to the cross-sectional differences in returns. By estimating firm-specific changes in expected cash-flows (CF) and discount rates (DR), this study finds that the price decrease of value stocks is driven by more negative CF and more positive DR shocks compared to the average firm. Negative CF revisions account for 70% of the difference in unexpected returns, and upwards revisions of future discount rates constitute the remaining 30%.

Momentum stocks follow an opposite pattern. By construction, those firms are selected to have the largest price increase, typically over the past year. We find that the price increase of winner stocks is due to significant positive CF shocks over the past four quarters. Those positive CF shocks make up for 85% of the unexpected return differential with the average stock. Negative DR shocks drive an extra 15% of the relative price decrease.

This study aims at documenting a salient contrast between value and momentum firms. The value anomaly manifests itself as a consistent price reversal. In particular, prior to the portfolio holding period, value stocks are subject to consistently higher DR news² which initially drive their price down but later translate into higher expected returns.

On the other hand, momentum captures a continuation phenomenon in the quarter following the prior twelve months of portfolio formation period. The price of high momentum firms continues to move in the direction of their past positive CF shocks. However, those firms do not cumulate high DR shocks as do value stocks, and subsequently momentum firms have lower expected returns as they experience a price reversal in the quarters and years after the positive CF shock, loosely consistent for example with Hong and Stein (1999) and Bloomfield et al. (2009), where uninformed traders drive long term price reversals.

We also document that the two anomalies also differ in their media coverage: value stocks tend to be subject to worse news *tone* than average, whereas momen-

¹See for example De Bondt and Thaler (1987), or Fama and French (1996).

²In this study we refer to “news” when talking in investor expectations about firm CF or DR. Equivalently, we might also refer to it as a news shock, CF shock, DR shock, CF update, or DR update. In contrast, when we talk about news *tone*, we refer to the polarity measure of media content as described in chapter II.

tum firms benefit from positive media coverage in the quarters leading to the holding period. Those media effects correlate with the unexpected returns patterns, highlighting the tight link between the two types of news: good (bad) news imply lower (higher) DR, and positive (negative) CF shocks. Below, we document how the literature finds pervasive inverse correlations between the two anomalies, to set the framework on how our can improve our understanding of those contrasts.

Indeed, the literature documents several oppositions between the two anomalies. Asness, Moskowitz, and Pedersen (2013) for example, find that the value and momentum factors are consistently and strongly negatively correlated, both across assets and internationally across markets. Daniel and Moskowitz (2016) also document this negative correlation and find that the premiums evolve in opposite directions during stressful market events. The correlation with the excess market return of the high-minus-low (HML) book-to-market portfolio increases following adverse market conditions. In contrast, the winner-minus-loser (WML) momentum portfolio starts correlating negatively with the market.

The challenge for a rational explanation is to tie both continuation and reversal patterns to priced risk, ideally in a joint setting (Cochrane, 2011); and provide an economic justification for why this risk should command a premium. A theory explaining value and momentum returns should simultaneously give a rationale for differences in mean return and justify why there is comovement.

Asness et al. (2013) propose to link both value and momentum to funding liquidity risk (see for example Pástor and Stambaugh (2003) and Sadka (2006)). Momentum firms, which represent the currently popular trades, have positive exposure to liquidity risk. A liquidity shock might push everyone towards the exit (Pedersen, 2009), forcing them to find liquidity in the contrarian view, i.e., value stocks. While value firms do indeed load negatively on liquidity risk, this factor fails to capture the full magnitude of the momentum and value anomalies.

Vayanos and Woolley (2013, 2012) propose an alternative explanation for the joint overreaction and reversal patterns of momentum and value based on fund flows, where an investor conflates cash-flow shocks with a manager's ability. Fama and French (1996) and Zhang (2005) are examples of economically motivated justifications for a positive premium on high book-to-market assets that account for both the differences in means as well as the correlation structure of returns. Nevertheless, Cochrane (2011) points out that such entirely rational and economically motivated explanations for the momentum anomaly may still be lacking.

On the other side of the spectrum, Asness et al. (2013) pose a high threshold

for behavioral explanations of the momentum anomaly, since they show that such theories will have to account for global momentum and value comovements, both across assets and markets.

Most explanations for momentum have some form of market friction or behavioral component as a starting point. Pedersen (2021) for example, provides a unifying framework for both momentum and reversal to fundamentals (i.e., a value effect) in a setting where different types of agents interact through social media. In this model, certain “stubborn” agents push back prices to fundamentals, but other agents get influenced by loud “fanatics” on social media, creating echo chambers and influencing other’s behavior. Bailey et al. (2018), Kuchler et al. (2020) and Cookson, Engelberg, and Mullins (2021) all provide examples how social media can shape agent’s investing behavior.

In behavioral models, value (reversal) and momentum (continuation) effects can arise without a link to fundamentals. Using sports bet as a clever setting where pricing does not depend on systematic risks or terminal values, Moskowitz (2021) shows that momentum arises in the form of delayed overreaction and that value effects emerge as long term reversals. Earlier models like those of Daniel, Hirshleifer, and Subrahmanyam (1998) or Hong and Stein (1999) assume that the irrationality of certain agents in the market is sufficient to give rise to momentum. Finally, survey evidence points out that investors form expectations by extrapolating past information (Greenwood and Shleifer, 2014), which can create overreaction and momentum (Barberis et al., 2015).

This paper aims at documenting the contrasting patterns in past CF and DR movements of value and momentum stocks. Conceptually, an increase in discount rates can mean one of two things: On one end of the spectrum, in a purely rational expectations setting, new information implies an unexpected increase in exposure to some priced risk factor(s) (Cochrane, 2011). On the other end of the spectrum, based on purely behavioral mechanisms, expected returns increase because of mispricing and prices which drop too low. We use the return decomposition framework of Campbell and Shiller (1988a) to estimate CF and DR shocks. However, this model is silent about the nature of contemporaneous changes in expected returns (rational or behavioral).

Nonetheless, a prediction of the risk-based framework for changes in DR implies that when a stock increases its exposure towards a risk factor, it will experience a positive revision in expected returns. The patterns for value firms are well in line with this prediction, as they gradually display increases in discount rates over

previous years. Furthermore, it is possible to back out a firm’s expected return component within the return decomposition framework by subtracting the unexpected return from the total realized return. Consistent with the prediction, value firms, or firms generally subject to big positive DR shocks over long periods (five to ten years), tend to exhibit significantly higher expected returns.

Our results for momentum stocks appear to be harder to fit into this framework. Their negative DR shocks, which only manifest in the four quarters of the portfolio formation period, should imply low subsequent expected returns. Indeed, we find that years after the holding period, the expected return component of winner stocks tends to be lower than average. Furthermore, it appears that stocks with the highest average CF shocks over the past four quarters consistently outperform those with the lowest CF revisions. A return continuation follows short-term upwards CF revisions for at least one more period, which seems to be the general pattern of momentum stocks. Those observations appear to be consistent with typical overreaction frameworks. However, potential limitations of the model must lead to caution when interpreting the results.

The baseline specification of this paper relies on the empirical approach of Lochstoer and Tetlock (2020) for the firm-level CF and DR return decomposition. It relies on price-momentum, book-to-market, and four other characteristics known to predict future returns³ to estimate the unexpected returns backed out from the Campbell and Shiller (1988a) decomposition. By construction, the model predicts that firms with high levels of positively priced characteristics (such as high momentum and high book-to-market) earn high expected returns during the holding period. To mitigate endogeneity concerns, we replicate the results using alternative decomposition frameworks that omit book-to-market and momentum. Revisions in expected returns must therefore capture changes in exposition to other sets of priced risk factors, not value and momentum directly⁴. Key results remain qualitatively similar to those initial robustness considerations: value firms still accumulate positive revisions in DR, while momentum stocks still present positive CF and negative DR shocks in

³Numerous so-called anomaly characteristics have been documented in top finance journals over the past three decades such that the hurdle required to consider a factor as meaningful and important needs to be high (Harvey et al., 2016; Mclean and Pontiff, 2016; Harvey and Liu, 2021). Thus, following this paper, the vector autoregression (VAR) model includes the logarithm of book-to-market for value, of twelve-month return excluding the most recent month for momentum, of 5-year change in market-cap for size, of revenue minus COGS, interest expenses and SG&A for profitability, of 5-year change in total assets for investment, and the past-period return, as explanatory variables.

⁴Note that alternatively, it may also be capturing mispricing and behavioral patterns, as discussed above.

the recent past.

This study relies on modeled revisions in investor expectations, i.e., CF and DR *News*. It is natural to assume that the financial media helps as an intermediary to present investors with information they can use for (or that reflects) updates in expectations about future firm cash-flows and discount rates. We use *tone* as an exogenous proxy of news⁵. The *tone* of a firm captures the average sentiment of the text within articles published by Thomson Reuters that relate to that particular company. For example, in the case of a firm subject to several good news with positive prospects (e.g. “exceeds expectations”, “signs new deal”, “performed well”), the text processing algorithm will probably assign a positive *tone* (tends to +1), whereas bad news (e.g. “layoffs”, “face litigation”, “ill-received product”) will get a negative *tone* (tends to -1).

An additional contribution of our paper is that positive (negative) *tone* correlates with positive (negative) CF shocks. Good (bad) news co-occur with negative (positive) revisions in expected return when it comes to discount rates. This can be interpreted in two possible ways: either good news imply a decrease in risk and a lowered exposure to priced risk factors, or good news imply lower expected returns due to an overreaction to the CF shock. On average, the impact of *tone* on CF is about five times larger than for *DR*.

News *tone* forecasts long-term changes in expected returns: Good (bad) news predicts low expected returns at least five years into the future. This suggests that the subsequent price reversal is either very slow or that the risk exposure has lastingly changed.

Furthermore, we find that CF shocks forecast future *tone*. If good news *tone* follows in the quarters after positive CF shocks, this can suggest that media reporting might either exhibit a delayed reaction or that past news heavily influences reporting over subsequent months.

Finally, *tone* provides additional insights into the reversal patterns of value and momentum stocks. Value firms are consistently subject to negative media coverage, even long before the portfolio holding period. Therefore, the bad news effect of value stocks can help as an additional piece of evidence in explaining their consistently high expected returns. Again, the opposite pattern plays out for momentum firms: winner firms start getting (very) good news only in the recent past, i.e., in the most recent three quarters. A reversal then follows this string of good news, as predicted by the above observations. News *tone* partly captures the differences in past DR and

⁵The news *tone* variable is extensively detailed in chapter III of the thesis. A literature review on the usage of proxies for news content is provided in chapter I

CF of value and momentum firms, and their differences in future expected returns. Nonetheless, the effects remain, and cross-sectional differences stand firm.

Overall, we document that the negative correlations between the value and momentum anomalies extend beyond what the literature had previously found. Not only do the value and momentum anomalies correlate negatively during stressful market conditions, but they are also driven by sharply different past news. This is true whether news are proxied through *tone* of media content or through revisions of investor expectations. The pervasiveness of those opposing patterns encourage us to echo previous calls from the literature: there is a need for a unified theory of value and momentum, which accounts for the different stylized facts.

This paper contrasts with chapter III of the dissertation by making the following different contributions. First, it shows what CF and DR news patterns lead to premia akin to value and momentum. Second, it focuses exclusively on value and momentum based on the recurrent contrasts documented in the literature and opposite observations that we make in this chapter. Third, it provides insights in long-term price dynamics of anomalies that significantly differ from random and from one another, and we are the first to our knowledge able to make this contribution about dynamics thanks to our quarterly news shock decomposition. And fourth, we relate CF and DR news to news *tone*, a first to our knowledge, and find that media content mainly impacts unexpected returns through the CF channel.

The rest of the paper is structured as follows. Section II discusses the data and the statistics of the estimated CF and DR news series. The results section is split into five parts: Section III.A presents the drivers of past price changes for value and momentum. Section III.C investigate how past CF and DR shocks forecast future returns. Sections III.E address robustness concerns Section IV investigates the correlation of media *tone* with unexpected returns. Finally, section V discusses theoretical frameworks that might reconcile our results and the joint patterns of value and momentum stocks, as well as potential limitations of our study. Section VI concludes.

II. Data: Quarterly CF and DR news series

As in chapter III we use the approximation of Campbell and Shiller (1988a) to define CF and DR shocks:

$$\begin{aligned}
 r_{t+1} - E_t r_{t+1} &\approx (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j} \\
 &\approx CF_{t+1} - DR_{t+1}
 \end{aligned} \tag{IV.1}$$

where E_t represents expectation at time t , r_t and d_t denote log-returns and dividends at time t respectively, and ρ is a parameter of linearization that depends on the long-term mean of the log dividend-price ratio. This equation suggests that higher unexpected returns are either due to higher than expected future cash flows, lower than expected future discount rates, or a combination of the two.

We detail in Chapter III, section III, how we estimate those components of unexpected returns. In essence, we follow the approach of Lochstoer and Tetlock (2020) to estimate firm-specific CF and DR shocks.

Our approach differs in that it uses data at quarterly instead of yearly frequency. This distinction is key for the design of our study: we investigate how CF and DR shocks of momentum firms evolve in the quarters prior to the portfolio holding period. Since the portfolio formation period is short (data from month $t-13$ up to $t-1$ is used to rank a stock along the momentum characteristic), being able to distinguish what happens with sufficient granularity in the months leading up to a stock becoming a winner or a loser stock is essential. To the best of our knowledge, we are the first to propose a firm-level return decomposition into CF and DR at quarterly level⁶.

The choice of data, which comes from the intersection of CRSP and Compustat, spans over the period from 1973 to 2020. The choice of variables, data availability, selection, and transformation are all detailed in chapter III. We estimate the VAR model and obtain the coefficient matrices Π from equations (III.9) and (III.10), and report them in table I of the chapter III.

With the estimated matrices Π from the VAR, it is possible to compute the DR and CF news series as detailed in equations (III.12) through (III.16). The parameter of linearization ρ is equal to $0.95^{3/12}$, which follows the discussion in chapter III

⁶Our approach is somewhat similar to Lyle and Wang (2015) who also rely on quarterly data. A non-exhaustive list of studies that estimate firm-specific CF and DR returns include Vuolteenaho (2002), Campbell et al. (2010), Chen et al. (2013), Maio and Santa-Clara (2015), all at yearly frequency.

and is thus adjusted to match quarterly frequency. Table I presents the estimated contributions to return variance of CF shocks, DR shocks, and their covariance. The first three lines present the drivers of firm-specific (i.e., the market adjusted), total, and aggregated (i.e., market) components of unexpected returns, respectively. At the firm level, CF shocks dominate, while updates in discount rates account for a large portion of market return variance. This is consistent with prior studies decomposing stock returns into CF and DR components (Vuolteenaho, 2002). Since the total return is the sum of the aggregated and market-adjusted components, both will influence the final CF and DR variances. Nonetheless, CF shocks still dominate total firm return variance (64%). DR shocks account for 11% of the variance, and the covariance component makes up the remaining 25%.

Given the focus on value and momentum, panel B reports unexpected return drivers of value and momentum portfolios. As in Lochstoer and Tetlock (2020) both the HML (value) and WML (momentum) portfolios are driven by CF shocks. Notice that chapter III of this thesis highlights some caveats about measuring common CF and DR shocks related to anomaly characteristics by building market-neutral long-minus-short portfolios. Nonetheless, there is no major apparent difference in CF and DR variance between value and momentum portfolios.

Table I - Variance contribution of CF and DR shocks: This table shows how much CF shocks, DR shocks, and the covariance component contribute to return variance. Panel A shows this decomposition for the three return components estimated in section III.C. Panel B reports the variance decomposition for the value-weighted portfolios along the two characteristics of interest: value and momentum.

Panel A: Return decomposition				
	Var(CF) %	Var(DR) %	-2Cov(CF,DR) %	Corr(CF,DR)
Firm-specific	71%	7%	22%	-0.49
Total joint return	64%	11%	25%	-0.48
Market-aggregate	33%	32%	35%	-0.54
Panel B: Portfolio variance contribution				
	Var(CF) %	Var(DR) %	-2Cov(CF,DR) %	Corr(CF,DR)
Value Portfolio	46%	20%	34%	-0.57
Growth Portfolio	42%	19%	39%	-0.69
HML Portfolio	69%	4%	27%	-0.78
Winner Portfolio	41%	21%	38%	-0.66
Loser Portfolio	46%	17%	37%	-0.67
WML Portfolio	64%	5%	31%	-0.89

Table II - Portfolio return correlations: This table shows the correlation between the CF and DR component of three value-weighted portfolios: The market portfolio (Market), the book-to-market portfolio (HML) and the momentum portfolio (WML). Standard errors of the estimates appear in parentheses.

		WML				Market				Market	
		CF	DR			CF	DR			CF	DR
HML	CF	-0.46*** (0.11)	0.36*** (0.05)	HML	CF	-0.04 (0.08)	-0.06 (0.08)	WML	CF	-0.06 (0.08)	0.1 (0.07)
	DR	0.46*** (0.04)	-0.50*** (0.11)		DR	-0.09 (0.08)	0.16** (0.06)		DR	0.06 (0.07)	-0.07 (0.08)

Table II reports correlations between the CF and DR of value (HML), momentum (WML) and market portfolios. There are significant negative correlations between value and momentum, both between their CF and DR components. CF shocks of the momentum portfolio tend to be high when value gets negative CF revisions. This finding echoes those of Asness et al. (2013) who find a strong negative correlation between value and momentum both across assets and across markets. The negative correlation between HML DR and WML DR also indicates that risk premia evolve in opposite directions over time. This is in line with Daniel and Moskowitz (2016) who document that it is following stressful market periods that value and momentum correlate the most negatively. Changes in DR might reflect those changes to underlying state variables related to adverse market conditions.

The correlation between value CF and momentum DR (and vice versa) is significantly positive. It could reflect the strong negative correlation structure between CF and DR, especially in long-minus-short portfolios. It suggests that CF shocks in one portfolio imply opposite DR shocks in the other portfolio, and strongly goes against the independence of CF and DR shocks.

Overall there are no significant correlations between market CF and DR shocks and either component of unexpected returns of HML and WML. There is one exception, however, since Market and HML are positively correlated at the 5% level.

News *tone* data is detailed at length in chapter II, section II. Relevant metrics for this paper are provided in section II.

III. Results

A. Drivers of value and momentum price changes

Following the seminal paper of Fama and French (1993), value is usually defined as the ratio of book-equity to market-equity. However, Gerakos and Linnainmaa

(2018) point out that it is possible to trace back a stock’s increase in book-to-market to two possible sources: there must either be a relative increase in book equity (BE), or a relative decrease in market equity (ME). Their key finding is that the value premium is linked to the change related to the drop in ME. Indeed, certain firms earn value-like returns, if they experienced a relative decrease in ME over the past five years. Other studies which seek to define value characteristics in different assets than stocks, where BE is not an observable metric, also rely on decreases in price over horizons of multiple years⁷.

In opposition to those patterns are momentum stocks, that by definition are those which experience the largest price increase over a short period of time (twelve months in this case). To understand why value stocks were subject to relative price decreases over past years and what has driven the price increase of momentum firms over the past year we run the following set of regressions:

$$\begin{aligned}
 Value_t &= \alpha_t + \beta_1 CF_{i,t-x} + \epsilon_{i,t} \\
 Growth_t &= \alpha_t + \beta_1 CF_{i,t-x} + \epsilon_{i,t} \\
 Value_t &= \alpha_t + \beta_1 DR_{i,t-x} + \epsilon_{i,t} \\
 Growth_t &= \alpha_t + \beta_1 DR_{i,t-x} + \epsilon_{i,t} \\
 Winner_t &= \alpha_t + \beta_1 CF_{i,t-x} + \epsilon_{i,t} \\
 Loser_t &= \alpha_t + \beta_1 CF_{i,t-x} + \epsilon_{i,t} \\
 Winner_t &= \alpha_t + \beta_1 DR_{i,t-x} + \epsilon_{i,t} \\
 Loser_t &= \alpha_t + \beta_1 DR_{i,t-x} + \epsilon_{i,t}
 \end{aligned} \tag{IV.2}$$

where $Value_t$, $Growth_t$, $Winner_t$, and $Loser_t$ are dummies characterizing firm-level observations as being part of those different portfolios during the current quarter t . To be classified into one portfolio, an observation must fall within the appropriate quintile: i.e., the quintile of firms with the highest book-to-market ratio go into the value portfolio, the quintile of stocks with the highest cumulative return over the months $t - 13$ to $t - 1$ go into the winner portfolio, and so on. α_t captures time fixed-effects, and $CF_{i,t-x}$ and $DR_{i,t-x}$ are the unexpected return shocks x quarters in the past.

Those equations seek to measure how much more likely a stock’s anomaly dummy is to be equal to one for each 1-point increase in its past CF or DR shock⁸. Table

⁷See for example Asness et al. (2013) for the case of currencies, commodity futures and government bonds, or Moskowitz (2021) for sports betting quotes.

⁸Note: This specification is based on a linear model with pooled panel data, including time fixed-

Table III - Drivers of past unexpected returns: This table presents the pooled regression estimates of equation (IV.2) for different lags up to forty quarters past. All regressions include date fixed effects and report clustered standard errors.

Panel A: Value, Past CF shocks										
	CF _{q-1}	CF _{q-2}	CF _{q-3}	CF _{q-4}	CF _{q-5}	CF _{q-6}	// CF _{q-8}	// CF _{q-10}	// CF _{q-20}	// CF _{q-40}
Value	-0.39*** (-14.90)	-0.36*** (-14.18)	-0.32*** (-13.70)	-0.30*** (-12.20)	-0.27*** (-10.61)	-0.23*** (-9.34)	-0.21*** (-8.11)	-0.18*** (-6.32)	-0.14*** (-4.74)	-0.08** (-2.20)
Growth	0.37*** (12.12)	0.37*** (12.02)	0.34*** (10.65)	0.33*** (10.72)	0.30*** (9.85)	0.28*** (9.51)	0.25*** (8.92)	0.22*** (7.44)	0.16*** (5.38)	0.09** (2.56)
N	173,184	161,599	153,439	146,948	140,837	135,085	124,589	115,095	79,166	43,984
Adjusted R ²	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.01	0.01

Panel B: Value, Past DR shocks										
	DR _{q-1}	DR _{q-2}	DR _{q-3}	DR _{q-4}	DR _{q-5}	DR _{q-6}	// DR _{q-8}	// DR _{q-10}	// DR _{q-20}	// DR _{q-40}
Value	1.49*** (14.79)	1.21*** (12.77)	1.10*** (11.76)	1.00*** (11.54)	0.97*** (12.00)	0.80*** (9.94)	0.67*** (8.82)	0.56*** (6.25)	0.41*** (3.75)	0.26** (2.46)
Growth	-1.14*** (-13.03)	-0.92*** (-11.04)	-0.81*** (-9.55)	-0.77*** (-9.34)	-0.74*** (-8.96)	-0.65*** (-8.05)	-0.53*** (-6.42)	-0.51*** (-5.65)	-0.30*** (-3.28)	-0.22** (-2.62)
N	173,184	161,599	153,439	146,948	140,837	135,085	124,589	115,095	79,166	43,984
Adjusted R ²	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.01

Panel C: Momentum, Past CF shocks										
	CF _{q-1}	CF _{q-2}	CF _{q-3}	CF _{q-4}	CF _{q-5}	CF _{q-6}	// CF _{q-8}	// CF _{q-10}	// CF _{q-20}	// CF _{q-40}
Winners	1.03*** (29.60)	0.97*** (33.26)	0.96*** (32.15)	1.05*** (33.43)	-0.05* (-1.72)	-0.02 (-0.57)	-0.02 (-0.66)	-0.01 (-0.44)	0.02 (0.58)	0.03 (1.14)
Losers	-1.08*** (-27.25)	-1.06*** (-31.73)	-1.07*** (-34.75)	-1.15*** (-34.55)	0.04 (1.27)	0.04 (1.24)	0.06** (2.02)	0.05* (1.74)	0.05* (1.68)	0.01 (0.40)
N	173,184	161,599	153,439	146,948	140,837	135,085	124,589	115,095	79,166	43,984
Adjusted R ²	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.01

Panel D: Momentum, Past DR shocks										
	DR _{q-1}	DR _{q-2}	DR _{q-3}	DR _{q-4}	DR _{q-5}	DR _{q-6}	// DR _{q-8}	// DR _{q-10}	// DR _{q-20}	// DR _{q-40}
Winners	-1.60*** (-16.64)	-1.83*** (-22.57)	-2.11*** (-25.57)	-2.27*** (-25.99)	0.10 (1.11)	0.05 (0.54)	0.12* (1.86)	0.10 (1.25)	-0.08 (-0.96)	0.01 (0.15)
Losers	1.54*** (11.23)	1.79*** (15.37)	2.25*** (21.18)	2.30*** (20.80)	-0.27*** (-3.04)	-0.30*** (-3.24)	-0.27*** (-3.08)	-0.29*** (-3.74)	-0.12 (-1.43)	-0.05 (-0.71)
N	173,184	161,599	153,439	146,948	140,837	135,085	124,589	115,095	79,166	43,984
Adjusted R ²	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.01

III reports the results on those regressions, looking up to 40 quarters back.

Panel A shows that stocks subject to negative (positive) CF revisions in the quarter prior ($q-1$) to the portfolio holding period are more likely to be classified as value (growth) firms in period q_0 . This probability slowly decays the further back in time the CF shock occurred. For example, an equivalent negative (positive)

effects and clustered errors by date and firm. We are conscious that since the dependent variable is a dummy in the space $[0,1]$, a non-linear model, such as a logistic regression might have been more appropriate. However, based on prior research designs (such as for example Engelberg et al. (2018)) and on the recommendations of Petersen (2009), the inclusion of the clustered errors seemed the most important in the design to avoid underestimation of standard errors and bias the resulting statistical inference. Due to technical difficulties, the latter was not possible with the logistic approach, and hence our choice for the linear model. As a result, our coefficient can not strictly speaking be treated as probabilities since it is possible for them to exceed the $[0,1]$ domain.

CF shock in q_{-4} is still 23% (11%) as likely to forecast a classification as a value (growth) stock in q_0 . Even if the CF shock took place five or ten years ago (q_{-20}) it still significantly forecasts the future book-to-market ranking.

Past DR shocks are also predictors of future classifications into value and growth portfolios. Firms with more positive revisions in expected returns (i.e., higher DR shocks) tend to subsequently be classified as value firms. Growth stocks on the other hand tend to be subject to lowered revisions expected returns. Those effects of DR shocks again seem to be long-lasting: DR shocks significantly impact book-to-market portfolio assignment even five or ten years ahead.

Momentum stocks exhibit very different patterns in their past unexpected returns. The key differentiation between winner and loser stocks occurs during the four quarters before the holding period, with a distinct break in q_{-5} . This matches with the portfolio formation period⁹. During those four quarters, stocks that were subject to positive (negative) CF shocks tend to subsequently become winner (loser) stocks. When looking five quarters ahead a sharp drop-off occurs and past unexpected returns fail to further forecast future momentum assignments.

The patterns of past DR shocks for momentum stocks, shown in panel D, may be the least anticipated ones. We find that stocks getting negative (positive) revisions in expected returns are significantly more likely to be become winner (loser) stocks, up to four quarters ahead. Therefore, despite having higher than average expected returns in q_0 ¹⁰, winner stocks were actually subject to more negative revisions in discount rates. This might come as a surprise, since negative DR shocks should subsequently lead to lower expected returns.

This result calls for a more thorough investigation of the dynamics of expected and unexpected returns of momentum stocks, which we do in section III.E. Indeed, the lower (higher) expected returns of winner (loser) stocks might only manifest multiple quarters after the initial DR shock. We find patterns consistent with this notion in figure IV.1: winner (loser) stocks have significantly lower (higher) expected returns in the quarters following the portfolio holding period. Those patterns fit particularly well with the notion that momentum goes hand-in-hand with a subsequent reversal.

Overall, value and momentum appear as polar opposites: one anomaly (value)

⁹In section III.E we seek to further investigate if the unexpected return patterns occur because of the choice of the formation period for momentum, or if momentum portfolio assignment is the result of past CF and DR shocks. We do this by removing momentum from the set of variables determining expected returns in the VAR.

¹⁰This is shown in the estimates of the VAR model (chapter III, section II, table I, panel A), where a high momentum rank forecasts higher than average next-period return ($p < 0.01$).

goes long stocks that were subject to the more negative CF revisions, but with sustained increases in expected returns; while the other (momentum) goes long firms that were subject to increases in expected cash-flows, but with lowered expected returns. The key difference between the two being the time-span over which the past relative CF and DR shocks are relevant for the two anomalies. Being able to measure quarterly movements in unexpected returns is crucial to understand the dynamics of the momentum anomaly in particular, given the sudden break in quarter $q-5$.

B. Relative contribution of CF and DR to past unexpected return differences

Next, we want to know the relative portion of unexpected price movements attributable to revisions in cash-flows and discount rates. This is possible to compute, since total unexpected returns are simply the difference of CF minus DR shocks (see eq. (III.7)). In practice, we measure by how much value and momentum stocks differed from the average firm's unexpected return over a period of one to five years as follows:

$$\begin{aligned}\overline{Ret}_{i,[t-x \rightarrow t]} &= \alpha_t + \beta_1 Value_{i,t} + \beta_2 Growth_{i,t} + \epsilon_{i,t} \\ \overline{Ret}_{i,[t-x \rightarrow t]} &= \alpha_t + \beta_1 Winner_{i,t} + \beta_2 Loser_{i,t} + \epsilon_{i,t}\end{aligned}\tag{IV.3}$$

where $\overline{Ret}_{i,[t-x \rightarrow t]}$ stands for the average over the period going from $t-x$ to t of the unexpected return component being studied. We consider the total unexpected return (UxRet), CF, and DR. The latter two sum up to UxRet by construction as per equation (III.7). Dummies and subscripts are the same as presented in equation (IV.2). The coefficients β_1 captures how much higher the unexpected return components are relative to other firms for long leg portfolio firms, whereas the β_2 measure this same relative difference for short portfolio stocks. For each time horizon, we are interested in the relative importance of the β coefficients for CF and DR relative to UxRet.

Gerakos and Linnainmaa (2018) show that decreases in relative prices are the main drivers behind the value premium. It seems therefore intuitive to expect that a sizeable portion of this price decline is driven by increases in discount rates, which would translate into higher expected returns for value stocks. Consistent with this hypothesis, table IV shows that a significant portion of value stock's relative price decline is driven by positive DR shocks. About 30% of the relative price decline can

be traced back to increases in expected returns, which should effectively translate into higher returns in subsequent periods.

Table IV - CF and DR contribution to average past unexpected returns: This table shows the differences in cumulated CF and DR shocks of value and momentum firms. β_1 and β_2 are the estimated coefficients from pooled regressions with the same specification as in equation (IV.2). The difference with table III is that there the independent variables were lagged returns, whereas here we use average log returns. The horizon (1-year, 2-year,...) indicates the length of the time period considered to compute the average return. “UxRet” is the unexpected return component and is equal to CF-DR. The percentages reported in brackets refer to the proportion of UxRet change caused to CF and DR.

		Book-to-Market				Momentum			
		Value		Growth		Winners		Losers	
		β_1	(%)	β_2	(%)	β_1	(%)	β_2	(%)
1-year	UxRet	-3.05		2.81		8.11		-8.10	
	CF	-2.16	(70.8%)	2.15	(76.3%)	6.79	(83.7%)	-6.87	(84.8%)
	(-)DR	-0.89	(29.2%)	0.67	(23.7%)	1.32	(16.3%)	-1.23	(15.2%)
2-year	UxRet	-2.53		2.40		3.93		-3.77	
	CF	-1.78	(70.3%)	1.82	(76.0%)	3.30	(83.8%)	-3.23	(85.5%)
	(-)DR	-0.75	(29.6%)	0.58	(24.0%)	0.64	(16.1%)	-0.54	(14.4%)
3-year	UxRet	-2.17		2.12		2.58		-2.27	
	CF	-1.51	(69.8%)	1.60	(75.3%)	2.16	(83.9%)	-1.97	(86.6%)
	(-)DR	-0.65	(30.2%)	0.52	(24.6%)	0.42	(16.2%)	-0.30	(13.4%)
4-year	UxRet	-1.91		1.94		2.02		-1.63	
	CF	-1.33	(69.7%)	1.44	(74.6%)	1.69	(83.9%)	-1.41	(86.7%)
	(-)DR	-0.58	(30.3%)	0.49	(25.4%)	0.32	(16.0%)	-0.22	(13.2%)
5-year	UxRet	-1.75		1.80		1.63		-1.27	
	CF	-1.21	(69.6%)	1.33	(73.9%)	1.37	(84.1%)	-1.11	(87.2%)
	(-)DR	-0.53	(30.4%)	0.47	(26.1%)	0.26	(15.9%)	-0.16	(12.8%)

The difference is driven mostly by CF shocks (almost 85%) for momentum stocks, which should be permanent in nature. When going further than one year in the past, momentum firms show no significant differences in unexpected returns, the entirety of the price differential beyond year(-1) being driven by the first four quarters¹¹. Conversely, as previously suggested by table III, value stocks were subject to consistent negative unexpected returns in the years past, going from -1.99% between year(-1) and year(-2), to -1.14% between year(-4) and year(-5). DR shocks consis-

¹¹Since the relative unexpected return differential (UxRet) is about 8% over the first year, 4% over the first two years, and 2% over the first 4 years, it follows arithmetically that beyond the first year no additional relative difference in UxRet is observed. This is expected given the finding in panels C and D of table III, where almost all the differences in unexpected returns for momentum stocks are concentrated in the first four quarters.

tently account for at least 30% of this price change, such that even at the five-year horizon value (growth) stocks have 0.53% higher (0.47% lower) DR returns than average.

C. Long term impacts of CF and DR shocks

This section investigates how changes in expected returns and cash-flows translate into future price changes. Can the consistently positive past DR shocks of value stocks be the driving force behind the premium of value stocks? Why do high momentum stocks get negative DR returns in the past four quarters yet still earn high returns during the holding period?

The first question has as a starting point that positive revisions about discount rates should forecast high expected returns in the future. Panel A of table V reports how much a 1% increase in DR shocks impacts future returns. We test the following specifications:

$$\begin{aligned} Ret_{i,t} &= \alpha_t + \beta \overline{DR}_{i,[t-x \rightarrow t]} + \epsilon_{i,t} \\ Ret_{i,t} &= \alpha_t + \beta \overline{CF}_{i,[t-x \rightarrow t]} + \epsilon_{i,t} \end{aligned} \tag{IV.4}$$

where $Ret_{i,t}$ is either the total log stock return (Totret), the expected return component (Eret, i.e., the difference between the total return and the unexpected return), the CF shock, or the DR shock, at time t for firm i . Indeed, as per equation (III.7) the total return can be decomposed into three components, i) the expected return, ii) the CF return, and iii) the DR return. In essence, the question those regression seek to answer is: how do past average unexpected returns impact the different components of future returns.

Positive DR shocks positively forecast future expected returns, especially when sustained over long horizons. A firm subject to 1% greater DR shocks over the past forty quarters has an expected return component in q_0 1.03% higher. This also translates in slightly significantly higher (+0.55%) total return at that horizon. This increased long-term return triggered by positive DR shocks is fully attributable to the expected return component. Sustained positive DR shocks over very long periods of 40 quarters also forecast a slightly significant decrease in future CF shocks (-0.45%).

This last result might suggest that the negative correlation between CF and DR shocks extends through time over multiple periods. However, notice that forecasting future CF shocks based on past DR shocks brings with it much more uncertainty than when forecasting future expected returns. As a result that the standard errors

Table V - Forecasting returns based on past CF and DR shocks: This table regresses cumulated DR shocks (panel A) and CF shocks (panel B) over determined periods x specified within the brackets, against the contemporaneous q_0 total return (Totret, first row of regressions), Expected Return (Eret, second row), CF and DR shocks. We report standard errors in parentheses to let the reader appreciate to which component of total returns the high standard errors of total returns can be attributed to. The regressions include date fixed-effects and the standard errors are clustered by date and firm.

Panel A: Impact of past DR shock										
$x=$	$[q-1,q-1]$	$[q-2,q-1]$	$[q-3,q-1]$	$[q-4,q-1]$	$[q-5,q-1]$	$[q-6,q-1]$	// $[q-8,q-1]$	// $[q-10,q-1]$	// $[q-20,q-1]$	// $[q-40,q-1]$
	(Totret$_{q_0}$)									
DR $_{q_x}$	0.06 (0.05)	0.10 (0.07)	-0.01 (0.09)	-0.08 (0.11)	0.02 (0.13)	0.09 (0.15)	0.09 (0.16)	0.17 (0.19)	0.24 (0.25)	0.55* (0.32)
	(Eret$_{q_0}$)									
DR $_{q_x}$	0.06*** (0.00)	0.06*** (0.01)	0.04*** (0.01)	0.01 (0.01)	0.08*** (0.01)	0.14*** (0.01)	0.25*** (0.01)	0.37*** (0.01)	0.72*** (0.01)	1.03*** (0.02)
	(CF$_{q_0}$)									
DR $_{q_x}$	-0.08* (0.04)	-0.15** (0.06)	-0.33*** (0.07)	-0.16* (0.09)	-0.16 (0.11)	-0.15 (0.12)	-0.19 (0.13)	-0.23 (0.15)	-0.53** (0.21)	-0.45* (0.27)
	(DR$_{q_0}$)									
DR $_{q_x}$	-0.08*** (0.01)	-0.19*** (0.02)	-0.28*** (0.02)	-0.08*** (0.03)	-0.10*** (0.03)	-0.10*** (0.03)	-0.03 (0.04)	-0.03 (0.04)	-0.06 (0.05)	0.02 (0.07)
N	173,184	161,599	153,439	146,948	140,837	135,085	124,589	115,095	79,166	43,984
date-FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Impact of past CF shock										
$x=$	$[q-1,q-1]$	$[q-2,q-1]$	$[q-3,q-1]$	$[q-4,q-1]$	$[q-5,q-1]$	$[q-6,q-1]$	// $[q-8,q-1]$	// $[q-10,q-1]$	// $[q-20,q-1]$	// $[q-40,q-1]$
	(TotRet$_{q_0}$)									
CF $_{q_x}$	0.00 (0.02)	0.00 (0.03)	0.05 (0.03)	0.08* (0.04)	0.06 (0.05)	0.05 (0.06)	0.06 (0.06)	0.02 (0.07)	0.03 (0.11)	-0.10 (0.12)
	(Eret$_{q_0}$)									
CF $_{q_x}$	0.00*** (0.00)	0.03*** (0.00)	0.05*** (0.00)	0.07*** (0.00)	0.07*** (0.00)	0.06*** (0.00)	0.04*** (0.00)	0.01*** (0.00)	-0.10*** (0.00)	-0.18*** (0.01)
	(CF$_{q_0}$)									
CF $_{q_x}$	-0.01 (0.01)	0.00 (0.02)	0.05* (0.03)	0.00 (0.03)	-0.01 (0.04)	-0.01 (0.05)	0.01 (0.05)	-0.02 (0.06)	0.12 (0.09)	0.07 (0.10)
	(DR$_{q_0}$)									
CF $_{q_x}$	-0.01 (0.00)	0.03*** (0.01)	0.05*** (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.02 (0.01)	-0.02 (0.02)	-0.01 (0.02)	-0.02 (0.03)
N	173,184	161,599	153,439	146,948	140,837	135,085	124,589	115,095	79,166	43,984
date-FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

of the estimated coefficients are big and reduce statistical significance. Therefore, those estimated errors suggest considering the observed negative correlation between past DR and present CF shocks with caution.

This table also highlights the strong link between contemporaneous DR shocks and expected returns, even multiple quarters ahead. This relationship is expected given the link between DR shocks and expected returns by construction of the model. It also fits with the positive premium of value stocks, which are more likely to have been subject to positive DR shocks over long periods of time of multiple years.

In panel B, still at long horizons of multiple years, we find that consistent nega-

tive CF shocks forecast high expected returns. Theoretically, CF shocks should be permanent in nature and have no impact on future expected returns. This pattern however aligns with the one observed for value stocks in table III. It might therefore result from the negative correlation between CF and DR shocks.

Based on the above results, to better understand the dynamics of momentum stocks we are interested in looking into short term impacts, of less than a year. Stocks with negative DR shocks over the recent four quarters appear to have slightly more positive CF and DR shocks. However, those two effects play against each other such that DR shocks do not forecast total returns at short horizons of one year. Therefore, the premium of momentum stocks at quarter q_0 is likely not the result of the most recent DR shocks.

Panel B also shows that positive CF shocks over the past four quarters forecast positive total returns. This aligns with the premium earned by momentum stocks in quarter q_0 and their high CF shocks in the preceding year.

D. Portfolios constructed on past CF and DR shocks

Table VI seeks to consolidate those results by moving away from the panel regression setting. The goal here is to measure the performance of portfolios formed based on past CF and DR shocks. As for the VAR model, the implicit assumption is that the transition matrix Π is known to the investor and remains constant over time. Panel B shows that a portfolio that goes long the top quintile of stocks sorted on average CF shocks over the past four quarters, and shorts those with the lowest CF revisions, earns a 5.24% return in excess of the market (t-stat = 2.37). This portfolio strongly correlates with the momentum factor (0.71) and is inversely related to the HML portfolio (cor=-0.21). As a result, when controlling for HML, the α of the portfolio goes up, whereas the momentum factor brings α down to zero.

Similarly, we can build portfolios based on past revisions in discount rates. Motivated by table V we would expect to see the stocks with the highest average DR shocks over the past forty years to earn high returns. Indeed, a portfolio that goes long into those stocks and shorts the quintile of firms with the most negative DR shocks over the same period earns a 5.26% return premium in excess of the market (t-stat = 2.14). This long-term DR portfolio has opposite exposures to value and momentum than the short-term CF portfolio. The correlation with HML is 0.68 and is -0.36 with WML, both significant at the one-percent level.

Those portfolios built on past unexpected returns again reflect the negative correlation between value and momentum documented by Asness et al. (2013) and Daniel

Table VI - Portfolios built on past CF and DR shocks: We build portfolios based on past average CF and DR shocks. The time span in brackets $[q-x, q-y]$ indicates over which period the averages for the classification are computed. We rank the portfolios based on those averages and go long the top quintile and short the bottom quintile. The portfolio joining CF and DR information described in section III.C is denoted with the union sign \cup . Panel A reports the correlation of those portfolio with the Market, HML and WML portfolios obtained from the website of Kenneth French. T-stats appear in parentheses. Panel B reports the intercept of a regression including those portfolios as controls and again reports the t-stat in parentheses.

Panel A: Correlations of portfolios built on past CF and DR shocks					
	$\overline{DR}_{[q-40, q-1]}$	$\overline{CF}_{[q-40, q-1]}$	$\overline{CF}_{[q-4, q-1]}$	$\overline{DR}_{[q-4, q-1]}$	$\overline{DR}_{[q-40, q-1]} \cup \overline{CF}_{[q-4, q-1]}$
Mkt.Rf	-0.23*** (-2.69)	0.30*** (3.66)	0.02 (0.23)	-0.06 (-0.79)	-0.19** (-2.21)
HML	0.68*** (10.66)	-0.56*** (-7.80)	-0.21*** (-2.85)	0.31*** (4.32)	0.37*** (4.56)
WML	-0.36*** (-4.44)	0.29*** (3.53)	0.71*** (13.06)	-0.65*** (-11.15)	0.36*** (4.46)

Panel B: Alpha of portfolios built on past CF and DR shocks					
	$\overline{DR}_{[q-40, q-1]}$	$\overline{CF}_{[q-40, q-1]}$	$\overline{CF}_{[q-4, q-1]}$	$\overline{DR}_{[q-4, q-1]}$	$\overline{DR}_{[q-40, q-1]} \cup \overline{CF}_{[q-4, q-1]}$
α - Mkt.Rf	5.26%** (2.14)	-5.00%** (-2.58)	5.24%** (2.37)	-1.50% (-0.80)	6.07%*** (2.91)
α - Mkt.Rf + HML	2.06% (1.10)	-3.02%* (-1.78)	6.63%*** (2.98)	-3.14%* (-1.73)	4.67%** (2.35)
α - Mkt.Rf + WML	8.81%*** (3.76)	-7.37%*** (-3.99)	-0.91% (-0.58)	3.43%** (2.31)	3.82%* (1.88)
α - Mkt.Rf + HML + WML	4.67%** (2.52)	-4.97%*** (-2.95)	-0.79% (-0.48)	2.43% (1.59)	1.22% (0.66)

and Moskowitz (2016), and their pervasive opposition with one another. They also note that this negative correlation between HML and WML implies that a portfolio on the efficient frontier will likely need exposure to both factors.


Motivated by these remarks, we use a scoring approach to build a portfolio that loads both on long-term DR and short-term CF shocks. Every quarter, each stock is assigned a percentile score from 1 to 100 along each of those two dimensions (highest average DR over 40 quarters and highest average CF over four quarters score 100). The sum of those scores defines the final rank along which the final portfolio goes long the top quintile and short the bottom quintile. This portfolio has a significant correlation with both HML (0.37) and WML (0.36). It earns a 6.07% return in excess of the market. Controlling for HML or WML in isolation does not eliminate the premium, but when including both together, the alpha of the portfolio becomes

insignificant.

E. Robustness concerns

At this stage, it is necessary to introduce robustness considerations. Indeed, the baseline VAR model relies on book-to-market and momentum as state variables to forecast future returns. The first column in panel A of table 1 (in chapter III) implies that for firms with high levels of book-to-market and price momentum in the previous quarter, the model will generate high expected returns. Consistent with this prediction, figures IV.1a and IV.1c show that both value and winner stocks have high expected returns in q_0 .

Figure IV.1. - Evolution of Expected Return Component: This set of figures reports the expected return component (Eret) for value and growth (blue and red lines) in the top row and winners and losers (blue and red) in the bottom row. The coefficients for each quarter are estimated from a pooled OLS regression as in equation (IV.2) where Eret is the dependent variable. The regressions include date fixed-effects and the standard errors are clustered by firm and date. The left column is for the baseline specification, whereas the right column reports the estimates for the VAR specification without value and momentum in the vector of state variables.



What happens when expected returns are not dependent on the level of book-to-market and price momentum? In that scenario, dynamics of expected returns change, at least for momentum. Consider a VAR specification from which momentum is omitted from the set of state variables. The consequence of this model, where the past twelve-month price change does not influence expected returns, is visible in figure IV.1d. It shows that the expected return component of winner and loser stocks does not spike in q_0 . Instead, the pattern that clearly subsists when compared to the default specification is that loser (winner) stocks see an increase in expected returns only after the holding the period. This behavior would be consistent with the positive (negative) DR shocks of loser (winner) stocks documented in section III.A above. This calls for a cautious robustness check of all previously documented results in the sections above.

Table VII investigates the robustness of the results across five different specifications. *LT20* refers to the baseline specification used throughout the paper and follows the approach of Lochstoer and Tetlock (2020). It is the preferred framework that leverages the most information about drivers of stock returns, including levels of momentum and book-to-market. *LT20-ex.Val-ex.Mom* is a modified version of *LT20*, where we remove momentum and book-to-market from the set of state variables. Similarly, *LT20-ex.Val* and *LT20-ex.Mom* are modified versions of the baseline VAR specification where only book-to-market and momentum, respectively, are omitted. The objective of those changes is to pinpoint which results are sensitive to their inclusion in the VAR. Finally, V02 is the last specification from which we estimate CF and DR by following the approach of Vuolteenaho (2002).

Panel A aims at completing the discussion raised by figure IV.1. Given the mechanics of the VAR model, variables that predict returns at time q_0 end up in the contemporaneous expected component of returns. High book-to-market and high momentum stocks have elevated expected returns. Since these characteristics capture the expected component of returns, they should not influence unexpected returns in q_0 . Indeed, we find no differences in CF and DR shocks in the *LT20* model. Neither along the value-growth nor along the winner-loser axis. In this framework, value and momentum stocks earn their premium because of positive expected returns, not because CF or DR shocks are systematically higher or lower. If the latter were the case, the model would be doing a poor job at its intended purpose.

Table VII - Robustness: This table reports how the previously documented results stand up to a set of four alternative specifications for the VAR model that estimates CF and DR shocks. The alternative specifications are detailed in section III.E. Panel A reports contemporaneous (q_0) average CF and DR shocks of value and momentum stocks. Panel B reports the forecast of returns for 1-year CF and 10-year DR shocks, as in table V. Panel C reports the contribution of CF and DR shocks to past unexpected returns at the one-year and five-year horizon, as in table IV. Finally, panel D reports the key results of table VI for the robustness specifications.

Panel A: CF and DR shocks in q_0						
		LT20	LT20- ex.Val- ex.Mom	LT20- ex.Mom	LT20- ex.Val	V02
DR q_0 (%)	Value	0.00 (0.03)	0.10** (2.28)	0.03 (0.53)	0.03 (0.53)	0.05 (0.40)
	Growth	0.01 (0.17)	-0.13** (-2.33)	0.01 (0.10)	-0.10 (-1.34)	-0.10 (-0.72)
	Winner	0.04 (0.60)	-0.28*** (-5.27)	-0.04 (-0.67)	-0.02 (-0.29)	0.03 (0.20)
	Loser	0.03 (0.36)	0.26*** (5.12)	0.16** (1.96)	-0.06 (-1.04)	0.03 (0.21)
CF q_0 (%)	Value	-0.34 (-1.67)	0.24 (-0.81)	-0.36* (-1.75)	0.41 (1.31)	-0.56 (-1.50)
	Growth	-0.03 (-0.11)	-0.70* (-1.77)	0.01 (0.03)	-0.89** (-2.22)	0.18 (0.36)
	Winner	-0.49* (-1.78)	0.53* (1.67)	0.23 (0.85)	-0.56* (-1.68)	1.49*** (3.75)
	Loser	-0.35 (-1.12)	-1.30*** (-3.26)	-1.05*** (-3.38)	-0.26 (-0.62)	-2.19*** (-5.30)

Table VIII - Continued.

Panel B: Decomposed return forecast to past CF/DR shocks						
		LT20	LT20- ex.Val- ex.Mom	LT20- ex.Mom	LT20- ex.Val	V02
$\overline{CF}_{[q-4,q-1]}$	TotRet	0.08* (0.04)	0.17*** (0.04)	0.08* (0.04)	0.16*** (0.04)	0.18*** (0.06)
	Eret	0.07*** (0.00)	0.01*** (0.00)	-0.01*** (0.00)	0.13*** (0.00)	-0.04*** (0.01)
	CF	0.00 (0.03)	0.14*** (0.04)	0.08** (0.03)	0.04 (0.04)	0.20*** (0.05)
	(-)DR	0.00 (0.01)	0.02*** (0.01)	0.01 (0.01)	-0.01** (0.01)	0.02 (0.02)
$\overline{DR}_{[q-40,q-1]}$	TotRet	0.55* (0.32)	1.32** (0.67)	0.58* (0.34)	1.00 (0.81)	1.16 (0.93)
	Eret	1.03*** (0.02)	1.01*** (0.05)	0.97*** (0.01)	1.16*** (0.08)	1.25*** (0.04)
	CF	-0.45* (0.27)	0.32 (0.59)	-0.30 (0.28)	-0.48 (0.81)	-0.11 (0.75)
	(-)DR	-0.02 (0.07)	-0.01 (0.18)	-0.09 (0.07)	0.32* (0.18)	0.03 (0.30)

Table VIII - Continued.

Panel C: CF/DR contribution to past shocks						
		LT20	LT20- ex.Val- ex.Mom	LT20- ex.Mom	LT20- ex.Val	V02
1-year UxRet Value	CF	-2.16	-2.40	-2.20	-2.38	-1.85
	(%)	(71%)	(87%)	(72%)	(92%)	(60%)
Value	DR	-0.89	-0.31	-0.85	-0.20	-1.26
	(%)	(29%)	(13%)	(28%)	(8%)	(40%)
5-year UxRet Value	CF	-1.21	-1.42	-1.23	-1.39	-1.24
	(%)	(70%)	(87%)	(72%)	(93%)	(57%)
Value	DR	-0.53	-0.18	-0.49	-0.10	-0.93
	(%)	(30%)	(13%)	(28%)	(7%)	(43%)
1-year UxRet Winners	CF	6.79	8.23	6.69	8.52	-6.92
	(%)	(84%)	(89%)	(81%)	(96%)	(77%)
Winners	DR	1.32	0.99	1.52	0.42	-2.11
	(%)	(16%)	(11%)	(18%)	(4%)	(23%)
5-year UxRet Winners	CF	1.37	1.74	1.10	1.82	1.41
	(%)	(84%)	(90%)	(84%)	(96%)	(79%)
Winners	DR	0.26	0.19	0.22	0.07	0.38
	(%)	(16%)	(10%)	(16%)	(4%)	(21%)

Table VIII - Continued.

Panel D: Past CF/DR portfolios						
		LT20	LT20- ex.Val- ex.Mom	LT20- ex.Mom	LT20- ex.Val	V02
$\overline{CF}_{[q-4, q-1]}$	α -Mkt.Rf	5.24%	10.78%	5.77%	10.80%	10.25%
	(<i>t</i>)	(2.37)	(2.98)	(2.58)	(3.04)	(2.59)
	Mkt cor	0.02	0.02	0.01	0.03	0.17
	HML cor	-0.21***	-0.11	-0.22***	-0.10	-0.15
$\overline{DR}_{[q-40, q-1]}$	WML cor	0.71***	0.76***	0.73***	0.75***	0.78***
	α -Mkt.Rf	5.26%	16.07%	5.56%	13.20%	18.01%
	(<i>t</i>)	(2.14)	(2.60)	(2.26)	(2.39)	(1.65)
	Mkt cor	-0.23***	-0.57**	-0.23***	-0.55**	-0.47
$\overline{DR}_{[q-40, q-1]}$	HML cor	0.68***	0.44*	0.70***	0.34	0.88***
	WML cor	-0.36***	-0.65***	-0.38***	-0.56**	-0.49

Omitting momentum from the VAR (*LT20-ex.Mom*) causes loser stocks to exhibit negative CF and positive DR shocks in q_0 . This disparity in unexpected returns

explains why the difference in expected returns in q_0 disappears when going from figure IV.1c to IV.1d. Removing book-to-market (*LT20-ex.Val*) from the state variables causes growth stocks to have more negative CF shocks than value firms. This accompanies the change in expected returns of value stocks observed when going from figure IV.1a to IV.1b: in the baseline specification, value stocks have significantly higher expected returns than growth stocks (about 1% difference), but in *LT20-ex.Val* this difference drops to about 0.3% because part of the value premium is “captured” by the difference in CF. When removing both book-to-market and momentum (*LT20-ex.Val.ex.Mom*), both effects documented above coincide. Finally, since *V02* does not include momentum and heavily relies on book-to-market, winner stocks have more positive CF shocks than losers, similar patterns as *LT20-ex.Mom*.

Panel B shows that the differences in drivers of expected returns across models carry over to the analysis from section III.C. Here we report how short-term (i.e., four quarters) CF shocks and long-term (i.e. forty quarters) DR shocks forecast future returns, as in table V. Across all specifications, high total returns follow short-run increases in CF. However, the source of this increase depends on the specification of the model. As shown before for *LT20*, positive CF shocks in the past four quarters coincide with a high momentum characteristic, which translates into high expected returns (see figure IV.1c). This observation also holds for *LT20-ex.Val*, which is the only other specification where momentum appears in the set of state variables. In specifications without momentum (*LT20-ex.Mom*, *LT20-ex.Val.ex.Mom* and *V02*), the forecasted increase in total return following 1-year high CF shocks is not driven by high expected returns. Instead, it is the consequence of positive CF shocks that follow the initial revisions.

Although the source of return forecastability varies across specifications, panel C highlights that forming portfolios ranked on past short-term CF shocks consistently outperform the market portfolio, with a t-stat ranging from 2.37 (*LT20*) to 3.04 (*LT20-ex.Val*). The implication here is that no matter the selected model, short-run increases in CF shocks are further followed by high returns. The link to momentum is found across all specifications since the correlation with the WML factor remains significant.

Long-term DR shock’s ability to forecast total returns, on the other hand, seems to draw more consistently from the high expected return that follows a string of positive DR shocks (see panel B). This appears to hold across all specifications, although standard errors vary and affect the statistical significance. In panel C, we again find that portfolios formed on past DR shorts earn a premium compared to

the market across all specifications. Note that the portfolio capturing this premium negatively correlates with the market and momentum factors but has a positive exposure to value, at least when including book-to-market in the VAR.

Finally, panel D checks the robustness of the initial results about drivers of value and momentum price changes in the years before the holding period. When book-to-market is included as a state variable in the VAR model, the price decline of value stocks is driven for a large part by positive DR shocks. They account for 28% to 43% of the difference in past unexpected returns, even at long horizons. This positive revision in expected returns likely reflects a progressive increase in exposure to the value factor. Indeed, when omitting book-to-market from the model, the proportion of past DR shocks drops by more than half but remains positive.

This again contrasts with momentum stocks, that are subject to negative DR shocks across all specifications, despite their strong positive CF revisions. The negative (positive) DR shocks of winner (loser) stocks is also consistent with the patterns of expected returns observed several quarters after the holding period in figure IV.1. Therefore, the negative DR shocks of momentum stocks imply lower returns, and thus a reversal, several years down the line.

IV. Tone correlates with the observed changes in CF and DR

As a follow-up to the discussion above, we investigate how *tone*, an exogenous measure of news that arrives to investors, relates to the patterns in CF and DR of value and momentum¹² stocks. We extensively detail the construction of this measure of news polarity in chapter II. The underlying premise is that the media reflects information that leads investors to revise their expectations about future cash flows and discount rates. Therefore, the content of media might help to pin down the nature of the differences in CF and DR updates documented in section III.

Table IX reports summary statistics about *tone* of value, growth, winner and loser stocks at quarterly frequency. We build firm-specific news *tone* over intervals of three months following equation II.2. We find that growth stocks get significantly

¹²Previous research has already highlighted the link between media coverage and momentum. Hillert et al. (2014) document that momentum is strongest among firms with high media coverage, while Da et al. (2014) bring more nuance to this argument, by suggesting that firms with continuous attention form a gradual and sustained momentum which contrasts with firms receiving infrequent and dramatic news.

Table IX - Tone of value and momentum: The left side of this table reports summary statistics on quarterly firm-specific news *tone*. “Value”, “Growth”, “Winners” and “Losers” refer to their specific sub-samples, “All” is the full sample. *tone* is scaled overall to have a mean of zero and a standard deviation of one. The differences (V-G, W-L) have the z-stat of the difference in mean reported in brackets. No measure of *tone* is available before January 2003. The right side of the table reports the estimated β coefficients for each sub-sample from a panel regression specified as: $ret_{i,t} = \alpha_t + \beta tone_{i,t}$, where *ret* is either the CF or DR shock. This regression includes date fixed effects and errors are clustered by date and firm. T-stats are reported in brackets. For the row measuring the differences (V-G, W-L), the significance of the difference in coefficients is measured by a z-test described in Clogg et al. (1995).

	μ	σ	Min	25 th _p	75 th _p	Max	% Obs. w. tone	% Obs. w. tone >= 2003	CF \sim tone	DR \sim tone
All	0.00	1.00	-4.61	-0.63	0.70	3.00	29%	65%	1.43*** [10.05]	-0.32*** [-8.91]
Value	-0.25	1.00	-4.61	-0.85	0.43	2.97	24%	54%	1.94*** [8.15]	-0.42*** [-6.67]
Growth	0.16	1.00	-4.61	-0.48	0.87	3.00	31%	70%	1.36*** [9.31]	-0.29*** [-10.33]
V-G [z-score]	-0.41*** [-31.29]	0.00	0.00	-0.37	-0.44	-0.03	-7%	-16%	0.58** [2.08]	-0.13* [-1.94]
Winners	0.15	0.97	-4.61	-0.48	0.84	3.00	29%	64%	1.38*** [7.04]	-0.32*** [-6.64]
Losers	-0.29	1.06	-4.61	-0.95	0.45	2.95	28%	62%	2.32*** [10.21]	-0.45*** [-8.23]
W-L [z-score]	0.44*** [32.78]	-0.09	0.00	0.47	0.39	0.05	1%	2%	-0.94*** [-3.16]	0.13* [1.76]

more positive news coverage than value firms on average. They also benefit from more news coverage overall. Winner stocks are also firms that significantly more positive media coverage than average. The difference with loser stocks is significant, who get very negative news. However, those firms do not significantly differ in the total amount of news coverage they are subject to.

Table IX also investigates the impact of *tone* on CF and DR shocks. An increase in *tone* positively correlates with CF shocks, which is expected if news in media reflect updates in expectations about future cash flows. *tone* also correlates negatively with DR shocks. Interpreting the latter might be somewhat more delicate. One possibility is that good (bad) news generally reduce (increase) the exposure to priced risk factors. Another possibility is that it reflects the negative correlation between CF and DR and that the negative impact of *tone* on DR shocks imply reversal due to overreaction. Overall, we find that good (bad) news primarily impact prices by improving (deteriorating) CF expectations and decreasing (increasing) risk and

Table X - Past tone: This table reports the average tone of value and momentum stocks at different lags up to 40. The table is identical to table III and is based on the same specification as in equation (IV.2), except that *tone* is the dependent variable. The regressions include date fixed effects and errors are clustered by date and firm.

Panel A: Value, Past <i>tone</i>										
	<i>tone</i> _{<i>q</i>-1}	<i>tone</i> _{<i>q</i>-2}	<i>tone</i> _{<i>q</i>-3}	<i>tone</i> _{<i>q</i>-4}	<i>tone</i> _{<i>q</i>-5}	<i>tone</i> _{<i>q</i>-6}	// <i>tone</i> _{<i>q</i>-8}	// <i>tone</i> _{<i>q</i>-10}	// <i>tone</i> _{<i>q</i>-20}	// <i>tone</i> _{<i>q</i>-40}
Value	-0.29*** (-18.75)	-0.26*** (-17.15)	-0.24*** (-13.92)	-0.22*** (-13.14)	-0.21*** (-13.15)	-0.19*** (-12.12)	-0.17*** (-11.83)	-0.16*** (-11.54)	-0.14*** (-7.19)	-0.11*** (-4.22)
Growth	0.18*** (15.29)	0.19*** (14.99)	0.19*** (13.85)	0.18*** (13.56)	0.17*** (13.08)	0.18*** (13.98)	0.18*** (15.02)	0.18*** (14.40)	0.20*** (12.08)	0.15*** (8.09)
<i>N</i>	56,696	54,324	53,010	50,712	48,519	46,439	42,580	38,998	24,701	9,679
Adjusted <i>R</i> ²	0.04	0.04	0.04	0.04	0.04	0.04	0.03	0.04	0.03	0.01

Panel B: Value Future <i>tone</i>										
	<i>tone</i> _{<i>q</i>0}	<i>tone</i> _{<i>q</i>+1}	<i>tone</i> _{<i>q</i>+2}	<i>tone</i> _{<i>q</i>+3}	<i>tone</i> _{<i>q</i>+4}	<i>tone</i> _{<i>q</i>+5}	<i>tone</i> _{<i>q</i>+6}	// <i>tone</i> _{<i>q</i>+8}	// <i>tone</i> _{<i>q</i>+10}	// <i>tone</i> _{<i>q</i>+20}
Value	-0.26*** (-17.25)	-0.22*** (-13.55)	-0.20*** (-11.98)	-0.18*** (-11.88)	-0.17*** (-10.90)	-0.15*** (-9.48)	-0.15*** (-9.76)	-0.10*** (-6.73)	-0.09*** (-6.80)	-0.03*** (-1.98)
Growth	0.15*** (14.05)	0.14*** (13.23)	0.13*** (12.00)	0.12*** (11.66)	0.10*** (10.08)	0.09*** (8.05)	0.08*** (6.77)	0.08*** (6.52)	0.08*** (5.61)	0.07*** (4.63)
<i>N</i>	59,871	56,252	53,463	51,738	50,092	48,495	46,967	44,122	41,464	30,717
Adjusted <i>R</i> ²	0.04	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03

Panel C: Momentum, Past <i>tone</i>										
	<i>tone</i> _{<i>q</i>-1}	<i>tone</i> _{<i>q</i>-2}	<i>tone</i> _{<i>q</i>-3}	<i>tone</i> _{<i>q</i>-4}	<i>tone</i> _{<i>q</i>-5}	<i>tone</i> _{<i>q</i>-6}	// <i>tone</i> _{<i>q</i>-8}	// <i>tone</i> _{<i>q</i>-10}	// <i>tone</i> _{<i>q</i>-20}	// <i>tone</i> _{<i>q</i>-40}
Winners	0.17*** (8.16)	0.14*** (6.14)	0.12*** (6.09)	0.07*** (4.05)	-0.05*** (-2.71)	-0.03* (-1.73)	-0.03 (-1.59)	-0.03 (-0.84)	0.01 (0.46)	0.07** (2.05)
Losers	-0.42*** (-15.86)	-0.35*** (-13.42)	-0.28*** (-11.60)	-0.17*** (-8.23)	-0.07*** (-3.99)	-0.04** (-2.48)	-0.01 (-0.62)	-0.00 (-0.13)	0.01 (0.32)	0.00 (0.12)
<i>N</i>	56,696	54,324	53,010	50,712	48,519	46,439	42,580	38,998	24,701	9,679
Adjusted <i>R</i> ²	0.06	0.05	0.04	0.03	0.02	0.02	0.02	0.03	0.02	0.00

Panel D: Momentum, Future <i>tone</i>										
	<i>tone</i> _{<i>q</i>0}	<i>tone</i> _{<i>q</i>+1}	<i>tone</i> _{<i>q</i>+2}	<i>tone</i> _{<i>q</i>+3}	<i>tone</i> _{<i>q</i>+4}	<i>tone</i> _{<i>q</i>+5}	<i>tone</i> _{<i>q</i>+6}	// <i>tone</i> _{<i>q</i>+8}	// <i>tone</i> _{<i>q</i>+10}	// <i>tone</i> _{<i>q</i>+20}
Winners	0.10*** (6.65)	0.06*** (4.46)	0.05*** (3.60)	0.04** (2.55)	0.01 (0.79)	0.00 (0.29)	-0.02* (-1.82)	-0.06*** (-4.51)	-0.05*** (-3.66)	-0.09*** (-3.76)
Losers	-0.33*** (-13.86)	-0.28*** (-12.10)	-0.25*** (-12.28)	-0.20*** (-12.64)	-0.18*** (-12.15)	-0.14*** (-9.71)	-0.11*** (-7.45)	-0.09*** (-4.60)	-0.09*** (-4.70)	-0.02 (-1.15)
<i>N</i>	59,871	56,252	53,463	51,738	50,092	48,495	46,967	44,122	41,464	30,717
Adjusted <i>R</i> ²	0.05	0.04	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.03

discount rates. We document a coefficient which is five times greater on CF than DR impacts.

Value stocks tend to be more sensitive than growth stocks to *tone*. The difference is significant both for the impact on CF and DR. At first, this result might come as a surprise when compared to the findings in chapter II, where we document that growth stocks are more sensitive to *tone* on EAD. On non-EAD, we did not find a significant between value and growth stocks sensitivity to news *tone*. Remember also that the increased sensitivity of growth stocks is specific to EAD with bad news. Here *tone* is aggregated over entire quarters, and therefore the movements we capture are of lower frequency. Similarly, loser stocks are also more sensitive than winner stocks.

Table XI - Past shocks, controlling for tone: This table is identical to table III, except that we include past *tone* as an additional independent variable.

Panel A: Value, Past CF controlling for past <i>tone</i>										
	CF _{q-1}	CF _{q-2}	CF _{q-3}	CF _{q-4}	CF _{q-5}	CF _{q-6}	// CF _{q-8}	// CF _{q-10}	// CF _{q-20}	// CF _{q-40}
Value	-2.16*** (-5.35)	-2.04*** (-6.53)	-1.57*** (-7.46)	-1.28*** (-6.51)	-1.08*** (-5.38)	-0.90*** (-4.79)	-0.78*** (-4.57)	-0.63*** (-3.31)	-0.33 (-1.24)	0.31 (1.03)
Growth	1.66*** (7.70)	1.57*** (7.32)	1.41*** (6.39)	1.31*** (6.03)	1.08*** (4.91)	1.03*** (4.72)	1.12*** (6.40)	0.92*** (4.96)	0.79*** (3.62)	0.36 (1.45)
<i>tone</i>	1.16*** (9.82)	1.14*** (9.54)	1.16*** (10.07)	1.10*** (9.92)	1.06*** (9.74)	1.03*** (9.48)	0.95*** (8.82)	0.93*** (8.70)	0.73*** (6.65)	0.53*** (3.16)
N	56,696	54,324	53,010	50,712	48,519	46,439	42,580	38,998	24,701	9,679
Adjusted R ²	0.22	0.22	0.21	0.21	0.21	0.21	0.21	0.21	0.23	0.14

Panel B: Value, Past DR controlling for past <i>tone</i>										
	DR _{q-1}	DR _{q-2}	DR _{q-3}	DR _{q-4}	DR _{q-5}	DR _{q-6}	// DR _{q-8}	// DR _{q-10}	// DR _{q-20}	// DR _{q-40}
Value	0.93*** (8.08)	0.74*** (7.24)	0.66*** (8.47)	0.53*** (7.89)	0.48*** (8.08)	0.44*** (7.45)	0.39*** (6.51)	0.32*** (4.82)	0.25*** (2.75)	0.07 (0.52)
Growth	-0.54*** (-6.76)	-0.41*** (-5.76)	-0.38*** (-5.18)	-0.37*** (-5.14)	-0.36*** (-4.80)	-0.30*** (-4.13)	-0.31*** (-4.47)	-0.28*** (-3.92)	-0.22*** (-2.84)	-0.13 (-1.44)
<i>tone</i>	-0.24*** (-7.82)	-0.25*** (-8.27)	-0.26*** (-9.17)	-0.26*** (-9.34)	-0.25*** (-8.94)	-0.24*** (-8.90)	-0.22*** (-8.34)	-0.21*** (-8.04)	-0.14*** (-6.00)	-0.06 (-1.46)
N	56,696	54,324	53,010	50,712	48,519	46,439	42,580	38,998	24,701	9,679
Adjusted R ²	0.62	0.62	0.63	0.63	0.64	0.64	0.65	0.66	0.67	0.50

Panel C: Momentum, Past CF controlling for past <i>tone</i>										
	CF _{q-1}	CF _{q-2}	CF _{q-3}	CF _{q-4}	CF _{q-5}	CF _{q-6}	// CF _{q-8}	// CF _{q-10}	// CF _{q-20}	// CF _{q-40}
Winners	5.32*** (19.51)	5.13*** (19.75)	5.11*** (18.82)	5.25*** (18.55)	-0.71** (-2.12)	-0.34 (-1.05)	0.38* (1.68)	0.02 (0.06)	-0.05 (-0.15)	0.39 (1.10)
Losers	-5.44*** (-11.05)	-5.36*** (-14.44)	-5.20*** (-14.11)	-5.37*** (-16.91)	0.62** (2.25)	0.56* (1.82)	0.27 (1.49)	0.07 (0.35)	0.32 (0.91)	0.41 (0.96)
<i>tone</i>	0.71*** (8.12)	0.79*** (8.64)	0.87*** (9.47)	0.95*** (10.22)	1.15*** (10.30)	1.10*** (10.03)	1.01*** (9.26)	0.98*** (9.01)	0.78*** (6.84)	0.54*** (3.18)
N	56,696	54,324	53,010	50,712	48,519	46,439	42,580	38,998	24,701	9,679
Adjusted R ²	0.32	0.31	0.31	0.32	0.20	0.20	0.20	0.21	0.23	0.14

Panel D: Momentum, Past DR controlling for past <i>tone</i>										
	DR _{q-1}	DR _{q-2}	DR _{q-3}	DR _{q-4}	DR _{q-5}	DR _{q-6}	// DR _{q-8}	// DR _{q-10}	// DR _{q-20}	// DR _{q-40}
Winners	-0.60*** (-8.06)	-0.86*** (-11.38)	-1.04*** (-14.54)	-1.17*** (-14.89)	0.12 (1.43)	0.05 (0.58)	-0.03 (-0.47)	0.06 (0.84)	-0.06 (-0.66)	0.01 (0.18)
Losers	0.66*** (4.56)	0.88*** (8.04)	1.17*** (10.58)	1.10*** (12.66)	-0.14** (-2.13)	-0.14* (-1.89)	-0.03 (-0.44)	-0.04 (-0.62)	-0.06 (-0.87)	-0.10 (-0.87)
<i>tone</i>	-0.23*** (-8.88)	-0.22*** (-8.37)	-0.22*** (-8.71)	-0.24*** (-9.75)	-0.28*** (-9.59)	-0.27*** (-9.56)	-0.25*** (-8.85)	-0.23*** (-8.39)	-0.15*** (-6.18)	-0.07 (-1.56)
N	56,696	54,324	53,010	50,712	48,519	46,439	42,580	38,998	24,701	9,679
Adjusted R ²	0.62	0.63	0.64	0.65	0.63	0.64	0.65	0.65	0.67	0.50

The patterns of media *tone* for value and momentum stocks closely follow those documented for past CF and DR shocks in table III. Table X shows that as value (growth) stocks get negative (positive) CF shocks over the past years, their average media coverage was also more negative (positive). It also coincides with the past increases in DR shocks. For momentum, *tone* also follows the CF and DR patterns: winner stocks start picking up positive news just four quarters prior to the holding period, then continue benefiting from positive media coverage for one more year. Loser stocks show opposite patterns, although their string of negative news *tone* is

more consistent and starts arising in q_{-6} .

Finally, table XI investigates if past *tone* can explain the past patterns in CF and DR shocks of value and momentum. We revisit equation (IV.2) by controlling for past *tone*. For momentum, the question it seeks to answer becomes: does the increase (decrease) in CF (DR) of winner stocks originate from the nature of the information released in the media about those firms? For value, the question is if the bad news *tone* observed the past quarters is responsible for the negative unexpected returns.

We can directly compare the coefficients of tables III and XI. It appears that the overall patterns of past CF and DR shocks still remain present for both value and momentum. *tone* positively impacts CF shocks and negatively DR shocks, but the cross-sectional difference in *tone* only accounts for a modest proportion of the differences in unexpected returns. Both for value and momentum, *tone* reduces the differences in average CF and DR shocks by somewhere between 10% and 40%.

V. Discussion

In the VAR model, a positive DR shock implies high future expected returns. In a rational framework, stocks with high expected returns earn a premium because of their exposure to priced risk factors. Risked factors are priced when investors want to hedge against those risks (for example, aggregate market shocks that could impact their level of consumption). Therefore, they demand a premium for holding that undesired risk in their portfolio. Following this interpretation, when a stock experiences a positive firm-specific DR shock, its exposure to some priced risk factor should increase (Ross, 1976).

The patterns we observe for the dynamics of discount rates of value stocks fit naturally into such framework. High book-to-market firms experience consistent increases in discount rates in the years before the holding period, translating into high expected returns. Those positive revisions in DR are also more substantial when including book-to-market as a state variable, consistent with the idea that the increase in expected returns of value firms is linked to this increase in exposure to the book-to-market factor.

Our results for momentum might be more challenging to reconcile in this risk factor setting. As stocks build up price momentum over the year before the holding period, their exposure to the momentum factor increases. If this factor implies greater risk, then stocks should experience an increase in discount rates in the period

leading to the holding period, followed by a drop in discount rates in subsequent periods once the risk exposure vanishes. We observe the opposite pattern, as past CF shocks drive the increase in price momentum, coupled with negative DR revisions, reflecting the negative correlation structure between CF and DR.

This observation could have multiple causes. First, price momentum might not be the “true” risk factor that would command a higher premium. The priced risk might be linked to other factors, like liquidity risk for example (Pástor and Stambaugh, 2003). Liu and Zhang (2014) propose a neoclassical explanation of momentum relying on investment and the higher marginal benefit of investment of winner stocks. Pazaj (2019) extends their framework in a q-theoretical set-up that allows to jointly explain the momentum anomaly and its negative correlation with book-to-market.

A risk explanation for the momentum anomaly which coincides negatively with momentum also coincides with our findings from chapter III. Indeed, we found that commonalities in CF shocks increase in loser and value stocks during recessions and aggregate CF stress. Conversely, commonalities in CF shocks increase in winner and growth stocks during periods of high investor sentiment and aggregate DR stress. This coincides with exposures to risk factors that materialize at different times and a negative correlation between the two anomalies. This kind of story fits naturally with the models of Avramov et al. (2016) or Wang and Xu (2015), where losers with the worst credit rating will drive the poor performance of momentum portfolios following recessions. It also fits with the model of Pazaj (2019), which predicts large betas of loser stocks during down markets.

A second cause could be that investors overreact to short-term updates in expected cash flows, as in the model of Barberis et al. (1998). Positive CF shocks being followed by positive CF shocks at first and a reversal, in the long run, is a hallmark prediction of behavioral models for the momentum effect. It would also go along with the negative correlation between CF and DR: an overreaction to CF shocks would imply lower expected returns.

A third potential explanation, might come from issues with the quarterly VAR model itself. Capturing price dynamics in a parsimonious model which relies on one-period lagged information might lack the precision to capture DR changes that occur at relatively high frequency, especially when most state variables are highly persistent and evolve slowly. This invites us to consider other approaches to model CF and DR shocks in future work to validate the robustness of the results. The approach of Chen et al. (2013) which seeks to model CF shocks directly using the implied cost of capital, might be a method worth considering for this endeavor.

Table VII panel D suggests that across all robustness specifications, high momentum stocks get negative DR shocks on average. One interpretation of this result is that high momentum stocks become less exposed to other priced risk factors, such as value or size. Indeed, this might be coherent with a fourth possible channel through which to reconcile the results. Conrad and Yavuz (2017) separate two types of stocks that get included into momentum portfolios: those who perform well during the holding period and are responsible for the premium, and those that revert. Momentum stocks that exhibit sustained reversal patterns are part of the latter group. On the contrary, firms that do generate the momentum premium do not revert and show a positive exposure to both value and size factors. Therefore, the negative DR shock of momentum that we document might be linked to the firms that revert and their lesser exposure to other priced factors. This type of story, where other risk factors drive momentum, resonates with the recent findings of Ehsani and Linnainmaa (2019), who argue that the traditional momentum factor is in reality driven by other factor’s momentum.

VI. Conclusion

This paper investigates the drivers of price changes leading to the formation of value and momentum portfolios. Value firms are long-term “losers” and are subject to consistent negative unexpected return shocks over the past five to ten years. Those surprises are driven by negative CF shocks on the one hand, which account for about 70% of the lower unexpected returns, and by positive DR shocks on the other hand, which make up the remaining 30% of the difference. Those positive DR shocks forecast high expected returns, even years into the future.

Momentum stocks follow a different pattern. By construction, those firms have the sharpest price increase over the past year and is the consequence of positive CF shocks, responsible for roughly 85% of the positive return, and negative DR shocks, which account for the remaining 15%.

In a rational framework, firm-specific increases in expected returns should reflect increases in exposure to some priced risk factor(s). It is straightforward to reconcile the long string of positive DR shocks of value stocks with increased exposure to priced risk, which ultimately commands a higher risk premium. Our results for momentum, however, pose more of a challenge. Despite getting negative DR revisions, winner stocks earn high returns, at least in the short term.

This CF return continuation pattern is not unique to momentum stocks: more

generally, stocks that were subject to the most positive CF revisions on average over the past four quarters earn higher returns. This finding is robust to different specifications of the VAR system that we use to model CF and DR shocks. Therefore, positive CF revisions forecast positive returns in the short-run and a reversal in the long run.

Those findings are consistent with usual conceptions of overreaction patterns. Furthermore, the consistent negative correlations with the value anomaly, echoing those documented in the literature, incite us to call for a unified framework accounting for the joint occurrence of both the value and momentum anomalies.

Finally, we investigate the role of media coverage in the patterns of CF and DR revisions of value and momentum stocks. Our first finding, is that negative news *tone*, which tends to be the case for value firms, correlates with negative CF news on the one hand and positive DR news on the other. The impact through the CF channel is roughly five times larger, suggesting that good (bad) news primarily impact prices by improving (deteriorating) CF expectations, and to a lesser extent by decreasing (increasing) risk and discount rates.

The news patterns we document coincides with the consistent negative CF and positive DR revisions of value stocks over multiple years. In the short run, momentum stocks are subject to positive news coverage, simultaneously with their one-year upward revisions in CF and negative DR shocks. News *tone*, however, does not suffice to explain the CF and DR patterns of value and momentum stocks, as those remain present even when controlling for past *tone*.

Chapter V

Private Equity and Employee Welfare

LAMBERT MARIE[‡] MORENO NICOLAS[‡] PHALIPOU
LUDOVIC[†]

APRIL 2022*

ABSTRACT

Using a comprehensive dataset made of over half a million ratings and reviews, this study documents the impact of corporate ownership changes on employee welfare. Overall, such organizational changes tend to have a negative impact on workplace satisfaction. However, we highlight important cross-sectional differences. Employees working at firms acquired through Leveraged Buy-Out (LBO) transactions experience the sharpest decline in welfare, especially if the company was previously publicly listed. Heterogeneity is also observed at the employee level; workers in non-managerial positions are more affected by such LBO transactions. Topic analysis of the written reviews highlights the LBO specific tension points: Post-transaction, it is complaints about layoffs, cost-cutting and lack of management care in particular that become more prevalent.

[‡]Marie Lambert and Nicolas Moreno are affiliated with HEC Liège, Management School of the University of Liège. Marie Lambert holds the Deloitte Chair in Sustainable Finance at HEC Liège. She is a Research Associate at EDHEC Risk Institute and a QMI research fellow.

[†]Ludovic Phalippou is affiliated to the University of Oxford Said Business School and is a fellow of the Queen's college.

*This version of the paper is based on the draft written by February 2021. A different introduction and conclusion was written by me for the purpose of this thesis. It is independent from the main version of the paper that is currently circulating, on which we later added a fourth co-author, Alexandre Scivoletto. We thank Gregory Besharov, John Gilligan, Elise Gourier, Thomas Hellman, Tim Jenkinson, Colin Mayer, Peter Morris, David Robinson, Sophie Shive, Sunil Wahal, Ayako Yasuda and seminar participants at HEC Paris, Oxford, and Wisconsin-Madison for useful comments and feedback.

Keywords: Big data, Crowdsourcing, ESG, Employee welfare, Private Equity, Leveraged Buy-Out, Merger & Acquisitions, Household Finance.

JEL Codes: G30, G34, G28, G50, I31, J21, J24, J31, J32.

I. Introduction

Today, companies owned by Private Equity Leveraged Buy-Out (PE/LBO)¹ funds employ a large portion of the US workforce: more than 11.7 million people work for such firms, generating \$1.4 trillion of gross domestic product, or 6.5% of the US GDP. It is also a fast-growing industry. The popularity of PE funds as investment vehicles rose quickly, going from just 1,000 to 20,000 funds over the past two decades. This growth is also reflected in the number of people directly affected by PE firms in their daily lives: in 2018, PE owned companies were employing 8.8 million people, an increase of 25% in just over two years. Additionally, another 7.5 million Americans work at supplier firms for PE owned companies².

As a result of its rise in prominence in our economic and everyday labor landscape, LBOs and PE funds have attracted increased interest and scrutiny from academics, professionals, news outlets, and the general public alike. In mainstream media, certain buyouts, having caused important layoffs³, have led PE firms to suffer from a generally poor reputation. Public officials have taken the chance to spin litigious situations⁴ related to PE firms as an opportunity to list their grievances with the industry. Accusing them of predatory practices, some politicians have even sought to regulate their field of action⁵.

Are these attacks about PE practices justified? This paper aims at answering this question through the lens of those often portrayed as the victims in LBO transactions: employees at target firms. Using about half a million written reviews about the companies they work for, we measure how PE firms affect their perceived welfare following LBO transactions.

Given the emphasis of politicians and media outlets on this particular outcome

¹We focus on Leveraged Buy-Outs (LBOs) because this is where the controversy lies. In an LBO, a fund acquires a controlling stake in a company and retains a significant oversight role until it sells its stake, three to five years later. Even though companies subject to an LBO are not technically “owned” by PE firms, we talk about PE ownership to refer to this form of ownership, for simplicity.

²www.investmentcouncil.org/economicimpact.

³Examples include the cases of “Shopko” (www.bloomberg.com/news/articles/2019-06-07/shopko-workers-latest-demanding-pe-firms-pay-up-as-chain-fails) and “Toys R Us” (www.ft.com/content/3d6ba4dc-ec6d-11e8-8180-9cf212677a57) that have received significant media coverage following thousands of job losses.

⁴See for example the recent dispute between singer Taylor Swift and private equity group Carlyle: www.ft.com/content/431f7d0d-f2a6-4f98-ae23-8dd877656307. Using this opportunity, house representative Alexandria Ocasio-Cortez went as far as claiming that “[...] leveraged buyouts have destroyed the lives of retail workers across the country, scrapping 1+ million jobs” (Twitter, @AOC, Nov 15, 2019, 4:49 PM).

⁵In July 2019, senator Elizabeth Warren led a bill entitled “Stop Wall Street Looting Act”. In October 2021, she reintroduced the bill: www.congress.gov/bill/117th-congress/senate-bill/3022/text

of LBOs, assessing how the key stakeholders that are employees fare in LBO transactions, has become an important concern in the literature. Overall, the picture depicted by academics appears significantly more nuanced and is driven by a general theme: not all employees are equally affected by PE firms.

Davis et al. (2014) find that LBO transactions lead to both job creation and destruction. Yet, their findings also suggest that this workforce reallocation goes hand in hand with a comparative increase in productivity and better wages. Underlining the efforts of PE firms to improve firm efficiency, Fang et al. (2021) also document that wage gaps decrease by cutting the salary of overly paid employees. Olsson and Tåg (2017) further support the notion that overall employment does not change, but that certain employees become more likely to lose their job than others: those performing routine and other easily offshorable tasks are twice as likely to be laid off by PE sponsors. Garcia-Gomez et al. (2020) also show that not all employees face equal consequences from LBO transactions, as those with poorer health conditions are more likely to lose their job, thus exacerbating pre-existing health issues. Finally, Davis et al. (2019a) document a slight increase in US jobs, which depends on the nature of the target firm, while Antoni et al. (2019) document a slight decrease in employment in Germany.

Not only do LBOs appear to lead to job creation at the expense of a high turnover, but the literature even documents improvements in practices that seem to benefit employees. For example, Agrawal and Tambe (2016) find that for employees at LBO target firms, both training and employability increases. LBOs also come with improvements in workplace safety (Cohn et al., 2021) and managerial practices (Edgerton, 2012; Bernstein and Sheen, 2016b).

We differ from all the studies above by investigating how employees perceive the change in ownership. We leverage data from a website called Glassdoor, where employees can post ratings and extensive reviews about their perceived experience at firms they work for. We obtain a rich dataset, spanning all types of company sizes, industries and ownership types. We collect over 700,000 ratings across multiple welfare metrics (perceived overall score, work-life-balance, career opportunities, opinion on senior management, future company outlook, etc.), almost half a million written reviews where employees detail the “Pros” and “Cons” about their workplace as well as “Advice to Management”, and over two million self-reported salary entries.

Using a difference-in-difference approach, we find that changes of ownership type lead to a decrease in employee satisfaction. On a scale of 1 to 5-star ratings, employees report significant drops in score of 0.13 following Initial Public Offerings (IPO)

and 0.10 for non-LBO Merger and Acquisitions (M&A).

Employees seem to dislike organizational changes following LBOs even more, but the pre-treatment type of ownership of the target firms matters: If the target was publicly listed, employees report a decrease in score of 0.26; if it was privately owned the score falls by 0.15; if it was already under PE control the drop in score aligns with other M&As at 0.10. To put those numbers into perspective, the average gap in score between employees in managerial positions and non-managers (excluding mid-level managers) is 0.29.

A main contribution of our study is that it provides unique insights into the nature of the problems that may arise for employees following LBOs. We achieve this by analyzing the content of the written reviews using tools borrowed from the field of machine learning. We use Latent Dirichlet Allocation (Blei et al., 2003) as a standard natural language processing technique to map the content of the hundreds of thousands of reviews along a parsimonious set of 25 topics.

Each topic captures a high-level representation of a different set of complaints that employees might express about the companies they work for. For example, consider the topic *“Layoffs & Cost-cutting”*. A review gets a quantitatively estimated relevance to the topic depending on how strongly it loads on the group of terms including {company, cut, employee, layoff, people, job, cost, business, leave, ... }.

This allows us to identify the main drivers of employee dissatisfaction, both across employees and following ownership changes. Overall, the complaints being the strongest predictors of poor employee ratings are *“Badly treated by management”* and *“Lack of care from management”*. Those two topics underscore the critical role of managers in the perceived well-being of their employees. A hypothetical review made up from text 100% drawn from those topics would get a score which is respectively 1.23 and 1.15 lower than the average review.

Those complaints are followed by *“Bad HR management”* (-0.92) and *“Poor upper management communication”* (-0.87) as the next most impacting topics, further cementing the role of managers in employee welfare. Complaints about *“Changes in leadership”* (-0.72) and *“Layoffs & cost-cutting”* (-0.47) are next in line; while on the other end of the spectrum, *“Fast changing and growing company”* (+0.94) and *“No complaints”* (+1.12) are the topics that are the most likely to appear in the “Cons” section when rating scores are above average. Overall, work load and work pace seem to be issues that matter only once core problems are out of the way, such as abuse from and conflict with management.

We observe that LBOs cause certain topics to become more prevalent, while

others appear less often post- transaction. *“Layoff and cost-cutting”* is the topic that is the most likely to become more prominent after a transaction. This is also true for M&As, but not for IPOs. Two more themes of complaint are particularly recurrent post-LBO: *“Changes in leadership”* and *“Lack of care from management”*. Both those increases are specific to LBO-related changes in ownership. Given that these three topics are all among those that negatively affect employee ratings the most, we are now better able to understand where the specific drivers of LBO dissatisfaction come from.

LBOs might, however, improve employee welfare in other dimensions. The content of reviews post-LBO hints towards improved operational processes. Complaints about *“Internal Politics & lack of diversity”* as well as *“Lack of communication & slow processes”* occur less frequently. Other gripes related to *“Benefits in kind”* and *“Overtime”* are also less frequent when a new PE sponsor takes over. Those findings resonate with the literature showing that operational processes at PE owned firms tend to get optimized (Edgerton, 2012; Davis et al., 2014; Bloom et al., 2015). Bernstein and Sheen (2016b) for example, show that operational processes brought in by PE firms in restaurants they own improve cleanliness and safety.

Beyond operational processes, a large body of literature contrasts with the bad reputation of PE firms documented above, by showing that target firms improve across a wide range of metrics: profitability increases (Kaplan, 1989; Guo et al., 2011; Cohn et al., 2016), total factor productivity improves (Davis et al., 2014), growth is bolstered (Cohn and Wardlaw, 2016; Bernstein et al., 2017; Boucly et al., 2019), access to external finance is facilitated (Boucly et al., 2019), product variety for consumers increases (Fracassi et al., 2019), resilience to economic downturns is strengthened (Bernstein et al., 2019), and innovation through patent activities thrives (Lerner et al., 2011).

Note that some papers investigated whether the improvements across all those performance metrics is achieved at the cost of other stakeholders or through negative externalities. For example, Shive and Forster (2020) find that companies managed by PE sponsors tend to pollute more. However, Bellon (2020) argues that PE firms try to mitigate toxic pollution and carbon emissions to maximize company value at exit⁶. In public sectors, such as education or nursing homes, PE-owned firms were found to maximize their profits at the expense of other stakeholders (Eaton et al., 2020; Gupta et al., 2021). Taxpayers and beneficiaries of those services (elder residents, students) appear to be on the losing side of the deal. Furthermore, Sheen

⁶This may also loosely align with the findings of Barber et al. (2021), who find that investors are willing to sacrifice performance in order to invest in “impact funds”.

et al. (2021) point to an increase in financial advisor misconduct after LBOs.

Above, we mention that not all employees fare equally in LBOs. Consistent with this idea, we find significant cross-sectional heterogeneity in complaints across employees in our reviews. We document that people in non-managerial positions are particularly likely to complain about “*Layoffs & cost-cutting*”, “*Changes in leadership*” and “*Lack of care from management*”. Those topics are all among the strongest predictors of negative reviews, and thus the dimensions that affect regular workers the most following LBOs are critical to their welfare. Managers, on the other hand, are more likely to write about topics such as “*No complaints*” post-LBO. This result highlights an increase in the welfare gap between the two groups, which is already pre-existing to PE interventions.

We find other cross-sectional differences, such as specific effects related to both industries and individual PE sponsors. The two sectors where the drop in ratings post-LBO are the most significant are “Software” and “Retail”, consistent with some stories that made it into the mainstream media. It is also among those industries that complaints about “*Layoffs & cost-cutting*” and “*Changes in leadership*” are the most frequent after a deal. There is a lot of heterogeneity across PE sponsors, too. Some tend to be associated with increases in ratings, whereas negative reviews tend to be especially concentrated in a few others, who tend to be specialists in the software and tech industry. Differences are also marked across LBO types. Public-to-PE transactions cause more issues with “*Layoffs & cost-cutting*”, but also “*Lack of care from management*”. Private-to-PE deals tend to have more problems with “*Changes in management*” and “*Upper leadership*”. Once again, PE-to-PE transactions are those where the observed effects are the weakest.

The cross-sectional differences across industries, firm characteristics, and employees that we document could also be reconciled with the findings of Gornall et al. (2021), who also leverage Glassdoor data to understand changes in employee satisfaction following LBOs. With a specialized dataset of theirs they find that firm performance and leverage are drivers of firm-level heterogeneity. However, our paper crucially differs from theirs in that we are uniquely able to understand the mechanisms driving employee dissatisfaction following LBOs thanks to our textual analysis. Gornall et al. (2021) also manage to replicate the finding that employee position plays a key role in changes of satisfaction following LBOs, but our procedure really shines in that we are able to explain the underlying drivers of those changes, with great nuance.

The remainder of the paper is structured as follows. Section II describes our

data sources and provides descriptive statistics. Section III contains panel regression analyses. Section IV is dedicated to the textual analysis of written reviews. Section V discusses the implications of our findings and concludes. The Appendix provides a discussion of the related literature, gives examples of reviews, and provides additional analyses.

II. Data and Descriptive Statistics

A. Glassdoor Website

Glassdoor is an employer review website launched in June 2008. Company ratings, reviews, and salaries are entered by employees and are displayed anonymously for all members to see. Most reviews are written by new users who need to submit information about their employer before accessing other people’s ratings, reviews, and salary benchmarks (see Appendix E and Green et al. (2019) for more details).

The website verifies that each review is genuine through checking of e-mail addresses, social networking accounts, various fraud-detection algorithms, and through screenings by a content management team.⁷ Green et al. (2019) provide a comprehensive description of the dataset, along with several external validity tests.

This dataset has been used in several academic studies, and the findings are that Glassdoor ratings are useful to predict key accounting-based information such as i) growth in sales, profitability, and net income; ii) Tobin’s q , and Return on Assets; and iii) earnings announcement surprises.⁸ In addition, similar to the finding of Edmans (2011) who used a different data source for employee welfare, Green et al. (2019) find that Glassdoor ratings predict subsequent stock returns.

Prior literature thus suggests that crowdsourced employee ratings contain important and relevant information, rather than being a collection of idiosyncratic opinions. In addition, we expect employees to provide honest evaluations due to the benefits associated with contributing to the public good (Lerner and Tirole, 2003). Examples of Glassdoor reviews are shown in Appendix F.

Finally, as Glassdoor effectively starts receiving reviews from 2012 onward, the reviews in our sample are from July 2012 to June 2020, and the transactions we use occurred between January 2013 and December 2019.

⁷In 2013, the company stated that it rejects about 20% of entries after screening. Source: <http://www.calgaryherald.com/business/Website+lets+workers+rate+their+bosses+anonymously/8221492/story.html>

⁸Green et al. (2019); Babenko and Sen (2014); Hales et al. (2018); Huang (2018b); Huang et al. (2015b)

Table I: Sample Selection - Panel A shows the number of companies selected in Capital IQ split across five categories: companies that experienced either an LBO, an IPO or an M&A (other than LBO) between 2013 and 2019, and companies that stayed private or public from 2011 to 2020. The next column shows how many of these companies could be found on glassdoor.com. Panel B shows the different filters we applied to the sample described in Panel A, to end up with our working sample. Panel C shows the number of LBOs as a function of the ownership before the transaction: company was publicly listed (Public-to-PE), was already under PE ownership (PE-to-PE), was privately held but not controlled by a PE firm (Private-to-PE).

Panel A: Initial Sample				
	Number of companies in Capital IQ	Of which, matched to Glassdoor	Number of Deals	Number of Ratings
<i>With ownership change between 2013 and 2019</i>				
Leveraged Buy-Outs (LBO)	3,706	2,143	2,302	205,603
Initial Private Offering (IPO)	1,166	595	595	125,889
Mergers & Acquisitions (M&A, not LBO)	8,705	2,887	2,934	280,948
	13,577	5,685	5,831	647,440
<i>Without ownership change between 2013 and 2019</i>				
Stayed Private	25,219	8,997	8,997	869,999
Stayed Public	3,182	2,174	2,174	985,515
Total	41,978	11,171	11,171	2,467,954

Panel B: Data Filters										
	Number of Deals					Number of Ratings				
	LBO	M&A	IPO	Private	Public	LBO	M&A	IPO	Private	Public
Initial sample	2,302	2,934	595	8,997	2,174	205,603	280,948	125,889	869,999	985,515
<i>Sample after removing</i>										
Former employees + Interns	2,120	2,642	556	8,152	2,095	114,822	149,774	74,201	465,584	524,229
Missing length of employment	2,006	2,383	540	7,576	2,024	76,468	93,852	50,883	289,687	341,182
Those who joined post transaction	1,832	2,247	457	7,576	2,024	49,100	73,332	27,650	289,687	341,182
Ratings falling outside a 6 year period centered on transaction day	1,722	2,185	438	7,576	2,024	37,266	53,551	20,561	289,687	341,182
<i>Working Sample:</i>										
Companies with more than two Ratings pre or post transaction	631	602	245	3,796	1,559	31,767	44,361	19,704	279,700	339,843

Panel C: Type of LBO		
	Number of LBOs	Number of Ratings
Public-to-PE	94	10,498
Private-to-PE	281	9,562
PE-to-PE	256	11,707
Total	631	31,767

B. Capital IQ

We use Capital IQ to generate a list of US-based companies. We separate the set of companies into five groups. In the first group are companies that were privately held throughout our sample period and without a change of ownership. As shown in Table I, this group contains 25,219 companies.⁹ In the second group are companies

⁹We require a minimal revenue of \$50 million. We use revenue because Enterprise Value (EV) is usually not available for private companies, and a \$50 million revenue coincides on average to a

which are publicly traded throughout our sample period. This group contains 3,182 companies.

In the third group are companies that experienced a Private Equity sponsored Leveraged Buy-Out transaction (simply referred to as LBO). To form that group, we follow Davis et al. (2014, 2019a) and obtain a sample of 3,706 LBO targets.¹⁰

In the fourth group are companies that went in the opposite direction: from being privately held to being publicly held. Finally, the fifth group contains companies that went through an M&A and are not present in either the LBO or the IPO sample. In this M&A sample, companies have also experienced a change in ownership, but have not experienced an LBO. The Initial Public Offering (IPO) sample contains 1,166 companies and the M&A sample contains 8,705 companies.¹¹ In total, we have 41,978 companies, which we seek to match with the Glassdoor dataset.

C. Working Sample

Table I - Panel A shows statistics when we merge Capital IQ and Glassdoor datasets. About half of the companies are matched based on their name and address.¹² In a similar exercise, Davis et al. (2014) match 65 percent of LBO-targets to the Census Bureau’s Business Register, which is the same rate we have for our sub-sample of LBOs. Most of the unmatched companies are small, which, by definition, are less likely to have a Glassdoor page.

Next, we apply several filters to this initial sample, as shown in Table I - Panel

\$100 million EV.

¹⁰We select M&A transactions with a PE firm as a financial sponsor, and which have one of the following features: “going private,” “leveraged buyout,” “management buyout,” or “platform.” We manually check each transaction to ensure sample integrity (e.g., making sure to exclude startup firms backed by venture capitalists, management buyouts that are not private equity sponsored, transactions without a change in control). See Davis et al. (2019a) for a thorough discussion on how to select LBOs in Capital IQ and why Capital IQ, over our sample period, is best suited for such an exercise.

¹¹We require the IPO to occur i) between January 2013 and December 2019, ii) on one of the three major US stock-exchanges, and iii) with a deal size of \$100 million or more. Reverse LBOs are not included. For example, Gardner Denver Holdings, which was publicly listed until 2013, a time at which it was subject to an LBO sponsored by KKR, and partially exited via IPO in 2017, is in the LBO sample from 2010 to 2016, and not in the IPO sample. We require transaction value (or deal value) of the M&A to be above \$100 million. Smaller companies are unlikely to have a page in Glassdoor. We only keep firms that went through a single M&A over our time period.

¹²For about 40% of the matched companies, we have an exact match on name (i.e., only one company entry exists in Glassdoor for this firm name). For the remaining companies, we have multiple possible Glassdoor matches and choose the one using HQ city, state of incorporation, country of incorporation, year of incorporation and website. The more matches we have on those meta-data, the greater our matching confidence. We found it reasonable to keep firms that matched on country and at least one other of the above criteria, given name similarity.

B. First, reviewers are required to state whether they are currently working for the company or not. 44% of the reviews are from people no longer employed at the company at the time they submit their review. As we cannot tell whether the review refers to the situation pre or post a given change of company ownership, we exclude all reviews from former employees in the main analysis. We also exclude ratings posted by interns.

Second, reviewers may select from a menu the number of years they have been working at the company: less than a year, more than a year, more than 3 years, more than 5 years etc. This information enables us to determine whether the reviewer joined the company before or after a transaction. We require this information; thereby losing one third of the ratings.

Third, we exclude observations from employees who joined the company post transaction as they cannot compare the situation before and after a transaction; thereby losing 21% of the remaining ratings. In the robustness section below, we show results when we add back the observations we took out.

Fourth, we keep only observations that fall outside the three years around the day of the transaction; and require at least three reviews pre- and post-transaction.

This last requirement generates a large drop in number of observations because we are often missing pre-transaction ratings. Nearly all divisional buy-outs do not have pre-LBO ratings because divisions are usually not treated as separate entities in Glassdoor. Moreover, company names may change pre- and post-transaction, which fails our matching process; and there are fewer ratings at the beginning of the sample and therefore LBOs made in years 2013-2015 are more likely to have an insufficient number of ratings pre-LBO. For companies with no change in ownership, we require at least six reviews over the entire sample period, but this requirement hardly affects the number of observations.

Our working sample counts 631 LBO transactions, for which 31,767 ratings were submitted. About as many companies are in the M&A sample, but with slightly more reviews on average. Fewer companies experience an IPO. In total, there are 1,478 companies that experience a change in ownership during our sample period.

We also have two samples of companies that did not experience a change in ownership. 1,559 companies are publicly listed throughout our sample period. On average, they have 218 reviews each, for a total of 339,843 reviews. 3,796 companies are privately held. They tend to be smaller, with 74 reviews per company on average, for a total of 279,700 reviews. Our overall dataset, i.e., the five samples described above together, consists of 715,375 employee reviews.

D. Type of LBOs

The social impact of LBOs may differ as a function of the type of ownership the company was under before the transaction occurred. Several studies indicate that Public-to-PE transactions might be those with the highest social cost. These transactions are more likely to result in job losses (Davis et al., 2019b), to go through bankruptcy procedures (Strömberg, 2009), and to face a higher debt burden (Axelson et al., 2013). In addition, privately-held firms are likely to benefit more from the relaxation of financial constraints and improvement of management practices that PE ownership brings (Boucly et al., 2019; Lerner et al., 2011). On the other hand, Cohn and Wardlaw (2016) find that workplace injury rates fall after Public-to-PE LBOs but not after Private-to-PE LBOs. They explain this result by public companies having more pressure for short-term performance.

Table I - Panel C shows the breakdown across the three types of LBOs. Only 15% of the LBOs are Public-to-PE, but represent 33% of the ratings, as they tend to be relatively large companies. We should also distinguish between companies that already had PE ownership before the transaction because employees are already used to PE ownership style. When the company targeted in an LBO was already under an LBO, the transaction is called a secondary Buy-Out (Arcot et al., 2015; Degeorge et al., 2016). In this paper, we refer to these as PE-to-PE; they represent 37% of our sample of LBOs.¹³

E. Glassdoor Ratings

Employees anonymously assign a one- to five-star rating for i) the Company, ii) Career Opportunities, iii) Compensation & Benefits, iii) Work/Life Balance, iv) Senior Management, and v) Culture & Values; and assign one of three ratings for i) Business Outlook, ii) Recommendation of the company and iii) approval of the CEO. In addition, reviewers enter in an open field the pros and cons of working for the company, and their recommendation to the management.¹⁴

The ratings are used in the regression analysis presented in section III. Table II shows related descriptive statistics. We observe that they are similar to those

¹³In many transactions, some of the previous owners are PE firms, but the development stage is venture capital rather than an LBO. For example, consider the acquisition of Acquia, sponsored by Vista Equity Partners, in September 2019 for \$1 billion. The sellers are venture capital firms (Sigma Partners, North Bridge Venture Partners, and Underscore Venture Capital). The preceding transactions are all standard venture capital funding rounds: \$55 million in September 2015, \$50 million in May 2014. This transaction is classified as Private-to-PE and not PE-to-PE because the preceding ownership type is not LBO.

¹⁴The open-field reviews are used in the textual analysis covered in Section IV.

reported in the literature¹⁵. The average score across the 4.8 million reviews is 3.49.

The sub-scores are available for nearly all the observations. They are widely distributed with the strongest dispersion observed for Recommendation, Outlook, Senior Management (SM), and Culture. As argued by Green et al. (2019), a wide dispersion is a sign that Glassdoor contains a full spectrum of views, and not only those of disgruntled employees. We also note that the average reviewer has a positive opinion of the CEO, of the outlook, and would recommend the company.

The correlations between the different ratings are high but are all below 80%, showing that people made a distinction between the different scores, rather than assigning the same rating to each category. The score for Career Opportunities is the most correlated with overall score. Furthermore, the extent to which employees recommend the firm is strongly correlated with their overall score. Work/Life Balance and CEO opinion are the least related to the overall score.

F. Glassdoor Job Titles & Salaries

On a separate page of the Glassdoor website, reviewers enter a salary along with a job title. They must enter salary information on their company in order to access salary benchmarks at this and other companies. This process is separate to the review process. We cannot directly link salaries to ratings, nor can we distinguish between former and current employees.

In addition, Glassdoor aggregates the information, and reports only the average salary for a given job title. That said, job titles are so granular that the information loss is probably minimal. For example, truck-drivers at “The Kraft Heinz Company” have an average salary of \$41,108.

We have 2,243,236 pairs of position-salary across the 21,817 companies in our sample. We then use textual analysis tools and the guidebook ‘Work in America’, page 597, – as detailed in Appendix I – to assign each job title to one of the following five job categories: i) Management; ii) White collars, which are middle management and professional service providers (consultant, researchers); iii) Purple collars, which are technical service providers; iv) Pink collars, which are support staff; and v) Blue collars. 11% of the job positions could not be classified into one of these categories.

Table III shows average salaries. Blue collars earn the least at \$42k and pink collars earn only slightly more at \$45k. There is a jump to \$66k for purple collars

¹⁵Our average Score is slightly higher than the one reported in other studies (e.g., Green et al. (2019)), because i) we exclude former employees from our sample, and these people give lower ratings and ii) Glassdoor ratings exhibit a positive trend over time, and that our sample contains more recent data.

Table II: Quantitative Ratings Given by Employees

There are nine different quantitative ratings given by employees to their company. The first six (Overall, WLB, CO, CB, SM, Culture) are either 1, 2, 3, 4 or 5; the next three (CEO, Outlook, Reco) are either 1, 3 or 5. Panel A shows descriptive statistics across our working sample. Panel B shows the pairwise coefficient of correlation between the score across our working sample.

Panel A: Descriptives									
	mean	std	Number of Scores						
Overall Score (Overall)	3.52	1.28	715,375						
Work-Life-Balance (WLB)	3.42	1.34	680,509						
Carreer Opportunities (CO)	3.52	1.41	661,274						
Compensation & Benefits (CB)	3.13	1.42	670,146						
Senior Management (SM)	3.31	1.35	678,432						
Culture	3.40	1.26	679,396						
CEO	3.79	1.48	566,093						
Outlook	3.70	1.58	607,512						
Recommended (Reco)	3.66	1.89	633,902						

Panel B: Correlations									
	Overall	WLB	CO	CB	SM	Culture	CEO	Outl	Reco
Overall	1.00								
WLB	0.63	1.00							
CO	0.78	0.60	1.00						
CB	0.78	0.60	0.76	1.00					
SM	0.75	0.51	0.67	0.69	1.00				
Culture	0.63	0.47	0.55	0.56	0.60	1.00			
CEO	0.62	0.44	0.60	0.61	0.53	0.45	1.00		
Outlook	0.67	0.46	0.61	0.62	0.59	0.47	0.58	1.00	
Reco	0.75	0.53	0.68	0.66	0.63	0.51	0.57	0.64	1.00

and a further increase to \$76k for white collars. For positions, we classified as management, the difference in salary is substantial at \$138k, which is nearly twice that of the average salary of white collars. Some reviewers enter a salary but a job title we could not classify (11%), the corresponding average salary is \$55k. A caveat is that no dates are entered. As salaries are unlikely to have changed much during our sample period, this should be a minor issue. This analysis gives comfort that our code to categorize job titles produces sensible results.

Table III: Job Categories

Employees may enter a job title which we then classify into management, middle management, white collar, purple collar, pink collar, or blue collar. For each company, Glassdoor has a page where it gives the average salary for different job positions. Panel A shows the average and inter-quartile across companies within each of the six job categories. 250,405 job titles could not be assigned to one of the six categories and are not included in these statistics. Panel B shows the average score across all the ratings given by employees in a given job category. 908,242 ratings were entered without a job title and 509,607 ratings were entered with a job title we could not classify. The average scores for these two 'not classified' categories are shown in the last columns. The number of reviews within each job category refers to all the ratings observed, not the number of employees or reviews.

Panel A: Salaries							
	TopMngt	MidMngt	WhiteC	PurpleC	PinkC	BlueC	NC
Avg. Salary	\$137,718	\$84,507	\$76,090	\$66,091	\$45,400	\$42,258	\$54,493
25 th percentile	\$103,222	\$54,708	\$50,400	\$40,409	\$23,520	\$21,840	\$26,880
75 th percentile	\$166,323	\$109,061	\$96,576	\$85,683	\$57,192	\$52,080	\$68,880
Number of Entries	170,217	409,384	751,303	156,342	326,432	177,433	250,405

Panel B: Scores								
	Mngt	MidMngt	WhiteC	PurpleC	PinkC	BlueC	NC	Anon
Overall Score	3.79	3.58	3.60	3.47	3.43	3.34	3.38	3.54
WLB	3.55	3.34	3.57	3.40	3.37	3.20	3.30	3.46
CO	3.79	3.59	3.62	3.43	3.46	3.29	3.37	3.52
CB	3.42	3.20	3.21	3.03	3.07	2.90	2.97	3.16
SM	3.63	3.46	3.40	3.23	3.17	3.09	3.17	3.31
Culture	3.69	3.48	3.49	3.36	3.24	3.22	3.28	3.44
CEO	4.07	3.87	3.90	3.67	3.65	3.48	3.64	3.85
Outlook	3.97	3.77	3.80	3.62	3.57	3.48	3.55	3.74
Reco	3.97	3.72	3.77	3.62	3.55	3.44	3.47	3.67
Number of Scores	234,473	783,328	1,290,180	346,426	733,528	406,346	509,607	908,242

We cannot study whether salaries decrease post LBO as we do not have information on the salary for each reviewer and the salary information is not time stamped. However, by matching the separate salary dataset with our sample of 715,375 reviews, we can study whether a decrease in ratings is stronger below an approximate salary threshold, which is more precise than using the average salary per job category. Among the 588,148 reviews for which we have the reviewer's job title, 357,814 reviews could be matched with a position-salary pair¹⁶. For the remaining reviews, we could not match the job position from the review with a salary-position pair.

¹⁶For 349,947 reviews we have an exact match on the reported position name within the company. For 7,867 we have a position name within the company that resembles very strongly the job position in the review, which we retrieve by fuzzy text matching.

We proceed as follows: (i) we assign the average salary of the corresponding collar category within the firm or the average salary of the “Not classified” job category if the job position can’t be classified into the five categories (205,805 reviews), (ii) we assign the average salary within the company in case salary information is missing for the corresponding collar category, (11,314 reviews) or (iii) we assign the average salary of employees within the same industry and collar category if the company has no salary entry in the salary dataset (13,215 reviews). Anonymous ratings were assigned the company average company salary or average industry salary in case of missing information.

Table III - Panel B shows the different ratings per job category. We see a perfect line up between job category and ratings. Job categories that command higher salaries show the highest ratings on each dimension including work-life-balance, although mid-management gives a relatively low rating to work-life-balance.

We observe that senior management is well rated by management, then mid-management and white collars give a similar lower rating, and the purple, pink and blue collars give a much lower rating. The CEO opinion is also directly related to job hierarchy.

G. Industries

A unique feature of Glassdoor is that each company is assigned to one of 121 industries, even those that are privately owned. 91 industries have more than 1000 ratings for LBOs. We therefore pool industries together into the following seven large industry categories: 1. Consumer Services (restaurants, leisure), 2. Business services (Finance, Insurance, Consulting, Staffing, Marketing), 2. Public services (healthcare & education), 4. IT services, 5. Industrial (manufacturing, pharmaceutical), 6. Retail (pet stores, department stores), 7. Software.

As shown in Table IV - Panel A, the numbers of ratings are well distributed across these seven industries. A partial exception is Software. Although Software is our most narrowly defined category, it is the one with the most deals (22% of the LBOs) and ratings (28% of the ratings). Software also stands out in terms of salaries; the \$93k average salary is nearly twice as much the average in the Retail industry, at the other side of the spectrum. The average rating is also the highest in the Software industry.

These statistics contrasts with the view that LBO targets are value companies, which is certainly the case pre-2010. Over the most recent decade, Tech LBOs have become increasingly common. Various data providers report that Tech LBOs make

Table IV: Company Characteristics

This table shows descriptive characteristics for three company characteristics: industry in which the company operates (Panel A), foundation year (Panel B) and number of employees (Panel C). Data on foundation year and size are from Glassdoor. Size is the number of employees reported on the website when we scrapped the data. We call post-war companies with foundation year between 1945 and 1999. Millennials are companies founded in 2000 or later. Industry categories are created from the 121 different industry classifications used by Glassdoor. We report the number of companies and scores observed in each category across the overall sample of companies and across the sub-set of companies subject to an LBO.

Panel A: Job Position						
	Number of Companies	Number of Ratings	Avg. Overall Score	Avg. Salary	Number of LBOs	Number of LBO Ratings
Consumer Services	635	72,284	3.55	\$45,824	39	2,485
Corporate Services	1,793	183,565	3.55	\$67,214	138	5,202
Public Services	420	37,942	3.49	\$56,902	77	3,332
IT Services	335	51,707	3.58	\$81,772	54	3,656
Industrial	2,146	161,275	3.46	\$71,582	119	3,062
Retail	500	89,048	3.34	\$36,983	43	4,939
Software	546	73,751	3.74	\$93,023	160	8,996
Overall	6,375	669,572	3.53	\$64,757	630	31,672

Panel B: Foundation year						
	Number of Companies	Number of Ratings	Avg. Overall Score	Avg. Salary	Number of LBOs	Number of LBO Ratings
Millenials	663	46,683	3.71	\$79,021	175	6,907
Post War	1,597	227,518	3.48	\$65,446	388	21,515
Pre War	674	150,601	3.41	\$67,314	44	2,142
Overall	2,934	424,802	3.53	\$70,594	607	30,564

Panel C: Size						
	Number of Companies	Number of Ratings	Avg. Overall Score	Avg. Salary	Number of LBOs	Number of LBO Ratings
Small (1-500)	1,598	32,327	3.81	\$61,255	168	3,439
Small (501-1000)	1,238	36,542	3.69	\$65,260	149	4,632
Mid (1001-5000)	2,395	144,774	3.58	\$65,954	237	12,673
Mid (5001-10000)	625	90,873	3.47	\$68,968	37	3,736
Big (10000+)	977	410,859	3.47	\$63,826	40	7,287
Overall	6,833	715,375	3.60	\$65,053	631	31,767

up one third of the deals in the 2010s, which is the proportion we have in our sample (adding up IT services and Software). Excluding Software, the other six industries have between 8% and 16% of all the ratings each, and have similar average ratings. Salary varies significantly, though. Salaries in IT services average \$82k and those in Consumer Services average \$69k.

H. Other variables

Glassdoor provides the year of foundation of the company. The 75th percentile is close to 2000; companies created after 2000 are labelled “Millennials”. The 25th percentile; companies created before 1945 are labelled “Pre-War”. Companies that are in the LBO sample tend to be older. There are only 28% of Millennials among LBO targets.

Glassdoor provides a range for the current number of employees, which we re-group into Small (less than 500 employees), Medium, and Large (more than 5,000 employees). More than half of the reviews are from large companies, even though there are about three times as many small companies (2,836) as there are large companies (977). Small companies have less reviews and make up for only 10% of the set of reviews.

III. Regression Analysis

In this section, we analyze how ratings change around LBO transactions. We detail our empirical strategy and test for pre-trends, study drivers of ratings, look at characteristics associated with a stronger LBO effect, and show robustness tests.

A. Empirical Strategy and Pre-trend Evaluation

Our econometric approach to estimate the effect of LBOs on employee welfare follows the recommendations of Petersen (2009): our panel is estimated by pooled OLS with both quarter fixed effects, and company fixed effects. Statistical inference is based on standard errors that are double-clustered at the company level and at the quarter level. Our main specification is as follows:

$$\begin{aligned} S_{r,c,d} = & \alpha_c + \alpha_{q(d)} + \beta * PostLBO_{c,d_e,d} \\ & + \theta_1 * PostM\&A_{c,d_c,d} + \theta_2 * PostIPO_{c,d_c,d} \\ & + \gamma * Z_r + \epsilon_{r,c,d} \end{aligned} \tag{V.1}$$

The dependent variable is the rating $S_{r,c,d}$ given by a reviewer r to its company c on a day d .¹⁷ $q(d)$ is the calendar quarter in which day d falls into. $PostLBO_{c,d_e,d}$

¹⁷Note that each rating/review is treated as being submitted by a separate reviewer. It is possible that the same person submits several reviews and ratings over time, but we cannot identify people.

takes a value of one if the company c has been subject to an LBO transaction that closed on a day $d_c < d$.

Company fixed effects (α_c) and quarter fixed effects ($\alpha_{q(d)}$) absorb the multitude of company invariant factors (e.g. industry) and time invariant factors (e.g. recessions). This specification is thus a within-company and within-quarter regression. The only control variables that are left are therefore the characteristics of the reviewer: Z_r .

The coefficient of interest in this model is β , which captures the relationship between PE ownership and employee ratings. Specifically, as there is a company fixed effect, β measures the incremental score given by employees following an LBO compared to the average score given to this company at any point in time. In addition, due to the time fixed effect, the scores are corrected for the average score given in that quarter across all companies. Having a quarter fixed effect is important because employee ratings are expected to vary significantly over business cycles. To sum up, the coefficient of interest, β , measures the rating given post LBO in excess of the average score for this company at any point in time and to the average score for any company in that quarter.

As pointed out in the literature and in particular by Davis et al. (2019b), LBO transactions are heterogeneous. Also, the contrast between the existing academic evidence and case studies covered in the media highlight that not all LBO transactions might have the same effect on employee welfare. A key contribution of this paper is to test whether some type of LBOs have systematically different effects on employee welfare than others and whether different types of employees react differently to LBOs. In addition, studying impact heterogeneity helps to address both potential sample selection biases and some endogeneity concerns.

To study impact heterogeneity, we split the *Post_LBO* variable in equation (1) into different types of LBOs, different types of companies, different types of employees etc. The split is achieved using a set of dummy variables labelled *SubTypes*:

$$S_{e,c,q} = \alpha_c + \alpha_q + \Sigma(\beta_s * PostLBO_{c,q} * SubTypes) + \gamma * Z_e + \epsilon_{e,c,q} \tag{V.2}$$

Where $\Sigma(SubTypes * PostLBO_{c,q}) = PostLBO_{c,q}$.

Since our setup is a standard difference-in-differences design, we must assume that both the companies that are subject to an LBO and the other companies, were on parallel trends before the LBO in order to interpret β as the causal effect of PE

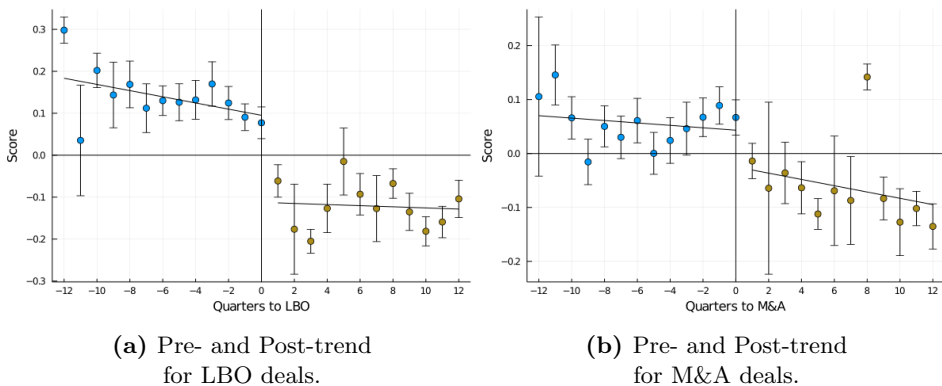


Figure V.1. Pre- and Post-trend

ownership.

We estimate a model that is similar to that shown in Equation 2. We replace the PostLBO dummy variable by a PreLBO dummy variables and the cross effects (SubType) are for each of the twelve quarters preceding the transaction; as shown in Equation 3 below:

$$\begin{aligned}
 S_{e,c,q} &= \alpha_c + \alpha_q + \Sigma(\beta_s * LBO_Quarter_{c,s}) \\
 &+ \gamma * Z_e + \epsilon_{e,c,q} \\
 s &= [-12, -11, \dots, -1]
 \end{aligned}
 \tag{V.3}$$

$LBO_Quarter_{c,s}$ is a dummy variable that is one if the rating for company c was submitted on a day that was between $d_c + 90 * s$ and $d_c + 90 * (s + 1)$, and is zero otherwise. As shown in Figure V.1, the time-series of β_s does not exhibit any time-trend at any horizon. On the figure we also plot the post LBO quarter fixed effects and observe no trend post LBO either.

B. The Cross-Section of Ratings and PE ownership

Table V shows the results from the estimation of the panel regression model in Equation V.1. The largest decline is observed for companies that were publicly listed before PE-ownership (-0.25). When the company was privately held, but was not already under LBO, we observe a decline that is smaller: -0.15. When the company was privately held but by PE, the decline is less: -0.12.

Companies that are subject to other ownership changes also experience a decrease

Table V: Determinants of Company Scores. This table shows the results from estimating pooled panel regressions with the (overall) score as the unit of observation and as the dependent variable. Standard errors are clustered at the company and quarter level. When a company is subject to two LBOs in our dataset, it is treated as two separate companies for the purpose of fixed effects and clustering of standard errors.

	Dependent variable: Overall Score			
	(1)	(2)	(3)	(4)
Public-to-PE	-0.26*** (-5.86)	-0.26*** (-5.97)	-0.26*** (-5.96)	-0.21*** (-4.81)
Private-to-PE	-0.15*** (-3.48)	-0.16*** (-3.59)	-0.16*** (-3.64)	-0.10*** (-2.21)
PE-to-PE	-0.12** (-2.48)	-0.13*** (-2.65)	-0.13*** (-2.69)	-0.07 (-1.47)
postIPO	-0.13*** (-3.55)	-0.14*** (-3.77)	-0.14*** (-3.82)	-0.06* (-1.83)
postMA	-0.10*** (-3.45)	-0.11*** (-3.55)	-0.11*** (-3.57)	-0.05* (-1.84)
Management		0.26*** (16.21)	0.21*** (12.96)	0.21*** (13.47)
Mid-Management		0.13*** (12.81)	0.11*** (10.87)	0.12*** (11.34)
White Collar		0.09*** (11.62)	0.09*** (12.12)	0.08*** (11.28)
Purple Collar		0.03** (2.56)	0.04*** (3.82)	0.03** (2.38)
Pink Collar		-0.01 (-0.77)	0.04*** (5.20)	0.03*** (4.35)
Blue Collar		-0.05*** (-4.88)	0.00 (0.39)	-0.01 (-0.57)
log(wage)			0.12*** (12.28)	0.14*** (14.41)
Tenure 0 < . < 1				0.25*** (13.67)
Tenure 1 < . < 3				0.08*** (8.28)
Tenure 3 < . < 5				0.01 (0.55)
Tenure 5 < . < 8				0.01 (1.01)
Tenure > 10				0.10*** (11.85)
Quarter Fixed-Effects	Yes	Yes	Yes	Yes
Company Fixed-Effects	Yes	Yes	Yes	Yes
<i>N</i>	715,375	715,375	715,375	715,375
<i>Adj - R</i> ²	0.12	0.12	0.12	0.13

in ratings. Ratings decrease by 0.10 and 0.12 for IPOs and M&As, respectively. Interestingly, this decrease is similar to that observed for SBOs, in which companies also experience a change of owner but not of ownership type. The finding that there is a decline in score around any change of ownership is consistent with Dahl (2011), who shows that there is an increase in uptake of stress-related medication for employees at organizations that change, especially those that undergo broad simultaneous changes along several dimensions, which is the case for IPOs, M&As and SBOs.

Since we already have a time and company fixed effect, the only control variables we can add are characteristics of reviewers: tenure, job position, and wage. All of the post-transaction effects decrease when we control for the highest score given by freshly arrived employees and those who have been in the company for more than ten years. Other control variables do not affect results.

With all the control variables, coefficients for SBOs, IPOs and M&As are not statistically significant. When companies experience an LBO for the first time, however, the decline is strong, especially if they were publicly listed before.

We confirm the decrease in ratings as we go down the hierarchy of jobs, which we saw in descriptive statistics. Management is most positive, followed by mid-management, then White Collars, Purple Collars, Pink Collars, and Blue Collars as the most dissatisfied. Wage plays an important role. The higher the salary, the higher the reported score. The relation is very strong. Controlling for salary naturally affect the score given by the different types of workers, but it does not affect the ranking. Hence, people that are higher up the hierarchy report a higher score than others, not just because they are better paid. They enjoy the hierarchical position per se.

C. Differences across PE sponsors

Our working sample is well distributed across PE sponsors (a.k.a. GPs); 49 GPs have four LBOs or more. We run the same regression as specification 4 in Table V, but instead of having cross-effects between industry and post-LBO, we have cross effects between each GPs and post-LBO. As the output from the regression would be too long to display in the format of Table V, we summarize the key results in Table VI.

There are nine GPs with an economically significant positive effect, i.e., ratings increase on average for their portfolio companies. Coefficients are between 0.20 and 0.56; six of these are statistically significant at a 10% level test. These GPs

Table VI: GP Fixed-Effects

We run the same regression as specification (4) in Table V, but instead of having cross-effects between industry and post-LBO, we have cross effects between each of the 47 GPs and post-LBO. This table summarizes the output from the regression analysis. Panel A shows the GPs with economically significant positive coefficients, which we set as being equal to 0.20 or above. Panel B shows the GPs with statistically significant coefficients at a 10% level test. Panel C shows the GPs with economically significant negative coefficients, which we set as being equal to -0.30 or below.

GP Name	Coef	T-stat	N Reviews	N Deals	% deals in Retail	% deals in Software	% deals in Other Indus	HQ
Summit Partners	-0.59	-2.33	85	4	0%	25%	75%	Boston
ABRY Partners	-0.55	-2.10	163	8	0%	13%	88%	Boston
TA Associates	-0.52	-2.97	166	4	0%	75%	25%	Boston
Providence Equity Partners	-0.47	-1.83	95	4	0%	25%	75%	Providence
Vista Equity Partners	-0.42	-27.45	1,932	29	0%	79%	21%	San Francisco
Marlin Equity Partners	-0.42	-2.04	177	4	0%	100%	0%	Hermosa Beach
Carlyle Group	-0.39	-1.25	893	14	0%	21%	79%	Washington
Platinum Equity	-0.33	-1.15	239	9	11%	11%	78%	Los Angeles
Centerbridge Partners	-0.33	-1.26	272	4	0%	25%	75%	New York
Thoma Bravo	-0.32	-1.33	2,038	28	0%	82%	18%	Chicago
GI Partners	-0.32	-1.55	137	6	0%	33%	67%	San Francisco
The Vistria Group	-0.29	-1.67	286	5	0%	0%	100%	Chicago
Silver Lake	-0.27	-2.25	317	5	0%	60%	40%	Menlo Park
Sycamore Partners	-0.26	-1.46	1,944	5	60%	20%	20%	New York
Vector Capital	-0.24	-0.80	341	5	0%	80%	20%	San Francisco
Francisco Partners	-0.24	-1.08	267	7	0%	71%	29%	San Francisco
Insight Partners	-0.23	-0.62	461	7	0%	86%	14%	New York
Madison Dearborn Partners	-0.21	-0.87	268	7	0%	14%	86%	Chicago
AEA Investors	-0.18	-0.69	416	7	0%	0%	100%	New York
KKR	-0.15	-0.55	1,246	14	7%	36%	57%	New York
Blackstone Group	-0.13	-0.40	253	4	0%	25%	75%	New York
Sentinel Capital Partners	-0.11	-0.43	210	6	17%	0%	83%	New York
Riverside Company	-0.09	-0.53	149	7	0%	14%	86%	New York
TPG	-0.07	-0.23	1,024	9	0%	33%	67%	Fort Worth
Cerberus Capital Management	-0.06	-0.25	118	5	0%	0%	100%	New York
Leonard Green & Partners	-0.04	-0.79	733	10	10%	0%	90%	Los Angeles
Siris Capital	-0.04	-0.18	493	7	0%	57%	43%	New York
Audax Group	-0.04	-0.14	67	5	0%	20%	80%	Boston
Bain Capital Ventures	-0.02	-0.07	769	11	18%	55%	27%	Boston
Genstar Capital Partners	0.00	-0.03	440	8	0%	38%	63%	San Francisco
Charlesbank Capital Partners	0.00	0.02	114	5	0%	0%	100%	Boston
Clayton Dubilier & Rice	0.02	0.08	323	7	14%	0%	86%	New York
H.I.G. Capital	0.02	0.14	866	13	8%	15%	77%	Miami
Hellman & Friedman	0.02	0.34	190	4	25%	50%	25%	San Francisco
AE Industrial Partners	0.05	0.33	241	5	0%	0%	100%	Boca Raton
Kohlberg & Company	0.06	0.28	151	4	0%	0%	100%	Mount Kisco
Apax Partners	0.07	0.27	823	7	29%	14%	57%	London
Stone Point Capital	0.08	0.39	527	6	0%	17%	83%	Greenwich
Investcorp	0.11	0.45	580	5	20%	0%	80%	Manama
Court Square Capital Partners	0.15	0.69	200	7	0%	14%	86%	New York
GTCR	0.15	0.95	221	8	0%	50%	50%	Chicago
Clearlake Capital Group	0.21	0.85	286	7	14%	71%	14%	Santa Monica
Golden Gate Capital	0.21	0.85	445	5	0%	20%	80%	San Francisco
Advent International	0.22	1.36	388	7	0%	14%	86%	Boston
American Securities	0.27	1.38	362	7	0%	14%	86%	New York
Warburg Pincus	0.45	2.77	155	5	20%	20%	60%	New York
Beecken Petty O'Keefe & Company	0.49	2.99	92	4	0%	0%	100%	Chicago

are listed in Panel A. They are all Generalists, i.e., invest across sectors, including Tech LBOs. Most of them discuss their ESG commitment on their website. For example, Golden Gate emphasizes on its website that they pay particular attention to all stakeholders and point out employees. Warburg Pincus has a long discussion

on ESG on its website. Their main headquarters are in different cities (New York, Boston, San Francisco).

Panel B lists thirteen GPs with a coefficient of -0.20 or lower. Only four GPs have a statistically significant fixed effect. The statistical significance, however, is sensitive to control variables, probably due to the large set of independent variables in this regression.

Nearly all of the GPs listed in panel B are specialized, and their specialty is in a single sector: Software.¹⁸ A partial exception is Summit, which does invest mostly in Tech LBOs but also sponsors LBOs in the Healthcare sector. Similarly, TA associates is a generalist but invests in Tech. Most of these GPs are based in the Silicon Valley and do not mention any ESG related content on their website.

D. Company characteristics

We have seen that Primary Buy-Outs are associated with a significant decline in ratings, especially when companies were publicly listed before; and that there is a strong GP fixed effect. We now study the effect of company and reviewer characteristics. Results are reported in Table VII.

Specification 1 shows that it is in the smaller companies that the effect is higher. The decrease in score is as much as 0.21 for small companies. The decrease is half as large for medium and large companies. This is important because as shown in the previous section our sample under-represents small companies.¹⁹

Mature companies also experience a smaller decrease in ratings than young companies (specification 2). We could have expected that it would be in mature companies that an LBO transition would represent the largest shock but it is not so.

The impact of LBO on ratings is different across job categories (specification 3). There are no changes in ratings for employees in managerial positions, and a slight decrease for employees in lower management positions. For all the other employee categories – i.e., all non managers – for find a similar coefficient and therefore pool them together in the Table for ease of readability. Employees in non managerial positions report much lower ratings.

¹⁸According to Pitchbook, the most active GPs in Tech LBOs are Vista, followed by ABRY, Providence and Thoma Bravo. <https://pitchbook.com/news/articles/in-a-high-tech-world-private-equity-wont-be-left-behind>

¹⁹While this issue is common in the literature, if the negative effect is concentrated in large companies, the observed decrease in ratings is upward biased. In addition, as our unit of observation is a review, we are over-weighting large companies compared to an analysis aggregated at the company level.

Table VII: Cross Effects. This table shows regression outputs generated as in Table V. The set of control variables is the same as in Table V specification (4). We show results when we add as an explanatory variable cross effects with the dummy variable 'Post LBO'. Panel A shows the results with reviewer and company characteristics: job category, size, foundation year, and industry. Panels B, C, and D of this table are presented in appendix L. Panel B is the same as Panel A, but with Post M&A instead of Post LBO in the cross effects. Panels C and D show results with specification (4) in panel A, to which we add relevant GP fixed effects (which are crossed with the post-LBO dummy variable).

Panel A: post LBO and Company Characteristics				
	Dependent variable: Overall Score			
	(1)	(2)	(3)	(4)
postLBO & Small	-0.20*** (-4.07)			
postLBO & Medium	-0.09** (-2.18)			
postLBO & Big	-0.11** (-2.15)			
postLBO & Mature		-0.09*** (-2.80)		
postLBO & Millenials		-0.28*** (-4.90)		
postLBO & Management			0.00 (0.05)	
postLBO & Mid Management			-0.08* (-1.69)	
postLBO & Not Management			-0.14*** (-5.84)	
postLBO & Other Services				-0.02 (-0.34)
postLBO & Corporate Service				-0.05 (-1.02)
postLBO & Industrial				-0.09 (-1.15)
postLBO & Retail				-0.22*** (-4.30)
postLBO & Software				-0.24*** (-4.52)
Control Variables	Yes	Yes	Yes	Yes
Quarter Fixed-Effects	Yes	Yes	Yes	Yes
Company Fixed-Effects	Yes	Yes	Yes	Yes
<i>N</i>	715,375	715,375	715,375	715,375
<i>Adjusted R</i> ²	0.13	0.13	0.13	0.13

The results on job positions are consistent with those in the literature. For example, Lichtenberg and Siegel (1990) finds that the decrease in number of jobs is concentrated in blue collars. Antoni et al. (2019) find that pink collars jobs are those reduced most. Olsson and Tåg (2017) who show that routine tasks tend to be

automated or offshored after an LBO.

For industries, we find that the Retail and Software sectors are the only ones with a significant decrease in ratings. The three services (IT, Consumer, and Public) have the same coefficient and we thus pool them together in specification 4 (OtherServices). The coefficient is negative but not significant, similar to Industrial. Notice that Business services has a positive, but not significant, coefficient.

In panel B, we show the same specifications but with M&A transactions instead of LBO transactions. We do not observe any of the above effects except for company size. For M&A transactions as well, there is a larger decrease in ratings in smaller companies. We do not find a difference across industries. Even job categories are unrelated to decrease in ratings post M&A. It is in fact management that has the largest decrease in ratings.

In non tabulated results, we estimate specification 4 in Table VII - Panel A, adding one of the GPs listed in table VI - Panel B at a time. We observe that the coefficient on Software hardly changes except when Vista is included. There are five other GPs that affect slightly the coefficient on software (Thoma Bravo, Marlin, TA Associates, GI Partners and Silver Lake); the other GP fixed effects do not.

In Panel C, we show specification 4 of Table VII panel A, when we add only Vista alone (specification 2), then the other five GPs just mentioned (specification 3) plus the GPs with significant (negative) fixed effect. Vista explains away one quarter of the decrease in ratings observed in software. The other GPs explain away another quarter, which indicates that the above results are primarily a GP effect and not a sector effect. When we add Thoma Bravo as well, software loses significance. Adding the other GPs only slightly lower the coefficient on the software sector.

To sum up, the decrease in ratings occurs with primary buy-outs, especially if companies were publicly traded before. There is no change in ratings for Management and a strong one for employees in non management positions. There is no significant industry fixed effects, except for Retail, once we control for GP fixed effects. There is wide heterogeneity across GPs. GPs associated with declines in ratings seem to have as a common point that they are located in the silicon Valley, specialize in Tech LBOs, and do not discuss ESG related matters on their website.

E. Robustness Checks

We make a number of changes to the empirical approach to gauge the robustness of our key results. Results are presented in Table VIII. Each line coincides with one change compared to our default specification, which is specification (4) in Table V,

and this base specification is reproduced in this table, as the first one, so we can benchmark the rest of the tests.

Glassdoor is sometimes perceived as a website where people rant about their employer. We have discussed and provided several pieces of evidence against this view. For example, we highlighted in section II.A that several studies showed that employee ratings are closely related to several traditional measures of company performance.

Yet, something else we can do is to remove reviews that are likely to be rants. Reviews with exclamation marks and with upper-cased words, can be seen as emotional, and long reviews (top quartile in number of words) may also be rants. We take out these 164,571 reviews, and note that results are similar except for management, whose score also goes down after the LBO. Another way to remove extreme views is to remove extreme scores. When the 248,917 scores of either 1 or 5 are taken out, the coefficient naturally changes but the t-statistics remain similar.

As the unit of observation in our panel analysis is a review, firms with very few reviews do not influence results much. We nonetheless increase the required number of reviews pre- and post-transaction to 5 instead of 3. Results are unchanged.

Transactions occurring early in the sample may be affected by the lack of coverage of Glassdoor. For example, LBOs occurring in 2013 do not have many pre-LBO reviews. It is therefore important to verify results hold if we remove early transactions. Results become slightly stronger when we remove the 30,816 reviews submitted between January 2013 and December 2015. We also verify that our results are similar if we take out recent reviews, i.e., those submitted between January 2017 and December 2019.




Figure V.2. Average Score and number of reviews per month.

Glassdoor uses advanced tools to detect fake reviews. However, a remaining concern is that companies may game the yearly Employees’ Choice Awards, also

known as the Best Places to Work Awards. This award is based on Glassdoor reviews as of the end of October. Companies that are among the best-scored ones may 'stuff' reviews before the deadline to earn the reward. Figure V.2 does show a spike in both the number of reviews and the average score in October.

Given our findings, it is unlikely that well-scored companies fighting for the top spot will affect our conclusions. In addition, if gaming these awards should improve the scores of companies that are up for sale then PE-owners would be particularly keen to push the scores up, whereas we find that scores are lower under PE-ownership. Yet, we can test for it by removing reviews submitted in October. Specification "No October Rev" of table VIII shows that doing so does not affect our results.

Table VIII: Robustness. This table shows the same specification as table V for the robustness different sub-samples.

	Score						
	(Baseline)	(No Emotional)	(No long review)	(No extreme Scores)	(min 5 Reviews)	(2016-2019 deals)	(2013-2016 deals)
Private2PE	-0.09** (-2.00)	-0.07 (-1.55)	-0.06 (-1.04)	-0.01 (-0.24)	-0.07 (-1.56)	-0.13** (-2.36)	-0.07 (-1.05)
Public2PE	-0.21*** (-4.74)	-0.19*** (-4.07)	-0.17*** (-4.08)	-0.08*** (-4.09)	-0.20*** (-4.76)	-0.17*** (-3.39)	-0.28*** (-4.14)
PE2PE	-0.07 (-1.33)	-0.05 (-1.12)	-0.02 (-0.36)	-0.04 (-1.41)	-0.05 (-1.02)	-0.10 (-1.61)	-0.01 (-0.19)
postIPO	-0.05 (-1.55)	-0.05 (-1.37)	0.00 (0.04)	-0.02 (-0.81)	-0.05 (-1.38)	-0.04 (-1.01)	-0.07 (-1.30)
postMA	-0.05* (-1.65)	-0.03 (-1.26)	-0.03 (-1.06)	0.01 (0.96)	-0.06** (-2.03)	-0.09** (-2.17)	0.01 (0.37)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	664,042	603,078	499,471	415,125	675,612	633,223	613,921
<i>R</i> ²	0.13	0.13	0.13	0.06	0.13	0.13	0.13

	Score				
	(No October Rev)	(No short Rev)	(min 1.000 Employees)	(No fake mention I)	(No fakse mention II)
Private2PE	-0.09** (-2.20)	-0.08* (-1.76)	-0.04 (-0.81)	-0.08* (-1.78)	0.02 (0.28)
Public2PE	-0.21*** (-5.16)	-0.23*** (-4.30)	-0.21*** (-4.70)	-0.20*** (-4.15)	-0.20* (-1.79)
PE2PE	-0.06 (-1.25)	-0.07 (-1.25)	-0.06 (-1.07)	-0.07 (-1.45)	-0.03 (-0.37)
postIPO	-0.06 (-1.60)	-0.06 (-1.58)	-0.06* (-1.83)	-0.04 (-0.92)	-0.01 (-0.12)
postMA	-0.05* (-1.70)	-0.07* (-1.92)	-0.04 (-1.33)	-0.04 (-1.43)	-0.04 (-0.89)
Controls	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes	Yes
<i>N</i>	601,765	500,952	632,691	615,306	162,350
<i>R</i> ²	0.13	0.15	0.12	0.13	0.17

Another test, albeit less direct, consists in removing very short reviews. If the goal of a review is to push up the score of a company, we expect the review to

contain the minimal amount of words necessary to submit a rating. Specification “No short Rev” shows that results are unaffected when the shortest quartile of reviews is excluded. Furthermore, it should be easier to manipulate the score of small companies. Specification “min 1,000 Employees” shows that results are also unaffected if we remove 31,351 reviews of small firms with less 1,000 employees.

Another potential concern is fake reviews. For this, we count the number of reviews of LBO targets that refer to firms engaging in the “fake-review” practice. We identify fake-review mentions with two dictionary-based approach, one restrictive and the other more relaxed²⁰. We only identify 26 reviews (39 with the more relaxed approach) which mention fake reviews post-LBO in our data-set. This concerns 44 (315 with the relaxed approach) firms post-LBO out of 586.

As an additional robustness consideration, we report the evolution of the number of reviews around LBO and other M&A transactions in appendix J. We find that the number of reviews post-LBO keep slightly increasing, with similar patterns for both current and former employees. This is reassuring, as it tends to hint that there is not a sudden shift in the composition of reviewers, which could have for example been triggered by important turnover levels post-LBO.

Finally, appendix K reports the main specification of table V along the different sub-rating categories, and we find very similar results as for the standard score. In panels B and C we further report our main cross-effect results for position and industry for the different sub-scores, and once again very similar results across all dimensions.

IV. Textual Analysis

In addition to the quantitative ratings, employees can write three reviews. Two reviews are mandatory: the cons of working for the company, which has 32 words on average, and the pros of working for their company, which has fewer words on average (24 words). A third review is optional – advice to management – and about half of the raters fill it in. In this section, we analyze the ‘cons’ reviews as they are the longest reviews, but we show results for the other two types of reviews in appendices. We have 611,000 cons reviews with sufficient information.²¹ As twenty

²⁰The list of words in the restrictive approach is made of: [“fake rev”, “HR push”, “human resources push”, “manipulating rev”, “manipulate rev”, “pushing to post”]. In the more relaxed approach, we also include reviews in which we find co-occurrences of “HR” or “human resources” with “review” inside the text.

²¹We use the same sample as in the previous section (e.g., people employed by the company they review at the time of the review, etc.). In addition, we require having at least five words in the

million words constitute too large a dataset to be analyzed in its raw form, we extract 25 common topics across these reviews using a standard technique called Latent Dirichlet Allocation. Each review is then given the proportion of each topic that it contains, and we then run regression analysis to establish which type of review is more or less likely to mention one of these topics.

A. *The Latent Dirichlet Allocation Approach*

The Latent Dirichlet Allocation (LDA) introduced by Blei et al. (2003) has been used extensively in computational linguistics. An influential early use is by Griffiths and Steyvers (2004) who study the content of 28,154 abstracts from the National Academy of Science. In Finance and Accounting, LDA has been applied to 10-K disclosures (Dyer et al., 2017), analyst discussions (Huang, 2018b), SEC comment letters (Dechow et al., 2015), firm disclosure in years surrounding fraud (Hoberg and Lewis, 2017), and to classify loans (Argyle et al., 2020). LDA is also used commercially. Newspapers such as the New York Times use LDA to recommend articles based on the topics of previously read articles (Spangher, 2015).

The LDA identifies statistical topics through groupings of terms, very much like in factor analysis. In our analysis, a term can be a unigram (single word), bi-gram (two successive words) or tri-gram. The LDA is particularly well suited to our setting because it allows for multiple topics to be present in a review and for any topic to occur in multiple reviews. In addition, as this computational linguistic technique is unsupervised, it is easily replicable and does not require assumptions about specific topics to be found in the document.

The LDA is a Bayesian technique; it assumes that there exists a posterior distribution based on hidden variables that generate the observed corpus of terms. The procedure seeks to infer i) the mixture distributions of terms $w = 1, \dots, N^w$ describing each topic $k = 1, \dots, N^k$, across the pooled set of reviews, and ii) the mixture distributions of topic $k = 1, \dots, N^k$ describing each review $r = 1, \dots, N^r$. Both distributions are Dirichlet; hence both have $[0,1]$ support, and

- i) $\sum_{w=1}^{N^w} \varphi_{k,w} = 1$, where $\varphi_{k,w}$ is the weight of each term w in topic k .
- ii) $\sum_{r=1}^{N^r} \theta_{r,k} = 1$, where $\theta_{r,k}$ is the weight of each topic k in review r .

review, after removing all non-informative words (e.g., 'the', 'a'), and keep only reviews written in English.

B. *Extracting topics*

The main parameter of choice is the number of topics across all reviews N^k .²² As detailed in appendix M, we maximize a coherence score to determine the optimal number of topics and the other parameters of choice needed for the LDA. This procedure indicates that we should work with 25 topics and the topics should be extracted separately for each type of reviews (cons, pros, advice). We detail our procedure to treat the textual data and estimate the LDA in appendix O.

As mentioned above, we focus on the ‘cons’ reviews. We gather a dozen of examples across the transaction types which display the highest weight on the topic. Based on these examples and the top 15 terms, we label each topic. Table IX reports successively the label, the top NGrams, one example of review drawn from the sample of post-LBO reviews and one example from the sample of M&A.

The list of topics appears to be natural, spanning the set of issues we would expect employees to complain about. Note that complaints about compensation are in different topics because some reviewers talk about the salary package offered by the company at entry such as benefits in kind or the package with regard to the industry, whereas others talk about the incentives and bonus or the pay raise in relations with their performance. We therefore label the topics separately, with on the one hand, benefits in kind and competitors salaries, and on the other hand, lack of incentives and low pay raise.

In the second column, we show the cumulative weight of these fifteen terms out of all the terms in the universe of reviews. A low cumulative weight indicates that the topic may lack dominant terms and is therefore more difficult to label.

²²There are some extra, but weak, additional assumptions in an LDA approach (see Dyer et al. (2017)). There are some other parameters to choose for this analysis.

Table IX- Topic description. Panel A: Corporate Strategy topics

Label	W_top15	Top Words	LBO example
Layoffs, Cost-cutting	0.18	compani, cut, employe, layoff, peopl, job, cost, busi, leav, term, profit, recent, moral, futur, due	1. Cost cutting due to Thoma Bravo acquisition 2. Skilled employees being laid off 3. Losing benefits (such as stock purchase program, stock bonus awa...
Changes in Leadership	0.23	manag, chang, leadership, constant, compani, cultur, senior, upper manag, direct, execut, decis, senior manag, constant chang, environ	Changes in senior leadership have started to change the culture of the company. Morale has taken a nosedive with the change in CEO and large-scale layo...
Fast changing and Growing company	0.26	chang, compani, grow, fast, lot, pace, challeng, growth, environ, sometime, time, pain, fast pace, move, quick	Sometimes it is hard to keep up with all the changes when a company is growing at such a fast pace.
Management style	0.24	team, manag, peopl, project, depend, compani, experi, manag team, lead, lot, cultur, don, skill, offic, leader	Executive team is a joke. We had an awesome, collaborative leadership team prior to LD stepping in. New team has no clue how to lead, motivate or man...
No complaint	0.35	con, compani, none, hard, don, time, job, negat, love, bad, peopl, experi, downsid, honest, isn	It just gets better with time...they are growing and changing and I can honestly say I don't see any cons with this job.
Bad HR Management	0.28	employe, manag, level, compani, hr, depart, lack, valu, leadership, senior, level manag, cultur, issu, feel, respect	- feedback from employees/lower management seems to be ignored by upper management - remote employees are frequently excluded from events/perks that p...
Competitor Salaries	0.36	salari, compani, pay, industri, competit, compans, low, compar, market, averag, lower, benefit, standard, red, tape	Compensation is not terrific, especially when compared with other companies in the industry. Career options are limited if you are working in one of ...
Work-life balance	0.45	life, balanc, life balanc, hour, time, lot, travel, hard, none, expect, sometime, stress, famili, workload, depend	Can be hard at certain times of the year on family life but you have to strike while the iron is hot in the busy seasons, there will be times to scale...
Career opportunities	0.39	opportun, career, growth, advanc, limit, move, compani, develop, path, littl, slow, career growth, opportun advanc, growth opportun, difficult	Poor career development. Lacking in education for future employment opportunities outside of academia.
Business development	0.15	product, compani, busi, develop, engin, softwar, support, market, system, resourc, technolog, issu, time, focus, process	This company has been unable to deliver on the vision for product management for at least five years. The reason for this is because it's not been run...
Internal Politics & Lack of diversity	0.23	polit, promot, manag, peopl, cultur, lot, senior, top, compani, level, leadership, base, offic, divers, boy	Not much growth due to rigid management structure Lots and lots of politics which tend to affect employees
Lack communication, Slow processes	0.24	process, lack, communic, slow, technolog, decis, depart, sometime, compani, lot, improv, organ, time, system, busi	Red tape and process can be cumbersome at times. Being an agency, it's frustrating when you put work into a solution for clients and they don't move f...
Benefits in kind	0.27	benefit, offic, insur, health, locat, expens, compani, employe, park, health insur, home, plan, offer, pay, cost	Healthcare benefits are expensive and not great. No education allowance, but they do provide some in house training options. Office space could use so...

Table IX- Topic description. Panel B: Working Conditions topics

Label	W top15	Top Words	LBO example
Lack of care from management	0.30	care, employe, compani, manag, don, doesn, peopl, money, job, care employe, don care, patient, line, custom, manag care	Company doesn't care about employees Bonuses are never what you deserve No set structure
Unrealistic sales goals	0.26	sale, custom, goal, sell, commiss, manag, servic, meet, product, call, pressur, expect, month, base, rep	Aggressive sales goals, focused more on profit than associates/customers
Poor upper management and communication	0.41	manag, poor, lack, communic, upper, train, upper manag, poor manag, support, micro, micro manag, staff, life, bad, balanc	Poor communication and lack of unified identity from upper management
Incentives	0.23	expect, employe, bonus, perform, rate, pay, turnov, manag, increas, compens, job, train, staff, base, incent	Weak bonus structure Average pay Short sighted objectives Frequent turnover
Badly treated by management	0.27	manag, employe, peopl, favorit, don, bad, worker, treat, store, horribl, terribl, supervisor, upper, upper manag, co	Bad managers make the job suck especially if you are a floor sales employee
Promotion & Hiring Process	0.26	promot, posit, hire, peopl, job, manag, move, compani, hard, rais, experi, advanc, pay, outsid, level	Worst place to look for advancement. They hire associates from outside the company, pay them more and then under pay their internal staff if they're...
Bad working schedule	0.27	hour, store, custom, time, shift, break, schedul, manag, week, sometim, lot, day, associ, rude, cut	Management shows a lot of favoritism, corporate is majorly cutting hours, crazy hours and managers get mad when you can't stay late or come in on a dal...
Day-to-day job	0.21	peopl, don, job, time, day, train, month, look, start, line, expect, tri	Consider this. Half of the people employed here will be gone in about two years. All of the work that they have has to be done by those who remain and...
Low pay, no raise	0.38	pay, low, rais, hour, low pay, wage, minimum, increas, littl, benefit, start, pay low, salari, minimum wage, rate	Low pay, fair distribution of work is questionable, merit increases do not match inflation to help with cost of living in the US, HR is not great a...
Stressful & Challenging	0.25	time, job, sometim, stress, train, difficult, lot, custom, client, busi, learn, servic, task, hard, littl	Very diverse portfolio of products, and customers sometimes expect you to know it all. This can be very stressful especially if you lack the proper t...
Other issues	0.09	call, compani, pay, month, center, time, money, phone, day, review, lr, call center, system, driver, employe	It can get loud as it is a Call Center, however there are dividers between the desks and the head phones provided cancel out surrounding sound.
Overtime	0.33	day, hour, time, week, schedul, holiday, shift, weekend, vacat, night, overtim, sick, flexibl, hour week, late	Under staffed, high volume of task, underpaid, Horrible hours, I can go to work at 4am, 12am, 12pm 6pm in the same work week! I have to use vacation t...
Benefits in kind	0.27	benefit, offic, insur, health, locat, expens, compani, employe, park, health insur, home, plan, offer, pay, cost	Healthcare benefits are expensive and not great. No education allowance, but they do provide some in house training options. Office space could use so...

We further categorize the topics into two groups: corporate strategy, and working conditions. As many aspects of working conditions are part of company strategy, the separation is not always clear-cut, though. For example, the topic of how compensation compares to competitors is assigned to corporate strategy, whereas complaints about low pay raise or incentives relating to performance are classified as working conditions. Another difficult topic to assign to one of these two groups is benefits. This topic is dominated by considerations like company health insurance policy, how well located the company is, the quality of offices, and expense policy. At the margin, we consider these to result from strategic corporate decisions. Similarly, work-life balance and career opportunities are grouped under corporate strategy. Topics we assign to working conditions are issues that we would not expect to be the direct result of the company strategy; e.g., overtime, lack of care from management, badly treated by management, bad working schedule, low pay raise, stressful and challenging environment, bad day-to-day working conditions, bad promotion and hiring processes, poor upper management and communication, unrealistic sales goals.

Table X - Panel A lists the twelve topics related to ‘corporate strategy’ (e.g., internal politics and lack of diversity, lack of resource for business development) in addition to a thirteenth topic, which is the ‘no complaint’ one. Panel B lists the twelve topics related to ‘working conditions.’

C. *Topic Analysis*

We estimate the same equation as equation (1) except for the dependent variable. In equation (4) below, the dependent variable is the weight of a given topic k in review r for company c written on day d (instead of the reviewer score for that company in equation (1)):

$$\begin{aligned} \theta_{r,c,d,k} = & \alpha_c + \alpha_{q(d)} + \beta * PostLBO_{c,d,c,d} \\ & + \theta_1 * PostM\&A_{c,d,c,d} + \theta_2 * PostIPO_{c,d,c,d} \\ & + \gamma * Z_r + \epsilon_{r,c,d} \end{aligned} \tag{V.4}$$

We estimate this equation separately for each of the 25 topics k , and show results in Table X. In Panel A, topics are those pertaining to corporate strategy, and in Panel B topics are related to working conditions. Topics are ordered from the one with the highest coefficient on post-LBO to the lowest in each Panel. Added to both Panels is a measure of the impact of each topic on the employee’s overall score, i.e., the

fixed effect of each topic on score.

Table X: Post-transaction Complaints

This table shows the results from estimating pooled panel regressions with the topic weight for the review as the dependent variable. Specification and controls are the same as in table V (4). Standard errors are clustered at the company and quarter level. For readability, we report coefficients multiplied by 100. T-stats appear in parentheses. The “Impact” coefficient from the first column, comes from a separate regression where the dependent variable is the “Overall Score”, but the only independent variable (except for the usual controls) is the topic weight.

Panel A: Corporate Strategy topics								
	Impact	postLBO	postIPO	postMA	Mngt	log(wage)	Controls	Adj-R2
Layoffs, cost-cutting	-0.47*** (-21.26)	1.01*** (2.80)	-0.46 (-1.18)	1.42*** (3.99)	-0.54*** (-4.36)	1.47*** (15.10)	YES	0.05
Changes in leadership	-0.72*** (-39.59)	0.36** (2.52)	0.20 (0.75)	-0.07 (-0.45)	-0.03 (-0.27)	1.20*** (16.00)	YES	0.03
Fast changing and growing company	0.94*** (45.75)	0.29 (1.15)	-1.40*** (-3.25)	-0.42 (-1.56)	0.97*** (4.36)	1.70*** (20.18)	YES	0.06
Management style	-0.32*** (0.65)	0.09 (0.95)	0.25 (-0.69)	-0.13 (-0.59)	-0.73*** (-5.93)	1.34*** (14.39)	YES	0.03
No complaint	1.12*** (56.59)	0.04 (0.16)	0.31 (0.81)	-0.48 (-1.61)	0.25 (1.28)	0.48*** (3.80)	YES	0.05
Bad HR management	-0.92*** (-38.22)	-0.02 (-0.15)	0.33* (1.72)	-0.29** (-2.27)	-0.90*** (-8.65)	0.27*** (5.98)	YES	0.04
Competitor salaries	0.29*** (19.88)	-0.10 (-0.54)	-0.16 (-0.63)	-0.24 (-1.16)	-0.66*** (-5.28)	1.09*** (12.75)	YES	0.04
Work-life balance	0.51*** (34.62)	-0.13 (-0.71)	-0.06 (-0.27)	0.21 (1.05)	0.15 (1.04)	0.38*** (3.65)	YES	0.05
Career opportunities	0.21*** (13.57)	-0.15 (-0.92)	-0.01 (-0.05)	0.44** (2.18)	-0.84*** (-6.32)	0.34*** (3.59)	YES	0.03
Business development	0.06*** (4.95)	-0.18 (-1.24)	-0.08 (-0.29)	-0.09 (-0.57)	0.24** (2.52)	1.64*** (29.05)	YES	0.04
Internal Politics & lack of diversity	-0.57*** (-25.23)	-0.26** (-2.36)	0.13 (0.65)	-0.38*** (-3.29)	0.01 (0.08)	1.26*** (16.71)	YES	0.03
Lack communication, slow processes	0.37*** (30.32)	-0.57*** (-2.91)	0.51* (1.85)	-0.13 (-0.68)	0.59*** (3.74)	2.48*** (23.71)	YES	0.04
Benefits in kind	0.59*** (41.41)	-0.62*** (-2.70)	-0.68** (-2.33)	-0.01 (-0.03)	-1.06*** (-7.98)	1.04*** (10.70)	YES	0.04

Table X: Continued.

Panel B: Working Conditions topics								
	Impact	postLBO	postIPO	postMA	Mngt	log(wage)	Controls	Adj-R2
Lack of care from management	-1.15*** (-35.27)	0.23* (1.68)	0.18 (1.09)	0.14 (0.82)	-0.25*** (-2.89)	-0.72*** (-10.06)	YES	0.03
Unrealistic sales goals	-0.29*** (-12.46)	0.20 (1.22)	0.17 (1.10)	0.08 (0.43)	1.30*** (10.94)	-1.80*** (-9.63)	YES	0.02
Poor upper management communication	-0.87*** (-40.47)	0.13 (0.64)	-0.00 (-0.02)	-0.13 (-0.57)	-0.52*** (-6.15)	-0.51*** (-5.22)	YES	0.02
Incentives	-0.51*** (-29.40)	0.12 (0.71)	-0.26 (-1.51)	0.01 (0.08)	-0.07 (-0.73)	-0.32*** (-5.53)	YES	0.02
Promotions & hiring Process	-0.27*** (-13.28)	0.11 (0.77)	0.37* (1.83)	-0.19 (-1.30)	-0.39*** (-4.32)	-0.95*** (-13.49)	YES	0.02
Badly treated by management	-1.23*** (-35.14)	0.09 (0.55)	0.22 (1.30)	0.15 (1.19)	-0.13 (-1.31)	-1.85*** (-21.71)	YES	0.04
Bad working schedule	0.20*** (16.78)	0.06 (0.26)	0.10 (0.42)	0.21 (0.91)	1.26*** (11.10)	-2.88*** (-16.92)	YES	0.02
Day-to-day job	-0.62*** (-25.87)	0.05 (0.27)	-0.16 (-0.81)	0.14 (0.77)	0.09 (1.14)	-0.88*** (-14.73)	YES	0.03
Low pay, no raise	-0.27*** (-16.61)	-0.08 (-0.40)	0.14 (0.71)	-0.01 (-0.09)	0.16 (1.46)	-2.35*** (-24.58)	YES	0.05
Other issues	0.06*** (4.95)	-0.11 (-0.81)	-0.05 (-0.23)	0.05 (0.33)	-0.14 (-1.60)	-0.54*** (-6.23)	YES	0.04
Stressful & challenging	0.67*** (30.68)	-0.15 (-0.70)	0.40 (1.56)	-0.16 (-1.36)	0.83*** (6.77)	-0.44*** (-5.41)	YES	0.02
Overtime	0.25*** (17.51)	-0.40** (-2.33)	0.02 (0.09)	-0.13 (-0.83)	0.40*** (3.87)	-1.47*** (-13.67)	YES	0.04

The most frequent complaints post LBO are related to layoffs and cost-cutting and changes in leadership, both related to corporate strategy. These two topics are also among the ones which cause significant drop in the overall score of employees, together with management style, bad HR management and internal politics and discrimination. Some corporate operating problems are less likely to occur post LBO such as bad processes, internal politics and discrimination and low benefits in kind.

Complaints about working conditions (Panel B) are mainly about lack of care from management (slightly significant), unrealistic sales goals, and poor upper management and communication, but the latter two are not statistically significant.

Topics related to the management-employee relationship drive the most important drop in employee score (i.e., lack of care from management, bad treatment from management, and poor upper management and communication). At the bottom of the Panel, are topics less likely to appear in a post-LBO review. Post LBO, employees are less likely to complain about i) stressful and challenging environment, ii) overtime or schedule issues.

If we look at post M&A, the main complaint is also about ‘layoffs and cost-cutting’, even more so than for post LBOs. As many M&As are motivated by synergies, we would indeed expect that they come with layoffs and cost-cutting, which is something not welcome by employees (O’Shaughnessy and Flanagan, 1998; Lee et al., 2018). Another complaint post M&A is a lack of career opportunities. We also observe less complaint about bad HR management. For post IPO, only a few topics are significant: there are fewer complaints about benefits and about company changes, and slightly more about bad HR management and bad processes.

Note that certain themes are less likely to appear in post-LBO reviews. Topics related to operational processes, such as ‘Internal Politics & lack of diversity’, ‘Lack of communication, slow processes’, or ‘Overtime’ seem to be less problematic. It is interesting to note that the last two in this list contrast with M&As, for which they are more likely to appear post-transaction.

Table XI: Post-LBO complaints characteristics

This table runs the same regression specification as in table X, with the difference that the post-LBO dummy is split in parts using interaction terms for employee position (panel A) and target firm industry (panel B). All controls and the sample are the same as before. This table is analogous to table VII, except that we use topic weights as independent variables. Coefficients were multiplied by 100 for readability. We report coefficients for the interaction terms and significance levels correspond to *** 10%, ** 5% and *1%.

Panel A: Employee position							
	postLBO & Mngt	postLBO & MidMngt	postLBO & NotMngt	Controls	Adj-R2		
<i>Corporate Strategy Topics</i>							
Layoffs, cost-cutting	1.402	1.186**	0.924**	YES	0.04		
Changes in leadership	-0.282	0.898***	0.305**	YES	0.03		
Fast changing and growing company	-0.915	1.477***	0.167	YES	0.03		
No complaints	3.211**	0.715	-0.183	YES	0.03		
<i>Working Conditions Topics</i>							
Lack of care from management	0.336	0.003	0.275*	YES	0.04		
Unrealistic sales goals	0.795	0.330	0.131	YES	0.02		
Poor upper management and communication	0.368	-0.106	0.207	YES	0.04		
Panel B: Target Firm Industry							
	postLBO & OtherInd	postLBO & Corporate Services	postLBO & Industrial	postLBO & Retail	postLBO & Software	Controls	Adj-R2
<i>Corporate Strategy Topics</i>							
Layoffs, cost-cutting	-0.303	0.368	-0.065	0.702	3.288***	YES	0.03
Changes in leadership	-0.037	0.096	0.114	0.771**	0.868**	YES	0.02
Fast changing and growing company	0.442	0.614	1.109	0.497	-0.29	YES	0.04
No complaints	0.587	0.800	-0.158	-0.616*	-0.223	YES	0.03
<i>Working Conditions Topics</i>							
Lack of care from management	0.037	-0.085	0.785	0.638**	0.186	YES	0.04
Unrealistic sales goals	0.133	-0.046	-0.173	0.881*	0.149	YES	0.03
Poor upper management and communication	-0.128	0.216	0.478	0.925***	-0.153	YES	0.04

Results on our proxy for employee wages are particularly interesting. This variable is positively related to all the corporate strategy topics (Panel A) and negatively related to all the working condition topics (Panel B). Employees who are paid more are more likely to complain about how the company is run: bad processes, fast changing and growing company, lack of resources for business development, internal politics and discrimination, management style, lay-offs and cost-cutting. Employees who are paid less are more likely to complain about their working conditions: working schedule, low pay raise, bad treatment from management, unrealistic sales goals, overtime etc. Employees classified as management complain less about management related variable (such as HR, compensation) and complain more about the challenge due to fast changing and growing environment, the lack of resource for business development, the bad internal processes. Regarding the working conditions, they are more likely to complain about unrealistic sales goals, overtime and scheduling problems, as well as the stressful environment. The clarity of these patterns is certainly reassuring and shows that these reviews do convey meaningful and logical information.

Overall, the topics mentioned post-LBO fail to show a similarly clear difference. Topics that are statistically significant cover both company operations and working conditions. Building on the results on employee wages, we refine the analysis by investigating differences post LBO according to the employee job position. Results in the previous section also highlighted peculiarities of LBOs conducted in the retail and software industries.

Table XI shows results when the post-LBO effect is split between the employee job position (management, mid-management and not-management) in Panel A, and between the industries in Panel B. We show clear differences between employees according to their job position. Management employees are more likely to express no complaint post LBO, while mid-management mostly complains about changes in corporate operating strategy, especially layoffs and cost-cutting, changes in leadership and fast change and growth. These also correspond to three most important complaints regarding corporate operating strategy post LBO. Employees which are not in a management position complain about both the changes induced in the corporate operating strategy but also the lack of care from management.

Panel B also highlights clear differences among industries. Problems induced by the corporate operating strategy are more likely to occur in the software industry (layoffs and cost-cutting and changes in leadership), while problems in the retail industry pertain mostly to the working conditions already developed above, i.e.,

Table XII: Post-LBO type complaints

This table reports the same regression specification as in table X. We split the post-LBO dummy in the three types of transactions, i.e., Public-to-PE, Private-to-PE and PE-to-PE. Specification and controls are again the same as in table V (4). Standard errors are clustered at the company and quarter level. For readability, we report coefficients multiplied by 100. T-stats appear in parentheses.

Panel A: Corporate Strategy topics					
	Public-to-PE	Private-to-PE	PE-to-PE	Controls	Adj-R2
Layoffs, cost-cutting	2.81*** (3.60)	0.22 (0.78)	0.00 (0.01)	YES	0.03
Changes in leadership	0.24 (0.85)	0.71** (2.22)	0.19 (0.72)	YES	0.02
Fast changing and growing company	-0.01 (-0.01)	0.28 (0.54)	0.57* (1.71)	YES	0.02
Managemnt style	-0.17 (-0.71)	-0.20 (-0.68)	0.56*** (2.62)	YES	0.02
No complaint	-0.45 (-0.95)	0.58 (1.21)	0.04 (0.08)	YES	0.02
Bad HR management	0.10 (0.44)	0.15 (0.56)	-0.28 (-1.01)	YES	0.04
Competitor salaries	-0.22 (-0.61)	0.36 (1.10)	-0.37* (-1.92)	YES	0.02
Work-life balance	0.14 (0.47)	-0.53** (-2.07)	-0.04 (-0.12)	YES	0.03
Career opportunities	-0.31 (-0.89)	-0.51** (-2.30)	0.30 (0.95)	YES	0.05
Business development	-0.28 (-1.63)	0.34 (0.96)	-0.52* (-1.82)	YES	0.04
Internal Politics & lack of diversity	-0.42 (-1.62)	-0.34 (-1.59)	-0.05 (-0.21)	YES	0.02
Lack communication, slow processes	-0.66* (-1.90)	-0.57* (-1.81)	-0.49* (-1.75)	YES	0.04
Benefits in kind	-1.27*** (-3.38)	-0.63** (-2.12)	0.00 (0.00)	YES	0.04

Table XII: Continued.

Panel B: Working Conditions topics					
	Public-to-PE	Private-to-PE	PE-to-PE	Controls	Adj-R2
Lack of care from management	0.75*** (3.11)	0.10 (0.44)	-0.14 (-0.63)	YES	0.03
Unrealistic sales goals	0.50 (1.43)	-0.22 (-0.88)	0.27 (1.54)	YES	0.02
Poor upper management communication	0.09 (0.26)	0.62*** (2.83)	-0.25 (-0.98)	YES	0.02
Incentives	-0.01 (-0.03)	0.23 (0.76)	0.14 (0.54)	YES	0.02
Promotions & hiring Process	-0.10 (-0.45)	-0.06 (-0.26)	0.44 (1.41)	YES	0.02
Badly treated by management	0.30 (0.87)	0.24 (0.78)	-0.22 (-1.18)	YES	0.04
Bad working schedule	-0.14 (-0.29)	-0.07 (-0.28)	0.35 (1.25)	YES	0.02
Day-to-day job	0.26 (0.81)	-0.06 (-0.19)	-0.04 (-0.17)	YES	0.03
Low pay, no raise	-0.34 (-1.16)	-0.17 (-0.65)	0.23 (0.63)	YES	0.05
Other issues	0.23** (2.21)	-0.11 (-0.38)	-0.43 (-1.34)	YES	0.04
Stressful & challenging	-0.80*** (-2.72)	0.12 (0.33)	0.22 (0.66)	YES	0.02
Overtime	-0.26 (-0.97)	-0.46 (-1.42)	-0.48* (-1.66)	YES	0.04

problems related to management-employee relationship (such as lack of care and poor upper management and communication) and to unrealistic sales goals.

Results in the previous section confirm those established in the literature, pointing out at important differences between types of LBOs: Public-to-PE, Private-to-PE and PE-to-PE.

Table XII - Panels A and B shows results when the post-LBO effect is split between the three types of LBOs exactly as done in Table IX. We observe that there are more significant topics, and they are LBO-type dependent. Complaints about layoffs and cost-cutting are more prevalent only in Public-to-PE, consistent with results in Lerner et al. (2011). Besides, there are more complaints about management lack of care. In these transactions, there are fewer complaints about stressful working conditions, bad processes and, benefits. One interpretation is that Public-to-PE transaction may improve how the company is run, but it comes with a strong negative shock: management ruthlessly laying off people and cutting costs.

In private-to-PE transactions, two complaints are more prominent: change in leadership and poor upper management and communication. Post private-to-PE, there are fewer complaints about bad processes, benefits, career opportunities and work-life balance. The coefficient on 'no complaint' is large, though not statistically significant. Finally, in PE-to-PE, complaints center around management style and company changes (due to growth). There are less complaints about the lack of resources for business development, processes, and salaries and working conditions.

The results with transaction type are not as clear-cut as those on job positions or industries, which indicates that some important heterogeneity of practices remain. That said, Public-to-PE transaction do seem to have a distinctively different effect on employees.

V. Discussion and Conclusion

Private equity funds often get a bad press in the media and from politicians alike. A recurrent reproach made at those funds is that they seek financial gain at the expense of employee welfare. In this paper, we document how changes in company ownership impact employee's perception of their workplace. Using a set of ratings and written reviews, spanning thousands of firms of different types, we find that employees report lower scores about their employers after changes of ownership.

A key finding of our study is the significant cross-sectional variation in employee ratings that we document following a change in company ownership. Employee ratings decline more after LBOs than in the aftermath of other types of organizational changes. The impact of LBO transactions is also dependent on the nature of the company's previous ownership type: Welfare concerns appear to be more pronounced in previously publicly held firms. In firms previously detained by another PE sponsor, employees report smaller decreases in ratings. We also find effects re-

lated to industries (retail and software industries are the most negatively affected) and individual PE sponsors.

More importantly, LBOs impact employees differently along the hierarchical structure: we find no decrease in ratings among managers, whereas non-managers are the most affected, thus widening a pre-existing welfare gap between the two groups.

We gain additional insights into the origins of those differences by analyzing the contents of the complaints written by the employees. Using a standard topic modelling tool from the field of Natural Language Processing, we map the written reviews along a parsimonious set of 25 topics. Following LBOs, non-managers seem to worry more about their position at the firm, as themes related to *“Layoffs & cost-cutting”* become more prevalent. They also complain more about management, as *“Changes in leadership”* and *“Lack of care from management”* are topics that appear more often post-LBO. Managers, on the other hand, are more likely to report *“No complaints”*.

Nonetheless, we find that certain complaints are less frequent post-LBO, suggesting that some processes improve, which allows us to draw a more nuanced picture. For example, problems related to *“Internal Politics & lack of diversity”*, *“Lack of communication & slow processes”*, and *“Overtime”* tend to decrease when a PE sponsor steps in.

We can sum up the study’s main contributions as follows. First, employees dislike organizational changes, as their self-reported welfare decreases. This is especially true for LBOs, in particular if the firm was previously publicly listed. Second, our textual analysis provides unique insights into the unconditional core issues workers face that most affect their welfare: Tensions with managers is what most affect their wellbeing, followed by problems linked to uncertainty caused by company changes. Third, LBOs are linked to cross-sectional differences in welfare changes. Non-managers report higher fears related to layoffs and more issues with management, coupled with a sharper decrease in overall welfare compared to managers.

Our findings have implications for decision-makers and managers of private equity firms buying out other companies. There are many reasons why managers would want to optimize the welfare of their employees. Positive reviews are associated with more profitable and wealth-generating companies (Huang, 2018a; Green et al., 2019; Chemmanur et al., 2019). Furthermore, with the rise of ESG and capital allocation increasingly favoring funds that behave positively in the social responsibility space, managers might start to increasingly consider employee welfare an important metric

of success.

Our study hints to where the main problems reside following LBOs and provide insights for managers to address potential problems proactively. Our results can also increase awareness about potential problems and thus increase pressure on managers to consider those as important. Finally, our findings also suggest that PE firms can put in place optimized procedures improving employee welfare, which could also be put forward by private equity firms to improve their image.

Finally, those results also have implications for regulators who desire to protect employees working for PE owned firms. Organizational changes cause dissatisfaction, but LBO changes even more so. Our study helps to identify the specific problems that employees perceive under PE management, but absent in other M&As: those are issues linked to changes in leadership and lack of care from management. While those are important issues, which heavily impact employee welfare, designing regulations aimed at alleviating those specific problems might prove challenging.

On the other hand, despite being a major concern, we find no greater fears about “*Layoffs & cost-cutting*”²³ in LBO reviews compared to M&A’s. Furthermore, echoing previous findings in the literature that point to the productivity benefits brought by PE firms, LBOs seem to improve certain operational practices. Issues related to overtime, internal politics, as well as slow processes, are all less likely to occur post-LBO; but are more present in complaints post-M&A. In their approach, regulators should keep in mind that it also benefits workers to preserve those positive aspects brought in by PE firms.

²³For example, in the “Wall Street Looting Act”, concerns about bankruptcies and ensuing layoffs are a central theme.

Chapter VI

Conclusion

Our informational landscape is changing. The amount of information available to economic agents is greater than ever before, and continues to increase at an exponential rate. This faces researchers in economics and finance to new challenges and opportunities, as increasingly powerful statistical and econometric tools allow them to make sense and leverage novel “Big Data” sets.

Consider the seminal presidential address of Roll (1988), dating over three decades back. Back then, he noted that even in retrospect, finance researchers are only able to explain a small fraction of the variation in stock returns. Worse, arrival of public information seems to have no effect on the proportion of unexpected news driving returns on a given day.

Overturing this finding has proven to be very difficult. To some extent, we might, however, be at the beginning of a shift of paradigm. Indeed, as we show in the literature review of chapter I, we are entering an area in which we are now able to quantify a tremendous amount of information, which previously was either way to complex to handle or simply non-existent. As we document, the literature is only starting to exploit those new possibilities.

Concerning the particular problem posed by Roll, Boudoukh et al. (2019) show that identifying the information which is truly relevant is a crucial first step into finding new insights contained in financial media. While Roll finds virtually no difference in his ability to explain the variation in stock returns between news days and non-news days, Boudoukh et al. (2019) find that on “relevant news” days, the R^2 of regressions explaining returns sharply decreases: a sign that unexpected returns play a much more significant role on those days.

This finding serves as an invitation for research to work on two fronts. First, it opens up the challenge of not only identifying the relevant news, but also to quantify its content. This should be done in order to, secondly, understand how the media,

or any other form of information, drives the unexpected components of returns, i.e., the part of the variation in stock returns that we cannot forecast. How important are public news, and how much do they really impact stock returns? How efficient are markets in assimilating this sort of information?

In this thesis, we seek to contribute towards those challenges, both by using proxies for the content of relevant news, and by linking them to updates in investor expectations, i.e., the unexpected component of returns. To proxy for media content, we rely on a measure of firm-specific news sentiment provided by Thomson Reuters. To proxy for unexpected returns, we rely on empirical econometric estimation methods widely used in the literature based on the return decomposition of Campbell and Shiller (1988a).

Recent findings in the literature have shown that news have a causal impact on anomaly returns (Da et al., 2014; Hillert et al., 2014; Engelberg et al., 2018). Armed with news sentiment and with unexpected cash flow and discount rate news, we seek to shed new light on the drivers of anomaly returns. In particular, we can go back to the four questions posed in chapter I, and discuss the implications of our findings.

1. Do value stocks earn a premium due to the way they react to earnings announcement news?

Engelberg et al. (2018) document that value firms earn a premium on EAD. We show that this premium is, in fact, a bad-news premium, as it occurs on EAD when sentiment in the media is negative. This happens because growth stocks are significantly more sensitive to the bad sentiment on EAD, whereas value firms are more resilient to negative news.

We fail to trace those findings back to systematic risk, in particular to dynamic exposures to a risk factor. We discuss potential alternative frameworks. In particular, theories arguing in favor of biased investor expectations or that rely on differences in investor attention (which match well with the low media attention of value firms that we document), do not contradict our findings. Nonetheless, we are not able to establish a definitive causal link and suggest avenues for future research.

In the broader context of this thesis, we document that beyond a causal link between news and market outcomes, it is possible to bring new insights to questions in asset pricing by also proxying for the content of the news themselves. In particular, this paper shows that they can present two useful applications: first, they can show characteristics of entity-date observations, and thus guide researchers in their theories to explain the observed phenomena. Second, they can serve to find

cross-sectional differences in sensitivity to the sentiment of news. Again, differing cross-sectional sensitivities can shed new light on previously poorly understood market phenomena, and in the context of asset pricing, naturally fit into the literature on market anomalies.

2. What do commonalities in anomaly portfolio returns tell us about theories seeking to explain market anomalies?

This paper aims at guiding the field in distinguishing valid theories for anomalies, by characterizing the empirical comovement patterns across anomaly characteristics. We find that anomaly returns are mainly driven by commonalities in CF news, but that differences exist across anomalies: the strongest CF comovement appears among value, small, and loser stocks in particular.

A distinctive feature of our paper is that we assess comovement conditionally on different states of the economy. We find that CF news of value and loser stocks comove more strongly during aggregate negative CF shock events, such as recessions. Conversely, CF news of growth and winner stocks exhibit stronger comovement in periods of DR stress, such as high levels of investor sentiment. Those results aim to guide theories for anomalies, and suggests the existence of different drivers of returns for the opposite legs of anomaly portfolios.

In the broader context of the dissertation, this paper focuses on a measure of individual perception which we cannot observe directly: revisions in investor expectations, i.e., CF and DR news. This paper shows that understanding the drivers of these news can help our understanding of pervasive market phenomena and guide researchers in their proposal for fitting theories.

3. What drives a firm to become a value and momentum stock?

The literature documents that the value and momentum anomalies share several opposing patterns. In this paper, we find that this opposition between value and momentum is pervasive in their CF and DR news as well. In particular, momentum stocks are driven by positive CF and negative DR shocks. In contrast, the past price decline of value firms is driven by positive DR shocks and negative CF news. More generally, we find that firms with high past CF shocks experience return continuation. Conversely, firms with high past DR shocks tend to have higher expected returns.

We link those patterns to news media sentiment. Good (bad) media news cor-

relate with positive (negative) CF shocks. By contrast, good (bad) news co-occur with negative (positive) revisions in discount rates. The consistent negative news to which value firms, and the positive news to which winner firms, thus correlate with their documented CF and DR patterns, but are insufficient to capture the full magnitude of the forecasting effects.

In the broader context of the thesis, this paper again highlights the link between news and revisions in investor expectations. In particular, it shows that sentiment captured from media tone tightly relate to patterns and CF and DR revisions, and may thus invite future research to improve estimates of updates in investor expectations using sentiment from news.

4. Do private equity firms operate at the expense of target firm employees?

PE sponsors often get a bad press in the media. This paper aims to helping understand if this bad reputation is deserved, by measuring how employee welfare is effectively affected when a PE firms takes over the company they work for. We find that employees dislike ownership changes, and even more so when a PE firm takes over a publicly listed company. More importantly, we document that not employees are affected equally by PE takeovers. Managers appear to be roughly unaffected, whereas non-managers report a significant drop in satisfaction, thus exacerbating a pre-existing welfare gap between the two groups.

To understand the cause of this divergence, we investigate the content of the complaints, by mapping the written reviews along a set of 25 topics using a Latent Dirichlet Allocation algorithm. We find that non-managers complain more about issues related to lack of care from management and cost-cutting, which are among the problems that cause the most severe issues in employee welfare unconditionally. Nonetheless, we also find that certain problems become less recurrent post-buyout, most notably complaints related to operational processes.

In the broader context of the thesis, this paper shows that research in finance can benefit from measures of sentiment outside the spectrum of news and financial media. Other economic agents have determinant roles in economic outcomes, and we are only scratching the surface of what social media can reveal about the opinion of key stakeholders.

Going forward

This thesis shows that in the context of economics and finance, not only can the new sources of information found in “Big Data” serve to solve questions related to financial markets, but it can also serve to guide decision-makers at all echelons of the economic cycle.

For managers, for example, this means that they can get access to new tools to assess the opinion of their customers about their products to match their demands. They can learn from their employees how to streamline and improve processes inside the company. They can get new tools to measure how they impact different stakeholders, knowing that the market will judge their performance based on a broad range of metrics that will ultimately impact their balance sheet.

For policymakers, sentiment extracted from the media can also be a valuable tool to guide them in their decisions. Chapter V is one such example on how social media can help us understand contentious situations better and to grasp all the nuance involved in a complicated issue.

In asset pricing, exponentially multiplying the amounts of information might lead to a change of paradigm, where researchers become able to explain a greater portion of unexpected returns. However, dealing with dramatically large amounts of data also brings forth potential concerns, such as data mining and model overfitting (Harvey et al., 2016). Progress to improve our understanding of financial markets will not only be driven by the usage of sentiment metrics affecting investor expectations across a broad range of media: sharp economic intuition and sound theoretical models will remain a cornerstone of academic research for many years to come.

This leads us to reflect back on the opening quote: “[...] *the relation between narratives and economic outcomes is likely to be complex and time varying. The impact of narratives on the economy is regularly mentioned in journalistic circles, but without the demands of academic rigor. [...] But, the advent of big data and of better algorithms of semantic search might bring more credibility to the field.*” (Shiller, 2017, p. 48).

We show that we can slowly approach and understand the complex and time-varying link between narratives and economic outcomes. However, we must acknowledge that we are still at the beginning of a big endeavor. Potential limitations must still be taken into account. Our measures might suffer from biases, for example related to the news coverage received by firms or by the employees who choose to report about their employers. This type of concern might be alleviated with time and the arrival of new data where selection bias will be diminished.

Furthermore, those narratives are inherently complex. We proxy those quantities with simple, aggregate measures. Yet this simplicity potentially comes at the cost of missing important nuance. The way in which our quantities are measured might also lead us to miss out on relevant information. We also discuss in the chapters describing our data, a set of potential improvements that could be considered. Despite existing limitations, as described in the introduction, novel methods of natural language processing can help us in the future to better understand the link between narratives coming in the form of soft information, and relevant economic questions. This thesis is about taking one step in that direction.

Finally, all those questions in finance and economics that can leverage “sentometrics” seem to rely on a common underlying denominator: subjective perceptions of economic agents. Here lies an important part of the complexities of the problem. It is also the starting point of new research possibilities offered by methods that can exploit and make sense of this type of data. Since our field of research falls under the umbrella of social sciences, it is unsurprising that a broad range of economic outcomes are shaped by subjective perceptions of individual and groups of agents. The good news is that sentiment and textual processing possibilities are progressively giving us a way to assess those innately complex quantities in a more and more reliable manner.

Bibliography

- Aboody, David, Reuven Lehavy, and Brett Trueman, 2010, Limited attention and the earnings announcement returns of past stock market winners, *Review of Accounting Studies* 15, 317–344.
- Agrawal, Ashwini, and Prasanna Tambe, 2016, Private Equity and Workers' Career Paths: The Role of Technological Change, *Review of Financial Studies* 29, 2455–2489.
- Agrawal, Shreyash, Pablo D. Azar, Andrew W. Lo, and Taranjit Singh, 2018, Momentum, Mean-Reversion, and Social Media: Evidence from StockTwits and Twitter, *The Journal of Portfolio Management* 44, 85–95, Publisher: Institutional Investor Journals Umbrella.
- Ahmad, Khurshid, JingGuang Han, Elaine Hutson, Colm Kearney, and Sha Liu, 2016, Media-expressed negative tone and firm-level stock returns, *Journal of Corporate Finance* 37, 152–172.
- Algaba, Andres, David Ardia, Keven Bluteau, Samuel Borms, and Kris Boudt, 2020, Econometrics Meets Sentiment: An Overview of Methodology and Applications, *Journal of Economic Surveys* 34, 512–547, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/joes.12370>.
- Allen, David E., Michael McAleer, and Abhay K. Singh, 2019, Daily market news sentiment and stock prices, *Applied Economics* 51, 3212–3235, Publisher: Routledge eprint: <https://doi.org/10.1080/00036846.2018.1564115>.
- Allen, Franklin, Xian Gu, and Julapa Jagtiani, 2020, A Survey of Fintech Research and Policy Discussion, SSRN Scholarly Paper ID 3622468, Social Science Research Network, Rochester, NY.
- Amess, Kevin, and Mike Wright, 2007, The Wage and Employment Effects of Leveraged Buyouts in the UK, *International Journal of the Economics of Business* 14, 179–195.
- Amess, Kevin, and Mike Wright, 2012, Leveraged buyouts, private equity and jobs, *Small Business Economics* 38, 419–430.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2009, High idiosyncratic volatility and low returns: International and further U.S. evidence, *Journal of Financial Economics* 91, 1–23.

- Antoni, Manfred, Ernst Maug, and Stefan Obernberger, 2019, Private equity and human capital risk, *Journal of Financial Economics* 133, 634–657.
- Antweiler, Werner, and Murray Z. Frank, 2004, Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards, *The Journal of Finance* 59, 1259–1294, [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2004.00662.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2004.00662.x).
- Appelbaum, Eileen, 2019, Hearing Entitled “America for Sale? An Examination of the Practices of Private Funds.”, *Working Paper, Center for Economic and Policy Research*.
- Arcot, Sridhar, Zsuzsanna Fluck, José Miguel Gaspar, and Ulrich Hege, 2015, Fund managers under pressure: Rationale and determinants of secondary buyouts, *Journal of Financial Economics* 115, 102–135.
- Argyle, Bronson, Taylor D Nadauld, and Christopher Palmer, 2020, Real effects of search frictions in consumer credit markets, Technical report, National Bureau of Economic Research.
- Arslan-Ayaydin, Özgür, James Thewissen, and Wouter Torsin, 2021, Disclosure tone management and labor unions, *Journal of Business Finance & Accounting* 48, 102–147.
- Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and Momentum Everywhere, *The Journal of Finance* 68, 929–985.
- Avramov, Doron, Si Cheng, and Allaudeen Hameed, 2016, Time-Varying Liquidity and Momentum Profits, *Journal of Financial and Quantitative Analysis* 51, 1897–1923, Publisher: Cambridge University Press.
- Axelson, Ulf, Tim Jenkinson, Per Strömberg, and Michael S. Weisbach, 2013, Borrow Cheap, Buy High? The Determinants of Leverage and Pricing in Buyouts, *The Journal of Finance* 68, 2223–2267, [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12082](https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12082).
- Azar, Pablo D., and Andrew W. Lo, 2016, The Wisdom of Twitter Crowds: Predicting Stock Market Reactions to FOMC Meetings via Twitter Feeds, *The Journal of Portfolio Management* 42, 123–134, Publisher: Institutional Investor Journals Umbrella.
- Babenko, Ilona, and Rik Sen, 2014, Money Left on the Table: An Analysis of Participation in Employee Stock Purchase Plans, *Review of Financial Studies* 27, 3658–3698.
- Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong, 2018, Social Connectedness: Measurement, Determinants, and Effects, *Journal of Economic Perspectives* 32, 259–280.

- Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor Sentiment and the Cross-Section of Stock Returns, *The Journal of Finance* 61, 1645–1680.
- Baker, Malcolm, and Jeffrey Wurgler, 2007, Investor Sentiment in the Stock Market, *Journal of Economic Perspectives* 21, 129–152.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis, 2016, Measuring Economic Policy Uncertainty, *The Quarterly Journal of Economics* 131, 1593–1636.
- Ball, Ray, and Philip Brown, 1968, An Empirical Evaluation of Accounting Income Numbers, *Journal of Accounting Research* 6, 159–178.
- Ball, Ray, and S. P. Kothari, 1991, Security Returns around Earnings Announcements, *The Accounting Review* 66, 718–738.
- Barbaglia, Luca, Sergio Consoli, and Sebastiano Manzan, 2021, Forecasting with Economic News, SSRN Scholarly Paper ID 3698121, Social Science Research Network, Rochester, NY.
- Barber, Brad M., Emmanuel T. De George, Reuven Lehavy, and Brett Trueman, 2013, The earnings announcement premium around the globe, *Journal of Financial Economics* 108, 118–138.
- Barber, Brad M., Adair Morse, and Ayako Yasuda, 2021, Impact investing, *Journal of Financial Economics* 139, 162–185.
- Barber, Brad M., and Terrance Odean, 2008, All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors, *The Review of Financial Studies* 21, 785–818.
- Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2015, X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1–24.
- Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2018, Extrapolation and bubbles, *Journal of Financial Economics* 129, 203–227.
- Barberis, Nicholas, and Andrei Shleifer, 2003, Style investing, *Journal of Financial Economics* 68, 161–199.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.
- Barberis, Nicholas, Andrei Shleifer, and Jeffrey Wurgler, 2005, Comovement, *Journal of Financial Economics* 75, 283–317.

- Bartov, Eli, Lucile Faurel, and Partha S. Mohanram, 2017, Can Twitter Help Predict Firm-Level Earnings and Stock Returns?, *The Accounting Review* 93, 25–57.
- Basu, Sanjoy, 1977, Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis, *The Journal of Finance* 32, 663–682.
- Bellon, Aymeric, 2020, Does private equity ownership make firms cleaner? the role of environmental liability risks, *The Role Of Environmental Liability Risks (May 18, 2020)* .
- Berk, Jonathan B, Richard C Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *The Journal of Finance* 54, 1553–1607.
- Bernard, Victor L., and Jacob K. Thomas, 1989, Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?, *Journal of Accounting Research* 27, 1–36.
- Bernstein, Shai, Josh Lerner, and Filippo Mezzanotti, 2019, Private Equity and Financial Fragility during the Crisis, *The Review of Financial Studies* 32, 1309–1373, Publisher: Oxford Academic.
- Bernstein, Shai, Josh Lerner, Morten Sorensen, and Per Strömberg, 2017, Private Equity and Industry Performance, *Management Science* 63, 1198–1213.
- Bernstein, Shai, and Albert Sheen, 2016a, The Operational Consequences of Private Equity Buyouts: Evidence from the Restaurant Industry, *The Review of Financial Studies* 29, 2387–2418, Publisher: Oxford Academic.
- Bernstein, Shai, and Albert Sheen, 2016b, The Operational Consequences of Private Equity Buyouts: Evidence from the Restaurant Industry, *Review of Financial Studies* 29, 2387–2418.
- Black, Fischer, 1986, Noise, *The Journal of Finance* 41, 528–543.
- Blei, David, Andrew Ng, and Michael Jordan, 2003, Latent Dirichlet Allocation, *Journal of Machine Learning Research* 3, 993–1022.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, 2015, Do Private Equity Owned Firms Have Better Management Practices?, *American Economic Review* 105, 442–446.
- Bloomfield, Robert J, William B Tayler, and Flora Zhou, 2009, Momentum, reversal, and uninformed traders in laboratory markets, *The Journal of Finance* 64, 2535–2558.
- Boucly, Quentin, David Sraer, and David Thesmar, 2019, Job Creating LBOs, *Journal of Financial Economics (forthcoming)* .

- Boudoukh, Jacob, Ronen Feldman, Shimon Kogan, and Matthew Richardson, 2019, Information, Trading, and Volatility: Evidence from Firm-Specific News, *The Review of Financial Studies* 32, 992–1033.
- Boudt, Kris, and James Thewissen, 2019, Jockeying for Position in CEO Letters: Impression Management and Sentiment Analytics, *Financial Management* 48, 77–115.
- Boudt, Kris, James Thewissen, and Wouter Torsin, 2018, When does the tone of earnings press releases matter?, *International Review of Financial Analysis* 57, 231–245.
- Boyer, Brian H., 2011, Style-Related Comovement: Fundamentals or Labels?, *The Journal of Finance* 66, 307–332.
- Bradshaw, Mark T., Scott A. Richardson, and Richard G. Sloan, 2006, The relation between corporate financing activities, analysts' forecasts and stock returns, *Journal of Accounting and Economics* 42, 53–85.
- Brav, Alon, and J. B. Heaton, 2002, Competing Theories of Financial Anomalies, *The Review of Financial Studies* 15, 575–606.
- Bybee, Leland, Bryan T. Kelly, Asaf Manela, and Dacheng Xiu, 2021, Business News and Business Cycles, Working Paper 29344, National Bureau of Economic Research, Series: Working Paper Series.
- Calomiris, Charles W., and Harry Mamaysky, 2019, How news and its context drive risk and returns around the world, *Journal of Financial Economics* 133, 299–336.
- Campbell, John, and Robert Shiller, 1988a, Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *The Review of Financial Studies* 1, 195–228.
- Campbell, John L., Hsinchun Chen, Dan S. Dhaliwal, Hsin-min Lu, and Logan B. Steele, 2014, The information content of mandatory risk factor disclosures in corporate filings, *Review of Accounting Studies* 19, 396–455.
- Campbell, John Y., 1991, A Variance Decomposition for Stock Returns, *The Economic Journal* 101, 157–179.
- Campbell, John Y., and John Ammer, 1993, What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns, *The Journal of Finance* 48, 3–37.
- Campbell, John Y., Christopher Polk, and Tuomo Vuolteenaho, 2010, Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns, *Review of Financial Studies* 23, 305–344.

- Campbell, John Y, and Robert J Shiller, 1988b, Stock prices, earnings, and expected dividends, *the Journal of Finance* 43, 661–676.
- Campbell, John Y, and Robert J Shiller, 2001, Valuation ratios and the long-run stock market outlook: An update.
- Campbell, John Y, and Tuomo Vuolteenhao, 2004, Bad Beta, Good Beta, *The American Economic Review* 94, 40.
- Carhart, Mark M., 1997, On Persistence in Mutual Fund Performance, *The Journal of Finance* 52, 57–82.
- Cassella, Stefano, and Huseyin Gulen, 2018, Extrapolation bias and the predictability of stock returns by price-scaled variables, *The Review of Financial Studies* 31, 4345–4397.
- Chambers, Anne E, and Stephen H Penman, 1984, Timeliness of reporting and the stock price reaction to earnings announcements, *Journal of accounting research* 21–47.
- Chapman, Kimball, 2018, Earnings notifications, investor attention, and the earnings announcement premium, *Journal of Accounting and Economics* 66, 222–243.
- Chari, Varadarajan V, Ravi Jagannathan, and Aharon R Ofer, 1988, Seasonalities in security returns: The case of earnings announcements, *Journal of Financial Economics* 21, 101–121.
- Chemmanur, Thomas J, Harshit Rajaiya, and Jinfei Sheng, 2019, How does online employee ratings affect external firm financing? evidence from glassdoor, *Evidence from Glassdoor (December 16, 2019)* .
- Chen, Long, Zhi Da, and Xinlei Zhao, 2013, What Drives Stock Price Movements?, *The Review of Financial Studies* 26, 841–876.
- Chen, Long, and Xinlei Zhao, 2009, Return Decomposition, *The Review of Financial Studies* 22, 5213–5249.
- Chomsky, Noam, 2009, Syntactic structures, in *Syntactic Structures* (De Gruyter Mouton).
- Chordia, Tarun, Amit Goyal, Gil Sadka, Ronnie Sadka, and Lakshmanan Shivakumar, 2009, Liquidity and the Post-Earnings-Announcement Drift, *Financial Analysts Journal* 65, 18–32.
- Clogg, Clifford C, Eva Petkova, and Adamantios Haritou, 1995, Statistical methods for comparing regression coefficients between models, *American journal of sociology* 100, 1261–1293.

- Cochrane, John H., 2011, Presidential address: Discount Rates, *The Journal of Finance* 1047–1108.
- Cohen, Daniel A., Aiysha Dey, Thomas Z. Lys, and Shyam V. Sunder, 2007, Earnings announcement premia and the limits to arbitrage, *Journal of Accounting and Economics* 43, 153–180.
- Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2003, The Value Spread, *The Journal of Finance* 58, 609–641.
- Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2009, The Price Is (Almost) Right, *The Journal of Finance* 64, 2739–2782.
- Cohn, Jonathan, Nicole Nestoriak, and Malcolm Wardlaw, 2021, Private equity buyouts and workplace safety, *The Review of Financial Studies* 34, 4832–4875.
- Cohn, Jonathan B., Nicole Nestoriak, and Malcolm Wardlaw, 2016, How Do Employees Fare in Leveraged Buyouts? Evidence from Workplace Safety Records .
- Cohn, Jonathan B., and M. I. Wardlaw, 2016, Financing Constraints and Workplace Safety, *The Journal of Finance* 71, 2017–2058.
- Conrad, Jennifer, and M Deniz Yavuz, 2017, Momentum and reversal: Does what goes up always come down?, *Review of Finance* 21, 555–581.
- Cookson, J Anthony, Joseph Engelberg, and William Mullins, 2021, Echo chambers, *Available at SSRN 3603107*, Working Paper.
- Cooper, Michael J., Huseyin Gulen, and Michael J. Schill, 2008, Asset Growth and the Cross-Section of Stock Returns, *The Journal of Finance* 63, 1609–1651.
- Da, Zhi, Umit G. Gurun, and Mitch Warachka, 2014, Frog in the Pan: Continuous Information and Momentum, *The Review of Financial Studies* 27, 2171–2218.
- Dahl, Michael S., 2011, Organizational Change and Employee Stress, *Management Science* 57, 240–256.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor Psychology and Security Market Under- and Overreactions, *The Journal of Finance* 53, 1839–1885.
- Daniel, Kent, and Tobias J. Moskowitz, 2016, Momentum crashes, *Journal of Financial Economics* 122, 221–247.
- Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam, 2001, Overconfidence, Arbitrage, and Equilibrium Asset Pricing, *The Journal of Finance* 56, 921–965.

- Das, Sanjiv R., 2019, The future of fintech, *Financial Management* 48, 981–1007, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/fima.12297>.
- Das, Sanjiv R., and Mike Y. Chen, 2007, Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web, *Management Science* 53, 1375–1388, Publisher: INFORMS.
- Davis, S, John Haltiwanger, Kyle Handley, Ben Lipsius, Josh Lerner, Javier Miranda, and J Zeitler, 2019a, The social impact of private equity over the economic cycle, Technical report, Working paper.
- Davis, Steven J., R. Jason Faberman, and John Haltiwanger, 2012, Labor market flows in the cross section and over time, *Journal of Monetary Economics* 59, 1–18.
- Davis, Steven J., John Haltiwanger, Kyle Handley, Ron Jarmin, Josh Lerner, and Javier Miranda, 2014, Private Equity, Jobs, and Productivity, *American Economic Review* 104, 3956–3990.
- Davis, Steven J, John C Haltiwanger, Kyle Handley, Ben Lipsius, Josh Lerner, and Javier Miranda, 2019b, The Economic Effects of Private Equity Buyouts, Working Paper 26371, National Bureau of Economic Research, Series: Working Paper Series.
- De Bondt, Werner F. M., and Richard H. Thaler, 1987, Further Evidence On Investor Overreaction and Stock Market Seasonality, *The Journal of Finance* 42, 557–581.
- De Bondt, Werner FM, and Richard Thaler, 1985, Does the stock market overreact?, *The Journal of finance* 40, 793–805.
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Noise Trader Risk in Financial Markets, *Journal of Political Economy* 98, 703–738.
- Dechow, Patricia M., Alastair Lawrence, and James P. Ryans, 2015, SEC Comment Letters and Insider Sales, *The Accounting Review* 91, 401–439.
- Degeorge, Francois, Jens Martin, and Ludovic Phalippou, 2016, On secondary buyouts, *Journal of Financial Economics* 120, 124–145.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 2018, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* .
- Dougal, Casey, Joseph Engelberg, Diego García, and Christopher A. Parsons, 2012, Journalists and the Stock Market, *The Review of Financial Studies* 25, 639–679.
- Dyer, Travis, Mark Lang, and Lorien Stice-Lawrence, 2017, The evolution of 10-K textual disclosure: Evidence from Latent Dirichlet Allocation, *Journal of Accounting and Economics* 64, 221–245.

- Eaton, Charlie, Sabrina T Howell, and Constantine Yannelis, 2020, When investor incentives and consumer interests diverge: Private equity in higher education, *The Review of Financial Studies* 33, 4024–4060.
- Edgerton, Jesse, 2012, Agency Problems in Public Firms: Evidence from Corporate Jets in Leveraged Buyouts, *The Journal of Finance* 67, 2187–2213.
- Edmans, Alex, 2011, Does the stock market fully value intangibles? Employee satisfaction and equity prices, *Journal of Financial Economics* 101, 621–640.
- Ehsani, Sina, and Juhani T. Linnainmaa, 2019, Factor Momentum and the Momentum Factor, Working Paper 25551, National Bureau of Economic Research, Series: Working Paper Series.
- Engelberg, Joseph, Matthew Henriksson, Asaf Manela, and Jared Williams, 2021, The Partisanship of Financial Regulators, SSRN Scholarly Paper ID 3481564, Social Science Research Network, Rochester, NY.
- Engelberg, Joseph, R. David Mclean, and Jeffrey Pontiff, 2018, Anomalies and News, *The Journal of Finance* 73, 1971–2001.
- Engelberg, Joseph E., and Christopher A. Parsons, 2011, The Causal Impact of Media in Financial Markets, *The Journal of Finance* 66, 67–97, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2010.01626.x>.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of financial economics* 33, 3–56.
- Fama, Eugene F., and Kenneth R. French, 1995, Size and Book-to-Market Factors in Earnings and Returns, *The Journal of Finance* 50, 131–155, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1995.tb05169.x>.
- Fama, Eugene F., and Kenneth R. French, 1996, Multifactor Explanations of Asset Pricing Anomalies, *The Journal of Finance* 51, 55–84.
- Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.
- Fang, Lily, Jim Goldman, and Alexandra Roulet, 2021, Private equity and pay gaps inside the firm, in *Available at NBER Conference f149062*.
- Fang, Lily, and Joel Peress, 2009, Media Coverage and the Cross-section of Stock Returns, *The Journal of Finance* 64, 2023–2052, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2009.01493.x>.

- Foucault, Thierry, Johan Hombert, and Ioanid Roşu, 2016, News Trading and Speed, *The Journal of Finance* 71, 335–382, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12302>.
- Fox, Isaac, and Alfred Marcus, 1992, The Causes and Consequences of Leveraged Management Buyouts, *The Academy of Management Review* 17, 62.
- Fracassi, Cesare, Alessandro Previtro, and Albert Sheen, 2019, Barbarians at the Store? Private Equity, Products, and Consumers, SSRN Scholarly Paper ID 2911387, Social Science Research Network, Rochester, NY.
- Francis, Jennifer, Ryan Lafond, Per Olsson, and Katherine Schipper, 2007, Information Uncertainty and Post-Earnings-Announcement-Drift, *Journal of Business Finance & Accounting* 34, 403–433.
- Garcia-Gomez, Pilar, Ernst G. Maug, and Stefan Obernberger, 2020, Private Equity Buyouts and Employee Health, SSRN Scholarly Paper ID 3601813, Social Science Research Network, Rochester, NY.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019a, Text as Data, *Journal of Economic Literature* 57, 535–574.
- Gentzkow, Matthew, Jesse M Shapiro, and Matt Taddy, 2019b, Measuring group differences in high-dimensional choices: method and application to congressional speech, *Econometrica* 87, 1307–1340.
- Gerakos, Joseph, and Juhani T. Linnainmaa, 2018, Decomposing Value, *The Review of Financial Studies* 31, 1825–1854.
- Glasserman, Paul, and Harry Mamaysky, 2019, Does Unusual News Forecast Market Stress?, *Journal of Financial and Quantitative Analysis* 54, 1937–1974, Publisher: Cambridge University Press.
- Goergen, Marc, Noel O’Sullivan, and Geoffrey Wood, 2014, The consequences of private equity acquisitions for employees: new evidence on the impact on wages, employment and productivity, *Human Resource Management Journal* 24, 145–158.
- Goldstein, Itay, Chester S. Spatt, and Mao Ye, 2021, Big Data in Finance, Working Paper 28615, National Bureau of Economic Research, Series: Working Paper Series.
- Gornall, Will, Oleg Gredil, Sabrina T Howell, and Xing Liu, 2021, Do employees cheer for private equity? the heterogeneous effects of buyouts on job quality, *The Heterogeneous Effects of Buyouts on Job Quality (August 23, 2021)* .

- Gotthelf, Nina, and Matthias W. Uhl, 2019, News Sentiment: A New Yield Curve Factor, *Journal of Behavioral Finance* 20, 31–41, Publisher: Routledge _eprint: <https://doi.org/10.1080/15427560.2018.1432620>.
- Green, T Clifton, Ruoyan Huang, Quan Wen, and Dexin Zhou, 2019, Crowdsourced employer reviews and stock returns, *Journal of Financial Economics* 134, 236–251.
- Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of Returns and Expected Returns, *The Review of Financial Studies* 27, 714–746.
- Griffith, John, Mohammad Najand, and Jiancheng Shen, 2020, Emotions in the Stock Market, *Journal of Behavioral Finance* 21, 42–56, Publisher: Routledge _eprint: <https://doi.org/10.1080/15427560.2019.1588275>.
- Griffiths, Thomas L., and Mark Steyvers, 2004, Finding scientific topics, *Proceedings of the National Academy of Sciences* 101, 5228–5235, Publisher: National Academy of Sciences Section: Colloquium.
- Guery, Lorin, Anne Stevenot, Geoffrey T. Wood, and Chris Brewster, 2017, The Impact of Private Equity on Employment: The Consequences of Fund Country of Origin-New Evidence from France, *Industrial Relations: A Journal of Economy and Society* 56, 723–750.
- Guo, Shourun, Edith S. Hotchkiss, and Weihong Song, 2011, Do Buyouts (Still) Create Value?, *The Journal of Finance* 66, 479–517.
- Gupta, Atul, Sabrina T Howell, Constantine Yannelis, and Abhinav Gupta, 2021, Does private equity investment in healthcare benefit patients? evidence from nursing homes, Technical report, National Bureau of Economic Research.
- Hales, Jeffrey, James R. Moon, and Laura A. Swenson, 2018, A new era of voluntary disclosure? Empirical evidence on how employee postings on social media relate to future corporate disclosures, *Accounting, Organizations and Society* 68-69, 88–108.
- Harter, James K, Frank L Schmidt, and Theodore L Hayes, 2002, Business-unit-level relationship between employee satisfaction, employee engagement, and business outcomes: a meta-analysis., *Journal of applied psychology* 87, 268.
- Harvey, Campbell, Heqing Zhu, and Yan Liu, 2016, ... and the Cross-Section of Expected Returns, *The Review of Financial Studies* 29, 5–58.
- Harvey, Campbell R., and Yan Liu, 2021, Lucky factors, *Journal of Financial Economics* 141, 413–435.

- Hendershott, Terrence, Dmitry Livdan, and Norman Schürhoff, 2015, Are institutions informed about news?, *Journal of Financial Economics* 117, 249–287.
- Heston, Steven L., and Nitish Ranjan Sinha, 2017, News vs. Sentiment: Predicting Stock Returns from News Stories, *Financial Analysts Journal* 73, 67–83, Publisher: Routledge .eprint: <https://doi.org/10.2469/faj.v73.n3.3>.
- Hillert, Alexander, Heiko Jacobs, and Sebastian Müller, 2014, Media Makes Momentum, *The Review of Financial Studies* 27, 3467–3501.
- Hirshleifer, David A., James N. Myers, Linda A. Myers, and Siew Hong Teoh, 2008, Do Individual Investors Cause Post-Earnings Announcement Drift? Direct Evidence from Personal Trades, *The Accounting Review* 83, 1521–1550.
- Hoberg, Gerard, and Craig Lewis, 2017, Do fraudulent firms produce abnormal disclosure?, *Journal of Corporate Finance* 43, 58–85.
- Hochreiter, Sepp, and Jürgen Schmidhuber, 1997, Long short-term memory, *Neural computation* 9, 1735–1780.
- Hong, Harrison, and Jeremy C. Stein, 1999, A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets, *The Journal of Finance* 54, 2143–2184.
- Hou, Kewei, Wei Xiong, and Lin Peng, 2009, A tale of two anomalies: The implications of investor attention for price and earnings momentum, *Available at SSRN 976394* .
- Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting Anomalies: An Investment Approach, *The Review of Financial Studies* 28, 650–705.
- Huang, Allen H, Amy Y Zang, and Rong Zheng, 2014a, Evidence on the information content of text in analyst reports, *The Accounting Review* 89, 2151–2180.
- Huang, Dashan, Fuwei Jiang, Jun Tu, and Guofu Zhou, 2015a, Investor Sentiment Aligned: A Powerful Predictor of Stock Returns, *The Review of Financial Studies* 28, 791–837.
- Huang, Jiekun, 2018a, The customer knows best: The investment value of consumer opinions, *Journal of Financial Economics* 128, 164–182.
- Huang, Jiekun, 2018b, The customer knows best: The investment value of consumer opinions, *Journal of Financial Economics* 128, 164–182.
- Huang, Kelly, Meng Li, and Stanimir Markov, 2020, What do employees know? evidence from a social media platform, *The Accounting Review* 95, 199–226.
- Huang, Minjie, Pingshu Li, Felix Meschke, and James P Guthrie, 2015b, Family firms, employee satisfaction, and corporate performance, *Journal of Corporate Finance* 34, 108–127.

- Huang, Xuan, Siew Hong Teoh, and Yinglei Zhang, 2014b, Tone management, *The Accounting Review* 89, 1083–1113.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, *The Journal of Finance* 48, 65–91.
- Jegadeesh, Narasimhan, and Di Wu, 2013, Word power: A new approach for content analysis, *Journal of Financial Economics* 110, 712–729.
- Jiang, Fuwei, Joshua Lee, Xiumin Martin, and Guofu Zhou, 2019, Manager sentiment and stock returns, *Journal of Financial Economics* 132, 126–149.
- Kaniel, Ron, Arzu Ozoguz, and Laura Starks, 2012, The high volume return premium: Cross-country evidence, *Journal of Financial Economics* 103, 255–279.
- Kaplan, Steven, 1989, The effects of management buyouts on operating performance and value, *Journal of Financial Economics* 24, 217–254.
- Kaplan, Steven N, and Per Strömberg, 2009, Leveraged Buyouts and Private Equity, *Journal of Economic Perspectives* 23, 121–146.
- Kearney, Colm, and Sha Liu, 2014, Textual sentiment in finance: A survey of methods and models, *International Review of Financial Analysis* 33, 171–185.
- Khimich, Natalya V., 2012, *Cash Flow and Discount Rate news estimation: which method to choose?*, Ph.D. thesis, UC Berkeley.
- Kogan, Leonid, and Dimitris Papanikolaou, 2013, Firm Characteristics and Stock Returns: The Role of Investment-Specific Shocks, *The Review of Financial Studies* 26, 2718–2759.
- Kogan, Shimon, Tobias J. Moskowitz, and Marina Niessner, 2021, Social Media and Financial News Manipulation, SSRN Scholarly Paper ID 3237763, Social Science Research Network, Rochester, NY.
- Kross, William, and Douglas A Schroeder, 1984, An empirical investigation of the effect of quarterly earnings announcement timing on stock returns, *Journal of accounting research* 153–176.
- Kuchler, Theresa, Yan Li, Lin Peng, Johannes Stroebel, and Dexin Zhou, 2020, Social proximity to capital: Implications for investors and firms, (*No. w27299*). *National Bureau of Economic Research* .
- La Porta, Rafael La, 1996, Expectations and the Cross-Section of Stock Returns, *The Journal of Finance* 51, 1715–1742.

- La Porta, Rafael La, Josef Lakonishok, Andrei Shleifer, and Robert Vishny, 1997, Good News for Value Stocks: Further Evidence on Market Efficiency, *The Journal of Finance* 52, 859–874.
- Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny, 1994, Contrarian Investment, Extrapolation, and Risk, *The Journal of Finance* 49, 1541–1578.
- Lamont, Owen, and Andrea Frazzini, 2007, The Earnings Announcement Premium and Trading Volume, Working Paper 13090, National Bureau of Economic Research.
- Larcker, David F, and Anastasia A Zakolyukina, 2012, Detecting deceptive discussions in conference calls, *Journal of Accounting Research* 50, 495–540.
- Larsen, Vegard H., Leif Anders Thorsrud, and Julia Zhulanova, 2021, News-driven inflation expectations and information rigidities, *Journal of Monetary Economics* 117, 507–520.
- Lawrence, Alastair, James Ryans, Estelle Sun, and Nikolay Laptev, 2018, Earnings announcement promotions: A Yahoo Finance field experiment, *Journal of Accounting and Economics* 66, 399–414.
- Lee, Charles MC, 1992, Earnings news and small traders: An intraday analysis, *Journal of Accounting and Economics* 15, 265–302.
- Lee, Kyeong Hun, David C. Mauer, and Emma Qianying Xu, 2018, Human capital relatedness and mergers and acquisitions, *Journal of Financial Economics* 129, 111–135.
- Lerner, Josh, Morten Sorensen, and Per Strömberg, 2011, Private Equity and Long-Run Investment: The Case of Innovation, *The Journal of Finance* 66, 445–477.
- Lerner, Josh, and Jean Tirole, 2003, Some Simple Economics of Open Source, *The Journal of Industrial Economics* 50, 197–234.
- Lettau, Martin, and Jessica A. Wachter, 2007, Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium, *The Journal of Finance* 62, 55–92.
- Lewis, Craig, and Steven Young, 2019, Fad or future? Automated analysis of financial text and its implications for corporate reporting, *Accounting and Business Research* 49, 587–615, Publisher: Routledge eprint: <https://doi.org/10.1080/00014788.2019.1611730>.
- Li, Feng, 2010, The information content of forward-looking statements in corporate filings—a naïve bayesian machine learning approach, *Journal of Accounting Research* 48, 1049–1102.
- Li, Feng, et al., 2010, Textual analysis of corporate disclosures: A survey of the literature, *Journal of accounting literature* 29, 143–165.

- Lichtenberg, Frank R., and Donald Siegel, 1990, The effects of leveraged buyouts on productivity and related aspects of firm behavior, *Journal of Financial Economics* 27, 165–194.
- Liu, Laura Xiaolei, and Lu Zhang, 2014, A neoclassical interpretation of momentum, *Journal of Monetary Economics* 67, 109–128.
- Lochstoer, Lars A, and Paul C Tetlock, 2020, What drives anomaly returns?, *The Journal of Finance* 75, 1417–1455.
- Lopez-Lira, Alejandro, 2020, Risk Factors That Matter: Textual Analysis of Risk Disclosures for the Cross-Section of Returns, SSRN Scholarly Paper ID 3313663, Social Science Research Network, Rochester, NY.
- Loughran, Tim, and Bill McDonald, 2011, When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks, *The Journal of Finance* 66, 35–65, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2010.01625.x>.
- Loughran, Tim, and Bill McDonald, 2016, Textual Analysis in Accounting and Finance: A Survey, *Journal of Accounting Research* 54, 1187–1230.
- Loughran, Tim, and Bill McDonald, 2020, Textual Analysis in Finance, *Annual Review of Financial Economics* 12, 357–375, eprint: <https://doi.org/10.1146/annurev-financial-012820-032249>.
- Lyle, Matthew R., and Charles C. Y. Wang, 2015, The cross section of expected holding period returns and their dynamics: A present value approach, *Journal of Financial Economics* 116, 505–525.
- Maio, Paulo, and Pedro Santa-Clara, 2015, Dividend Yields, Dividend Growth, and Return Predictability in the Cross Section of Stocks, *Journal of Financial and Quantitative Analysis* 50, 33–60, Publisher: Cambridge University Press.
- Manela, Asaf, and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.
- Mclean, R. David, and Jeffrey Pontiff, 2016, Does Academic Research Destroy Stock Return Predictability?, *The Journal of Finance* 71, 5–32.
- Mendenhall, Richard R., 2004, Arbitrage Risk and Post-Earnings-Announcement Drift, *The Journal of Business* 77, 875–894.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean, 2013, Distributed representations of words and phrases and their compositionality, in *Advances in neural information processing systems*, 3111–3119.

- Moskowitz, Tobias J., 2021, Asset Pricing and Sports Betting, *Available at SSRN 2635517*, *The Journal of Finance*, Forthcoming.
- Muscarella, Chris J., and Michael R. Vetsuypens, 1990, Efficiency and Organizational Structure: A Study of Reverse LBOs, *The Journal of Finance* 45, 1389.
- Ng, Jeffrey, Tjomme O Rusticus, and Rodrigo S Verdi, 2008, Implications of transaction costs for the post-earnings announcement drift, *Journal of Accounting Research* 46, 661–696.
- Odean, Terrance, 1998, Volume, volatility, price, and profit when all traders are above average, *The Journal of Finance* 53, 1887–1934.
- Olsson, Martin, and Joacim Tåg, 2017, Private Equity, Layoffs, and Job Polarization, *Journal of Labor Economics* 35, 697–754.
- O’Shaughnessy, KC, and David J Flanagan, 1998, Determinants of layoff announcements following m&as: An empirical investigation, *Strategic management journal* 19, 989–999.
- Patton, Andrew J., and Michela Verardo, 2012, Does Beta Move with News? Firm-Specific Information Flows and Learning about Profitability, *The Review of Financial Studies* 25, 2789–2839.
- Pazaj, Elisa, 2019, The Value-Momentum Correlation: An Investment Explanation, SSRN Scholarly Paper ID 3476844, Social Science Research Network, Rochester, NY.
- Pedersen, Lasse Heje, 2009, When everyone runs for the exit, (*No. w15297*). *National Bureau of Economic Research* .
- Pedersen, Lasse Heje, 2021, Game On: Social Networks and Markets, *Available at SSRN 2635517*, *The Journal of Finance*, Forthcoming.
- Penman, Stephen H, 1984, Abnormal returns to investment strategies based on the timing of earnings reports, *Journal of Accounting and Economics* 6, 165–183.
- Pennington, Jeffrey, Richard Socher, and Christopher D Manning, 2014, Glove: Global vectors for word representation, in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Peress, Joel, 2014, The Media and the Diffusion of Information in Financial Markets: Evidence from Newspaper Strikes, *The Journal of Finance* 69, 2007–2043, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12179>.
- Petersen, Mitchell A., 2009, Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches, *The Review of Financial Studies* 22, 435–480.

- Picault, Matthieu, and Thomas Renault, 2017, Words are not all created equal: A new measure of ECB communication, *Journal of International Money and Finance* 79, 136–156.
- Price, S McKay, James S Doran, David R Peterson, and Barbara A Bliss, 2012, Earnings conference calls and stock returns: The incremental informativeness of textual tone, *Journal of Banking & Finance* 36, 992–1011.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2003, Liquidity Risk and Expected Stock Returns, *Journal of Political Economy* 111, 642–685.
- Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al., 2019, Language models are unsupervised multitask learners, *OpenAI blog* 1, 9.
- Ranco, Gabriele, Darko Aleksovski, Guido Caldarelli, Miha Grčar, and Igor Mozetič, 2015, The Effects of Twitter Sentiment on Stock Price Returns, *PLOS ONE* 10, e0138441, Publisher: Public Library of Science.
- Rehurek, Radim, Petr Sojka, et al., 2011, Gensim—statistical semantics in python, *Retrieved from genism.org* .
- Renault, Thomas, 2017, Intraday online investor sentiment and return patterns in the U.S. stock market, *Journal of Banking & Finance* 84, 25–40.
- Roll, Richard, 1988, R2, *The Journal of Finance* 43, 541–566.
- Ross, Stephen, 1976, The arbitrage theory of capital asset pricing, *Journal of economic theory* 13, 341–360.
- Röder, Michael, Andreas Both, and Alexander Hinneburg, 2015, Exploring the Space of Topic Coherence Measures, in *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, WSDM '15, 399–408 (Association for Computing Machinery, New York, NY, USA).
- Sadka, Ronnie, 2006, Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk, *Journal of Financial Economics* 80, 309–349.
- Santos, Tano, and Pietro Veronesi, 2010, Habit formation, the cross section of stock returns and the cash-flow risk puzzle, *Journal of Financial Economics* 98, 385–413.
- Savor, Pavel, and Mungo Wilson, 2016, Earnings Announcements and Systematic Risk, *The Journal of Finance* 71, 83–138.
- Sheen, Albert, Youchang Wu, and Yuwen Yuan, 2021, Private equity and financial adviser misconduct, *Available at SSRN 3986875* .

- Shiller, Robert J., 2017, Narrative Economics, *American Economic Review* 107, 967–1004.
- Shive, Sophie A, and Margaret M Forster, 2020, Corporate governance and pollution externalities of public and private firms, *The Review of Financial Studies* 33, 1296–1330.
- Shleifer, Andrei, and Lawrence Summers, 1988, *Corporate Takeovers: Causes and Consequences*. (University of Chicago Press, Chicago).
- Shleifer, Andrei, and Robert W Vishny, 1997, The limits of arbitrage, *The Journal of Finance* 52, 35–55.
- Skinner, Douglas J, and Richard G Sloan, 2002, Earnings surprises, growth expectations, and stock returns or don't let an earnings torpedo sink your portfolio, *Review of accounting studies* 7, 289–312.
- Smales, Lee A., 2014, News sentiment and the investor fear gauge, *Finance Research Letters* 11, 122–130.
- Smales, Lee A., 2015, Time-variation in the impact of news sentiment, *International Review of Financial Analysis* 37, 40–50.
- Spangher, Alexander, 2015, Building the next new york times recommendation engine, *The New York Times* .
- Strömberg, Per, 2009, The economic and social impact of private equity in europe: summary of research findings, *Available at SSRN 1429322* .
- Tetlock, Paul C., 2007, Giving Content to Investor Sentiment: The Role of Media in the Stock Market, *The Journal of Finance* 62, 1139–1168, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2007.01232.x>.
- Tetlock, Paul C., 2011, All the News That's Fit to Reprint: Do Investors React to Stale Information?, *The Review of Financial Studies* 24, 1481–1512.
- Tetlock, Paul C., Maytal Saar-Tsechansky, and Sofus Macskassy, 2008, More Than Words: Quantifying Language to Measure Firms' Fundamentals, *The Journal of Finance* 63, 1437–1467, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2008.01362.x>.
- Uhl, Matthias W., and Milos Novacek, 2021, When it Pays to Ignore: Focusing on Top News and their Sentiment, *Journal of Behavioral Finance* 22, 461–479, Publisher: Routledge _eprint: <https://doi.org/10.1080/15427560.2020.1821375>.
- Uhl, Matthias W., Mads Pedersen, and Oliver Malitius, 2015, What's in the News? Using News Sentiment Momentum for Tactical Asset Allocation, *The Journal of Portfolio Management* 41, 100–112, Publisher: Institutional Investor Journals Umbrella.

- Vayanos, Dimitri, and Paul Woolley, 2012, A theoretical analysis of momentum and value strategies, *Working Paper, London School of Economics* .
- Vayanos, Dimitri, and Paul Woolley, 2013, An Institutional Theory of Momentum and Reversal, *The Review of Financial Studies* 26, 1087–1145.
- Vuolteenaho, Tuomo, 2002, What Drives Firm-Level Stock Returns?, *The Journal of Finance* 57, 233–264.
- Wang, Kevin Q., and Jianguo Xu, 2015, Market volatility and momentum, *Journal of Empirical Finance* 30, 79–91.
- Wright, Mike, Steve Thompson, and Ken Robbie, 1992, Venture capital and management-led, leveraged buy-outs: A European perspective, *Journal of Business Venturing* 7, 47–71.
- Zhang, Lu, 2005, The Value Premium, *The Journal of Finance* 60, 67–103.
- Zhou, Guofu, 2018, Measuring investor sentiment, *Annual Review of Financial Economics* 10, 239–259.

Chapter VII

Appendix

Appendix A. Detailed TRNA sentiment computation process

This appendix aims at detailing the entire text processing procedure employed by TRNA to compute their measure of news sentiment. First, a text pre-processing is put in place. Here, sentences are split, and individual words are tokenized (i.e., split from each other). Next, those words are stemmed, with the objective to reduce to a common basic form. This basic form aims to encompass all words with the same root and with equivalent meaning for the classification purpose. An example would be “argue, argued, argues, arguing”, which would all be reduced to the stem “argu”. Finally, the stemmed tokens are parsed with a basic part-of-speech tagger, which gives additional annotations to label tokens as being verbs, nouns, adjectives, etc.

As discussed in chapter I, one of the key challenges of Natural Language Processing, is dealing with the high dimensionality of the input text data, which can be problematic when fitting the data to a prediction model. Therefore, one of the key features of TRNA is to perform a thoughtful dimensionality reduction before the actual training of the classifier algorithm. This step focuses on feature extraction, i.e., the representation of the tokens along a reduced set of “atomic features”.

This step is achieved through manual classification and given the time-consuming effort it represents, this is where a commercial database with the necessary means can shine. Over 16,000 words and 2,500 phrases were hand-annotated by human analysts. Those words and phrases were then assigned to a variety of features based on the average human consensus annotation. Features used as dimensions include the impact on sentiment polarity, whether words act as intensifiers (e.g., very, hardly, etc.), or whether they act as negations (e.g., not good = bad). Other characteristics

which are manually annotated at this stage include disambiguation features based on the part-of-speech tags. For example, “fine” as a noun could mean “penalty”, whereas as an adjective, it could mean “good”. Verb resolution is also implemented to increase accuracy and to understand the specificity of words: in the example “ACME disappointed the market.”, disappointed is negative for the subject, ACME, but neutral for the object, the market. Once again, the part-of-speech tags are leveraged for this task. At the end of the process, all those abstractions allow representing the text in a lower-dimensional space. For example, “magnificent” and “fantastic” would have the same representation in this space.

With this low-dimensionality input, the model is finally trained based on 5,000 triple-annotated news articles. The classification model outputs the probability that the given news article is positive, neutral, or negative for the given entity (i.e., company) as follows:

$$\mathbb{P}(pos) + \mathbb{P}(neg) + \mathbb{P}(neut) = 1 \quad (\text{A1})$$

The classification algorithm takes the form of a three-layer back-propagation Neural Network with weight relaxation. The low input dimensionality achieved thanks to the steps described above prevents the need for more complicated models requiring more layers.

Figure (VII.1)¹ gives a good idea of the classification precision of TRNA’s algorithm. Generally, the automated classification is close to the one a human analyst would make. When forcing the classification of any news into either positive, neutral or negative, TRNA agrees with the average human assessment 75% of the time. This compares closely to the performance of a real analyst: On average, when three humans read a particular news, they agree on the classification in 82% of the cases. When recall is lowered, for example when only 50% of the most certain outputs is included, 87% of the classifications match with the average human assessment. Finally, cross-validation tests are performed in ten-fold batches. On average, removing one tenth of the training set in that way reduces classification accuracy by two percent.

Finally, note that Thomson Reuters provides additional meta-data, such as target audience and topic tags as selected by the author of the story². It also includes the

¹Source: TRNA – White Paper. Document version 1.1

²Topic codes are tags provided by the author of the story. They are based on a series of topics hand-picked by Reuters. Examples of topics include, Merges & Acquisitions, IPOs, Results, Litigation, Bankruptcy, etc. The exhaustive list of topics can be found at: <https://liaison.reuters.com/tools/topic-codes>. Audiences are the news services the articles are




Figure VII.1. Precision and Recall of the TRNA news classification algorithm: The plot should be read as follows: the x-axis reports the percentage of news on which a prediction is made. The y-axis reports the number of news that were accurately classified in the same category (either positive, neutral or negative) as the average human assessor. Thus, when considering any news, the algorithm outputs the right sentiment class $\sim 75\%$ of the time. Moving right along the x-axis displays the classification precision when increasing the confidence interval. For example, sentiment class predictions are correct 87% of the time when considering the half of news with most marked probabilities. When considering 20% of the news with the highest classification confidence, the algorithm achieves a $\sim 95\%$ precision level.

source of the news item. In addition, two extra metrics are provided: News *Volume*, which counts the number of stories in which the entity was mentioned over the past 24-hour cycle, and news *Novelty*, which counts the number of news in the past 24-hour cycle about that entity that also treat about a subject directly related in the current story. Stories are considered as linked when a topic similarity algorithm exceeds a given threshold³.

directed towards. Those can for example include “Money International Services”, “US General News Services”, etc.

³Not disclosed by TRNA.

Appendix B. News Tone Surprise

A potential limitation of our study relates to the nature of news *tone*. The implicit assumption behind the use of this variable is that the news it seeks to capture represents novel information for investors, i.e., unexpected news. In the table below, we seek to forecast firm-specific news *tone*. The rationale is as follows: if *tone* can be forecasted, then it is at least partly expected.

The table above summarizes all the relevant variables to forecast *tone*. We find that we can capture up to 16% of the variation in daily news *tone* with lagged variables. Including a firm's past day's news *tone* is the strongest predictor of current *tone*: the R^2 does not exceed 3% without it, as shown in column (7). Other significant variables include past returns, both at one day (*ret_1.1*) and over longer horizons of up to three months (*ret_2.90*). It implies that as prices rise, news about the firm in the media tend to subsequently be more positive. *tone* tends to be lower the day following an EAD, and thus they forecasts sentiment negatively. Past *tone* and return have a stronger predictive impact if the previous was an EAD. However, we find that traditional measures of standardized earnings surprises (SUE)⁴ as usually used in the literature are poor predictors of news *tone*.

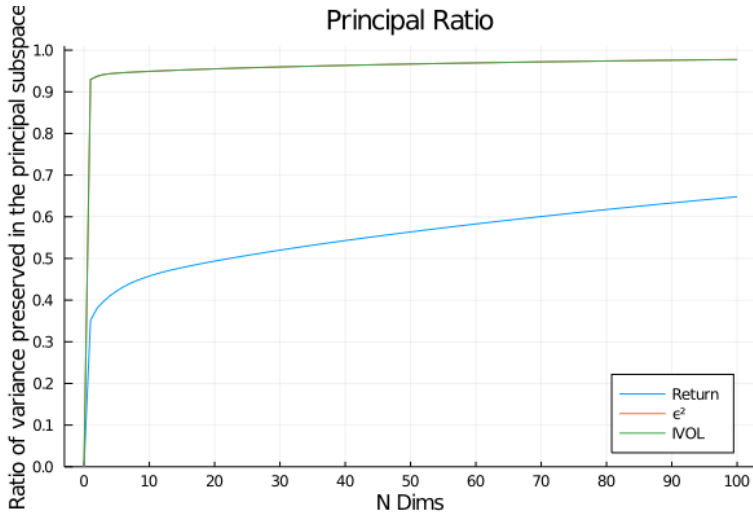
Our robustness measure for news *tone*, which we call *surp* for "surprise" is therefore the difference between the realized news *tone* and the expected news *tone* as computed as in specification (6), which focuses on all the significant variables.

⁴In this case SUE is the difference between the average analyst estimates and the actual earnings realization, as per I/B/E/S data. This measure is the same as the one used in Engelberg et al. (2018). Using alternatives measures of SUE based on expected earnings computed with accounting measures yield similar results which are not significant.

	SENT						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Intercept)	0.020*** (10.27)	0.021*** (11.04)	0.023*** (12.84)	0.078*** (63.95)	0.081*** (70.73)	0.080*** (77.14)	0.302*** (304.78)
TONE_1_1	0.127*** (99.14)	0.130*** (102.91)	0.131*** (106.94)				
TONE_2_90	0.702*** (231.33)	0.699*** (237.67)	0.697*** (254.98)				
ret_1_1	0.603*** (49.60)	0.587*** (50.30)	0.559*** (49.77)	0.554*** (50.12)	0.535*** (50.73)	0.488*** (49.31)	0.490*** (46.30)
ret_2_90	0.058*** (32.02)	0.055*** (31.85)	0.054*** (32.73)	0.044*** (31.92)	0.041*** (31.44)	0.038*** (31.66)	0.102*** (79.88)
EAD1_BM	-0.001*** (-5.40)	-0.001*** (-4.97)	-0.001*** (-5.35)	-0.002*** (-14.12)	-0.002*** (-13.89)	-0.002*** (-15.26)	-0.007*** (-62.98)
EAD1_Size	-0.001*** (-3.70)	-0.001*** (-4.39)	-0.001*** (-5.60)	-0.005*** (-45.15)	-0.006*** (-50.86)	-0.005*** (-52.06)	-0.017*** (-164.70)
EAD	-0.006** (-2.02)	-0.006** (-2.14)	-0.007** (-2.32)	-0.027*** (-13.61)	-0.032*** (-17.37)	-0.039*** (-23.34)	-0.127*** (-91.06)
EAD1_TONE	0.004** (2.30)	0.004** (2.57)		-0.007*** (-5.04)	-0.008*** (-6.30)		
EAD1_sue2	-0.005 (-1.49)			-0.004 (-1.43)			
TONE_1_1 & EAD	-0.088*** (-12.84)	-0.093*** (-13.95)	-0.096*** (-14.99)				
TONE_2_90 & EAD	-0.474*** (-32.41)	-0.470*** (-33.66)	-0.395*** (-31.32)				
ret_1_1 & EAD	0.295*** (3.95)	0.285*** (4.00)	0.391*** (5.82)	0.020 (0.37)	-0.025 (-0.52)	0.057 (1.26)	0.074 (1.53)
ret_2_90 & EAD	0.073*** (7.65)	0.070*** (7.62)	0.076*** (8.79)	0.046*** (7.58)	0.034*** (6.18)	0.046*** (9.05)	-0.001 (-0.17)
EAD1_TONE & EAD	0.148*** (16.09)	0.151*** (17.09)		0.226*** (41.02)	0.242*** (47.31)		
EAD1_sue2 & EAD	0.017 (0.79)			0.011 (0.68)			
TONE_1_90				0.806*** (405.28)	0.801*** (423.32)	0.779*** (467.33)	
TONE_1_90 & EAD				-0.630*** (-82.39)	-0.626*** (-89.33)	-0.489*** (-81.52)	
Estimator	OLS	OLS	OLS	OLS	OLS	OLS	OLS
N	653,980	679,484	718,324	1,264,685	1,354,772	1,519,910	1,521,685
Adjusted R ²	0.15	0.16	0.16	0.15	0.16	0.15	0.03

Appendix C. Principal Ratio

Figure VII.2. Principal Ratio: This figure illustrates how much of the variance is captured by the N first dimensions in the Principal Component Analysis. In blue, the line for returns shows that almost 70% of all the variation is captured by the first 100 dimensions.



Appendix D. Alternative VAR model estimations

Table II - S2 - VAR coefficient estimates: This table reports the estimated coefficient matrix Π for equations (III.9) in panel A and (III.10) in panel B. Calendar-quarter fixed are included in the aggregate specification. Errors are clustered by date and firm in the panel regression. The chosen sample restrictions correspond to S2, detailed in section III.C

	lnRealRet_v3	lnBM_cslag_v3	lnProf_v3	lnInv5_v3	lnME_D5_v3	lnMom_v3	lnROE_v3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Intercept)	-0.01*** (-4.36)	0.01*** (3.94)	0.00 (1.18)	0.00 (0.33)	-0.01* (-1.77)	-0.01*** (-4.57)	-0.00*** (-6.08)
lag_lnRealRet_v3	-0.02 (-0.84)	0.01 (0.60)	-0.00 (-0.63)	-0.00 (-0.22)	0.00 (0.01)	0.30*** (13.06)	-0.01*** (-8.48)
lag_lnBM_cslag_v3	0.01*** (3.10)	0.97*** (264.63)	0.00 (0.76)	-0.00*** (-4.32)	0.00 (0.32)	0.00 (0.95)	-0.01*** (-15.56)
lag_lnProf_v3	0.04*** (3.18)	-0.01 (-0.86)	0.24*** (3.21)	0.00*** (2.58)	0.05*** (3.65)	0.03** (2.23)	0.03*** (4.40)
lag_lnInv5_v3	-0.12** (-2.36)	0.12** (2.14)	-0.06*** (-2.60)	0.91*** (202.78)	0.21** (2.35)	-0.27*** (-3.86)	-0.02*** (-2.72)
lag_lnME_D5_v3	-0.01** (-2.10)	0.01*** (4.52)	0.00 (1.36)	0.00*** (18.51)	0.92*** (168.83)	-0.00 (-1.05)	0.00*** (5.73)
lag_lnMom_v3	0.03** (2.38)	-0.01 (-0.56)	0.02*** (4.15)	-0.00 (-0.59)	0.10*** (7.25)	0.69*** (57.24)	0.01*** (15.00)
N	154,030	154,030	154,030	154,030	154,030	154,030	154,030
R^2	0.01	0.94	0.06	0.94	0.89	0.58	0.14
Adjusted R^2	0.01	0.94	0.06	0.94	0.89	0.58	0.14

	lnRealRet_agg	lnBM_cslag_agg	lnProf_agg	lnInv5_agg	lnME_D5_agg	lnMom_agg	lnROE_agg
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
lag_lnRealRet_agg	0.26*** (2.73)	-0.20** (-2.16)	-0.02 (-0.50)	-0.00 (-0.05)	0.48*** (2.84)	0.52*** (3.97)	-0.01 (-1.06)
lag_lnBM_cslag_agg	0.05* (1.87)	0.95*** (38.04)	-0.01 (-0.83)	0.00 (1.27)	0.04 (0.83)	0.04 (0.98)	0.00 (1.21)
lag_lnProf_agg	-0.25 (-1.19)	0.45** (2.27)	-0.18** (-2.27)	0.01*** (2.93)	-0.72* (-1.95)	-0.32 (-1.12)	0.03** (2.19)
lag_lnInv5_agg	-2.65** (-2.05)	3.18** (2.62)	0.21 (0.44)	0.95*** (45.82)	-3.47 (-1.53)	-1.55 (-0.88)	0.24*** (3.09)
lag_lnME_D5_agg	0.02 (0.78)	-0.03 (-1.23)	-0.00 (-0.49)	0.00 (1.60)	0.93*** (19.63)	0.01 (0.15)	0.00 (0.87)
lag_lnMom_agg	-0.10** (-2.00)	0.10** (2.07)	0.00 (0.06)	0.00 (0.34)	-0.08 (-0.93)	0.64*** (9.28)	0.00 (1.15)
Q1	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Q2	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Q3	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	135	135	135	135	135	135	135
R^2	0.15	0.96	0.17	0.96	0.87	0.60	0.17
Adjusted R^2	0.07	0.95	0.09	0.95	0.85	0.56	0.10

Table III - S3 - VAR coefficient estimates: Same as previous. The chosen sample restrictions correspond to S3, detailed in section III.C

	<u>lnRealRet_v3</u>	<u>lnBM_cslag_v3</u>	<u>lnProf_v3</u>	<u>lnInv5_v3</u>	<u>lnME_D5_v3</u>	<u>lnMom_v3</u>	<u>lnROE_v3</u>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Intercept)	-0.04*** (-5.09)	0.04*** (4.36)	-0.01*** (-7.63)	0.00 (0.15)	-0.04*** (-2.72)	-0.01** (-2.34)	-0.01*** (-4.79)
lag_lnRealRet_v3	0.06 (1.33)	0.03 (0.72)	0.04*** (5.07)	0.01* (1.77)	0.34*** (6.98)	-0.01 (-0.37)	0.08*** (11.60)
lag_lnBM_cslag_v3	0.03*** (2.99)	0.94*** (87.21)	-0.01*** (-2.80)	-0.01*** (-6.52)	0.00 (0.36)	0.02** (2.57)	-0.03*** (-8.98)
lag_lnProf_v3	0.13*** (3.42)	0.03 (0.54)	0.70*** (20.19)	-0.00 (-0.13)	0.16*** (2.67)	0.09*** (3.50)	0.29*** (12.93)
lag_lnInv5_v3	-0.09** (-2.40)	0.09* (1.84)	-0.03*** (-3.64)	0.66*** (20.83)	0.02 (0.46)	-0.01 (-0.31)	-0.03** (-2.45)
lag_lnME_D5_v3	-0.02** (-2.36)	0.05*** (4.71)	-0.00 (-1.43)	0.02*** (9.25)	0.74*** (49.37)	-0.01* (-1.92)	0.01*** (4.12)
lag_lnMom_v3	0.05 (1.33)	-0.03 (-0.67)	0.03*** (4.86)	-0.01*** (-4.08)	0.02 (0.27)	0.06** (2.21)	0.02*** (2.67)
<i>N</i>	51,240	51,240	51,240	51,240	51,240	51,240	51,240
<i>R</i> ²	0.02	0.82	0.42	0.73	0.63	0.01	0.24
Adjusted <i>R</i> ²	0.02	0.82	0.42	0.73	0.63	0.01	0.24

	<u>lnRealRet_agg</u>	<u>lnBM_cslag_agg</u>	<u>lnProf_agg</u>	<u>lnInv5_agg</u>	<u>lnME_D5_agg</u>	<u>lnMom_agg</u>	<u>lnROE_agg</u>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Intercept)	0.05 (0.22)	-0.07 (-0.35)	0.05* (1.91)	-0.03* (-1.94)	-0.14 (-0.33)	0.06 (0.40)	0.03 (1.23)
lag_lnRealRet_agg	-0.01 (-0.04)	-0.01 (-0.08)	0.02 (0.75)	-0.03 (-1.66)	0.29 (0.78)	-0.12 (-0.89)	0.01 (0.31)
lag_lnBM_cslag_agg	0.04 (0.47)	0.90*** (12.01)	0.01 (1.01)	-0.00 (-0.48)	-0.01 (-0.06)	0.07 (1.13)	0.00 (0.06)
lag_lnProf_agg	1.42 (1.32)	-1.24 (-1.35)	0.80*** (5.93)	0.21** (2.55)	3.14 (1.57)	0.86 (1.17)	0.41*** (3.16)
lag_lnInv5_agg	-2.16* (-1.96)	2.69*** (2.86)	-0.03 (-0.24)	0.78*** (9.11)	-3.48* (-1.69)	-1.42* (-1.87)	0.04 (0.30)
lag_lnME_D5_agg	-0.05 (-0.53)	-0.02 (-0.28)	0.01 (0.50)	0.01 (1.29)	0.62*** (3.83)	0.03 (0.49)	-0.00 (-0.03)
lag_lnMom_agg	-0.04 (-0.13)	0.20 (0.82)	-0.07* (-1.79)	0.02 (0.94)	-0.32 (-0.60)	-0.06 (-0.30)	-0.03 (-0.93)
<i>N</i>	45	45	45	45	45	45	45
<i>R</i> ²	0.14	0.89	0.61	0.84	0.49	0.14	0.32
Adjusted <i>R</i> ²	0.00	0.87	0.55	0.82	0.41	0.00	0.21

Table IV - S4 - VAR coefficient estimates: Same as previous. The chosen sample restrictions correpsond to S4, detailed in section III.C

	<u>lnRealRet_v3</u>	<u>lnBM_cslag_v3</u>	<u>lnROE_v3</u>
	(1)	(2)	(3)
(Intercept)	-0.06*** (-4.75)	0.06*** (4.38)	-0.01*** (-4.85)
lag_lnRealRet_v3	0.06* (1.74)	0.05 (1.44)	0.10*** (13.07)
lag_lnBM_cslag_v3	0.05*** (3.75)	0.88*** (57.35)	-0.04*** (-11.32)
lag_lnROE_v3	-0.01 (-0.20)	0.09** (2.49)	0.19*** (11.06)
<i>N</i>	38,846	38,846	38,846
<i>R</i> ²	0.01	0.77	0.19
Adjusted <i>R</i> ²	0.01	0.77	0.19

	<u>lnRealRet_agg</u>	<u>lnBM_cslag_agg</u>	<u>lnROE_agg</u>
	(1)	(2)	(3)
(Intercept)	-0.07 (-0.29)	0.03 (0.13)	0.05** (2.11)
lag_lnRealRet_agg	-0.21 (-1.07)	0.15 (0.76)	-0.05** (-2.34)
lag_lnBM_cslag_agg	-0.04 (-0.47)	1.00*** (11.33)	-0.01 (-1.11)
lag_lnROE_agg	1.36 (0.85)	-0.52 (-0.32)	0.65*** (4.04)
<i>N</i>	27	27	27
<i>R</i> ²	0.07	0.87	0.47
Adjusted <i>R</i> ²	-0.05	0.85	0.40

Table V - S5 - VAR coefficient estimates: Same as previous. The chosen sample restrictions correpsond to S5, detailed in section III.C

	<u>lnRealRet_v3</u>	<u>lnBM_cslag_v3</u>	<u>lnROE_v3</u>
	(1)	(2)	(3)
(Intercept)	-0.05*** (-3.47)	0.05*** (2.68)	-0.02*** (-11.28)
lag_lnRealRet_v3	0.04 (1.24)	0.06* (1.94)	0.11*** (8.86)
lag_lnBM_cslag_v3	0.00 (0.36)	0.96*** (82.50)	-0.02*** (-6.67)
lag_lnROE_v3	-0.07 (-1.43)	0.14*** (2.89)	0.26*** (7.98)
<i>N</i>	30,303	30,303	30,303
<i>R</i> ²	0.00	0.86	0.16
Adjusted <i>R</i> ²	0.00	0.86	0.16

	<u>lnRealRet_agg</u>	<u>lnBM_cslag_agg</u>	<u>lnROE_agg</u>
	(1)	(2)	(3)
(Intercept)	0.87** (2.89)	-0.69*** (-4.64)	0.03 (0.93)
lag_lnRealRet_agg	0.24 (1.00)	-0.27** (-2.26)	-0.00 (-0.04)
lag_lnBM_cslag_agg	0.29 (1.25)	0.74*** (6.45)	-0.02 (-0.87)
lag_lnROE_agg	-4.91** (-2.53)	4.32*** (4.48)	0.59** (2.52)
Estimator	OLS	OLS	OLS
<i>N</i>	16	16	16
<i>R</i> ²	0.38	0.88	0.37
Adjusted <i>R</i> ²	0.22	0.85	0.22
Within- <i>R</i> ²			

Table VI - S6 - VAR coefficient estimates: Same as previous. The chosen sample restrictions correspond to S6, detailed in section III.C

	lnRealRet_v3	lnBM_cslag_v3	lnProf_v3	lnInv5_v3	lnME_D5_v3	lnMom_v3	lnROE_v3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Intercept)	-0.01*** (-3.54)	0.01*** (3.76)	0.00 (1.63)	0.00 (0.61)	-0.01 (-1.28)	-0.01*** (-4.38)	-0.00*** (-7.03)
lag_lnRealRet_v3	-0.03* (-1.80)	0.03* (1.94)	-0.00 (-0.21)	0.00 (0.26)	-0.00 (-0.26)	0.30*** (15.79)	-0.01*** (-7.35)
lag_lnBM_cslag_v3	0.01*** (4.22)	0.97*** (248.43)	0.00 (1.49)	-0.00*** (-5.17)	0.00 (0.33)	0.00 (1.53)	-0.01*** (-16.37)
lag_lnProf_v3	0.05*** (3.06)	-0.03 (-1.36)	0.27*** (3.52)	0.00*** (2.60)	0.06*** (4.25)	0.04*** (3.54)	0.03*** (4.15)
lag_lnInv5_v3	-0.14*** (-2.84)	0.13** (2.45)	-0.03 (-1.38)	0.91*** (210.31)	0.19** (2.42)	-0.29*** (-4.58)	-0.01 (-0.73)
lag_lnME_D5_v3	-0.01** (-2.47)	0.01*** (4.50)	0.00 (0.45)	0.00*** (18.91)	0.92*** (180.16)	-0.00 (-0.63)	0.00*** (6.15)
lag_lnMom_v3	0.03*** (3.02)	-0.01 (-0.89)	0.02*** (5.77)	-0.00 (-0.37)	0.09*** (8.18)	0.69*** (66.85)	0.02*** (19.43)
<i>N</i>	212,381	212,381	212,381	212,381	212,381	212,381	212,381
<i>R</i> ²	0.01	0.93	0.08	0.93	0.89	0.59	0.14
Adjusted <i>R</i> ²	0.01	0.93	0.08	0.93	0.89	0.59	0.14

	lnRealRet_agg	lnBM_cslag_agg	lnProf_agg	lnInv5_agg	lnME_D5_agg	lnMom_agg	lnROE_agg
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
lag_lnRealRet_agg	0.26*** (2.73)	-0.20** (-2.17)	-0.02 (-0.48)	-0.00 (-0.08)	0.48*** (2.85)	0.52*** (3.97)	-0.01 (-1.09)
lag_lnBM_cslag_agg	0.05* (1.87)	0.95*** (38.05)	-0.01 (-0.84)	0.00 (1.29)	0.04 (0.83)	0.04 (0.98)	0.00 (1.23)
lag_lnProf_agg	-0.25 (-1.18)	0.45** (2.28)	-0.17** (-2.19)	0.01*** (2.94)	-0.72* (-1.95)	-0.32 (-1.12)	0.03** (2.19)
lag_lnInv5_agg	-2.63** (-2.03)	3.18** (2.61)	0.21 (0.43)	0.95*** (45.93)	-3.48 (-1.53)	-1.56 (-0.88)	0.24*** (3.05)
lag_lnME_D5_agg	0.02 (0.77)	-0.03 (-1.23)	-0.01 (-0.51)	0.00 (1.63)	0.93*** (19.60)	0.01 (0.16)	0.00 (0.90)
lag_lnMom_agg	-0.10** (-2.01)	0.10** (2.09)	0.00 (0.04)	0.00 (0.34)	-0.09 (-0.96)	0.64*** (9.27)	0.00 (1.16)
Q1	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Q2	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Q3	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	135	135	135	135	135	135	135
<i>R</i> ²	0.15	0.96	0.16	0.96	0.87	0.60	0.17
Adjusted <i>R</i> ²	0.07	0.95	0.09	0.95	0.85	0.56	0.10

Appendix E. Scraping Glassdoor

To match our list of US-based companies from CapitalIQ with employee reviews from www.glassdoor.com, we set up a web scraping algorithm. The web crawler we use is based on "Scrapy", an open-source python library under BSD license⁵. We proceed as follows: (i) the web-scrafer starts by sending a request to the Glassdoor servers with the name of the target company. If Glassdoor returns multiple results for a given name we go for the preferred result with matching geographic location. (See Figure A1). (ii) Then, for the selected company, the web-scrafer gathers all the firm-related data in the overview panel. (see Figure A2). To ensure accuracy of the matching we force to have at least one of these additional features matched with CapitalIQ data: "City Headquarter", "Foundation" (with a margin of -1,+1 year.) or "Size/Number of employees" (grouped in bins of size 1 to 7.); (iii) As a final step, we download all the reviews for valid matches only.


Figure A3 depicts a typical review. The web crawler gathers the data from the title tag and saves what the reviewer left as a comment in the Pros, Cons (which are mandatory) and Advice to Management sections. At this step it also collects the overall Score (number of stars) and "recommended", "outlook" and "CEO" opinions. Importantly it also registers the date at which the review was posted and parses out i) the employment status (current or former employee) ii) the location and iii) for how long the reviewer worked (or had been been working) at the company. Finally, the crawler registers the scores related to Work/Life Balance, Culture, Career Opportunities, Compensations and Benefits and Senior Management (see Figure A4).

⁵We used Python 3.7 with Scrapy version 1.8.0.


petco Companies Location

Petco Reviews

Showing 1-6 of 6 Companies





Petco
www.petco.com / San Diego, CA, United States [+ Add a Review](#)


3.0 
48% recommended to a friend

3.1k Reviews	7.1k Salaries	643 Interviews
-----------------	------------------	-------------------


"Okay"
Coworkers were cool and I like being around animals [See all Reviews](#)

Sponsored Jobs


 Manager, Marketing Sciences Resolution Media – New York, NY	 Senior Back End Engineer (Golang) – AdTech Dailymotion – New York, NY
---	---




Petco (Iran)
www.pumpturbine.ir / Tehran, Tehran, Iran [+ Add a Review](#)

5.0 
100% recommended to a friend

1 Reviews	-- Salaries	-- Interviews
--------------	----------------	------------------



Petco Trading Labuan
Kuala Lumpur, Malaysia [+ Add a Review](#)

5.0 
100% recommended to a friend

1 Reviews	1 Salaries	-- Interviews
--------------	---------------	------------------

Figure A1. Searching a company on Glassdoor based on its name.

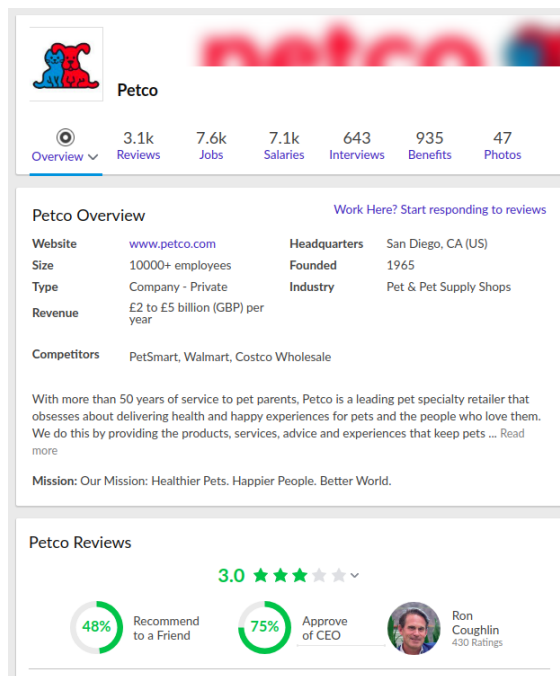



Figure A2. Overview of a company given by Glassdoor.

 **"Love and hate....the good and bad"**
 Current Employee - Partner in Knoxville, TN

Doesn't Recommend Neutral Outlook No Opinion of CEO

I have been working at Petco part-time for less than a year

Pros

- Watching the animals you personally work with, find a good home.
- Having first dibs on new products
- Sharing your knowledge about your favorite animal to a customer
- Supporting adoption from local shelters

Cons

- *Short staffed
- * You either work 5 to 6 days a week or you're just lucky to work 2 days a week
- *Training new people is non existent. You are shown once and then tossed right into the job.
- *Animal care is lacking
- *For certain price matching or discounts, you have to call up management to type in their password for the price to change. It is annoying, upsets the customer because they don't understand why a manger...

[Show More;](#)

Advice to Management

The higher ups insure we don't get to stay long to finish important tasks that effect animal health. Petco's higher ups are only concerned about payroll. So if a sick animal is in the wellness room but you cannot care properly for it because they want you to clock out on exact time. This actually causes a lot of animal deaths....but hey, at least we didn't go over payroll right? So what, is a poor small animal is...

[Show More;](#)



 [Helpful \(1\)](#)

Figure A3. Example of a review.

 **"Love and hate....the good and bad"**
 Current Employee - Partner in Knoxville, TN

Work/Life Balance	Neutral Outlook
Culture & Values	
Career Opportunities	
Compensation and Benefits	
Senior Management	

part-time for less than a year

personally work with, find a good home.

ducts

ut your favorite animal

Supporting adoption from local shelters

Cons

- *Short staffed
- * You either work 5 to 6 days a week or you're just lucky to work 2 days a week

Figure A4. Detailed scores along different welfare dimensions.

Appendix F. Glassdoor anecdotal evidence

Table A1 This Table presents example of reviews. We search on 'private equity', 'LBO', 'buyout', 'PE owners', 'private ownership', 'going private', 'Leveraged buy-out' for private equity firm names such as TPG, Apollo, Apax etc.

Date of review	Quotes
2009	At the end of February, 2 private equity firms bought our company and took it private. The nature of this transaction has resulted in several mid-level supervisors, managers and assistants being laid off. Currently, this leaves the employees with a dark uncertainty as to what will happen to them and their jobs in the near future, thus many are more stressed out than usual.
2009	I'll be checking the "ownership" and financial results of any company I join in the future.
2009	Private Equity firm pigs Apollo and TPG have eviscerated this company through massive layoffs and cuts. Customer service went out the window in the name of reducing expenses.
2009	I would never work with any of the management team that "managed" Mervyn's during it's private equity days. They ran the company into the ground.

Table A1 continued.

2009	For the last few years under Apollo Management, there was no corporate direction and the company did not invest and plan for the long term. The debt load taken on when they went private killed the company. (...)
2013	Capital constraints and private equity ownership continuously create operating challenges and unrealistic expectations
2013	Due to the financial downturn and private equity owner mentality, increasing the organization focuses on cost cutting and tends to dismiss the long-term implications of these decisions. Universally there is a lack of accountability, lack of revisiting the business case of approved projects, and holding senior and junior management accountable for poor outcomes.
2014	Public to private transition is complete and dust has settled. PE firm is hands off - Excellent work life balance
2014	Since the company went private last year there have been a lot of changes for better or worse. Some positives I can note on the company are a strong focus on safety and quality.

Table A1 continued.

2015	Solenis is recently purchased from Ashland Inc. by private equity Clayton, Dublier, & Rice. So far in the first year of new company, employee morale is terrible
2015	Wonderful people to work with. Prior to buy-out it was a company that cared about their employees.
2015	Entrepreneurial spirit, however, that is always being tested by current private equity ownership. New leadership has a vision and is moving rapidly to achieving it. Often do nice things for the associates;
2015	Management became too corporate with the buyout a couple of years ago. Treat your employees like they make a difference and they will!
2015	Private equity owners are just extracting rent now. My friends who still work there say innovation and best practices are being stifled, and the company feels gutted now.
2015	Since the leveraged buyout the company has changed direction. Negative environment and lost core competencies. They used to say people are their best asset. Now they a cost factor. Its all about Margins now.

Table A1 continued.

2015	Riverbed used to be a decent place to work, despite a core of people who acted like a "fun" cult. But since before the buyout they have completely lost direction. There are too many useless managers, and too many fulltime engineers working 30 hour weeks.
2015	Beginning of the year, there was some uncertainty as folks cashed out after the Private Equity buyout. When all your shares vest, it gives the \$\$-cushion to try something new. I get that. But that's all over and done now.
2015	Company going through changes. Acquisition by private equity firm has cut (or pushed out) a lot of the long time employees. Lets see what the future holds...
2015	Going private is slowly destroying the company. The primary focus now is cost ; cost reduction and associated metrics at the expense of the core business.
2015	Since the company went private last year there have been a lot of changes for better or worse. Some positives I can note on the company are a strong focus on safety and quality.

Table A1 continued.

2016	BMC is a company in transition. The corporate culture and morale used to be really great but since the company was taken private a few years ago, morale has fallen really low. Where there used to be a feeling of "let's all work together to make BMC successful", it's turned into an "every man for himself" culture with the primary focus being on the bottom line. So much so that management is more than willing to sacrifice..
2016	It all started falling apart last year when it went private and laid off 10-15% of its workforce. After all this, people lost confidence and almost all the people that made the company once great left. It all happened within an year
2018	After the Micro Focus buyout, all those accolades disappeared
2019	Acquisitions and buyouts leads to disgruntled employees

Appendix G. Additional Literature Review

A longstanding literature studies the impact of Leveraged Buy-Out (LBO) on human resources. Shleifer and Summers (1988) and Fox and Marcus (1992) discuss potential wealth transfers from current employees to new owners as LBOs are used as an opportunity to renegotiate employment contracts. Lichtenberg and Siegel (1990) find a decrease in non-production jobs. However, Davis et al. (2014) have shown that LBOs result in modest net job losses but large increase in gross job creation due to the exit of less productive establishment and greater entry of highly productive ones. This evidence finds explanations into the rationalization of jobs with for instance replacement of routine tasks by machines, offshoring and disappearance of middle wage workers (Olsson and Tåg (2017)) as well as the disposal of non-core parts of the business (Davis et al. (2014, 2012), Amess and Wright (2012)). Antoni et al. (2019) have shown that the decline employment is mainly found in the administrative staff. Wright et al. (1992) find an increase in employment and Davis et al. (2012) find an increase in greenfield jobs post MBIs. Antoni et al. (2019) show that the hiring of new people usually takes place in the first years after buyout, while cuts in jobs might occur later to improve the profitability of the deal.

Several studies highlight the heterogeneity of results across skilled and non-skilled workforce, countries, types of LBOs (corporate orphans, management buy-outs etc.), and in the presence or not of union. Employment effects have been shown to be more adverse in MBIs and LBOs due to an external management team and due to the fact that target companies are more likely to underperform (Amess and Wright (2007)). With regard to MBOs, IBOs usually prioritize fund returns and financial engineering over human resource policies which might hurt the well being of employees themselves (see Ludkin, interview 2008 in Goergen et al. (2014)). Local investors usually have a greater commitment than foreign investors to their social community. This usually leads to a more modest reduction in employment (Guery et al. (2017)). PE

effects found in divisional buyouts differ also from full LBOs with a likely increase in employment in divisional LBOs (Lichtenberg and Siegel (1990), Muscarella and Vetsuypens (1990)). Finally, we usually observe different impact on jobs from internal versus external management buyouts. On the one hand, it might be more difficult for external management team to value the current human workforce. On the other hand, it might be easier for them to break current working contracts (Goergen et al. (2014)). It has also been shown that a lower reduction in employment is expected if employment rights are stronger as protected by unions or worker collectives) (Goergen et al. (2014)). Public-firm buyouts are more likely to be accompanied by employment reductions (Davis et al. (2014)). Closely related to the LBO market, Cohn et al. (2016) show that higher leverage and more financial pressure lead to an increase in injury rates and decrease investments in worker safety.

Other aspects should also be taken into consideration when analyzing the impact buyout has on employees. First, economic conditions are likely to affect private equity activity, investment policies and operating performance (Kaplan and Strömberg (2009)). On the other hand, PE activity might induces some economic effects on industries. Bernstein et al. (2017) show that private equity activity leads to higher industry growth (for PE-backed or non PE-backed firms) without introducing more sales cyclical and business risk. Boucly et al. (2019) show that LBOs lead to important operating improvements and strong growth for targeted firms. Second, Agrawal and Tambe (2016) give evidence of a positive knowledge transfer induced by private equity ownership to existing employees. Bloom et al. (2015) examine management practices of PE-backed versus non-PE backed firms. Management quality in PE-backed firms is shown to be superior overall and especially with regard to setting objectives and monitoring. No significant differences are however found in incentives such as compensation and benefits given to employees. Yet, a recent study by Appelbaum (2019) retrieves testimonials on private equity buyout showing ev-

idence on how buyout hurt companies' financial and ultimately employee welfare. Third, Lerner et al. (2011) investigates whether private equity ownership relieves the management from being short term focus. They show that patent quality and activity improve for firms under PE ownership.

Appendix H. Case Study: Vista

Vista⁶ buys only software. It can be software in a variety of sectors but the product is a software. Examples from Vista's portfolio: Accelya provides software for transportation management, Kazoo is a HR software. Vista Equity Partners: 2016 buyout of Marketo, a cloud-based innovator in marketing automation software. Marketo was losing money. Took it private for \$1.78 billion, 64% premium, 8x revenue, a huge Valuation. Vista installed an experienced CEO, who focused the company's sales effort on large deals in the enterprise space. Rapid growth eventually produced positive EBITDA, and just two years later, Vista sold the company to Adobe for \$5 billion.

Employees of acquired companies and candidates for hiring must submit to tests. A personality test aims to determine which of them are suited to which jobs. Salespeople are better off being extroverted, and software developers more introverted. A proprietary cognitive assessment, similar to an IQ test, includes questions on logic, pattern recognition, vocabulary, sentence completion and math. The test inspires consternation and fear among existing employees, according to former employees. The goal of the Austin, Texas-based firm, which is 18 years old, is to transform business-software companies into profit machines. Behind its approach is Mr. Smith's belief that certain aspects of the companies Vista buys are interchangeable. "Software companies taste like chicken," he said at a conference in New York a few years ago. "They're selling different products, but 80% of what they do is pretty much the same."

Former employees say cost cutting is critical to Vista's model. Some of the companies Vista takes over are located in markets with a high cost of living, such as Southern California or New York City. To tamp down wages and other costs, Vista will relocate part or all of the company to a less-expensive city such as Dallas.

⁶<https://www.wsj.com/articles/billionaires-secret-buyout-formula-110-instructions-and-an-intelligence-test-1531151197>

Many employees won't make the move, allowing Vista to hire cheaper replacements. Vista often keeps a company's headquarters in place and encourages it to expand in lower-cost markets.

Appendix I. Details on position classification

Employees report their position in the company in an open field, which provides guided suggestions as you type, which helps for uniformization and to avoid typos. We go through the 500 most frequent unique entries and classify them manually using a specialized guidebook⁷. Based on this, we use the following rules to categorize all employees as follows:

- Manager (Mngt) : Each job title containing the words "director" or "vice president" are in this category (e.g. senior director, associate director, vice president, senior vice president) unless the words assistant or sales are also present.
- Middle Management (MidMngt) : Each job title containing the words manager or leader are in this category (account manager, project manager, store manager, team leader, store leader).
- White Collars (WhiteC) : Each job title containing the words consultant, executive, assistant-manager, analyst, specialist are in this category.
- Purple Collar (PurpleC) : Each job title containing the words engineer or software are in this category.
- Pink Collar (PinkC) : Each job title containing the words sales or administrative or assistant are in this category. Teachers and marketing are also included here.
- Blue Collar (BlueC) : Each job title containing the words technician or driver are in this category. Cashier and servers are also included here.

⁷'Work in America', page 597, ISBN.9781576076767.

Appendix J. Number of Reviews around transaction date




Figure A1. Current employees. Number of reviews in quarter around deal, with quarter fixed-effects. Current employees, keeping those who might have joined post-deal. We first compute a panel where we count the number of reviews for each quarter between -12 and 12 from the deal. Using this panel, we then run the following regression: $Nreviews = \alpha + \beta * DistanceQ + QFE + \epsilon$. We report the coefficients and error bars of coefficient β . Here we use our final working sample.

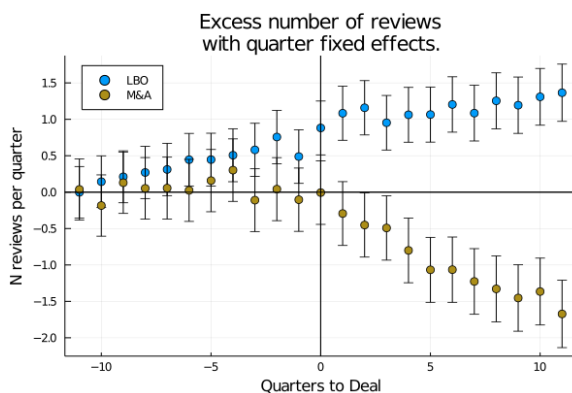


Figure A2. Former employees. Identical procedure as A1. The sample consists of former employees.

Appendix K. Sub-scores

Table A2: Sub-scores

This table shows the same specification as table V - specification 4 for the different sub-score categories.

Panel A: Sub-scores baseline									
	Score	WLB	CO	CB	SM	Cult	CEO	Outl	Reco
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Public2PE	-0.21*** (-4.81)	-0.19*** (-3.32)	-0.21*** (-4.02)	-0.13*** (-3.58)	-0.18*** (-3.86)	-0.15*** (-4.02)	-0.21** (-2.07)	-0.29*** (-4.19)	-0.27*** (-3.35)
Private2PE	-0.10** (-2.21)	-0.10** (-2.38)	-0.11** (-2.49)	-0.09*** (-2.90)	-0.10** (-2.49)	-0.12*** (-3.14)	-0.10* (-1.74)	-0.11* (-1.78)	-0.10* (-1.91)
PE2PE	-0.07 (-1.47)	-0.06 (-1.02)	-0.02 (-0.38)	-0.02 (-0.36)	-0.01 (-0.21)	-0.09** (-2.17)	-0.04 (-0.75)	-0.05 (-0.76)	-0.02 (-0.39)
postIPO	-0.06* (-1.83)	-0.07 (-1.61)	-0.03 (-1.02)	-0.03 (-1.07)	-0.06 (-1.60)	-0.04 (-1.37)	-0.04 (-1.10)	-0.07 (-1.50)	-0.07* (-1.84)
postMA	-0.05* (-1.84)	-0.08** (-2.51)	-0.07** (-2.42)	-0.03 (-1.13)	-0.09*** (-3.63)	-0.05* (-1.91)	-0.16*** (-4.22)	-0.10** (-2.56)	-0.10** (-2.53)
Mngt	0.21*** (13.47)	0.27*** (16.52)	0.24*** (16.15)	0.08*** (5.61)	0.25*** (15.05)	0.17*** (12.48)	0.19*** (10.98)	0.27*** (14.27)	0.20*** (12.39)
MidMngt	0.12*** (11.34)	0.14*** (12.44)	0.15*** (13.09)	-0.02** (-1.99)	0.20*** (18.12)	0.12*** (10.95)	0.13*** (11.28)	0.16*** (12.13)	0.12*** (10.66)
WhiteC	0.08*** (11.28)	0.10*** (14.09)	0.11*** (14.96)	0.12*** (16.97)	0.10*** (13.14)	0.05*** (7.07)	0.08*** (10.66)	0.13*** (12.04)	0.08*** (10.96)
PurpleC	0.03** (2.38)	0.00 (0.45)	0.02 (1.53)	0.02** (2.24)	0.01 (0.93)	-0.02 (-1.55)	-0.03*** (-2.58)	0.07*** (4.27)	0.01 (0.51)
PinkC	0.03*** (4.35)	0.04*** (4.04)	0.04*** (4.63)	0.07*** (7.23)	0.00 (0.63)	0.02*** (2.74)	0.00 (0.32)	0.06*** (5.34)	-0.01 (-0.55)
BlueC	-0.01 (-0.57)	-0.05*** (-3.90)	-0.04*** (-3.60)	0.00 (0.15)	-0.03*** (-3.00)	0.02* (1.84)	-0.10*** (-6.57)	0.01 (0.79)	-0.05*** (-3.89)
Tenure_0X	0.25*** (13.67)	0.44*** (16.26)	0.31*** (15.28)	0.28*** (14.71)	0.34*** (17.15)	0.17*** (11.19)	0.26*** (13.96)	0.33*** (13.90)	0.29*** (15.81)
Tenure_1X	0.08*** (8.28)	0.19*** (13.04)	0.12*** (12.01)	0.15*** (13.09)	0.12*** (11.05)	0.03*** (2.74)	0.15*** (12.86)	0.15*** (9.50)	0.15*** (12.04)
Tenure_3X	0.01 (0.55)	0.05*** (4.47)	0.03*** (3.46)	0.07*** (7.63)	0.03*** (3.73)	-0.02*** (-2.68)	0.07*** (7.18)	0.06*** (4.28)	0.06*** (4.65)
Tenure_5X	0.01 (1.01)	0.03** (2.23)	0.02** (2.42)	0.04*** (4.47)	0.02** (2.48)	-0.00 (-0.12)	0.04*** (3.20)	0.04*** (3.36)	0.03** (2.47)
Tenure_10X	0.10*** (11.85)	0.10*** (7.90)	0.10*** (11.05)	0.05*** (5.88)	0.10*** (11.71)	0.12*** (16.37)	0.04*** (2.93)	0.10*** (6.17)	0.08*** (5.37)
log(wage)	0.14*** (14.41)	0.13*** (11.97)	0.11*** (9.94)	0.13*** (11.28)	0.18*** (16.45)	0.21*** (15.87)	0.17*** (13.72)	0.16*** (12.10)	0.16*** (14.65)
Quarter Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	715,375	670,146	661,274	680,509	678,432	679,396	566,034	633,902	607,512
<i>Adj - R</i> ²	0.13	0.14	0.13	0.11	0.13	0.15	0.13	0.09	0.11

Panel B: Management Cross Effects									
	Score	CB	CO	WLB	SM	Cult	CEO	Reco	Outl
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
POSTLBOdeal & Mngt	0.00	0.10	0.05	-0.04	0.05	-0.04	0.07	-0.03	-0.02
	(0.05)	(1.40)	(0.70)	(-0.80)	(0.65)	(-0.60)	(0.94)	(-0.38)	(-0.26)
POSTLBOdeal & MidMngt	-0.08*	-0.05	-0.07	-0.04	-0.02	-0.07**	-0.04	-0.10*	-0.08
	(-1.69)	(-1.01)	(-1.35)	(-0.95)	(-0.37)	(-2.06)	(-0.65)	(-1.81)	(-1.33)
POSTLBOdeal & NotMngt	-0.14***	-0.14***	-0.13***	-0.08***	-0.12***	-0.14***	-0.14***	-0.16***	-0.15***
	(-5.84)	(-4.72)	(-4.61)	(-4.22)	(-4.75)	(-6.14)	(-3.62)	(-4.54)	(-3.87)
postIPO	-0.06*	-0.06	-0.03	-0.03	-0.05	-0.04	-0.04	-0.06	-0.07*
	(-1.82)	(-1.60)	(-1.00)	(-1.07)	(-1.59)	(-1.37)	(-1.08)	(-1.49)	(-1.83)
postMA	-0.05*	-0.08**	-0.07**	-0.03	-0.09***	-0.05*	-0.16***	-0.10**	-0.10**
	(-1.83)	(-2.51)	(-2.42)	(-1.13)	(-3.63)	(-1.91)	(-4.22)	(-2.56)	(-2.53)
Quarter Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	715,375	670,146	661,274	680,509	678,432	679,396	566,034	633,902	607,512
<i>Adj - R</i> ²	0.13	0.14	0.13	0.11	0.13	0.15	0.13	0.09	0.11

Panel C: Industry Cross Effects									
	Score	CB	CO	WLB	SM	Cult	CEO	Reco	Outl
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
POSTLBOdeal & OtherInd	-0.02	-0.02	-0.01	-0.02	-0.01	-0.07	-0.04	0.00	0.02
	(-0.34)	(-0.29)	(-0.23)	(-0.49)	(-0.11)	(-1.62)	(-0.66)	(0.07)	(0.30)
POSTLBOdeal & CorporateService	-0.05	-0.05	-0.07	-0.04	-0.04	-0.11*	-0.06	-0.01	-0.00
	(-1.02)	(-0.90)	(-1.35)	(-0.75)	(-0.71)	(-1.71)	(-0.81)	(-0.08)	(-0.06)
POSTLBOdeal & Industrial	-0.09	-0.08	-0.09	-0.10*	-0.07	-0.09	-0.03	-0.18	-0.11
	(-1.15)	(-0.94)	(-1.06)	(-1.88)	(-0.83)	(-1.21)	(-0.26)	(-1.63)	(-0.98)
POSTLBOdeal & Retail	-0.22***	-0.20***	-0.23***	-0.18***	-0.11**	-0.10***	-0.16	-0.27***	-0.30***
	(-4.30)	(-3.66)	(-3.58)	(-3.50)	(-2.23)	(-3.01)	(-0.97)	(-4.51)	(-4.76)
POSTLBOdeal & Software	-0.24***	-0.21***	-0.17***	-0.08**	-0.23***	-0.20***	-0.22***	-0.30***	-0.26***
	(-4.52)	(-2.92)	(-2.98)	(-2.49)	(-3.96)	(-5.59)	(-3.19)	(-3.76)	(-2.97)
postIPO	-0.06*	-0.07	-0.03	-0.03	-0.05	-0.04	-0.04	-0.06	-0.07*
	(-1.82)	(-1.60)	(-1.01)	(-1.07)	(-1.59)	(-1.37)	(-1.09)	(-1.49)	(-1.83)
postMA	-0.05*	-0.08**	-0.07**	-0.03	-0.09***	-0.05*	-0.16***	-0.10**	-0.10**
Quarter Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	715,375	670,146	661,274	680,509	678,432	679,396	566,034	633,902	607,512
<i>Adj - R</i> ²	0.13	0.14	0.13	0.11	0.13	0.15	0.13	0.09	0.11

Appendix L. Cross-effects

The following three panels are the continuation of table VII, including panels B, C and D detailed in said table caption.

Panel B: Post M&A Cross Effects				
	Dependent variable: Overall Score			
	(1)	(2)	(3)	(4)
postMA & Small	-0.08** (-2.03)			
postMA & Medium	-0.08 (-1.41)			
postMA & Big	-0.03 (-0.65)			
postMA & Mature		-0.02 (-0.69)		
postMA & Millenials		-0.17** (-2.06)		
postMA & Management			-0.09 (-1.57)	
postMA & Mid Management			-0.06 (-0.91)	
postMA & Not Management			-0.05* (-1.84)	
postMA & Other Services				-0.09 (-1.61)
postMA & Corporate Service				-0.04 (-0.88)
postMA & Industrial				-0.05 (-0.70)
postMA & Retail				0.02 (0.78)
postMA & Software				-0.02 (-0.34)
Control Variables	Yes	Yes	Yes	Yes
Quarter Fixed-Effects	Yes	Yes	Yes	Yes
Company Fixed-Effects	Yes	Yes	Yes	Yes
<i>N</i>	715,375	715,375	715,375	715,375
<i>Adjusted R</i> ²	0.13	0.13	0.13	0.13

Panel C: Software vs. GP Fixed-Effects

	Score										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
POSTLBOdeal	-0.04 (-1.22)	-0.04 (-1.22)	-0.04 (-1.22)	-0.04 (-1.16)	-0.04 (-1.50)	-0.03 (-0.97)	-0.03 (-0.95)	-0.03 (-1.02)	-0.03 (-0.41)	-0.03 (-0.35)	-0.03 (-0.33)
POSTLBOdeal & Retail	-0.18*** (-2.77)	-0.18*** (-2.77)	-0.18*** (-2.77)	-0.18*** (-2.80)	-0.18* (-1.89)	-0.18*** (-2.86)	-0.18*** (-2.88)	-0.19 (-0.94)	-0.19 (-1.03)	-0.19 (-1.07)	-0.19 (-1.22)
POSTLBOdeal & Software	-0.20*** (-3.01)	-0.20*** (-2.90)	-0.21*** (-2.91)	-0.18*** (-2.71)	-0.17*** (-2.75)	-0.09* (-1.82)	-0.08* (-1.70)	-0.08 (-0.42)	-0.09 (-0.54)	-0.09 (-0.37)	-0.09 (-0.45)
POSTLBOdeal & Marlin Equity Partners		-0.20 (-1.15)	-0.20 (-1.15)	-0.23 (-1.31)	-0.23 (-1.54)	-0.31* (-1.76)	-0.32* (-1.82)	-0.33 (-1.21)	-0.32*** (-10.93)	-0.32*** (-8.91)	-0.32*** (-6.96)
POSTLBOdeal & Insight Partners			0.10 (0.48)	0.07 (0.35)	0.07 (0.37)	-0.01 (-0.05)	-0.02 (-0.09)	-0.02 (-0.10)	-0.01 (-0.04)	-0.02 (-0.05)	-0.01 (-0.05)
POSTLBOdeal & Thoma Bravo				-0.15 (-0.85)	-0.15 (-0.96)	-0.23 (-1.34)	-0.24 (-1.38)	-0.24 (-1.60)	-0.23 (-1.06)	-0.22 (-0.82)	-0.22 (-0.69)
POSTLBOdeal & Vector Capital					-0.08 (-0.79)	-0.14 (-1.20)	-0.15 (-1.25)	-0.16 (-0.92)	-0.15 (-1.37)	-0.15 (-1.29)	-0.15 (-1.15)
POSTLBOdeal & Vista Equity Partners						-0.31*** (-3.33)	-0.31*** (-3.46)	-0.32*** (-3.15)	-0.31* (-1.94)	-0.31* (-1.63)	-0.31* (-1.73)
POSTLBOdeal & TA Associates							-0.39 (-1.08)	-0.39*** (-4.96)	-0.43*** (-4.90)	-0.43*** (-2.26)	-0.43** (-2.26)
POSTLBOdeal & Francisco Partners								-0.16*** (-3.12)	-0.15*** (6.61)	-0.15*** (6.06)	-0.15 (2.25)
POSTLBOdeal & Clearlake Capital Group									(-2.19)	(-2.03)	(-1.06)
POSTLBOdeal & Silver Lake									0.32*** (6.61)	0.31*** (6.06)	0.31** (2.25)
POSTLBOdeal & Siris Capital										-0.23 (-1.28)	-0.23*** (-3.95)
Quarter Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375
Adj. R ²	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13

Panel D: Retail vs. GP Fixed-Effects

	Score										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
POSTLBOdeal	-0.04 (-1.22)	-0.04 (-1.21)	-0.05 (-1.39)	-0.05 (-1.41)	-0.05 (-1.48)	-0.06* (-1.66)	-0.06* (-1.70)	-0.06* (-1.68)	-0.06*** (-2.65)	-0.06*** (-2.37)	-0.06*** (-2.23)
POSTLBOdeal & Retail	-0.18*** (-2.77)	-0.16* (-1.87)	-0.16* (-1.90)	-0.16* (-1.91)	-0.16** (-2.04)	-0.16*** (-1.99)	-0.16*** (-2.00)	-0.16*** (-2.00)	-0.16*** (-2.45)	-0.16*** (-2.63)	-0.17*** (-2.59)
POSTLBOdeal & Software	-0.20*** (-3.01)	-0.20*** (-3.00)	-0.20*** (-2.91)	-0.20*** (-2.93)	-0.20*** (-2.90)	-0.19*** (-2.82)	-0.20*** (-2.87)	-0.20*** (-2.87)	-0.21** (-2.48)	-0.20*** (-2.48)	-0.21*** (-3.01)
POSTLBOdeal & Sycamore Partners	-0.06 (-0.69)	-0.05 (-0.69)	-0.05 (-0.62)	-0.05 (-0.59)	-0.04 (-0.52)	-0.04 (-0.50)	-0.04 (-0.48)	-0.04 (-0.47)	-0.04 (-0.44)	-0.03 (-0.38)	-0.03 (-0.45)
POSTLBOdeal & Apex Partners	0.15* (1.08)	0.15* (1.08)	0.15* (1.08)	0.15* (1.08)	0.16* (1.17)	0.16* (1.80)	0.16* (1.81)	0.16* (1.83)	0.17 (0.52)	0.17 (0.52)	0.17 (0.72)
POSTLBOdeal & Hellman & Friedman	0.19 (0.51)	0.19 (0.51)	0.19 (0.51)	0.19 (0.51)	0.19 (0.52)	0.19 (0.52)	0.19 (0.53)	0.20 (0.53)	0.20 (0.55)	0.20 (0.56)	0.20 (0.51)
POSTLBOdeal & Warburg Pincus	0.58 (1.20)	0.58 (1.20)	0.58 (1.21)	0.57 (1.20)	0.57 (1.20)	0.58 (1.21)	0.58 (1.21)	0.58 (1.21)	0.58*** (4.45)	0.58*** (4.34)	0.58*** (2.54)
POSTLBOdeal & Investcorp	0.17 (1.07)	0.17 (1.07)	0.17 (1.07)	0.17 (1.07)	0.17 (1.07)	0.17 (1.07)	0.17 (1.07)	0.17 (1.07)	0.17 (0.62)	0.17 (0.65)	0.17 (0.94)
POSTLBOdeal & Bain Capital Ventures	0.09 (0.49)	0.09 (0.49)	0.09 (0.49)	0.09 (0.49)	0.09 (0.49)	0.09 (0.49)	0.09 (0.49)	0.09 (0.49)	0.10 (0.71)	0.10 (0.54)	0.10 (0.38)
POSTLBOdeal & Sentinel Capital Partners	0.02 (0.10)	0.02 (0.10)	0.02 (0.10)	0.02 (0.10)	0.02 (0.10)	0.02 (0.10)	0.02 (0.10)	0.02 (0.10)	0.02 (0.13)	0.02 (0.17)	0.02 (0.13)
POSTLBOdeal & Clearlake Capital Group	0.40** (2.38)	0.40** (2.38)	0.40** (2.38)	0.40** (2.38)	0.40** (2.38)	0.40** (2.38)	0.40** (2.38)	0.40** (2.38)	0.40** (2.38)	0.40** (2.29)	0.40** (2.08)
POSTLBOdeal & Clayton Dubilier & Rice	0.14 (1.00)	0.14 (1.00)	0.14 (1.00)	0.14 (1.00)	0.14 (1.00)	0.14 (1.00)	0.14 (1.00)	0.14 (1.00)	0.14 (0.71)	0.14 (0.71)	0.14 (1.00)
POSTLBOdeal & Platinum Equity	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)	-0.23 (-1.06)
Quarter Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375	715,375
Adj. R ²	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13

Appendix M. Fitting LDA models: Topic Coherence

As discussed, the LDA model uses a Bayesian approach that relies on three parameters that determine the shape of distribution of topics: a hyper-parameter α , a hyper-parameter β and the ex-ante choice on number of topics desired. The α parameter controls the shape of the Dirichlet distribution of documents on topics. Lower (higher) values of α will cause each document to be composed of fewer (more) different dominant topics. The β parameter controls the shape of the Dirichlet distribution of words on topics. Lower (higher) values of β will cause each topic to be composed of a smaller (bigger) set of dominant words.

We follow the work of Röder et al. (2015) to identify the optimal parameters. The issue in choosing appropriate parameters is in evaluating ex-post what is an appropriately fitted model. The suggestion is to compute a coherence measure of the topics generated by the model, which are based on probabilities that words inside a same topic effectively co-occur inside a document. We use the "gensim" package in python to effectively compute the coherence score for different values of α , β and topic number.

We report the results for the 'Cons' sample in table A3. To maximize topic coherence we seek the highest possible score. Since those computations are time-consuming we can not performing an exhaustive search of the entire space of parameters. Nonetheless, our results help us provide an objective measure on why the topics are of greater quality when going down to 25 topics. We can also observe that coherence is relatively robust to the choice of α and β . To ease our high-level understanding of topics it is desirable to go for lower values of both α and β . This helps us label topics when there are fewer words that capture the idea of the topic. We therefore end up at settling for 25 topics, $\alpha = 0.1$ and $\beta = 0.05$.

Table A3- Coherence Score.

This table reports coherence scores (based on an implementation of (Rehurek et al., 2011)) computed for different values of α (rows) and β (columns) hyperparameters and for different number of topics.

Panel A: 150 topics							
α/β	0.00005	0.0001	0.001	0.005	0.02	0.05	0.1
0.0005	-5.02	-5.14	-4.97	-5.03	-5.21	-5.03	-5.01
0.001	-5.19	-5.08	-5.10	-5.12	-5.13	-5.14	-4.98
0.025	-5.26	-5.43	-5.30	-5.19	-5.43	-5.49	-5.25
0.005	-5.05	-5.02	-5.07	-5.10	-5.07	-5.09	-5.00
0.01	-5.16	-5.13	-5.02	-5.23	-5.20	-5.07	-5.14
0.05	-5.73	-5.68	-5.59	-5.8	-5.71	-5.73	-5.63

Panel B: 100 topics							
α/β	0.00005	0.0001	0.001	0.005	0.02	0.05	0.1
0.0005	-4.44	-4.42	-4.44	-4.64	-4.47	-4.37	-4.49
0.001	-4.48	-4.42	-4.45	-4.48	-4.39	-4.43	-4.42
0.025	-4.50	-4.53	-4.57	-4.49	-4.60	-4.63	-4.70
0.005	-4.43	-4.46	-4.42	-4.48	-4.54	-4.37	-4.44
0.01	-4.42	-4.42	-4.49	-4.54	-4.49	-4.40	-4.46
0.05	-4.80	-4.73	-4.71	-4.65	-4.68	-4.77	-4.65

Panel C: 50 topics							
α/β	0.00005	0.0001	0.001	0.005	0.02	0.05	0.1
0.0005	-3.93	-3.86	-3.85	-3.95	-3.91	-3.86	-3.85
0.001	-3.85	-3.86	-3.87	-3.81	-3.88	-3.85	-3.83
0.025	-3.91	-3.9	-3.83	-3.88	-3.81	-3.89	-4.00
0.005	-3.83	-3.84	-3.89	-3.94	-3.87	-3.85	-3.90
0.01	-3.93	-3.86	-3.97	-3.83	-3.84	-3.91	-3.80
0.05	-3.83	-3.94	-3.97	-3.86	-3.86	-4.06	-4.03

Panel D: 25 topics							
α/β	0.00005	0.0001	0.001	0.005	0.02	0.05	0.1
0.0005	-3.80	-3.68	-3.71	-3.83	-3.7	-3.74	-3.69
0.001	-3.76	-3.67	-3.77	-3.67	-3.69	-3.7	-3.63
0.025	-3.81	-3.73	-3.68	-3.69	-3.60	-3.69	-3.89
0.005	-3.63	-3.66	-3.73	-3.81	-3.66	-3.73	-3.79
0.01	-3.78	-3.73	-3.77	-3.71	-3.66	-3.71	-3.66
0.05	-3.74	-3.79	-3.8	-3.73	-3.70	-3.94	-3.84

Appendix N. Detailed LDA topics

Table A4
Detailed LDA Topics: Cons Reviews

Cons Topics	Top Ngrams	LBO example	Other example
<i>Poor Upper Management - Lack of communication</i>	manag; poor; lack; communicate; upper; train; upper manag; poor manag; support; micro; micro manag; staff; life; bad; balance; level; lack communication; poor communicate; turnover; littl manag; chang; leadership; constant; compani; cultur; senior; upper; upper manag; direct; excute; decisi; senior manag; constant chang; environ; middl; corpor; frequent; ceo; busi	Disconnect between upper management and other staff members, high turnover rates Change in the management focus and objective due to change in the senior management. Flow of information from the top management can be better	benefits, poor management, really bad upper management
<i>Changing management</i>	manag; chang; leadership; constant; compani; cultur; senior; upper; upper manag; direct; excute; decisi; senior manag; constant chang; environ; middl; corpor; frequent; ceo; busi	Over the last couple of years Company seems to have lost it's way. A lot of mismanagement by the executives put the company in a place where they had to sell out to venture capital firm Vista. The layoffs and cost cutting have begun morale is bad and getting worse. I not sure what the future holds for us	Management direction changes constantly and frequently, leading to constant churn Company feels like it is circling the drain. No stability or job security. Lots of wasted talent in the company-- morale is ruined by corporate financial scandals, a botched merger, and constant reductions in force. 10% layoffs in the last month resulted in almost 200 people gone. With a cloudy business outlook and no buyout in sight, there is a sinking ship feeling in the ranks.
<i>Lag-offs</i>	compani; cut; employe; layoff; peopl; job; cost; busi; leav; term; profit; recent; moral; futur; due; ceo; lay; stock; focus; secur	The company has really taken a turn for the worse since the new CEO, Eric Jungbluth came into the company. He cares nothing for the employees (has even been heard saying that he does not value them and they can be replaced easily). This company has a collaborative mindset. It's team-based. Lone wolves who are not open to learning and growing might struggle. Internal processes are often lacking, so you need good problem solving skills to get things done. Communication at an executive level needs work. Too many priorities create unnecessary risk.	Walmart don't care about family or you just make money no matter is you are we them for 1 or 20 yrs they don't show any respect.
<i>Management not caring about employees</i>	care; employe; compani; manag; dow; down; peopl; manag; job; care; employe; dow; care; partent; live; custom; manag; care; staff; bottom; corpor; down; care; manag; down	The company has really taken a turn for the worse since the new CEO, Eric Jungbluth came into the company. He cares nothing for the employees (has even been heard saying that he does not value them and they can be replaced easily). This company has a collaborative mindset. It's team-based. Lone wolves who are not open to learning and growing might struggle. Internal processes are often lacking, so you need good problem solving skills to get things done. Communication at an executive level needs work. Too many priorities create unnecessary risk.	I feel that my Management team has started heading the company in a less than ideal direction for many people.
<i>Management lead of team</i>	team; manag; peopl; project; depend; compani; experi; manag team; lead; lot; cultur; dow; skill; offic; leader; person; technic; depart; feel	This company has a collaborative mindset. It's team-based. Lone wolves who are not open to learning and growing might struggle. Internal processes are often lacking, so you need good problem solving skills to get things done. Communication at an executive level needs work. Too many priorities create unnecessary risk.	Being a multinational company, the structure is a bit complex sometimes slowing some processes.
<i>Slow and poor internal processes</i>	sometim; compani; lot; improv; organ; time; system; busi; intern; chang; structur; direct; littl	Internal processes are often lacking, so you need good problem solving skills to get things done. Communication at an executive level needs work. Too many priorities create unnecessary risk.	Sometimes the customers can give you a hard time, but if you handle the situation properly everything will be fine.
<i>Time pressure</i>	time; job; sometim; stress; train; difficult; lot; custom; client; busi; learn; servic; task; hard; littl; environ; help; feel; make; challeng	Not enough payroll to custom service and complete tasks all at the same time. A lot of tasks give at once.	The health insurance benefits are less than impressive. I would hope that as we become a larger company our health insurance coverage would improve as well.
<i>Benefits (healthcare)</i>	benefit; offic; insur; health; locat; expend; compani; employe; park; health insur; home; plan; offer; pay; cost; option; live; medic; commat; match	Healthcare benefits are expensive and not great. No education allowance, but they do provide some in house training options. Office space could use some updates.	Little opportunity for advancement unless you move departments
<i>Lack of opportunity / Career growth</i>	compani; develop; path; littl; slow; career growth; opportun; advance; growth opportun; difficult; hard; career path; progress; grow; advance opportun	The lack of opportunities to develop or progress are limited and often not available once you reach a certain level.	

Table A5
Detailed LDA Topics: Pros Reviews

Pros Topics	Top Ngrams	LBO Example	Other Example
<i>Good pay / Benefits</i>	pay; benefit; decent; pay benefit; decent pay; health; competi; decent benefit; ok; competi pay; benefit pay; peopl; bonus; insur; pay decent; premi; fat; share; profit; health benefit	Not bad pay, decent health benefits, immediate peers are friendly	-Good Pay -Incentive Plans -Good Health Benefits -40k matching up to 6%
<i>Team support / Training</i>	team; projects; support; technolog; develop; client; train; manag; equip; techn; product; process; resource; skill; knowledge; technic; tool; excel; basi; staff	Excellent team and great group of individuals good local management good training available	Develops great relationships with client. Always makes the client first and works to develop relationships with entire project team such as architects, building engineers, and consultants.
<i>Growth opportunities</i>	opportun; compani; growth; career; potent; global; basi; world; intern; grow; travel; lot; develop; path; person; locat; industri; profession; divers; project	The people in this company are great to work for. And the growth opportunities are endless especially if you are able to relocate. This review is directed to you, the job seeker researching Alcanai as a potential employer. My goal is to provide some additional insight, or at least a perspective I believe to be different than what other reviews will help you make the most informed decision on your behalf. I'm sure you're already aware of what I did. And if I decide to join Alcanai, my goal is to help explain to you what you can expect in the year and a half or so that you'll be with us. At the time of t his review, Alcanai hold an uncorrected rating of 2.0 stars here on Glassdoor. [-] Company is growing quickly, which is giving lots of employees meaningful opportunities to grow into bigger roles. The perks are as good as anywhere and management is open and approachable.	Good growth opportunities, challenging global projects, great diverse workforce
<i>Time & MISC</i>	job; time; don; compani; re; ve; press; day; peopl; hire; start; posit; month; manag; call; pay; look; leave; review; hard		Not all that different from other airlines I've worked at (Virgin and Emirates), but you keep your head down and do your job. The perks are good, but I don't care to come with you, enjoy the job itself, it's a good place to be.
<i>Company growth</i>	compani; grow; move; opportun; chang; lot; posit; ve; grow compani; peopl; look; opportun move; move compani; depart; promot; compani grow; career; hard; fast; time	Growing company, treat employees well, management takes an interest in employees and value feedback, pay	There are a lot of opportunities to move up in the company for those that are hard working and capable.
<i>Good community & employee treatment</i>	employe; compani; valu; communiti; improv; focus; treat; provid; cultur; respect; apprec; custom; thi; feel; invol; benefit; engag; encourag; care; environ		They really care about their employees and are always looking to improve their work environment. They value their employees opinions.
<i>Great culture/people</i>	peopl; cultur; environ; fun; smart; lot; amuz; benefit; awesom; smart peopl; team; challeng; friend; compani; talent; collegu; collabor; offic; perk; atmospher	- Great culture - Great and smart people - A lot of autonomy	Great people, fun environment. Lots of good perks.
<i>Fast & meaningful environment</i>	peopl; help; love; fast; pace; day; fast pace; environ; enjoy; job; will; team; fun; challeng; hard; pace environ; make; fast pace environ; will help; feel	Learn so much as it is a fast paced environment. Work with the best people who help grow you and are motivated to help save lives with what Alcanai does.	Great team of people to work within a positive, hardworking yet fun environment. Everyone seems genuinely engaged, excited about what we are doing, and willing / wanting to help everyone succeed and thrive.

Table A6
Detailed LDA Topics: Management Advice Reviews

MngtAdv Topics	Top Ngrams	LBO Example	Other Example
<i>Pay employees what they deserve</i>	people; pay; money; employe; rais; compani; don; leav; job; stop; re; deserv; hard; time; turnov; stay; worth; bonus; live; worker	Don't be cheap with your employees. People work really hard and take pride in their work, and they should be paid accordingly. Also, when something changes or goes wrong, COMMUNICATE about it instead of stonewalling.	Think about the employees instead of your pockets. A lot of good people have come and gone from this company and that could have been avoided if they got a little more incentive to stay.
<i>Listen and care about your employees</i>	people; care; hire; employe; actual; listen; hire peopl; don; care employe; stop; job; instead; listen peopl; listen employe; tri; care peopl; talk; start; take; one pay; employe; attent; increas; pay attent; wage; benefit; rais; level; hour; offer; time; staff; lower; pay employe; bonus; health; insur; rate; paid	Pay your people a real wage and stop with the MMR. When issues are brought to you please actually listen instead of pretending to care. Pay attention to the lower level employees with great potential.	Hire more people in the office as needed instead of loading down current employees with more work.
<i>Pay attention to your employees</i>	people; account; promot; valr; hold; perform; base; leader; lead; leadership; cultur; practic; manag; posit; compani; exampl; divers; team; core; don	-Learn how to be managers and hold your people accountable -Talk to each other	Pay raises should be based on performance. How much the employee has to offer
<i>Be accountable</i>			Keep consistent when holding people accountable. Trust me it shows and people take notice when you don't.

Appendix O. Details on code and implementation of LDA procedure

The below code shows the part of the LDA estimation procedure after having garnered all the reviews in a Dataframe and having estimates optimal parameters α and β . Essentially, the approach is the following:

- Drop non-english reviews.
- Create and clean tokens (i.e., remove corrupt characters, lowercase everything, remove tags, remove punctuation, numbers and non-textual characters, remove extra whitespaces).
- Eliminate articles and stop-words using pre-imported lists.
- Stem all tokens, i.e. set all words in their root form, such that in the document-term matrix, all plural forms, and other declined forms (including certain misspellings) of a word are grouped in a single token (e.g., "thought", "thinking", "thinks", would all get the same root "think"). We rely on the snowball stemmer⁸ for this procedure.
- Create an object with all possible unigrams, bigrams and trigrams.
- Remove n-grams occurring too often, or not enough, to minimize noise.
- Eliminate suspicious n-grams or protect certain terms that appear "too" often but that we desire to keep.
- Create the final corpus and term-document matrix.
- Estimate the LDA with the specified parameters.

```
using JLD2, FileIO, DataFrames, Dates, TextAnalysis,  
    ↪ Languages  
  
#==== helpfcts.jl : The functions below are invoked in  
the main LDA estimation file hereafter. =====#
```

⁸<http://snowball.tartarus.org/>

```

#=# This function takes as input a string of text "x", and
  ↪ outputs all
unigrams, bigrams, and trigrams derived from it, after
  ↪ cleaning
the text, stemming and removing stop words. =#
function my_clean_NGram(x; ngram=3)
  if !isnothing(x) && !ismissing(x) && x!="NA, NA"
    && x!="missing, missing" # Remove missing and
      ↪ corrupted reviews.
    sd = StringDocument(x) #=# Convert review in
      ↪ string format to
      "StringDocument". This step produces tokens under
      ↪ the hood. =#
    remove_corrupt_utf8!(sd) # Remove corrupt
      ↪ characters
    remove_case!(sd) # Set all to lower case
    remove_html_tags!(sd) # Remove tags
    prepare!(sd, strip_non_letters) #=# Remove
      ↪ numerical, punctuation,
      and other non-text characters =#
    sd = strip(sd.text) #=# Make sure to remove extra
      whitespaces around tokens.=#
    prepare!(sd, strip_articles) # Remove articles
    prepare!(sd, strip_stopwords) # Remove stop words
    stem!(sd) # Reduce all terms to root form
    return NGramDocument("$sd", 1,2,3) #=# Create all
      ↪ the
      combinations of unigrams, bigrams, and trigrams
      ↪ =#
  else
    return missing
  end
end

#=# This function returns true if the review is written in
  ↪ English =#
detector = LanguageDetector()
function isEnglish(x)
  l=0
  try
    l, t, c = detector(x)
  catch err
    l = 0
  end
  return l==Languages.English()

```

```

end

#= This function takes a corpus set "U" as input and
↳ outputs the final
filtered corpus from which noise-generating n-grams are
↳ removed. =#
function createCrps(U, mingrams, minoccurUni, minoccurBi,
↳ minoccurTri,
maxoccur; protectedTerms = ["manag", "chang", "sale", "
↳ lack", "hard",
"expect", "team"], termsToRemove = [""])

    maxDocs = round(length(U)*maxoccur)
    crps = Corpus(U)
    update_lexicon!(crps)
    update_inverse_index!(crps)

    for (ngram, cdocs) in crps.inverse_index
        if !(ngram in protectedTerms)
            #= Remove all unigrams that occur too few
            ↳ times,
            those that occur too often, those
            in the manual removal list, and those
            being suspiciously long (>=13 characters) =#
            if length(split(ngram, " "))==1
                if length(cdocs)<minoccurUni || length(
                    ↳ cdocs)>maxDocs
                    || (ngram in termsToRemove) || length(
                        ↳ ngram)>11
                        for doc in cdocs
                            delete!(crps.documents[doc].
                                ↳ ngrams, ngram)
                        end
                    end
                # Removal of bigrams following same logic
                elseif length(split(ngram, " "))==2
                    if length(cdocs)<minoccurBi || (ngram in
                        ↳ termsToRemove)
                        || length(ngram)>18
                            for doc in cdocs
                                delete!(crps.documents[doc].
                                    ↳ ngrams, ngram)
                            end
                        end
                    # Removal of trigrams following same logic

```



```

elseif length(split(ngram, " "))==3
  if length(cdocs)<minoccurTri || (ngram in
    ↪ termsToRemove)
    || length(ngram)>25
      for doc in cdocs
        delete!(crps.documents[doc].
          ↪ ngrams, ngram)
      end
    end
  end
end
end
end
end
en_rows = [length(ngrams(x))>=mingrams for x in crps.
  ↪ documents]
crps.documents = crps.documents[en_rows]
update_lexicon!(crps)
update_inverse_index!(crps)
summ = sort(collect(zip(values(lexicon(crps)),
  keys(lexicon(crps)))), rev=true)
return crps, summ, en_rows
end

function mergeRevs(x,y,z)
  return "$(x) $(y) $(z)"
end

##### mainLDAestim.jl : Estimation of LDA model #####
@load "/home/nicolas/Data/LB0/AllReviews.jld2"

### Manual Parameter settings ###
minoccurUni, minoccurBi, minoccurTri = 300,35,15
pass=100; xgrams=3; maxoccur=0.25
mingrams = 3
 $\alpha$  = 0.1;  $\beta$  = 0.05; minoccur = 50;
protectedTerms = ["chang", "sale", "lack", "hard", "team"
  ↪ ]
nTopicList = [25]

for nTopics in nTopicList
  # Loop over the set of chosen reviews
  for revs in [:Cons] #, :all, :Cons, :Pros, :MngtAdv
    # Merge Cons, Pros, and Mngt Advice if desired.
    if revs==:all
      Y[:, :all] .= ""
    end
  end
end

```

```

        for r in 1:size(Y,1)
            Y[r,:all] = mergeRevs(Y[r,:Cons], Y[r,:
                ↪ Pros],
                Y[r,:MngtAdv])
        end
    end

# Keep english rows
en_rows = @time isEnglish.(Y[!,revs])
X = convert(Array{String}, Y[en_rows, revs])

# Compute clean Ngrams
U = @time my_clean_NGram.(X, ngram=xgrams)

# Create final corpus serving as LDA input
crps, summ, en_rows2 = @time createCrps(U,
    ↪ mingrams, minoccurUni,
minoccurBi, minoccurTri, maxoccur; protectedTerms
    ↪ =protectedTerms)

# Create Document-Term matrix
m = @time DocumentTermMatrix(crps)

# Estimate LDA with specified parameters
 $\Phi$ ,  $\theta$  = @time lda(m, nTopics, pass,  $\alpha$ ,  $\beta$ )
    ↪ alpha$,  $\beta$ )

# Save results
@save "/home/nicolas/Data/LB0/
LDA_FormerJoinerRevs$(revs)$(nTopics).jld2"
en_rows crps  $\Phi$   $\theta$  revs mingrams
    ↪ minoccurUni minoccurBi
minoccurTri pass xgrams maxoccur  $\alpha$   $\beta$ 
    ↪ nTopics m en_rows2
end
end

```