

Detecting Equity Style Information Within Institutional Media



Cédric Gillain, Ashwin Ittoo, and Marie Lambert

Abstract This study examines the detection of information related to small and large equity styles. Using a novel database of magazines targeting institutional investors, the institutional media, we compare the performance of dictionary-based and supervised machine learning algorithms (Naïve Bayes and support vector machine). Our three main findings are (1) restricted word lists are the most efficient approach, (2) bigram term frequency matrices are the best weighting scheme for algorithms, and (3) Naïve Bayes exhibits overfitting while support vector machine delivers encouraging results. Overall, our results provide material to construct small-cap and large-cap coverage indexes from specialized financial media.

Keywords Textual analysis · Machine learning · Style investing

1 Introduction

In their stock allocation process, some investors first allocate their funds between several categories before selecting individual stocks. Those categories are called “styles,” and this investment behavior is called “style investing” (Barberis & Shleifer, 2003). Empirical studies (Froot & Teo, 2008; Kumar, 2009) have documented the importance of investor preference shifts between styles across time (small vs. large, value vs. growth). Change in style preferences provokes important inflow into one style to the detriment of its counterpart. Experts’ recommendations within media are one potential major variable influencing the choice of investors between extreme styles. In this study, we document the most appropriate textual methodology to detect small and large style information within specialized advising magazines.

C. Gillain (✉) · A. Ittoo · M. Lambert (✉)
HEC Liège, Management School of the University of Liège, Liège, Belgium
e-mail: cedric.gillain@uliege.be; Ashwin.Ittoo@uliege.be; Marie.Lambert@uliege.be

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
P. Alphonse et al. (eds.), *Essays on Financial Analytics*, Lecture Notes in Operations Research, https://doi.org/10.1007/978-3-031-29050-3_8

Investors cannot find directly or explicitly information about equity style investing into traditional financial newspapers (*The Wall Street Journal*, Dow Jones Newswires) or social network platforms (Twitter, StockTwits). Information in these media sources is often related at stock level or with a broader market vision. Therefore, we explore a new source of information to address our problem, namely, an institutional media database, which we created. It is composed of nine magazines targeting institutional investors, asset managers, advisers, and financial consultants. These magazines provide information focusing on different asset classes with an emphasis on strategic allocation. This way of relating financial market information is in accordance with style investing foundation, i.e., investors shifting preferences between assets sharing opposite characteristics. To our knowledge, no other publication suggests this type of analysis with a corpus targeting professional investors.

Whether existing textual analysis methodologies in finance are transposable to our problem is worth being addressed. The first and most widespread application is sentiment analysis, where researchers translate a qualitative piece of information into a score indicating polarity conveyed. Many sentiment indexes have been developed to study different asset classes (stock, housing, commodities, etc.) using dedicated source of information (newspapers, social media, company filings, etc.). These sentiment indicators exist with an important variation of methodologies applied in different context. For example, measuring sentiment should be different between a more formal newspaper content and informal social network language. In this study, we summarize major methodologies applied in sentiment analysis. We find that existing literature mainly uses word lists (i.e., dictionaries) or supervised machine learning algorithms. But many existing publications lack comparison between the two main solutions leaving a room for discussion. One comparison within an identical sample was documented by Henry and Leone (2016). They find similar performance between word lists and Naïve Bayes algorithms applied to earnings announcement. We formulate our detection problem as a classification task distinguishing between news with style content (i.e., style news) and news without style content (i.e., neutral news). We document the limits of dictionary-based, Naïve Bayes, and support vector machine classifiers, compare their performance, and provide recommendations on the optimal methodology to detect small and large style content.

We encounter three main challenges to detect style information in institutional media. First, we lack labelled data and have to proceed to manual annotation by ourselves. This task is required to train supervised machine learning algorithms so that they learn from data. In contrast to our difficulty, recent studies on social network now find annotated data available (Oliveira et al., 2016; Renault, 2017). This increasing amount of self-generated data improves algorithm efficiency. Second, style information does not necessarily constitute news headlines but mainly represents secondary information. The way that information is presented increases the difficulty of manual annotation. Third, style content is disseminated across various financial topics such as financial markets, retirement planning, or mutual funds industry-related news. In order to adapt our annotation process, we manually select neutral news and use dictionary detection for style news. This semi-automatic

labelling process is only possible since dictionary-based approach delivers high precision: news classified within style class contains in vast majority style content.

We find promising results from support vector machine algorithms, especially with a term document matrix composed of bigrams. This methodology delivers balanced performance, limiting false detection (neutral news classified as style news) and missed detection (style news classified as neutral news). Naïve Bayes fails to limit false detection and is extremely sensitive to the size of annotated sample. Naïve Bayes is also more susceptible to overfitting, strongly modeling labelled data but weakly generalizing to an extended corpus such as the overall institutional media database. Compared to supervised algorithms, dictionary-based approach delivers the best performance. Restricted unigram and bigram lists with implicit stock content outperform supervised algorithms. The absence of existing labelled data is detrimental to the training phase of such algorithms.

Our research contributes to the literature on textual analysis in finance. We provide a critical analysis of the capacity of current methodologies to detect style information. We compare the use of human word lists with supervised machine learning algorithms. The results of our work provide style coverage indexes with applications focused on investor behavior, fund flows, and the importance of expert advices in the financial press.

The remainder of the paper is organized as follows. Section 2 discusses main methodologies applied in the financial domain. Section 3 introduces the institutional media database. Section 4 describes the different approaches selected to detect style investing-related content and the evaluation of their performance. Section 5 reports the empirical results. Section 6 concludes.

2 Literature Review

In this section, we present a review of existing classification methodologies applied to financial applications. As such, a classification methodology consists of identifying the class of a piece of information (word, sentence, or text), i.e., reducing the content of information to one or several categories. Most publications covered in our work concern sentiment analysis where the information is classified into positive or negative content. Table 1 summarizes publications covered in this review.

The usage and creation of expert dictionaries in financial publications is important. Tetlock (2007) uses the Harvard IV-4 psychosocial dictionary (hereafter H4N) to study the role of newspapers in the stock market. He counts the number of word occurrences from the 77 predetermined H4N categories. He finds that the “negative,” “weak,” “fail,” and “fall” (i.e., pessimism factor) categories mostly influence stock market prices. In their seminal paper, Loughran and McDonald (2011, hereafter LM) develop financial word lists to better capture domain-specific information. They prove that many words from Harvard General Inquirer are misclassified in the financial context. A concrete example lies in the title of their paper, “When is a liability not a liability?”. The word liability is erroneously

Table 1 Literature review of classification methodologies in finance

Approach	Publication	Application	Description of methodology	Sample	
Dictionary-based	Tetlock (2007)	Tone content of stock market in newspapers (sentiment)	Selection of most pertinent lists within Harvard IV-4 psychosocial dictionary with a principal component factor analysis	<i>The Wall Street Journal</i> , “Abreast of the Market” daily news (3709 observations), from 1984 to 1999	
	Loughran and McDonald (2011)	Tone content of company earnings filings (sentiment)	Creation of word lists dedicated to the financial domain (negative and positive, uncertainty, litigious, strong and weak modal)	37,287 MD&A section of 10-K filings (8341 different firms), from 1994 to 2008	
	Larcker and Zakolyukina (2012)	Detection of deceptive narrative in conference calls	Selection of pertinent word lists from LIWC and WordNet	29,663 transcripts for US companies, from September 2003 to May 2007	
	Garcia (2013)	Tone content of stock market in newspapers (sentiment)	Estimation of the fraction of positive and negative words classified by Loughran and McDonald (2011) dictionary	55,307 articles from <i>The New York Times</i> “Financial Markets” and “Topics in Wall Street” columns, from January 1905 to December 2005	
	Bodnaruk et al. (2015)	Firm financial constraint	Creation of a word list referring to financial constraint	51,533 10-K filings (3607 different firms, excluding regulated financial firms and utilities), from 1996 to 2011	
	Baker et al. (2016)	Economic policy uncertainty index	Use of three manually defined word lists (economy, uncertainty, and policy)	Articles in 10 US leadings newspapers, from 1985 to 2012	
	Loughran et al. (2019)	Oil tone index	Creation of related oil word lists with inclusion of modifier words	41,432 Dow Jones Energy Services (DJES) oil-related articles from 2000 to 2016	
	Soo (2018)	Housing sentiment index	Adaptation of four lists (increase, rise, decrease, and fall) from Harvard IV-4 psychosocial dictionary	37,357 articles in major newspapers of 34 US cities, from January 2000 to December 2013	
					(continued)

Table 1 (continued)

Approach	Publication	Application	Description of methodology	Sample
Machine learning	Antweiler and Frank (2004)	Sentiment analysis of message boards	Naïve Bayes trained on 1000 manually annotated messages	1.5 million messages from Yahoo Finance and Raging Bull over 45 firms in the Dow Jones Industrial Average and Dow Jones Internet Commerce Index during 2000
	Das and Chen (2007)	Sentiment analysis of message boards	Five distinct intuitive algorithms trained on 1000 manually annotated messages	145,110 messages from Yahoo Finance over 24 tech stocks between July and August 2001
	Li (2010)	Tone and content of forward-looking statements (FLS)	Naïve Bayes algorithms trained on 30,000 sentences. Multi-class into “tone” and “content” categories	140,000 FLS (13 million sentences) between 1994 and 2007
	Jegadeesh and Wu (2013)	Tone content of company earnings filings	Regression-based approach of tone content on market reaction	45,860 filings from 7606 firms, between 1995 and 2010
	Oliveira et al. (2016)	Sentiment analysis of StockTwits posts	Unigram and bigram classification using co-occurrence measures with a training set representing 75% of overall sample	350,000 messages on StockTwits platform, from June 2010 to March 2013
	Manela and Moreira (2017)	News implied volatility index (NVIX)	Support vector regression trained on a period ranging from 1996 to 2009	Title and abstract of <i>The Wall Street Journal</i> front page, from July 1889 to December 2009
	Renault (2017)	Sentiment analysis of StockTwits posts	Sentiment weight of unigrams and bigrams using 375,000 bearish and 375,000 bullish posts	Nearly 60 million messages from StockTwits platform, from January 2012 to December 2016

This table provides several key information about publications covered in our literature review. We provide a general description of the financial application and the methodology used. We give an overview of sample size and the related period of analysis

considered as negative by common dictionaries while exhibiting no valence in a financial context. Negative and positive word lists created by LM better capture tone content in financial information such as earnings filings. Using the LM dictionary, Garcia (2013) extends Tetlock (2007) findings: both positive and negative tones in media have predictive power on market returns, especially during recession (high uncertainty period). He confirms the importance of selecting an adapted dictionary to the financial domain.

Other word lists have been created subsequent to LM financial dictionary. Larcker and Zakolyukina (2012) analyze conference call transcripts to detect deceptive narrative from CEO and CFO. They manually form several word lists including general knowledge, emotions, hesitations, personal references, and shareholder value. False narratives contain more extreme positive emotion and references to general knowledge but less shareholder value terms. Soo (2018) creates an index for housing sentiment by quantifying tone in different local newspapers in the United States. He adapted four listed categories from Harvard IV-4 psychosocial dictionary: “Increase” and “Rise” (resp. “Decrease” and “Fall”) for positive (resp. negative) sentiment. He manually removes misclassified words and expands lists with all inflections of remaining words. He shows that housing sentiment has higher potential to explain house price variations than a cluster of economic variable (such as rents, employment, or interest rates).

Baker et al. (2016) search terms related to economy (“economic” or “economy”), uncertainty (“uncertain” or “uncertainty”), and policy (“Congress,” “deficit,” “Federal Reserve,” “legislation,” “regulation,” or “White House”) within US newspapers. To define the policy list, authors conduct an audit consisting of 12,009 article readings. They compare human classification with automatic word detection from an enlarged word list (15 terms). They form the policy final list by minimizing coding misclassification (in comparison to manual labelling). The audit process was highly time-consuming, involving the construction of an extensive annotation guide and the participation of a research team. This work leads to the construction of an economic policy uncertainty index allowing further research about policy implications on financial markets. For example, authors show that policy uncertainty affects stock price volatility and reduces investment and employment in certain sectors. Finally, Loughran and McDonald extend their approach to other word list applications. Bodnaruk et al. (2015) study firm financial constraint with a list of related words. Based on their word list, they are able to predict firm liquidity events (such as dividend omission) from financial variables, proving again that textual analysis helps develop new pertinent (and low correlated) variables of interest. Loughran et al. (2019) manually create lists affecting oil prices. Their approach goes beyond the creation of two positive and negative lists. They define a third list of keywords whose effect on oil prices depends on an accompanied modifier word (such as “fall” or “surge”). They find investor overreaction to oil news.

Statistical algorithms are mainly used in absence of a reference dictionary or to alleviate limited size of restricted word lists. Antweiler and Frank (2004) were the first to construct a sentiment index with a Naïve Bayes methodology. Their model correctly classify more than 80% of 1000 messages manually labelled as {Buy, Sell,

Hold}. They show that talk from social media has an effect on trading volume and volatility, opening path for applications in event studies and insider trading. Das and Chen (2007) insist on web talk ambiguity and the difficulty to correctly classify message boards (even for humans). They suggest the combination of five different intuitive algorithms computationally efficient in such a way that they do not require important optimization. They find that a majority voting among algorithms delivers higher sentiment accuracy and reduces false positive.¹ Li (2010) annotated manually 30,000 sentences from 13 million included in forward-looking statements from firms' earnings filings. He trained a multi-class Naïve Bayes algorithms to categorize sentence tone ("positive," "neutral," "negative," and "uncertain") and content ("revenue," "cost," "profits," "operations," "liquidity," "investing," "financing," "litigation," "employees," "regulation," "accounting," "others"). By applying an N-fold cross-validation test, the algorithm correctly predicts sentence classes with an accuracy above 60%. He advocates that in absence of an adapted dictionary, researchers should rely on machine learning classifier instead of using general word lists.

From these pioneering studies, we learn that there is a compromise between sample training size and time spent for annotations. Moreover, existing dictionaries do not adjust to the informal content of Internet messages. But the recent evolution of social networks brings one green light to researchers: self-reported annotations. Oliveira et al. (2016) automatically construct dictionaries dedicated to StockTwits. From 350,000 users' labelled posts, they identify unigrams and bigrams co-occurring with bullish and bearish stance. They estimate co-occurrence with information gain, pointwise mutual information, and TF-IDF measures. Their automated methodology produces high classification rate outperforming classical dictionaries, which are not suitable to microblogging data. Renault (2017) attributes a sentiment weight to unigrams and bigrams by counting their occurrences in bearish and bullish posts from StockTwits. His methodology produces higher out-of-sample accuracy than LM and H4N dictionaries. Exploring nearly 60 million posts, he shows that sentiment is a predictor of intraday stock return, mainly driven by novice traders. These findings provide a direct empirical evidence of noise trading.

All previous studies require subjective interpretation: lists of words are constructed from experts' knowledge, while automatic labelling from microblogs still requires manual deletion or reclassification to improve algorithm efficiency.² Some researchers avoid intentionally any manual classification in their methodology. Jegadeesh and Wu (2013) use a regression-based approach to estimate word positive and negative weightings using market reaction to 10-K filings. They depart from the hypothesis that all words are equally relevant and suggest that market reaction should objectively determine this relative importance. They regress word content of 10-K filings on firms' abnormal returns during the releasing day: word coefficients

¹ Renault (2017) corrected manually misclassification using Oliveira et al. (2016) methodology. He finds, for example, that the words "further" or "commodity" are classified as negative.

² Ibid

representing a positive or negative power. They find that word power weightings differ significantly from proportional and TF-IDF weightings. They proved that the word power approach is robust while using combination of dictionaries and omitting some word from these dictionaries. Their approach gets rid of word classification subjectivity but requires to identify a market reaction to the information studied. Manela and Moreira (2017) apply a support vector regression (hereafter SVR) procedure where unigrams and bigrams from *The Wall Street Journal* were regressed on implied volatility index (VIX). The advantage of SVR is the reduction of term matrix into its most impactful constituents permitting a regression setting (otherwise impossible with OLS). They find that high news implied volatility index (NVIX) predicts above-average stock returns. Using WordNet to classify word in categories, they find that wars and government are driving the variation of equity risk premium. This last example demonstrates that all research tend usually to blend objective and subjective components.

3 Institutional Media Corpus

Several media groups deliver financial information through their portfolio of magazines. Their general mission statement includes production of trusted information targeting financial decision-makers (such as investment managers, advisers, pension trustees, and financial intermediaries). We identified nine different pension funds and institutional magazines from five well-known financial media groups: they form our institutional media corpus. Table 2 presents a description of institutional media and the related magazines. Although it might differ among our magazine sample, most of the readers are institutional investors, asset managers, advisers, and consultants. For instance, the Global Fund Media reach 43,000 active readers, and more than 60% have position in buy-side industry (fund managers, institutional investors, and advisers). The remaining readers occupy various positions in financial services. PLANSPONSOR magazine has 35,000 subscribers from corporate and plan pension-related positions.

The structure of information within our institutional media corpus is similar to newspapers content. News released in institutional media cover a broad range of topics such as macro-economics, market analysis, expert insight, and strategic portfolio allocation. Equity-related information is essentially discussed at category-based level, i.e., emerging market and developed market equities, equity investment style, or industry sectors. This departs from the information provided from conventional newspapers such as *The Wall Street Journal*, *Dow Jones News*, or other leading newspapers. Those media are featuring news on individual firms rather than opinions, recommendations, or performance analysis at a more aggregate level regarding equity style investment. One may have aggregated information from individual companies with regard to equity style portfolio. But two caveats are in order: first, this would assume investors are able to process all this information and, second, style rankings among companies are time-varying.

Table 2 Institutional media corpus

<p>Media group</p> <p><i>Euromoney Institutional Investor</i> PLC (“Euromoney”) is a global, multi-brand information business which provides critical data, price reporting, insight, analysis, and must-attend events to financial services, commodities, telecoms, and legal markets. Euromoney is listed on the London Stock Exchange and is a member of the FTSE 250 share index</p>	<p>Magazines</p> <p><i>Euromoney</i>: <i>Euromoney</i>, founded in 1969 to chart the liberalization of cross-border capital flows, is the leading publisher on the world’s banking and financial markets. Our coverage provides unrivalled insight into the finance houses at the heart of global finance through our privileged access to their senior leaders</p> <p><i>Institutional Investor</i>: For 50 years, <i>Institutional Investor</i> has built its reputation on providing must-have information for the world’s most influential decision-makers in traditional and alternative asset management</p>
<p><i>FTAdvisor</i>: FTAdvisor.com is dedicated to the financial intermediary market covering investments, mortgages, pensions, insurance, regulation, and other key issues. The strength of FTAdvisor.com comes from dedicated up-to-the-minute news articles and in-depth commentary written by the FTAdvisor.com team, combined with the expertise of <i>Financial Advisor</i> and <i>Money Management</i> magazines, whose content feeds directly into the site</p>	<p><i>Financial Advisor</i>: The premier weekly newspaper for the UK’s financial intermediary community, <i>Financial Advisor</i> was launched in 1988 after the Financial Services Act 1986 defined for the first time the role of the independent financial advisor</p> <p><i>Financial Advisor</i> offers comprehensive and in-depth coverage of the retail finance landscape</p>
<p><i>Global Fund Media</i>: Founded in 2002, GFM Ltd. is the most targeted digital news publisher serving institutional investors/wealth managers and their investment managers/advisers across all asset classes with seven daily global newswires and real-time news-driven websites</p>	<p><i>AlphaQ</i>: Compendium of investment ideas, skills, and talent across all asset classes</p> <p><i>Institutional Asset Manager</i>: Institutional investors/pension funds and their managed funds/investment managers</p> <p><i>Wealth Adviser</i>: Private client/wealth managers, family offices, trustees, and their investment advisers</p>

(continued)

Table 2 Institutional media corpus

<p>Media group</p> <p><i>IPE International Publishers Ltd.</i></p>	<p>Magazines</p> <p><i>Investment & Pensions Europe (IPE)</i>: IPE is the leading European publication for institutional investors and those running pension funds. It is published by IPE International Publishers Ltd., an independently owned company founded in July 1996</p>
<p><i>PLANSPONSOR</i>: <i>PLANSPONSOR/PLANADVISER</i>, with its reputation for editorial integrity, objectivity, and leadership, is the trusted information and solutions resource for America's retirement benefits decision-makers. With its powerful array of customer-driven marketing programs, <i>PLANSPONSOR/PLANADVISER</i> offers industry providers an unparalleled ability to reach this influential audience. With all of the changes within the retirement industry, plan sponsors and advisers rely on <i>PLANSPONSOR/PLANADVISER</i> magazine to help them stay informed of crucial issues and important new innovative solutions</p>	<p><i>PLANSPONSOR</i>: Since 1993, <i>PLANSPONSOR</i> has been the nation's leading authority on retirement and benefits programs and has been dedicated to helping employers navigate the complex world of retirement plan design and strategy</p> <p><i>PLANADVISER</i>: Over the past 10 years, retirement plan advisers have reshaped the face of retirement benefits programs, and <i>PLANADVISER</i> has been there every step of the way—providing deep insight into the most pressing retirement plan challenges and strategies facing this specialized group. Our mission, through diverse media channels, is to identify and explore the most critical selling and servicing strategies and tactics facing retirement plan advisers and their clients</p>

This table presents description of pension funds and institutional magazines forming our corpus

We use *Scrapy* library from Python³ to collect news content from institutional media websites. Each news collected is converted in a plain text file with the following information: name of magazine, date, title, author, section, and textual content. We gather 108,638 news from January 1996 to June 2018 from the 9 identified magazines. PLANSPONSOR is the first contributor followed by *Wealth Adviser*, *Euromoney* and *Institutional Investor*. *AlphaQ*, *Institutional Asset Manager*, and *Financial Advisor* provide only recent releases but were nonetheless integrated in the corpus. More than 70% occurrences of news are collected within the period ranging from 2009 to 2017. Table 3 present descriptive statistics of our database.

Figure 1 illustrates a news from institutional media corpus after collection. The title “What to expect from the markets in 2018” does not refer directly to equity styles. The author relates vision of managers for the coming year. While the article may certainly hold strategic content, it is not guaranteed that the author will speak about small or large styles. In this case, they are mentioned only in one sentence. With this example, we highlight that style information is often secondary topic within institutional media. In addition to the absence of existing annotation and the broad diversity of topics within institutional media, the way information is presented makes the detection of style information a challenge.

4 Methodology

We reformulate the detection of style information from institutional media as a classification problem. We intend to classify news into two categories: news containing style-related information (style class) or no style information (neutral class). One news potentially cover any one or two different style information. Since we want to disaggregate the information content by style, we will perform this classification for small- and large-cap styles separately. We investigate two approaches (1) dictionary-based and (2) a selection of machine learning methods. This methodology section will present our methodologies and the evaluation of their performance. Finally, we will discuss how we handle the construction of manually labelled sample to train machine learning algorithms.

4.1 Dictionary Approach

Dictionary-based method detects words (unigrams) and group of two words (bigrams) defining style investment strategies. Dictionary entries for style information include:

³ <https://scrapy.org/>

Table 3 Descriptive statistics of institutional media corpus

	2000–2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018 (*)	Total by magazines
Euromoney Institutional Investor												
<i>Euromoney</i>	5902	887	930	1285	1779	1071	900	748	730	677	365	15,274
<i>Institutional Investor</i>	0	328	736	3997	1278	871	1079	1422	1314	917	0	11,942
FTAdviser												
<i>Financial Advisor</i>	0	0	0	0	0	163	542	1158	2252	2618	1268	8001
Global fund Media												
<i>AlphaQ</i>	0	0	0	0	0	0	0	25	139	129	93	386
<i>Institutional Asset Manager</i>	0	0	0	0	0	0	0	0	722	1858	814	3394
<i>Wealth Adviser</i>	0	100	1583	1416	1448	1654	1789	2259	2133	2152	1102	15,636
IPE International Publishers												
<i>Investment & Pensions Europe</i>	163	118	114	272	239	254	756	853	645	1230	16	4660
PLANSPONSOR/PLANADVISER												
PLANSPONSOR	20,058	3309	3203	3520	2335	2227	1890	1456	1382	1308	653	41,341
PLANADVISER	1251	834	514	881	748	1018	541	441	718	688	370	8004
<i>Total by year (*) until June 2018</i>	27,374	5576	7080	11,371	7827	7258	7497	8362	10,035	11,577	4681	108,638

This table presents the distribution of news across time and magazines within our institutional media corpus. We group first years from 2000 to 2008 in one column (2000–2008). We only collect news until July for 2018

```

magazine : FTAdviser
section : Investments
date : 20180119
title : What to expect from the markets in 2018
author :Simoney Kyriakou
url : https://www.ftadviser.com/investments/2018/01/19/what-to-expect-from-the-markets-in-2018/

Cautious optimism: this seems to be the catchphrase as we enter 2018.

Fund managers have claimed there are many pockets of opportunity globally, whether this is in the rise of disruptive technology or clean energy, Japanese large-caps or domestic small-caps in the US.

But with global political issues on the horizon, such as the Brexit negotiations and ongoing regulation in the financial world, what sort of conversations should advisers be having with their clients?

What might be the potential pitfalls to watch out for this year - and avoid - and where might the new investment opportunities be hiding?

According to Guy Stephens, technical investment director for Rowan Dartington: "When a market correction comes, it will be sudden and savage as real fear returns."

Yet for others, such as Russ Mould of AJ Bell, the danger signs that usually presage a downturn in the markets, such as overly high valuations, a sharp increase in volatility and a steep drop in copper prices, are just not flashing on the investment dashboard right now.

Peter Harrison, group chief executive of Schroders, commented: "The main risk we see lies in deflation, as governments turn to lower taxes and higher infrastructure spending to stimulate economies, which could lead to overheating and unexpected rises in inflation and interest rates.

"Overall, we carry a spirit of cautious optimism into 2018, albeit that caution may start to overwhelm optimism as the year wears on."

The message, according to the contributors to this report, seems to be stay invested but stay alert and stay diversified.
    
```

Fig. 1 Example of news from institutional media corpus

- *Small style*: “microcap(s),” “micro cap(s),” “smallcap(s),” “small cap(s),” “mid-cap(s),” “mid cap(s)”
- *Large style*: “large cap(s),” “mega cap(s)”

The occurrence of each dictionary term is searched for in the news. If one news holds one (or several) terms in the lists above, the news belongs to the referred style class. The news is otherwise classified as neutral. While applying dictionary-based detection, we choose to restrain pre-processing steps to (1) removing {urls, special characters, and numbers}, (2) tokenization, and (3) lowering words. Since dictionary terms are nouns, we search for singular and plural forms.

4.2 Machine Learning Approaches

We apply supervised learning algorithms to detect style investment news. They are specifically adapted to classification problems when the number of classes is well

defined. In supervised learning, the algorithms require a training phase, using a sub-sample of manually labelled data called training sample. Once algorithms are trained, they can perform classification on the overall corpus.

We select two different supervised methodologies for our work: Naïve Bayes (hereafter NB) and support vector machine (hereafter SVM). We choose NB since it is often used as the reference methodology in the literature (Antweiler & Frank, 2004; Das & Chen, 2007; Li, 2010; Henry & Leone, 2016). We complement the research by investigating SVM potentiality. Each method aims to solve the following problem: determine the class i of a news j containing n words. Considering each investment style separately, this is a binary classification problem where a news belongs to the style class or neutral class.

NB approach estimates the probability score of each word to belong to each class. It then predicts the best class based on the probability that words in the news belong to the class:

$$P(\text{class}_i | \text{news}_j) = P(w_1 | \text{class}_i) * P(w_2 | \text{class}_i) * \dots * P(w_n | \text{class}_i)$$

where w_1, w_2, \dots, w_n are words in the news j and class i refers to style or neutral news. This method makes two fundamental assumptions: (1) conditional independence (words occur independently from each other) and (2) positional independence (words have equal probabilities of occurring at all positions).

The SVM approach projects all news j in a space with n dimensions corresponding to each word. In this configuration, news j represents a point in the multi-dimensional word space. The algorithm then creates a boundary to separate the instances of each class i with the following rule: maximize the distance between the boundary and points from each class. SVM is an optimization method that produces the optimal boundary between components of each class. This boundary is constructed using linear or more complex kernel functions. We apply a linear kernel in this study.

We perform more in-depth pre-processing for machine learning algorithms than for the dictionary approach. We restrain vocabulary size to maximize algorithms efficiency. Our pre-processing steps include (1) removing {urls, special characters, and numbers}, (2) tokenization, (3) POS-tagging to select only {noun, verb, and adjective}, (4) lowering, (5) removing stop words, and (6) deleting term with less than three characters. Our pre-processing reduces around 40% of vocabulary size from 24,000 to 15,000 words.

4.3 Performance Assessment

Performance measurements are computed by comparing the actual class of news to the classification provided by the different methods. In our binary classification problem, we get four different results:

- True positive (TP): Style news is correctly classified within style class.
- False positive (FP): Neutral news is wrongly classified within style class.
- True negative (TN): Neutral news is correctly classified within neutral class.
- False negative (FN): Style news is wrongly classified within neutral class.

We estimate precision, recall, and F-score for dictionary-based and machine learning methods:

- Precision: $p = \frac{TP}{TP+FP}$
- Recall: $r = \frac{TP}{TP+FN}$
- Balanced F -score: $F = \frac{2 * p * r}{p + r}$

Precision is the fraction of actual “style” news among news classified within “style” class by the method. Recall is the fraction of “style” news correctly retrieved among all existing style news. For example, suppose there exists 20 “style” news within a sample and the method classifies 10 news within style class. If eight of ten news classified within style class are really “style” news, precision is 0.8 (or 80%). And as the method has correctly retrieved 8 of 20 “style” news but fail to find the 12 remaining ones, recall is 0.4 (40%). We estimate the overall performance of the method with a balanced F -score.

We will promote the method exhibiting the best F -score. To construct the most efficient style coverage indexes, we intend to retrieve almost all style information included in our corpus, i.e., define a tolerable recall threshold. Moreover, we want to be sure that the retrieved information concerns style investing, i.e., we will maximize precision for accepted recall threshold.

4.4 Data Labelling, Training Sample, and Cross-Validation

The main benefit of dictionary-based approach is its high precision, attributed to the fact that terms are manually crafted. However, this method often exhibits poor recall since it is unlikely that a dictionary would contain all possible relevant terms. Machine learning methods, on the other hand, alleviate the issue of low recall since they can learn new instances of a given class, if properly trained. It is well-known that increasing the training data size results in more accurate methods.

We started our labelling task by parsing news titles in search of equity-related information. As such, we try to maximize manual detection of style news: we find more than 100 news referring to each style. We complete our sample with at least 100 neutral news. This initial labelled sample informed us on the performance of dictionary-based approach. We observe that style information represents a minority class within our corpus, i.e., there exists important imbalance between classes where neutral information is dominant. While increasing labelled sample size, we encounter high difficulty to find additional style news. Our manual labelling task became a tedious process.

To overcome this issue, we propose a semi-automatic labelling approach: we use dictionary detections as labelled data in order to avoid searching scarce style information. As previously mentioned, dictionaries exhibit high precision, i.e., detected style news are mainly true positives. Dictionary detections represent adequate candidates for labelled style data. We do not encounter difficulty to manually label neutral information since it represents the dominant class.

We annotate 800 neutral news and select 800 style news from our dictionary detections to form a balanced training sample. We document performance improvement related to the increasing size of training sample from 400 to 1600 news (equally weighted between style and neutral classes). We estimate performance scores of each algorithm with a fivefold cross-validation. We divide the training sample into five equally sized partitions. We use four partitions to train algorithms and estimate performance scores with the remaining fifth partition. We repeat the process five times by alternatively using each partition to assess performance (see Fig. 2). We compute average performance scores from these five estimations.

Finally, we extend our research to different integration of textual content, considering different term weighting estimations. First, we use word frequency as a base case. Second, we train algorithms with term frequency-inverse document frequency estimation. This evaluation adjusts word frequency by its number of observation across news. Such weighting scheme integrates the fact that words are more common than others. Third, we use bigram frequency to explore potential performance improvements as documented in Oliveira et al. (2016) and Renault (2017).

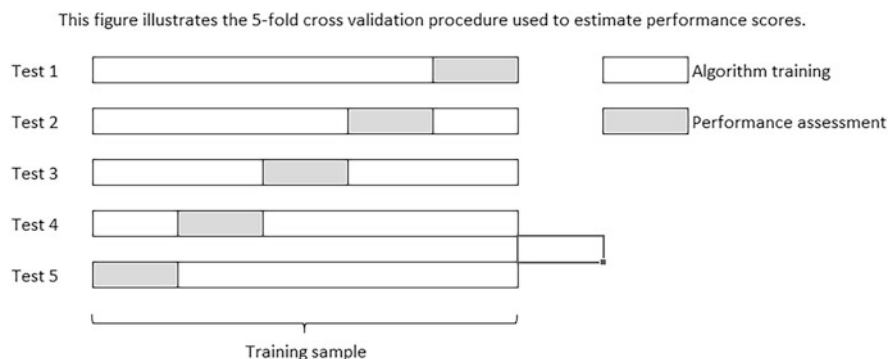


Fig. 2 Fivefold cross-validation. This figure illustrates the five-fold cross validation procedure used to estimate performance scores

5 Results

5.1 Performance of Style Dictionaries

We proceed to a performance assessment of our dictionary approach described in Sect. 4.1. Our results are summarized in Table 4. Small and large dictionaries have, respectively, balanced F-scores of 98 and 97%. They exhibit high precision and recall. Small-cap and large-cap terms defined in our dictionaries are so specific to the investments in style equity that we only find one false detection for each style (FP):

- “The fund invests in high-yield (‘junk’) bonds, foreign securities, emerging markets, liquidity risk, and small- and mid-cap *issuers*” [bond information].
- Non-US equities saw a median return of -4.80% , slightly ahead the Russell Developed ex-US Large Cap Index result of -5.79% [ex term, excluding large caps].

While we do not totally discard false detections, they remain exceptions. Second, small and large dictionaries detect the majority of related news. We find, respectively, four (resp. six) false negatives (FN) for small (resp. large) style. Dictionaries mainly fail to detect the information for two reasons: the presence of the other style in the sentence and the presence of the word “company” replacing the term “cap” from dictionary bigrams.

- “. . . investing in both *small and large-cap* based investment trusts” [fail to detect small information]
- “. . . as well as exposure to international and domestic *large- and mid-cap* equities” [fail to detect large information]
- “. . . dividends paid by *large companies* are more secure than those paid by *small and medium sized companies*, broadly speaking” [fail to detect small and large information]

We do not suggest the integration of terms such as “small company” or “large company” in dictionaries. The informational content is not specific to investments related to stocks and adds many false detections (FP). We do not suggest either adding extended n-gram to detect the opposite style. “Small” and “large” terms have too broad meanings. The power of small-cap and large-cap terms lies in their implicit stock content and their restricted size (unigram or bigram). To be exhaustive, we provide the reader with an Appendix presenting all false-positive and false-negative dictionary detection.

Table 4 shows that small-cap and large-cap terms are so specific to style investments that they are almost present when there is related style information within news. While we construct very restricted lists, our dictionaries are strongly efficient to detect small and large style information. We, respectively, find 2319

Table 4 Performance assessment of dictionary approach

Style	Initial sample			Dictionary classification				Performance measures		
	Total news	Style class	Neutral class	TP	FP	TN	FN	Precision	Recall	F-score
Small	233	121	112	117	1	111	4	0.99	0.98	0.98
Large	233	131	102	125	1	101	6	0.99	0.95	0.97

This table reports the performance scores of dictionary methodology for small and large equity style investments. We estimate performance scores from a manually labelled data sample. We refer the reader to section performance for more detail about dictionary classification and performance measures

small news and 1705 large news within our institutional media corpus of 108,638 news. This represents 1 to 2% of news confirming our previous observation: style information is relatively scarce among institutional media. The main drawback of our approach is the selection and size of our manual labelled sample. In the next section, we test the robustness of our results by investigating if our dictionaries fail to detect style news classified as such by machine learning algorithms.

5.2 Performance of Machine Learning Algorithms

In this section, we focused all tests and results on small style detection. Table 5 summarizes performance of Naïve Bayes and support vector machine approaches. We observe that SVM are more efficient than NB with term frequency weighting: F1 scores reach 84% for SVM compared to 76% for NB. Precision of NB is relatively low (from 62 to 66%): this algorithm tends to wrongly classified neutral news within style class. This is not the case with SVM which exhibits higher precision scores. We find that increasing sample size slightly increases precision of NB term frequency algorithm. Surprisingly, SVM precision is decreasing from 85 to 81% with increasing sample size. We attribute this observation to two possible reasons: the random selection of our training sample and the difficulty to construct a frontier clearly separating the two classes. First, while constructing our training sample, we randomly select 800 news among dictionary detection. This sample differs between each algorithm and therefore impacts their performances. Second, SVM frontier between classes becomes more difficult to construct with increasing number of news present in the multi-dimensional space. If SVM frontier is sensitive to increasing size, we should find similar trend regarding recall scores. However, recall scores seem more stable. We conclude that precision score discrepancy is more related to our random sampling selection.

Considering alternative term document matrices, bigram construction highlights important performance improvements. NB and SVM precision increases strikingly, while recall decreases slightly, leading to higher balanced F1 scores. These results indicate that bigrams add information to NB and SVM models. They are in accordance with observations in Sect. 5.1: bigrams hold more information and are more discriminant than unigrams. Oliveira et al. (2016) confirm our findings with their approach to automatically create sentiment lexicons. They advocate that bigrams present a better sentiment score than unigrams. We do not find the same evidence using term frequency-inverse document frequency weighting algorithms. NB precision, already its weakness, is decreasing to 60% leading to more false positives.

Table 5 Performance of machine learning algorithms

Training settings		In-sample performance scores					Out-of-sample: small news
Method	Sample size	Term document matrix	Precision	Recall	F1 score	Holdout accuracy	
Naïve Bayes	400	TF	0.62	0.89	0.73	0.90	32,094
	800	TF	0.63	0.89	0.73	0.93	28,855
	1600	TF	0.66	0.90	0.76	0.90	24,129
	1600	BF	0.77	0.84	0.80	0.89	22,098
	1600	TF-IDF	0.60	0.95	0.73	0.99	43,356
	400	TF	0.85	0.84	0.84	0.81	15,179
Support vector machine	800	TF	0.84	0.84	0.84	0.81	14,075
	1600	TF	0.81	0.83	0.81	0.85	15,875
	1600	BF	0.94	0.81	0.87	0.80	4569
	1600	TF-IDF	0.85	0.83	0.84	0.85	15,855

This table presents in-sample and out-sample performance assessment of Naïve Bayes and support vector machine approaches. We perform training phase with different sample sizes and term document matrices (*TF* term frequency, *BF* bigram frequency, *TF-IDF* term frequency-inverse document frequency). We compute in-sample performance scores following a fivefold cross-validation. The out-sample measures represent the number of news classified as small within the overall institutional media corpus




Fig. 3 Training procedure for Naïve Bayes and support vector machine algorithms. This figure illustrates training and performance evaluation for NB and SVM algorithms. We train algorithms with balanced class samples selected from our semi-automatic approach. We compute different in-sample performance scores: precision, recall, balanced F1-score and holdout accuracy

We document an additional performance measure, i.e., small class holdout accuracy. This measure represents the fraction of small news detected by the dictionary-based methodology and correctly classified as small by the algorithms. To compute the measure, we ask the algorithm to classify small news that were not used during the training phase. The overall procedure is presented in Fig. 3. Algorithms correctly classify between 80 and 99% of small news detected by our dictionary approach. As such, NB approach is more efficient to classify small news than SVM. As such, NB tends to retrieve more small news (higher recall) but with less precision.

While all measurements discussed indicate the best in sample performing method, they still not guarantee that algorithms will correctly retrieve style news in a sizable sample such as our institutional media corpus. It is possible that trained algorithms model too precisely the training data while not generalizing to new data. This problem is known as overfitting. We investigate out-of-sample performance by asking algorithms to classify all institutional media corpus news and record the number of small news afterwards. We expect this measure to be similar to the number of small news classified by dictionary-based methodology, i.e., 2319 small news detected. We find that NB approach classifies too much news as small in the institutional media corpus (from 22,098 to 43,356). The number of small

news classified is 10–18 times more important than dictionary-based method. This suggests that we have an overfitting problem: term frequency algorithms model too precisely the training sample but fail to generalize to the institutional media corpus. SVM is more immune to overfitting, and it classifies between 4569 and 15,855 news as small depending on the algorithm.

The most promising algorithm is the SVM with bigram weighting scheme: it is more effective in out-of-sample classification than its NB counterpart. SVM trained with bigrams classify 4569 news (NB, respectively, 22,098 news) within a small class in the overall corpus. This algorithm is a potential substitute of our dictionary-based methodology. While dictionary methodology exhibits higher performance, measures are estimated from a smaller sample of manually annotated data (see Table 4. in Sect. 5.1). This manual annotation process may be biased. We may have inadvertently selected small news which contains dictionary terms and fail to select small news that do not contain those terms. In this case, we artificially inflate dictionary performance. We investigate this potential problem by comparing dictionary methodology and SVM using bigram term document matrix. We try to find if dictionary-based methodology fails to capture small news. We randomly read 100 news classified as small by SVM but classified as neutral by dictionary methodology. If the majority of news are really small news, our manual annotation process is biased, and the SVM approach is the best proposed methodology. On the contrary, if those news are mainly neutral ones, we will prefer the dictionary methodology. None of the news we read concern small-cap investments: dictionary-based methodology is still the best approach. We summarize our results in three points. First, best machine learning methodology (bigram SVM) captures 80% of small news detected by dictionary-based methodology (holdout accuracy). Second, news only detected as small by SVM are mainly false positive. Third, performance in dictionary-based section is reliable. We conclude that dictionary-based methodology will produce the most accurate small coverage indexes within institutional media.

6 Conclusions

This paper presents an original classification problem, detecting style content within a novel source of information, the institutional media corpus. We observe three main difficulties in our task: (1) the absence of labelled data, (2) an important class imbalance between neutral and style content, and (3) a diversity of topics within the corpus. This combination of problems makes the manual labelling task difficult for the minority style class.

We find that restricted word lists are the most promising methodology to detect small- and large-cap style information within institutional media corpus. Dictionaries composed of only five different terms (micro-, small-, mid-, large-, and mega-cap) efficiently capture style information. The explicit stock content of these terms avoids false detection (high precision) but also limits failed detection (acceptable recall). In comparison, supervised machine learning algorithms exhibit high recall but low precision scores. This comparison confirms that subjective handed crafted lists are more precise than high recall machine learning algorithms. One drawback of algorithms is their failure to generalize to the overall institutional media corpus. However, support vector machine trained with a bigram term frequency matrix presents encouraging results in alleviating this overfitting problem. One possible improvement for this method is the use of an alternative to linear kernel.

We delimit our research to the comparison of main methodologies applied in sentiment analysis. We do not cover any advanced techniques such as unsupervised learning (clustering, LDA, etc.) or deep learning (neural network), leaving a room for further comparison. Moreover, we do not extend our research to other equity style investments (growth vs. value, momentum vs. contrarian) due to manual labelling constraint. We expect that new source of data will fill the gap in the future, such as it exists user self-reported sentiment on social media. Data augmentation using generative model and semi-supervised learning are also alternative solutions.

Finally, this study opens a path for further research around media coverage and attention related to equity style investments. As such, style detections from dictionaries can be aggregated to construct small-cap and large-cap coverage indexes. With an appropriate fund flow analysis, researchers can investigate the potential impact of style information on institutional investors' behavior.

Appendix: Dictionary Approach: False Positives and False Negatives (Exhaustive List)

Style	Error	ID	Date	Title	Magazine	Sentence
Small	FP	54636	29/10/2012	Prudential launches mutual fund	PLANADVISER	The fund invests in high-yield ("junk") bonds, foreign securities, emerging markets, liquidity risk, and small- and mid-cap issuers.
	FN	51644	21/05/2012	European investors regain a wary confidence	Institutional investor	What is becoming clear to investors is that the best hopes for the next year or so lie in Europe's multinational corporations and smaller-cap companies that know how to build market share abroad, especially in the rapidly growing emerging markets of Asia and Latin America, and to some extent Russia, Eastern Europe and Africa.
	FN	62613	12/5/2014	Europe is moving back into favour	FT adviser	"Given the emerging economic stability appearing more broadly across the continent—including the previously struggling economies of Portugal, Ireland, Italy, Greece and Spain (the 'PIIGS')—investors are increasingly looking to capture a recovery in the region, following a long period of under-allocation to the region, investing in both small and large-cap based investment trusts." ⁷
Large	FN	91234	7/6/2017	Special report: large caps in uncharted waters	FT adviser	Despite the well publicised plight of Pearson, which earlier this year announced its intention to cut its dividend payout after issuing the latest in a string of profit warnings, dividends paid by large companies are more secure than those paid by small and medium sized companies, broadly speaking.
	FP	65724	#####	Equities hurt master trust returns in Q3	PLANSPONSOR	Non-U.S. equities saw a median return of -4.80%, slightly ahead the Russell Developed ex U.S. Large Cap Index result of -5.79%.
	FN	29505	12/5/2009	Turner launches fund with access to long/short strategies	PLANADVISER	Turner Long/Short Equity strategy, directed by Christopher McHugh, lead portfolio manager, investing in stocks in all stock-market sectors, with a focus on mid-sized and large companies globally;

(continued)

(continued)

Style	Error	ID	Date	Title	Magazine	Sentence
	FN	53659	28/09/2012	Third avenue Nabs banker David Resnick	Institutional investor	With banks pulling back, Resnick believes many of the opportunities for Third Avenue will be found among midcap companies that don't have larger companies' ability to refinance.
	FN	76228	20/01/2016	Vanguard makes advisory changes for two active equity funds	Wealth adviser	The large- and mid-cap growth fund has been multi-managed since 1990, and following the transition, the allocation is expected to be as follows: Wellington, 48%; Jennison, 22%; Frontier, 15%; and Vanguard, 15%.
	FN	85585	14/12/2016	The universe of small-cap stocks is shrinking	Institutional investor	In the decades since professors Eugene Fama and Kenneth French published their famous 1992 paper on company size being one of the three sources of equity returns, the edict that small stocks outperform their larger brethren has become conventional wisdom.
	FN	99186	15/03/2018	RBC global asset management strengthens US investment management team	Institutional asset manager	In his previous roles at Rutabaga Capital Management and MFS Investment Management, he had a focus on small- and micro-cap equities, as well as exposure to international and domestic large- and mid-cap equities.
	FN	100636	11/5/2018	Solactive releases media and communications index	Wealth adviser	Starting from the Solactive GBS Developed Markets Large & Mid Cap Index as the founding universe, the index implements a sector-based filtering strategy selecting companies operating in telecommunications, media and publishing services, communication equipment, web-based data, and game software.

References

- Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? The information content of internet stock message boards. *Journal of Finance*, 59(3), 1259–1294. <https://doi.org/10.1111/j.1540-6261.2004.00662.x>
- Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131, 1593–1636.
- Barberis, N., & Shleifer, A. (2003). Style investing. *Journal of Financial Economics*, 68(2), 161–199. [https://doi.org/10.1016/S0304-405X\(03\)00064-3](https://doi.org/10.1016/S0304-405X(03)00064-3)
- Bodnaruk, A., Loughran, T., & McDonald, B. (2015). Using 10-K text to gauge financial constraints. *Journal of Financial and Quantitative Analysis*, 50, 623–646.
- Das, S. R., & Chen, M. Y. (2007). Yahoo! For Amazon: Sentiment extraction from small talk on the web. *Management Science*, 53(9), 1375–1388. <https://doi.org/10.1287/mnsc.1070.0704>
- Froot, K. A., & Teo, M. (2008). Style investing and institutional investors. *Journal of Financial and Quantitative Analysis*, 43, 883–906.
- Garcia, D. (2013). Sentiment during recessions. *Journal of Finance*, 68(3), 1267–1300. <https://doi.org/10.1111/jofi.12027>
- Henry, E., & Leone, A. J. (2016). Measuring qualitative information in capital markets research: Comparison of alternative methodologies to measure disclosure tone, 91(1), 153–178. <https://doi.org/10.2308/accr-51161>
- Jegadeesh, N., & Wu, D. (2013). Word power: A new approach for content analysis. *Journal of Financial Economics*, 110, 712–729.
- Kumar, A. (2009). Dynamic style preferences of individual investors and stock returns. *Journal of Financial and Quantitative Analysis*, 44(03), 607. <https://doi.org/10.1017/S0022109009990020>
- Larcker, D. F., & Zakolyukina, A. A. (2012). Detecting deceptive discussions in conference calls. *Journal of Accounting Research*, 50, 495–540.
- Li, F. (2010). The information content of forward-looking statements in corporate filings – A naïve Bayesian machine learning approach. *Journal of Accounting Research*, 48(5), 1049–1102.
- Loughran, T. I. M., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks, 66(1), 35–65.
- Loughran, T., McDonald, B., & Pragidis, I. (2019). Assimilation of oil news into prices. *International Review of Financial Analysis*, 63, 105–118.
- Manela, A., & Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, 123, 137–162.
- Oliveira, N., Cortez, P., & Areal, N. (2016). Stock market sentiment lexicon acquisition using microblogging data and statistical measures. *Decision Support Systems*, 85, 62–73. <https://doi.org/10.1016/j.dss.2016.02.013>
- Renault, T. (2017). Intraday online investor sentiment and return patterns in the U.S. stock market. *Journal of Banking and Finance*, 84, 25–40. <https://doi.org/10.1016/j.jbankfin.2017.07.002>
- Soo, C. K. (2018). Quantifying sentiment with news media across local housing markets. *Review of Financial Studies*, 31, 3689–3719.
- Tetlock, P. C. (2007). *Giving content to investor sentiment: The role of media in the stock market* published by: Wiley for the American Finance Association Stable URL: <http://www.jstor.org/stable/4622297>. The role of media in the stock market giving content to investor sentiment, 62(3), 1139–1168.

Ashwin Ittoo is Full Professor at HEC Liège, Management School of the University of Liège. Marie Lambert is Full Professor at HEC Liège, Management School of the University of Liège where she holds the Deloitte Chair in Sustainable Finance.

Marie Lambert is a fellow of Quantitative Management Research Initiative (QMInitiative.org).