Norwegian School of Economics
Bergen, Spring 2020

Norwegian Mutual Fund Newsletters

Newsletter Information and its Relation to Performance,
Fund Manager Efforts, and Investor Flows

Charlotte Leikanger Baade

Amalie Dimmen Øvrelid

Supervisors: André Wattø Sjuve & Andreas Ørpetveit

Master thesis, Economics and Business Administration

Major: Financial Economics

NORWEGIAN SCHOOL OF ECONOMICS

Acknowledgements

This thesis marks the conclusion of our Master of Science in Economics and Business Administration at the Norwegian School of Economics (NHH).

We would like to extend our gratitude to our supervisors, André Wattø Sjuve and Andreas Ørpetveit, for their availability, valuable input, and feedback on our work. Moreover, we would like to thank them for access to data and advice on the choice of a feasible topic.

Norwegian School of Economics

Bergen, June 2020

Charlotte Leikanger Baade

Amalie Dimmen Øvrelid

Abstract

By analysing the sentiment and similarity of monthly Norwegian mutual fund newsletters, we find that more positive sentiment and higher similarity is associated with higher performance in the relevant month. We find no proof of an overall connection between newsletter characteristics and fund manager efforts, measured by active share. Further, by constructing portfolios based on the most recent newsletters available to investors, we find that a high-sentiment portfolio does not outperform a low-sentiment portfolio. This suggests that the rationale behind distributing newsletters is in part marketing. Investors do however not respond to sentiment-related information, indicating a lack of successfulness for newsletters as a marketing tool. Lastly, a "non-changer" portfolio outperforms a "changer" portfolio for two out of three similarity measures, and cannot solely be explained by common risk factors. As investors do not respond to newsletter changes, this signals the existence of a potential market inefficiency.

Keywords – mutual funds, active management, investor flows, textual analysis, sentiment analysis, document similarity

Contents

List of Figures

List of Tables

1 Introduction

Why do mutual funds distribute monthly newsletters? Are newsletters shared to keep investors informed about fund performance, or are they mainly intended as marketing? Under the assumption of profit maximisation, mutual funds have two options: increase the expense ratio, or increase fund inflows. Because of price competition, an increase in costs might lead to investors leaving the fund. Thus, the most viable option is maximising net flows. This is primarily done through increased efforts, leading to better performance, or alternatively, through advertising. Many papers have studied the performance aspect of mutual funds. Literature on increased net flows as a result of advertising, on the other hand, is more limited. On this topic, Gallaher et al. (2006) find that advertising is one of multiple strategic decisions that significantly affect flows. In this paper, we use textual data analysis to examine whether newsletter information is reflected in actual performance, or if this is just marketing to attract investors.

The way we perceive textual information plays a role in the assessment of value in financial markets, and automated processes are able to more quickly analyse relevant text and detect patterns. This paper contributes to the field of textual analysis by analysing a relatively unexplored source of information: mutual fund newsletters. Using Norwegian data, we also contribute to the limited Norwegian-based literature. To illustrate the relevance of textual analysis, see table 1.1. This table includes the funds with the highest and lowest sentiment – or "tone" – in our August 2016 newsletter sample.

Table 1.1: Sentiment in August 2016 Newsletters

	Highest Sentiment		
	Sentiment(t)	Alpha(t)	Flow(t+1)
Pareto Aksje Norge	1.76	-1.00%	2.21%
Delphi Norge	1.55	-0.56%	5.03%
FORTE Norge	1.22	-0.96%	9.22%
	Lowest Sentiment		
	Sentiment(t)	Alpha(t)	Flow(t+1)
Holberg Norge	0.18	0.53%	1.00%
C WorldWide Norge	0.06	-0.49%	0.51%
Alfred Berg Gambak	0.00	0.00%	8.44%

This table reports data from the August 2016 mutual fund newsletter sample. It includes the three mutual funds with the highest and lowest sentiment scores in the associated newsletters. Gross alphas (benchmark-adjusted returns) for August 2016 as well as net flows for September 2016 are also reported.

Note that the three most positive funds underperform relative to their benchmarks in August. Meanwhile, only one out of three low-sentiment funds underperform. Moreover, the high-sentiment funds have higher average inflows in the following month, although their performance was worse recently. This may suggest that fund managers communicate a more positive tone in their newsletters when performance is weak in order to avoid investor outflows in the near future.

However, if this is the case, one may question why Alfred Berg Gambak has the lowest sentiment in August, and one of the highest inflows in September. This particular fund has an overall neutral tone in their newsletters over time, and reports positive and negative contributors in an objective manner. For this reason, the content of fund newsletters may not matter as much for their investors. In addition to this, it is worth noting that even the low-sentiment funds in August 2016 are classified as positive or neutral[1].

What exactly do positive funds write about, compared to less positive funds? Table 1.2 illustrates the difference, using extracts from high-sentiment and low-sentiment newsletters. We note that the highly positive fund, which underperformed relative to its benchmark, focuses on future expected performance. This further strengthens the hypothesis that some newsletters primarily function as marketing. On the contrary, the underperforming low-sentiment fund seems to more objectively summarise recent happenings.

Table 1.2: Extracts from August 2016 Newsletters

FORTE Norge	**C WorldWide Norge**
The company (BW Offshore) has a long contract backlog, and the financial risk is now minimised until 2020. However, there is still a non-negligible technical risk in the large new-building of the production vessel Catcher, which will be operating in the English sector next year. If it goes smoothly, the stock is a clear doubling candidate from today's level. [...] There are several indicators signalling that the Norwegian economy has now bottomed out and that GDP growth will pick up in line with rising oil prices. The interest rate level will still be record low and drive more and more capital into the stock market in search for better returns than in the fixed income market. This means that the Oslo Stock Exchange can become a very nice place to be in the time ahead. We strive to continue to provide excess returns.	The biggest positive surprises in our portfolio, compared to consensus estimates, were Aker Solutions, Subsea7, Schibsted, Borregaard, Det Norske Oljeselskap and Orkla. It should be mentioned that our expectations for several of these companies were higher than market expectations, and that we had positioned ourselves according to this. [...] What affected the portfolio negatively in July was the sharp fall in oil prices. Brent oil fell from USD50 per barrel at the beginning of the month to USD42 per barrel at the end of the month, down 16%. The portfolio is directly exposed to the oil price fall through our investments in the oil companies Statoil and Det Norske Oljeselskap, as well as the oil service companies Subsea7 and Aker Solutions. The fall in oil prices has first and foremost been strong on the spot price of oil, while for longer delivery periods the fall has been more moderate.

This table includes extracts from two mutual funds' August 2016 newsletters. The extracts have been translated to English. For the original text in Norwegian, see appendix A1.

[1] A score around zero indicates neutral sentiment, and a higher score indicates positive sentiment.

The newsletter information we examine includes sentiment as well as newsletter similarity. These two characteristics are analysed in relation to fund performance, fund manager efforts, and investor behaviour. Further, we test whether high-sentiment funds outperform low-sentiment funds. Lastly, inspired by Cohen et al. (2020)'s work on American firm disclosures, we test whether high-similarity funds outperform low-similarity funds. High-performing funds may put less energy in writing and updating their newsletters, letting the results "speak for themselves".

Our findings suggest that both increased newsletter sentiment and increased similarity is related to increased performance in the relevant month. However, only recent high-performers seem to sufficiently adjust their newsletter sentiment. We find no overall connection between sentiment or similarity and fund manager efforts. The exception is funds with low active share, where more positive sentiment is related to higher fund manager efforts. Further, we find that sentiment and similarity information is not reflected in investor flows, perhaps in part because of difficulties in accessing newsletters. Moreover, we find that a portfolio of high-sentiment funds does not outperform a portfolio of low-sentiment funds, strengthening the advertising hypothesis. With regard to similarity, on the other hand, we find signs of a market inefficiency, as a "non-changer" portfolio outperforms a "changer" portfolio, and this outperformance is robust to common risk factors.

The remainder of this paper is structured as follows. Section 2 reviews related literature and defines our research questions. Section 3 presents the data sample and textual analysis procedure. Section 4 presents the methodology and gives the main results. Section 5 concludes.

2 Background

In this section, we review related literature and present our four research questions. There are three main areas that this section explores: textual analysis applied to mutual funds, fund manager efforts, and investor behaviour in mutual funds.

2.1 Textual Analysis in Accounting and Finance

Textual analysis in accounting and finance is an emerging field, and Loughran and McDonald (2016)'s survey provides an extensive review of related literature showing that stock market investors incorporate more than just quantitative data in their valuations.

Sentiment analysis, a central topic of this paper, has been performed on a wide range of text in accounting and finance, including annual (10-K) and quarterly (10-Q) reports, earnings press releases, analyst reports, IPO prospectuses, and perhaps most extensively on newspaper articles. With regard to mutual funds, Solomon et al. (2014) study how media coverage of mutual fund holdings affect investors' allocation of money across funds. They find that investors chase funds with high past returns only if the funds' holdings were recently featured in the media.

Hillert et al. (2016) study the shareholder letter section of annual (N-CSR) and semi-annual (N-CSRS) shareholder reports, and find that mutual fund investors react to the writing style of shareholder letters, in which a less negative tone and a more personal writing style lead to higher net flows. They also find that writing style predicts changes in fund managers' risk taking and investment styles, and that personal writing styles predict better fund performance.

Another central topic of this paper is document similarity. In this area, our paper is perhaps most closely related to Cohen et al. (2020). They show that changes to the language and construction of financial reports have strong implications for firms' future returns and operations, in which a portfolio that shorts "changers" and buys "non-changers" earns a significant alpha. Furthermore, Cohen et al. (2020) find no announcement effect, suggesting investor inattention to subtle changes in financial reports.

In this paper, we analyse monthly Norwegian mutual fund newsletters, which differ from the N-CSR and N-CSRS forms filed with the SEC. The newsletter sample that we study has yet to be analysed. In general, previous literature analysing mutual fund text is limited. This is also the case for Norwegian text, especially in the areas of accounting and finance.

We start by measuring the informativeness of newsletters. This is done by studying whether newsletter sentiment or similarity corresponds to the relevant month's mutual fund performance. For newsletters to be informative, there should be a connection between sentiment and/or changes to newsletters and actual performance. In relation to sentiment, the tone should be more positive when performance is high, and more negative when performance is low. For similarity, we hypothesise that mutual funds with higher performance change their newsletters less from month to month. Thus, we test for a relationship similar to the findings of Cohen et al. (2020) on firm disclosures. A lack of connection between performance and textual features may signal that newsletters primarily function as a marketing tool to attract investors. This can for instance be reflected through an overly positive newsletter tone. We present our first research question:

- **Are mutual fund newsletters informative?**

2.2 Active Management and Fund Manager Efforts

The data sample that we present in section 3 consists of actively managed mutual funds. The conventional wisdom of active management shows that the average fund underperforms after fees, that the performance of the best funds does not persist, and that some fund managers are skilled, but rarely in excess of costs (Jensen, 1968; Sharpe, 1991; Carhart, 1997). However, more recent literature has challenged this view. For instance, Cremers and Petajisto (2009) find that funds with high active share[2] tend to outperform their benchmark index. They also show that the performance of funds with low active share drives previous results, explaining why the average mutual fund underperforms.

Jin et al. (2015) study overconfidence among mutual fund managers using active share, and find that fund managers tend to boost their confidence following superior past performance.

[2]Active share is the sum of absolute deviations from the benchmark index.

In this paper, we study whether this increased confidence is partly due to fund managers' *perception* of superior past performance, and not only due to an objective criteria such as observed returns. This is done by analysing the sentiment of text written by fund managers.

We expect a positive relationship between fund managers' tone in newsletters and active share, linked to previous literature on fund manager overconfidence: if fund managers have been positive in recent newsletters, this might signal that they perceive themselves as skilled, and thus increasingly deviate from the benchmark index in the following month. Further, we also test for potential relationships between newsletter similarity and fund manager efforts. For instance, active fund managers may also be active writers who make larger changes to their newsletters. We present our second research question:

- **Is newsletter sentiment or similarity connected to fund manager efforts?**

2.3 Investor Inattention and Behaviour in Mutual Funds

Sirri and Tufano (1998) study flows[3] into and out of mutual funds, and find that investors purchase disproportionately more in funds with superior past performance. Jin et al. (2015) find an irrational investor reaction to fund manager overconfidence in the form of higher inflows as a reward for good performance, but little penalty for poor performance.

While multiple papers study the effect of performance on flows, literature on the marketing aspect of mutual funds is more limited. Gallaher et al. (2006) examine the effect of mutual fund families' strategic decisions on investor flows. They find that beyond performance, decisions such as advertising have significant effects on flows. We contribute to this literature by testing mutual fund newsletters' success as a marketing tool. This is done by analysing how newsletter sentiment or similarity affects mutual fund inflows and outflows. If newsletters are successful as a marketing tool, or if newsletters convey useful information, we expect investors to respond to newsletter information upon release. We present our third research question:

[3]Net flow is the net growth in mutual fund assets beyond reinvested dividends.

- **Do investors respond to the information conveyed in mutual fund newsletters?**

If investors do *not* respond to useful information, or do respond to useless information, some sort of market inefficiency may exist. Our fourth and final research question tests for potential inefficiencies. As previously mentioned, Cohen et al. (2020) find proof of investor inattention on the release of firm disclosures, as investors are missing subtle changes in 10-Ks that predict large negative returns in the future.

We construct portfolios based on sentiment and similarity scores. If mutual fund newsletters were entirely objective, the funds with the highest sentiment in their newsletters should also have high performance. Thus, we test whether highly positive mutual funds obtain significant alphas in the following month, and whether they outperform a portfolio of the least positive funds. Furthermore, inspired by Cohen et al. (2020), we test whether a portfolio of "non-changer" funds obtain significant alphas in the following month, and whether they outperform "changer" funds. Thus, our fourth and final research question is presented as follows:

- **Do mutual funds with high newsletter sentiment or similarity outperform mutual funds with low sentiment or similarity?**

3 Data

This section describes the data that our analyses are based on. We start by describing the mutual fund newsletters and associated data in subsections 3.1 and 3.2, respectively. Next, we explain the textual analysis procedure in subsection 3.3. Finally, subsection 3.4 reports summary statistics on sentiment and similarity measures.

3.1 Mutual Fund Newsletters

We obtain mutual fund newsletters for 15 equity funds in the Norwegian market, with a total of 1,149 observations. The newsletters are obtained from our thesis supervisors, and in some cases collected by ourselves. Table 3.1 summarises the fund newsletters, including the number of newsletters in total, the first and last newsletter observation, as well as the number of newsletters missing within the specified time frame. The newsletters in our sample are sent out to investors on a monthly basis. Mutual fund newsletters vary in scope and focus, but typically include a comment or update from the fund manager, information about the fund and its mandate, and in some cases an update on the general market situation. This paper analyses the fund manager comment section of the newsletters.

The limited sample is due to difficulties in collecting newsletters. Ideally, we would perform an analysis of the entire market of Norwegian equity funds. Not all mutual funds are willing to share their history of newsletters, leading to a somewhat hand-picked dataset. This can cause problems if the mutual funds in our sample have special traits. For instance, if mutual funds are not willing to share their newsletters because they did not perform well enough historically, this can lead to sampling bias as we might be left with “high-performers”. In addition to this, our sample includes multiple mutual funds within the same fund families. These funds are likely to have more in common, which might make our sample less representative for the overall market.

Table 3.1: Mutual Fund Newsletter Summary Statistics

	Obs.	First newsletter	Last newsletter	Missing
Alfred Berg Aktiv	43	2014-09	2018-05	2
Alfred Berg Gambak	43	2014-09	2018-05	2
Alfred Berg Norge Classic	43	2014-09	2018-05	2
Arctic Norwegian Equities	68	2012-10	2018-05	0
C WorldWide Norge	144	2006-01	2018-05	4
Delphi Norge	142	2006-08	2018-05	0
Delphi Vekst	86	2006-08	2013-09	0
DNB Norge Selektiv (III)	68	2010-06	2016-12	11
DNB SMB	77	2010-06	2016-12	2
Fondsfinans Norge	96	2010-01	2018-05	5
FORTE Norge	76	2012-01	2018-05	1
FORTE Trønder	63	2013-02	2018-05	1
Holberg Norge	62	2013-01	2018-05	3
Pareto Aksje Norge A	101	2010-01	2018-05	0
Pareto Investment Fund A	37	2015-05	2018-05	0

This table displays the number of newsletter observations for each mutual fund, year and month of the first and last newsletter, as well as the number of newsletters missing within the specified time frame.

3.2 Mutual Fund Data

Monthly data on Norwegian equity mutual funds is obtained from our thesis supervisors, including alphas, fund returns, benchmark returns, expense ratios, active share, and assets under management (AUM). The alphas included in our dataset are benchmark-adjusted returns.

In order to analyse investor behaviour later in this paper, we calculate net flows according to Sirri and Tufano (1998),

$$FLOW_{i,t} = \frac{AUM_{i,t} - AUM_{i,t-1} \cdot (1 + R_{i,t})}{AUM_{i,t-1}} \tag{3.1}$$

where $AUM_{i,t}$ is mutual fund i's AUM in month t, and $R_{i,t}$ is the fund's return over the past month.

Table 3.2 summarises the 15 funds' respective annualised alphas, returns, expense ratios, active share, AUM, and net flows. Note that the measures are based on varying time frames, as reported in table 3.1. Due to this, special events like the financial crisis of 2007-08 and the 2014 oil price collapse could affect some of the mutual funds' average returns. We should therefore use caution when comparing the estimates. The high averages of certain funds could for example be explained by missing data from the financial crisis.

Table 3.2: Mutual Fund Data Summary Statistics

	Alpha	Return	ER	AS	AUM	Flow
Alfred Berg Aktiv	7.02	18.68	0.13	44.20	929	53.31
Alfred Berg Gambak	8.83	20.81	0.15	56.48	2,585	44.16
Alfred Berg Norge Classic	4.10	15.31	0.10	31.66	1,218	6.49
Arctic Norwegian Equities	2.42	15.96	0.13	38.18	2,198	28.93
C WorldWide Norge	2.08	12.35	0.10	28.75	460	-10.22
Delphi Norge	4.25	14.08	0.17	54.73	742	1.27
Delphi Vekst	2.43	9.71	0.22	69.71	121	-11.63
DNB Norge Selektiv (III)	0.77	12.92	0.07	32.23	3,867	-2.76
DNB SMB	7.50	11.94	0.16	51.22	1,061	-7.25
Fondsfinans Norge	2.38	14.45	0.08	57.29	1,487	1.63
FORTE Norge	3.35	18.84	0.17	62.38	63	56.42
FORTE Trønder	7.70	21.87	0.17	87.62	112	76.77
Holberg Norge	3.14	17.14	0.13	69.74	845	6.77
Pareto Aksje Norge A	-0.89	10.81	0.13	64.56	7,068	-16.27
Pareto Investment Fund A	5.47	17.46	0.15	81.39	1,662	48.72

This table reports average gross alphas, gross returns, expense ratios (ER), active share (AS), assets under management (AUM), and net flows for the mutual funds. The measures are based on the relevant funds' time frame reported in table 3.1. Alphas, returns, and flows are annualised. Alphas, returns, ER, AS, and flow are reported in percent. AUM is reported in NOK million.

3.3 Textual Data

3.3.1 Pre-Processing

We are interested in fund manager comments and updates on our relevant mutual funds. The length and focus of the newsletters vary greatly across funds, and they typically do not keep the same newsletter layout throughout the relevant time frame. This complicates the task of extracting relevant information. We automate the process through R programming. Fund newsletters are obtained in pdf or doc/docx format. We read the newsletters in R, extract relevant information based on specific characteristics such as pages or subsections, and save this information to new text files. We base further analysis on text files as they require limited space and only contain relevant text. The pdf/docx files include text with different formatting, in addition to tables, graphs and figures, which for our purpose – performing textual analysis – is considered noise.

Next, we perform the following pre-processing steps:

- Remove punctuation, numbers, and line separators
- Lower-case all words because of case sensitivity
- Change special characters that are not recognised (*æ*, *ø*, and *å* to *ae*, *oe* and *aa*)
- Remove stop words[4] that do not affect sentiment classification

The choice of pre-processing steps can affect the results, and sentiment analysis is particularly sensitive to stop word removal. One example is the sentence "he is not happy". Without stop word removal, this sentence gives negative sentiment, because of the negator *not* which alters the sentiment of *happy* from positive to negative. When we remove stop words, this sentence turns into "happy", because both *he*, *is* and *not* are considered stop words. This leads to positive sentiment, and the underlying meaning of the sentence is no longer correct. Thus, stop words should only be removed if they do not add any new information (Manning et al., 2008). We solve this by creating a new list of stop words, in which words that can alter the underlying meaning of a sentence are excluded.

3.3.2 Sentiment Lexicon

Sentiment resources have been limited for the Norwegian language. We use the first version of the SANT[5] project's Norwegian sentiment lexicon published in October 2019 (Barnes et al., 2019). They machine translate and manually correct the English lexicon by Hu and Liu (2004), resulting in a comprehensive lexicon consisting of 601 positive and 3,917 negative words. We use their full-form lexicon with a total of 6,103 positive and 14,839 negative words. Neutral words are omitted as they do not change the sentiment in either direction.

[4]Stop words are common words that generally do not add value to the analysis.

[5]The SANT project (Sentiment Analysis for Norwegian Text) is a collaboration between the Language Technology Group (LTG) at the University of Oslo's Department of Informatics and the three media outlets NRK/P3, Schibsted Media Group, and Aller Media.

We give all positive words +2 points, and all negative words -2 points. A more nuanced scoring system would give more precise results. However, given the available resources at the time of writing, the scope of the thesis, and time limitations, we have chosen not to go into more detail in this area. A more nuanced scoring system created by us would also be exposed to our subjective opinions, which could potentially bias the results.

Loughran and McDonald (2011) criticise the increasing use of “off-the-shelf” lexicons to measure sentiment in American annual reports. Lexicons developed for other fields of study generally misclassify words commonly used in financial text. To avoid this, we adjust the sentiment lexicon. For instance, we find *value* to be classified as positive, and *cost* to be classified as negative. In financial context, these are considered neutral terms.

Multiple words are initially included in both the positive and negative list, perhaps because these words can be classified as positive or negative depending on context. We go through all duplicates and manually decide where they fit in financial context. We remove words that can easily be misclassified.

As a final step in adjusting the lexicon, we browse through (1) the 300 most commonly used words in the newsletter sample, and (2) all positive and negative words used at least 10 times in the newsletter sample. This allows us to check if relevant words are included, or if they have been misclassified. Although we do not have the capacity to check how all words are categorised, this step assures a correct treatment of frequently used words.

Keep in mind that only words specified in the lexicon are assessed. Abbreviations, slang, spelling errors and sarcasm is ignored in lexical-based sentiment analysis.

See appendix A2 for a list of common words in our newsletter sample, as well as a summary of the changes we make to the lexicon by Barnes et al. (2019).

3.3.3 Valence Shifters

Valence shifters are words that affect the meaning of another word. Common valence shifters are amplifiers, de-amplifiers, and negators. Amplifiers are strengthening, while de-amplifiers are weakening. Negators infer the opposite polarity. An example of how valence shifters can change sentiment is comparing the sentence “it has been a *very good* month” to “it has *not* been a *good* month”. *Very* is an amplifier, while *not* is a negator.

Without valence shifters, only *good* is evaluated, resulting in identical sentiment for the two sentences.

We use the English standard list from the *sentimentr*[6] package in R as a starting point to develop a list of Norwegian valence shifters. We implement amplifiers, de-amplifiers, and negators. The sentiment function we apply does not allow any valence shifters to also be included in the sentiment lexicon. In cases where we find overlapping words, we manually classify where we best see them fit. See appendix A2 for our list of valence shifters.

3.3.4 Sentiment Score

We use the *sentimentr* package to assess sentiment. See appendix A2.4 for an in-depth explanation of the sentiment function. The function searches for polarised words included in the lexicon. These are then tagged as either positive or negative. A number of words before and after a polarised word form context clusters, and are classified as either neutral, amplifiers, de-amplifiers or negators. Neutral words hold no value, but affect the word count.

The polarised words are first weighted depending on their score in the sentiment lexicon, and then weighted by the valence shifters surrounding them. Negators flip the sentiment sign of the polarised word as long as the context cluster contains an odd number of negators. The odd number rule is added to account for possible double negatives. Amplifiers increase polarity, as long as there is no negators or an odd number of negators. If there is an odd number of negators, amplifiers become de-amplifiers (e.g., "it has *not* been a *very good* month"). De-amplifiers decrease the polarity score, but has a lower bound to avoid turning sentiment negative.

Finally, the weighted context clusters yield the unbounded polarity score for the respective mutual fund newsletter.

[6]https://cran.r-project.org/web/packages/sentimentr/

3.3.5 Similarity Score

We quantify the similarity of mutual funds' month-to-month newsletters by using three lexical-based measures: Cosine similarity, Jaccard similarity, and Levenshtein distance.

Cosine similarity measures the angle between two document term vectors[7] on a unit sphere (Hanley and Hoberg, 2010), and is defined as the dot product of two document term vectors normalised by their vector lengths,

$$C(A,B) = \frac{A \cdot B}{||A|| \cdot ||B||} \tag{3.2}$$

where A and B are the document term vectors of newsletters A and B, respectively. $||A||$ and $||B||$ is the Euclidean norm. The output is a measure in the range $[0,1]$, where 0 represents no common terms, and 1 represents identical newsletters. In other words, a higher measure signals greater similarity between two consecutive newsletters.

Jaccard similarity is defined as the size of the intersection divided by the size of the union of two term frequency sets A and B (Niwattanakul et al., 2013).

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|} \tag{3.3}$$

The output is a measure in the range $[0,1]$, where 0 represents no common terms, and 1 represents identical newsletters. As Jaccard similarity is binary, each word is counted only once in a newsletter. This differs from cosine similarity, which also includes the frequency of words.

Levenshtein distance is the minimum number of word-level insertions, deletions and substitutions required to transform one newsletter into the next newsletter (Boytsov, 2011). Because a low distance infers high similarity, the interpretation is the opposite to that of cosine and Jaccard similarity. In order to make the Levenshtein distance more comparable across funds, we normalise by dividing this measure by the maximum possible distance between two newsletters of given word counts. The normalisation also takes into account the interpretation relative to cosine and Jaccard similarity, putting all three

[7]Document term vectors describe the frequency of terms that occur in a newsletter.

similarity measures on a range from low to high.

This can be expressed as follows

$$1 - \frac{L(A, B)}{\max(|A|, |B|)} \tag{3.4}$$

where $L(A, B)$ is the Levenshtein distance from newsletter A to B, $|A|$ is the length (word count) of newsletter A, and $|B|$ is the length of newsletter B. This results in a Levenshtein distance in the range $[0, 1]$, where 0 represents completely different newsletters, and 1 represents identical newsletters.

The main conceptual difference between the three measures is that cosine and Jaccard similarity compare unordered sets, and thus measure similarity in terms of words used, whereas the Levenshtein distance considers the order or sequence of words. This implies that cosine and Jaccard similarity tend to provide similar results, whereas the Levenshtein distance can be quite different, as they measure two different types of document similarity.

Changes are generally more easily detected using Levenshtein distance, as simply moving a sentence to another part of the newsletter is considered change, although the newsletter may still convey the exact same information. If we instead use cosine or Jaccard similarity, no change is detected. All three measures have pros and cons, but cosine and Jaccard similarity may be better suited for newsletter analysis. We regard changes in words used as more informative than overall insertions, deletions, and substitutions.

To illustrate how the similarity measures are calculated in practice, consider the following sentences:

A. "We are very positive to the market."

B. "We are still positive to the market."

The union is

$$A \cup B = [\text{we, are, very, positive, to, the, market, still}]$$

The term frequency vectors are

$$A = [1, 1, 1, 1, 1, 1, 1, 0]$$
$$B = [1, 1, 0, 1, 1, 1, 1, 1]$$

Cosine similarity is

$$C(A, B) = \frac{1 \cdot 1 + 1 \cdot 1 + 1 \cdot 0 + 1 \cdot 1 + 1 \cdot 1 + 1 \cdot 1 + 1 \cdot 1 + 0 \cdot 1}{(\sqrt{1^2 + 1^2 + 1^2 + 1^2 + 1^2 + 1^2 + 1^2}) \cdot (\sqrt{1^2 + 1^2 + 1^2 + 1^2 + 1^2 + 1^2 + 1^2})} = 0.86$$

Jaccard similarity is

$$J(A, B) = \frac{|[\text{we, are, positive, to, the, market}]|}{|[\text{we, are, very, positive, to, the, market, still}]|} = \frac{6}{8} = 0.75$$

Sentence A turns into sentence B by replacing “very” with “still”. Since only one operation is required, we have $L(A, B) = 1$. As the length of both sentence A and B is 7, the normalised Levenshtein distance is

$$1 - \frac{1}{7} = 0.86$$

3.4 Sentiment and Similarity Data

Table 3.3 summarises sentiment and similarity scores for the 15 mutual funds’ newsletters. Again, note that time frames are different across our sample of funds, such that the average sentiment and similarity scores are not directly comparable.

A sentiment score around zero indicates neutral sentiment. A higher score indicates positive sentiment, and a lower score indicates negative sentiment. Thus, we observe that most mutual funds in our sample have positive sentiment on average. In addition to this, sentiment does not seem to directly correspond to the obtained alphas observed in table 3.2. This is perhaps a result of different writing styles, as well as differences in the content of fund managers’ text.

Table 3.3: Sentiment and Similarity Summary Statistics

	Sentiment	Cosine	Jaccard	Levenshtein
Alfred Berg Aktiv	0.21	0.70	0.49	0.47
Alfred Berg Gambak	0.02	0.73	0.53	0.50
Alfred Berg Norge Classic	0.20	0.70	0.48	0.47
Arctic Norwegian Equities	0.72	0.71	0.42	0.37
C WorldWide Norge	0.91	0.35	0.19	0.09
Delphi Norge	0.55	0.40	0.18	0.09
Delphi Vekst	0.28	0.41	0.18	0.10
DNB Norge Selektiv (III)	-0.01	0.42	0.23	0.14
DNB SMB	0.35	0.43	0.29	0.20
Fondsfinans Norge	0.52	0.59	0.37	0.23
FORTE Norge	1.75	0.57	0.20	0.07
FORTE Tronder	1.71	0.61	0.20	0.07
Holberg Norge	1.01	0.65	0.45	0.40
Pareto Aksje Norge A	1.10	0.54	0.32	0.27
Pareto Investment Fund A	1.30	0.44	0.19	0.11

This table reports average sentiment, cosine similarity, Jaccard similarity, and Levenshtein distance scores for the mutual funds. A sentiment score around zero indicates neutral sentiment. A higher score indicates positive sentiment, and a lower score indicates negative sentiment. Cosine similarity, Jaccard similarity, and Levenshtein distance is measured in the range $[0, 1]$, where higher values indicate higher similarity.

Similarity scores differ depending on the measure used. Jaccard similarity is typically lower than cosine similarity as it only counts each word once. Levenshtein distance is typically the lowest measure, as all insertions, deletions, and substitutions are counted.

Lastly, we observe that mutual funds within the same fund family tend to have similar sentiment and similarity scores. This is not surprising as newsletter set-ups are often identical within fund families. Some of them may be more adapted for fund managers' subjective opinions, while others aim for more objective, standardised monthly comments.

4 Empirical Results

This section analyses our research questions and presents the results. Subsection 4.1 studies the informativeness of mutual fund newsletters, subsection 4.2 studies fund manager efforts, and subsection 4.3 studies investor behaviour in mutual funds. Finally, in subsection 4.4, we construct quintile portfolios based on textual measures and test for significant alphas.

4.1 Are Mutual Fund Newsletters Informative?

To analyse the informativeness of newsletters, we test whether there is a connection between monthly newsletter characteristics and associated performance. Gross alpha is used as a measure of fund performance, defined as the fund return before costs in excess of the benchmark return.

When choosing the preferred estimation method for our panel data, we consider the assumptions behind random effects and fixed effects estimation. Random effects assumes that fund-specific effects are uncorrelated with the independent variables, while fixed effects assumes correlation. If omitted variables are correlated with the independent variables, random effects can cause biased estimates. In our case, the difference in alphas varies greatly between funds. It is reasonable to believe that unobserved factors such as fund manager skills explain at least parts of the alpha, and not including a skill measure can cause bias. We assume that this effect is constant over time, and we therefore use fixed effect estimation to account for it[8].

More specifically, we account for time-invariant fund-specific effects as well as time (year and month) fixed effects. Month effects account for patterns such as a decreased focus on fund newsletters during the summer months. Year effects are included to account for general trends or fluctuations in the economy. We regress textual measures on performance, measured by gross alpha. As alpha is simply defined as the benchmark-adjusted return, fund performance can be strong relative to benchmark performance. Textual measures are expected to reflect market performance as well as fund alphas. For instance, sentiment may have been low during the financial crisis because of the market situation. Our independent

[8]This choice has been further confirmed by running Hausman specification tests.

variables may not be able to explain this variation if funds outperformed their benchmark and thus obtained positive alphas. For this reason, time effects account for newsletter descriptions of special market situations that do not reflect fund-specific performance.

The regression equation is defined as follows

$$Score_{i,t} = \alpha + \beta Alpha_{i,t} + \gamma C_{i,t} + f_i + \delta_t + \epsilon_{i,t} \tag{4.1}$$

The dependent variable, $Score_{i,t}$, is the respective sentiment or similarity score for fund i in month t. $Alpha_{i,t}$ is fund i's gross alpha (performance) in month t. $C_{i,t}$ is a vector including the following fund-level control variables: expense ratio, active share, size (AUM), and net flow. We include fund f_i and time δ_t (year and month) fixed effects. The results are presented in table 4.1.

Table 4.1: Performance Regressions

	Dependent variable:							
	Sentiment		Cosine similarity		Jaccard similarity		Levenshtein distance	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Gross alpha	0.09***	0.08***	0.003**	0.003*	0.004*	0.005*	0.004*	0.004*
	(3.51)	(3.89)	(2.49)	(1.84)	(1.94)	(2.03)	(2.02)	(2.01)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,149	1,145	1,105	1,103	1,105	1,103	1,105	1,103
Adjusted R^2	0.24	0.26	0.47	0.51	0.44	0.45	0.47	0.48

This table reports results of fixed effect regressions of textual measures on mutual fund gross alphas. We have included the following control variables: expense ratio, active share, the logarithm of AUM as a measure of fund size, and net flow. Full regressions, including controls, are located in appendix A3. Gross alpha is multiplied by 100. Fund and time (year and month) fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund and year. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

We note that performance is highly significant and positive for newsletter sentiment. Thus, there seems to be some credibility in what is written in mutual fund newsletters, driving sentiment in a direction that is consistent with funds' actual performance. This does however not imply a 1:1 relationship. For instance, one may expect funds to exaggerate the positive newsletter tone when performance has increased recently, but not sufficiently lower the positive tone when performance has decreased. The effect of performance on newsletter sentiment depending on whether performance in the recent past has been strong or weak will be examined more closely later in this subsection.

Next, we observe that performance is positive and significant at the 10% level for all three similarity measures. Although the coefficients are low, the observed consistency across measures indicates that higher performance is associated with higher newsletter similarity. In other words, an increase in performance seems to be linked to less changes made to newsletters, and vice versa. When the observed results speak for themselves, there should be no need for fund managers to use newsletter text to convince investors to stay. Our findings on similarity are consistent with Cohen et al. (2020)'s findings on firm disclosures, in which significant changes are related to lower returns. Thus, we find that a similar relationship may exist for Norwegian equity mutual funds and their newsletters.

Overall, there seems to be a relationship between observed performance and the newsletter characteristics studied in this paper. This signals that mutual fund newsletters are informative, at least to some extent, although our findings do not imply that the information is capable of predicting performance in the following month. Such relationships will be tested in subsection 4.4.

As a result of the strong sentiment significance, we test whether a change in sentiment is associated with a change in gross alpha in the relevant month. Simultaneously, we test whether the change in sentiment is different depending on recent performance. We use the past three months' alphas ($t-3$, $t-2$, and $t-1$) to categorise funds as high or low performers. If the average three-month alpha is at or below zero, the fund is categorised as a low performer, while if alpha is higher than zero, the fund is categorised as a high performer. In other words, we categorise by whether funds recently outperformed their benchmark index or not. We regress the change in current month's sentiment ($t-1$ to t) on the change in current month's alpha ($t-1$ to t) for both groups. The results are presented in table 4.2. Controls and fixed effects are equivalent to those used in table 4.1.

Table 4.2: Sentiment and Performance Group Regressions

	Dependent variable:			
	ΔSentiment			
	Low performance		High performance	
	(1)	(2)	(3)	(4)
ΔGross alpha	0.08	0.07	0.10***	0.10***
	(1.69)	(1.48)	(4.34)	(4.37)
Controls	No	Yes	No	Yes
Fund FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Observations	439	439	632	632
Adjusted R^2	0.004	0.003	0.02	0.02

This table reports results of fixed effect regressions of changes in sentiment on changes in gross alphas. Change is defined as the change from month $t-1$ to t. The sample is divided in two groups by the average alpha over the previous three months: $t-3$, $t-2$, and $t-1$. Low performance includes three-month alphas at or lower than zero, while high performance includes three-month alphas higher than zero. We have included the following control variables: expense ratio, active share, the logarithm of AUM as a measure of fund size, and net flow. Full regressions, including controls, are located in appendix A3. Gross alpha is multiplied by 100. Fund and time (year and month) fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund and year. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

We observe that for low performers, newsletter sentiment does not seem to be sensitive to changes in performance. High performers' newsletter sentiment, on the other hand, is sensitive to changes in performance. If performance increases for this group, sentiment increases significantly, and vice versa. Thus, sentiment seems to be a reflection of recent performance only when performance has been high in the recent past. This is interesting considering that one may expect low performers to communicate positive changes in performance to avoid investor outflows. Our findings suggest that recent low performers are generally more restrained in their newsletter communication.

4.2 Fund Manager Efforts

In this subsection, we study whether newsletter characteristics affect fund manager efforts, measured by active share. Observed alpha is an objective measure of past performance, while newsletter text and its sentiment can be subjective. We hypothesise that fund managers may be increasingly willing to deviate from the benchmark index when they have written more positively about the fund in recent times. This may signal that their subjective view of the fund's performance is increasingly positive. Further, we also study potential relationships between newsletter similarity and fund manager efforts. Here, results in either direction could be plausible: fund managers who spend less time changing or rewriting their newsletters may have more active investment strategies, or perhaps active fund managers are also active writers who make sure to update investors on new active positions.

We use fixed effects estimation for reasons specified in subsection 4.1. We account for time-invariant fund-specific effects only, as active share is generally stable over time. We run regressions on previous month's sentiment and the similarity from previous month's to current month's newsletter. A lagged sentiment measure is used to study the effect on active positions in the following month. The regression equation can be expressed as follows

$$AS_{i,t} = \alpha + \beta Score_{i,t-1,t} + \gamma C_{i,t-1} + f_i + \epsilon_{i,t} \tag{4.2}$$

The dependent variable, $AS_{i,t}$, is fund i's active share in month t. $Score_{i,t-1,t}$ is the respective measure for fund i from month $t-1$ to t. For sentiment, this is the sentiment score in month $t-1$. For similarity, this is the similarity score between month $t-1$ and t. $C_{i,t-1}$ is a vector of the following lagged fund-level control variables: gross alpha, expense ratio, size (AUM), and net flow. We include fund f_i fixed effects. The results are presented in table 4.3.

Table 4.3: Active Share Regressions

	Dependent variable:							
	Active share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sentiment$[t-1]$	0.38	0.38						
	(1.02)	(1.28)						
Cosine$[t-1,t]$			3.06	3.01				
			(0.49)	(0.57)				
Jaccard$[t-1,t]$					−0.61	−1.15		
					(−0.15)	(−0.25)		
Levenshtein$[t-1,t]$							0.04	−0.34
							(0.01)	(−0.08)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No	No	No
Observations	1,134	1,130	1,105	1,100	1,105	1,100	1,105	1,100
Adjusted R^2	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84

This table reports results of fixed effect regressions of active share on the four textual measures: sentiment, cosine similarity, Jaccard similarity, and Levenshtein distance. Previous month's measure is used for the sentiment variable, while the similarity from previous month's newsletter to the current month's newsletter is used for the similarity variables. We have included the following lagged control variables: gross alpha, expense ratio, the logarithm of AUM as a measure of fund size, and net flow. Full regressions, including controls, are located in appendix A4. Active share is multiplied by 100. Fund fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

We find no significant relationships between sentiment and active share, or similarity and active share. Thus, we cannot conclude that recent newsletter sentiment affects fund managers' decision to take active positions. Moreover, we cannot conclude that fund managers who change their newsletters less are more or less active. With regard to sentiment, the lack of findings may arise because some funds simply summarise recent updates, such that the information is not able to predict fund manager efforts. Alternatively, the lack of findings may be linked to the idea that newsletters are distributed for marketing reasons, and thus, newsletters are not related to fund manager efforts. The communicated tone does not seem to reflect fund managers' perception or opinions of their mutual fund, at least not to the extent that it affects active share.

As we find no overall connection, we further study fund manager efforts and its relation to sentiment in more detail. We divide our fund sample in three groups by active share: low, mid, and high. Low includes all observations where active share is at or lower than 40%. High includes all observations where active share is at or higher than 60%. Mid includes observations with active share in the range 40% to 60%. These numbers are based on ESMA (2016), where funds with an active share of less than 50% are classified as potentially being closet-indexers[9] in relatively small equity markets. As the mean and median of our sample is close to 50%, and many observations are clustered around this level, we use three groups to more clearly separate "closet-indexers" from truly active funds. The results are presented in table 4.4. Controls and fixed effects are consistent with those used in table 4.3.

Table 4.4: Active Share Group Regressions

	Dependent variable:					
	Active share					
	Low		Mid		High	
	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment$[t-1]$	0.53	0.57**	−0.17	−0.33	−0.39**	−0.23
	(1.90)	(2.63)	(−0.63)	(−1.22)	(−2.47)	(−0.99)
Controls	No	Yes	No	Yes	No	Yes
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No
Observations	288	288	427	426	419	416
Adjusted R^2	0.20	0.22	0.53	0.54	0.66	0.67

This table reports results of fixed effect regressions of active share on previous month's newsletter sentiment. The sample of mutual funds is divided in three groups by active share: low (AS at 40% or below), mid (AS between 40% and 60%), and high (AS at 60% or above). We have included the following lagged control variables: gross alpha, expense ratio, the logarithm of AUM as a measure of fund size, and net flow. Full regressions, including controls, are located in appendix A4. Active share is multiplied by 100. Fund fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

From the active share group results, we observe that sentiment is positively correlated with active share when active share is low. When active share is high, on the other hand, sentiment is negatively correlated with active share, although in this case, the significance

[9] A closet index fund is an actively managed mutual fund that stays close to its benchmark index, and therefore does not justify the higher fees.

vanishes with controls. A plausible explanation to these results may be that "truly active" funds have a set active share target, and are therefore less likely to be affected by other factors. "Closet indexers", on the other hand, might vary more over time in terms of activeness. Based on our results, managers of funds with low active share seem to be increasingly willing to deviate from the benchmark index when the newsletter sentiment has increased in recent times. This is perhaps linked to increased confidence.

4.3 Investor Flows

This subsection studies whether investors respond to information conveyed in mutual fund newsletters. To analyse this question, we test whether information from the most recent available newsletter is reflected in the following month's net flows. In an efficient market, investors are expected to react to available, useful information. We expect sentiment to affect flows as changes to newsletter sentiment should be relatively easy for investors to notice. Overall newsletter similarity, on the other hand, might be more difficult for investors to detect.

Net flows are regressed on one of our four textual measures. We use fixed effects estimation for reasons specified in subsection 4.1. We account for time-invariant fund-specific effects as well as time (year and month) fixed effects. Year effects are included to account for general trends or fluctuations in the economy. Monthly effects are included as patterns in investor activity may exist, such as the January effect. The regression equation is defined as

$$Flow_{i,t} = \alpha + \beta Score_{i,t-1} + \gamma C_{i,t-1} + f_i + \delta_t + \epsilon_{i,t} \tag{4.3}$$

where the dependent variable, $Flow_{i,t}$, is fund i's net flow in month t. $Score_{i,t-1}$ is the respective sentiment or similarity score for fund i in month $t-1$. We use lagged textual measures to analyse whether previous month's newsletter – the most recent newsletter available to investors – affects flows in the following month. $C_{i,t-1}$ is a vector of the following lagged fund-level control variables: gross alpha, expense ratio, active share, and size (AUM). We include fund f_i and time δ_t (year and month) fixed effects. The results are presented in table 4.5.

Table 4.5: Net Flow Regressions

	Dependent variable:							
	Net flow							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sentiment[$t-1$]	0.37	0.32						
	(0.94)	(0.87)						
Cosine[$t-2, t-1$]			1.69	2.38				
			(0.64)	(0.92)				
Jaccard[$t-2, t-1$]					−1.03	−1.02		
					(−0.37)	(−0.35)		
Levenshtein[$t-2, t-1$]							1.81	1.73
							(0.62)	(0.59)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,132	1,132	1,089	1,089	1,089	1,089	1,089	1,089
Adjusted R^2	0.15	0.16	0.15	0.16	0.15	0.16	0.15	0.16

This table reports results of fixed effect regressions of mutual fund net flows on previous month's sentiment and similarity scores. We have included the following lagged control variables: gross alpha, expense ratio, active share, and the logarithm of AUM as a measure of fund size. Full regressions, including controls, are located in appendix A5. Net flow is multiplied by 100. Fund and time (year and month) fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund and year. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

We find no relationship between newsletter sentiment or newsletter similarity and flows. Our results suggest that sentiment as well as overall changes in monthly newsletters go unnoticed. Our findings on sentiment are in contrast to the findings of Hillert et al. (2016), as they find that a less negative tone in shareholder reports (forms N-CSR and N-CSRS) is related to higher net flows. However, these shareholder reports are not directly comparable to the monthly newsletters that we analyse. Our findings are likely affected by the fact that many of the newsletters in our sample are not publicly available, while some funds make only their most recent newsletter available. In other words, the information may be difficult to obtain, and far from all investors are able to act on this information. In addition to this, it might be that investors with access to newsletters do not read them, or perhaps do not find the newsletters to be helpful for investment choices. On the other hand, if it turns out that newsletter information predicts future performance, and investors do not respond to the information, this may signal a market inefficiency.

4.4 Quintile Portfolios

We have shown that newsletter sentiment and similarity seem to be informative, and that investors do not seem to act on newsletter information. In this subsection, we study whether the newsletter information has been predictive for the following month's performance, and not only consistent with current month's performance. This is done by constructing quintile portfolios and testing their performance using simple t-tests and factor models.

We create five portfolios for each of the four measures: sentiment, cosine similarity, Jaccard similarity, and Levenshtein distance. Portfolios in month t are based on information from the newsletter in month $t-1$. This way, portfolios are based on information from newsletters that are available to investors. All monthly observations from January 2010 to May 2018 are included in one of the quintile portfolios. Observations prior to this are excluded to ensure a sufficient number of funds within our portfolios. Q1 includes the observations with the lowest respective measure score, while Q5 includes the observations with the highest measure score. The quintile portfolios are equally-weighted, and rebalanced on a monthly basis, for each new newsletter made available.

Note that this subsection analyses fund performance after costs, as we base t-tests on net alphas and factor models on net returns. In the following subsection, net alpha is defined as benchmark-adjusted returns after subtracting costs. In relation to factor models, net alpha is defined as the net return that cannot be explained by the model.

4.4.1 Benchmark-Adjusted Returns

Table 4.6 reports monthly net alphas, including t-statistics from one-sided t-tests, to study whether the alphas are significantly higher than zero. We include a long-short (Q5-Q1) alpha for each of the four measures, and include t-statistics to show whether the Q5 alpha is significantly different from the Q1 alpha. For sentiment, we test whether a high-sentiment portfolio performs significantly better than a low-sentiment portfolio. For the similarity measures, we run high-low similarity to test whether a "non-changer" portfolio performs significantly better than a "changer" portfolio, equivalent to Cohen et al. (2020). Note that the presented long-short portfolios are hypothetical, as it is not possible to short the mutual funds included in our analysis.

Table 4.6: Quintile Portfolio Net Alphas

Sentiment					
Q1	Q2	Q3	Q4	Q5	Q5-Q1
0.23**	0.15	-0.02	0.13	0.04	-0.19
(1.71)	(1.16)	(-0.19)	(0.84)	(0.21)	(-1.19)
Cosine Similarity					
Q1	Q2	Q3	Q4	Q5	Q5-Q1
-0.07	0.23*	0.07	0.07	0.23*	0.30**
(-0.50)	(1.52)	(0.51)	(0.43)	(1.51)	(1.68)
Jaccard Similarity					
Q1	Q2	Q3	Q4	Q5	Q5-Q1
0.00	0.15	-0.02	0.05	0.35***	0.35**
(-0.03)	(0.90)	(-0.13)	(0.35)	(2.41)	(1.89)
Levenshtein Distance					
Q1	Q2	Q3	Q4	Q5	Q5-Q1
0.20*	-0.06	0.14	0.01	0.24*	0.03
(1.37)	(-0.49)	(0.96)	(0.08)	(1.37)	(0.17)

This table reports benchmark-adjusted returns after costs and one-sided t-tests for quintile and long-short portfolios on the four textual measures. We test whether benchmark-adjusted returns are significantly higher than zero. Q1 is the quintile with the lowest measure, while Q5 is the quintile with the highest measure.

We find that a portfolio that goes long high-sentiment funds and shorts low-sentiment funds does *not* obtain a significant net alpha. In fact, we observe a positive and significant alpha for the low-sentiment (Q1) portfolio, while the high-sentiment (Q5) portfolio alpha is insignificant. This strengthens the hypothesis that newsletters function as a marketing tool intended to attract investors.

Long-short portfolios based on cosine and Jaccard similarity obtain positive and significant alphas at the 5% level. This suggests that high similarity predicts high performance in the following month. The portfolios based on Levenshtein distance, which is calculated in a very different manner, moderates our confidence in this conclusion. However, as noted in section 3, we consider cosine and Jaccard similarity to be the most suitable measures when analysing newsletter similarity, as they are based on words and not sequences.

Figure 4.1 illustrates cumulative net alphas over time for Q1 and Q5 portfolios. The holding period is January 2010 to May 2018, as previously mentioned. We observe that low sentiment (Q1) performs better, but not significantly better than high sentiment (Q5) over time, as seen in table 4.6. We note high cumulative alphas for the Q5 portfolios based on cosine and Jaccard similarity. There is no clear difference over time between Q1 and Q5 Levenshtein distance.

Figure 4.1: Cumulative Quintile Portfolio Alphas

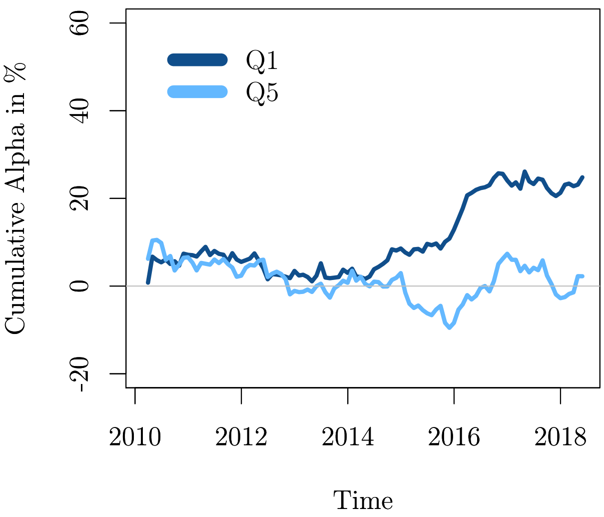

(a) Sentiment

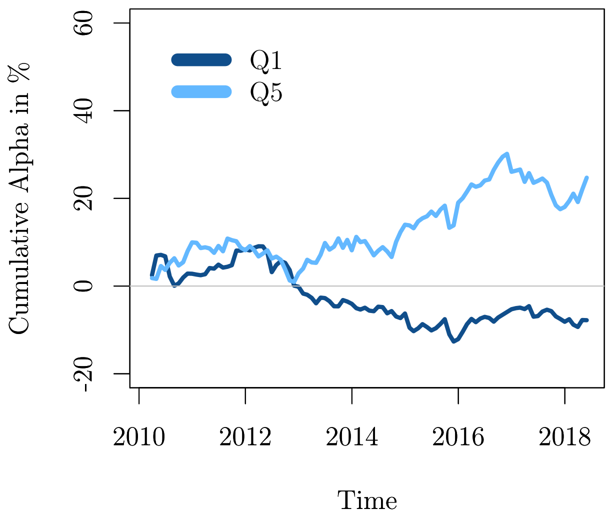

(b) Cosine Similarity

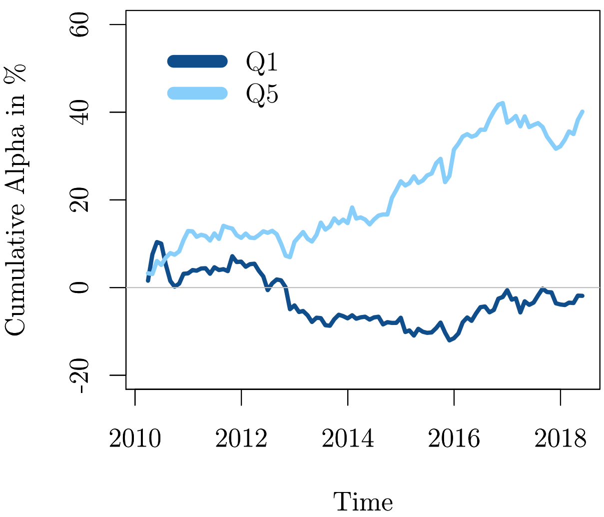

(c) Jaccard Similarity

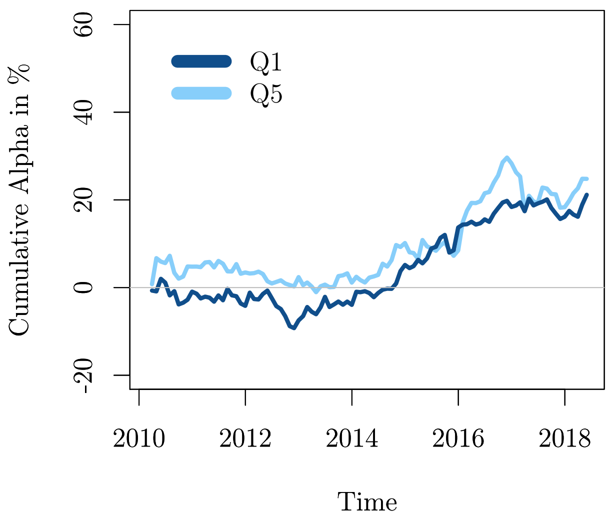

(d) Levenshtein Distance

Figure 4.2 plots the four long-short portfolio alphas from January 2010 to May 2018. We observe a negative 23% alpha for the sentiment portfolio, a positive 33% alpha for the cosine similarity portfolio, a positive 42% alpha for the Jaccard similarity portfolio, and a positive 4% alpha for the Levenshtein distance portfolio in May 2018. As noted in table 4.6, only the cosine and Jaccard similarity long-short alphas are significant.

Figure 4.2: Cumulative Long-Short Portfolio Alphas

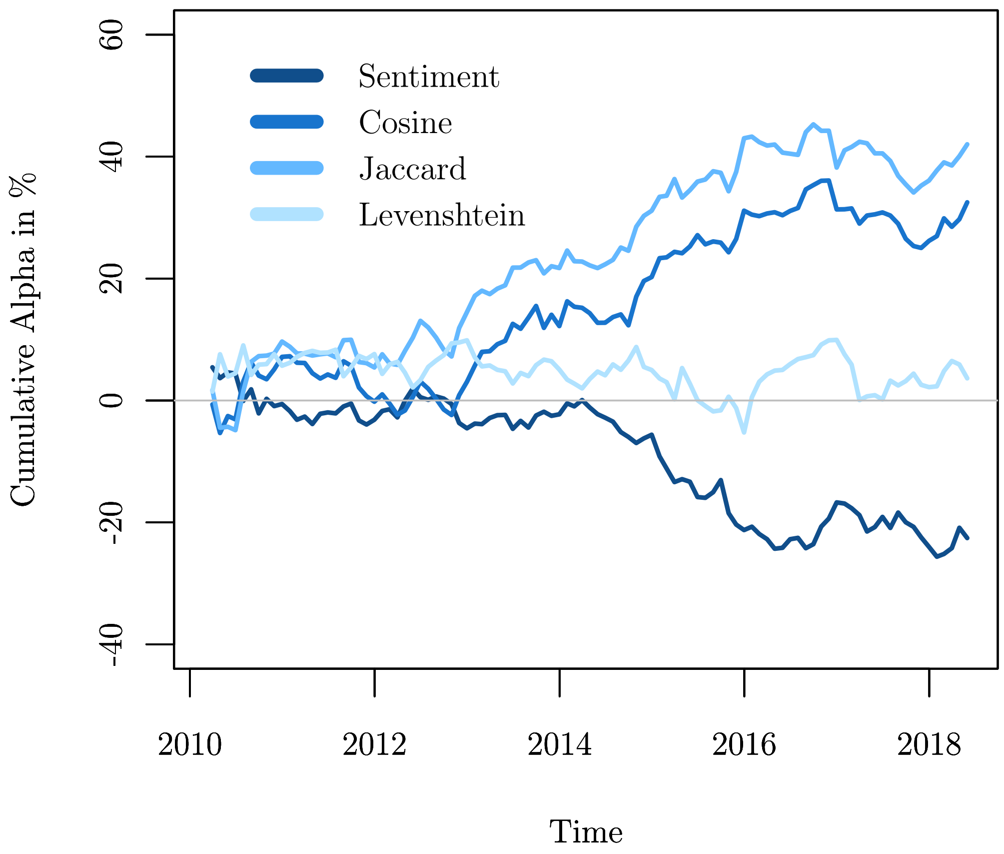

4.4.2 Factor Regressions

Next, we test quintile and long-short portfolios for each of the four textual measures against common factor loadings, to analyse whether returns can be explained by these factors. Factor portfolio alphas are presented in table 4.7. For full regressions including factor loadings, see appendix A6.

We run the following three regressions per portfolio: CAPM (market), the traditional three-factor model (market, SMB, HML), and an extended five-factor model (market, SMB, HML, UMD, LIQ). The three regression equations are defined as follows

$$r_t - r_t^{RF} = \alpha + \beta_1(r_t^{MKT} - r_t^{RF}) + \epsilon_t \tag{4.4}$$

$$r_t - r_t^{RF} = \alpha + \beta_1(r_t^{MKT} - r_t^{RF}) + \beta_2 SMB_t + \beta_3 HML_t + \epsilon_t \tag{4.5}$$

$$r_t - r_t^{RF} = \alpha + \beta_1(r_t^{MKT} - r_t^{RF}) + \beta_2 SMB_t + \beta_3 HML_t + \beta_4 UMD_t + \beta_5 LIQ_t + \epsilon_t \tag{4.6}$$

The dependent variable is the monthly net return on portfolio r_t minus the monthly risk-free rate r_t^{RF}. We use OSEBX[10] as a proxy for the market, a broad index relevant to our sample of Norwegian equity mutual funds. Most of our funds are general active funds that mainly invest in the Norwegian equity market without sector-specific exposure in their mandates. Exceptions, including a SMB fund, are expected to be identified by one of the included pricing factors. Pricing factors and risk-free rates based on the Norwegian market are obtained from Bernt Arne Ødegaard's webpage[11]. Fama-French size (SMB) and value (HML) factors are calculated according to Fama and French (1998). The momentum factor (UMD) is calculated according to Carhart (1997), and the liquidity factor (LIQ) is constructed according to Næs et al. (2009).

[10] Oslo Børs Benchmark Index, comprising the most traded shares listed on Oslo Børs.
[11] http://finance.bi.no/~bernt/financial_data/ose_asset_pricing_data/index.html

Table 4.7: Quintile Portfolio Regressions

	Sentiment						*Cosine Similarity*					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
CAPM alpha	0.26	0.25	0.18	0.437	0.21	-0.05	0.01	0.31	0.21	0.31	0.42	0.41**
	(0.80)	(0.74)	(0.54)	(1.11)	(0.61)	(-0.32)	(0.04)	(0.92)	(0.64)	(0.91)	(1.30)	(2.12)
3-factor alpha	0.43	0.40	0.38	0.53*	0.33	-0.10	0.15	0.48	0.39	0.49	0.56*	0.41**
	(1.43)	(1.25)	(1.21)	(1.68)	(1.00)	(-0.60)	(0.45)	(1.53)	(1.28)	(1.51)	(1.81)	(2.05)
5-factor alpha	0.46	0.41	0.42	0.55*	0.30	-0.16	0.19	0.53	0.40	0.49*	0.54*	0.35*
	(1.63)	(1.39)	(1.47)	(1.95)	(1.01)	(-0.96)	(0.59)	(1.84)	(1.43)	(1.72)	(1.85)	(1.69)

	Jaccard Similarity						*Levenshtein Distance*					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
CAPM alpha	0.09	0.33	0.02	0.27	0.56*	0.47**	0.34	0.10	0.14	0.18	0.52	0.18
	(0.27)	(0.90)	(0.07)	(0.82)	(1.69)	(2.61)	(0.95)	(0.29)	(0.42)	(0.52)	(1.60)	(0.93)
3-factor alpha	0.25	0.48	0.19	0.43	0.71**	0.47**	0.51	0.22	0.31	0.36	0.66**	0.15
	(0.76)	(1.40)	(0.64)	(1.41)	(2.29)	(2.55)	(1.52)	(0.70)	(0.99)	(1.14)	(2.15)	(0.76)
5-factor alpha	0.31	0.52	0.21	0.40	0.70**	0.39**	0.57*	0.29	0.31	0.34	0.63**	0.06
	(1.05)	(1.62)	(0.78)	(1.44)	(2.44)	(2.04)	(1.87)	(1.00)	(1.03)	(1.21)	(2.19)	(0.26)

This table reports quintile portfolio net alphas. The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. For each textual measure, we construct quintiles based on previous month's measure scores. Q1 is the quintile with the lowest measure score, while Q5 is the quintile with the highest measure score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM alphas (market), three-factor alphas (market, size, value), and five-factor alphas (market, size, value, momentum, liquidity). Full regressions, including factor loadings, are located in appendix A6. t-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

We note that the positive and significant alpha for Q1 sentiment vanishes in the factor regressions. The lack of significance for the long-short sentiment portfolio again proves that fund managers with a very positive tone do not perform significantly better in the following month compared to the least positive fund managers. "Blindly" trusting fund managers' text and investing in funds with highly positive sentiment has not led to alphas higher than what other sentiment portfolios would have obtained.

The cosine similarity long-short portfolio obtains a positive alpha significant at the 5% level in the three-factor model, and at the 10% level in the five-factor model, implying that the alpha cannot solely be explained by traditional risk factors. The Jaccard similarity long-short portfolio obtains a positive and significant alpha at the 5% level, even in the five-factor model. Our results on cosine similarity and Jaccard similarity are consistent with the results of Cohen et al. (2020) on American firm disclosures: going long in a "non-changer" portfolio and shorting a "changer" portfolio obtains a significant alpha. Thus, we provide evidence of a similar relationship for equity mutual funds' monthly newsletters.

Again, the Levenshtein distance long-short portfolio obtains no significant alpha. Both Q1 and Q5 are positive and significant in the five-factor model. As previously mentioned, the lack of similarity consistency when taking Levenshtein distance portfolios into consideration is likely due to differences in computation. For our purpose, we consider cosine and Jaccard similarity to be the preferred similarity measures as they measure changes in words instead of sequences.

Note that costs are subtracted in this analysis. Thus, our findings hold even after taking costs into consideration. The comprehensive, existing literature on active management of mutual funds shows that it is difficult for active mutual funds to beat the market after costs over time. The significant results on cosine and Jaccard similarity are in contrast to this.

As previously mentioned, these portfolios have been constructed using information in previous month's newsletter (i.e., $t-1$), which is the most recent newsletter available to investors. In appendix A7, we construct equally-weighted portfolios for each quintile based on the three previous newsletters, i.e., $t-3$, $t-2$, and $t-1$. For each new month, the oldest newsletter is excluded and a new newsletter is included. This allows us to analyse whether the observed results persist over a longer period of time.

The sentiment long-short portfolio alpha is, as expected, still insignificant. Further, both the Q1 and Q5 sentiment alphas are significant at the 10% level in the five-factor model, with the Q1 alpha only marginally higher than the Q5 alpha. For cosine and Jaccard similarity, we find that the high-similarity portfolios (Q5) are still positive and significant at the 10% level in the three-factor and five-factor models. However, the long-short significance vanishes, except for Jaccard similarity at a 10% level when controlling for the market. Overall, there seems to be limited persistence in the results.

4.5 Discussion and Summary

In this subsection, we review and summarise our findings on newsletter sentiment and similarity.

In subsection 4.1, we find that current month's performance, as measured by gross alpha, is strongly significant for the sentiment of current month's newsletter: increased performance is linked to higher newsletter sentiment, and vice versa. These findings suggest that the sentiment aspect of mutual fund newsletters is informative and gives a correct representation of fund performance. However, by delving deeper into performance groupings, we find that only recent high-performs seem to sufficiently adjust their newsletter sentiment when performance changes.

Sentiment does not seem to be related to active share, as observed in subsection 4.2. After dividing funds in groups by degree of active management, however, we find a positive and significant relationship for the "closet-index" group. Thus, managers of such funds seem increasingly willing to deviate from the benchmark when their fund has had higher sentiment recently.

In subsection 4.3, our results indicate that investors do not act on sentiment-related information conveyed in newsletters. There might be multiple reasons for this, but the most evident explanation is perhaps limited newsletter accessibility or that investors who receive newsletters do not read them or find the information useful. Not acting on newsletter sentiment does however seem to be a good choice, as our quintile portfolio analysis in subsection 4.4 shows that a high-sentiment portfolio (Q5) does not outperform a low-sentiment portfolio (Q1). In fact, the low-sentiment portfolio obtains positive and significant benchmark-adjusted returns after costs, while the high-sentiment portfolio does not. Thus, our results indicate that sentiment is not consistent with performance in the following month, strengthening the hypothesis that newsletters are intended as marketing. Alternatively, if funds instead objectively summarise the current month, and do not focus on future performance, there may not be a relationship because fund managers do not attempt to predict the following month. As sentiment is related to performance in the relevant month, as seen in subsection 4.1, newsletters likely serve as more than just marketing.

For informativeness based on the similarity between two consecutive months' newsletters, we find a positive relationship between similarity and performance in subsection 4.1. High-performing funds seem to change their newsletters less from month to month, letting the results speak for themselves. This relationship holds for all three similarity measures. We find no relationship between newsletter similarity and fund manager efforts in subsection 4.2.

In subsection 4.3, we find that changes to newsletters go unnoticed by investors, as there is no significant relationship between similarity and net flows. Overall changes to newsletters may be difficult for investors to discover, or perhaps investors do not believe that changes can predict fund performance. This signals a form of investor inattention and market inefficiency. Again, the accessibility of newsletters is perhaps affecting the results. Nonetheless, we find that a "non-changer" portfolio performs significantly better than a "changer" portfolio in subsection 4.4 based on cosine and Jaccard similarity. This signals that investors are missing meaningful information. The lack of significance for Levenshtein distance moderates our confidence in the conclusion, but we consider cosine and Jaccard similarity to be the preferred measures when applied to newsletters. The significant alphas on cosine and Jaccard similarity portfolios, even after costs, are not solely explained by common risk factors, and is consistent with the findings of Cohen et al. (2020) on firm disclosures.

We would like to make a note on the accessibility of monthly mutual fund newsletters. Making investment decisions based on such information is dependent on investors having access to a variety of newsletters. Some mutual funds make their most recent newsletters publicly available, while others only give access to their investors. Thus, acting on newsletter information is perhaps only possible for larger investors with access to all available information, in contrast to small, uninformed investors. In addition to this, all newsletters are not necessarily published or distributed to investors at the very end of the month, but may instead be made available at some point in the following month. Regardless, our suggested long-short portfolios are primarily constructed to illustrate inefficiencies, and are not meant for trading purposes.

Lastly, the restricted sample size is a limitation to our analysis, particularly in the number of mutual funds. This arises from difficulties in obtaining mutual fund newsletters. Newsletters within a particular fund family typically have similar traits, implying that some of our funds have similar characteristics, which can introduce bias. We do however include newsletters from multiple fund families in Norway, which reduces this problem. There may be reasons why certain funds do not wish to hand out their history of newsletters, such as underperformance. An analysis based on a larger part of the Norwegian equity mutual fund market is necessary in order to conclude about the overall relationship, as we cannot be certain that our sample is representative for the entire market.

5 Conclusion

In this thesis, we analyse the sentiment and similarity of Norwegian equity funds' monthly newsletters, and relate these textual features to mutual fund performance, fund manager efforts, and investor flows. To conclude, we will summarise our findings on sentiment and similarity, followed by thesis limitations and suggested areas of further research.

We find increased performance to be associated with increased newsletter sentiment and similarity for the relevant month. However, changes in alphas reflect changes in sentiment only for recent high-performers. Further, we find no overall relationship between sentiment or similarity and fund manager efforts. The exception is that increased sentiment seems to be related to increased fund manager efforts in the following month for closet-index funds.

We find no overall proof that investors respond to newsletter information. We do however hypothesise that this is partly due to the limited accessibility of newsletters, making it difficult for newsletters to function as an effective marketing tool. Our quintile portfolio analysis shows that high-sentiment funds do not outperform low-sentiment funds. Thus, newsletter sentiment does not correspond to the following month's performance, and investors should not chase high-sentiment funds. Finally, we find that a "non-changer" portfolio outperforms a "changer" portfolio for two out of three similarity measures, and that this is robust to common risk factors. This signals the existence of a market inefficiency that is not caught by investors.

The most evident limitation to our thesis is the sample size. Potential bias may be introduced as multiple fund families are not willing to share their history of newsletters. In addition to this, a limited sample makes it difficult to further test the robustness of our results by for instance testing different time periods. Due to few mutual funds and relatively few years of newsletter observations, this is not possible at the time of writing.

In addition to this, textual analysis is still an emerging field in accounting and finance. There is a need for further development of textual analysis technology to assure its applicability. For analyses on Norwegian language in particular, available resources are limited. For instance, there are few sentiment lexicons available, and these are not tailored to financial data. For this reason, conducting analyses on specific fields of study becomes more demanding. It is also worth noting that lexical-based approaches do not guarantee measurement of the true, underlying semantics of the text. In our case, the varying content and writing styles in newsletters may also complicate the process. A larger sample size would likely diversify away problems with potential outliers.

Textual analysis has been performed on a wide variety of text. Considering the limited availability of analyses conducted on the Norwegian language, and Norwegian financial text in particular, we suggest this as an area of further research that deserves increased attention. This could for instance be related to Norwegian media articles, or Norwegian annual and quarterly reports. With regard to mutual funds, mutual fund holdings can be tested, and linked to relevant text, such as annual reports.

References

Barnes, J., Touileb, S., Øvrelid, L., and Velldal, E. (2019). Lexicon information in neural sentiment analysis: a multi-task learning approach. In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 175–186.

Boytsov, L. (2011). Indexing methods for approximate dictionary searching: Comparative analysis. *Journal of Experimental Algorithmics*, 16:1.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance*, 52(1):57–82.

Cohen, L., Malloy, C., and Nguyen, Q. (2020). Lazy prices. *The Journal of Finance.*

Cremers, K. J. M. and Petajisto, A. (2009). How active is your fund manager? a new measure that predicts performance. *The Review of Financial Studies*, 22(9):3329–3365.

ESMA (2016). Public statement - supervisory work on potential closet index tracking. https://www.esma.europa.eu/sites/default/files/library/2016-165_public_statement_-_supervisory_work_on_potential_closet_index_tracking.pdf. Accessed: 2020-06-10.

Fama, E. F. and French, K. R. (1998). Value versus growth: The international evidence. *The Journal of Finance*, 53(6):1975–1999.

Gallaher, S., Kaniel, R., and Starks, L. T. (2006). Madison avenue meets wall street: Mutual fund families, competition and advertising. *Competition and Advertising (January 2006).*

Hanley, K. W. and Hoberg, G. (2010). The information content of ipo prospectuses. *The Review of Financial Studies*, 23(7):2821–2864.

Hillert, A., Niessen-Ruenzi, A., and Ruenzi, S. (2016). Mutual fund shareholder letters: Flows, performance, and managerial behavior. *Performance, and Managerial Behavior (March 29, 2016).*

Hu, M. and Liu, B. (2004). Mining and summarizing customer reviews. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 168–177.

Jensen, M. C. (1968). The performance of mutual funds in the period 1945–1964. *The Journal of Finance*, 23(2):389–416.

Jin, L., Eshraghi, A., Taffler, R., and Goyal, A. (2015). Fund manager active share, overconfidence and investment performance. *Working Paper, Warwick University, University of Edinburgh, and University of Lausanne.*

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *The Journal of Finance*, 66(1):35–65.

Loughran, T. and McDonald, B. (2016). Textual analysis in accounting and finance: A survey. *Journal of Accounting Research*, 54(4):1187–1230.

Manning, C. D., Raghavan, P., and Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press.

Næs, R., Skjeltorp, J. A., and Ødegaard, B. A. (2009). What factors affect the oslo stock exchange?

Niwattanakul, S., Singthongchai, J., Naenudorn, E., and Wanapu, S. (2013). Using of jaccard coefficient for keywords similarity. In *Proceedings of the International Multiconference of Engineers and Computer Scientists*, volume 1, pages 380–384.

Sharpe, W. F. (1991). The arithmetic of active management. *Financial Analysts Journal*, 47(1):7–9.

Sirri, E. R. and Tufano, P. (1998). Costly search and mutual fund flows. *The Journal of Finance*, 53(5):1589–1622.

Solomon, D. H., Soltes, E., and Sosyura, D. (2014). Winners in the spotlight: Media coverage of fund holdings as a driver of flows. *Journal of Financial Economics*, 113(1):53–72.

Appendix

A1 Introductory Example

Table A1.1: Extracts from August 2016 Newsletters – Norwegian

FORTE Norge	**C WorldWide Norge**
Selskapet (BW Offshore) har en lang kontraktreserve og den finansielle risiko er nå minimert frem til 2020. Det gjenstår likevel en ikke ubetydelig teknisk risiko ved det store nybygget av produksjonsskipet Catcher som skal i drift på engelsk sektor til neste år. Går det problemfritt er aksjen en klar doblingskandidat fra dagens nivå. [...] Det er flere indikatorer som tyder på at norsk økonomi nå har passert bunnen og at BNP-veksten vil ta seg opp i takt med økte oljepriser. Rentenivået vil fortsatt være rekordlavt og i tillegg drive stadig mere kapital inn i aksjemarkedet i søk etter bedre avkastning enn man får i rentemarkedet. Dette betyr at OSLO Børs kan bli et veldig hyggelig sted å være i tiden fremover. Vi skal stå på for å fortsette å gi meravkastning.	De største positive overraskelsene i vår portefølje, i forhold til konsensusestimatene, var Aker Solutions, Subsea7, Schibsted, Borregaard, Det Norske Oljeselskap og Orkla. Det skal nevnes at våre forventninger til flere av disse selskapene var høyere enn markedets forventninger og at vi dermed hadde posisjonert oss i forhold til dette. [...] Det som påvirket porteføljen negativt i juli var det kraftige oljeprisfallet. Brent olje falt fra USD50 per fat ved inngangen til måneden til USD42 per fat ved slutten av måneden, en nedgang på hele 16 %. Porteføljen er direkte eksponert mot oljeprisfallet gjennom våre investeringer i oljeselskapene Statoil og Det Norske Oljeselskap, samt oljeserviceselskapene Subsea7 og Aker Solutions. Oljeprisfallet har først og fremst vært kraftig på spotprisen på olje, mens for lengre leveringstidsperioder har fallet vært mer moderat.

This table includes extracts from two mutual funds' August 2016 newsletters in Norwegian.

A2 Sentiment Analysis

A2.1 Common Words

Words commonly used in the mutual fund newsletters are reported in table A2.1.

Table A2.1: Common Words

Positive	anbefaler, ansvarlig, attraktiv, attraktive, attraktivt, bedre, bedret, bedring, bekrefter, bekreftet, best, beste, billig, billige, bra, børsoppgangen, dyktighet, enighet, enkelt, fart, fin, fordel, fornuftig, fornøyd, fred, fremvoksende, garanti, gevinst, gjennomført, gjennomførte, gledelig, god, gode, godt, gunstig, gunstige, hensyn, hjulpet, hyggelig, hyggelige, håp, høy, høye, høyere, høyest, høyeste, høyt, integrert, interessant, interessante, interesse, klarer, klarte, kursoppgang, kursoppgangen, kvalitet, lede, ledende, ledet, lett, lettet, like, liker, likt, lovende, lover, lykkes, lys, løft, løfte, løftet, lønnsom, markedsleder, nyter, opp, oppgang, oppgangen, oppnå, oppnådd, oppnådde, oppnår, oppnås, oppover, opprettholde, opprettholder, opprettholdt, oppside, oppsiden, optimisme, optimistiske, overbevisende, overgikk, passerte, pen, pene, pluss, positiv, positive, positivt, potensial, potensiale, potensialet, raskere, raskt, rekordhøy, rekordhøye, ren, riktig, rimelig, robust, robuste, sikkert, sikre, sikrer, sikret, skape, skaper, skapt, skapte, slo, slår, slått, solid, solide, stabil, stabile, stabilt, steg, steget, stige, stigende, stiger, stimulerende, styrke, styrket, støtte, suksess, takket, tildelt, tilfredsstillende, tillit, topp, toppen, trolig, trygge, tydelig, tålmodighet, utbytte, utmerket, utvikle, utvikler, utviklet, vant, vellykket, verdsatt, verdt, viktig, viktige, viktigste, vinner, vinneren, vinnerne, vokse, voksende, vokser, vokst, øke, økende, øker, øket, økning, økningen, økt, økte, ønsker
Negative	bekymret, bekymring, bekymringer, bunn, dessverre, drøy, drøye, drøyt, dyrere, dårlig, dårlige, dårligere, fall, falle, fallende, faller, fallet, falt, faren, feil, fell, felt, frykt, frykten, frykter, fryktet, følsom, gjeld, grådig, hard, hardt, irritert, knapp, knappe, konsekvens, konsekvensene, krevende, kursfall, kursfallet, kursnedgang, kursnedgangen, kutte, kuttet, lav, lave, lavere, lavest, laveste, lavt, manglende, midlertidig, mindre, minst, minste, minus, ned, nedgang, nedgangen, nedjustert, nedover, nedtur, nedturen, nedturer, negativ, negative, negativt, oppbremsing, press, presset, problemene, problemer, rammet, reduksjon, redusere, redusert, reduserte, resesjon, risikable, risiko, risikoen, skuffelse, skuffelser, skuffende, skuffet, skyldes, sliter, strenge, svak, svake, svakere, svakest, svakeste, svakt, svekke, svekkelse, svekket, svingninger, svingningene, tap, tapene, taperen, taperne, tapet, tapt, tapte, tilbakeslag, tunge, tungt, turbulent, tvil, uro, uroen, urolig, urolige, urolighetene, uroligheter, usikkerhet, usikkerheten, usikre, utfordrende, utfordringene, utfordringer, utsatt, uventet, vanskelig, vanskelige, verre, verste, vond

This table reports words included in our sentiment lexicon that are used at least 10 times in our sample of mutual fund newsletters. All relevant word endings are reported.

A2.2 Sentiment Lexicon

We use the full-form sentiment lexicon by Barnes et al. (2019), and make the changes reported in table A2.2.

Table A2.2: Sentiment Lexicon Changes

	Positive
Added:	bedring, fart, fremgang, himmelferd, høy, kvalitet, lyspunkt, lønnsom, markedslede, oppadgående, oppgang, oppjuster, oppover, oppside, opptur, pengemaskin, potensial, rekordhøy, spektakulær, suksessful, suksessrik, synergi, utbytte
Removed:	absolutt, alvor, autoritær, balanse, bar, bedra, begjærlig, beskjed, bestem, betydelig, bidrag, bløt, bære, demp, dominer, dristig, drøy, eksempel, ekstra, ekstrem, energi, enkelte, enorm, fantasifull, fengsl, fikk, flaut, flyt, fløt, forbaus, forbløff, fordele, forenkl, forfin, fornem, forsikring, forsiktig, forskjønn, fort, fortjenest, forundr, fremstå, følsom, få, gi, gjelde, grep, hele, helle, hels, helt, hevde, idealiser, indirekte, intens, irritabel, jevn, klar, korreksjon, korriger, kraftig, kupp, lakk, langvarig, led, lek, let, lik, listig, lokk, markant, minsk, mør, målløs, nervepirrende, ny, nytte, næring, nåd, omfatt, oppsiktsvekkende, overflod, overrask, overveld, pris, rett, roer, rolig, rør, sammenheng, sannsynlig, sinn, skarp, skatt, skikkelig, slås, slåtte, smal, spar, spenn, spent, spesiell, still, stor, strev, stør, støt, tank, tapp, tilfeldig, tiltal, tro, triviel, tyn, tøf, uberørt, uforlignelig, under, urokkelig, usedvanlig, utrolig, utvikling, varig, varsom, vekst, vel, veldig, verdi, verk, vesentlig, vid, vint, virk, vis, ydmyk, økonomisk
	Negative
Added:	bunn, dessverre, destabiliser, fall, følsom, helsvart, innstramming, mislighold, ned, nedgang, nedjuster, nedover, nedside, nedtur, oppbremsing, overoppheting, rekordlav, skakkjørt, stup, svart, svingning, tøf, ulønnsom
Removed:	abnorm, alvor, ansvarlig, avhengig, avslutt, avvik, bank, bedr, begjærlig, begrens, bekymringsløs, besitt, billig, blend, bløt, brent, bulk, by, del, dominer, dristig, evn, fantasifull, fart, fengsl, fint, flir, flyt, fløt, forbaus, forbløff, forenkl, forestå, forfin, forfør, forlokk, forskjønn, forundr, frigjor, frigjør, fôr, full, gammel, gjelde, glis, godt, heng, høy, idealiser, isoler, jul, kapitalist, kjenn, knøttlit, knøttsmå, konkurrent, konsesjon, kontroller, korthet, kortsiktig, kostnad, kraftig, krev, kupp, lakk, langvarig, led, lek, let, lis, list, listig, lit, liter, litt, lokk, løs, marginal, massiv, mektig, merk, minsk, motsatt, motset, myk, målløs, nervepirrende, never, omfatt, oppbring, oppgi, oppsiktsvekkende, ordinær, overdriv, overflod, overlege, overrask, overvekt, overveld, rakk, re, redd, ris, roe, rustet, rykk, sen, skarp, skift, slo, sinn, slutt, slå, smal, smiger, smigr, spenn, spent, spiss, sterk, stol, streb, styr, støtt, sure, sørger, sørget, tapp, tilfeldig, tiltal, triviel, tross, true, tynn, tynt, uberørt, ubeskjeden, uerstattelig, uforlignelig, ulik, underliggende, undervurder, unngikk, unngå, unormal, uovervinnelig, urokkelig, utløp, utnytt, utrolig, uvanlig, uvirkelig, variabel, varier, varsl, varsom, ve, vekt, vell, velt, verk, ydmyk, ørlit, ørsmå

This table reports the additions and deletions we make to the sentiment lexicon by Barnes et al. (2019). Note that the table reports only the stem of words. As we use a full-form lexicon, additions and deletions include all word endings.

A2.3 Valence Shifters

Table A2.3 reports the amplifiers, de-amplifiers, and negators included in the sentiment analysis.

Table A2.3: Valence Shifter List

Amplifiers:	absolutt, ualminnelig, betraktelig, betydelig, desidert, ekstra, ekstrem, enorm, enda, historisk, klar, kraftig, langt, maksimal, markant, markert, massiv, meget, skikkelig, soleklar, spesiel, sterk, stor, svært, unormal, usedvanlig, utrolig, uvanlig, veldig, vesentlig
De-amplifiers:	begrens, delvis, demp, forholdsvis, ganske, knøttlit, knøttsmå, liten, litt, marginal, minimal, noe, ørlit, ørsmå
Negators:	aldri, ikke, ingen, snudd, unngå, uten

This table reports our list of valence shifters: amplifiers, de-amplifiers, and negators. Note that the table only reports the stem of words. As we use a full-form lexicon, additions and deletions include all word endings.

A2.4 Sentiment Function

The *polarity score* is the sum of all weighted context clusters divided by the square root of the word count, and can be summarised as follows

$$\delta = \frac{c'_{i,j}}{\sqrt{w_{i,jn}}}$$

where

$$(1)\ c'_{i,j} = \sum([(1 + z(w_a - w_d)] \cdot w^p_{i,j,k} \cdot (-1)^{w_n})$$

$$(2)\ w_a = \sum(w_{n_1} \cdot w^a_{i,j,k})$$

$$(3)\ w_d = \max(w'_d, -1)$$

$$(4)\ w'_d = \sum(w_{n_2} \cdot w^a_{i,j,k} + w^d_{i,j,k})$$

$$(5)\ w_n = (\sum w^n_{i,j,k}) \bmod(2)$$

$$(w_n = 1 \Rightarrow w_{n_1} = 0, w_{n_2} = 1)\ \vee\ (w_n = 0 \Rightarrow w_{n_1} = 1, w_{n_2} = 0)$$

$w_{i,j,k}$ is word k in paragraph i and sentence j. $w^n_{i,j,k}$ is negators, $w^a_{i,j,k}$ is amplifiers, $w^d_{i,j,k}$ is de-amplifiers, $w^p_{i,j,k}$ is the polarised word score, z is the amplifier/de-amplifier weight (default 0.8), $c'_{i,j}$ is the sum of weighted context clusters, $w_{i,jn}$ is the word count (where jn is the number of words in a sentence), and δ is the polarity score.

(1) Sum of all weighted context clusters: The weight of amplifiers is added to the sentiment factor (as they increase sentiment), while the weight of de-amplifiers reduce the factor (as they reduce sentiment). Amplifier and de-amplifier weight is multiplied by z. The polarised word score from the lexicon is multiplied by the sentiment factor. The sum is then multiplied by either 1 (if there is an even number of negators) or -1 (if there is an odd number of negators).

(2) Weight of amplifiers: If there is an even number of negators ($w_n = 0$), amplifiers remain amplifiers. The weight from amplifiers is then equal to the total times amplifiers are used in the context cluster.

(3) Restriction on the weight of de-amplifiers: The weight from de-amplifiers has a lower bound of -1. A bound lower than -1 would turn sentiment negative.

(4) Weight of de-amplifiers: If there is an odd number of negators ($w_n = 1$), amplifiers become de-amplifiers. The weight of de-amplifiers is then the total times amplifiers and de-amplifiers are used in the context cluster. If there is an even number of negators ($w_n = 0$), amplifiers remain amplifiers. The weight of de-amplifiers is then only the total times de-amplifiers are used.

(5) Check if there is an even or odd number of negators: If there is an even number of negators, then $w_n = 0$. If there is an odd number of negators, then $w_n = 1$. w_{n1} and w_{w2} are defined depending on whether the number of negators is even or odd. This is used in (2) and (4).

A3 Performance Regressions

Table A3.1: Performance Regressions

	Dependent variable:							
	Sentiment		Cosine similarity		Jaccard similarity		Levenshtein distance	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Gross alpha	0.09***	0.08***	0.003**	0.003*	0.004*	0.005*	0.004*	0.004*
	(3.51)	(3.89)	(2.49)	(1.84)	(1.94)	(2.03)	(2.02)	(2.01)
Expense ratio		−0.85		−0.16		0.19		0.14
		(−0.47)		(−0.36)		(0.48)		(0.29)
Active share		0.01		0.002*		−0.0002		0.001
		(1.06)		(1.89)		(−0.17)		(0.33)
log(AUM)		0.14		0.06**		0.03		0.02
		(1.74)		(2.23)		(1.04)		(0.76)
Net flow		0.01*		0.001		0.0001		0.001
		(2.16)		(0.94)		(0.08)		(0.96)
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,149	1,145	1,105	1,103	1,105	1,103	1,105	1,103
Adjusted R^2	0.24	0.26	0.47	0.51	0.44	0.45	0.47	0.48

This table reports results of fixed effect regressions of textual measures on mutual fund gross alphas. Gross alpha, expense ratio, active share and flow are multiplied by 100. Fund and time (year and month) fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund and year. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A3.2: Sentiment and Performance Group Regressions

	Dependent variable:			
	ΔSentiment			
	Low performance		High performance	
	(1)	(2)	(3)	(4)
ΔGross alpha	0.08	0.07	0.10***	0.10***
	(1.69)	(1.48)	(4.34)	(4.37)
Expense ratio		0.86		1.73
		(0.68)		(0.91)
Active share		−0.004		−0.003
		(−0.35)		(−0.55)
log(AUM)		−0.05		0.03
		(−0.56)		(0.33)
Net flow		0.02*		0.003
		(2.14)		(0.59)
Fund FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Observations	439	439	632	632
Adjusted R^2	0.004	0.003	0.02	0.02

This table reports results of fixed effect regressions of changes in sentiment on changes in gross alphas. Change is defined as the change from month $t-1$ to t. The sample is divided in two groups by the average alpha over the previous three months: $t-3$, $t-2$, and $t-1$. Low performance includes three-month alphas at or lower than zero, while high performance includes three-month alphas higher than zero. Gross alpha, expense ratio, active share, and flow are multiplied by 100. Fund and time (year and month) fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund and year. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

A4 Active Share Regressions

Table A4.1: Active Share Regressions

	Dependent variable:							
	Active share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sentiment$[t-1]$	0.38	0.38						
	(1.02)	(1.28)						
Cosine$[t-1, t]$			3.06	3.01				
			(0.49)	(0.57)				
Jaccard$[t-1, t]$					−0.61	−1.15		
					(−0.15)	(−0.25)		
Levenshtein$[t-1, t]$							0.04	−0.34
							(0.01)	(−0.08)
Gross alpha$[t-1]$		0.32*		0.32*		0.33*		0.33*
		(2.01)		(2.12)		(2.10)		(2.10)
Expense ratio$[t-1]$		36.95**		36.85**		36.50***		36.39***
		(2.84)		(2.64)		(3.05)		(2.98)
log(AUM$[t-1]$)		0.29		0.20		0.42		0.39
		(0.12)		(0.09)		(0.18)		(0.16)
Net flow$[t-1]$		−0.07*		−0.07*		−0.06		−0.06
		(−1.87)		(−1.79)		(−1.75)		(−1.75)
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No	No	No
Observations	1,134	1,130	1,105	1,100	1,105	1,100	1,105	1,100
Adjusted R^2	0.84	0.84	0.84	0.84	0.84	0.84	0.84	0.84

This table reports results of fixed effect regressions of active share on the four textual measures: sentiment, cosine similarity, Jaccard similarity, and Levenshtein distance. Previous month's measure is used for the sentiment variable, while the similarity from previous month's newsletter to the current month's newsletter is used for the similarity variables. Active share, gross alpha, expense ratio, and flow are multiplied by 100. Fund fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A4.2: Active Share Group Regressions

	Dependent variable:					
	Active share					
	Low		Mid		High	
	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment[$t-1$]	0.53	0.57**	−0.17	−0.33	−0.39**	−0.23
	(1.90)	(2.63)	(−0.63)	(−1.22)	(−2.47)	(−0.99)
Gross alpha[$t-1$]		0.26		0.24***		0.08
		(1.16)		(3.53)		(0.68)
Expense ratio[$t-1$]		115.75		2.94		32.98**
		(1.70)		(0.26)		(2.74)
log(AUM[$t-1$])		0.73		1.37*		−0.07
		(0.72)		(1.88)		(−0.06)
Net flow[$t-1$]		−0.20**		0.01		−0.08
		(−3.06)		(0.39)		(−1.48)
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No
Observations	288	288	427	426	419	416
Adjusted R^2	0.20	0.22	0.53	0.54	0.66	0.67

This table reports results of fixed effect regressions of active share on previous month's newsletter sentiment. The sample of mutual funds has been divided in three groups by active share: low (AS at 40% or below), mid (AS between 40% and 60%), and high (AS at 60% or above). Active share, gross alpha, expense ratio, and flow are multiplied by 100. Fund fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

A5 Net Flow Regressions

Table A5.1: Net Flow Regressions

	Dependent variable:							
	Net flow							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sentiment$[t-1]$	0.37	0.32						
	(0.94)	(0.87)						
Cosine$[t-2, t-1]$			1.69	2.38				
			(0.64)	(0.92)				
Jaccard$[t-2, t-1]$					−1.03	−1.02		
					(−0.37)	(−0.35)		
Levenshtein$[t-2, t-1]$							1.81	1.73
							(0.62)	(0.59)
Gross alpha$[t-1]$		0.21		0.26		0.27		0.26
		(0.72)		(0.83)		(0.87)		(0.82)
Expense ratio$[t-1]$		28.65**		30.54**		30.35**		29.84**
		(2.65)		(2.78)		(2.59)		(2.75)
Active share$[t-1]$		0.03		0.04		0.05		0.05
		(0.44)		(0.68)		(0.74)		(0.72)
log(AUM$[t-1]$)		−1.05		−1.33		−1.16		−1.22
		(−0.69)		(−0.83)		(−0.71)		(−0.75)
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,132	1,132	1,089	1,089	1,089	1,089	1,089	1,089
Adjusted R^2	0.15	0.16	0.15	0.16	0.15	0.16	0.15	0.16

This table reports results of fixed effect regressions of mutual fund net flows on previous month's sentiment and similarity scores. Flow, gross alpha, expense ratio, and active share are multiplied by 100. Fund and time (year and month) fixed effects are included. t-Statistics are reported below the estimates, based on standard errors clustered by fund and year. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

A6 Quintile Portfolios with Factor Loadings

Table A6.1: Sentiment Quintile Portfolio Regressions

	Dependent variable:																	
	Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.26	0.25	0.18	0.37	0.21	−0.05	0.43	0.40	0.38	0.53*	0.33	−0.10	0.46	0.41	0.42	0.55*	0.30	−0.16
	(0.80)	(0.74)	(0.54)	(1.11)	(0.61)	(−0.32)	(1.43)	(1.25)	(1.21)	(1.68)	(1.00)	(−0.60)	(1.63)	(1.39)	(1.47)	(1.95)	(1.01)	(−0.96)
RMRF	0.79***	0.80***	0.76***	0.69***	0.71***	−0.08*	0.69***	0.70***	0.65***	0.59***	0.63***	−0.06	0.52***	0.51***	0.44***	0.38***	0.43***	−0.08*
	(8.82)	(8.51)	(8.07)	(7.49)	(7.60)	(−1.77)	(8.18)	(7.81)	(7.42)	(6.74)	(6.84)	(−1.22)	(6.08)	(5.73)	(5.24)	(4.55)	(4.79)	(−1.70)
SMB							−0.44***	−0.45***	−0.48***	−0.42***	−0.33***	0.10**	−0.13	−0.11	−0.10	−0.03	0.05	0.18***
							(−4.61)	(−4.49)	(−4.84)	(−4.20)	(−3.21)	(2.00)	(−1.21)	(−0.98)	(−0.96)	(−0.32)	(0.47)	(2.93)
HML							0.01	−0.05	0.03	0.001	−0.004	−0.01	0.05	−0.01	0.07	0.05	0.04	−0.01
							(0.11)	(−0.55)	(0.28)	(0.02)	(−0.04)	(−0.29)	(0.58)	(−0.14)	(0.89)	(0.58)	(0.47)	(−0.14)
UMD													−0.05	−0.04	−0.06	−0.05	−0.01	0.04
													(−0.78)	(−0.60)	(−0.91)	(−0.70)	(−0.14)	(1.09)
LIQ													−0.47***	−0.52***	−0.58***	−0.58***	−0.59***	−0.11**
													(−4.86)	(−5.10)	(−6.00)	(−6.07)	(−5.68)	(−2.02)
N	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Adj R^2	0.44	0.42	0.39	0.36	0.36	0.02	0.53	0.51	0.51	0.45	0.42	0.04	0.62	0.61	0.64	0.60	0.56	0.08

This table reports sentiment quintile portfolio regression results. We construct quintiles based on previous month's sentiment scores. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. t-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A6.2: Cosine Similarity Quintile Portfolio Regressions

	Dependent variable:																	
	Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.01	0.31	0.21	0.31	0.42	0.41**	0.15	0.48	0.39	0.49	0.56*	0.41**	0.19	0.53*	0.40	0.49*	0.54*	0.35*
	(0.04)	(0.92)	(0.64)	(0.91)	(1.30)	(2.12)	(0.45)	(1.53)	(1.28)	(1.51)	(1.81)	(2.05)	(0.59)	(1.84)	(1.43)	(1.72)	(1.85)	(1.69)
RMRF	0.81***	0.76***	0.77***	0.75***	0.66***	−0.15***	0.72***	0.66***	0.66***	0.66***	0.57***	−0.15***	0.51***	0.46***	0.48***	0.44***	0.39***	−0.12*
	(8.35)	(8.19)	(8.47)	(7.96)	(7.32)	(−2.85)	(7.60)	(7.49)	(7.83)	(7.26)	(6.54)	(−2.72)	(5.47)	(5.32)	(5.74)	(5.12)	(4.55)	(−1.90)
SMB							−0.40***	−0.44***	−0.47***	−0.40***	−0.40***	0.002	−0.04	−0.09	−0.13	0.01	−0.07	−0.03
							(−3.80)	(−4.43)	(−4.92)	(−3.96)	(−4.13)	(0.03)	(−0.32)	(−0.85)	(−1.22)	(0.08)	(−0.66)	(−0.44)
HML							−0.03	0.01	−0.01	0.05	−0.04	−0.003	0.01	0.05	0.03	0.10	0.001	−0.01
							(−0.33)	(0.10)	(−0.07)	(0.52)	(−0.39)	(−0.05)	(0.12)	(0.62)	(0.42)	(1.20)	(0.02)	(−0.16)
UMD													−0.06	−0.07	−0.04	−0.03	−0.01	0.05
													(−0.80)	(−1.04)	(−0.56)	(−0.51)	(−0.15)	(0.99)
LIQ													−0.56***	−0.53***	−0.52***	−0.63***	−0.50***	0.06
													(−5.26)	(−5.37)	(−5.43)	(−6.42)	(−5.03)	(0.90)
N	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Adj R^2	0.41	0.40	0.42	0.39	0.35	0.07	0.48	0.49	0.53	0.47	0.43	0.05	0.59	0.61	0.63	0.62	0.55	0.04

This table reports cosine similarity quintile portfolio regression results. We construct quintiles based on previous month's cosine similarity scores. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. *t*-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A6.3: Jaccard Similarity Quintile Portfolio Regressions

	Dependent variable:																	
	Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.09	0.33	0.02	0.27	0.56*	0.47**	0.25	0.48	0.19	0.43	0.71**	0.47**	0.31	0.52	0.21	0.40	0.70**	0.39**
	(0.27)	(0.90)	(0.07)	(0.82)	(1.69)	(2.61)	(0.76)	(1.40)	(0.64)	(1.41)	(2.29)	(2.55)	(1.05)	(1.62)	(0.78)	(1.44)	(2.44)	(2.04)
RMRF	0.77***	0.78***	0.80***	0.76***	0.64***	−0.13***	0.68***	0.68***	0.71***	0.66***	0.54***	−0.13**	0.46***	0.47***	0.52***	0.48***	0.36***	−0.10*
	(8.21)	(7.76)	(9.07)	(8.39)	(7.02)	(−2.64)	(7.48)	(7.01)	(8.43)	(7.70)	(6.25)	(−2.56)	(5.27)	(4.89)	(6.31)	(5.68)	(4.18)	(−1.79)
SMB							−0.41***	−0.44***	−0.40***	−0.44***	−0.42***	−0.01	−0.04	−0.08	−0.06	−0.09	−0.07	−0.03
							(−4.05)	(−4.08)	(−4.24)	(−4.57)	(−4.29)	(−0.14)	(−0.34)	(−0.62)	(−0.55)	(−0.81)	(−0.62)	(−0.40)
HML							−0.002	−0.02	0.04	−0.02	−0.01	−0.01	0.04	0.02	0.08	0.02	0.03	−0.01
							(−0.02)	(−0.23)	(0.40)	(−0.19)	(−0.11)	(−0.14)	(0.52)	(0.22)	(0.98)	(0.29)	(0.37)	(−0.24)
UMD													−0.08	−0.06	−0.05	−0.01	−0.02	0.06
													(−1.15)	(−0.80)	(−0.72)	(−0.12)	(−0.26)	(1.36)
LIQ													−0.58***	−0.57***	−0.52***	−0.54***	−0.53***	0.04
													(−5.79)	(−5.13)	(−5.60)	(−5.60)	(−5.46)	(0.66)
N	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Adj R^2	0.40	0.37	0.45	0.41	0.33	0.06	0.48	0.46	0.53	0.51	0.43	0.04	0.61	0.57	0.64	0.62	0.56	0.04

This table reports Jaccard similarity quintile portfolio regression results. We construct quintiles based on previous month's Jaccard similarity scores. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. *t*-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A6.4: Levenshtein Distance Quintile Portfolio Regressions

	Dependent variable:																	
	Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.34	0.10	0.14	0.18	0.52	0.18	0.51	0.22	0.31	0.36	0.66**	0.15	0.57*	0.29	0.31	0.34	0.63**	0.06
	(0.95)	(0.29)	(0.42)	(0.52)	(1.60)	(0.93)	(1.52)	(0.70)	(0.99)	(1.14)	(2.15)	(0.76)	(1.87)	(1.00)	(1.03)	(1.21)	(2.19)	(0.26)
RMRF	0.77***	0.79***	0.81***	0.74***	0.64***	−0.13**	0.67***	0.70***	0.72***	0.63***	0.55***	−0.12**	0.45***	0.49***	0.55***	0.42***	0.38***	−0.07
	(7.82)	(8.58)	(8.84)	(7.76)	(7.15)	(−2.32)	(7.09)	(7.84)	(8.16)	(7.07)	(6.38)	(−2.05)	(4.88)	(5.64)	(6.13)	(4.97)	(4.38)	(−1.12)
SMB							−0.44***	−0.39***	−0.40***	−0.50***	−0.39***	0.04	−0.06	−0.03	−0.07	−0.11	−0.05	0.004
							(−4.14)	(−3.86)	(−4.04)	(−4.99)	(−4.05)	(0.70)	(−0.48)	(−0.31)	(−0.64)	(−1.02)	(−0.48)	(0.05)
HML							0.01	−0.04	0.04	−0.02	−0.01	−0.03	0.06	0.004	0.08	0.03	0.02	−0.03
							(0.12)	(−0.42)	(0.43)	(−0.17)	(−0.15)	(−0.43)	(0.67)	(0.05)	(0.92)	(0.35)	(0.30)	(−0.57)
UMD													−0.08	−0.08	−0.02	−0.02	−0.001	0.08*
													(−1.15)	(−1.22)	(−0.34)	(−0.30)	(−0.01)	(1.66)
LIQ													−0.59***	−0.55***	−0.50***	−0.59***	−0.51***	0.08
													(−5.65)	(−5.49)	(−4.92)	(−6.05)	(−5.25)	(1.08)
N	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Adj. R^2	0.38	0.42	0.44	0.37	0.34	0.04	0.46	0.49	0.51	0.49	0.42	0.03	0.59	0.61	0.61	0.63	0.54	0.05

This table reports Levenshtein distance quintile portfolio regression results. We construct quintiles based on previous month's Levenshtein distance scores. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. t-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

A7 3-Month Rolling Average Analysis

Table A7.1: 3-Month Rolling Average Alphas

	Sentiment						*Cosine Similarity*					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
CAPM alpha	0.32	0.26	0.26	0.24	0.43	0.12	0.19	0.30	0.29	0.30	0.43	0.24
	(0.90)	(0.76)	(0.77)	(0.65)	(1.30)	(0.66)	(0.59)	(0.91)	(0.81)	(0.78)	(1.24)	(1.44)
3-factor alpha	0.43	0.41	0.40	0.38	0.55*	0.12	0.33	0.44	0.43	0.41	0.56*	0.24
	(1.29)	(1.35)	(1.29)	(1.07)	(1.72)	(0.65)	(1.10)	(1.40)	(1.32)	(1.12)	(1.74)	(1.40)
5-factor alpha	0.54*	0.41	0.39	0.41	0.53*	-0.01	0.41	0.43	0.51*	0.42	0.51*	0.10
	(1.78)	(1.44)	(1.42)	(1.24)	(1.88)	(-0.03)	(1.58)	(1.47)	(1.71)	(1.25)	(1.75)	(0.59)
	Jaccard Similarity						*Levenshtein Distance*					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
CAPM alpha	0.15	0.31	0.28	0.26	0.51	0.37*	0.25	0.05	0.13	0.02	0.39	0.14
	(0.43)	(0.88)	(0.85)	(0.73)	(1.42)	(1.84)	(0.72)	(0.13)	(0.38)	(0.06)	(1.14)	(0.61)
3-factor alpha	0.29	0.45	0.41	0.39	0.63*	0.33	0.40	0.19	0.25	0.18	0.48	0.08
	(0.93)	(1.37)	(1.32)	(1.18)	(1.82)	(1.65)	(1.23)	(0.59)	(0.78)	(0.53)	(1.44)	(0.34)
5-factor alpha	0.37	0.48	0.44	0.40	0.59*	0.21	0.45	0.24	0.30	0.19	0.43	-0.02
	(1.35)	(1.62)	(1.53)	(1.34)	(1.82)	(1.02)	(1.54)	(0.82)	(1.00)	(0.62)	(1.39)	(-0.08)

This table reports quintile portfolio net alphas. The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. For each textual measure, we construct quintiles based on the three previous months' measure scores, equally weighted. Q1 is the quintile with the lowest measure, while Q5 is the quintile with the highest measure. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM alphas (market), three-factor alphas (market, size, value), and five-factor alphas (market, size, value, momentum, liquidity). t-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A7.2: 3-Month Sentiment Regressions

	Dependent variable: Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.32	0.26	0.26	0.24	0.43	0.12	0.43	0.41	0.40	0.38	0.55^{*}	0.12	0.54^{*}	0.41	0.39	0.41	0.53^{*}	−0.01
	(0.90)	(0.76)	(0.77)	(0.65)	(1.30)	(0.66)	(1.29)	(1.35)	(1.29)	(1.07)	(1.72)	(0.65)	(1.78)	(1.44)	(1.42)	(1.24)	(1.88)	(−0.03)
RMRF	0.77^{***}	0.76^{***}	0.79^{***}	0.79^{***}	0.72^{***}	−0.05	0.67^{***}	0.65^{***}	0.68^{***}	0.70^{***}	0.64^{***}	−0.04	0.45^{***}	0.46^{***}	0.47^{***}	0.49^{***}	0.42^{***}	−0.03
	(7.82)	(8.01)	(8.26)	(7.53)	(7.68)	(−0.94)	(7.11)	(7.45)	(7.68)	(6.88)	(6.98)	(−0.73)	(4.90)	(5.41)	(5.63)	(4.80)	(4.90)	(−0.52)
SMB							-0.41^{***}	-0.47^{***}	-0.47^{***}	-0.37^{***}	-0.37^{***}	0.04	−0.06	−0.14	−0.11	0.001	0.01	0.07
							(−3.91)	(−4.86)	(−4.84)	(−3.28)	(−3.69)	(0.67)	(−0.52)	(−1.29)	(−1.03)	(0.005)	(0.14)	(1.04)
HML							−0.03	0.03	−0.01	0.07	−0.001	0.03	0.01	0.07	0.03	0.11	0.04	0.03
							(−0.31)	(0.37)	(−0.13)	(0.66)	(−0.01)	(0.55)	(0.15)	(0.88)	(0.37)	(1.19)	(0.51)	(0.53)
UMD													−0.11	−0.02	−0.02	−0.05	−0.02	0.09^{**}
													(−1.63)	(−0.35)	(−0.39)	(−0.70)	(−0.30)	(2.17)
LIQ													-0.57^{***}	-0.53^{***}	-0.58^{***}	-0.59^{***}	-0.61^{***}	−0.03
													(−5.50)	(−5.43)	(−6.08)	(−5.16)	(−6.25)	(−0.52)
N	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96
Adj. R^2	0.39	0.40	0.41	0.37	0.38	−0.001	0.46	0.52	0.53	0.43	0.45	−0.02	0.59	0.63	0.66	0.55	0.61	0.02

This table reports 3-month sentiment quintile portfolio regression results. We construct quintiles based on the three previous months' sentiment scores, equally weighted. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. *t*-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A7.3: 3-Month Cosine Similarity Regressions

	Dependent variable: Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.19	0.30	0.29	0.30	0.43	0.24	0.33	0.44	0.43	0.41	0.56*	0.24	0.41	0.43	0.51*	0.42	0.51*	0.10
	(0.59)	(0.91)	(0.81)	(0.78)	(1.24)	(1.44)	(1.10)	(1.40)	(1.32)	(1.12)	(1.74)	(1.40)	(1.58)	(1.47)	(1.71)	(1.25)	(1.75)	(0.59)
RMRF	0.82***	0.75***	0.78***	0.73***	0.73***	−0.09**	0.72***	0.66***	0.69***	0.64***	0.63***	−0.09*	0.50***	0.48***	0.47***	0.41***	0.43***	−0.07
	(9.05)	(8.04)	(8.01)	(6.87)	(7.51)	(−2.01)	(8.51)	(7.40)	(7.38)	(6.12)	(6.83)	(−1.88)	(6.37)	(5.39)	(5.21)	(4.02)	(4.85)	(−1.31)
SMB							−0.44***	−0.39***	−0.41***	−0.43***	−0.43***	0.01	−0.08	−0.06	−0.05	−0.04	−0.06	0.02
							(−4.68)	(−3.90)	(−3.94)	(−3.79)	(−4.21)	(0.18)	(−0.85)	(−0.52)	(−0.49)	(−0.29)	(−0.52)	(0.39)
HML							0.003	0.04	0.06	−0.06	0.01	0.01	0.05	0.08	0.10	−0.01	0.05	0.001
							(0.03)	(0.48)	(0.62)	(−0.52)	(0.08)	(0.10)	(0.63)	(0.97)	(1.21)	(−0.12)	(0.57)	(0.03)
UMD													−0.09	−0.02	−0.09	−0.04	0.01	0.10***
													(−1.57)	(−0.30)	(−1.32)	(−0.53)	(0.16)	(2.63)
LIQ													−0.58***	−0.52***	−0.57***	−0.63***	−0.58***	−0.002
													(−6.50)	(−5.14)	(−5.54)	(−5.49)	(−5.75)	(−0.03)
N	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96
Adj. R^2	0.46	0.40	0.40	0.33	0.37	0.03	0.56	0.48	0.48	0.41	0.46	0.01	0.69	0.59	0.61	0.54	0.60	0.06

This table reports 3-month cosine similarity quintile portfolio regression results. We construct quintiles based on the three previous months' cosine similarity scores, equally weighted. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. *t*-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A7.4: 3-Month Jaccard Similarity Regressions

	Dependent variable:																	
	Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.15	0.31	0.28	0.26	0.51	0.37^{*}	0.29	0.45	0.41	0.39	0.63^{*}	0.33	0.37	0.48	0.44	0.40	0.59^{*}	0.21
	(0.43)	(0.88)	(0.85)	(0.73)	(1.42)	(1.84)	(0.93)	(1.37)	(1.32)	(1.18)	(1.82)	(1.65)	(1.35)	(1.62)	(1.53)	(1.34)	(1.82)	(1.02)
RMRF	0.77^{***}	0.77^{***}	0.82^{***}	0.72^{***}	0.74^{***}	−0.03	0.67^{***}	0.68^{***}	0.73^{***}	0.63^{***}	0.65^{***}	−0.02	0.44^{***}	0.45^{***}	0.54^{***}	0.42^{***}	0.45^{***}	0.01
	(8.06)	(7.77)	(8.84)	(7.36)	(7.30)	(−0.50)	(7.45)	(7.12)	(8.22)	(6.66)	(6.58)	(−0.33)	(5.24)	(4.98)	(6.14)	(4.56)	(4.63)	(0.22)
SMB							-0.44^{***}	-0.43^{***}	-0.42^{***}	-0.42^{***}	-0.39^{***}	0.04	−0.07	−0.04	−0.10	−0.05	−0.03	0.03
							(−4.41)	(−4.08)	(−4.34)	(−4.00)	(−3.60)	(0.69)	(−0.64)	(−0.34)	(−0.95)	(−0.44)	(−0.26)	(0.44)
HML							0.04	0.03	−0.01	0.02	−0.01	−0.05	0.08	0.08	0.02	0.06	0.02	−0.06
							(0.40)	(0.35)	(−0.15)	(0.16)	(−0.14)	(−0.85)	(1.04)	(0.94)	(0.29)	(0.67)	(0.27)	(−0.96)
UMD													−0.09	−0.05	−0.05	−0.04	0.002	0.09^{*}
													(−1.47)	(−0.76)	(−0.77)	(−0.58)	(0.02)	(1.96)
LIQ													-0.60^{***}	-0.62^{***}	-0.51^{***}	-0.58^{***}	-0.56^{***}	0.04
													(−6.31)	(−6.09)	(−5.19)	(−5.62)	(−5.07)	(0.50)
N	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96
Adj. R^2	0.40	0.38	0.45	0.36	0.36	−0.01	0.50	0.47	0.53	0.45	0.42	−0.01	0.65	0.62	0.63	0.58	0.54	0.01

This table reports 3-month Jaccard similarity quintile portfolio regression results. We construct quintiles based on the three previous months' Jaccard similarity scores, equally weighted. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. *t*-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Table A7.5: 3-Month Levenshtein Distance Regressions

	Dependent variable: Excess net return																	
	CAPM						3-factor model						5-factor model					
	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1	Q1	Q2	Q3	Q4	Q5	Q5-Q1
Alpha	0.25	0.05	0.13	0.02	0.39	0.14	0.40	0.19	0.25	0.18	0.48	0.08	0.45	0.24	0.30	0.19	0.43	−0.02
	(0.72)	(0.14)	(0.38)	(0.06)	(1.14)	(0.61)	(1.23)	(0.59)	(0.78)	(0.53)	(1.44)	(0.34)	(1.54)	(0.82)	(1.00)	(0.62)	(1.39)	(−0.08)
RMRF	0.72***	0.83***	0.82***	0.72***	0.73***	0.003	0.63***	0.74***	0.73***	0.61***	0.65***	0.02	0.40***	0.52***	0.53***	0.40***	0.45***	0.05
	(7.35)	(8.61)	(8.65)	(6.94)	(7.51)	(0.05)	(6.71)	(8.01)	(8.00)	(6.26)	(6.78)	(0.26)	(4.51)	(5.89)	(5.90)	(4.16)	(4.83)	(0.74)
SMB							−0.41***	−0.44***	−0.40***	−0.49***	−0.34***	0.07	−0.02	−0.09	−0.07	−0.12	0.03	0.04
							(−3.95)	(−4.36)	(−3.96)	(−4.51)	(−3.21)	(0.97)	(−0.17)	(−0.80)	(−0.61)	(−0.98)	(0.22)	(0.49)
HML							0.07	0.02	−0.001	0.02	−0.04	−0.10	0.11	0.06	0.04	0.07	0.002	−0.11
							(0.69)	(0.18)	(−0.01)	(0.23)	(−0.37)	(−1.52)	(1.36)	(0.71)	(0.45)	(0.73)	(0.02)	(−1.62)
UMD													−0.07	−0.07	−0.06	−0.04	0.01	0.08
													(−1.01)	(−1.04)	(−0.94)	(−0.58)	(0.16)	(1.42)
LIQ													−0.63***	−0.57***	−0.53***	−0.59***	−0.57***	0.06
													(−6.24)	(−5.66)	(−5.20)	(−5.43)	(−5.33)	(0.70)
N	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96	96
Adj. R^2	0.36	0.44	0.44	0.33	0.37	−0.01	0.45	0.53	0.51	0.45	0.42	0.01	0.61	0.64	0.62	0.57	0.55	0.01

This table reports 3-month Levenshtein distance quintile portfolio regression results. We construct quintiles based on the three previous months' Levenshtein distance scores, equally weighted. Q1 is the quintile with the lowest score, while Q5 is the quintile with the highest score. Portfolios are equally-weighted, and rebalanced on a monthly basis. We report CAPM results (market), three-factor results (market, size, value), and five-factor results (market, size, value, momentum, liquidity). The dependent variable, net return in excess of the risk-free rate, is multiplied by 100. All factors are multiplied by 100. *t*-Statistics are reported below the estimates. Statistical significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.