Norwegian School of Economics
Bergen, Fall 2020

Shareholder Letters on Oslo Børs

A textual analysis of shareholders letters from companies listed on Oslo Børs

Joakim Knudsen & Even Kolsgaard

Supervisor: Tommy Stamland

Master thesis, Economics and Business Administration

Major: Financial Economics

NORWEGIAN SCHOOL OF ECONOMICS

Acknowledgements

We would like to thank our supervisor Tommy Stamland for giving us helpful guidance and valuable feedback. We would also like to thank our friends and family for supporting us throughout the process of writing this master's thesis.

Norwegian School of Economics

Bergen, December 2020

Joakim Knudsen

Even Kolsgaard

Abstract

This study looks at the characteristics of the shareholder letter from companies listed on Oslo Børs and the relationship between sentiment and future financial performance. Regarding the characteristics, we have investigated the informativeness and the use of impression management by looking at sentiment development, similarity, writing style and the correlation between sentiment and past performance. Quintile portfolios sorted on the sentiment were created to investigate the relationship between sentiment and future performance. Our analysis shows that the shareholder letter usually follows the same structure and that the financial performance is reported honestly, to a certain degree. We also find evidence of impression management techniques, such as increased sentiment at the end of the letter. Regarding our second objective, do we not find any apparent relationship between the sentiment and future stock price. Therefore, we conclude that the shareholder letter may be a worthwhile read, but to blindly use the sentiment of the shareholder letter as an indicator of future performance, does not seem to be a good idea.

Keywords – Shareholder letter, textual analysis, sentiment analysis, impression management

Contents

List of Figures

List of Tables

1 Introduction

Textual analysis in finance is a relatively new and intriguing research field. With the development in computing power and available information, the field of textual analysis in finance has seen a dramatical growth since the early 2000s. However, most of the research focuses on corporate filings of companies listed in the United States, predominately because of the data availability. A much less researched area is corporate filings of companies listed on Oslo Børs. In this paper, we will contribute to the unexplored area of textual analysis on shareholder letters from companies listed on Oslo Børs.

Listed companies on Oslo Børs publish their annual report once a year. Of content, we often find the shareholder letter, which is a voluntary segment. In contrast to other parts of the annual report, the shareholder letter is written by one of the top executive's and do not follow any formal requirements. It is a segment where the management can convey whatever information they want to the readers. This autonomy creates endless possibilities on how the shareholder letter can be written, but are the shareholder letters that different?

The shareholder letter also contributes to the immense amount of information that is available to the investors. Since the 1960s, the financial market has become considerably more informative due to growing institutional ownership (Bai et al., 2016). The growth of information has given rise to an associated increase in observable data. A small portion of this data is numerical and easily interpretable by computers. However, most of the data, such as the shareholder letter, is textual data and have historically been challenging to use effectively. Computers, and their ability to handle data collection and data processing, have simplified the process. As a result, the field of textual analysis started to grow in the early 2000s, with the most prominent research produced in the last decade (Kearney and Liu, 2014).

One approach to exploit the textual information is to quantify the sentiment expressed in the text. This method is called sentiment analysis and determines if the text is positive, negative or neutral (Kearney and Liu, 2014). Based on the computed sentiment, it has been shown that it can be used to predict future firm performance (Che et al., 2020). The question is whether the same is true for shareholder letters in Norway. There is no doubt that the top executives have more knowledge about the company's performance than the typical investor. However, is this knowledge reflected in the shareholder letter's sentiment? More precisely, can the sentiment in the text be used to predict future performance?

Our thesis thus focuses on two main themes within textual analysis on shareholder letters from companies listed on Oslo Børs. The first one being the structure of the shareholder letter, where we look at sentiment development, the relationship between sentiment and past performance, text similarities and writing styles. The second theme is the prediction value of sentiment, and we investigate if it can be used to create excess returns.

The remainder of this thesis is organised as follows: Section 2 introduces the relevant literature surrounding textual and sentiment analysis as well as introducing our research questions. Section 3 describes the data used in our analysis. Section 4 presents the methodology as well as the limitations. In section 5, we provide the results from our analysis. Finally, section 6 concludes our main findings.

2 Background and related literature

To provide a brief overview of the fundamentals of textual analysis, we review related literature in this field. Further we introduce the relevant aspects of shareholder letters and impression management. At last, we present our research questions.

2.1 Sentiment and textual analysis in finance

Kearney and Liu (2014) claims that qualitative information analyzed by researchers in the field of textual analysis in finance comes from three sources: corporate disclosures, media articles, and internet messages. In our paper, we will focus exclusively on corporate disclosures. In this area of research, Feldman et al. (2008) investigated 10-Ks and 10-Qs filings using changes in the sentiment of the MD&A section. They found a significant relationship between changes in the tone and short window contemporaneous returns around the SEC filing date. Concerning research on shareholder letters, Boudt and Thewissen (2014) have provided inspiring insights. They hypothesize using a word's position in the text to determine the informativeness for future firm performance. Key findings from the paper are that a word's location contains information value, which increases the prediction accuracy for future firm performance. Additionally, Che et al. (2020) used the sentiment from CEO letters along with advanced machine learning, to anticipate corporate financial performance. Using a logistic regression approach, they predicted financial performance with an accuracy of 70,46%. Thus, concluding that information in CEO letters is a crucial factor for anticipating financial performance, and is a valuable area of research as it improves stakeholder's decision making.

In addition to sentiment being shown to be predictive of future performance, document similarity has also been shown with the same properties. Cohen et al. (2020) researched similarity in quarterly and annual filings from U.S. corporations. They found that changes to the language and construction of financial reports have substantial implications for firms' future returns and operations. Additionally, they obtained an alpha of over 22% yearly by creating a portfolio that shorted changers and bought non-changers.

2.2 Introducing the shareholder letter

Introducing the contents of shareholder letters is vital in understanding the motivation behind our research questions. The mandatory annual reports are formal, regulated and not exposed to large fundamental changes from one year to another, (Stanton and Stanton, 2002). This is true for most parts of the annual reports, the exception being the shareholder letters. Not subjected to regulations and voluntary, the shareholder letters may be the only part of the annual report where the top management can express themselves freely. As there are no guidelines for writing a shareholder letter, there are large differences in how they are written. Among subjects that appear multiple times, we often find a walkthrough of the past year with the highlights provided, share price development, the company's financial status, and outlook. Below, we show two excerpts from the shareholder letter of (Orkla, 2010, p. 4-5) and (Arcus, 2018, p. 8). They are good examples of what is being discussed. They mostly focus on their operations, as Arcus shows, but also discuss market events when this is appropriate, Orkla being a good example. Interestingly, the Orkla shareholder letter is impacted by negative wording, even though the financial performance seems to have improved.

Table 2.1: Examples of shareholder letters

Orkla, annual report 2009 (p. 4-5)	Arcus, annual report 2017 (p. 8)
Strengthened position in a challenging year. 2009 was a challenging year. It began in the shadow of the financial crisis, with plummeting economic activity and a recession in the US and Europe. (...) Increasing cash flow from operating activities and other measures to free up capital have been high priorities in 2009. All the business areas helped to boost cash flow from operating activities to NOK 5.8 billion, an improvement of NOK 3.3 billion on 2008. (...) The Orkla share ended the year at NOK 56.85, up from NOK 45.45 at the start of 2009. (...) The return on shareholders' investment for the year was thus 31%.(...) Experience has taught us that challenging economic conditions can also present opportunities for value-creating industrial initiatives and improvement measures. Orkla's financial position is strong, and the Group intends to fully exploit its frontier of opportunity in the future.	**2017: Record result for Arcus.** Arcus achieved its best result ever in 2017. Operating profit (adjusted EBITDA) was NOK 361 million, an increase of 8 per cent from 2016. Revenue was NOK 2,575 million, compared to NOK 2,582 million last year. (...) Based on the Group's financial results, Arcus' Board of Directors recommends a dividend of NOK 1.66 per share, an increase of 13 per cent from the previous year. The dividend rate will be approved by the annual general meeting on 11 April 2018. The growth plans for 2018 are in place. Activities in countries without monopolies will be strengthened. Early in the year, Gammel Opland, Norway's most popular aquavit, will be relaunched with a new bottle and design, aimed to strengthen sales in Norway and abroad. Our wine portfolio is continuously being renewed, to meet current and future trends

2.3 Impression management

Impression management in the annual report refers to "the process by which management convey their specification of reality through narrative, quantitative and visual disclosures in the annual report both to manage the corporate image and to report the performance of management" Cooper and Slack (2015, p. 802). Although the shareholder letter is not mandated, it is undoubtedly a vital part of the annual report. Personalized and signed, the letter is representing the top management's understanding of significant corporate events and achievements. The shareholder letter is a crucial sense-giving instrument (Conaway and Wardrope, 2010), capable of affecting the reader through selectivity and bias of the narrative's content (Leung et al., 2015). Among the observed strategies related to impression management is the use of personal pronouns. According to Wang et al. (2012), personal pronouns blurs the line between different stakeholders. Using personal pronouns, the letter fosters sympathy, builds corporate identity, and resonates better with the audience.

Another strategy in impression management is to structure the text strategically. Boudt and Thewissen (2019) hypothesized that CEO letters' narrative structure is used as a vehicle for impression management. In their report, they found strong evidence of a U-shape in the intertextual dynamics of the CEO's positive sentiment. Most notably, they discovered that net sentiment started low and ended high, referring to a "right-sided smirk" in their 2019 paper.

2.4 Research questions

A reoccurring subject is the non-regulated properties of the shareholder letter, and it begs the questions about the shareholder letter common characteristics. The excerpts from Orkla (2010) and Arcus (2018) showed that thematic and operational subjects are being discussed. They both had some differences in content, but there were obvious similarities. Operations were discussed in both letters, and the structural presentation of the content also shared some similarities. This leads us to the first research question.

Q1: What are the common characteristics of the shareholder letter?

The question is broad and general, which requires further limitation. Our focus will stay on the informativeness and aspects associated with impression management. The first area of interest is the development of sentiment. Boudt and Thewissen (2014) found evident characteristics in the sentiment development in CEO letters for companies listed on the Dow Jones Industrial Average (DJIA). To study whether there are strategic ways that information gets portrayed in the letter for Norwegian companies, we present our first sub-question:

SQ1: How does the sentiment develop through the shareholder letters?

Regarding the informativeness of the shareholder letter, there should be a connection between sentiment and performance. If the letter is informative, the reader will get an unbiased assessment of past performance. Therefore, sentiment should be positive when performance is good and negative when performance is poor. A lack of connection between performance and textual features may signal that shareholder letters are un-informative or that the letter is primarily used to influence the reader by presenting a biased overview.

To investigate the relation between past performance and sentiment, we present our second sub-question:

SQ2: Do the sentiment reflect past performance?

Delving further into shareholder letters' characteristics, we are interested in studying the uniqueness of the shareholder letter. If a shareholder letter is unique, then the textual content changes from year to year, and the similarity is low. On the contrary, using boilerplate text, the letter will display a high degree of similarity from year to year, which may reduce the overall informativeness of the text. This motivates us to investigate whether the companies are using boilerplate text, but also how similarity affects informativeness. If companies use boilerplate text, this may imply that the shareholder letter is merely a formal addition to annual reports, suggesting that the shareholder letter is not enclosing crucial information. Our third sub-question is thus:

SQ3: Do different companies use boilerplate text, and how does similarity affect the informativeness

Our last area of interest is how the writing style varies. In line with impression management, writing style can be used as a tool to affect the reader. Therefore, we are interested in how the writing style is affected by past performance. We present our final sub-question:

SQ4: How is the writing style affected by past performance?

After establishing the common characteristics of the shareholder letters, we want to analyze whether there exists any predictive power in the sentiment from shareholder letters. We do believe that top management is better informed compared to outsiders. This knowledge might be reflected in the shareholder letter. Given this, the sentiment could be used as a decision variable in an investment portfolio. Our second research question is presented:

Q2: Is the shareholder letter sentiment a viable decision variable for an investment strategy?

3 Data

This chapter describes the data used in our analysis. We start by explaining our sample selection and the retrieval of the shareholder letters. 3.2 gives a description of the company data and stock prices before we introduce dictionaries. At last summary statistics are presented.

3.1 Shareholder letters

Before detailing the sample selection, we must define what we mean by the shareholder letter. We define it as a letter written by one of the firm's top executives, that addresses the shareholder directly and provides information about the company's operations. This means that we include all sorts of CEO letters, but also letters from the chairman when there are no CEO letter. We do not consider similar text that is on behalf of the whole management group, because it is likely no one from the management group has written it.

The shareholder letters are extracted from the annual reports of the companies. Annual reports are downloaded from the respective company's website. Refinitiv Eikon is used to download annual reports if the company itself does not provide the complete set. Refinitiv Eikon is also used to download annual reports for delisted companies. We do not look at annual reports older than 1999, mainly because of the poor data quality and availability.

Our initial sample is every shareholder letter in every annual report for every company listed on Oslo Børs from 1999 to 2020. However, the shareholder letter is not required to be included in the annual report. It is up to each company to decide, which leads to a lot of different practices. Some companies always include a shareholder letter, some have never done it, and others write them sometimes.

We are also limiting this study to shareholder letters written in English. This limits the sample further since it is sufficient for Norwegian companies to submit the annual report in Norwegian, even though Oslo Stock Exchange encourages the companies to submit in both languages (Oslo Børs, 2019). It is possible to also include Norwegian shareholder letters in the sample, but to do textual analysis on two different languages can be problematic since the analysis depends on dictionaries, which differs from language to language. Trying to create two equivalent dictionaries would introduce a significant source of error.

The layouts of annual reports in Norway are not required to be standardized, like in the US, which leads to a lot of different layouts. The lack of standardization makes it hard to automize the data collection and ensure correctness. The Corporate Financial Information Environment developed a researcher tool for extracting text from annual reports in the UK (El-Haj et al., 2020). Layouts of annual reports in the UK varies in the same way as annual reports in Norway. However, the tool often fails, and it cannot ensure correctness even if it succeeds.

To ensure correctness, the shareholder letters are manually copied from each annual report and saved as a new text document. Only the text is saved, pictures and tables are excluded since our focus is on the text. The text is also looked through to ensure that it will be parsed correctly in later stages. Every sentence, including headings, is made sure is ending with a dot. This is done to make it easier to separate the text into sentences automatically.

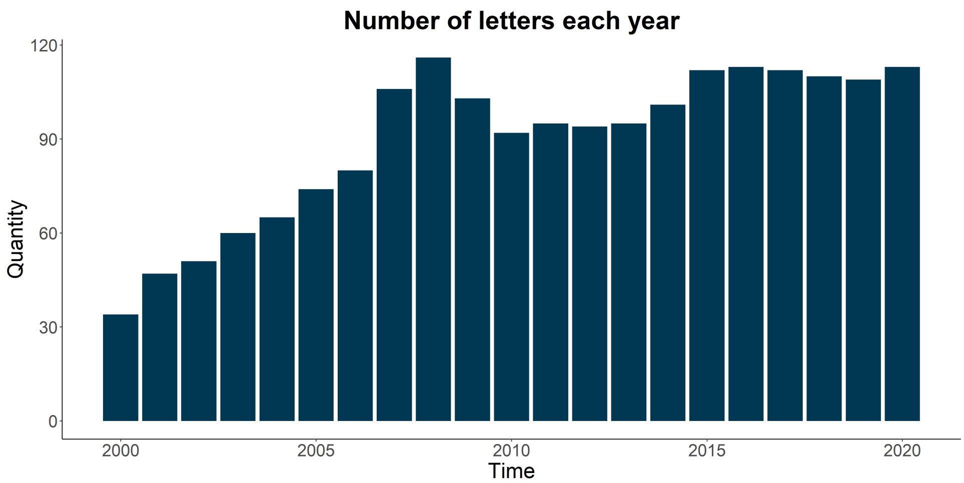

Figure 3.1: The number of shareholder letters each year

The resulting sample is of 1890 shareholder letters from 278 different companies listed on Oslo Børs in the period 1999-2019. Figure 3.1 shows the number of annual reports each year. Note that the x-axis is the year the annual report is published. It reveals that the number of shareholders at the start of the period is significantly lower than later, even though the number of shares listed on Oslo Børs has been relatively stable. One reason is that data availability is considerably worse at the start of the period. Some annual reports are lost, and others may be impossible to read due to the available copy's quality. Another reason is the share of annual reports written in English, which is higher today than it was 20 years ago. A final reason is that it has become more popular to write shareholder letters over the years, which figure 3.2 shows. It shows how many of the 20 largest Norwegian companies have a shareholder letter included in the annual report, and the trend has been positive. There is a significant decline in 2009 and 2010. The main reason for this decline is that the top 20 list changed, and the new companies did not have a shareholder letter in their annual report. Some companies did also stop writing the shareholder letter, which increased the decline.

Figure 3.2: The development of shareholder letters in the top 20 largest Norwegian companies listed on Oslo Børs

3.2 Company data and stock prices

Accounting data and filing dates are obtained from Refinitiv Eikon (2020), and all stock related information is obtained from Børsprosjektet (2020). The downloaded stock prices are adjusted for stock splits and dividends.

3.3 Dictionaries

Central to our analysis is the use of dictionaries. A dictionary in our context is a group of words that can be is assigned some common value. The most crucial dictionary used is the sentiment dictionary, which associates words with a sentiment. Since words can have different meanings in a different context, the choice of the sentiment dictionary could significantly impact the result, especially in a financial context (Loughran and McDonald, 2011). Therefore, have we chosen to use the sentiment dictionary made by Loughran and McDonald (2018). It is developed specifically for financial texts and is thus suitable for our needs. The dictionary introduces seven different sentiments, but we only use the positive and negative dictionary.

Another essential dictionary we use is the stop word dictionary. Stop words are words that do not add any significant meaning to a sentence. As with the sentiment dictionary, the choice of stop word dictionary could impact the results. We use the stop word dictionary made by Loughran and McDonald (2020), but we have modified it to exclude valance shifters. The modifications can be viewed in Appendix A.1

The last dictionary we use is for valence shifters. Valence shifters are lexical items that change the positive or negative attitude (valence) of a term (Polanyi and Zaenen, 2006). The most common valance shifter are negators, amplifiers, and de-amplifiers. Negators flip the valance of a term, amplifiers intensify it, and de-amplifiers reduces the valance. For example, in the sentence "This was not a good year", "not" is a negator and flips the sentence's valance from positive to negative. Adversative conjunctions are also a type of valence shifter that we use in our analysis. These are terms that express a contrast between two statements (EnglishGrammar, 2019). Examples of common adversative conjunctions are 'but', 'still', and 'however'. We use the valance shifter dictionary from the R-package sentiment. The complete list of valence shifters can be viewed in Appendix A.2.

3.4 Summary statistics

Our concern about the data sample variations is confirmed by the summary statistics that table 3.1 provides. The average sentiment is as expected positive, and the variance is large. A large variance is a favourable characteristic because it helps differentiate the letters. A less desirable trait is the significant variance in the word count. A difference of 5,219 words between the shortest and longest letter shows that there may be no standard length on these letters. The cosine similarity shows the same type of variance, with the minimum being close to zero and the maximum close to 1. These are extreme values, but they are correct. The minimum cosine similarity is possible because of the shortest letters. When there are so few words, it is possible to write a new one with almost entirely different words. The maximum may also seem too high, but some companies reuse the same text with minor changes. The similarity and sentiment change are only measured between shareholder letters from the same company in time *t* and time *t-1*. This explains why there are only 1 538 observations for these measures.

Sentiment change shows the absolute change in sentiment, and it is close to zero on average, implying that the sentiment might not change that much. One hundred twenty-five observations of sentiment change are between 0.01 and 0.01, which is a relatively small portion. Investigations of the distribution of sentiment change further reassure that the sentiment has a reasonable variability, even though the changes are not massive. However, small changes are expected as it is limited how many times the sentiment for a company should change drastically.

Table 3.1: Descriptive statistics for the shareholder letters

Measure	N	Mean	St. Dev.	Min	Max
Sentiment	1,890	0.184	0.108	−0.394	0.480
Sentiment change	1,538	−0.003	0.116	−0.473	0.413
Similarity	1,538	0.536	0.130	0.035	0.975
Word count	1,890	880	500	92	5,311

The table shows summary statistics for the sample. Sentiment change and similarity is dependent on a preceding text, which explains the lower amount of observations. Sentiment change is the absolute change in sentiment from one year to another.

4 Methodology

The following chapter describes the methodology in our analysis. We start by detailing the textual analysis steps, from pre-processing to the computation of different textual measures. This is followed by a description of the statistical test we rely on and quintile portfolios. Lastly, we discuss some limitations to the chosen methodology.

4.1 Textual analysis

4.1.1 Pre-processing

Pre-processing is the process of cleaning and preparing the text for further analysis (Haddi et al., 2013). A lot of the noise is already reduced since the texts have been extracted manually. As a result, the amount of pre-processing necessary is reduced, but there are still some steps that need to be done to quantify the textual data. We perform the following steps:

- **Split the text into sentences.** Each text is split into sentences to be able to get the context of the words.
- **Remove numbers.** All numbers are removed from the text since it does not add anything in textual analysis.
- **Convert all characters to lowercase.** The textual analysis is case sensitive, which makes this step important.
- **Stop words removal.** Stop words removal. Stop words are only removed before calculating the similarity. We do not remove them before the sentiment calculation because the impact on sentiment is debatable (Vallantin, 2019), and it can do more harm than good. We have also tried both possibilities, and the results are the same.

After the pre-processing steps are done, the different textual measures are ready to be computed.

4.1.2 Sentiment computation

We are using a score-based approach to quantify each text's sentiment, which means that the texts are processed through an algorithm and given a sentiment score. The R-package *sentiment* is used to compute this score. *sentiment* is a sentiment analysis tool that is highly customizable to suit the need of the user. The following paragraphs describe how the sentiment of a sentence is computed. The first step is identifying the positive and negative words in the sentence. Based on the sentiment dictionary introduced in chapter 3 are positive words identified and assigned the value 1, while negative words are assigned the value -1. Neutral words have a value of 0. The next step is to consider possible valance shifters in the sentence that can affect the sentiment of a word. This is done by forming subsets based on the positioning of the polarized words. Four words before and two after a polarized word forms a subset. Valance shifters in a subset will affect the score. An odd number of negators flip the sentiment, while an even number does nothing. Amplifiers increase the value while de-amplifiers decrease it, but the de-amplifiers' effect is limited to -1. Adversative conjunctions affect the score if it occurs either before or after a polarized word, weighing it up when it occurs before and down when it occurs after. At last, the sentiment value is divided by the square root of the number of words in the sentence.

> *With reference to our cost ambition, our 2017 performance was not good enough.* (Grieg Seafood, 2018, p. 2-3)

The first step is to locate the polarity words. 'good' is a positive word and given the value 1. This is also the only polarized word which results in one subset.

[2017,performance,was,not,good,enough]

The subset is looked through for valance shifters, and the only one is 'not'. 'not' is a negator, which means that the value changes from 1 to -1. Finally, the value is scaled by the length of the sentence.

$$\textit{Sentiment value} = \frac{-1}{\sqrt{12}} = -0.29$$

The procedure is applied to every sentence in a text and results in n sentiment scores where n is the number of sentences. An average is computed to get a sentiment value for the whole text. The average function we use downweighs sentences that have a score of zero. We use this average function because shareholder letters often contain much text with a neutral sentiment. When this is the case, the sentiment will get downplayed using a standard average. Downweighing neutral sentences could reflect the 'true sentiment' better. The complete sentiment calculation formula is available in Appendix A.3.

4.1.3 Similarity

Another textual measure we use in our analysis is a measure of similarity between two texts. There exist different methods to measure the similarity, but we have chosen to use the cosine similarity. It is efficient and the most widely used method in textual analysis (Li and Han, 2013).The cosine similarity between two texts is defined as:

$$sim(A,B) = \frac{A \cdot B}{||A||\,||B||}$$

$$A = \textit{Vector of term frequency in text A}$$

$$B = \textit{Vector of term frequency in text B}$$

The term frequencies express how many times each term appears in the text, and the cosine similarity is a measure of the angle between the two vectors. The result is between zero and one, where zero means that the two texts do not share any words, and one indicates that the two texts are identical.

To better understand how the similarity score works, consider the two following statements from the CEO of Medistim, Kari Eian Krogstad, in the annual report of 2018 and 2019[1]

> *In 2010, TTFM was included in the European guidelines for coronary revascularization.* (Medistim, 2019, p. 2-3)

> *TTFM is a proven technology and was included in the European guidelines for coronary revascularization in 2010.* (Medistim, 2020, p. 2-3)

The full set of words is:

[1]The example sentences are not cleaned for stop words and numbers, as they are in our analysis. This is just to make the example easier to follow.

[in,2010,TTFM,is,a,was,included,proven,technology,and,the,European,guidelines, for,coronary,revascularization]

Term frequencies is therefore as follows:

$$A = [2,1,1,0,0,1,1,0,0,0,1,1,1,1,1,1]$$

$$B = [2,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1]$$

The result is 0.86 and expresses a high degree of similarity. The same procedure is conducted for each shareholder letter for company i at time t and compared to company i at time *t-1*. This means that the cosine similarity for "Medistim 2019", is the similarity between the shareholder letter from 2019 and 2018.

4.1.4 Writing style

The final textual measure that we use in our analysis is measures of writing style. Hillert et al. (2016) define how to classify a text as personal or not, which we adapt. If a text is written in first person singular or plural, it is classified as personal. The requirement is that the fraction of first-person singular pronouns (I, me, myself, my, mine) and first-person plural pronouns (our, ours, ourselves, us, we) in a text must be higher than the median of the full sample. We also look at the direct percentage of personal and plural pronouns in the texts.

4.2 Correlation tests

Spearman's rank correlation coefficient is computed when testing for correlations. It does not make any assumption of the distribution, which is essential because the sentiment scores are not normally distributed. It is a rank correlation, which means it looks at the correlation between the rankings between two variables, and not the variables themselves (Corder and Foreman, 2009). This is not a problem for our analysis since the rank may be more interesting than the direct relationship.

4.3 Quintile portfolios

Quintile portfolios is a portfolio approach where the assets are first sorted based on some characteristic, and then the top 20% is chosen as the portfolio (Zhou and Palomar, 2020). In our case, the main sorting characteristic is the sentiment score. Hence, we first take all annual reports that are published between time t and *t-1*. The letters are then sorted based on the sentiment score, and the stocks corresponding to the 20% most positive letters are chosen as the portfolio. We want the strategy to consider the newest available information at all times, and since the publication date varies for each company and year, the portfolio is rebalanced each month. As an example, consider a portfolio that starts 01.01.2020. All annual reports published between 01.01.2019 and 01.01.2020 are sorted based on the sentiment of the letter. The stocks corresponding to the 20% most positive letters are held for one month before the portfolio is updated to consider annual reports published in the holding period.

After deciding which stocks to keep in the portfolio, it must be decided how the different stocks should be weighted. We use two different weighting methods, equal weighting and weighting based on the inverse proportion of the stock volatility. The reason to use multiple weighting approaches is that the choice of weighting scheme can heavily influence the results (Plyakha et al., 2014). Using multiple schemes is going to give a better understanding of the validity of the results.

To get a better perception of the portfolio's performance, are the results also controlled for known factors. The portfolios are tested against two different factor models: the market model and the Fama-French three-factor model (FF3). The rationale behind adding the two additional risk-factors, small minus big (SMB) and high minus low (HML), is that value stocks tends to outperform growth stocks and small-cap stocks tend to outperform large-cap stocks (Fama and French, 1993).

The following equations describe the models:

$$r_t^i - rf_t = \alpha_t^i + \beta_{MKT}(r_t^{MKT} - rf_t) + \epsilon_t$$

$$r_t^i - rf_t = \alpha_t^i + \beta_{MKT}(r_t^{MKT} - rf_t) + \beta_{SMB}SMB_t + \beta_{HML}HML_t + \epsilon_t$$

where:

α_t^i : Excess return on portfolio i at time t

r_t^i : The return on portfolio i at time t

rf_t : The risk-free rate at time t

r_t^{MKT} : Market return at time t

SMB_t : Size premium (small minus big)

HML_t : Value premium (high minus low)

$\beta_{HML,MKT,SMB}$: factor coefficients

Pricing factors and risk-free rates are obtained from Ødegaard (2019). The risk-free rates are monthly, forward-looking estimates based on Norwegian government securities and NIBOR. Oslo Børs All Share Index (OSEAX) is used as the proxy for the market since all the companies in our sample are listed on Oslo Børs, making OSEAX a natural choice. SMB and HML are calculated according to Fama and French (1998).

4.4 Limitations

Quantifying the sentiment of a sentence with a high degree of accuracy is difficult. There can be many interdependencies in a sentence that are hard to get right. Our approach does consider valance shifters, but there are still other lexical items that complicate the sentiment computation. For example, sentences where the sentiment is dependent on the preceding sentence. However, our approach has been used on similar data before, and it appears to be a viable method.

Another limitation is that our method does not consider the context when assessing the sentiment. Our method reports the sentiment in each sentence but disregards what is being discussed. This is an evident limitation because it limits the information we are able to extract. For instance, if a CEO discusses a recent acquisition in a neutral sentiment, then this may be a negative signal. However, in our method, it is not recognized.

5 Analysis

In this chapter, we will describe the results of our analysis. The first section analyzes the structure and content of the shareholder letters, where we try to answer our first question about the characteristics of the shareholder letters. The second section looks at the shareholder letters' prediction value by looking at different quintile portfolios sorted on the sentiment value.

5.1 The informativeness and structure of shareholder letters

In this first part, we are going to look at the informativeness and structure of shareholder letters. The analysis is broken into smaller parts, and the structure follows our sub-question presented in chapter 2. The first sub-section looks at how the sentiment develops through the texts. This is followed up by an examination of the correlation between past performance and the sentiment score. At last, both similarity and writing style are analyzed.

5.1.1 Sentiment development

There seem to be common characteristics in how the sentiment develops through the shareholder letters. Figure 5.1a shows the average development of sentiment for the whole sample, and it shows that the average sentiment tends to increase towards the end of the letter. One reason for this may be that most letters start by addressing the past before ending the letter with some thoughts about the future, and the increased sentiment at the end may indicate the author's optimism. This sentiment development is the same as Boudt and Thewissen (2014) found for shareholder letters of a DJIA company sample.

Dividing the sample into gainers and losers supports the theory of the writer's optimism about the future. Figure 5.1b shows the sentiment development of letters where the company's stock price increased 20% or more and when the stock price decreased by 20% or more. We observe that the starting sentiment is lower when it has been a rough year compared to a positive year, which is evidence that the start of the letter is used to describe the past year. In comparison, the sentiment at the end is almost the same for both groups, indicating that the letters may have an overly optimistic outlook.

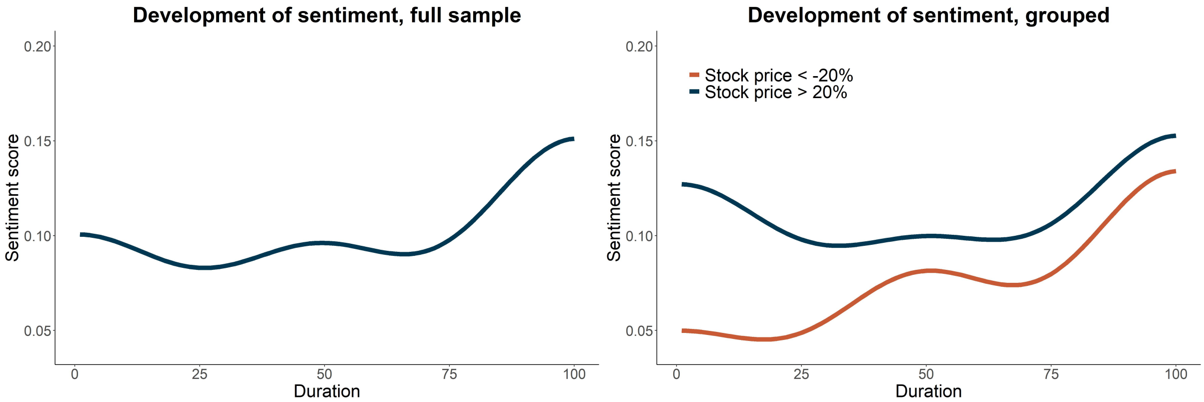

(a) Full sample **(b)** Sample grouped by stock gain

Figure 5.1: Figures of the sentiment development. Full sample and grouped by stock price movement

5.1.2 Sentiment and past performance

Spearman's rank correlation is computed between the sentiment score and a set of financial ratios, as well as the stock return, corresponding to the year the annual report describes. Table 5.1 shows the coefficients. Seven of the nine coefficients are statistically significant, and the sign for these coefficients are also as expected. For instance, a higher sentiment score is correlated with a higher ROE, which is reasonable. A higher ROE is an indication of better financial performance and, subsequently, more positive reporting.

Table 5.1: Correlation coefficient for correlation between the sentiment value and financial measures

Measure	Coefficient
Asset turnover	**0.174**
Assets/Equity	0
Cash cycle	0
Debt/equity	**-0.131**
Operating margin	**0.167**
Reinvestment rate	**0.176**
ROA	**0.237**
ROE	**0.242**
Stock price movement	**0.225**

The table shows the computed rank correlation between the sentiment score and the financial measure. Coefficients in bold are statistically significant.

Table 5.2: Correlation coefficients for correlation between change in sentiment and change in financial measures

Measure (change)	Coefficient
Asset turnover	0.052
Assets/Equity	**-0.081**
Cash cycle	**-0.061**
Debt/Equity	**-0.081**
Operating margin	**0.204**
Reinvestment rate	**0.147**
ROA	**0.196**
ROE	**0.181**
Stock price movement	**0.193**

The table shows the computed rank correlation between the change in sentiment score and the change in a financial measure. Coefficients in bold are statistically significant.

Similar results are obtained when correlating change in the measures with the change in the sentiment. Table 5.2 reports these results, and we observe some differences. Most notable is that both Assets/Equity and Cash cycle now have a significant negative correlation, the opposite of our expectations. Even though the coefficients are significant, they are small and less significant than the other significant coefficients.

The correlation coefficients indicate that the shareholder letter contains information about the previous year and that it is reported honestly to a certain degree. Even when dividing the sample into gainers (>20%) and losers (<-20%), remains the coefficients strength for both groups, indicating that the sentiment expressed is honest. However, the rank correlation between sentiment and stock gain differs between the two groups. The coefficient loses its significance for the gainers, but it remains the same for the losers, implying that the rank relationship disappears within the gainer group. In the loser group, the sentiment is still correlated with the stock gain. This means that companies experiencing a stock price fall close to 20% have a more positive sentiment than the companies with stock prices falling more, which seems reasonable. A stock price movement of -20% is not ideal, but it is also not too bad and may be possible to avoid addressing in the shareholder letter. However, if the stock price falls, for example, 50%, something is wrong and most likely discussed in the letter.

Dividing the letters into four parts of equal lengths increases the support of our perceived letter structure. Figure 5.2 shows how the correlation coefficient varies for the different parts of the letters. The figure shows that the first part has a larger correlation coefficient with past performance. It should be noted that some of the coefficients are still statistically significant for the last part. One contributing factor is that a shareholder letter does not have to end with an outlook and could just as well end with a summary of the year.

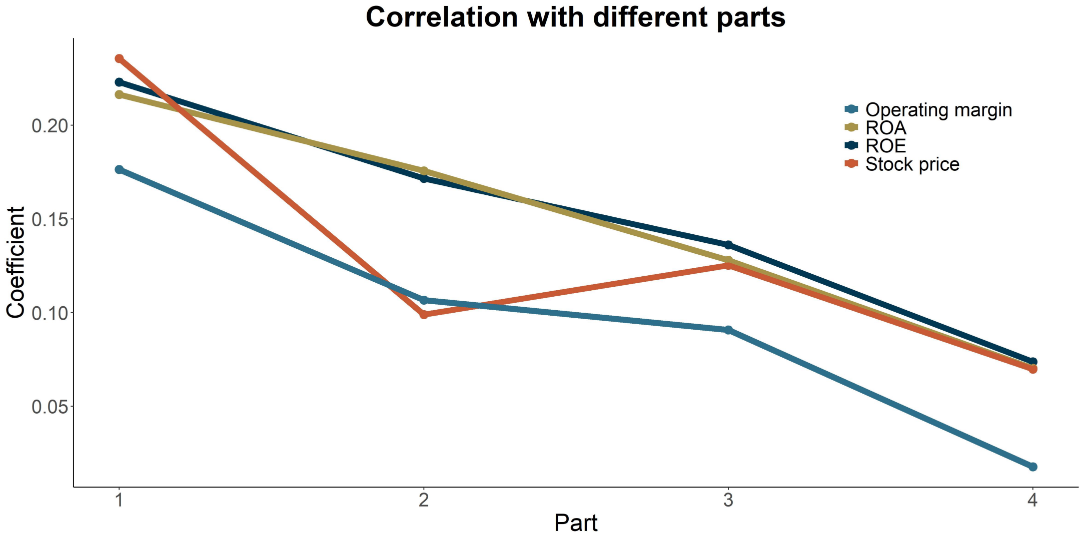

Figure 5.2: Correlation coefficient for different parts of the shareholder letter

Although a significant correlation, there are some letters that do not follow this pattern. One example is the 101 letters where the stock price fell more than 50%, even though the first two parts of the text are positive. It could be that the sentiment scores are wrong, but a quick read-through does not support this conjecture. For example, consider the shareholder letter from the annual report of SeaBird Exploration for 2017. In this letter, the CEO discuss changes in the market and opportunities for the company. It does not give the impression that it went too well in 2017, but it certainly does not give the impression that the stock price fell 98%. The problem is that there are no quantifiable measures that can be used to identify such letters. Vague descriptions and focus on new opportunities are typical characteristics but challenging to identify and quantify automatically. Consequently, extracting clear signals about the future can prove challenging.

5.1.3 Similarity

A high degree of similarity between subsequent letters may indicate the use of boilerplate text. However, there is no pre-defined level of similarity that would automatically indicate this. An approximation to this is to look at the similarity when a new person is writing the company's shareholder letter, e.g., a new CEO. When this is the case, we would expect a less similar text. If the similarity stays the same as before, it can be an indication of boilerplate text. The mean cosine similarity for letters written by the same person is 0.544 and 0.490 when it is not, with the difference being statistically significant. Evidentially, the shareholder letter is affected by the author. Further investigations of the sample reveal that few companies rely on boilerplate texts, but the letter's structure tends to be the same.

Arranging the texts based on similarity changes little when looking at the sentiment development and correlations with past performance. The sample is divided into three groups based on the similarity quantiles. Group one is the first quantile with the least similar texts, group two is the second and third quantile, and the last group is the fourth quantile. One minor difference is observed between group one and the two others. The sentiment development is slightly different, where the least similar has a lower starting sentiment and a slightly higher ending sentiment. Hence, the starting sentiment is lower for the least similar texts, indicating that the letter is more likely to change if it has been a rough year. Further investigation of the correlation between similarity and performance measures confirms this conjecture. A higher similarity is correlated with better firm performance, and the correlations are statistically significant.

5.1.4 Writing style

The average sentiment value between personal letters and formal letters are small and not significant. However, letters that use a higher proportion of personal pronouns have a higher sentiment value. This connection between the use of personal pronouns and sentiment value is further confirmed when looking at the direct correlation. A higher proportion of personal pronouns is correlated with a higher sentiment. The same is also true for the correlation with financial measures, and the relationships remain the same when looking at changes and not the absolute values. These correlations could indicate that the writer is comfortable using more personal pronouns in good times while more hesitant when the performance has been less satisfactory, maybe to avoid mentioning his role in such times.

This theory is further strengthened when investigating the correlation between the use of plural pronouns and financial measures. The proportion of plural pronouns is negatively related to most of the financial measures, indicating that plural pronouns are often used more when performance have been poor. This connection is also present when looking at the change in the use of plural pronouns and change in the financial measures.

5.2 Quintile portfolios

Shareholder letter sentiment seems to correlate with past performance, begging the question whether it is possible to predict future performance. This section investigates the prediction power of the shareholder letter by using the sentiment value as the investment criteria. We first look at quintile portfolios sorted on the full sentiment and change in sentiment. At last, we look at alternative sorting schemes based on our knowledge of the shareholder letter.

5.2.1 Sentiment portfolios

We start by creating six portfolios sorted on the sentiment value and six portfolios sorted on the relative change in sentiment. Portfolio Q5 is composed of the stocks corresponding to the top 20%, Q1 for the bottom 20%, and LS is long Q5 and short Q1. The Q5, Q1 and LS portfolio are created with both equal-weighting and weighting based on the inverse of stock volatility. All portfolios start in 2003 to ensure a large enough sample size to get a sizeable portfolio of stocks.

Table 5.3: Monthly return for portfolios sorted on sentiment

	Weighting scheme	
Portfolio	EW	IVW
Q5	0.926	0.614
Q1	0.798	0.017
LS	-0.136	0.373
Q5>Q1	t=0.194	t=1.067

The table shows the monthly return of the portfolios. Q5 is the portfolio with the highest sentiment, Q1 the lowest and LS is long Q5 and short Q1. The t-value is from a Welch's t-test, testing whether the Q5 portfolio have significantly higher monthly returns than the Q1 portfolio.

Table 5.4: Monthly return for portfolios sorted on sentiment change

	Weighting scheme	
Portfolio	EW	IVW
Q5	0.661	0.372
Q1	0.386	-0.306
LS	0.120	0.426
Q5>Q1	t=0.397	t=1.168

The table shows the monthly percentage return of the portfolios. Q5 is the portfolio with the highest sentiment, Q1 the lowest and LS is long Q5 and short Q1. The t-value is from a Welch's t-test, testing whether the Q5 portfolio have significantly higher monthly returns than the Q1 portfolio.

Table 5.3 shows the monthly return for the different portfolios when sorting on the sentiment value. It shows that investing in the equal-weighted Q5 portfolio has given a monthly return of 0.926% since 2003. On the other hand, the inverse volatility-weighted has only produced a monthly return of 0.614%, which shows how vital the weighting scheme is for the return. The difference between Q5 and Q1 is positive for both weighting schemes, but the difference is not statistically for either. Furthermore, the LS portfolios are noteworthy. The return is negative for equal-weighting, even though Q5 outperforms Q1 in the long run. This is because the Q1 portfolio has two periods with exceptional gain, making the LS portfolio suffer, even though the Q1 portfolio has overall lower performance. For the inverse volatility-weighting is the difference in performance between Q5 and Q1 more stable across the whole period, and the LS is positive.

Results are similar when sorting based on the relative change in sentiment, as table 5.4 shows. The Q5 portfolio still outperforms Q1, and the differences between the weighting schemes remain. One interesting difference is that the return is lower for both Q5 and Q1, whereas the LS portfolio performs slightly better.

Controlling the sentiment portfolio returns for known factors shows that none of the portfolios can create a significant positive alpha, as can be seen in table 5.5 and 5.6. Table 5.5 shows the alphas when the portfolio is sorted on sentiment, and table 5.6 shows when it is sorted on the relative change in sentiment. The Q5 portfolio creates a negative alpha in each scenario, and it is even significant when sorting on relative change and controlling for the FF3. Strong evidence that using the positive sentiment in a text is not a good predictor of future stock performance.

Table 5.5: Alphas for portfolios sorted on sentiment

	Alphas for different portfolios			
	Market		FF3	
	EW	IVW	EW	IVW
Q5	-0.026	-0.211	-0.272	-0.441*
	(0.281)	(0.253)	(0.286)	(0.258)
Q1	-0.236	-0.749**	-0.636*	-0.941***
	(0363)	(0.351)	(0.364)	(0.359)
LS	0.012	0.339	0.165	0.300
	(0.431)	(0.415)	(0.448)	(0.429)

The table shows the monthly alphas for different portfolios sorted on sentiment when controlled for the market model and the FF3 model. Standard deviation is reported in parentheses. All numbers are in percentages.

Table 5.6: Alphas for portfolios sorted on sentiment change

	Alphas for different portfolios			
	Market		FF3	
	EW	IVW	EW	IVW
Q5	-0.380	-0.361	-0.813***	-0.618*
	(0.305)	(0.347)	(0.293)	(0.355)
Q1	-0.613*	-1.222***	-0.981***	-1.441***
	(0.341)	(0.243)	(0.343)	(0.246)
LS	0.035	0.663*	-0.032	0.625
	(0.389)	(0.383)	(0.404)	(0.400)

The table shows the monthly alphas for different portfolios sorted on sentiment change when controlled for the market model and the FF3 model. Standard deviation is reported in parentheses. All numbers are in percentages.

The Q1 portfolios create a significant negative alpha in almost every scenario. The equal-weighted portfolio sorted on sentiment controlled for the market is the only Q1 alpha that is not significant, but it is negative. Interestingly is the significance greater for the Q1 portfolios that are sorted on the relative change in sentiment. There are no apparent reasons for this, but we see that also the Q5 portfolios have larger coefficients when sorting on the relative change in sentiment. This raises the question if the absolute change, positive or negative, can be used as a sorting criterion. It turns out it does not work any better than the other ones, and no new insight is provided.

These results do not give any clear indication of how the sentiment can be used to allocate a portfolio. The analysis does show evidence that negative sentiment is better at predicting negative stock price movements than positive is for predicting positive movement. Suggesting that negative sentiment or negative change in sentiment can be used to short such stocks. However, the Q1 portfolios had some periods with exceptional gain, for instance, in 2006 and 2007, which would make shorting the Q1 portfolio costly. Additionally, some of the Q5 portfolios do also have significant negative alphas, even though the coefficients are smaller compared to the Q1 portfolios. This weakens our confidence in how viable the short Q1 strategy is. Taking a short position in Q1 must also be controlled for the other factors if the negative alpha is to be utilized.

5.2.2 Alternative sorting schemes

One modification to the previous portfolios is only to consider the part of the letter where the outlook is most likely discussed, hence only using the sentiment in the last part of the shareholder letter. Twelve new portfolios are created, six based on the absolute sentiment value and the other six based on the relative change in sentiment, with both groups only considering the 25% last sentence of the shareholder letter. Table 5.8 show the alphas when sorting on the sentiment change in the final part of the letter. We observe that the results are similar to what we have seen before, and none of the alphas are positive. Table 5.7 shows when the absolute sentiment value in the final part of the letter.

One significant difference is the relationship between Q5 and Q1, which is now the opposite. When the whole letter was considered, the Q5 portfolios performed better than Q1. Whereas now, when only the last part is considered, then Q1 outperforms Q5. One interpretation is that companies with the most positive outlooks are way too optimistic, and such overly optimistic outlooks can act as a red flag. This is also the case for the relative change in sentiment, where the portfolio with the most positive sentiment change did worse than the portfolio with the most negative change in sentiment. Indicating that huge positive shifts in the sentiment may not be a particularly positive signal. This also related to being overly optimistic.

Table 5.7: Alphas for portfolios sorted on the ending sentiment

	Alphas for different portfolios			
	Market		FF3	
	EW	IVW	EW	IVW
Q5	-0.396 (0.336)	-0.475** (0.287)	-0.818*** (0.333)	-0.822*** (0.286)
Q1	-0.024 (0.325)	-0.129 (0.320)	-0.447 (0.321)	-0.413 (0.325)
LS	-0.571 (0.416)	-0.545 (0.392)	-0.571 (0.435)	-0.609 (0.409)

The table shows the monthly alphas for different portfolios sorted on sentiment, in the final part of the text, when controlled for the market model and the FF3 model. Standard deviation is reported in parentheses. All numbers are in percentages.

Table 5.8: Alphas for portfolios sorted on change in ending sentiment

	Alphas for different portfolios			
	CAPM		FF3	
	EW	IVW	EW	IVW
Q5	-0.432 (0.327)	-0.612** (0.284)	-0.853*** (0.322)	-0.773*** (0.293)
Q1	-0.121 (0.336)	-0.359 (0.313)	-0.458 (0.340)	-0.568* (0.322)
LS	-0.510 (0.417)	-0.452 (0.452)	-0.594 (0.435)	-0.404 (0.444)

The table shows the monthly alphas for different portfolios sorted on sentiment change, in the final part of the text, when controlled for the market model and the FF3 model. Standard deviation is reported in parentheses. All numbers are in percentages.

Other schemes with different textual measures are also tried but do not generate any significant alpha nor any new insights. Similarity showed a correlation with past performance, but the prediction value is not existent, and the alphas are insignificant and negative for every portfolio. Transformations of the sentiment values, restricting the sample based on similarity and writing style and using other parts of the texts are also tried. We do not find alphas of larger absolute value in any of these additional strategies.

6 Conclusion

Our first objective was to determine the typical characteristics of the shareholder letter. Looking at the sentiment development showed that the ending sentiment was higher than the starting sentiment, an indication of an overly optimistic outlook. Our results are thus similar to Boudt and Thewissen (2019), who coined this observation a "right-sided smirk" in intertextual development. Further, the relationship between sentiment and past performance indicated that the performance is reflected in the sentiment, suggesting that results are discussed and reported, to a certain degree, honestly. It also confirmed that most companies structure the shareholder letter by starting with thoughts on the past, before ending with a future outlook. With regards to similarity, we concluded that boiler-plate text is seldom used, but the shareholder letter of a company tends to follow the same structure each year. The analysis of writing style suggested that the writer tends to use more personal pronouns when things went well, and plural pronouns when it did not.

The sentiment letter thus seems to have, on average, some common characteristics. Findings of shared characteristics are not surprising, because it is natural that it follows some industry standard. Regarding the content, the letter seems to be quite informative, but with so much available information, it would be difficult to lie about the performance without being caught. One possibility would be to not talk about the performance at all, which a minority have done. With this in mind, it is not surprising that the ending sentiment seemed to be too high. Future speculation cannot be fact-checked, and with no consequences of being overly-optimistic, the writer can be as positive as he wants. By affecting the reader through the bias of the narrative's content, this poses as an impression management technique used to convince the reader of the companies positive future outlook.

The second objective was to investigate whether the sentiment can be used to generate excess returns. Creating quintile portfolios and regressing the results on known factors uncovered that we were not able to create a significant positive alpha, neither for absolute sentiment nor change in sentiment. However, we did obtain negative alpha values. Most notably for the Q1 portfolios, indicating that the lower sentiment may have an information value in predicting negative results. However, a simple short Q1 strategy would not be profitable; the other factors of the portfolio must be considered. Only when the "sentiment-factor" is isolated, can the alpha be exploited. We obtained similar results when only using the last part of the letter, but with one major difference; now it was the Q5 portfolio that had significant negative alphas—indicating that the most optimistic outlooks may be a bad sign. Other strategies were also tried, but none were successful.

These results are understandable when the characteristics of the shareholder letter are considered. The variety in the data sample is large, which makes it harder to use the shareholder letter as a framework to predict future performance. Additionally, the overly optimistic outlook inflates the sentiment value and makes it challenging to use it.

Regarding the significant negative alphas, there are some concerns about its validity. For the total sentiment, are there significant negative alphas for both Q1 and Q5, which begs the question whether something else than the sentiment is responsible for the alphas. Only when using the sentiment for the last part of the letter is the negative alpha consistent in the Q5 portfolio and only significant in one Q1 portfolio.

To conclude, the shareholder letter seems to be a useful read to get an overview of the company's performance. However, the matters discussed should be taken with caution, and the whole annual report should be considered when evaluating the future of the company. To blindly trust the shareholder letter about the future does not seem like the best idea.

In addition to the limitations discussed in chapter 4, there are also some limitations regarding our data. Since the shareholder letter is a voluntary segment of the annual report, we do get a selection bias. A selection bias may distort our analysis by excluding data for some companies. Using a different segment in the annual report, which is mandatory for all companies would mitigate this limitation. Using the Board of Director's report instead would pose as an interesting research field for future studies. However, this is a more formal part, lacking variation, which may make sentiment analysis more challenging. Another limitation of the study is that we have not included trading costs. However, including it would not have any impact on the overall conclusions.

More data could also be included to improve the analysis. First of all, considering additional statements from the top management, in the news and other corporate filings, could be an improvement. Another interesting addition would be to include the author's stake in the company and see if "skin in the game" could improve the viability of the sentiment as an investment criterion. As a final note, textual analysis is an emerging field, and there may be developed new methods for sentiment analysis which can be better suited to answer our research questions.

References

Arcus (2018). Annual report 2017. https://s3-eu-west-1.amazonaws.com/arcus-upload/\%C3\%85rsrapport-2017.pdf.

Bai, J., Philippon, T., and Savov, A. (2016). Have financial markets become more informative? *Journal of Financial Economics*, 122(3):625–654.

Boudt, K. and Thewissen, J. (2014). Not all words are equal: Sentiment dynamics and information content within ceo letters. *Available at SSRN.*

Boudt, K. and Thewissen, J. (2019). Jockeying for position in ceo letters: Impression management and sentiment analytics. *Financial Management*, 48(1):77–115.

Børsprosjektet (2020). Amadeus. https://bors.nhh.no/amadeus/index.php.

Che, S., Zhu, W., and Li, X. (2020). Anticipating corporate financial performance from ceo letters utilizing sentiment analysis. *Mathematical Problems in Engineering*, 2020.

Cohen, L., Malloy, C., and Nguyen, Q. (2020). Lazy prices. *The Journal of Finance*, 75(3):1371–1415.

Conaway, R. N. and Wardrope, W. J. (2010). Do their words really matter? thematic analysis of us and latin american ceo letters. *The Journal of Business Communication (1973)*, 47(2):141–168.

Cooper, S. and Slack, R. (2015). Reporting practice, impression management and company performance: a longitudinal and comparative analysis of water leakage disclosure. 45(6-7):801–840.

Corder, G. W. and Foreman, D. I. (2009). *Nonparametric statistics for non-statisticians: A step by step approach.* John Wiley & Sons, Inc.

Eikon (2020). Refinitiv Eikon [Online]. Available at: Subscription Service (Accessed: October 2020).

El-Haj, M., Alves, P., Rayson, P., Walker, M., and Young, S. (2020). Retrieving, classifying and analysing narrative commentary in unstructured (glossy) annual reports published as pdf files. *Accounting and Business Research*, 50(1):6–34.

EnglishGrammar (2019). Kinds of co-ordinating conjunctions. https://www.englishgrammar.org/kinds-coordinating-conjunctions/.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33:3–56.

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

Feldman, R., Govindaraj, S., Livnat, J., and Segal, B. (2008). The incremental information content of tone change in management discussion and analysis.

Grieg Seafood (2018). 2017 annual report. https://investor.griegseafood.com/reports-&-presentations.

Haddi, E., Liu, X., and Shi, Y. (2013). The role of text pre-processing in sentiment analysis. *Procedia Computer Science*, 17:26–32.

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)*.

Kearney, C. and Liu, S. (2014). Textual sentiment in finance: A survey of methods and models. *International Review of Financial Analysis*, 33:171–185.

Leung, S., Parker, L., and Courtis, J. (2015). Impression management through minimal narrative disclosure in annual reports. *The British accounting review*, 47(3):275–289.

Li, B. and Han, L. (2013). *Distance weighted cosine similarity measure for text classification.*

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. (2018). Loughranmcdonald_sentimentwordlists_2018 [data set]. https://drive.google.com/file/d/15UPaF2xJLSVz8DYuphierz67trCxFLcl/view.

Loughran, T. and McDonald, B. (2020). Stopwords_genericlong.txt [data set]. https://drive.google.com/file/d/0B4niqV00F3msSktONVhfaElXeEk/view.

Medistim (2019). Annual report 2018. https://medistim.com/2019/03/annual-report-2018/.

Medistim (2020). Annual report 2019. https://medistim.com/2020/03/annual-report-2019/.

Orkla (2010). Annual report 2009. http://hugin.info/111/R/1512328/447863.pdf.

Oslo Børs (2019). The oslo børs code of practice for ir. https://www.oslobors.no/ob_eng/Oslo-Boers/Listing/Shares-equity-certificates-and-rights-to-shares/Oslo-Boers-and-Oslo-Axess/Code-of-Practice-for-IR.

Plyakha, Y., Uppal, R., and Vilkov, G. (2014). Equal or value weighting? implications for asset-pricing tests. *Implications for Asset-Pricing Tests (January 15, 2014)*.

Polanyi, L. and Zaenen, A. (2006). *Contextual valence shifters*. Springer.

Stanton, P. and Stanton, J. (2002). Corporate annual reports: research perspectives used. *Accounting, Auditing & Accountability Journal.*

Vallantin, L. (2019). Why is removing stop words not always a good idea. https://medium.com/@limavallantin/why-is-removing-stop-words-not-always-a-good-idea-c8d35bd77214.

Wang, H., Li, L., and Cao, J. (2012). Lexical features in corporate annual reports: a corpus-based study. *European Journal of Business and Social Sciences*, 1(9):55–71.

Zhou, R. and Palomar, D. P. (2020). Understanding the quintile portfolio. *IEEE Transactions on Signal Processing*, 68:4030–4040.

Ødegaard, B. A. (2019). Asset pricing data at ose. http://finance.bi.no/~bernt/financial_data/ose_asset_pricing_data/index.html.

Appendix

A1 Stop words

We use the stop word dictionary to Loughran and McDonald (2020) with some minor modifications. The additions and deletions we do can be viewed in table A1.1.

Table A1.1: Additions and deletions in the stop word dictionary

Additions:	nok
Deletions:	ain't, almost, although, aren't, best, better, but, can't, cannot, cant, certain, certainly, couldn't, definitely, didn't, doesn't, don't, especially, few, hadn't, hasn't, haven't, however, isn't, least, most, much, neither, never, no, none, nor, not, particular, particularly, quite, really, serious, seriously, shouldn't, truly, very, wasn't, weren't, whereas, will, won't, wouldn't

A2 Valence shifters

Table A2.1: List of different valance shifters

Adversative conjunctions:	although, but, despite all that, despite all this, despite that, despite this, however, that being said, whereas
Amplifiers:	absolutely, acute, acutely, certain, certainly, colossal, colossally, considerably, decidedly, definite, definitely, enormous, enormously, especially, extreme, extremely, greatly, heavily, heavy, high, highly, huge, hugely, immense, immensely, incalculable, incalculably, majorly, massive, massively, more, most, much, particular, particularly, purpose, purposely, quite, real, really, serious, seriously, severe, severely, significant, significantly, surely, totally, truly, uber, vast, vastly, very
De-amplifiers:	almost, barely, faintly, few, hardly, incredibly, kind of, kinda, least, little, only, partly, rarely, seldom, slightly, somewhat, sort of, sorta, sparsely, sporadically, very few, very little
Negators:	ain't, aren't , can't, cannot, couldn't, daren't, didn't, doesn't, don't, hadn't, hasn't, haven't, isn't, mightn't, mustn't, needn't, neither, never, no, nobody, none, nor, not, oughtn't, oughtn't, shan't, shouldn't, wasn't, , weren't, won't, wouldn't

A3 Sentiment computation

Text sentiment

The sentiment of a text, δ_i, is the average of each sentence. The average function downweighs neutral sentences, and is defined as:

$$\delta_i = \frac{\sum \delta_{i,j}}{\delta_i^1 + \sqrt{\ln 1 + \delta_i^0}}$$

where:

δ_i^0 is the number of elements of $\delta_{i,1}$, ..., $\delta_{i,n}$ that are equal to 0

δ_i^1 is the number of elements of $\delta_{i,1}$, ..., $\delta_{i,n}$ that are different from 0

$\delta_{i,j}$ is the sentiment for sentence j in text i

Sentence sentiment

The sentiment for each sentence, $\delta_{i,j}$, is computed as follows:

$$\begin{aligned}
\delta_{i,j} &= \frac{c'_{i,j}}{\sqrt{w_{i,j,n}}} \\
c'_{i,j} &= \sum((1 + w_{amp} + w_{deamp}) \cdot w_{i,j,k^p} \cdot (-1)^{2+w_{neg}} \\
w_{neg} &= (\sum w_{i,j,k^n}) \mod 2 \\
w_{amp} &= \sum(w_{neg} \cdot (0.8 \cdot w_{i,j,k^a}) + w_{adv_b}) \\
w_{deamp} &= \max(w'_{deamp}, -1) \\
w'_{deamp} &= \sum(0.8 \cdot (-w_{neg} \cdot w_{i,j,k^a} + w_{i,j,k^d}) + w_{adv_a}) \\
w_{adv_a} &= -0.25 \cdot w_a \\
w_{adv_b} &= 0.25 \cdot w_b
\end{aligned}$$

where:

w_{i,j,k^0} : Neutral word	w_{i,j,k^a} : Amplifier
w_{i,j,k^d} : De-amplifier	w_{i,j,k^n} : Negator
w_{i,j,k^p} : Polarized word	$w_{i,j,n}$: Number of words in sentence j in text i

w_a: The number of adversative conjunctions after w_{i,j,k^p}

w_b: The number of adversative conjunctions before w_{i,j,k^p}

A4 Portfolio regressions with factor loadings

Table A4.1: Portfolio regressions for portfolios sorted on the absolute sentiment value

	Portfolio											
	E_Q5	E_Q1	E_LS	E_Q5	E_Q1	E_LS	IV_Q5	IV_Q1	IV_LS	IV_Q5	IV_Q1	IV_LS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MKT	0.83***	0.97***	-0.14*	0.93***	1.14***	-0.20**	0.69***	0.68***	0.01	0.79***	0.75***	0.05
	(0.05)	(0.07)	(0.08)	(0.06)	(0.08)	(0.10)	(0.05)	(0.06)	(0.08)	(0.06)	(0.08)	(0.09)
SMB				0.27***	0.44***	-0.16				0.26***	0.20*	0.06
				(0.09)	(0.11)	(0.14)				(0.08)	(0.11)	(0.13)
HML				-0.05	-0.16*	0.11				-0.03	-0.24**	0.20*
				(0.07)	(0.09)	(0.11)				(0.07)	(0.09)	(0.11)
Alpha	-0.03	-0.24	0.01	-0.27	-0.64*	0.17	-0.21	-0.75**	0.34	-0.44*	-0.94***	0.30
	(0.28)	(0.36)	(0.43)	(0.29)	(0.36)	(0.45)	(0.25)	(0.35)	(0.41)	(0.26)	(0.36)	(0.43)
Observations	203	203	203	203	203	203	203	203	203	203	203	203
R^2	0.56	0.52	0.02	0.58	0.55	0.03	0.52	0.36	0.0001	0.55	0.39	0.02
Adjusted R^2	0.56	0.51	0.01	0.58	0.55	0.01	0.52	0.35	-0.005	0.54	0.38	0.004

The table shows the result from the factor regression when the portfolios are sorted on the absolute sentiment value. The alphas are reported as percentages. The first part of the portfolio name indicates the weighting scheme, E is equal weighting and IV is inverse stock volatility. The second part of the portfolio name gives the portfolio. Q5 is the portfolio consisting of the stocks with the correspondingly most positive shareholder letter. Q1 is the least positive and LS is long Q5 and short Q1.

Table A4.2: Portfolio regressions for portfolios sorted on the change in sentiment

	Portfolio											
	E_Q5	E_Q1	E_LS	E_Q5	E_Q1	E_LS	IV_Q5	IV_Q1	IV_LS	IV_Q5	IV_Q1	IV_LS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MKT	0.97***	0.93***	0.05	1.14***	1.08***	0.06	0.69***	0.76***	-0.06	0.79***	0.84***	-0.05
	(0.06)	(0.06)	(0.07)	(0.06)	(0.07)	(0.09)	(0.06)	(0.04)	(0.07)	(0.08)	(0.05)	(0.09)
SMB				0.47***	0.41***	0.06				0.28**	0.24***	0.04
				(0.09)	(0.11)	(0.12)				(0.11)	(0.08)	(0.12)
HML				-0.26***	-0.11	-0.16				-0.17*	-0.12*	-0.05
				(0.07)	(0.09)	(0.10)				(0.09)	(0.06)	(0.10)
Alpha	-0.38	-0.61*	0.03	-0.81***	-0.98***	-0.03	-0.36	-1.22***	0.66*	-0.62*	-1.44***	0.62
	(0.31)	(0.34)	(0.39)	(0.29)	(0.34)	(0.40)	(0.35)	(0.24)	(0.38)	(0.35)	(0.25)	(0.40)
Observations	203	203	203	203	203	203	203	203	203	203	203	203
R^2	0.60	0.52	0.002	0.66	0.56	0.01	0.37	0.59	0.003	0.40	0.61	0.01
Adjusted R^2	0.60	0.52	-0.003	0.66	0.55	-0.0004	0.37	0.59	-0.001	0.39	0.61	-0.01

The table shows the result from the factor regression when the portfolios are sorted on the change in sentiment value. The alphas are reported as percentages. The first part of the portfolio name indicates the weighting scheme, E is equal weighting and IV is inverse stock volatility. The second part of the portfolio name gives the portfolio. Q5 is the portfolio consisting of the stocks with the correspondingly most positive change in the shareholder letter. Q1 is the least positive change and LS is long Q5 and short Q1.

Table A4.3: Portfolio regressions for portfolios sorted on the absolute sentiment value in the last part

	Portfolio											
	E_Q5_4	E_Q1	E_LS	IV_Q5	IV_Q1	IV_LS	E_Q5_4	E_Q1	E_LS	IV_Q5	IV_Q1	IV_LS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MKT	0.98***	1.00***	-0.02	0.70***	0.74***	-0.04	1.16***	1.18***	-0.02	0.85***	0.86***	-0.005
	(0.06)	(0.06)	(0.08)	(0.05)	(0.06)	(0.07)	(0.07)	(0.07)	(0.09)	(0.06)	(0.07)	(0.09)
SMB							0.47***	0.47***	-0.003	0.39***	0.31***	0.08
							(0.10)	(0.10)	(0.13)	(0.09)	(0.10)	(0.13)
HML							-0.12	-0.09	-0.03	-0.08	-0.13	0.05
							(0.08)	(0.08)	(0.11)	(0.07)	(0.08)	(0.10)
Alpha	-0.40	-0.02	-0.57	-0.48*	-0.13	-0.54	-0.82**	-0.45	-0.57	-0.82***	-0.41	-0.61
	(0.34)	(0.32)	(0.42)	(0.29)	(0.32)	(0.39)	(0.33)	(0.32)	(0.43)	(0.29)	(0.32)	(0.41)
Observations	203	203	203	203	203	203	203	203	203	203	203	203
R^2	0.56	0.58	0.0002	0.47	0.44	0.001	0.60	0.63	0.001	0.52	0.47	0.004
Adjusted R^2	0.56	0.58	-0.005	0.47	0.44	-0.004	0.60	0.62	-0.01	0.51	0.47	-0.01

The table shows the result from the factor regression when the portfolios are sorted on the absolute sentiment value for the last part of the text. The last part is the last 25% of the sentences. The alphas are reported as percentages. The first part of the portfolio name indicates the weighting scheme, E is equal weighting and IV is inverse stock volatility. The second part of the portfolio name gives the portfolio. Q5 is the portfolio consisting of the stocks with the correspondingly most positive final part of the shareholder letter. Q1 is the least positive and LS is long Q5 and short Q1.

Table A4.4: Portfolio regressions for portfolios sorted on sentiment change in the letter's last part

	Portfolio											
	E_Q1_4_c	E_Q5	E_LS	E_Q1	E_Q5	E_LS	IV_Q1	IV_Q5	IV_LS	IV_Q1	IV_Q5	IV_LS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MKT	1.05***	0.99***	-0.06	1.20***	1.16***	-0.03	0.80***	0.79***	-0.002	0.88***	0.86***	-0.02
	(0.06)	(0.06)	(0.08)	(0.07)	(0.07)	(0.09)	(0.06)	(0.05)	(0.08)	(0.07)	(0.06)	(0.10)
SMB				0.37***	0.46***	0.09				0.23**	0.17*	-0.06
				(0.11)	(0.10)	(0.13)				(0.10)	(0.09)	(0.14)
HML				-0.10	-0.15*	-0.06				-0.09	-0.10	-0.01
				(0.09)	(0.08)	(0.11)				(0.08)	(0.07)	(0.11)
Alpha	-0.12	-0.43	-0.51	-0.46	-0.85***	-0.59	-0.36	-0.61**	-0.45	-0.57*	-0.77***	-0.40
	(0.34)	(0.33)	(0.42)	(0.34)	(0.32)	(0.43)	(0.31)	(0.28)	(0.42)	(0.32)	(0.29)	(0.44)
Observations	203	203	203	203	203	203	203	203	203	203	203	203
R^2	0.59	0.57	0.003	0.62	0.62	0.01	0.49	0.54	0.0000	0.51	0.55	0.001
Adjusted R^2	0.59	0.57	-0.002	0.61	0.61	-0.01	0.49	0.53	-0.005	0.50	0.54	-0.01

The table shows the result from the factor regression when the portfolios are sorted on the sentiment change in the final part of the shareholder letter. The alphas are reported as percentages. The first part of the portfolio name indicates the weighting scheme, E is equal weighting and IV is inverse stock volatility. The second part of the portfolio name gives the portfolio. Q5 is the portfolio consisting of the stocks with the correspondingly most positive change in the final part of the shareholder letter. Q1 is the least positive change and LS is long Q5 and short Q1.

Table A4.5: Portfolio regressions for portfolios sorted on similarity

	Portfolio											
	E_Q5	E_Q1	E_LS	IV_Q5	IV_Q1	IV_LS	E_Q5	E_Q1	E_LS	IV_Q5	IV_Q1	IV_LS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
MKT	1.04***	0.94***	0.11	0.89***	0.74***	0.15***	1.17***	1.06***	0.11	0.97***	0.81***	0.16**
	(0.06)	(0.06)	(0.07)	(0.04)	(0.04)	(0.05)	(0.07)	(0.07)	(0.08)	(0.05)	(0.05)	(0.07)
SMB							0.35***	0.34***	0.01	0.22***	0.21***	0.02
							(0.10)	(0.10)	(0.12)	(0.07)	(0.08)	(0.09)
HML							-0.19**	-0.16*	-0.03	-0.09	-0.20***	0.11
							(0.08)	(0.09)	(0.10)	(0.06)	(0.06)	(0.08)
Alpha	-0.41	-0.54	-0.07	-0.43	-0.49	-0.05	-0.34	-0.42	-0.04	-0.34	-0.44	-0.08
	(0.35)	(0.35)	(0.36)	(0.24)	(0.24)	(0.29)	(0.33)	(0.34)	(0.38)	(0.24)	(0.25)	(0.31)
Observations	203	203	203	203	203	203	203	203	203	203	203	203
R^2	0.59	0.54	0.01	0.67	0.58	0.04	0.62	0.57	0.01	0.69	0.61	0.05
Adjusted R^2	0.59	0.53	0.01	0.67	0.57	0.03	0.62	0.56	-0.001	0.69	0.60	0.03

The table shows the result from the factor regression when the portfolios are sorted on cosine similarity. The alphas are reported as percentages. The first part of the portfolio name indicates the weighting scheme, E is equal weighting and IV is inverse stock volatility. The second part of the portfolio name gives the portfolio. Q5 is the portfolio consisting of the stocks with the correspondingly most similar shareholder letter, compared with the preceding one. Q1 is the least similar, and LS is long Q5 and short Q1.