Trading on LinkedIn

Expressed sentiment and stock returns

Master's thesis supervised by Professor Thierry Foucault, written by Tiphaine Boutin and Camille Obeid.

ABSTRACT

The development of natural language processing (NLP) models has significantly expanded the alternative data available for understanding stock market evolution. By leveraging sentiment analysis, data scientists can now monitor investor sentiment and company disclosures more effectively. This paper explores the impact of sentiment expressed on social media, particularly LinkedIn, on stock returns. We utilize the FinBERT model to analyze financial sentiment in company posts, investigating whether this sentiment explains same-day or subsequent days' abnormal returns.

Our findings reveal a positive and statistically significant relationship between financial sentiment in LinkedIn posts and same-day stock returns, with a tendency for market overreaction. Interestingly, the effects are consistent across both financial and non-financial posts, indicating hidden information in general corporate communications. This overreaction is evidenced by a correction that happens in the following days. However, for smaller capitalization, this correction is largely explained by the standard reversal. Hence, sentiment expression is a signal that has value for larger capitalizations, where reversal does usually not apply.

We also examine the influence of stock market performance on subsequent post sentiment, finding a positive correlation, although not robust when isolating days with positive or negative returns. This suggests. that companies do not try to over leverage (resp. undermine) their positive (resp. negative) stock performance.

Finally, we develop two trading strategies to test the robustness of our main reversal findings by creating long-short portfolios that capitalize on market corrections following overreactions. Furthermore, we separate the performance of a basic reversal with that of our findings by creating a double sort long-short portfolio that shorts (resp. longs) the companies with the best (resp. worst) sentiment amongst the companies with best (resp. worst) day performance. This strategy yields 18% annualized returns for large capitalizations, which compares favorably to the performance that a reversal could produce. For small capitalization however, the standard reversal remains superior.

INTRODUCTION

With the recent development of natural language processing models (NLP), allowing data scientists to extract sentiment from text written in plain language, alternative data available to understand stock market evolution has expanded. In particular, since excess volatility in the market can partly be attributed to behavioral causes such as over- or under-reaction to news, sentiment analysis constitutes a powerful tool that allows one to monitor both what investors think about a given company, and what the company itself choses to disclose.

Furthermore, social media is taking an increasingly important place in the information space, partly replacing conventional media and creating a direct link between companies with customers or investors. Subsequently, this can be leveraged in combination with NLP to monitor sentiment on specific populations, regarding certain topics.

On the one hand, a literature examining the wisdom of crowds has emerged, showing that sentiment expressed on social media can have explanatory power regarding the stock market. At the market level, a working paper of the Federal Reserve Board by Adams and al. (2023)[1] shows, among other findings, that overnight Twitter financial sentiment helps predict next day stock market returns, showing that social media is an effective reflection of investors' sentiment. At the stock level, a causal relationship was exhibited by Wang et al.(2022)[2] between sentiment expressed in social media and stock returns. They conducted a controlled experiment on China's largest and most active stock investment forum and showed that investors' opinions posted on social media had an influence on same-day stock returns. Using a more basic sentiment metric, Sul and al. (2017)[3] examined the dynamics of sentiment diffusion in social media, and showed in particular that return predictability through social media sentiment was strongest when the tweets were originated from users with low following. In this case, there was a positive relationship between sentiment and subsequent days' returns.

On the other hand, although investors' opinions matter for stock returns, both due to herding effects and because some investors might be informed, a substantial driver of the stock market movements remains the discovery of firm fundamentals. Therefore, opinions expressed by company employees ought to get incorporated in prices. Green and al. (2019)[4] show that employee reviews on Glassdoor are a predictor of earnings announcement surprises, confirming thereby that they reveal fundamental information about the firm.

However, we found that the literature was lacking analysis of how information disclosed by employees on social media could impact daily stock returns. Indeed, companies have official accounts on social media where they can potentially post financially-relevant information. We hypothesized that the latter could have explanatory power over stock returns, as sentiment disclosed through social media could be traded on by investors, inducing a positive relationship between expressed sentiment and short-term stock returns. However, it could also be that information would travel from the stock market to information channels, in which case sentiment would reflect past returns.

To test these hypothesis, we collect a history of employee-sourced posts from a widely used social media channel. Given the limitations on scraping imposed by the websites, particularly the recent policy put in place for X (formerly Twitter), we chose to work on LinkedIn, the largest business and employment-focused social media platform worldwide, used by companies mainly as a self-promotion channel. Boasting 930 million monthly active users, it ranks amongst the top 10 largest social media platforms. Therefore, companies regularly disclose various news on their feed, such as corporate events, business developments or financial updates.

We then analyze financial sentiment in each of these posts using the NLP model FinBERT introduced by Araci (2019)[5] and used in Adams and al. (2023). As the model is trained specifically on financial data, it is supposedly able to discriminate against positive sentiment of general marketing content and emphasize financial sentiment.

This alternative data allows us to answer four main questions in our paper.

The first question we address is whether company social media post sentiment is a predictor of same-day or subsequent days' returns. Our analyses uncover a positive and statistically significant relation between financial sentiment expressed in posts and same-day returns. In addition, further analysis suggests that the market would overreact to the positive news, as the relationship reverses in the following days. Indeed, subsequent cumulative abnormal returns remain significantly and negatively impacted by a positive sentiment for about 3 weeks.

Then, we examine how these dynamics differ depending on company size, proxied by market cap. We find that the effect of sentiment expressed in LinkedIn posts is much stronger for smaller companies, and more persistent.

In addition, we test whether financial posts exhibit specific patterns and hold additional information compared to regular posts. Indeed, we find that these effects are the same in all sorts of posts, which means that there is information hidden in non-financial communications. In fact, the correlation between same day abnormal returns and LinkedIn post sentiment is not significantly improved with respect to non-financial posts.

Moreover, we explore whether stock market returns influence financial sentiment expressed in subsequent posts. In fact, we show that the performance of a company's stock on a certain day is positively correlated with subsequent posts sentiment after market hours. However, the fact that these findings don't hold when we isolate days with positive and negative abnormal returns, shows that these findings aren't robust. This suggests that companies don't consciously react and try to manipulate market performance through their posts but are rather prone to unconscious mood swings in their posting patterns due to past market performance.

Finally, we test the robustness of our analysis by leveraging the findings above to create trading strategies. To take advantage of the correction following market overreaction, we build a long-short portfolio by selecting every day the companies whose posts expressed the most extreme sentiment, and then taking respectively a long position in the negative companies and a short position in the positive. Furthermore, we separate the performance of a basic reversal with that of our findings by creating a double sort long-short portfolio that shorts (resp. longs) the companies with the best (resp. worst) sentiment amongst the companies with best (resp. worst) day performance. We find that these strategies produce significant returns when restricting the universe to larger companies, for which a standard reversal does not apply.

METHODOLOGY

Influence of sentiment expressed on stock market returns
How does the stock market react to sentiment expressed in social media?

Our hypothesis is that sentiment expressed by management contains insider information on firm fundamentals and should therefore have predictive power over the stock market. However, the timeliness of this information and the speed at which it gets incorporated is a determinant of the practicability of this information for investors. To test at which time horizon the information is relevant, we run three sets of tests.

Sentiment expressed and same-day returns
Do markets consider posts' financial sentiment as novel information?

To capture immediate market reaction to sentiment expression, we test the effect of sentiment expressed during a day on stock returns in the same timespan. We run the following panel regression:

$$AR_{i,j} = \alpha + \beta \times S_{i,j} + \beta_{rev} \times AR_{i,j-1} + \beta_X X_{i,j} \quad (1)$$

where:
- $AR_{i,j}$ are the abnormal returns for stock i on day j,
- $S_{i,j}$ is the mean sentiment expressed between the close of trading day $j-1$ and the open of trading day j,
- $X_{i,j}$ are the control variables

Given that with a daily timestep, the reversal has strong explanatory power, we control for previous day's returns. Indeed, we expect a negative and significant β_{rev} coefficient. Furthermore, similarly as in the previous test, we control for seasonality and industry-specific biases.

The sign and significance of the β coefficient measures the effect of sentiment expression on stock returns. Indeed, a positive β would indicate that sentiment expressed in LinkedIn posts is perceived as novel information by markets and quickly incorporated into prices. On the contrary, a negative β would suggest that either investors question the credibility of the information disclosed, or they have already accessed the information through a different channel and overreacted to it.

To run our tests, we standardize previous days' returns, and rescale by multiplying by the standard deviation of mean sentiment. The goal is to be able to compare the magnitude of β and β_{rev}, as we expect the reversal to constitute the main effect.

Sentiment expressed and subsequent days' returns
How quickly is sentiment information incorporated in prices?

To understand how quickly sentiment information gets incorporated into prices, we estimate the effect of LinkedIn financial sentiment on subsequent cumulative abnormal returns, for multiple time horizons. We run the panel regression:

$$MAR_{i,j+1\to j+n} = \alpha_n + \beta_n \times S_{i,j} + \beta_{rev,n} \times AR_{i,j} + \beta_X X_{i,j} \quad (2)$$

where $MAR_{i,j+1\to j+n}$ are the mean abnormal returns for stock i from day $j+1$ to day $j+n$ included. Like in the previous regression, we also estimate the effect of the previous day's returns and rescale the regressor so that the sentiment effect β_n is comparable to the reversal effect $\beta_{rev,n}$.

We expect to see the significance and/or magnitude of the two effects to fade over time, potentially with different speeds.

We expect the reversal to be the main effect, its sign to be negative and its significance to fade in a couple of days. Furthermore, as long as the sentiment effect is positive, it translates a slow incorporation of the news into prices. On the contrary, if it becomes negative, it adds to the reversal and translates an overreaction of the market to expressed sentiment. Finally, no significance of the sentiment effect would mean that LinkedIn news do not influence stock price variations.

Practicability of results
Can a trading strategy on LinkedIn sentiment produce superior returns?

As a robustness check of our findings, as suggested by Sul and al. (2017), we build a trading strategy based on sentiment data. We use a double-sort methodology to control for reversal effects, and we benchmark our returns against a standard reversal.

The goal is to test whether selecting stocks on sentiment can add to the reversal, i.e. refine the selection on past returns. We also test whether selecting stocks on sentiment can do better than the reversal, i.e. replace the selection on past returns.

Hence, we build three series of long-short portfolios, one for the standard reversal, used as benchmark, one for the sentiment reversal and one for the double-sort. We then compare returns and alphas over the Fama-French 3-factor model.

For the standard reversal, we divide each day our set of stocks into quantiles based on the previous day's abnormal returns. We then initiate a long position on the bottom subset and a short position on the upper subset. The position is kept for a predefined holding period, ranging from 1 to 15 trading days. We make sure that leverage is kept constant and equal to 1, both on the short and long leg, by scaling the positions at initiation as a function of the number of names in each quantiles.

For the sentiment reversal, the method is the same, except that we select on the previous day's sentiment instead of returns.

For the double-sort, we first divide our set of stocks into quantiles based on the previous day's returns, as for the simple reversal long-short explained above. Then, we divide again the upper and lower subset in quantiles, based on the previous day's expressed sentiment. We finally select the upper sentiment subset in the upper return subset, and the lower sentiment subset in the lower return subset. We initiate equal-weight short positions in the former and long positions in the latter.

In terms of parameters, we chose the quantiles such that the same number of companies is held in the benchmark and test portfolios.

We expect that, with holding periods for which sentiment is shown to be significant in the previous regressions, performance of the double-sort can exceed performance of the standard reversal. However, if the effect of the reversal still dominates, we do not expect the simple sentiment reversal to beat the standard reversal.

Figure 1: Portfolio construction for the three long-short series.

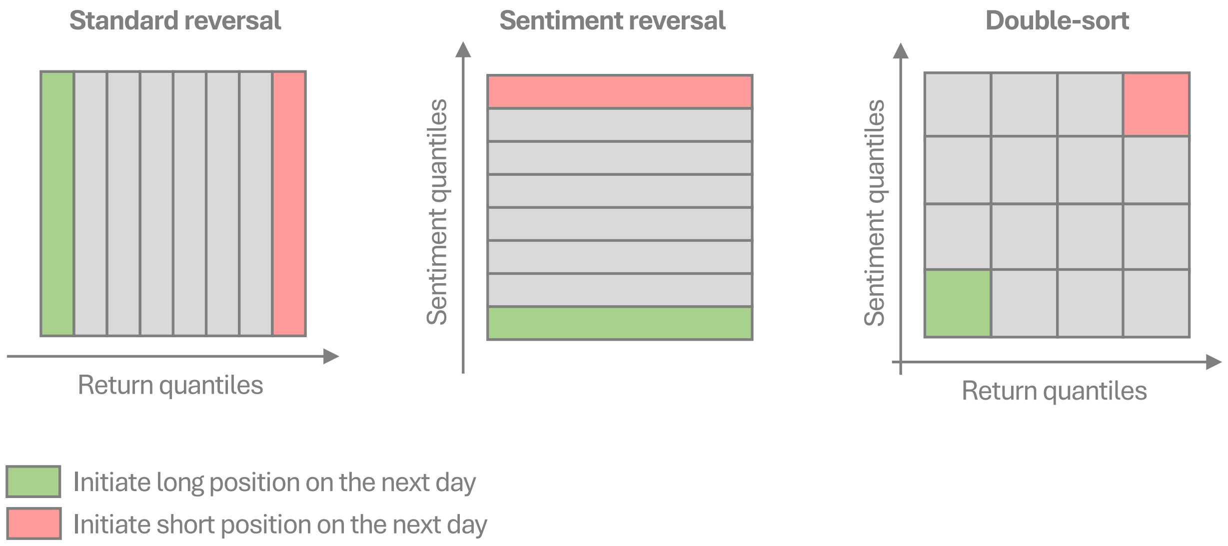

Size and social media information

How do financial disclosure on social media and stock market dynamics vary depending on the size of the company?

Since stock market dynamics can be highly dependent on company size, as it conditions analyst coverage and liquidity, we test the validity of our conclusions on two subsamples of our universe. We divide the companies into a high revenue and a low revenue group, with the median revenue as a cutoff point. Then, we run the all the tests above on these subsamples.

As the reversal is supposedly stronger on smaller companies, we expect to find that the lower revenue group has a larger beta than the upper revenue group on the reversal for the two regressions, and that its reversal long-short performs better . Then, we also expect the effect of sentiment to be larger for less liquid stocks, leading to larger betas on sentiment expressed, and better performance of the sentiment long-short.

However, it is unclear how the magnitude and dynamics of standard reversal and sentiment effect would compare, both in the lower and upper revenue group. In particular, we test the performance of the sentiment reversal and double-sort in settings where the standard reversal is shown to fail, i.e. for the upper revenue group.

Posting dynamics and financial information

Do financial posting dynamics exhibit specific patterns and hold additional information compared to regular posting dynamics?

In order to understand posting dynamics for financial content, we first need to define what constitutes a financial post. To this end, we use a methodology similar to Calomiris et al. (2020)[6] who defined a dictionary of root words relating to regulation to determine whether a text is related to financial regulation. In our case, we define a dictionary of roots and words relating to financial disclosure (cf. Table 1).

Table 1: financial disclosure roots and words.

earning	report	revenue	profit	margin	debt	expense	gross	stock
market	income	share	analyst	return	dividend	gain	8K	10K
merg	acqui	equity	fisc	quarter	result	ESG	financ	

Note that some words aren't roots. That is because we want to be stricter in the financial classification of our posts. Indeed, FinBert, our sentiment analysis model, is already tuned to determine financial sentiment, which discriminates between financial and non-financial posts to some extent. This added dictionary-based measure hence has a goal of identifying posts that are purely related to big financial events (e.g. earnings announcements, M&A activity etc.) in order to verify if these big market movers are the ones driving the correlation between sentiment and returns or if the nuanced posts have a predictive power themselves.

With that in mind, we classified the posts in the following three categories:
- Non-financial posts: posts that include less than two of the financial dictionary words
- Somewhat financial posts: posts that include three of the financial dictionary words
- Financial posts: posts that include more than three of the financial dictionary words

This classification is based on manual checks and observations throughout random samples of the data, keeping in mind that the primary objective is to separate financial disclosure posts with definitive market moving potential from the rest.

Classifying the posts in terms of financial nature will allow us to better understand posting dynamics and the financial sentiment per type of post before moving to the regressions. Most notably, we will have a better understanding on when financial posts are posted, which we expect to be before the market open and close as per financial disclosure customs, and we will see whether financial posts have more skewed sentiments, which we expect. This could lead to more explainability of abnormal returns through the following regression with interaction terms between mean sentiment and financial score:

$$AR_{i,j} = \alpha + \beta \times S_{i,j} + \beta_{sw\_fin} \times S_{i,j} \times \mathbb{1}_{sw\_fin} + \beta_{fin} \times S_{i,j} \times \mathbb{1}_{fin} + \beta_X X_{i,j} \quad (3)$$

where:
- $S_{i,close,j \to j+1}$ is the mean sentiment expressed on trading day j , for company i,
- $AR_{i,j}$ are the abnormal returns on day j, i.e. from the close of trading day $j-1$ to the close of trading day j,
- $\mathbb{1}_{sw\_fin}$ (resp. $\mathbb{1}_{fin}$) is an indicator function equal to 1 if the post is somewhat financial (resp. financial)
- $X_{i,j}$ controls for seasonality and industry-specific biases by including a dummy variable for each day of the week, and for each industry.

Influence of stock market returns on sentiment expressed

How are company posts' financial sentiments influenced by recent stock market returns?

To test the influence of stock returns on subsequent financial communication, we first observe if all posts whose timestamp is between the close of trading day j and the open of trading day $j + 1$ are influenced by close-to-close returns of day j.

Therefore, the potential effect of stock returns on subsequent managers' communication would be uncovered by running the following panel regression:

$$S_{i,close,j \to j+1} = \alpha + \beta \times AR_j + \beta_X X_{i,j} \tag{4}$$

where:
- $S_{i,close,j \to j+1}$ is the mean sentiment expressed after the close of trading day j and before the open of trading day $j + 1$,
- AR_j are the abnormal returns on day j, i.e. from the close of trading day $j - 1$ to the close of trading day j,
- $X_{i,j}$ controls for seasonality and industry-specific biases by including a dummy variable for each day of the week, and for each industry.

In addition to running the above regression on the whole sample, we test the effect separately on both small and big firms. Indeed, the effect might be more prominent for small firms as their communications might have a stronger impact on their less renowned and studied stocks, whereas bigger firms can have less impact through their LinkedIn communication as analysts and institutional investors are watching them more closely.

A positive β would indicate that managers relay the stock performance they witness. On the contrary, a negative β would suggest they downplay the price variations.

Furthermore, we test whether the effect is different if the close-to-close stock market performance is positive or negative. Indeed, we might see different patterns for those two scenarios as managers might try to capitalize on the market's positive sentiment but downplay or even reverse the downward trend if the market's sentiment is negative.

Indeed, all the above scenarios are plausible at this point. Therefore, to better understand the dynamics, we must run a variety of regressions for this specific setting.

DATA

Definition of the universe

To define our trading universe, we use the S&P Compustat database. To get a sample that best represents the entirety of the American stock market, we retrieve the list of all stocks quoted in USD on the NYSE and Nasdaq (exchange codes 11 and 14), excluding financial tickers such as funds and ETFs (naics starting with 5259), between January of 2018 and April of 2024. We end up with 6,246 different company names, that constitute the base of our universe.

Alternative data

From the initial list of tickers, we automate a Google search to find potential LinkedIn usernames associated with the official company names provided by Compustat. In practice, for each ticker we enter into the Google search bar the keyword `LinkedIn` followed by the company name, and we check whether the first link provided is a regular expression of the form `^https://\w+.linkedin.com/company/[\w-]+$`, which would indicate a company page. With this method, we get a database of 4,396 usernames. We also retrieve the number of followers and employees for each username.

Then, we automate scraping of post history for each username by loading the respective post page, then retrieving each post's content, timestamp, reactions, comments and repost. At this point, we make two observations relevant to the subsequent analysis. First, LinkedIn caps post history to 500 posts; when the limit is reached, the oldest posts are deleted. Hence, we do not include accounts whose posts history reaches 500 posts, because they may be incomplete. Second, posts are visible for a year, after which they disappear unless they are pinned by the user. Therefore, the start date for our analysis is set exactly one year prior to post collection, i.e. on March 22nd, 2023. All posts whose timestamp is prior to this date are disregarded.

We then compute an array of sentiment scores for each post using FinBERT, a natural language processing (NLP) model created by Araci (2019)[5]. In fact, Adams et al. (2023)[1] argue that this model, can successfully detect financial sentiment. Indeed, FinBERT takes the BERT NLP model created by Google Language (Devlin et al. (2019)[7]) as a base and is then further trained on the financial text corpus *TRC2-Financial* and on a Financial Phrase Bank from Malo et al. (2014)[8]. On the one hand, *TRC2-Financial* which consists of 1.8 million news articles that were published by Reuters between 2008 and 2010. This improves the BERT model's classification ability through a first pre-training in the target domain of finance. On the other hand, the Financial Phrase Bank consists of 4845 English sentences randomly selected from financial news and analyzed by 16 people with backgrounds in finance or business. Each analyst determines whether the sentence has a positive, negative, or neutral financial sentiment. The results are aggregated to give us the phrase bank. Note that we chose the bank where at least 50% of analysts agreed on the final rating of the sentences (instead of 66%, 75% or 100%), because we wanted to keep the highest number of sentences and we are using the model on posts that are mostly of non-financial nature. We would've liked to explore the difference in our results if we had chosen another bank, but we lacked the computing power to perform the training more than once.

With that in mind, we then ran FinBERT on all the LinkedIn posts we gathered, which gave us a financial sentiment score between -1 (negative) and 1 (positive) for each sentence of each post. Looking at the sentiments and phrases we notice that all sentences with absolute sentiment score of less than 0.1 are what we call "financially useless sentences", notably hashtags, and greetings. Moreover, over 75% of posts have 4 sentences or less but, the remaining posts do have a sentence count that can reach 58 (cf. Figure 2) amongst which is a high proportion of "financially useless sentences" that actually bring out the overall "financial uselessness" of the post.

In order to aggregate the sentence sentiments for each post we hence used a weighted average of the sentiments where all sentences with a sentiment with absolute value less than 0.1 are aggregated and counted as one sentence in the overall mean computation. This method allows us to keep the information in the presence of “financially useless sentences” while mitigating their effect on the score of more meaningful sentences.

Figure 2: Distribution of posts sentence count.

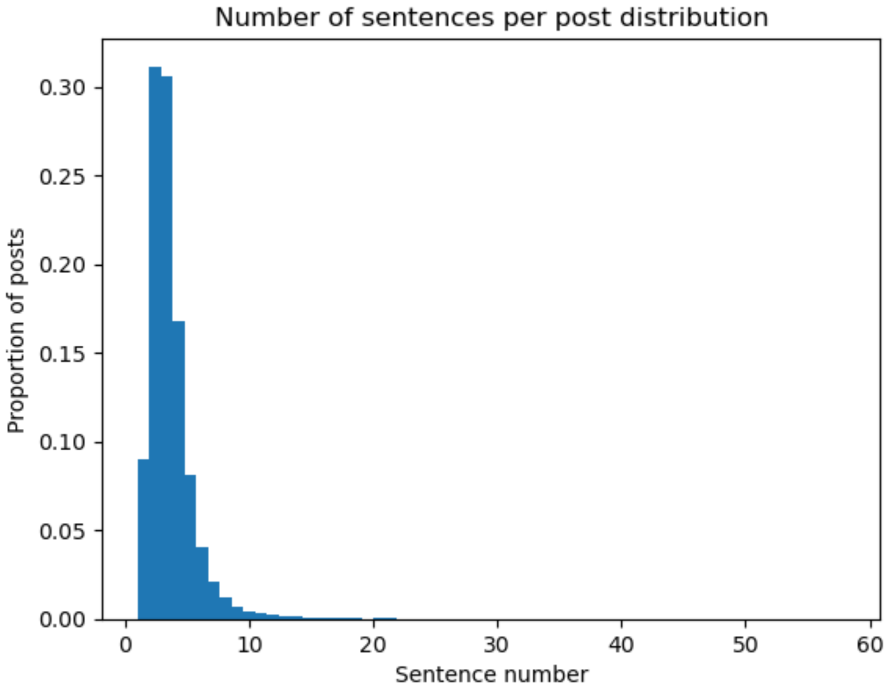

Finally, we aggregate sentiment data for each company at the daily level by grouping and averaging the sentiment score available for all posts on each trading day. In other words, we consider that posts belong to the same trading day if they fall within the close of the last weekday and today’s close. Concisely, all weekend posts belong to Monday, and all after close posts belong to the next day’s average. We chose this strategy because we wanted to provide insights that are actionable. Furthermore, this setup works well with the market data where all computations are from close-to-close.

Market data

With that in mind, we retrieve **market data** from the **CRSP** (Center for Research in Security Prices) library; for each ticker from the universe, we get **close-to-close daily returns** until **December 29th, 2023**. Furthermore, we download data for the **Fama-French 3-factor model** on US stocks from Kenneth French’s website. Indeed, the 3-factor model developed in Fama and French (1993)[9] is widely used to compute abnormal returns.

Then, for each stock, following the methodology in Wang et al. (2022), we compute the historical beta on each factor, by running the regression:

$$R_{i,t} = \alpha_i + \beta_{Mkt,i} R_{Mkt-RF,t} + \beta_{SMB,i} R_{SMB,t} + \beta_{HML,i} R_{HML,t} + \epsilon \tag{5}$$

where $R_{i,t}$ is the return of stock i on day t, $R_{Mkt-RF,t}$, $R_{SMB,t}$, $R_{HML,t}$ are respectively the daily returns of the mimicking portfolios for the market, small-minus-big and high-minus low risk premia, and $\beta_{Mkt,i}$, $\beta_{SMB,i}$, $\beta_{HML,i}$, α_i are the parameters to estimate. We use historical data for the 4 years prior to the observation period for the post, i.e. March 22nd, 2019, to March 21rst, 2023.

Finally, we compute abnormal returns in the observation period by plugging in the betas obtained on historical data:

$$AR_{i,t} = R_{i,t} - \left(\beta_{Mkt,i}R_{Mkt-RF,t} + \beta_{SMB,i}R_{SMB,t} + \beta_{HML,i}R_{HML,t}\right) \quad (6)$$

Company statistics

We merge the sentiment and market data on ticker and date. Furthermore, we keep only the tickers for which abnormal return data is available on all **196 trading days** between the start date (March 22nd, 2023) and the end date (December 29th, 2023).

Considering that our market of interest is that of North American companies quoted in USD on the NYSE or Nasdaq, excluding funds, our universe covers 75% of the names available on Compustat, and 81% of the total market capitalization.

Table 2: Summary statistics of the universe

	Universe	Total market	
Number of companies	3,915	5,217	75%
Market capitalization ($T)	33.08	40.67	81%

These companies are relatively large as their revenues are more than 100 million dollars and their employee count is over 1000 (cf. Figure 3). Furthermore, less than 10% of companies have less than 1,000 followers, whereas over 60% of companies have over 10,000 followers which means that most company posts reach a large audience. This is good to note as it assures us that the information relayed by a large majority of companies is seen by a large enough audience to potentially have an impact on markets, which supports the potential significance of our studies.

Figure 3: Summary statistics of the companies in the universe

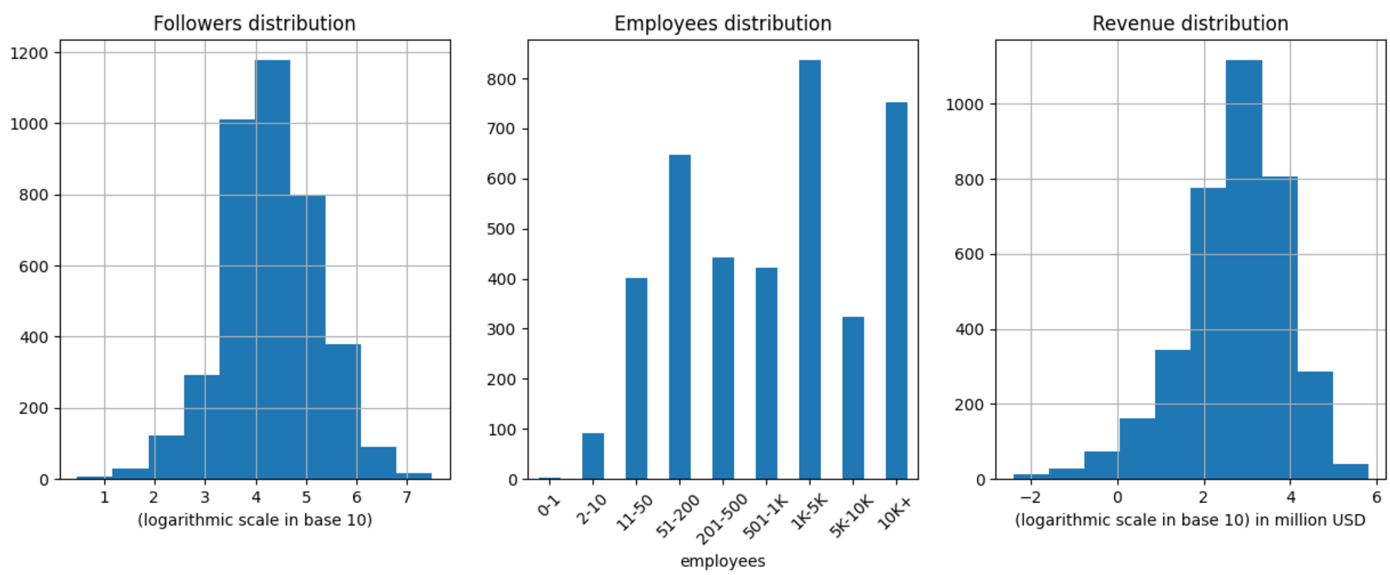

Posts and sentiment statistics

Looking at posts data, we have a total of 368,136 posts in our 196 trading days study period. To put things into perspective, this boils down to 1878 posts per day, and 0.5 posts per trading day per company. This gives us a large enough daily dataset to conduct our long short portfolio strategy to check the robustness of our regression findings. Furthermore, as shown on the right of figure 4, the posts are evenly distributed over the 5 trading days in a week with a small dip on Friday as managers have a higher probability of leaving early or taking a long weekend. On the other hand, the left and center graphs of figure 4 show us that our trading day convention does have an impact on the overall explainability of our studies as a non-negligeable proportion of posts happens after the close, whether that's on weekdays after 4pm ET or weekends.

Figure 4: Summary statistics of the posting times.

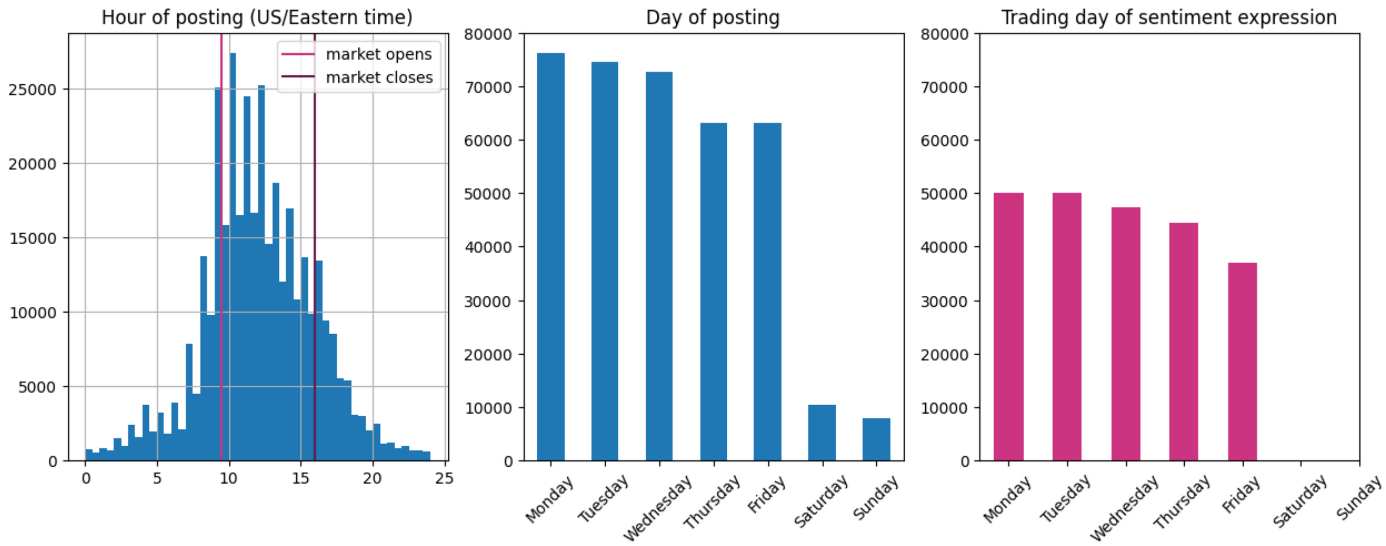

On the other hand, we notice that the number of posts is proportional to the company size, both in terms of revenue and employees. However, as shown is figure 5, this relation is not linear but more logarithmic, which in some sense is good as it is harder to find explainable abnormal returns (arbitrages) for large companies, so the larger the dataset, the more information we can extract for a potentially improved explainability.

Figure 5: Posts count to company size relation.

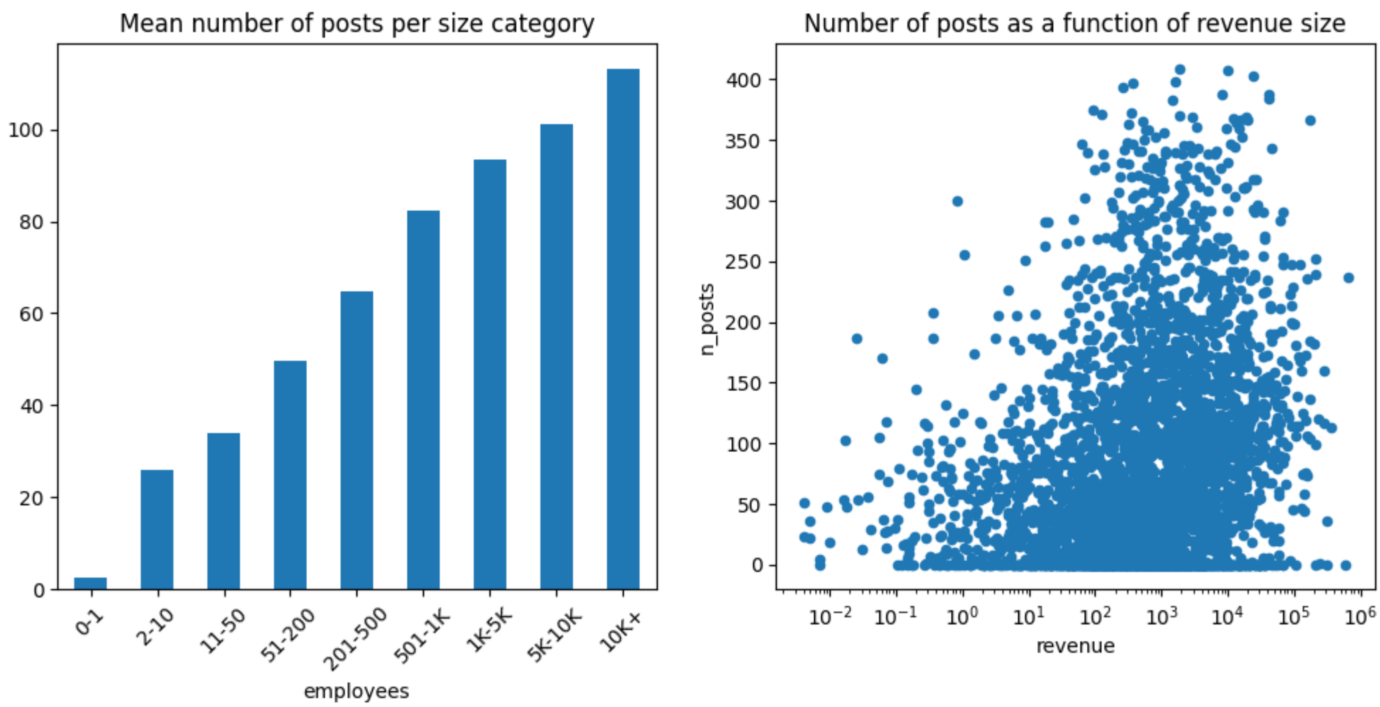

Finally, we look at the sentiment scores to better understand how they work, what they mean and what are the main statistics behind them. First, table 3 shows us a few examples of posts with their sentiment breakdown per sentence and for the overall post using different aggregation methods.

Table 3: Post scoring examples.

Post text	Sentiment per sentence	Simple average	Weighted average	Excluding "financially useless sentences"	Maximum sentiment (absolute)
We're proud to announce we've won the Chicago Innovation award for the Eterna™ SCS System. We are honored to be recognized for this award. Learn more about Eterna by visiting.\n\nImportant Safety Info:	0.92 0.72 0.03 0.02	0.42	0.56	0.82	0.92
Karooooo Limited, NASDAQ: KARO, announces FY23 Q4 and Full Year results reporting solid earnings and record free cash flow generation. Earnings per share increased 27%. These strong results extend the group's track record of growth at scale, profitability and cash generation.	0.93 0.93 0.93	0.93	0.93	0.93	0.93
According to Newmark's 3Q23 Capital Markets Report, CRE debt origination activity has dropped 48% YoY, with a 26% decrease in active lenders. Additionally, $792 billion of potentially troubled debt will be maturing between 2023 and 2025.\n\nThese challenging times call for careful navigation and strategic decision-making.\n\nDownload our report to learn more.	-0.96 0.06 0.09 0.03	-0.2	-0.45	-0.96	-0.96
Vote for AAR as your Top Shop 145! 🏆\n\nUse this link to support our Component Repair - Amsterdam team:\n\nThank you! 💙	0.15 0.88 0.00	0.34	0.34	0.52	0.88

First, we can see that FinBERT has a good accuracy in terms of determining the sentence sentiments, especially since the scores become more extreme if the sentence has a financial nature. Most notably, we can see that the 0.1 threshold is a good candidate for what we defined as "financially useless phrases". This accuracy is overwhelmingly present in the database, which is why FinBERT is often referred to in other research papers, as aforementioned in our methodology.

On the other hand, we can see that the weighted average combines the downplaying advantage of the average when there are a lot of "financially useless phrases" (example 4) while keeping their impact limited in case of truly valuable information (example 1 and 3). Furthermore, when posts are short, which is in most cases, and in the most extreme sentiment cases (example 2) we can see that all aggregating methods yield the same result as the sentiment is clear across the post. The above observations support the use of the weighted average method.

With that in mind, we end up with the sentiment score distribution shown on the shown in figure 6 (top left). This distribution has a higher variance than the regular average but keeps the variance within reason compared to the maximum score and low sentiment exclusion methods as shown in figure 6.

Figure 6: Post scoring distributions using different aggregation methods.

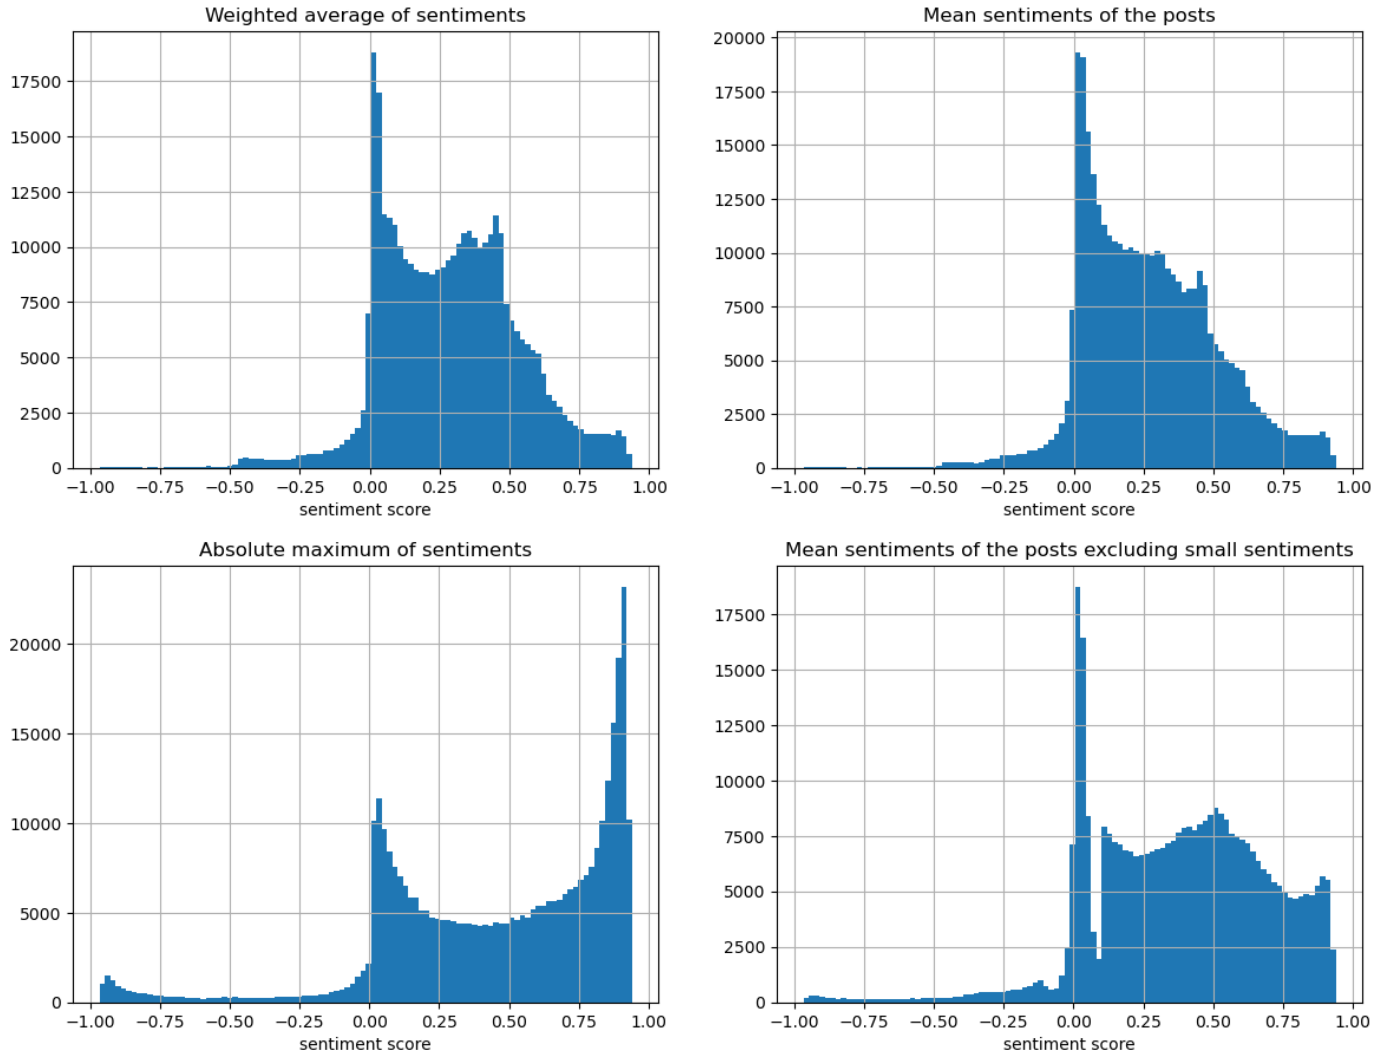

Finally, on a more general note, we can see that the post sentiments have a positive skew, which is normal considering that these posts are done by the company themselves which have little to no incentive to post negative news about their own company. In fact, many negative posts do not necessarily target the company itself they are generally negative about a certain market or the global outlook. Therefore, the fact that the overwhelming majority of posts have positive sentiment is good as the regression results will be mostly based on these datapoints, which generally have a higher chance of talking about the company itself. However, one good thing about FinBERT that was previously mentioned is that it discriminates well between posts that do and don't have financial news. Therefore, posts that signal bad financial news for the company are often on the lower side of the sentiment spectrum which gives them more significance.

RESULTS

Influence of sentiment expressed on stock market returns
How does the stock market react to sentiment expressed in social media?

We find evidence that sentiment expressed in LinkedIn posts has a positive impact on same-day stock returns, and a negative impact on subsequent days' returns. Thus, we conclude that markets consider expressed sentiment as novel information, although the previous section shows that part of this sentiment is a mere reflection of past returns. This overreaction phenomenon is followed by a correction, that cannot be subsumed by standard reversal.

Sentiment expressed and same-day returns
Do markets consider posts' financial sentiment as novel information?

We find a significant and positive relation between expressed sentiment and same-day returns, for the whole sample as well as the two subsamples on market capitalization. Table 4 displays the coefficients β and β_{rev} in equation (1), quantifying respectively the effects of expressed sentiment and past returns on same-day returns.

Table 4: Regression results of same-day returns on expressed sentiment and previous day's returns.
$AR_{i,j} = \alpha + \beta \times S_{i,j} + \beta_{rev} \times AR_{i,j-1} + \beta_X X_{i,j}$

	Full sample	Lower capitalization group	Upper capitalization group
$\boldsymbol{\beta}$	45.95***	109.29***	5.86*
(t-value)	(5.25)	(4.65)	(1.88)
$\boldsymbol{\beta_{rev}}$	-56.36***	-66.29***	12.41
(t-value)	(-4.33)	(-2.73)	(1.47)
in bps per std	-13.03	-15.32	2.87

Notes: The results are interpreted as follows. If sentiment increases by 1 point, which corresponds to going from neutral to as positive as the index goes, expected abnormal returns on the next day increase by β (in bps). Furthermore, if rescaled abnormal returns increase by 1 point, expected abnormal returns on the next day increase by β_{rev} (signed, in bps). We include the coefficient β_{rev} in bps per std, representing the change in expected abnormal returns the next day for every 1 standard deviation change in abnormal returns.
$^{*}p \leq 10\%, ^{**}p \leq 5\%, ^{***}p \leq 1\%.$

We decide to present the β coefficients in bps per sentiment index point. Indeed, as shown in the data section, the distribution of mean sentiment is far from normal, as it showcases two peaks around 0 and 0.4, translating a high share of respectively neutral and somewhat positive posts.

As expected, the reversal effect is strong (p-value lower than 1%) and negative on the lower capitalization group, but unsignificant (p-value higher than 10%) and positive on the upper capitalization group. Thus, it remains significant and positive on the full sample. This difference in reversal is well documented in the literature; smaller and more illiquid stocks present strong reversal patterns, whereas larger and more liquid stocks are better arbitraged.

Furthermore, for the lower capitalization group, the sentiment effect is highly significant, and close to twice the magnitude of the reversal. On the contrary, for the upper capitalization group, the sentiment effect has weaker significance, with a p-value around 6%, and weaker magnitude. In fact, the sentiment effect is close to 20 times stronger for the lower capitalization group compared to the upper capitalization group. Subsequently, the former drives up the significance and magnitude of the effect on the full sample. As a result, the sentiment effect on the full sample is of comparable magnitude to the reversal, of high significance (p-value lower than 1%), and of opposite sign.

These findings are in line with the more general principle that large stocks tend to be more informationally efficient. Indeed, the lesser effect observed on large capitalizations is likely to be due to the fact that larger companies would already have communicated on most news through other channels, potentially more finance-specific, and that information communicated through sentiment would be regarded as less relevant. On the contrary, as small capitalizations tend to have less media coverage, communication on LinkedIn seems to be an important source of information for investors, who react on social media posts within the same trading day.

To conclude, mean sentiment expressed on LinkedIn has a strong positive effect on same-day abnormal returns, suggesting that investors would trade on this alternative data, especially when it comes to small capitalizations.

Sentiment expressed and subsequent days' returns

How quickly is sentiment information incorporated in prices?

We find a significant and negative effect of sentiment expression on subsequent days' abnormal returns. It is the strongest on the trading day following sentiment expression and persists for 3 weeks. The magnitude is comparable for the lower and upper capitalization groups, but the significance is higher for larger companies. Furthermore, for small stocks the effect vanishes after 3 days, whereas for large stocks it remains significant for up to 50 days. Figure 7 below shows the evolution of coefficients β_n and $\beta_{rev,n}$ from equation (2) with the time horizon, for the three universes. The first coefficients are then displayed with their significance levels in Table 5.

Figure 7: Evolution of the sentiment and reversal effect.

For each subsample, the left graph shows the full evolution of the sentiment effect, measured by β_n (grey line), from a 1-day to a 90-day time horizon. It also displays the same-day sentiment effect β computed in the previous section (black cross). The right graph focuses on the 1-day to a 50-day time horizon and displays the significance levels for β_n (yellow stars). The reversal effect, quantified by $\beta_{rev,n}$ is also plotted (light grey line) for comparison. Note that the datapoints are not evenly spaced in the time horizon range, with a higher frequency for smaller values of n.

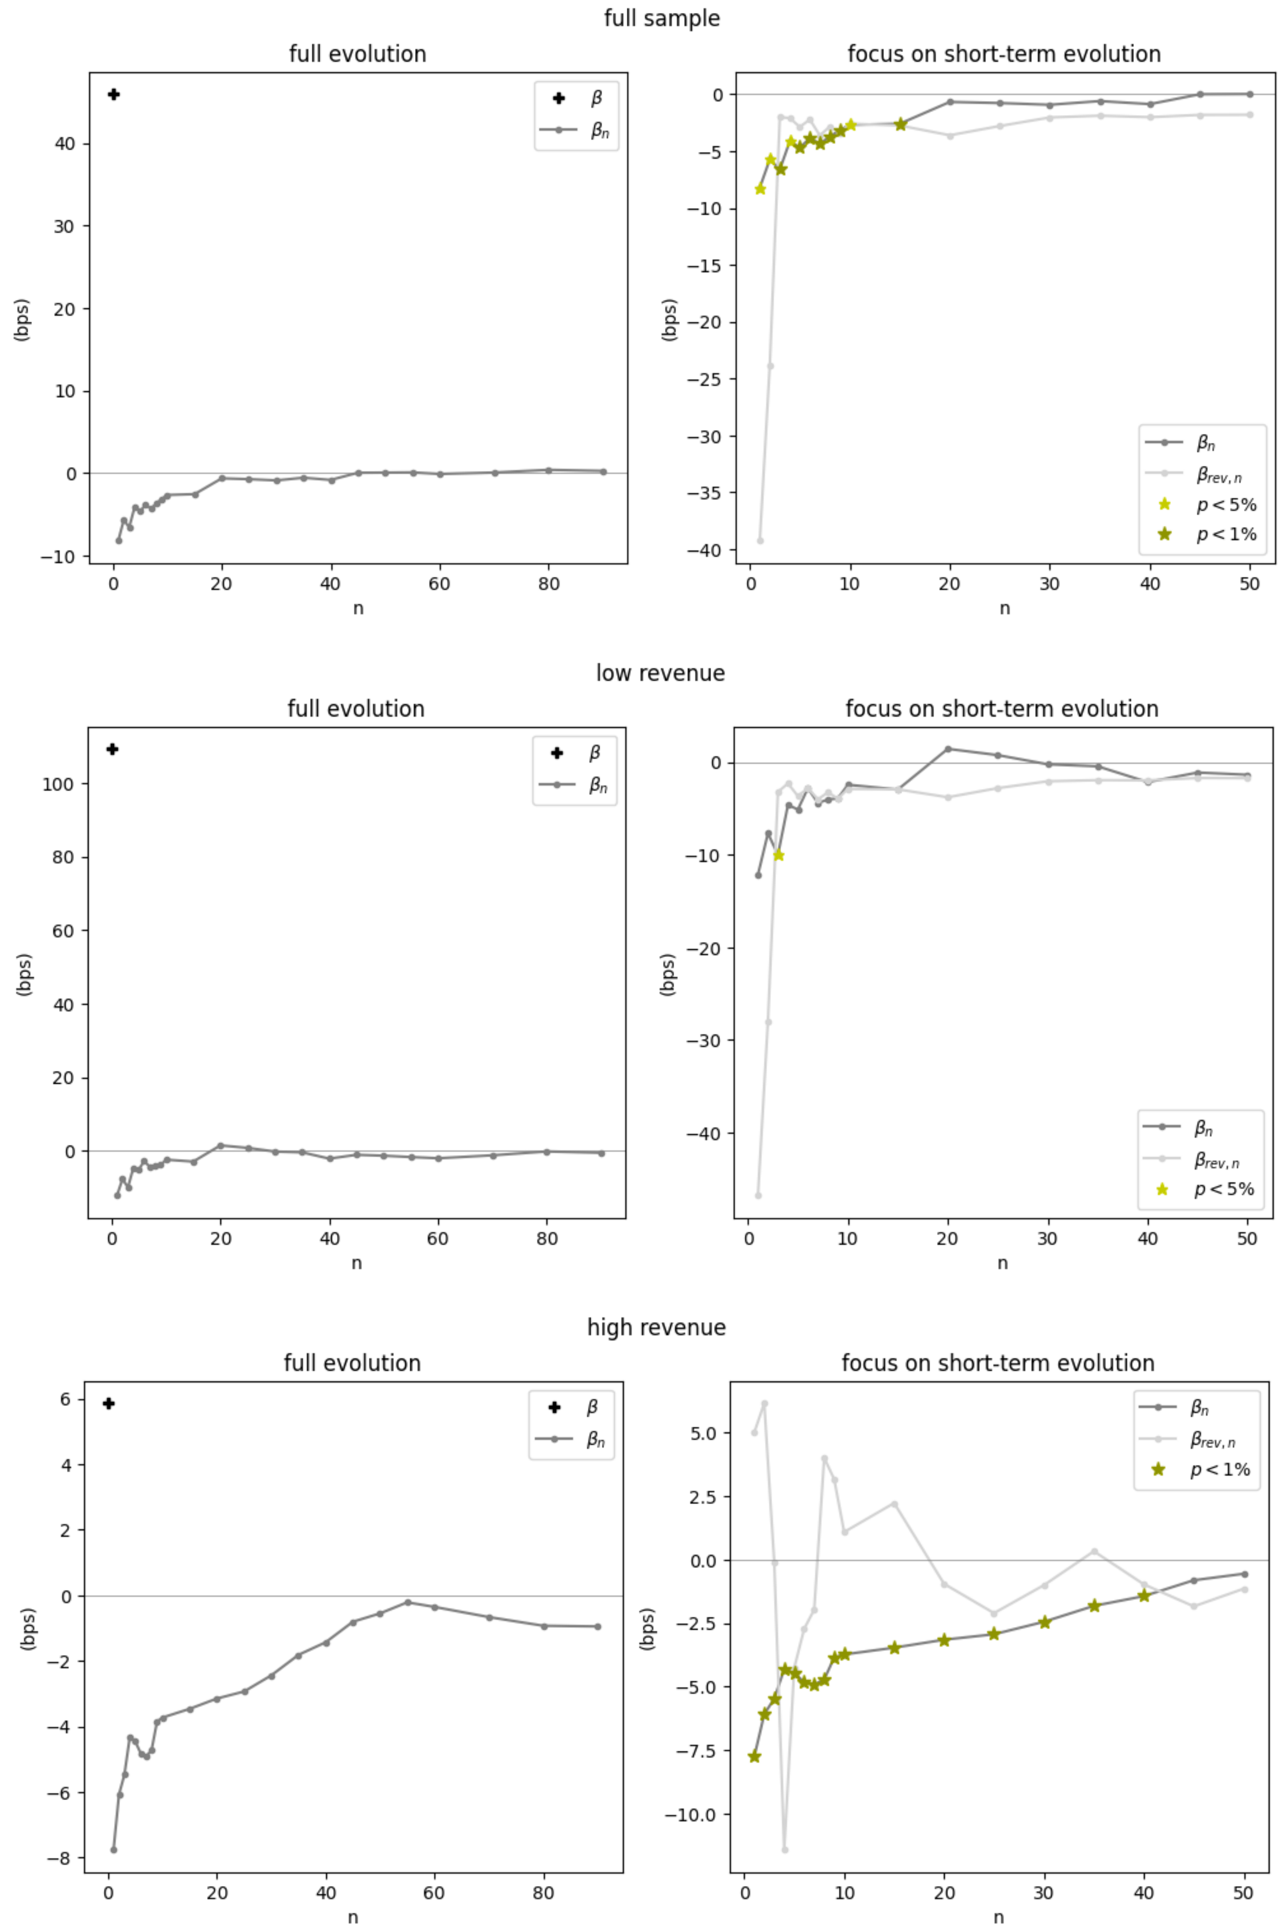

full sample
full evolution
focus on short-term evolution
(bps)
n
β
β_n
β_rev,n
p < 5%
p < 1%
low revenue
full evolution
focus on short-term evolution
(bps)
n
β
β_n
β_rev,n
p < 5%
high revenue
full evolution
focus on short-term evolution
(bps)
n
β
β_n
β_rev,n
p < 1%

Table 5: Regression of subsequent mean abnormal returns on mean sentiment and previous day's abnormal returns.

$MAR_{i,j+1\to j+n} = \alpha_n + \beta_n \times S_{i,j} + \beta_{rev,n} \times AR_{i,j} + \beta_X X_{i,j}$

n	**Full sample**		**Lower capitalization group**		**Upper capitalization group**	
	β_n (t-value)	$\beta_{rev,n}$ (t-value) in bps per std	β_n (t-value)	$\beta_{rev,n}$ (t-value) in bps per std	β_n (t-value)	$\beta_{rev,n}$ (t-value) in bps per std
1	-8.21**	-39.26***	-12.15	-46.77***	-7.76***	5.0
	(-2.35)	(-16.79)	(-1.53)	(-14.08)	(-2.63)	(0.67)
		-11.52		-13.72		1.47
2	-5.72**	-23.82***	-7.69	-27.97***	-6.08***	6.16
	(-2.27)	(-14.14)	(-1.31)	(-11.41)	(-2.89)	(1.16)
		-6.99		-8.2		1.81
3	-6.62***	-2.0	-10.0**	-3.23	-5.47***	-0.09
	(-3.24)	(-1.47)	(-2.1)	(-1.62)	(-3.2)	(-0.02)
		-0.59		-0.95		-0.03
4	-4.12**	-2.09*	-4.64	-2.29	-4.31***	-11.42***
	(-2.35)	(-1.79)	(-1.14)	(-1.35)	(-2.91)	(-3.06)
		-0.61		-0.67		-3.35
5	-4.68***	-2.87***	-5.11	-3.67**	-4.46***	-4.27
	(-2.92)	(-2.69)	(-1.36)	(-2.35)	(-3.36)	(-1.28)
		-0.84		-1.08		-1.25
6	-3.86***	-2.21**	-2.8	-2.81**	-4.83***	-2.72
	(-2.65)	(-2.29)	(-0.83)	(-1.99)	(-3.98)	(-0.89)
		-0.65		-0.82		-0.8
7	-4.34***	-3.62***	-4.42	-4.02***	-4.92***	-1.97
	(-3.23)	(-4.07)	(-1.42)	(-3.11)	(-4.37)	(-0.7)
		-1.06		-1.18		-0.58
8	-3.71***	-2.84***	-4.03	-3.27***	-4.72***	4.01
	(-2.95)	(-3.41)	(-1.38)	(-2.71)	(-4.45)	(1.51)
		-0.83		-0.96		1.18
9	-3.15***	-3.45***	-3.91	-3.94***	-3.85***	3.17
	(-2.65)	(-4.4)	(-1.42)	(-3.47)	(-3.84)	(1.26)
		-1.01		-1.16		0.93
10	-2.68**	-2.58***	-2.44	-2.89***	-3.73***	1.09
	(-2.38)	(-3.48)	(-0.94)	(-2.69)	(-3.9)	(0.45)
		-0.76		-0.85		0.32

For each subsample, the table displays coefficients β_n and $\beta_{rev,n}$ and associated t-statistics, for time horizons spanning from 1 day to 90 days. The coefficients are in bps per sentiment index unit for comparison. The results read: for one unit increase in respectively mean sentiment and rescaled abnormal returns, mean abnormal returns on the following n days are expected to increase by β_n and $\beta_{rev,n}$. For the reversal effect, we also included the normalized effect in bps per standard deviation for interpretation. The results hence read: for one standard deviation increase in abnormal returns, mean abnormal returns on the following n days are expected to increase by the standardized $\beta_{rev,n}$.
$^{*}p \leq 10\%$, $^{**}p \leq 5\%$, $^{***}p \leq 1\%$.

We observe that the reversal effect is significant in the two first days following sentiment expression for the lower capitalization group and the full sample, whereas is it not significant for the upper capitalization group. On the contrary, the sentiment effect is negative and significant for the upper capitalization group and the full sample, but not for the lower capitalization group.

Our explanation is that for larger capitalization, markets react to expressed sentiment although they had already priced in the information. As a result, they slowly correct this overreaction in the following days. On the contrary, for small capitalization the information provided on social media is novel, and gets correctly priced in.

Practicability of results

Can a trading strategy on LinkedIn sentiment produce superior returns?

The long-short portfolios confirm that trading on sentiment can produce superior returns, predominantly when it comes to large capitalizations. Indeed, performing a double-sort tends to increase returns and alpha compared to a simple reversal. However, sentiment information is better in combination with past returns information, as opposed to in replacement of it, as trading solely on sentiment delivers lower returns than trading on past returns.

Table 6 below displays the performance of the reversal, the sentiment reversal and the double-sort, for the three subsamples. Figure 8 shows cumulative returns for these portfolios.

Table 6: Performance of long-short portfolios formed on past returns and/or expressed sentiment.

annualized returns (%) alpha (%)	**Reversal**	**Sentiment reversal**	**Double-sort**
Full sample	40 39	20 24	34 38
Lower capitalization group	58 55	45 53	39 45
Upper capitalization group	2 2	5 6	18 21

The portfolios are formed with a holding period of 5 trading days. Stocks are selected if they are in on of the 2.5% extreme quantiles for single -sorts. For the double-sort, we use a quantile of 25% for past returns, and 10% for sentiment, which esures that the same number of stocks is selected.

We observe that the reversal does better than the two portfolios integrating sentiment information for the full sample and lower capitalization group. However, for the upper capitalization group, both the single-sort on sentiment and the double-sort perform better than the reversal.

The results are in line with the previous section, suggesting that sentiment reversal is strongest and most actionable when it comes to larger capitalizations. In the latter case, and estimating 10 bps of trading fees, a double-sort portfolio formed on past returns and sentiment would produce 8% returns.

Figure 8: Cumulative returns of the long-short portfolios.

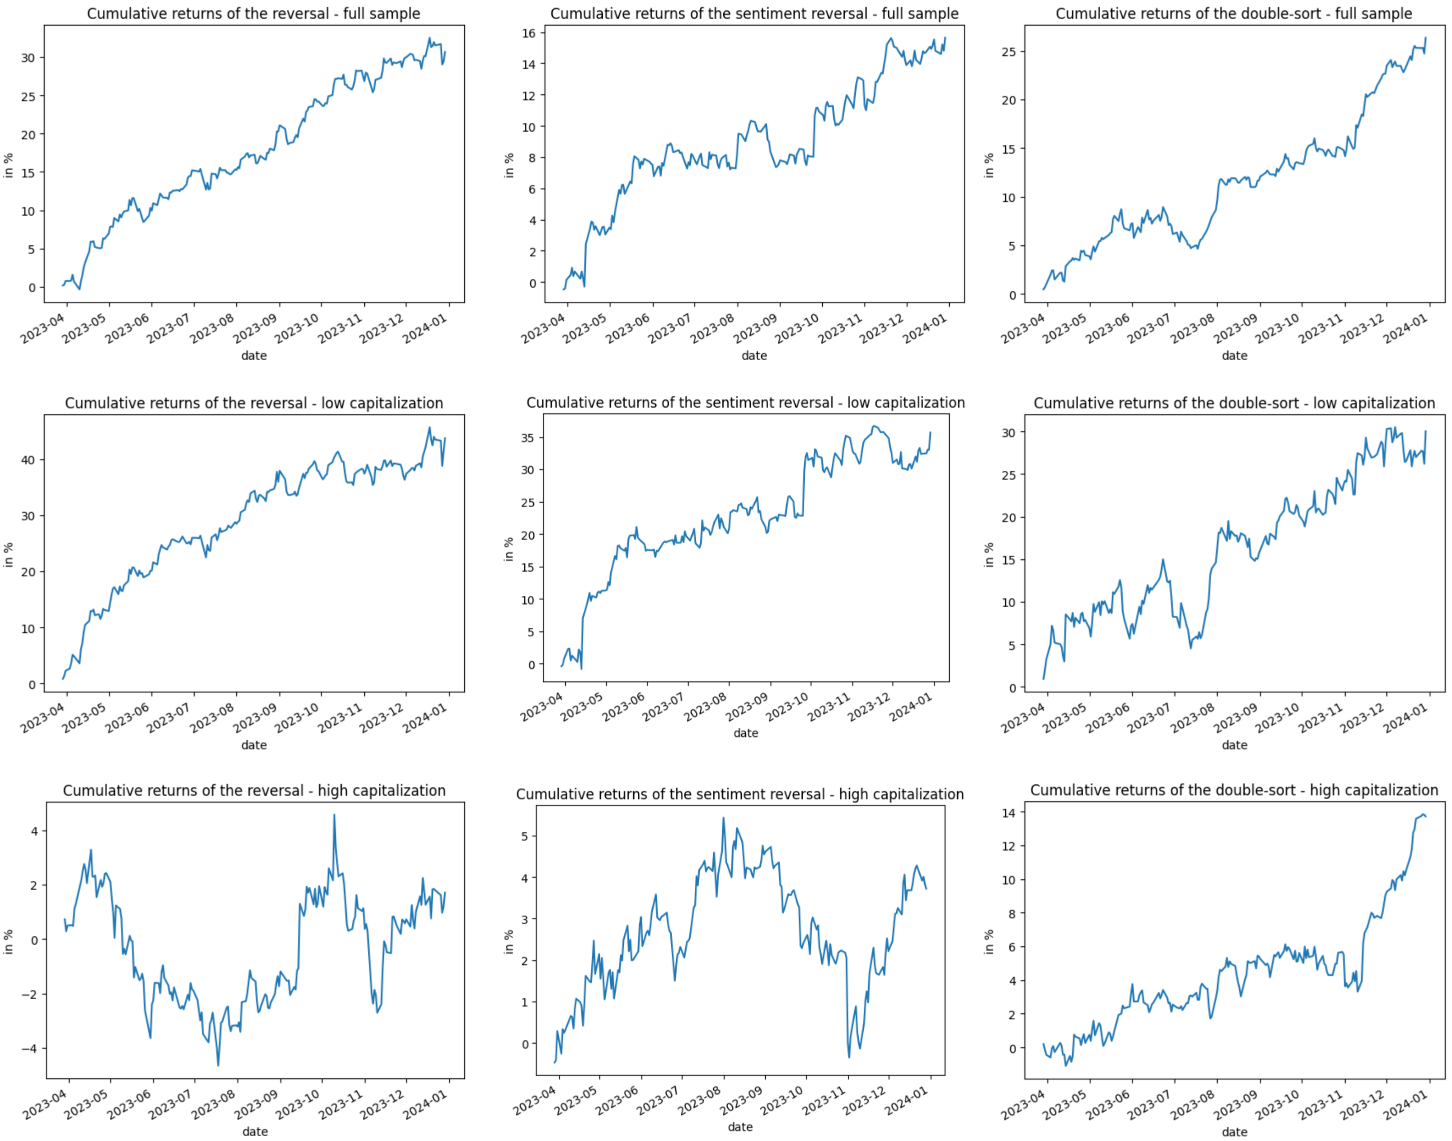

Posting dynamics and financial information

Do financial posting dynamics exhibit specific patterns and hold additional information compared to regular posting dynamics?

To check for any differences in posting patterns we start by looking at the posting times of non-financial somewhat financial and financial posts shown in figure 9. We can clearly see that financial posts have a higher tendency to be posted right after the close. This is normal as many earning announcements come after market hours, as this gives time for investors to read through the announcement's details before trading the next day. This observation is important as it confirms the legitimacy of our trading day convention in keeping the after close posts to the next trading day so that positive or negative earning announcements aren't leaked to past market performance.

Figure 9: Posting time as a function of the post's financial nature.

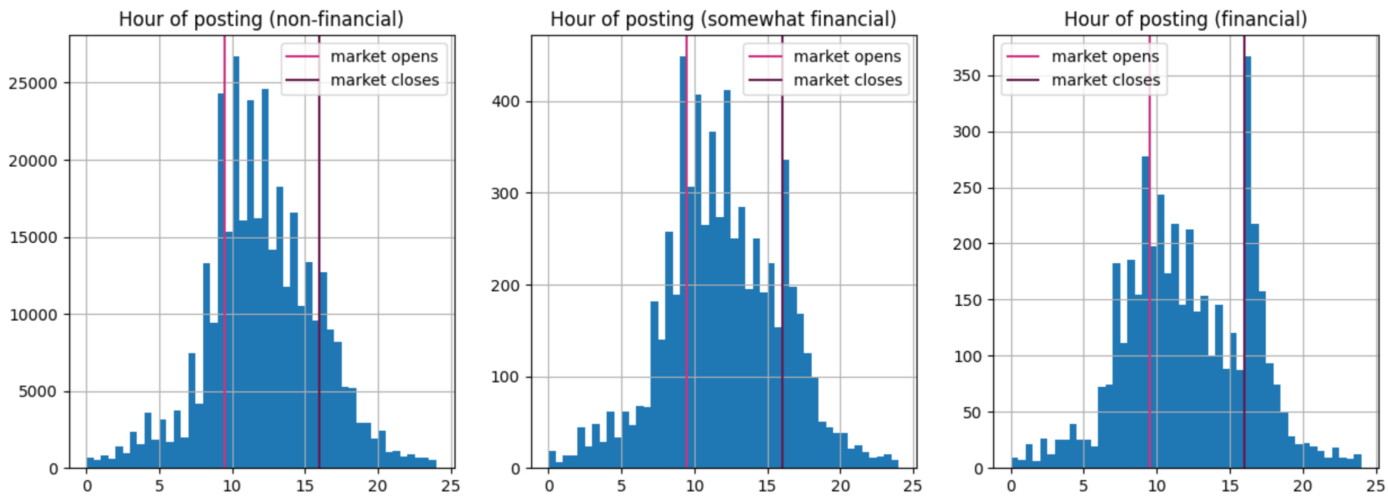

Furthermore, we notice that financial posts tend to have more extreme sentiment scores, as shown in figure 10. This increased variance isn't very significant but shows nonetheless that FinBERT allocates more extreme measures to posts that have a financial meaning.

Figure 10: Sentiment distribution for posts with different financial nature.

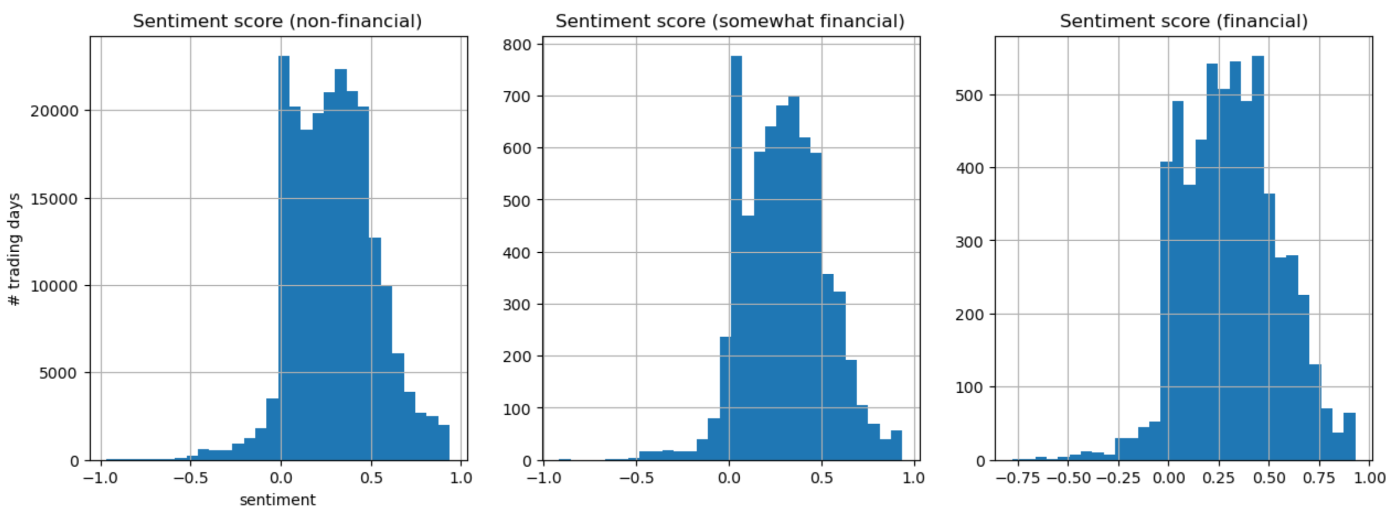

With that in mind, we now check if financial posts exhibit a higher explainability of same day and subsequent days abnormal returns. Table 7 displays the coefficients β, β_{sw_fin}, and β_{fin} in equation (3), quantifying respectively the main effect of sentiments on same day abnormal returns, as well as the additional effects of sentiments of posts that are somewhat financial and financial.

Table 7: Regression results of same-day returns on expressed sentiment and previous day's returns.
$AR_{i,j} = \alpha + \beta \times S_{i,j} + \beta_{sw\_fin} \times S_{i,j} \times \mathbb{I}_{sw\_fin} + \beta_{fin} \times S_{i,j} \times \mathbb{I}_{fin} + \beta_X X_{i,j}$

	β	β_{sw_fin}	β_{fin}
Value	42.79***	67.36**	40.43
t-value	4.85	2.09	1.25

Notes: The results are interpreted as follows. If sentiment of non-financial posts increases by 1 point, which corresponds to going from neutral to positive as the index goes, expected abnormal returns on the next day increase by β (in bps). On the other hand, if we see the same increase for a post that is somewhat financial then we can expect the abnormal returns to be $\boldsymbol{\beta} + \boldsymbol{\beta_{sw\_fin}}$ higher. The same reasoning applies for financial posts.
$*p \leq 10\%$, $**p \leq 5\%$, $***p \leq 1\%$.

Looking at the above results we notice that the beta of the cross-interaction term for somewhat financial posts is statistically different from 0 but the t-stat is much higher than the cutoff of 5% significance, whereas financial posts beta (β_{fin}) is statistically insignificant. Considering that somewhat financial posts represent 7.7% of all posts, we consider that their large count has contributed to the slight statistical significance of the beta. In other words, we do not find the results of this regression satisfying enough to conclude that financial posts have a stronger predictive power over same day abnormal returns.

This conclusion is interesting because this means that the statistically significant factors we get in subsequent regressions, are not really driven by financial posts. This means that the sentiment of posts with non-financial information is also related to their market performance. We think that this might be due to managers relaying positive internal sentiment through positive posts, financial or not, on their social media accounts. Furthermore, the low significance of financial post sentiments beta might be due to the fact that most financial information is already priced into the stock price. On the other hand, non-financial information might be more overlooked which is why it has more predictive power. To put things into perspective, this effect can be compared to that of the overperformance of analysts after they visit a company site (Cheng et al. (2016)[10]) or the predictive power of employee reviews on Glassdoor (Green et al. (2019) [4]).

Influence of stock market returns on sentiment expressed

How are company posts' financial sentiments influenced by recent stock market returns?

Running the regression of equation (4), we find evidence that post sentiment after market hours on a certain trading day and prior to next day's open, are positively correlated with the performance of the company stock on that day, as shown in table 8. In other words, companies tend to underline the performance of their stock in their communications. This can be explained by two factors:

- Company managers are happy about their stock performance and relay their sentiment online, in a conscious or unconscious manner
- Company managers try to leverage this positive performance to further boost their performance through this initial momentum. This can only be done consciously.

Table 8: Regression results of after close sentiment on past performance.
$S_{i,close,j\rightarrow j+1} = \alpha + \beta \times AR_j + \beta_X X_{i,j}$

	Full sample	Lower capitalization group	Upper capitalization group
$\boldsymbol{\beta}$	0.014**	0.009	0.026
t-value	2.04	1.04	1.62

Table 9: Regression results of after close sentiment on past performance split by past abnormal returns

	Full sample	Negative past abnormal returns	Positive past abnormal returns
$\boldsymbol{\beta}$	0.014**	0.025*	-0.003
t-value	2.04	1.70	-0.34

Notes: The results are interpreted as follows. If abnormal returns on day j increase by one standard deviation of mean sentiment, since the abnormal returns have been rescaled to the scale of mean sentiment scores, then the sentiment of posts after the market closes and before next day's open will increase by β. This applies to all regression splits as the only difference is the reference dataframe that is accordingly filtered for each different regression.

$^{*}p \leq 10\%, {}^{**}p \leq 5\%, {}^{***}p \leq 1\%.$

However, it is important to note that this result isn't robust enough to apply in all scenarios. Indeed, when we separate large companies from small companies, the effect is no longer statistically significant, as shown in columns 2 and 3 of table 8. Furthermore, the effect disappears when we split the regression between days with positive abnormal returns and days with negative abnormal returns, shown in table 9. If our initial finding was robust, we could have at least expected the beta to remain positively significant for one group of companies or for days when abnormal returns are positive. Indeed, it would have been understandable that companies do not relay and exacerbate negative performance but the fact that positive performance days don't have a statistically significant effect on after hours sentiment is a sign of weak robustness. This pushes us to assume that the initial correlation we had might be due to unconscious bias rather than managerial will of leveraging results.

Furthermore, this could mean that insider trading isn't significant in the studied markets. Indeed, if insider trading was widespread, we would have a significant positive correlation between abnormal returns and sentiment even after we split for positive and negative performance as insiders anticipate the upcoming positive/negative announcements. This clearly isn't the case as the correlation is statistically significant.

In a nutshell, performance does have an overall positive correlation with after hour trading sentiment, but this correlation isn't robust enough to work in all settings, which pushes us to assume that there is no conscious will of leveraging performance to boost company marketing. In addition, this is also a sign that insider trading isn't measurably significant in this market.

CONCLUSION

In conclusion, this study investigates the influence of financial sentiment expressed in social media posts by companies, particularly on LinkedIn, on daily stock returns. By utilizing the FinBERT NLP model, we analyze the sentiment of company disclosures and its predictive power concerning stock market movements.

Our findings indicate that the sentiment expressed in LinkedIn posts has a significant positive relationship with same-day stock returns. This suggests that the market reacts promptly to sentiment-laden information, with a tendency to overreact to positive news. Such overreaction is evidenced by the subsequent negative impact on cumulative abnormal returns in the following weeks. This dynamic is more pronounced in larger companies, where the effect of social media sentiment on stock returns is more persistent and cannot be subsumed by simple reversal.

Additionally, our research demonstrates that the information conveyed through non-financial posts can be just as impactful as financial disclosures. This highlights the presence of valuable insights in general corporate communications, which investors may leverage for decision-making.

Interestingly, we observe that stock market performance influences subsequent social media sentiment. However, this correlation is not robust when isolating days with extreme abnormal returns, suggesting that companies' social media activities are more reflective of general mood swings rather than deliberate attempts to influence market performance.

To test the practical implications of our findings, we develop a trading strategy that capitalizes on market overreactions by forming double-sort portfolios based on extreme sentiment expressions and extreme past returns. This strategy produces significant returns for large capitalization, where a standard reversal does not apply.

Overall, our study underscores the importance of incorporating social media sentiment into financial analysis. It highlights the potential of alternative data sources in enhancing the understanding of stock market behavior and developing effective trading strategies. Future research could further explore the mechanisms driving these dynamics and extend the analysis to other social media platforms and geographic markets.

REFERENCES

[1] Adams, T., Ajello, A., Silva, D., Vazquez-Grande, F. (2023). More than Words: Twitter Chatter and Financial Market Sentiment. Finance and Economics Discussion Series 2023-034. Washington: Board of Governors of the Federal Reserve System. ISSN 1936-2854.

[2] Wang, X., Xiang, Z., Xu, W., Yuan, P. (2022). The causal relationship between social media sentiment and stock return: Experimental evidence from an online message forum. Economics Letters. Volume 216, 110598.

[3] Sul H., Dennis A., Yuan L. (2017). Trading on Twitter: Using Social Media Sentiment to Predict Stock Returns. Decision Sciences, a journal of the decision sciences institute. Volume 48, issue 3.

[4] Green, C., Huang, R., Wen, Q., Zhou, D. (2019). Crowdsourced employer reviews and stock returns. Journal of Financial Economics 134 (2019) 236–251.

[5] Araci, D. (2019). FinBERT: Financial Sentiment Analysis with Pre-trained Language Models. University of Amsterdam.

[6] Calomiris, C. W., Mamaysky, H., Yang, R. (2020). Measuring the Cost of Regulation: A Text-Based Approach. National Bureau of Economic Research, working paper 26856.

[7] Devlin, J., Chang, M.-W., Lee, K., Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Google AI Language, arXiv:1810.04805.

[8] Malo, P., Sinha, A., Korhonen, P., Wallenius, J., Takala, P. (2014). Good Debt or Bad Debt: Detecting Semantic Orientations in Economic Texts. .Journal of the association for information science and technology, 65(4):782–796.

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

[10] Cheng, Q., Du, F., Wang, X., Wang, Y. (2016). Seeing is believing: analysts' corporate site visits. Review of Accounting Studies 21.