Norwegian School of Economics
Bergen, Spring 2021

Predicting Takeover Targets on Oslo Stock Exchange

An Extension to the Prediction Literature on the Norwegian M&A Market

Haavard Meyer Traa & Magnus Tønsberg

Supervisor: Karin S. Thorburn

Master thesis, Economics and Business Administration

Majors: Financial Economics & Business Analytics

NORWEGIAN SCHOOL OF ECONOMICS

Acknowledgements

This thesis was written as part if our Master of Science in Economics and Business Administration at The Norwegian School of Economics (NHH).

First, we want to express our deepest gratitude to our supervisor, Karin S. Thorborn, for the discussions and counseling leading us in the right direction. Her expertise in the fields of Corporate Finance and M&A have proven invaluable in the writing of this thesis.

Further, we would like to thank Magnus Lundstrøm at Oslo Stock Exchange for providing us with all the required stock exchange announcement data for our analysis.

Finally, we would like to thank Samfunns- og næringslivsforskning AS (SNF) and Børsprosjektet ved NHH for access to data on ownshership structure and stock-market data.

Norwegian School of Economics

Bergen, June 2021

Haavard Traa

Magnus Tønsberg

Abstract

The purpose of this research is to extend the previous literature on prediction of takeover targets on Oslo Stock Exchange. We aim to improve the prediction model developed by Khan and Myrholt (2018) by introducing new variables related to intellectual property and target-firm announcements, in addition to applying a gradient boosting algorithm (XGBoost). Furthermore, we seek to obtain positive abnormal returns from the predicted target portfolios. We examine the marginal contribution of predictive power from the new variables with logistic regression, using the predictors from Khan and Myrholt (2018) as control variables. Using the area under the Reciever Operating Characteristics curves (AUC-ROC) and Precision-Recall curves (AUC-PR) as performance metrics, we examine the difference in prediction skill between logstic regression and XGBoost. Furthermore, we evaluate the systematic risk adjusted returns of the target portfolios using the Fama-French 3-factor model.

We find that companies without a patent portfolio are significantly more likely to receive a takeover bid. Further, we observe that companies with a more negative tone in their announcements are more likely to be the target of a takeover bid. However, we only find significance at a 10% level. By applying the XGBoost algorithm, we observe better AUC-ROC and AUC-PR values. However, we cannot conclude that the improvements are significant. We are not able to achieve positive abnormal returns for the target portfolios individually. Moreover, in line with the performance metrics, we find no significant abnormal returns when comparing the portfolios from XGBoost to logistic regression.

Keywords – Intellectual property, Stock Exchange Announcements, Natural Language Processing, Mergers & Acquisitions, Takeover Target Prediction, Predictive Analytics, Machine learning, Fama French

Contents

List of Figures

List of Tables

1 Introduction

This section will firstly present our thesis's motivational background, both why the research is of interest and what we know from previous research. Secondly, we present our research questions, how these are analyzed, and how our thesis contributes to previous research on target prediction. Lastly, we give an outline of how the paper is organized.

According to the market efficiency hypothesis (Fama, 1970), the price of a stock should always reflect all publicly available information. In the event of a takeover bid announcement, new information is provided. Previous literature on takeover announcement returns[1] show that target shareholders on average receive a significant premium when a takeover bid is announced. Several studies have tried to capitalize on this knowledge by predicting takeover targets, with Palepu (1986) laying the foundation for the most commonly used predictors and prediction methodologies. The prediction of potential takeover targets can be used as a trading strategy. Furthermore, it is relevant to investment banks, institutional investors, and companies with an M&A-based growth strategy.

Khan and Myrholt (2018) are the first to predict takeover targets on Oslo Stock Exchange. They use the predictors from the hypothesis suggested by Palepu (1986), in addition to including macro factors relevant to the Norwegian market such as oil price and interest rates. Furthermore, they utilize logistic regression to create their prediction models. Khan and Myrholt (2018) manage to create a target portfolio consisting of 10% actual takeover targets, compared to the average takeover rate of 6.4% for the entire holdout sample. They find positive abnormal annualized returns of 0.9% (maximum target) and 1.8% (minimum classification) from their target portfolio using a yearly rebalancing strategy over four years. However, the abnormal returns are not statistically significant. This suggests that it is, to some extent, possible to predict takeover bids. There is currently no evidence indicating that it is possible to capitalize on the knowledge of takeover announcement premiums in the Norwegian market.

Following the empirical study of Khan and Myrholt (2018), our main research question is whether we are able to increase predictive power by extending their model. We first test

[1] See Jensen and Ruback (1983), Eckbo & Sollibakke (1991),Goergen and Renneboog (2004), Betton et al. (2008), Martynova and Renneboog (2011), Khan and Myrholt (2018), among others.

the marginal contribution of including variables created from patent data and target-firm announcements with logistic regressions, using Khan and Myrholt's predictors as control variables. Second, we introduce a gradient boosting algorithm (XGBoost) and compare its predictive skill to logistic regression. Finally, we create target portfolios based on the best-performing prediction models and test whether we are able to achieve positive abnormal returns adjusted for systematic risk factors.

Out thesis contributes to previous research by introducing new variables related to intellectual property and target-firm announcements. Additionally, we introduce a different methodology by utilizing machine learning algorithms able to detect more complex relationships between the features in our data.

The paper is organized as follows: In section 2, we present the literature used to motivate the construction of our research hypotheses. In section 3, we formulate our research hypotheses. Section 4 provides an extensive overview of how data is collected and processed, in addition to descriptive statistics. In section 5, we present the methods used to test our research hypotheses. In section 6, the empirical results of our analyses are presented. Finally, in section 7, we conclude our research findings and provide recommendations for further research.

2 Literature Review

In the following section, we present literature on intellectual property and stock exchange announcements in relation to M&A transactions to motivate the research on these topics to predict takeover targets. Furthermore, we discuss potential areas within these topics that may provide insightful data and methodologies that can be used to create variables included in our predictions. Following, we present studies on target takeover announcement returns, previous attempts to predict takeover targets, and lastly, the performance of portfolios assembled by target takeover prediction models.

2.1 Intellectual Property and Patents in M&A

Intellectual property has not always been a noteworthy part of deal negotiations and valuations when considering potential acquisitions. Especially in terms of valuation, one would assume that the value of intellectual property was included in the forecast of future cash flow (Mousavi, 2011). However, the valuation of intellectual property is both complex and highly uncertain. Thus, the intellectual property part of deal negotiations was often neglected (Mousavi, 2011).

"As intellectual property – particularly patents – is increasingly viewed as an asset with distinct and separable value from the businesses of targets, we are likely to continue to see the role of intellectual property evolve and gain relevance in driving, facilitating and financing M&A transactions for the benefit of both targets and acquirers." (Mousavi, 2011, p. 7)

According to Mousavi (2011), three important elements coherently have created the foundation for the evolving role of intellectual property in M&A transactions: value of intangible assets, patent trading markets, and valuation of patents. First, Ocean Tomo (2020) shows that from 1975 to 2020, the average value of intangible assets for S&P500 companies has increased from 17% to 90% of total company values. This development holds for the S&P350 Europe index to a certain degree and reaches a relative market value of 74% in 2020 (Ocean Tomo, 2020). The growth of intangible asset values is in large related to intellectual property, including patents, trademarks, and copyrights. Secondly, the market for trading patents has become more liquid, leading to managers and other stakeholders not necessarily viewing the patents as intrinsically tied to the target firm. The increased trading calls for more accurate valuations of the patents themselves. Lastly, more robust and sophisticated methodologies for patent valuation have been developed.

The study of intellectual property and especially patents in relation to M&A activity is scarce, despite the described evolving importance of the asset type. Frey and Hussinger (2006) study the role of technology in cross-border- and domestic M&A deals in Europe between 1994 and 2000. They show that in cross-border acquisitions, the target patent portfolio is of high value to the acquirer if the patents concern innovative technology in technology fields related to the operations of the acquiring firm. However, they do not find evidence of patents being attractive to the acquirer in domestic deals or if there is no technical proximity between the merging partners. In 2018, 41% of Norwegian targets were acquired by foreign investors (Hellebust and Sandvik, 2020). This is 50% more than the share of cross-border deals worldwide in the same year (Statista, 2020b), supporting the potential relevance of patents in the Norwegian M&A market.

In another more recent study, Alimov and Officer (2017) find a significantly positive impact of intellectual property rights protection on cross-border M&A flows. In cross-border deals, the acquirer is often driven by either transferring its intellectual property to the target company to improve the target's production processes or gaining access and ownership to the targets' intellectual property portfolio (Markusen, 1995). The study indicates that the intellectual property of a target company is of importance to the acquirer in cross-border deals when the target is in an industry that uses intellectual property intensively. At the same time, the evidence is weaker for domestic transactions (Alimov and Officer, 2017).

In more specific research on the use of patents to identify and evaluate potential takeover targets, several studies are analyzing the technological relatedness between potential targets and acquirers (Huang et al., 2013; Ma et al., 2017; Park et al., 2013). These papers have an acquiring company's point of view when searching for targets. The studies are based on previous papers suggesting that the maximum benefits from an acquisition can be realized when the technological assets of the target are related to the acquiring firm (Hussinger, 2010). It is crucial to understand how an acquirer searches for targets when developing methods to include the patent variable in a takeover prediction model.

In order to identify potential targets, Park et al. (2013) collect all the patents of a technological industry broadly covering the technological fields wherein the acquiring business is operating. A natural language processing method converts each patent's abstract (description text) into SAO structures. SAO stands for subject-action-object and is a way to convert unstructured textual data to structural data comprised of subjects, actions, and objects (Kim et al., 2020). The patents are placed on a two-dimensional map based on the SAO structures' technological proximity and then divided into technological areas using the K-nearest neighbor classification method. Based on the density of the acquiring company's patents within each technological area, the areas are divided into three categories: enhancement of core technology, enhancement of sub or minor technology, or entry into new technology areas. Thus, based on the acquiring company's M&A strategy, potential targets can be found by searching for companies possessing patents within the clusters labeled with the desired growth strategy category.

Huang et al. (2013) exploit the method introduced by Park et al. (2013). In addition to the advanced SAO patent mapping, they also introduce a formula comparing the IPC-classes of a company's patents to a potential target company to quantify the overall technological similarity. The International Patent Classification is a means to obtain an internationally uniform classification of patent documents (WIPO, 2020). The system's primary purpose is to ease the searching process for intellectual property offices and other users when retrieving patent documents (WIPO, 2020). Thus, the patent classification is structured to make it easier to check for similarities. However, the system is not perfect, and countries tend to interpret the classification system somewhat differently. In Appendix A1.1, an excerpt from the "Guide to the IPC 2020" shows how a patent class code is assembled (WIPO, 2020). Anyhow, this is a less resource-demanding method compared to the SAO structure and can serve as a great pre-processing tool for raw data to collect potential targets. Huang et al. (2013) show that the IPC-class method ranks the technological similarity between the investigated companies similar to the more advanced SAO structure method, indicating that the use of IPC-classes is a good proxy.

The two studies above require a prediction of potential acquirers for the models to be incorporated in target prediction. Ma et al. (2017) add to the research, introducing a target-firm side indicator. They look at the technological quantity- and quality of a firm based on the number of granted patents and the average citation frequency of the patent portfolio. In addition, they also look at a firm's R&D capability in the form of the number of patent inventors in the firm and their R&D productive efficiency based on the average patent records per inventor.[2]

To our knowledge, only one study has incorporated patent analysis in a takeover prediction model. Bourne et al. (2019b) introduces a variable they call "technological drift". The economic rationale for the variable is that a company exploring new research sectors is likely to be in a transitory period of its business. They evaluate the year-to-year change of a company's technological profile by comparing the proportion of patents of each IPC-class granted in the last five years with the similar proportion calculation one year prior. The formula will be described in detail in the methodology section (4.1.4.2). The study finds that companies with a high technological drift are significantly less likely to be acquired.

Studies on target prediction have yet to explore the general patent distribution of acquired targets using IPC-class codes. We will test whether the IPC-classes and quantity of patents are related to takeover probability. Additionally, we will test for the relevance of the technological drift variable. One issue with this method is that a patent application is published 18 months after the filing date. In addition, it takes time to invent the patented products. Thus, the described transitory period has likely begun a long time before the patent is published. Consequently, there may be another rationale for the observed statistical significance from that study.

2.2 Textual Analysis in Finance

In recent years, different applications of textual analysis with the use of machine learning methods such as natural language processing (NLP) have been widely researched in finance literature. Loughran and McDonald (2020, p.1) state that "Textual analysis, implemented at scale, has become an important addition to the methodological toolbox of finance.". NLP has many interesting applications, but this thesis will limit its focus to sentiment analysis due to its computational intensity.

Sentiment analysis aims to determine the emotional tone of a text by utilizing a pre-defined dictionary containing words assigned a continuous sentiment score or words classified as,

[2]These are all four highly interesting factors that could be added to our primary model. However, the source of our patent data does not include the number of citations nor the patent applicant (inventor).

e.g., positive, negative, or uncertain. A widely used dictionary in sentiment analysis is the Harvard IV-4 Psychological Dictionary. However, as demonstrated by Loughran and McDonald (2011) in the paper "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks", the Harvard Dictionary is not suitable for all categories of text. When analyzing over 50 000 10-K forms, they found that 73.8% of the words classified as negative in the Harvard Dictionary are typically not negative in a financial context. Loughran and McDonald (2011) construct alternative word lists specifically tailored for texts in a corporate or financial context. This dictionary is known as Loughran and McDonald's Master Dictionary and is widely used when analyzing the sentiment of corporate texts.

Graffin et al. (2016) found evidence that acquiring firms tend to release positive but unrelated announcements in the period leading up to and surrounding the event of a takeover announcement. When managers anticipate shareholders and potential investors to react negatively to the announcement of an acquisition, they use this impression offsetting technique to 'inhibit' the observers' negative perception of an event by distracting them with good news. Graffin et al. (2016) find evidence that impression offsetting reduces the negative market reaction to acquisition announcements by over 40%.

Takeover prediction literature based on textual analysis is limited. Routledge et al. (2017) examine the effects of combining variables from textual analysis with more traditional financial measures when predicting acquirers and takeover targets in the period 1995-2011. Their acquirer prediction model found that combining traditional financial measures and text regression outperforms their baseline model, focusing solely on financial measures, increasing the pseudo-R^2 from 0.0696 (baseline) to 0.1022 (combined model). The findings from the target prediction model are not as promising, with a pseudo-R^2 of 0.0294 from the best combined model and a pseudo-R^2 of 0.0262 from the baseline model.

Katsafados et al. (2021) seek to identify U.S. listed banks that participate in a merger transaction in the subsequent year based on textual analysis of the banks' annual reports filed between 1997 and 2015. Their sentiment calculations are based on Loughran and McDonald's lists of positive and negative words from the Master Dictionary. Using logistic regression, they examine the likelihood of the bank becoming a target or acquirer. In both cases, they find statistically significant coefficients for the sentiment variables after controlling for financial characteristics of the banks, indicating potential predictive power.

The existing literature applying NLP in finance focuses mainly on the U.S. markets. To the best of our knowledge, there are very few studies focusing on stock exchange announcements in the Norwegian market. Medby and Nordgård (2019) apply NLP to stock exchange announcements from Oslo Stock Exchange, attempting to estimate stock price changes. Using their best-performing model in a trading application with a long/short strategy, they found that it is possible to achieve excess returns utilizing NLP.

2.3 Takeover Announcement Returns

The abnormal return in the event of a takeover announcement is a widely studied phenomenon. According to the efficient market hypothesis, share prices should reflect all available information (Fama, 1970). Thus, when a takeover is announced, new information is incorporated into the stock price. Table 2.1 shows the main results of the studies discussed in this section, while we go more into detail on findings we find relevant to our research below.

One of the earliest celebrated studies on abnormal return for target firm shareholders in corporate takeovers was conducted by Jensen and Ruback (1983). Preceding popular event studies on this issue, such as Mandelker (1974), Ellert (1976), and Langetieg (1978), used the final approval from target shareholders as their event day. However, Jensen and Ruback (1983), and subsequent studies show that most of the abnormal return is earned already on the first public announcement date of the potential takeover. Jensen and Ruback (1983) collect and summarize the results from 13 studies where target-firm stockholder returns are researched. They find that, on average, successful tender offers generate an abnormal return of 29.1% in the month or two surrounding the first announcement date. The paper also shows that unsuccessful acquisitions gain approximately the same abnormal returns in the months surrounding the first announcement date. However, when extending the holding period until the bid is declined, the shareholders experience a negative abnormal return of -2.88% on average.

Another more recent meta-analysis study on takeover announcements in the U.S. can be found in Betton, Eckbo, and Thorburn's (2008) "Corporate takeovers". They find an abnormal return of 6.8% in the runup period [-41, -2] to the announcement, with a CAAR of 14.6% for the announcement event [-1, 1]. This is lower than Jensen and Ruback (1983), indicating a negative trend in takeover abnormal returns. However, Betton et al. (2008) compare the CAARs' of the period 1991-1995 with 1996-2000, where they find a higher abnormal return for the later period. When examining the transaction market in the U.S. closer, we see that in the period 1955-1975, when most of the studies examined by Jensen and Ruback (1983) are conducted, there is a surge in M&A activity, peaking around 1970 (Dieudonne et al., 2014). While the transaction activity was relatively low at the beginning of the '90s, it steadily increased during the dot com bubble (Dieudonne et al., 2014). Seemingly, the magnitude of the average abnormal returns is correlated with the overall transaction activity in the market, with higher abnormal returns in periods with higher activity.

Goergen and Renneboog (2004) research the M&A market returns in the U.K. and continental Europe. In total, they find significant announcement effects of 9.01% for the target firm on the event day. Including a two-month runup period before the announcement date, they observe an average cumulative abnormal return of 23.1%, which is in line with the Betton et al. (2008) study from the same period. We observe that the runup period CAAR [-41, -2] in the Betton et al. (2008) study constitutes around 1/3 of the total runup plus announcement CAAR, while for the Goergen and Renneboog (2004) study, the same runup CAAR [-40, -2] constitutes around 3/5 of the total runup plus announcement CAAR. Many factors can explain this dissimilarity; one potential reason is a higher pre-announcement deal leakage in Europe.

Goergen and Renneboog (2004) find a slightly higher CAAR for domestic deals than cross-border deals (22.7% vs 19.8%). Additionally, they observe a significantly higher return for hostile acquisitions than friendly (30% vs. 22%, only tested for the U.K.). The authors suggest that the effect comes from the opposition against bids, leading to a revision of the offer and a higher bid premium. Lastly, the study also finds that post-announcement abnormal returns are not statistically different from zero based on all the observations and significantly negative for continental Europe, in line with the Jensen and Ruback (1983) study.

Renneboog extends the research further with Martynova (2011), including a significantly higher number of observations from the fifth takeover wave (1993-2001) in Europe. They find a significant abnormal return of 11.49% in the runup period [-40, -1] and 9.13% for the event day [T = 0]. The slightly lower abnormal returns compared to the Goergen and Renneboog (2004) paper can be explained by the lower density of U.K. transactions in this study. Martynova and Renneboog (2011) report an abnormal return of 14.2% runup and 13.66% on the event day for U.K. transactions compared to 7.47% and 4.45% respectively for Continental Europe. This suggests that the U.K. market is more similar to the U.S., while Continental Europe, including Norwegian targets, yield lower takeover returns. In line with Betton et al. (2008), Martynova and Renneboog (2011) also find significantly higher returns closing up to the burst of the dot-com bubble compared to the early '90s.

To our knowledge, there have been conducted two studies on takeover returns in Norway. Eckbo and Solibakke (1991) find a cumulative abnormal return of 4.88% in the runup period [-40, -2] and 3.9% for the event period [-1, 0]. The merger activity in the market was low during this period (1983-1989) compared to what we see nowadays, which might explain the low abnormal returns.

Khan and Myrholt (2018) conduct a study on takeover returns in Norway in their paper on target takeover prediction. The authors find a non-statistically significant CAAR of 2.8% in the runup period [-50, -1], with a statistically significant CAAR of 12.4% for the event period [-1, 1]. However, excluding the natural resources sector, the study finds a significant CAAR of 6% for the runup period and a significant CAAR of 12.8% for the event period. They suggest that the negative impact from natural resources comes from the industry consolidating during cyclical downturns but showing a significant abnormal return for the event period. As the runup- and event period overlap, the exact returns of the runup relative to the event period are uncertain, especially as there is often observed a significant abnormal return during the T-1 event day. Nevertheless, it is reasonable to estimate that around 1/3 of the CAAR stems from the runup period excluding natural resources. Compared to the other studies we have reviewed, the relative runup CAAR is closer to the findings in the U.S. than in the U.K. and continental Europe. Khan and Myrholt (2018) suggest that a more concentrated ownership structure can explain this in Norway, with fewer parties involved in deals, reducing the probability of pre-announcement information leakage.

To conclude, the findings of zero or negative abnormal returns in the period from the day after the announcement and beyond have implications for creating an investment strategy for the target portfolio, suggesting that a position should be closed rapidly after an acquisition announcement is made. In general, the studies mentioned above suggest that we are likely to observe lower abnormal returns in Norway and continental Europe compared to the U.K. and the U.S. due to fewer hostile bids (Aabø-Evensen, 2021). The high percentage of cross-border transactions in Norway also suggests lower abnormal returns. Moreover, the observable abnormal return from the runup period before an announcement substantiates the importance of predicting targets at an early stage. However, the studies show that this is less important in Norway and continental Europe compared to the U.K and the U.S. The last thing we take from these studies is that we are likely to observe higher abnormal returns in periods when deal activity is high.

Table 2.1: Previous Studies on Takeover Announcement Returns

Study	CAR	Event window	Sample size	Sample period	Country coverage
Jensen & Ruback	29.1%	[T +- 30]	13 studies 88 <N <302	1956-1981	US
Betton, Eckbo & Thorburn	14.6% 6.8%	[-41, -2] [-1, 1]	16 studies N >1000	1980-2005	US
Goergen & Renneboog	23.1% 9.01%	[-40, 0] [-1, 0]	136	1993-2000	UK & Continental Europe
Martynova & Renneboog	11.49% 9.13%	[-40, -1] [T = 0]	760	1993-2001	UK & Continental Europe
Eckbo & Sollibakke	4.88% 3.9%	[-40, -2] [-1, 0]	193	1983-1989	Norway
Khan & Myrholt	6%* 12.8%*	[-50, -1] [-1, 1]	136	1995-2012	Norway

Table 2.1 shows an overview of the aforementioned studies. It is undoubtedly well supported in academia both historically and in recent times that target shareholders receive substantial positive abnormal returns in the event of a takeover announcement and during the runup period before the announcement. However, as we will review in the following chapters, creating successful strategies to capitalize on this knowledge as an outside investor is less clear.

2.4 Target Prediction & Portfolio Construction

There have been several studies dedicated to predicting takeover targets over the last decades. In this section, we review some of the most influential and cited studies to form the basis for our hypothesis. We found “Predicting takeover targets” written by Palepu

(1986) to be particularly interesting. Palepu points out issues in the methodology used in prior studies and argues that these flaws cause biases that affect the results. After being published, Palepu's (1986) study has been frequently referenced in the literature. To avoid the methodological flaws discussed in the study, we dedicate a subsection of this literature review to these methodological flaws. This section contains three subsections: Palepu's methodological flaws, empirical evidence on predictors, and investment strategy.

2.4.1 Palepu's Methodological Flaws

Palepu (1986) dedicates a section to discuss methodological issues in the acquisition prediction literature. In this section, he discusses how these flaws impact the results of studies and suggests alternative methodologies to avoid them.

The first issue raised is how samples are typically drawn in prior studies. First, Palepu (1986) focuses on the sample used for the model estimation, stating that the samples were typically drawn so that the number of targets and non-targets were approximately equal, referred to as state-based samples. This is usually done due to the large number of non-targets in the population compared to the very few targets. A true random sample would likely be heavily imbalanced to the side of non-targets, giving the model estimation little information and resulting in imprecise parameter estimates. In the paper, Palepu (1986) acknowledges that the state-based sample method can be more efficient, require a smaller sample, and can be justified for model estimation through econometry. However, he argues that earlier studies fail when using state-based samples in conjunction with estimators assuming random sampling, causing biases and incorrect inferences. To avoid these issues, Palepu (1986) suggests modifying the estimators employed. Instead of using simple maximum likelihood (MLE), stated to be the practice in acquisition prediction literature, it is referred to two alternatives discussed by Manski and McFadden (1981): the conditional maximum likelihood estimator (CMLE) and the weighted maximum likelihood estimator (WMLE).

The second issue Palepu (1986) raises regarding sampling is the use of state-based samples when testing the estimations. Unlike in model estimation, where the significant imbalance of targets and non-targets justifies using a state-based sample to enhance the information content, state-based samples cannot be validly justified in test samples. When evaluating a model's forecasting ability, the expected error rate of the model's forecasts is usually the metric of choice. As discussed above, a state-based sample is not truly random and does not accurately represent the distribution of targets and non-targets in the population. Thus, using state-based samples in testing model predictions leads to error rate inferences not generalizable for the population. In other words, the difficulty of correctly predicting targets is not accurately presented. To avoid this issue, Palepu (1986) suggests using the entire population in a period as a test sample. He argues that using a large sample for prediction testing is not unrealistic, as the model parameters are already estimated, and the computational cost of prediction testing is relatively low.

The final methodological issue raised is the use of arbitrary cut-off probabilities in the prediction testing of a model. As stated in the paper, prediction testing is not necessary if the intention of the model is to solely test the statistical significance of variables in relation to a firm's acquisition probability (Palepu, 1986). However, for models intended to predict potential targets, prediction testing is vital for model development. When determining potential targets, the estimated takeover probability is compared to a pre-defined cut-off probability. The issue Palepu (1986) raises is that the prior acquisition prediction studies use an arbitrary cut-off probability, failing to state or consider the decision context for which the model is built. This makes the prediction tests hard to interpret as it is unclear what the observed prediction accuracy indicates. To avoid this issue, the paper suggests defining a decision context for the model's results and deriving the optimal classification scheme according to standard decision theory methodology. The use of cut-off probabilities will be explored further in the following sub-section (2.4.3).

2.4.2 Sample Selection, Independent Variables & Methodologies

Palepu's (1986) study has become widely celebrated in acquisition prediction literature and has formed the basis for many subsequent studies of the topic. Thus, Palepu's (1986) proposed target prediction model using nine independent variables was a natural starting point when discussing empirical evidence.

Palepu's (1986) choice of variables, was based on the following six hypotheses: inefficient management, growth-resource mismatch, industry disturbance, size, market-to-book ratio, and price-earnings ratio. Many subsequent studies have proposed additional hypotheses and variables in attempts to achieve better predictions. As the rationale behind the added hypotheses and variables are extensively described in the literature, this paper will not go into the finer details of each variable but rather give an overview of the development in target prediction literature.

Ambrose and Megginson (1992) added the ownership structure hypothesis. By examining institutional and insider shareholdings measures, Ambrose and Megginson (1992) find no significant effects on the takeover probability when using absolute levels. However, the percentage change in institutional shareholdings in the quarter preceding a takeover bid is found to be statistically significant. Although the paper includes hypotheses and variables regarding takeover defenses, these are not included in our thesis. Most takeover bids in Norway end up being recommended by the board of directors (Aabø-Evensen, 2021). Thus, the time-consuming process of gathering information about takeover defenses is expected to yield relatively low reward in terms of prediction accuracy.

The remaining firm-specific variables, focus on the leverage and liquidity of the firm. Arguing that financial distress increases a company's probability of becoming a target, Brar et al. (2009) extend Palepu's 1986 model to include a leverage hypothesis. However, examining both a company's level of debt in the balance sheet and the change in debt over the year leading up to the acquisition announcement, they find no support for these hypotheses in their data. In the same paper, Brar et al. (2009) extend the model further by adding a liquidity hypothesis, finding significantly lower cash to total capital ratio for the targets in their data compared to the non-targets.

Khan and Myrholt (2018) test the firm-specific hypotheses discussed above in their analysis of the Norwegian market, in addition to several macroeconomic factors shown in the literature to affect M&A activity. In selecting variables for the prediction model (presented in section 5.1), Khan and Myrholt (2018) refer to frequently cited studies of macroeconomic conditions' effects on deal activity and extensively discuss the rationale for each variable. As our thesis mainly focuses on the model's predictive power and not the causality of predictors, the rationale for each variable will not be discussed. However, one macroeconomic factor, perhaps particularly important when studying the Oslo Stock Exchange, is oil prices. EY (2017) argues that both low- and appreciating oil prices have positive effects on M&A activity. On one side, it is argued that high debt ratios, lower valuation, and restructurings drive M&A activity when the oil price is believed to have reached a minimum. On the other side, in periods with appreciating oil prices, companies' focus on expansion and growth is argued to be the driver for deal activity.

Like Palepu (1986), all the subsequent studies discussed thus far have opted to employ logistic regression in their takeover prediction models but differ in the structure and periods used to form their samples. Common for most of the studies is that they employ a relatively small number of independent variables. This thesis intends to test whether more advanced algorithms can produce better predictions. In many cases, more advanced algorithms can detect and consider more complex and subtle patterns in the data (Shmueli et al., 2010). Thus, we have chosen to include all the variables presented above in our model building, even if they have not been found to have a statistically significant effect in the logistic regression models.

Considering that our thesis will employ a different methodology than the studies discussed so far, it is relevant to review newer acquisition literature applying machine learning methods in prediction models. Bourne et al. (2019a) published a report from a research project predicting U.S. listed, publicly-traded, takeover targets. They find that by employing a Random Forest model, a Neural Network model, and an ensemble of the two with quarterly data, they achieve a portion of targets in their predicted portfolio six times higher than expected from a random estimator. The ensemble method is shown to be the best performing model, correctly classifying 10.8% of predicted targets in their portfolio. The variable selection for one of our models is heavily influenced by this study and is shown in Table 4.1.

2.4.3 Cut-off Probabilities & Portfolio Construction

Many studies developing acquisition prediction models are motivated to test whether the results can be used as an investment strategy, capitalizing on the abnormal returns following an acquisition announcement. In this section, we focus on empirical evidence of forming investment strategies based on acquisition prediction models.

Assuming the markets are efficient, according to the efficient market hypothesis, the market's expectations of a firm's takeover probability and value creation from the potential deal should already be reflected in the stock price. Thus, for an investment strategy based on an acquisition prediction model to be successful, both the model's predictions and the user's evaluation of deal accretion must be different and more accurate than the market's beliefs.

The final methodological flaw presented in Palepu's (1986) paper focuses on using arbitrary cut-off probability in the prediction testing of the model. Although the cut-off probability does not affect the model's results, it plays a central part when constructing a portfolio based on the results as it determines which companies to include. As mentioned, Palepu (1986) suggests deriving the optimal classification scheme according to standard decision theory methodology, where the cut-off probability is set to the probability yielding the highest accuracy rate in the training sample. Palepu's (1986) determination of the optimal cut-off probability assumes that the expected cost of Type I (target incorrectly classified as non-target) and Type II (non-target incorrectly classified as target) errors are equal. Powell (2001) argues that the penalty of Type II errors is significantly smaller than the pay-off from avoiding Type I errors and suggests that the optimal portfolio selection criterion should be to maximize the proportion of target firms in the portfolio. This is done by dividing the estimation sample into deciles sorted in descending order of takeover probability. The decile containing the highest concentration of targets sets the optimal cut-off probability equal to the first takeover probability in the selected portfolio. We believe the method suggested by Powell (2001) is the optimal way, further justified in the following chapter.

2.5 Target Portfolio Performance

Many of the investigated studies are, to some extent, successful in identifying takeover targets, and the literature on positive abnormal returns for takeover targets is extensive. However, the papers report limited success in achieving abnormal returns when exploiting this information to create an investment strategy with a portfolio of potential takeover targets.

Palepu (1986) reports a non-statistically significant negative abnormal return of -1.63% from his optimal takeover portfolio. However, there are two obvious flaws with this result that needs to be addressed. Firstly, when using the minimum misclassification method for portfolio selection, less than 4% of his portfolio consists of actual targets. Thus, it is unlikely that the takeover return will significantly impact the portfolio return. Secondly, the model is only tested for one year, with one portfolio. An investment strategy must be highly superior to see statistically significant returns from one year of testing.

Powell (2001) targets the issue related to the minimum misclassification on an imbalanced data set when the goal is to earn abnormal returns. He argues that the abnormal returns of the targets will be diluted by the large proportion of non-targets when investing in such a sizeable portfolio. However, testing for abnormal returns from the portfolios, constructed using both minimum misclassification and maximizing the proportion of targets, Powell (2001) finds statistically significant negative returns of respectively -11% and –5% for a 12-month buy-and-hold period. Hence, suggesting that "developing statistical models to predict takeover targets is unlikely to result in a profitable investment strategy" (Powell, 2001, p. 16). However, it is notable that also Powell only tests his model for one year.

Brar et al. (2009) report an average yearly market-adjusted return of 8.5% from their best model when rebalancing the portfolio monthly, which is seemingly an extraordinary result compared to previous studies. However, due to the split where every year, 50% of the targets and non-targets are used for training, while the remaining 50% are used for testing, we believe constructed models to be forward-looking biased. The models will likely capture future changes in trends unknown at the actual time of the portfolio rebalancing. In order to create a true out-of-sample result, one must base the training solely on known historical factors or forecasts. Additionally, they have selected targets from European deals, while the non-targets are collected from the S&P/Citigroup broad market index, consisting of global companies (excluding the U.S.). Hence, their models might be trained to capture both differences between targets and non-targets and differences between European companies and the global market, making it easier to predict the targets in the test set. Lastly, this non-random selection of data creates a higher proportion of targets relative to non-targets (27.2%) compared to any actual market, thus yielding a misrepresentative proportion of predicted targets and abnormal returns.

Khan and Myrholt (2018) use the maximum proportion of targets method for their optimal portfolio due to the disruptive number of companies in the portfolio using minimum misclassification. They train their model on data from 1995 to 2012 and test it on a test set from 2013 to 2016, finding a four-year non-statistically significant CAAR of

0.9% over the test period. Interestingly, they also observe that the non-target portfolio outperforms the target portfolio for the total sample, showing the importance of timing when acquiring targets. In conclusion, they doubt their takeover prediction model's ability to outperform the market.

To our knowledge, no studies using the correct method to solve issues related to look-ahead bias have obtained positive abnormal returns from a target portfolio. All studies use the traditional Capital Asset Pricing Model (CAPM) model when estimating abnormal returns. Utilizing the Fama French factors to find the abnormal return would be interesting, as this would also give a better overview of the risk factors the target portfolio imposes. Timing is of the essence when investing in potential targets. All the studies above use yearly data and rebalancing strategies for their target portfolio, potentially bringing other variables into play. Higher frequency in the data and rebalancing may improve target predictions and takeover returns. Additionally, re-estimating the model yearly by including the years already predicted to predict the following year may allow the model to capture potential changes in the market over the estimation years without compromising real-world investment strategy testing.

3 Hypotheses

To the best of our knowledge, Khan and Myrholt (2018) is the only takeover target prediction study conducted on the Norwegian M&A market, and thus, a natural reference for our thesis. In our study, we examine the marginal contribution of predictive ability from adding two new hypotheses to Khan and Myrholt's (2018) model, in addition to testing whether an algorithmic prediction model performs better than the traditionally used logit model. To test our hypotheses, we use Khan and Myrholt's (2018) logistic regression model as a benchmark.

3.1 Intellectual Property Hypothesis

Firstly, we examine the patent portfolio of the companies in our data. As discussed in the literature review section 2.1, evidence suggests that the importance of intellectual property has increased during the later years (Mousavi, 2011), with patents being one of the most important publicly available indicators of a company's intellectual property portfolio. We examine the companies' overall patent portfolio in terms of quantity, technological change, and technological profile. Our hypothesis is the following:

H0.1: The likelihood of receiving a takeover bid is unrelated to the target company's patent portfolio.

H1.1: The likelihood of receiving a takeover bid increases with improvements in the target company's patent portfolio.

We test this hypothesis in three ways. First, we use logistic regression to test if the size of a company's patent portfolio, measured as either the total number of patents or a dummy indicating whether a company has patents, affects the takeover probability. Second, we use a variable defining the drift or change in the technological profile of a company as a proxy for a restructuring of the patent portfolio. Following Bourne et al. (2019b), this is assumed to have a temporary negative effect on a company's takeover probability. The variable is created by comparing a company's patent applications in the year before the acquisition observation to the four preceding years. We include control variables from Khan & Myrholt's (2018) study in order to test the marginal contribution of the patent variables.

Third, we examine the predictive power of a target company's technological profile based on the IPC-class distribution of their patent portfolio. The IPC-class defines the area of technology to which a given patent pertains. Evidence suggests that companies are more likely to acquire targets with a similar technological profile (Hussinger, 2010). However, there is no renowned research on the technological profile of takeover targets. The technological profile of targets is also likely to change over time. By examining the methods proposed for acquiring companies to search for potential targets (Park et al., 2013; Huang et al., 2013; Ma et al., 2017) we find the IPC-class distribution of a company's patent portfolio to be a good proxy for the technological profile. A company's IPC-class distribution in a given year is defined by the number of patents owned within the different IPC-classes. As each of the 118 IPC-classes requires its own coefficient, it is not appropriate to test the predictive power of technological profile similarly to size and change. Instead, we test by training one model on the benchmark variables, including the technological profile variables, and comparing the model performance to the benchmark model. We can not statistically prove that the takeover probability is related to the technological profile using this method. However, a better-performing model will be a good indication of the importance of the technological profile.

3.2 Stock Exchange Announcement Hypothesis

As our second expansion of Khan and Myrholt's (2018) model, we examine the marginal contribution to predictive power from the number and tone of target firm announcements in the period before an acquisition announcement. As discussed in the literature review section 2.2, evidence suggests that combining textual analysis with traditional financial measures improves merger participant predictions (Routledge et al., 2017). Existing finance literature based on textual analysis focuses mainly on U.S. companies with filings accessible through the EDGAR system. Since there is no similar system in place in Norway, allowing for easy extraction of the Management's Discussion & Analysis (MD&A) section from the annual reports, we apply textual analysis to stock exchange announcements from Oslo Stock Exchange.

H0.2: The likelihood of receiving a takeover bid is unaffected by target-

company stock exchange announcements.

H1.2: The likelihood of receiving a takeover bid increases with target-company stock exchange announcements.

We test this hypothesis in two ways. First, we use logistic regression to test for statistical significance of the number of stock exchange announcements published by each company, both in total and within each announcement category. Second, we use logistic regression to test for statistical significance in the average yearly sentiment of a company's announcements. The sentiment variable, indicating the tone of the announcements, is calculated based on dictionaries constructed specifically for text analysis in a corporate or financial setting. In both cases, Khan and Myrholt's (2018) independent variables are used as control variables.

3.3 Algorithmic Approach Hypothesis

Shmueli et al. (2010) argue that statistical models mainly fall into three categories based on their purpose: predictive models, explanatory models, and descriptive models. As the names suggest, a predictive model is built to predict an outcome given a collection of predictors. In contrast, an explanatory model seeks to explain how an outcome changes based on changes in explanatory variables. A descriptive model has the purpose of giving a compact summary or representation of the data structure and the association between the dependent variable and the independent variables (Shmueli et al., 2010).

In the target prediction literature, researchers have mainly used explanatory models to evaluate what variables cause a response in the dependent variable and make statistical inferences of the included variables. As seen in the literature review, hypotheses have been based on impacts of financial ratios, management inefficiencies, industry disturbances, and macroeconomic factors. Potential predictors have been deemed insignificant due to the non-observed statistical significance from the models they are tested in. However, insignificant parameters in a smaller model do not imply insignificant parameters in a larger model (Heinze and Dunkler, 2017). The explanatory-focused approach may inhibit the ability to capture subtle and more complex relationships in the data due to the requirement of interpretable statistical models (Shmueli et al., 2010).

Our thesis intends to develop a target prediction model without restricting the choice of models due to interpretability. We test how the results from our model perform as the basis of an investment strategy compared to the more traditionally used logit models. Prioritizing the predictions rather than model transparency allows us to utilize a broader range of algorithmic models and methods. Although the algorithmic approach may not explain the underlying causal effects, it can capture complex associations and potentially generate more accurate predictions (Shmueli et al., 2010). Breiman et al. (2001) argue that the commitment to stochastic data models in the statistical community has led to questionable conclusions, irrelevant theory, and restricted the problems worked on by statisticians. Encouraging researchers to depend less on data models and utilize the diverse set of tools provided by algorithmic modeling, he stated that: "Using complex predictors may be unpleasant, but the soundest path is to go for predictive accuracy first, then try to understand why" (Breiman et al., 2001, p.208). In light of this reasoning, our third hypothesis is as follows:

H0.3: Applying a gradient boosting prediction technique do not improve the predictive power of the model in Khan & Myrholt (2018).

H1.3: The predictive power of the Khan & Myrholt (2018) model is improved by applying a gradient boosting prediction technique.

To test this hypothesis, we compare the predictive power of logistic regression models to gradient boosting models using four sets of independent variables. First, we test the benchmark variables from Khan and Myrholt (2018), presented in section 5.1. Second, we add our extensions to the Khan and Myrholt (2018) model with the patent and announcement variables independently and combined. This allows us to isolate the marginal effects of the added variables and the effects of introducing an algorithmic prediction method. Finally, we include all the 188 prognostic variables, shown in Table 4.1, in a gradient boosting model. By introducing a high number of predictors in conjunction with an algorithmic prediction approach, the assumption is that we will obtain better predictive ability and higher abnormal returns from the takeover target portfolio.

The models are all subjected to a recursive model evaluation technique following Danbolt et al. (2016). Furthermore, the predictive power of the models is compared using two model evaluation metrics; the precision-recall curve and the operating characteristics curve. The recursive technique and evaluation metrics are described in sections 5.1 and 5.2, respectively.

Following Powell (2001), the target portfolios are created using the lowest cut-off probability from the decile yielding the maximum proportion of targets in the training data for each model. The overall performance and the difference in the performance of the target portfolios are based on monthly returns over the whole test period using the Fama French three-factor model, where the factor portfolios are based on Oslo Stock Exchange companies (Ødegaard, 2021).

4 Data

In the following chapter, we present the data sources and methods used to create the variables for our prediction models. The subsections are divided by accounting data, stock-market data, ownership structure, macro variables, intellectual property (patents), and stock exchange announcements. The stock-market data will also be used to measure the target portfolios' abnormal returns.

Second, we present the collection and definition of deals included in the data set. As mentioned, to avoid look-ahead bias in our model, all the data is from the calendar year before the takeover announcement. Thus, a firm's takeover probability is based on publicly available data at the end of year t, with the takeover bid dummy during the year t+1 as the dependent variable.[3]

Lastly, we present descriptive statistics for the whole dataset. Additionally, we divide the observations into training data and test data and look at the initial differences between the two data sets. The data include most Norwegian publicly traded companies between 2000 and 2019. Banks, insurance, and real estate investment trusts (GICS 40 & GICS60) are excluded from the sample since many financial ratios do not make sense in these branches.

4.1 Collection of Predictors

A significant part of this research study is collecting, structuring, and preparing data used in the machine learning algorithms. The raw data necessary to create the independent variables in our data set were collected from several sources and required substantial preprocessing before including it in our data set. We provide an extensive explanation of the data sources and how we created and assembled the variables. Table 4.1 presents a complete overview of the predictors included in our final extensive model.

[3] Preferably, the quarter before the takeover announcement should be used. However, the available data for Norwegian companies is not sufficient for this research.

Table 4.1: Overview of All Predictors Included in the Dataset

Profitability and growth	**Returns**
Revenue (NOK)	ROCE (EBIT/Capital employed)
Revenue growth 1-5y	ROCE growth
COGS	Invested Capital (NOK)
COGS growth	Invested Capital growth
EBIT (NOK)	Return on Equity
EBIT growth	Return on Equity growth
EBIT margin	
EBIT margin growth	Return on Assets
	Return on Assets growth
EBITDA (NOK)	
EBITDA growth	**Capital Structure/Liquidity**
EBITDA margin	
EBITDA margin growth	Total Debt (NOK)
	Total Debt growth
Net Income (NOK)	
Net Income growth	Book value of Equity (NOK)
Net Income margin	Book value of Equity growth
Net Income margin growth	
	Total Debt % of EV
Basic EPS (NOK) (Net income/shares issued)	Total Debt % of EV growth
Basic EPS growth	
	debt ratio (multiple) Debt/equity
Dividend yield (div/market cap)	debt ratio growth
Dividend yield growth	
	Total assets (NOK)
Days inventory outstanding	Total assets growth
Days inventory outstanding growth	
Days sales outstanding	Non-Controlling interests (NOK)
Days sales outstanding growth	Non-Controlling interests growth
Days payables outstanding	
Days payables outstanding growth	Interest Coverage ratio (EBIT/Total interest expense)
CCC	Interest Coverage ratio growth
CCC growth	Quick Ratio (current assets - inventory)/current liabilities
	Quick Ratio growth
CAPEX (NOK)	Current Ratio (current assets/current liabilities)
CAPEX growth	Current Ratio growth
CAPEX as % of sales	Cash/Total assets
CAPEX as % of sales growth	Cash/Total assets growth
CAPEX as % of PP&E	
CAPEX as % of PP&E growth	Asset turnover (revenue/total assets)
	Asset turnover growth
Levered FCF (NOK)	
Levered FCF growth	Goodwill + Intangibles (NOK)
Levered FCF margin	Goodwill + Intangibles growth
Levered FCF margin growth	Goodwill + Intangibles % of Total Assets
	Goodwill + Intangibles as % of Total Assets growth

Security Pricing/Volume

1-5 year annualized average return
1-5 year variance of daily returns
Market Beta (OSEAX) - daily 1-3 year
1-3 year abnormal return (CAPM)

Distance from 52 week high/low
Distance from 3y high/low

Relative volume (1 daily relative to 3y daily)
Average volume (daily 1y)

Trading Multiples

Market Cap (NOK)
Market Cap growth

EV (NOK)
EV growth

Cash/Market Cap
Cash/Market Cap growth

EV/Revenue
EV/Revenue growth

EV/EBITDA
EV/EBITDA growth

EV/EBIT
EV/EBIT growth

P/E
P/E growth

Market Cap/Book value of Equity
Market Cap/Book value of Equity growth

Ownership/Share Classes

Single/dual class stock system
Ownership concentration

Industry market trends

Industry classification (dummy)
Industry 1-3 year index price change
Industry 1-3 year volatility
Industry 1-3 year beta

Market/M&A trends

LTM aqcuisitions Oslo stock exchange
LTM aqcuisitions by industry

Overnight lending rate
Norwegian government bonds 3 y
Norwegian government bonds 5 y
Norwegian government bonds 10 y
1-3 year annualized risk free rate

CPI (Consumer price index)
Corporate Tax rate
OSEAX 1-3 year annualized average return
Private Consumption expentditure Norway
Oil price

Intellectual property

Technological Drift
Patent dummy

Company announcements

Publishing frequency total
Publishing frequency by category
Sentiment variables

Company specific feature

Years since establishment
Growth-resource mismatch

4.1.1 Accounting Data

The accounting data is retrieved from the Compustat Global database through Wharton Research Data Services (WRDS). This data is used as a base for companies and years included in the study. All companies with the stock exchange code 228 (Oslo Stock Exchange) between 1995 and 2019 are selected, excluding companies listed on Euronext Growth (previously Merkur Market). We use this data to complete the variables under the categories "Profitability & Growth", "Returns" and "Capital Structure & Liquidity" in Table 4.1. The companies' filings are initially stated in different currencies. Thus, they are converted to NOK using yearly average historic rates assembled from the Central Bank of Norway (2020a) and DNB Markets (2020). The growth variables are based on a one-year change unless stated otherwise.

4.1.2 Stock-market Data & Ownership Structure

The raw stock-market data is extracted from Amadeus 2.0 (the client for NHH Børsprosjektet). From the companies' daily stock prices on the Oslo Stock Exchange between 1995-2019, we calculate the average annualized total return and variance (1-5 years) for each company each year. We then calculate the market beta ($COV_{r,m}/VAR_m$, 1-3 years) relative to the OSEAX index based on excess daily returns. Subsequently, we estimate the annualized abnormal return (1-3 years) based on CAPM using the OSEAX index as the market portfolio. We also compute the return compared to 52 week- and three-year high/low. Lastly, we estimate the average daily turnover volume and the relative volume compared to the three previous years.

After we have created all the stock market data defined under "Security pricing/volume" in Table 4.1, we merge the new variables with the accounting data variables. The Compustat Global database and Amadeus 2.0 are not consistent with company names. Thus, the matching process requires some manual work but is mainly done using regular expressions. We exclude observations for all companies or years in which we cannot retrieve stock-market and accounting data. The Compustat Global database and Amadeus 2.0 are used as sources when creating the "trading multiple" variables.

The ownership structure variables are retrieved from the database made available by Samfunns- og næringslivsforskning AS (SNF). The retrieved variables: max ownership share (max_eiera) and the Herfindahl Index (Ownership Concentration), are joined directly to the data set. The matching of the SNF data to our data set was done using a list of company names and organization numbers from Proff Forvalt and manually looking up organization numbers through Brønnøysundregistrene.

4.1.3 Industry Data & Macro Variables

The raw stock-market data segregated by industry is extracted from Amadeus 2.0. We calculate returns and variance from the GICS industry index prices, in addition to industry betas relative to the OSEAX index for the three years before the takeover observation.

Furthermore, we have collected several macro variables from different sources. Overnight lending rate and government bond yields come from the Central Bank of Norway (2020b). Consumer price index (2020a), private consumption expenditure (2020b), and corporate tax rates (2020c) come from Statistics Norway (SSB). The yearly average annual brent crude prices are gathered from Statista (2020a). Furthermore, the M&A trends both on Oslo Stock Exchange in total and by industry are collected from Reuters' SDC Platinum platform. This will be further detailed in chapters 4.2 and 4.3.

4.1.4 Patent data

4.1.4.1 Data Collection and Cleaning

The unprocessed patent data is retrieved from the Orbis platform owned by Bureau van Dijk. Each company's patent portfolio history is given with title, status, publication number (including date published), and IPC class. There is no way of knowing if the company has sold its patents, as we only observe the patent owner at the publication date. Consequently, we assume that a patent is held for its lifetime. Following the Norwegian Patents Act (2020) section 40, a patent may be obtained for up to 20 years from the date of filing the patent application. Thus, we assume that a patent is held valid for 20 years. However, the actual value of a patent is often diminished earlier as new technologies are introduced.

We do not differ by granted or pending patent applications as we cannot do this for historical patent data where the patents are marked as expired. Additionally, as stated by Park et al. (2013, p. 888); "Every patent, whether or not it is granted, and whether or not it has commercial value, is a result of R&D activity and thus includes technological insight that can offer inspirations or hints to subsequent developments in technology.". This is important for the technological change variable. Moreover, we count patents from different countries on a single invention as several patents, as all patents are assumed to create value for the firm.

We extract all patents published in the period from 1979-2018 from companies on the Oslo Stock Exchange, totaling 66300 applications. Each year, we create a companies' patent portfolio by including all their patent applications for the last 20 years up to the year before the acquisition observation. Figure 4.1 shows the development of the total number of patents owned by companies listed on the Oslo Stock Exchange, including the average number of patents owned by companies with patents and the number of public companies owning patents a given year. We observe that the number of patents has increased 2.5 times in the last 20 years. The average number of patents per company with patents was reduced from 300 to 230 between 2003 and 2007. However, this can be explained by the number of companies with patents increasing from 63 to 102 in the same period.

Figure 4.1: Overview of total number of patents, average number of patents per company with patents and number of companies with patents

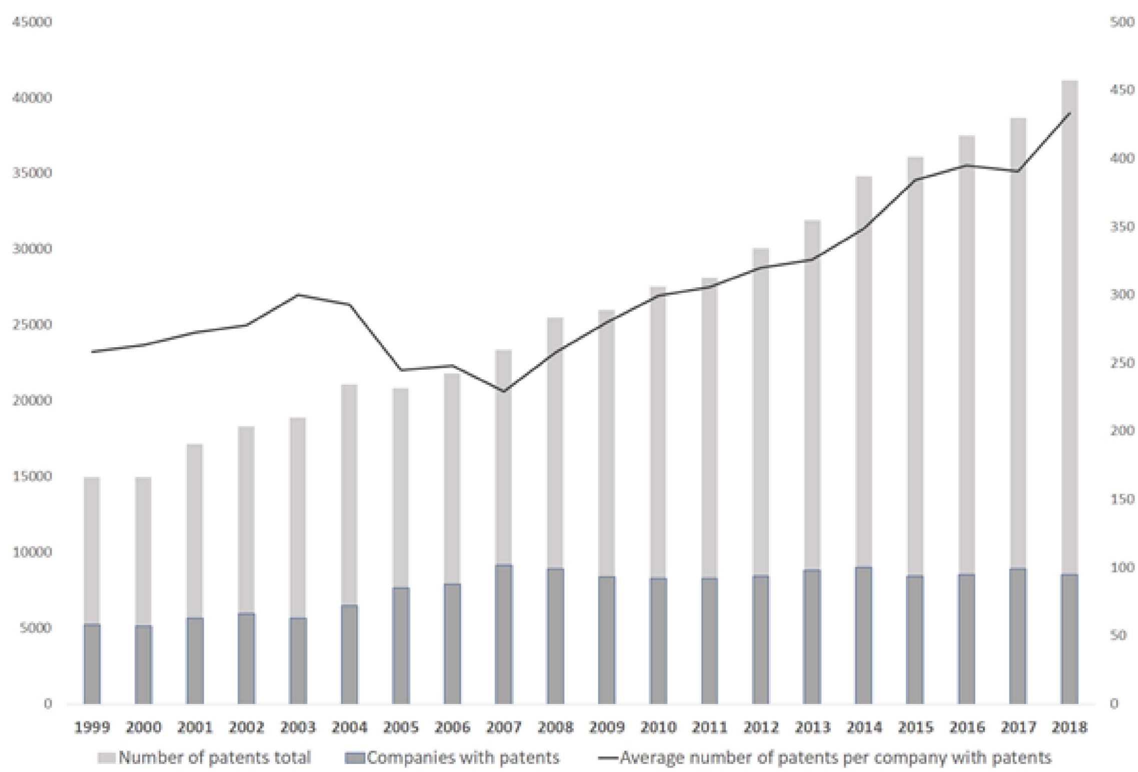

In Figure 4.2 we show the relative proportion of companies with a patent portfolio compared to companies without a patent portfolio on Oslo Stock Exchange. This has increased steadily from approximately 40% in 1999 to over 60% in 2018.

Figure 4.2: Proportion of companies with and without patents on Oslo Stock Exchange

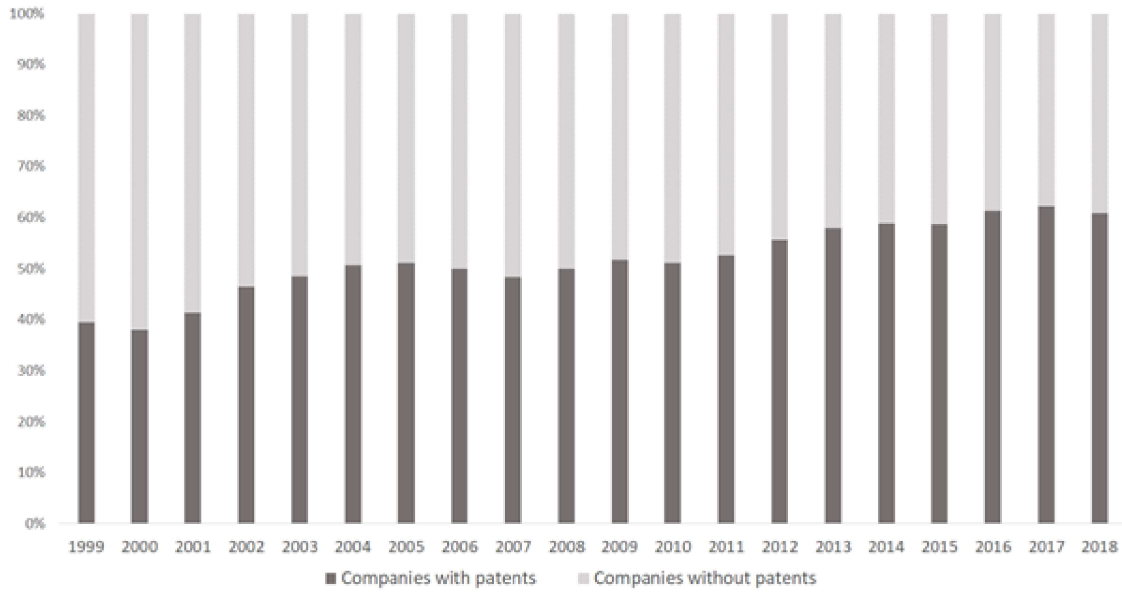

4.1.4.2 Variable Construction

We extract the total patent portfolio of a company constricted by the 20-year validity of a patent for each year. Thereafter, we create a vector with all the individual IPC-classes and count the number of patents within each IPC-class. This operation is repeated for all the observations (company, year). Lastly, we combine the vectors of all observations to create a large matrix constituting all IPC classes observed on the column vector and each company each year on the row vector. In total, we observe 118 different IPC classes. An example displaying the data before and after preprocessing is shown in Appendix A1.3. We use the IPC class matrix and sum up row-wise to find the total size of a company's patent portfolio in any year. Additionally, we create a dummy variable defining whether a company has a patent portfolio.

The technological change is based on the change in the proportion of patents within a given IPC class, accounting for newly added IPC-classes and changes in the proportion of patents within existing IPC-classes. The metric created by Bourne et al. (2019b) quantifies the year-on-year change in a company's technological profile and can be obtained by using the following formula:

$$Techdrift_{it} = \frac{1}{2} * \sum_{c=0}^{n} |PATENTPROP_{ic,t} - PATENTPROP_{ic,t-12}| \tag{4.1}$$

A company's technological drift each year is the sum of the absolute values of the number of patents within a given IPC-class published in the last five years, minus the number of patents within the same IPC-class published in the last five years starting from 1 year earlier. Thus, the higher the value obtained, the higher the technological change.

4.1.5 Stock Exchange Announcements

4.1.5.1 Data Collection and Cleaning

To retrieve the stock exchange announcement data, we contacted Oslo Stock Exchange and received all the announcements published from the beginning of 2000 to the end of 2019. The initial data set contained 348 783 announcements. As textual analysis is computationally intensive, unnecessary observations were removed, keeping only the announcements published by companies represented in our data set. Matching the announcements to our data set was primarily done using regular expressions.

Utilizing Google's pre-trained machine learning translation models, we detected the language of the announcements. We removed duplicated announcements published in both English and Norwegian, opting to keep the English announcements to improve the robustness of the sentiment analysis. The resulting data set contains 165 425 announcements. All the announcement categories are represented, with "Innsideinformasjon" being the most common category. The announcement categories and the number of observations within each category are shown in Appendix A2.1.

4.1.5.2 Variable Construction

The quantity variables for the stock exchange announcements are created by counting each company's observations each year. We include both a company's total number of announcements in a year and the number of announcements within each stock exchange announcement category.

The sentiment variables are constructed by conducting a sentiment analysis, a technique used to determine the emotional tone in a text. This is done by utilizing a pre-defined dictionary, as described in the literature review. Loughran and McDonald (2011) show that the choice of dictionary can significantly impact the results of sentiment analyses. In our thesis, the analysis is applied to corporate announcements published through Oslo Stock Exchange. Thus, we opted to use Loughran and McDonald's Master Dictionary (henceforth referred to as the LM Dictionary) and Henry's Dictionary (henceforth referred to as the HE Dictionary). The LM Dictionary is the most frequently used in the finance literature, and thus, our preferred dictionary. However, as both dictionaries are developed explicitly for financial and corporate texts and readily available through the "SentimentAnalysis" package in R, we include both.

The sentiment analysis was conducted using the "analyzeSentiment" function from the "SentimentAnalysis" package in R. A detailed explanation of the text pre-processing for the analysis can be found in Appendix A2.1. The resulting data frame contains the sentiment score, positivity-, negativity-, and uncertainty ratio based on the LM Dictionary, and the sentiment score, positivity-, and negativity ratio based on the HE Dictionary for every announcement. The sentiment variables are constructed by calculating the mean of each of these variables for each company each year.

4.2 M&A Activity

The acquisition data is from the Refinitiv SDC Platinum database from Reuters. We include unsuccessful and withdrawn takeover bids, as previous studies show that the stock price in the event of a takeover bid announcement also experiences abnormal returns similarly to successful acquisitions (Jensen and Ruback, 1983). We define takeover bids as bids in which the controller if the acquisition is completed, will gain control over more than 40% of the target company prior to 01.01.2008 and 33% following 01.01.2008. Our definition is based on The Norwegian Securities Trading Act Chapter 6, section 6-1 (1), which states that:

"Any person who through acquisition becomes the owner of shares representing more than 1/3 of the voting rights of a Norwegian company the shares of which are quoted on a Norwegian regulated market is obliged to make a bid for the purchase of the remaining shares in the company" (The Financial Supervisory Authority of Norway, 2020).

The law entered into force 01.01.2008. Before the beginning of 2008, the mandatory bid obligation threshold was 40% of the voting rights (Ministry of Finance, 1997). As ownership above these boundaries entails a mandatory bid for all the remaining shares, the company is deemed a potential in-play target. In Figure 4.3 we present an overview over yearly takeover bids.

Figure 4.3: Number of Takeover Bids in Our Sample Period

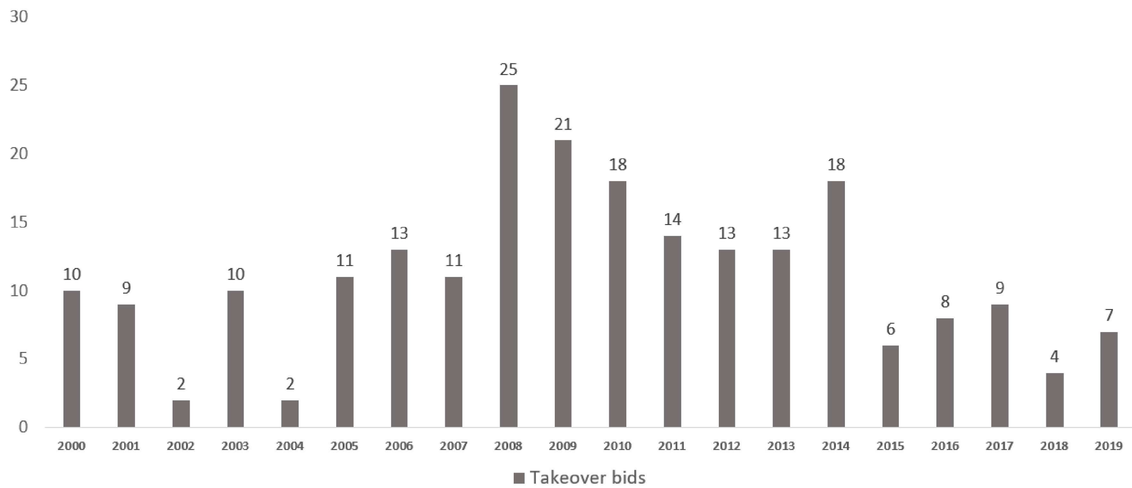

4.2.1 Imputing Missing Values with MICE

Unfortunately, our data sample initially contains missing values. This is mainly the case for the data retrieved from the Compustat Global database for the earlier years in our sample period. The variables relevant to Khan & Myrholt's (2018) model are central to our analysis. Thus, time was dedicated to obtaining the missing values by manually searching through annual reports and Proff Forvalt. We were able to complete all the Khan and Myrholt (2018) variables except the ownership concentration. As we do not have the required information to compute the Herfindahl-Hirschman Index ourselves and have not found an alternative data source to retrieve the missing values, we impute them using the machine learning algorithm MICE. In total, 716 out of 3339 values of the Herfindahl-Hirschman Index are imputed. 132 out of 393 companies have one or more imputed values of the variable. However, in the univariate regression shown in Appendix A3.1, the takeover probability is shown to be unrelated to ownership concentration. Thus, the imputations should not impact the results of our analysis. Neither the patent data nor the announcement data contain imputed values.

In our final model, we include 188 potential predictors in an algorithmic prediction model. As we do not spend time on feature engineering and examining the rationale for each of these variables, this model is not central to our analysis. For the 144 variables (beyond the variables from Khan and Myrholt (2018), patent- and announcement variables), 23 735 out of 507 528 data points are imputed (4.67%). Our decision to impute the missing values, rather than deleting observations containing missing values, is based on the findings of Lall (2016). The study reanalyzes the results of several quantitative studies using multiple imputation rather than listwise deletion. Lall (2016) argues that multiple imputation generates more efficient inferences and is less biased than listwise deletion, concluding that "inferences produced by listwise deletion are likely to always be (to some extent) biased as well as inefficient" (Lall, 2016, p. 31).

4.3 Descriptive Statistics

The final data set, which consists of companies traded on the Oslo Stock Exchange in the years 2000-2019, includes 387 unique companies, with a total of 3339 observations. In total, there are 224 acquisitions or proposed acquisitions over the period. Figure 4.4 shows the yearly number of takeover bids and the following takeover rate.

Figure 4.4: Number of Takeover Bids and Non-Targets with Yearly Takeover Rate

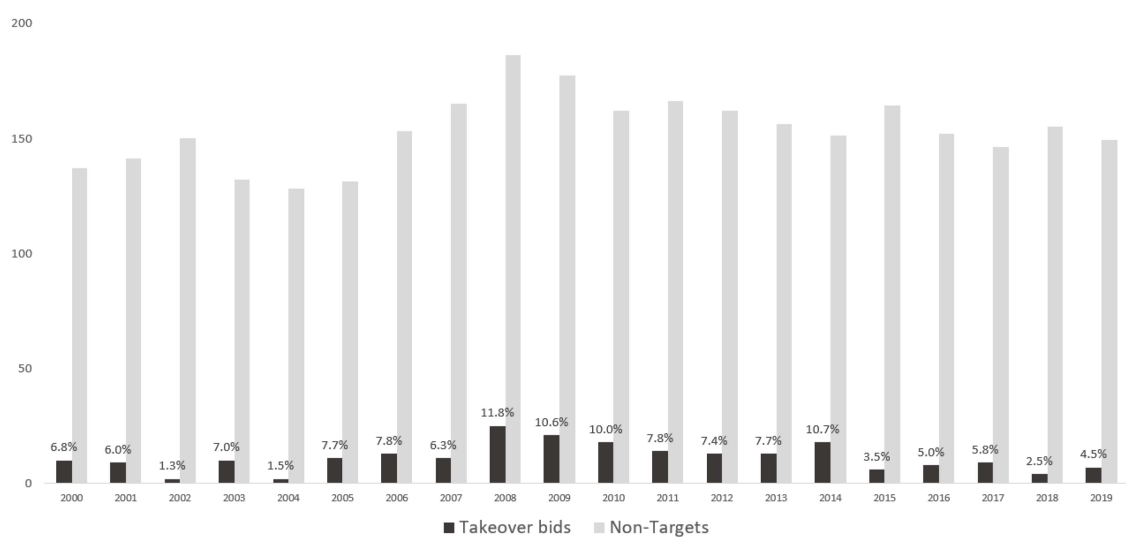

Table 4.2 gives an overview of targets and non-targets divided by sector. We observe

that the consumer discretionary, consumer staples, and information technology sectors generally have a higher average takeover rate. At the same time, it is lower for utilities, health care, and materials.

Table 4.2: Entire Sample - Sector Distribution of Takeover Bids

G-Sector	Total Observations	Non Targets	Takeover Targets	Takeover Rate
Energy	970	905	65	6.70%
Materials	191	184	7	3.66%
Industrials	811	762	49	6.04%
Consumer Discretionary	228	206	22	9.65%
Consumer Staples	239	216	23	9.62%
Health Care	209	201	8	3.83%
Information Technology	491	448	43	8.76%
Communication Services	114	108	6	5.26%
Utilities	34	33	1	2.94%

4.3.1 Training Set

In the initial estimation sample (acquisitions between 2000 and 2012, with observation data between 1999 and 2011), we have 159 observations classified as targets. This corresponds to an average of 7.4% of the companies receiving a takeover offer in any one year. The takeover rate is somewhat higher than the 6.8% reported by Khan and Myrholt (2018). The positive difference is related to the inclusion of unsuccessful takeover bids. In contrast, the exclusion of the fifth takeover wave (1995-1999) with high activity and demanding 40% instead of 33% ownership until 2007 takes the takeover rate in the opposite direction.

Table 4.3: Initial Training Sample - Sector Distribution of Takeover Bids

G-Sector	Total Observations	Non Targets	Takeover Targets	Takeover Rate
Energy	591	549	42	7.11%
Materials	119	112	7	5.88%
Industrials	547	516	31	5.67%
Consumer Discretionary	178	159	19	10.67%
Consumer Staples	152	137	15	9.87%
Health Care	122	117	5	4.10%
Information Technology	359	324	35	9.75%
Communication Services	68	63	5	7.35%
Utilities	13	13	0	0.00%

In Table 4.3 we present the takeover bid rate by industry for the training set observations. As the training set constitutes a major part of the total data set, the takeover rates are unsurprisingly similar to that of the entire data set.

4.3.2 Test set

The total test sample (2013-2019) consists of 1138 observations, with 65 observations (5.7%) classified as takeover targets. We employ a recursive model evaluation technique in the takeover prediction model, meaning multiple models will be trained to predict multiple test sets. This will be discussed in further detail in section 5.1.

Table 4.4: Total Test Sample - Sector Distribution of Takeover Bids

G-Sector	Total Observations	Non Targets	Takeover Targets	Takeover Rate
Energy	379	356	23	6.07%
Materials	72	72	0	0.00%
Industrials	264	246	18	6.82%
Consumer Discretionary	50	47	3	6.00%
Consumer Staples	87	79	8	9.20%
Health Care	87	84	3	3.45%
Information Technology	132	124	8	6.06%
Communication Services	46	45	1	2.17%
Utilities	21	20	1	4.76%

Interestingly, in Table 4.4 we observe that the takeover rates are very different in the test set compared to the training set. This changing trend in takeover rate by industry suggests that industry independently is an unfavorable variable to use when predicting takeover targets.

5 Methodology

In this section, we present the methodologies used to test the hypotheses defined in section 3. The following section consists of four stages. First, we describe the methodology for determining the marginal contribution from including intellectual property and announcement data in the takeover prediction model. Second, we introduce an algorithmic prediction method for the takeover target prediction. Third, we present the performance metrics used to compare the performance of the different models. Finally, we present the method used to construct and evaluate the portfolios based on the target prediction models.

5.1 Logit Models

The aim of our two first hypotheses, H1 and H2, is to test the marginal contribution from adding patent and target announcement data to Khan & Myrholt's (2018) model. To test the hypotheses and the marginal effects of the additional variables, Khan & Myrholt's (2018) predictors are included as control variables when comparing targets and non-targets. These predictors will also lay the foundation for the benchmark model in our research. Khan and Myrholt's (2018) predictors are: Growth (2-year sales), the natural logarithm of sales, cash to total capital, current ratio, ownership concentration (based on the Herfindahl-Hirschman index), industry disturbance (quantity of acquisitions in the same sector in the previous calendar year), oil price, and 10-year Norwegian government bond yield. To diminish the effects of potential differences in data quality, we replicate Khan & Myrholt's (2018) model, our benchmark model, using our data. In their thesis, their model is given by:

$$logit(p(x)) = log\frac{p(x)}{1-p(x)} = \alpha + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + ... + \beta_k x_k \tag{5.1}$$

In the logistic regression (Equation 5.1), p(x) denotes the conditional probability $\Pr\langle Y = 1|X = x\rangle$, where Y is a dummy variable indicating targets and x is the explanatory variables. Following Heldal and Fosen (2006), Khan and Myrholt (2018) use the logistic transformation, $log\frac{p}{1-p}$, to ensure that the values p(x) are between 0 and 1. As shown by

Khan and Myrholt (2018), solving Equation 5.1 for p(x), the takeover probability is given by:

$$p(x) = \frac{e^{\alpha+\beta_1 x_1+\beta_2 x_2+\beta_3 x_3+\ldots+\beta_k x_k}}{1+e^{\alpha+\beta_1 x_1+\beta_2 x_2+\beta_3 x_3+\ldots+\beta_k x_k}} \tag{5.2}$$

Both the samples in our and Khan and Myrholt's (2018) theses span over multiple periods, and the independent variables, x, vary over time. To find the functional relationship between the independent variables and the takeover likelihood in a given period, they apply the following equation:

$$p_{i,t} = \frac{1}{1+e^{-\beta x(i,t)}} \tag{5.3}$$

In Equation 5.3, p(i, t) denotes the probability of firm i receiving a takeover bid in period t. The independent variables are represented by vector x(i, t), and β is a vector of the estimated parameters.

As mentioned in the hypothesis section, all models are subjected to a recursive evaluation technique following Danbolt et al. (2016). The models are firstly constructed from a training set of observations in 1999-2011 (acquisitions from 2000 to 2012) and tested on a test set containing observations in 2012 (acquisitions in 2013). Thereafter, the model is re-estimated for 1999-2012 and tested on observations in 2013. This process is repeated for all the years in our data sample, with the final test set containing observations for 2018 (acquisitions in 2019). As mentioned, a time lag is imposed in the dependent variable so that all predictors obtained in year t are used to predict the outcome of the dependent variable in year t+1.

5.2 Algorithmic Approach

Our third hypothesis aims to test whether introducing an algorithmic prediction model produces higher accuracy than the traditionally used logistic regression models. To test this hypothesis, we employ a logistic regression model and an Extreme Gradient Boosting (XGBoost) model to predict our test set. Both models are estimated using Khan and Myrholt's (2018) independent variables with, and without, the addition of intellectual property and announcement data to isolate the effects of introducing the XGBoost model. Thereafter, we compare the performance metrics described in the following section to determine the differences in performance between the models. Finally, we include an XGBoost model using all 188 potential predictors in our data set. We examine whether a high number of predictors introduced in conjunction with the XGBoost algorithm result in better predictive power than models based on Khan & Myrholt's (2018) independent variables with the addition of intellectual property and announcement data.

5.2.1 Choice of Algorithmic Model

Researching potential algorithmic models to be compared to the logistic regression model, we decided on two main alternatives: XGBoost and Artificial Neural Networks (ANN). Both algorithms are widely used and well-known for classification problems with a proven ability to generate accurate predictions. However, due to the relatively limited number of observations on our data set, we were not convinced that we would fully utilize the ANN algorithm's potential. Our initial testing of both algorithms confirmed our suspicion, as XGBoost consistently outperforming the ANN. Thus, XGBoost is our algorithmic model of choice.

Extreme Gradient Boosting, or XGBoost, is an ensemble method further developing the gradient boosting framework. Gradient boosting of decision trees was first introduced by Friedman (2001) and is often referred to as Gradient Boosting Machines. The technique involves iteratively fitting separate decision trees to modified versions of the original data. The decision trees are fitted sequentially, and every new tree is fitted using information from the previously fitted trees. Thus, gradient boosting algorithms can be thought of as slowly learning (James et al., 2013). Conceptually, XGBoost is similar to GBM. For a detailed explanation of the XGBoost algorithm and how it differs from Friedman's (2001) GBM, we refer to "Tree Boosting with XGBoost" by Nielsen (2016). A notable strength of the boosted decision trees, highly relevant to our thesis, is their ability to capture high-order interactions between features in the data and to automatically perform feature selection by reducing the weight of less important features (Nielsen, 2016).

5.2.2 Hyperparameter Tuning

Hyperparameters are specifiable parameters used to control the learning process of a machine learning model. To tune the hyperparameters in the fitting of our models, we used an approach known as random search. This method involves splitting the training set further into a training set and a validation set. We opted to split the training set 80-20, with 80% of the original training data being kept as training data, while the remaining 20% are used as validation data. The idea behind a validation set is to keep the test set hidden from the model during the learning process while still being able to test the model on out-of-sample data to avoid overfitting.

After creating the new training and validation sets, we randomly generated 20 000 values for each hyperparameter. The performance of the hyperparameters was tested by training a model with each set of randomly generated hyperparameters and using the models to predict the validation set. The combination of hyperparameters resulting in the best prediction of the validation set was used when fitting the model to the original training data to predict the test set. This process repeated for each prediction year, following the recursive evaluation technique.

5.3 Performance Metrics

In addition to examining the number of actual targets predicted to each quantile, we will utilize two model performance metrics often used to evaluate the overall performance of a machine learning model. Firstly, we look at the area under the receiver operating characteristics (ROC) curve. In our predictions, each observation obtains a probability of acquisition. The ROC curve is created by a continuous probability threshold calculating the true-positive rate and the false-positive rate. The true-positive rate is defined as the number of acquisitions correctly classified as a target of all the targets in the test set at any given probability threshold. The false-positive rate is the number of non-acquisitions wrongly classified as targets of all the non-targets in the test set at any given probability threshold. The values are plotted with the true-positive rate on the y-axis and the respective false-positive rate on the x-axis. Thereafter, the area under the ROC curve (AUC-ROC) is calculated. The AUC-ROC takes on a value between 0 and 1, where the value indicates how well the model can distinguish between targets and non-targets. For example, an AUC value of 0.7 suggests a 70% probability that the model will distinguish between a target and a non-target. Thus, a value below 0.5 indicates that the model performs worse than a random classification.

Additionally, we will calculate the area under the precision-recall curve. The method to obtain this curve is similar to that of the ROC curve. The precision-recall curve places higher importance on the correct prediction of the minority class in an unbalanced data set, in our case, the correct classification of targets. Thus, the area under the precision-recall curve is, in our case, a more desirable metric, especially considering that we are interested in obtaining the highest possible proportion of targets in our portfolio. The precision value is calculated by the rate of actual targets in the predicted portfolio divided by the total number of companies in the portfolio for a given probability threshold. The recall value is calculated by the rate of actual targets in the predicted portfolio divided by the total number of actual targets in the whole test set. The values are plotted with the precision ratio on the y-axis and the respective recall ratio on the x-axis. The area under the developed precision-recall curve is then calculated. Again, the higher the AUC, the better the model. However, the interpretation of the value will be different based on the hypothesis in question.

5.4 Target portfolio returns

Following Powell (2001), we use the maximum target strategy as described in chapter 2.4.3 when finding the cut-off probabilities to create the target portfolios. When evaluating the portfolio performance for each model independently, we use a long-only approach with the derived target portfolio.

Firstly, we calculate the holding period return over the years from 2013-2019 for both the target portfolios in addition to the OSEAX index. The returns are updated monthly. However, the portfolio is balanced yearly. All companies begin with the same portfolio weight, while well-performing companies will obtain a larger weight over the year. The portfolio is rebalanced on the first of January each year with the newly predicted targets. The cost of rebalancing is not included in the calculations of returns.

From the literature review, we saw that most of the abnormal returns from takeovers are generated in the runup weeks before and including the announcement day. Moreover, Jensen and Ruback (1983) show that in the subsequent month of the announcement, target firms yield a negative abnormal return on average. With this in mind, we implement a strategy where if a takeover is announced, the company's stock will be sold, and the cash will not be reinvested before the next period. The selling price is set to the closing price on the announcement day or the first trading day after the announcement. This depends on whether the announcement is published before or during trading hours or after the market has closed.

We further evaluate the risk-adjusted returns of the portfolios using the Fama French 3-Factor model. Khan and Myrholt (2018), and their reference studies, in general, use the CAPM when evaluating the abnormal return of the target portfolios. Using the Fama French 3-factor model, we seek to capture more than just the systematic market-risk defined using CAPM. We also try to observe whether the portfolios are exposed to other types of systematic risk, which may be the cause as we are searching for companies with specific features. The two additional factors added in addition to the market risk factor are known as "Small-Minus-Big" (SMB) and "High-Minus-Low" (HML). SMB refers to the historical returns of a portfolio of small-capitalization companies in excess of a portfolio of large-capitalization companies. The same holds for HML with a portfolio of companies with high book-to-market value minus low book-to-market value.

The rationale behind the inclusion of the two factors is that smaller companies tend to outperform larger ones in the long term and that high book-to-market companies tend to outperform low book-to-market companies (Fama and French, 1998). Thus, if the returns of a portfolio are correlated with one of these factors, positive and negative excess returns can be explained by the changed exposure to these systematic risk factors. The risk factors will help us get a better picture of the stock picker performance of the portfolios. This is especially important as the test period we examine consists of an uninterrupted positive market trend, where risk is often rewarded.

From the descriptive statistics of the Fama factors in Table 5.1, we observe that the market has performed well in the period with an average monthly excess return of 0.8%. As expected, small-cap stocks have outperformed large-cap stocks. Interestingly, we see that the high book-to-market portfolio has underperformed compared to the low book-to-market portfolio, in conflict with the rationale of Fama and French (1998). However, the results are not surprising, knowing that growth stocks have been more popular than value stocks since the financial crisis. Additionally, the period of observation is small. The HML is also by far the most volatile portfolio indicated both by the standard deviation and the maximum and minimum values.

Table 5.1: Fama French Factor Portfolio Returns on Norwegian Data from 2013 to 2019

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Allshare	84	0.009	0.031	−0.071	−0.005	0.030	0.074
market_excess	84	0.008	0.031	−0.072	−0.005	0.029	0.073
SMB	84	0.007	0.029	−0.060	−0.013	0.026	0.072
HML	84	−0.007	0.048	−0.196	−0.030	0.024	0.139
RF	84	0.001	0.0003	0.001	0.001	0.001	0.002

The explicit calculation to construct the portfolios is found in "Value versus Growth: The International Evidence" (Fama and French, 1998). The factor portfolios used in our evaluation are factor portfolios as calculated by Fama and French (1998) with Norwegian data. The portfolio data is retrieved from the website of Ødegaard (2021). We subtract the monthly risk-free rates from the target portfolios. Thereafter we regress the excess target returns on the three-factor portfolio returns using the following formula (5.4):

$$R_{it} - R_{ft} = \alpha_{it} + \beta_1(R_{Mt} - R_{ft}) + \beta_2 SMB_t + \beta_3 HML_t + \epsilon_{it} \tag{5.4}$$

It is assumed that these factors capture the systematic risk exposure of the target portfolios. Thus, the remaining intercept (alpha) indicates the excess return adjusted for the non-diversifiable risk, only prone to idiosyncratic (firm specific) risk. [4]

Lastly, we look at the performance of the portfolios against each other using a long-short strategy by calculating the difference between the monthly returns of two portfolios before regressing the long-short portfolio returns on the 3-factor model. This ables us to evaluate the difference in risk. Additionally, we can determine whether there is a significant difference in the returns of the portfolios.

[4] In later years Fama and French (2015) have proposed two new systematic risk factors on profitability and investment strategy of firms, known as the 5-factor model. However, today there are not portfolios created for these factors using Norwegian data. Thus, we exclude these factors from the study.

6 Empirical Results

In the following section, we present our empirical evidence. In section 6.1, we present the empirical evidence to test the first hypothesis on intellectual property. In section 6.2, we perform the same analysis on the stock exchange announcement variables. Furthermore, in section 6.3, we create the models described in the hypothesis chapter and evaluate their predictive power. Lastly, in section 6.4, we present the target portfolio returns of the best-performing models and examine their exposure to systematic risk factors.

6.1 Patent data

The following section aims to test whether patent variables can marginally contribute to the predictive ability of the benchmark model replicated using Khan Myrholt's (2018) variables in their optimal prediction model. Firstly we examine the size and technological change variables. In Table 6.1 we present the coefficient estimates. The number of patents coefficient (model 1) is negative and significant at a 5% level when regressed independently on the dependent variable. However, when including the patent dummy variable (model 2), we observe that the significance of the model comes from whether a company has patents, not the number of patents in the portfolio. The patent dummy variable coefficient is negative and significant at a 1% level. This suggests that companies without a patent portfolio are less likely to receive a takeover bid. The results are surprising, considering the increased importance of intellectual property as discussed in the literature review. This could be the result of a tangible asset-heavy Oslo Stock Exchange compared to the S&P500.

We observe the technological change coefficient (model 3) moving towards statistically significant at a 10 percent level with a negative sign, suggesting that the theory that companies with high internal technological drift are less likely to be acquired holds (Bourne et al., 2019a). However, all companies without patents will yield a technological change of zero, which will affect the results. When adjusting for the patent dummy (model 4), we observe no statistical significance. Without knowing more about the study, it may be the case that Bourne et al. (2019a) experienced the same results while reaching a different conclusion as they did not have a patent versus no-patent variable in their model.

The predictors from the benchmark model are added to the patent dummy in model 5 to test for marginal contribution. The patent dummy is still significant at a 5% level after the control variables are included (p-value of 0.011). Thus, we can reject the null hypothesis stating that the likelihood of receiving a takeover bid is unrelated to the target company's patent portfolio. The results are not in line with our alternative hypothesis, stating that improving a company's patent portfolio increases the likelihood of receiving a takeover bid. Nevertheless, the likelihood is definitely related to the patent portfolio. All else equal, the univariate takeover probability of companies with patents is 5.2% compared to 8.4% for companies without patents.

Table 6.1: Regression of Patent Variables

	Dependent variable:							
	dependent_variable							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Patent Total	-0.001^{**}	-0.0004						
	p = 0.038	p = 0.134						
Patent Dummy'		-0.426^{***}		-0.531^{***}	-0.378^{**}	-0.414^{***}	-0.428^{***}	-0.415^{***}
		p = 0.004		p = 0.001	p = 0.011	p = 0.006	p = 0.005	p = 0.006
Technologic Drift			-0.775	0.181				
			p = 0.193	p = 0.774				
Growth					0.028^{*}	0.028^{*}	0.029^{*}	0.028^{*}
					p = 0.081	p = 0.081	p = 0.070	p = 0.094
Return on Equity					0.018	0.017	0.017	0.016
					p = 0.308	p = 0.349	p = 0.354	p = 0.451
Revenue					-0.00002^{*}	-0.00002^{*}	-0.00002^{**}	-0.00002^{*}
					p = 0.054	p = 0.056	p = 0.048	p = 0.054
Cash/Total Assets					-1.748^{***}	-1.633^{***}	-1.567^{***}	-1.666^{***}
					p = 0.0001	p = 0.0003	p = 0.001	p = 0.0003
Current Ratio					0.003^{**}	0.003^{**}	0.003^{**}	0.003^{**}
					p = 0.046	p = 0.028	p = 0.031	p = 0.026
Ownership Concentration					0.004	0.142	0.119	0.132
					p = 0.988	p = 0.561	p = 0.627	p = 0.590
Industry Disturbance					0.015	0.025	0.023	0.023
					p = 0.593	p = 0.386	p = 0.466	p = 0.432
Brent Crude Oil Price					0.009^{***}	0.009^{***}	0.009^{***}	0.008^{***}
					p = 0.002	p = 0.003	p = 0.004	p = 0.004
10 Year Government Bond Yield					0.112^{*}	0.003	0.115^{*}	0.123^{**}
					p = 0.058	p = 0.983	p = 0.055	p = 0.037
(Godwill+Intangible)/Total Assets						0.808^{***}	0.787^{***}	0.789^{***}
						p = 0.0005	p = 0.001	p = 0.001
Year						-0.039		
						p = 0.343		
G-Sector 10							-0.136	
							p = 0.426	
G-Sector 35							-0.463	
							p = 0.238	
G-Sector 55							-0.955	
							p = 0.353	
Growth-Resource Mismatch								0.157
								p = 0.294
Price/Earnings								-0.000
								p = 0.990
Market-to-Book								-0.003
								p = 0.876
Debt Ratio								0.001
								p = 0.955
Constant	-2.562^{***}	-2.388^{***}	-2.582^{***}	-2.388^{***}	-3.249^{***}	74.833	-3.377^{***}	-3.514^{***}
	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.364	p = 0.000	p = 0.000
Observations	3,339	3,339	3,339	3,339	3,339	3,339	3,339	3,339
Log Likelihood	-822.370	-818.094	-825.832	-819.999	-798.704	-792.690	-791.559	-792.573
Akaike Inf. Crit.	1,648.740	1,642.188	1,655.665	1,645.998	1,619.409	1,611.380	1,613.118	1,617.146
Note:								*p<0.1; **p<0.05; ***p<0.01

The patent dummy is examined further to find other control variables and look for potential reasoning for the observed results. We test the correlation of the patent dummy variable against all the 188 predictors in our data set. The highest correlation is naturally found in the goodwill plus intangible assets variable. Interestingly, companies with a high proportion of goodwill plus intangible assets relative to total assets are significantly more likely to be acquired (model 6). This is more in line with the development of intellectual property presented by Mousavi (2011), suggesting that companies with a high proportion of goodwill and intangible assets without patent portfolios are desirable targets. Additionally, the goodwill plus intangible assets proportion of total assets may lead the patent applications, as companies first invest in R&D and record intangible assets to receive a patent subsequently. This can suggest that potential acquirers are made aware of, and manages to acquire, companies with valuable patents or patents with promising technologies before the publication of the patent.

The lower takeover probability of companies with patents could be explained by the higher number of acquisitions in the earlier years and the larger proportion of companies with patents in the later years of our sample. However, when including a pure time variable (model 6), we still find the patent dummy variable highly significant.

The patent variable is positively correlated with size variables, indicating that larger companies are more likely to have patents. Consequently, we also find a slightly positive correlation with average daily volume and a negative correlation with 3–5-year variance because the stock price of large companies often, on average, is less volatile over time but more traded. Nevertheless, the size variable should be accounted for with the revenue factor from our benchmark model. We also find a positive correlation with macro variables such as CPI and household consumption. This is likely because these factors and the number of companies with patents have increased over time. Government bonds are negatively correlated with the patent dummy, with the opposite explanation as the yield has generally decreased over the last 20 years. Again, these effects are already accounted for in models 5 and 6.

We find a negative correlation with the Energy- (g-sector 10) and the Utilities sector (g-sector 55), and as expected, a positive correlation to the Health Care sector (g-sector 35). From the descriptive statistics section, we know that these sectors deviate somewhat from the average takeover probability. However, when including dummy variables for the mentioned sectors (model 7), we still find the patent variable highly significant. Lastly, we include the variables that were excluded from Khan and Myrholt's (2018) optimal model selection (model 8), but these variables show no significance. Due to the imputed data on ownership structure, we replicate all the models without this variable (Appendix A4.1). However, the results do not change notably. This is also done for the models in the following section on stock exchange announcements with the same results.

In conclusion, we do not believe that the lack of patents has a causal effect on a company's increased takeover probability. However, it does seem to capture a strong relationship that is yet to be found using the explanatory variables in our data set.

Secondly, we examine the patent profile of companies by testing for marginal increased predictive power when adding the IPC class distribution variables to the benchmark model. Both models are estimated using the XGBoost algorithm to avoid differences in the prediction method performance. The XGBoost algorithm is used to more efficiently find patterns in combinations of IPC classes in the portfolio and the quantity or appearance of patents within IPC classes associated with targets or non-targets. Unsurprisingly, considering the unexpected results from the patent dummy variable, we find no apparent signs of the IPC-class distribution having any predictive power. The benchmark model returns an AUC-ROC of 0.1805 and an AUC-PR of 0.6765. The model with the IPC-class variables returns an AUC-ROC of 0.6539 and an AUC-PR of 0.1568. Thus, the model performs worse compared to the model with the benchmark variables. This can be explained by changes in desirable patent profiles from the training set to the test set. There are only 21 actual targets in the test data with a patent portfolio. Consequently, the patent profile of targets would have had to be extraordinarily apparent for the predictive power to increase significantly.

6.2 Stock exchange announcements

The focus of this section is to investigate whether potential variables found by analyzing announcements from the companies on Oslo Stock Exchange have predictive power and can add value on the margin to the benchmark prediction model from Khan and Myrholt (2018). Due to poor data quality in our sample between 1999 and 2002, we exclude these years from the data set when examining the announcement variables. Firstly, we examine variables related to the number of stock exchange announcements, both in total and by category. Secondly, we analyze the sentiment variables constructed in chapter 4.1.5.

6.2.0.1 Number of announcements

Similar to the patent analysis, in Table 6.4 (model 1), we observe a negative and statistically significant coefficient for the number of announcements published in the calendar year before a takeover bid. When adjusting for a dummy variable stating whether a company has published an announcement within the last calendar year, we see that the observed significance is related to this effect, not the number of announcements published (model 2).

In theory, all companies should be obliged to publish some announcements. However, a significant number of observations are companies that have not published any announcements over the year of observation. As seen in Figure 6.2, nearly all the observations stem from before 2008. This can either be related to the new updated “Lov om verdipapirhandel” (Ministry of Finance, 1997) or the adapted “Børsloven” (Ministry of Finance, 2007). Another potential explanation is the demand for more transparency from shareholders following the financial crisis and more follow-up from the Norwegian Financial Supervisory Authority (Finanstilsynet). Nevertheless, this potential variable will have no predictive power in our models, as all companies today publish announcements. Anyhow, it is still an interesting finding. We are uncertain as to why this variable has such a high correlation with acquisitions. However, we suspect it can be related to the storage of information for companies prior to the financial crisis that no longer exists. Nevertheless, the number of published announcements within each category is highly correlated with the announcement dummy variable. Thus, we add it as a control variable to account for this effect.

Table 6.2: Number of Companies Not Publishing Announcements

Year	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2017	2018	
Companies	58	45	46	50	40	47	29	11	2	0	0	0	0	0	0

Since there are 25 different types of announcements, we run a backward elimination method in conjunction with economic rationale when searching for significance of the categorical variables. This means that we regress all the size announcement variables coherently on the dependent variable. Each round, the category variable with either the highest p-value or an unreasonable economic rationale is eliminated. This procedure is repeated until the variables in model 3 remain. When controlled for by the benchmark predictors and the announcement dummy, the only variable that was statistically significant is the variable named "ÅRSOVERSIKT". The variable should, in theory, not affect the takeover likelihood of a firm. However, the variable excelled during the elimination procedure. Table 6.3 shows an overview of the yearly number of "ÅRSOVERSIKT" announcements published.

Table 6.3: Number of Companies Publishing 'Årsoversikt'

2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018
0	0	0	121	147	233	205	186	201	200	0	0	0	0	0	0

The "ÅRSOVERSIKT" announcements were only published in the years between 2006 and 2012. The announcement category came as a result of a new law regulation stating that "issuers of transferable securities listed on a regulated market within the EU/EEA are required to publish an annual overview of information made available to the public." (OSLO BØRS ASA, 2007). The regulation was canceled in 2012. The number of takeover bids was especially high during this period, which is likely why we observe the significance.

Furthermore, we observe that the "EKS.DATO" variable is close to significant at a 10 percent level with a negative coefficient. This announcement is published after a dividend is paid and indicates whether a company pays dividends. Thus, the results can indicate that companies that pay dividends are less likely to receive a takeover bid. The remaining category variables in model 3 do not provide predictive power on the margin when controlling for the benchmark variables.

Table 6.4: Regression of Number of Announcements

	Dependent variable:		
	dependent_variable		
	(1)	(2)	(3)
Total Announcements	−0.010***	−0.0004	
	p = 0.00002	p = 0.847	
Announcement Dummy		−1.794***	−1.865***
		p = 0.000	p = 0.000
INNSIDEINFORMASJON			0.003
			p = 0.391
‘MELDEPLIKTIG HANDEL FOR PRIMÆRINNSIDERE‘			−0.006
			p = 0.344
ÅRSOVERSIKT			0.211**
			p = 0.033
EKS.DATO			−0.071
			p = 0.164
‘BØRSPAUSE / HANDELSPAUSE‘			−0.036
			p = 0.470
Growth	0.017	0.015	0.013
	p = 0.321	p = 0.450	p = 0.482
Return on Equity	0.009	0.012	0.012
	p = 0.654	p = 0.504	p = 0.506
Revenue	−0.00001	−0.00001	−0.00001
	p = 0.301	p = 0.157	p = 0.172
Cash/Total Assets	−2.043***	−2.249***	−2.254***
	p = 0.00003	p = 0.00002	p = 0.00002
Current Ratio	0.003**	0.004**	0.004***
	p = 0.037	p = 0.012	p = 0.008
Ownership Concentration	−0.005	0.024	0.056
	p = 0.986	p = 0.925	p = 0.829
Industry Disturbance	0.026	0.026	0.025
	p = 0.393	p = 0.394	p = 0.411
Brent Crude Oil Price	0.009***	0.012***	0.010***
	p = 0.006	p = 0.0003	p = 0.007
10 Year Government Bond Yield	0.195***	−0.019	−0.065
	p = 0.007	p = 0.814	p = 0.475
Constant	−3.167***	−1.687***	−1.329***
	p = 0.000	p = 0.0003	p = 0.009
Observations	2,744	2,744	2,744
Log Likelihood	−665.907	−641.290	−636.969
Akaike Inf. Crit.	1,353.814	1,306.581	1,305.938
Note:			*p<0.1; **p<0.05; ***p<0.01

6.2.0.2 Sentiment of announcements

Moreover, we examine the predictive power of the sentiment variables. Observations where a company has not published announcements in the last calendar year, will obtain a sentiment of 0 for all sentiment types. Thus, the announcement dummy variable is used as a control variable to account for this effect when analyzing the impact of the sentiment variables.

Initially, from Table 6.5 we observe that the average yearly sentiment using the LM dictionary (model 2) is closing up to statistically significant at a 5% level with a p-value at 0.057. Thus, we can reject our null hypothesis with some degree of certainty, saying that the likelihood of receiving a takeover bid is unaffected by target-company stock exchange announcements. The coefficient is negative, suggesting that companies with a positive or less negative tone, on average, in their announcements are less likely to be acquired. Furthermore, the positive sentiment using the LM dictionary (model 6) is also statistically significant at a 10% level with a negative coefficient. This suggests that companies publishing announcements with a high proportion of words included in the positive LM word list is less likely to be acquired in the following year. The negative sentiment variable from the LM Dictionary (model 4) is also moving towards statistically significant at a 10% level. The coefficient is positive, with the same rationale that companies publishing announcements with a high proportion of words included in the negative LM word list is more likely to be acquired in the following year. There is an obvious pattern in the observable variables. A more negative- or less positive sentiment in the announcements can, in theory, have an indirect causal effect on the probability of a company receiving a takeover bid. Acquiring firms and managers are more likely to find potential opportunities in companies where one observes negative announcements more frequently or less positive announcements in general.

Table 6.5: Regression of Sentiment Variables

	Dependent variable:							
	dependent_variable							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Yearly_SentimentHE	−15.681							
	p = 0.543							
Yearly_SentimentLM		−15.313*						−13.061
		p = 0.057						p = 0.108
Yearly_NegativeHE			−77.716					
			p = 0.311					
Yearly_NegativeLM				13.014				
				p = 0.115				
Yearly_PositiveHE					−26.843			
					p = 0.321			
Yearly_PositiveLM						−70.449*		
						p = 0.092		
Yearly_RatioUncertaintyLM							−4.219	
							p = 0.861	
Announcement Dummy	−1.770***	−2.066***	−1.749***	−2.070***	−1.710***	−1.599***	−1.804***	−2.053***
	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.000
3-Year Market Returns								−0.407*
								p = 0.089
Growth	0.014	0.015	0.014	0.016	0.014	0.013	0.015	0.020
	p = 0.453	p = 0.424	p = 0.468	p = 0.422	p = 0.463	p = 0.486	p = 0.453	p = 0.282
Return on Equity	0.012	0.012	0.013	0.012	0.012	0.012	0.012	0.014
	p = 0.510	p = 0.520	p = 0.479	p = 0.520	p = 0.500	p = 0.492	p = 0.501	p = 0.433
Revenue	−0.00001	−0.00001	−0.00001	−0.00001	−0.00001	−0.00001	−0.00001	−0.00001
	p = 0.118	p = 0.109	p = 0.116	p = 0.112	p = 0.116	p = 0.116	p = 0.121	p = 0.141
Cash/Total Assets	−2.279***	−2.314***	−2.332***	−2.285***	−2.322***	−2.366***	−2.264***	−2.255***
	p = 0.00001	p = 0.00001	p = 0.00001	p = 0.00001	p = 0.00001	p = 0.00001	p = 0.00001	p = 0.00001
Current Ratio	0.004**	0.004***	0.004**	0.004***	0.004**	0.004**	0.004**	0.004***
	p = 0.012	p = 0.007	p = 0.012	p = 0.007	p = 0.014	p = 0.014	p = 0.011	p = 0.009
Ownership Concentration	0.017	0.025	0.025	0.029	0.010	0.009	0.029	0.024
	p = 0.946	p = 0.921	p = 0.923	p = 0.909	p = 0.970	p = 0.972	p = 0.909	p = 0.925
Industry Disturbance	0.026	0.017	0.027	0.018	0.026	0.026	0.026	0.009
	p = 0.403	p = 0.586	p = 0.384	p = 0.556	p = 0.404	p = 0.396	p = 0.393	p = 0.772
Brent Crude Oil Price	0.012***	0.011***	0.012***	0.011***	0.012***	0.012***	0.012***	0.011***
	p = 0.0003	p = 0.002	p = 0.0003	p = 0.001	p = 0.0003	p = 0.0003	p = 0.0003	p = 0.002
10 Year Government Bond Yield	−0.024	0.015	−0.027	0.010	−0.028	−0.022	−0.023	0.012
	p = 0.770	p = 0.857	p = 0.737	p = 0.905	p = 0.734	p = 0.792	p = 0.778	p = 0.881
Constant	−1.656***	−1.732***	−1.638***	−1.728***	−1.629***	−1.654***	−1.670***	−1.698***
	p = 0.0003	p = 0.0002	p = 0.0004	p = 0.0002	p = 0.0004	p = 0.0003	p = 0.0003	p = 0.0002
Observations	2,744	2,744	2,744	2,744	2,744	2,744	2,744	2,744
Log Likelihood	−641.119	−639.520	−640.749	−640.089	−640.791	−639.762	−641.294	−638.011
Akaike Inf. Crit.	1,306.238	1,303.041	1,305.497	1,304.179	1,305.582	1,303.524	1,306.587	1,302.023
Note:								*p<0.1; **p<0.05; ***p<0.01

The total sentiment variable using the LM Dictionary is 97% correlated with the negative sentiment variable. This indicates that nearly all the information in the average sentiment variable stems from the negative sentiment dictionary. However, including the information from the positive sentiment dictionary is what makes the variable statistically significant. This makes sense as the total sentiment variable should capture more information than the two other variables. The reason why the average sentiment is so correlated with the negative sentiment variable comes from the construction of the positive and negative word lists created by Loughran and McDonald (2011). The negative- and positive LM word lists contain 2355 and 354 words, respectively. Evidence suggests that negative corporate news is often conveyed using positive words (e.g., "did not benefit"), while positive corporate news is rarely announced using negative words (Loughran and McDonald, 2011). The rationale behind the imbalance of positive and negative words is to diminish the effects of positive phrasing of negative news by only including positive words deemed not to be "easily compromised" (Loughran and McDonald, 2011).

The HE Dictionary variables do not show any significance. Additionally, we observe that the negative sentiment variable from the HE Dictionary (model 3, Table 6.5) has the opposite sign of the LM Dictionary (model 4, Table 6.5). Henry's Dictionary was the first published finance-specific word list, containing 85 negative- and 105 positive words (Loughran and McDonald, 2015). As none of the most frequently used negative LM words in their sample of over 50,000 10-K forms are included in the HE negative word list, Loughran and McDonald (2015) argue that the HE list is not an exclusive list of negative words used by managers to describe the operations of a company. The large difference in the number of included words in the LM- and HE negative word lists is likely the reason for the conflicting results. As previously stated, we focus mainly on the sentiment results from the LM Dictionary.

To test the average sentiment variable more thoroughly, we try to find correlating predictors in the data set that may explain the observed marginal contribution to the benchmark model. Initially, we observe that the average sentiment variable is negatively correlated with macro variables that have increased over time, such as the consumer price index and the household consumption index. At the same time, it is negatively correlated with macro variables that have decreased in our time frame, such as government bond yield. This indicates that the average sentiment in the market has decreased over time.

Figure 6.1: Yearly Average Sentiment of All Stock Exchange Announcements

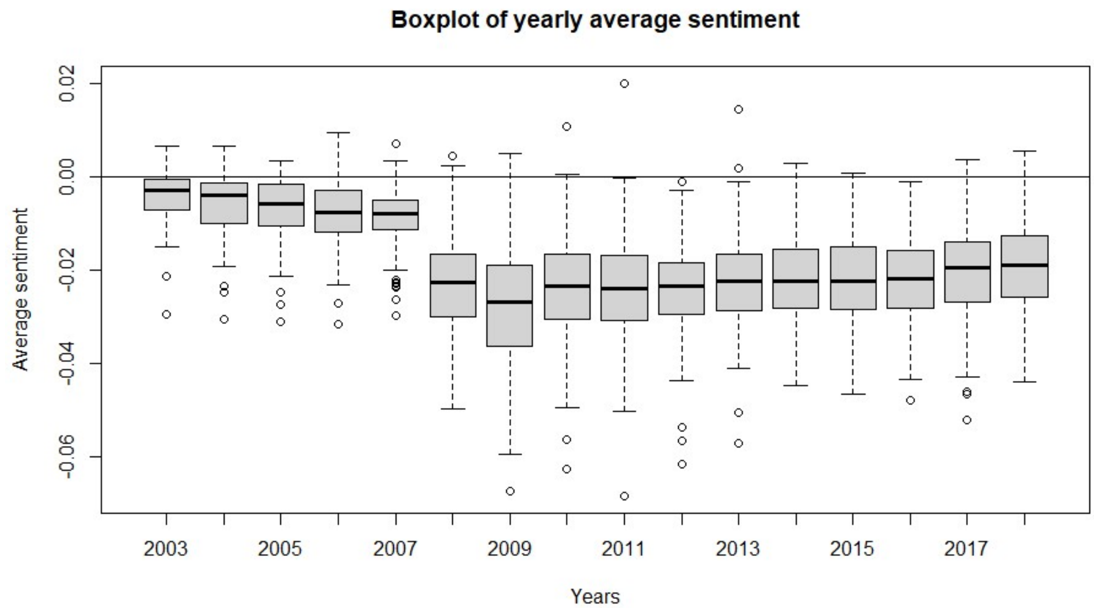

In Figure 6.1 the average sentiment over the years for all companies having published one or more announcements is displayed. Following the discussion on the construction of LM Dictionaries and the correlation between the average sentiment and the negative sentiment variables, the effect is evident in the plot. There are very few observations with positive sentiment on average. We observe a downward spike in sentiment in the years around the financial crisis, which is seemingly very natural. However, the sentiment stabilizes at a much lower level in the following years as well. Again, this can indicate the increased demand for transparency in the markets following the market dissatisfaction in the aftermath of the financial crisis and higher pressure on the authorities to follow up on questionable practices regarding shareholder information. This will cause a natural correlation with the aforementioned macro variables, with a steady increase or decline in the time frame of our observations. However, this correlation is already controlled for in the model with the 10-year government bond yield.

Additionally, we observe that the sentiment variable is highly correlated with the oil price. Surprisingly, the correlation between the variables is negative. That is likely due to lower oil prices on average in the years before the financial crisis compared to the years after, where we observe the change in sentiment level. However, we observe a spike of negative sentiment during the oil crisis (Appendix A5.1) when the oil price is low. This suggests that if we were to aggregate the pre-financial crisis sentiment level to the level after, we would likely observe a positive correlation, where low oil prices lead to more negative sentiment. Nevertheless, this variable is also controlled for in the benchmark model.

In terms of company-specific control variables beyond the benchmark model, we find the average sentiment variable to positively correlate with the stock return variables from 2-5 years. Meaning that the better the returns are, the better the sentiment tone we observe on average in the company's announcements over the year. These variables are another indication of the firm's and management's performance, in addition to market beliefs. Thus, the correlation is expected. Controlling for a 3-year return variable (model 8, Table 6.5), we observe that the return variable obtains a negative coefficient with a p-value of 0.089 while the average sentiment variable is moving towards statistically significant at a 10% level with a p-value of 0.108. Both variables express the same type of publicly readily available information about the performance of a company. Thus, these are likely two variables that potential acquirers are looking at in conjunction with each other when evaluating targets.

In conclusion, we see that the number of published announcements both total and by category does not likely influence the takeover likelihood of a company. However, in 2002-2009, we observe that companies that did not publish announcements were significantly more likely to be acquired than those that did. Nevertheless, this variable will be useless in our prediction model, as all companies publish announcements today. In conclusion, we observe a significant negative correlation between the average sentiment of a company's announcements and the dependent variable, which adds to the predictive ability of the benchmark model on the margin.

6.3 Algorithmic approach

This section aims to examine whether the introduction of an algorithmic prediction approach, XGBoost, improves the predictive power of the models. Using both logistic regression and XGBoost, we fit models using the benchmark variables and the patent and announcement extensions separately and combined. Additionally, we fit an XGBoost model using all the 188 variables shown in Table 4.1. We do not fit a logistic regression model with all 188 variables as this would likely lead to problems such as overfitting without further examination of each variable. The final XGBoost model (188 variables) is included mainly to test whether introducing many potential predictors in combination with XGBoost's ability to perform feature selection automatically outperforms the models using the predictors established by Khan and Myrholt, with and without our extensions. As this model does not have a logistic regression counterpart, it will not be the focus of our discussion.

The table below (Table 6.6) shows the frequency of actual targets in each decile for each model. Decile 1 contains the companies with the lowest estimated takeover probability, while decile 10 contains the companies with the highest estimated takeover probability. Comparing the XGBoost models to their respective logistic regression counterparts, we observe that the XGBoost models consistently place a higher number of actual targets in the top three deciles. At the same time, we do not observe a lower frequency of actual targets in the bottom three deciles in the XGBoost models compared to the logistic regression models. This may indicate that the models are able to detect the most obvious takeover targets. However, they seem to struggle once the most obvious targets are accounted for. We observe the biggest improvement when introducing the XGBoost model to the benchmark variables from Khan and Myrholt (2018), doubling the number of actual targets in the tenth decile compared to the logistic regression benchmark model. The best-performing model seems to be the XGBoost model with 188 variables. We discuss the differences in model performance further by examining the performance metrics.

Table 6.6: Distribution of Actual Targets Predicted to Each Decile

	Logistic Regression									
Decile	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Benchmark (KM)	2	4	7	7	10	5	8	6	9	7
KM + Patent	1	5	4	9	5	6	7	8	10	10
KM + Announcement	3	6	7	9	6	5	8	7	7	7
KM + Patent & Announcement	2	6	8	7	5	6	7	4	10	10
	XGBoost									
Decile	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Benchmark (KM + XGBoost)	4	5	4	5	7	6	4	10	5	15
KM + Patent	4	5	4	7	5	7	4	9	7	13
KM + Announcement	1	4	0	8	7	4	9	10	12	10
KM + Patent & Announcement	6	5	0	8	5	7	6	7	8	13
XGBoost w/188 variables	0	2	2	4	6	8	8	12	13	10

The performances of the models are evaluated using the AUC-ROC and the AUC-PR for each model. A no-skill model, choosing targets at random, would produce an AUC-ROC of 0.5 and an AUC-PR of approximately 0.056. A perfect prediction would result in an AUC-ROC of 1 and an AUC-PR of 1. For the models to show predictive skill, the values must be between 0.5 and 1 for the AUC-ROC and 0.056 and 1 for the AUC-PR. Figure 6.2 shows the ROC curve for the logistic regression benchmark model. The diagonal line indicates the performance of a no-skill model with an AUC-ROC of 0.5. The model's ROC curve indicates that the model performs slightly better than random. However, it is not a big improvement to a random prediction. Figure 6.2 also shows the PR curve of the same model. The horizontal dashed line represents the performance of a no-skill model as a constant equalling the percentage of targets in the entire test sample (5.6%). Again, we observe that the model seemingly performs slightly better than random. However, the curve is very close to that of a random estimator. The ROC- and PR curves for all the models discussed in this section are shown in the Appendix (section A6).

Figure 6.2: ROC and Precision-Recall from Model with Benchmark Variables using Logistic Regression

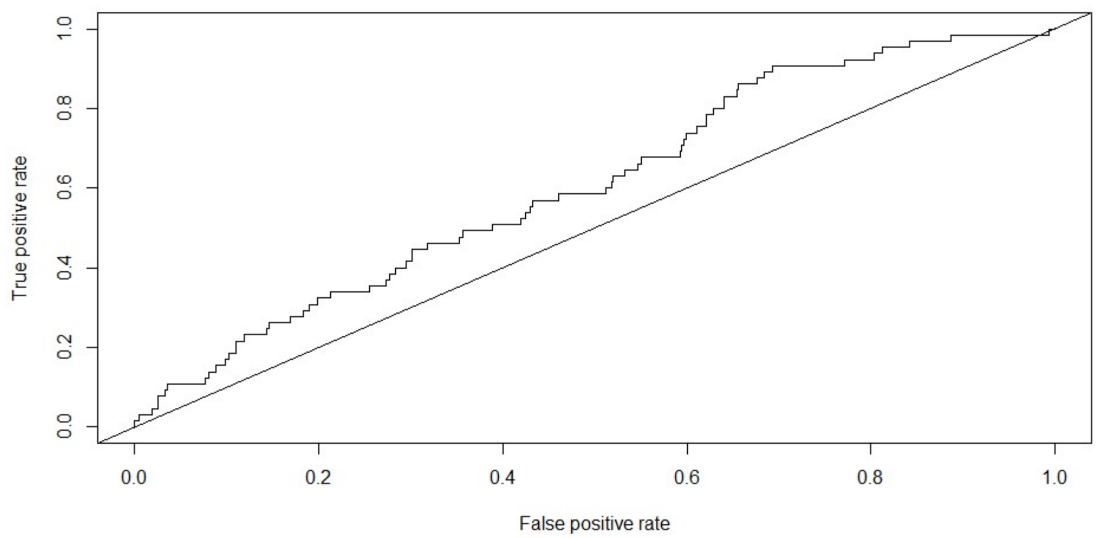

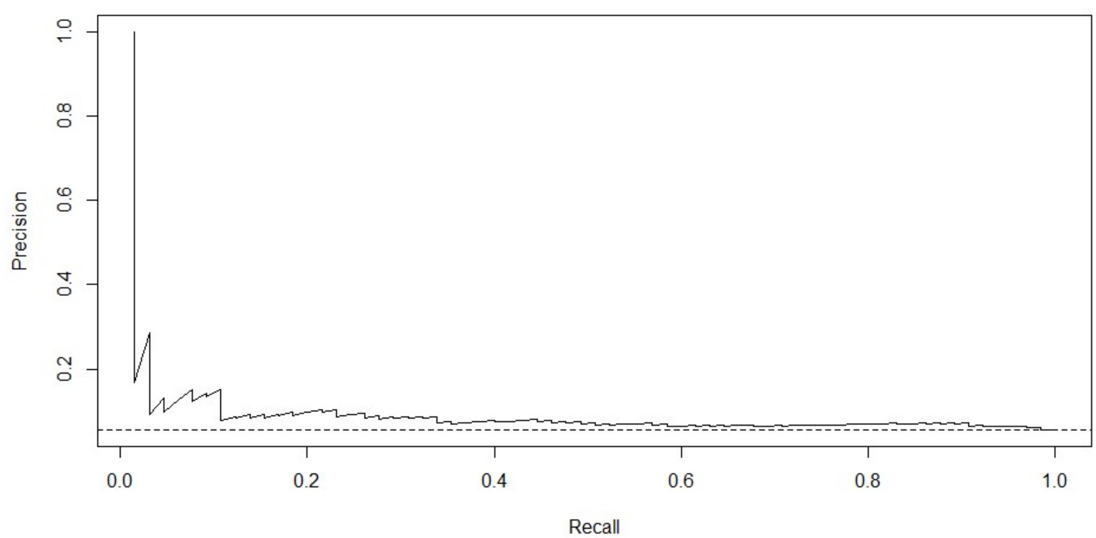

As we cannot accurately determine the difference in performance solely based on visual inspection of the ROC- and PR curves, we calculate the area under the curves (AUC), shown in Table 6.7. In the table, KM denotes Khan and Myrholt's (2018) independent variables, Patent indicates the inclusion of the patent dummy variable, and Announcement indicates the inclusion of the announcement dummy- and LM average yearly sentiment variables.

Table 6.7: AUC of ROC and Precision-Recall Curves for All Models

	Logistic Regression		**XGBoost**	
	PR	ROC	PR	ROC
Benchmark (KM)	0.1054	0.6072	0.1805	0.6765
KM + Patent	0.1183	0.6223	0.1550	0.6723
KM + Announcement	0.1018	0.5913	0.1406	0.5991
KM + Patent & Announcement	0.1050	0.5988	0.1887	0.6413
All Variables	-	-	0.1446	0.6931

The performance metrics show that none of the models perform very well, highlighting the difficulty of predicting takeover targets. However, the XGBoost models consistently outperform their logistic regression counterparts, both in terms of AUC-ROC and AUC-PR. The performance metrics confirm our observation from the decile table that the largest overall improvement is observed when using the benchmark variables in the XGBoost model compared to logistic regression.

The AUC-PR values all seem to improve when introducing the XGBoost algorithm. To our knowledge, there are no acknowledged tests for comparing PR curves. Thus, we cannot conclude whether the improvements are statistically significant. The AUC-PR is a better measure than AUC-ROC when evaluating a model's ability to predict the minority class of an imbalanced data sample, in this case, the actual targets. Hence, the XGBoost models are tuned to optimize the AUC-PR in the validation set when training the models. Examination of the PR curves from the models, included in section A6 of the Appendix, shows high precision levels before rapidly converging towards the line indicating a no-skill prediction. This strengthens our observation that the models are able to identify the most obvious takeover targets but struggle once the most obvious targets are identified.

To test whether the improvement in AUC-ROC from introducing the XGBoost algorithm is statistically significant, we use DeLong's (1988) test for comparing ROC curves. We test each XGBoost model against its respective logistic regression counterpart. Additionally, we test the final XGBoost model containing 188 independent variables against the logistic regression benchmark model. The results of the tests are shown in Table 6.8, below.

Table 6.8: Comparison of AUC-ROC between XGBoost Models and Logistic Regression Models

	Logistic Regression AUC-ROC	**XGBoost** AUC-ROC	**P-Value** (DeLong)
Benchmark (KM)	0.6072	0.6765	0.07412
KM + Patent	0.6223	0.6723	0.1464
KM + Announcement	0.5913	0.5991	0.8609
KM + Patent & Announcement	0.5988	0.6413	0.3883

	Benchmark	**XGBoost (All variables)**	**P-Value**
XGBoost w/188 variables	0.6072	0.6931	0.0632

Using the Khan and Myrholt (2018) independent variables in the XGBoost model, we observe an AUC-ROC of 0.6765 compared to the 0.6072 from the logistic regression. This results in a p-value from the DeLong-test of 0.074, which is significant at a 10% level. Although we observe minor improvements in the AUC-ROC for the other models, we cannot conclude that they are statistically significant. A potential reason why the XGBoost algorithm seemingly performs slightly better than logistic regression is its ability to capture complex relationships between the features in the data. Additionally, XGBoost's ability to reduce the weight of less important features may diminish the effects of noise in the data, potentially affecting the prediction results.

From the tables, we observe that the highest AUC-ROC (0.6931) is achieved by the final XGBoost model with all the variables in Table 4.1. Comparing the AUC-ROC of the final XGBoost model to the logistic regression benchmark model results in a p-value of 0.0632, moving towards being statistically significant at a 5% level. However, compared to the XGBoost model with Khan and Myrholt's (2018) independent variables, the improvement is minimal and not significant. Additionally, we observe a lower AUC-PR value. There could be numerous reasons why the extensive model does not significantly outperform the XGBoost model with the benchmark variables. Although target firms have recognizable features, making it possible to predict better than a random model, predicting the timing of the acquisition announcement has proven difficult. A firm may possess all the characteristics of an acquisition target. However, without any interested acquirers, there will never be a takeover bid.

Additionally, the financial market is an ever-changing environment, where desirable features in a potential target one year may not be desirable later. These changes can lead to the algorithm being trained to recognize patterns in the training set that no longer hold when the predictions are made on the test set. Including several independent variables without spending time on feature engineering increases the risk of these pitfalls.

In conclusion, we observe minimally higher AUC-ROC and AUC-PR values when extending Khan and Myrholt's (2018) logistic regression model with the patent dummy variable. Extending the Khan and Myrholt (2018) model by including announcement data on its own and by including announcement data in combination with patent data both result in minimally lower AUC-ROC and AUC-PR values. When introducing the XGBoost algorithm, we observe consistently higher AUC-ROC and AUC-PR values compared to the logistic regression models. However, without an acknowledged statistical test for comparing PR curves and only one model showing improvement statistically significant at a 10% level, we cannot conclude that XGBoost significantly outperforms logistic regression. The final XGBoost model results in the highest AUC-ROC and is the closest to being significantly better than the logistic regression benchmark model at a 5% level. However, we are not able to reject our null hypothesis (H0.3), stating that: "Applying a gradient boosting prediction technique does not improve the predictive power of the model in Khan & Myrholt (2018)".

6.4 Target portfolio performance

In the final section, we compare the returns of the best algorithmic approach models against the benchmark portfolio obtained from the previous section. This includes the benchmark variables using logistic regression (benchmark model), the benchmark variables using XGBoost (xgboost1), the benchmark variables plus the patent dummy and average yearly sentiment using XGBoost (xgboost2), and the "all variable" model using XGBoost (xgboost3). The portfolios are created following Powell (2001) maximum target strategy as described in section (literature review). The proportion of actual targets within each portfolio is relatively low. Table 6.9 presents the number of predicted targets and the number of actual targets within each portfolio per year.

Table 6.9: Number of Predicted Targets and Number of Actual Targets in the Takeover Portfolios

	2013	2014	2015	2016	2017	2018	2019	Total	Target proportion
Benchmark model:									
Portfolio size	134	44	108	16	15	16	16	349	
Targets in portfolio	12	5	4	1	1	0	0	23	6.59%
XGBoost with benchmark variables:									
Portfolio size	90	22	117	16	159	8	127	539	
Targets in portfolio	10	4	5	1	9	0	7	36	6.68%
XGBoost with benchmark variables, patents and announcements:									
Portfolio size	20	17	7	16	6	16	16	98	
Targets in portfolio	2	3	3	2	1	1	1	13	13.3%
XGBoost with "All variables" :									
Portfolio size	18	6	22	12	13	16	10	97	
Targets in portfolio	1	0	2	1	2	1	2	9	9.28%

We see that the two models using the benchmark variables predict more companies to the portfolio than the portfolios with more variables added. This is likely as more variables increase the model fit in the training set, which leads to the decile with the highest proportion of targets being the decile with the highest takeover probabilities. Consequently, this causes a higher cut-off probability and a lower number of predicted companies to the portfolio. Looking at the precision-recall curves in section A6 of the Appendix, we observe that the models are likely to have a higher proportion of targets with a smaller portfolio, as the models manage to capture the more obvious targets. In contrast, after 10% to 20% of the actual targets are predicted, the models do not perform better than a no-skill predictor. We would observe a higher proportion of actual targets if we were to invest in the highest predicted deciles instead of the cut-off probabilities yielding the maximum number of targets in the training set. This can be observed in Table 6.6, looking at the proportion of targets in the upper deciles.

In Figure 6.3 we plot the cumulative returns over the test period for all the target portfolios and the all-share index. We observe relatively similar patterns in the target portfolios. This indicates that the portfolios consist of companies with the same features. Initially, we observe that the portfolios created using the XGBoost models outperform the portfolio using the logistic regression model. However, we can not say much more about the performance of the portfolios before examining the risk-adjusted returns.

Figure 6.3: Cumulative Returns of the Target Portfolios from 2013 to 2019

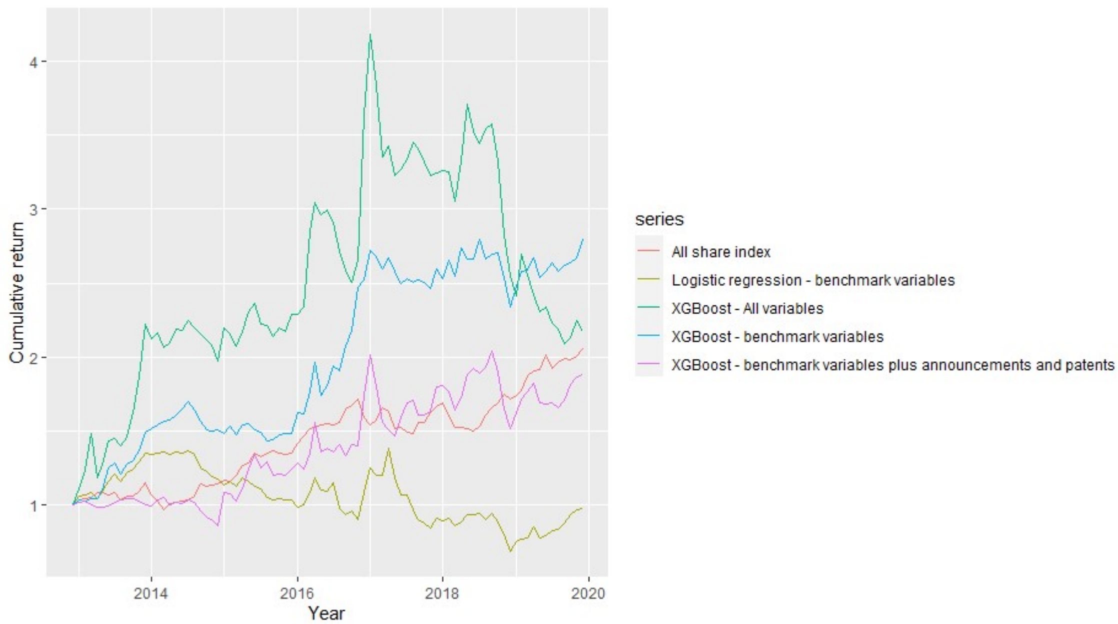

The portfolio of the predicted targets is expected to consist of companies with a certain type of features. Thus, it is interesting to observe how these portfolios perform relative to systematic market risk factors independently and in relation to each other. Due to the low proportion of actual targets in the portfolios, not much of the returns can be attributed to the announcement returns. However, the stock price of a potential target can often be driven by the expectations of a takeover bid, as a bid for the company often involves a premium. Much of the abnormal return from an acquisition stems from the period before the announcement. This is due to anticipation of a takeover bid or leakage of inside information. If the price of a potential target that in the end is not acquired in a given calendar year act similarly, an efficient market should, in theory, push the prices up due to anticipation of a bid that in most cases involves a premium. Thus, even though the target portfolios consist of a low proportion of targets, it should be possible to achieve abnormal returns based on market anticipation for the non-targets in the portfolio. We examine the portfolios further by looking at risk-adjusted returns. The four portfolios are tested individually against the Fama-French three-factor model using a long-only strategy. In Table 6.10 we present the result of the regressions.

Table 6.10: Target Portfolio Market Performance Evaluation

	Dependent variable:				
	benchmark_excess	xgboost1_excess	xgboost2_excess	xgboost3_excess	only_targets_excess
	(1)	(2)	(3)	(4)	(5)
market_excess	0.940^{***}	0.585^{***}	0.787^{***}	1.102^{***}	0.279
	p = 0.00003	p = 0.0005	p = 0.003	p = 0.0003	p = 0.118
SMB	0.596^{**}	0.615^{***}	0.483^{*}	1.228^{***}	-0.392^{*}
	p = 0.015	p = 0.002	p = 0.096	p = 0.0004	p = 0.056
HML	0.373^{***}	0.286^{***}	0.454^{***}	0.595^{***}	-0.010
	p = 0.010	p = 0.010	p = 0.009	p = 0.004	p = 0.931
Constant	-0.009	0.005	-0.001	-0.004	0.017^{***}
	p = 0.168	p = 0.311	p = 0.923	p = 0.696	p = 0.005
Observations	84	84	84	84	84
R^2	0.257	0.228	0.174	0.257	0.098
Adjusted R^2	0.229	0.199	0.143	0.230	0.064
Residual Std. Error (df = 80)	0.057	0.044	0.069	0.079	0.048
F Statistic (df = 3; 80)	9.231^{***}	7.865^{***}	5.633^{***}	9.248^{***}	2.886^{**}
Note:					*p<0.1; **p<0.05; ***p<0.01

Initially, we observe that the models fit the data rather poorly, with the adjusted-R^2 ranging between 0.13 to 0.24, telling us that the three factors explain between 13% to 24% of the portfolios' monthly returns. Neither of the portfolios achieves a significant monthly abnormal return (alpha). Thus, the observed excess returns in Figure 6.10 can be explained by higher risk exposure. As expected, we observe the target portfolios to be exposed to similar factor risks to a greater or lesser extent. The interpretation of the factor coefficients is not always straightforward. A study conducted by Chen and Bassett (2014) shows that a portfolio consisting of 80% large-cap and 20% small-cap companies yielded a significantly positive SMB coefficient of 0.13, probably due to a dominance of large-cap stocks in the SMB factor portfolios. This might be different for the factor portfolios created by Ødegaard (2021) on Norwegian data. Nevertheless, a positive coefficient does not necessarily mean an overweight of small-cap companies in the portfolio.

Firstly, we look at the market risk factor, also known as the more traditional Beta (β) of a portfolio. All the portfolios yield a $\beta < 1$, indicating that the portfolios are less volatile than the rest of the market. This is unexpected based on what we see in Figure 6.3. However, in some periods, the portfolios yield a negative β, which likely equalizes the higher volatility in general.

Secondly, we consider the SMB factor. The benchmark variable model portfolios (model 1 model 2) are correlated with the SMB portfolio with a coefficient of around 0.5. Thus, it is likely that the portfolios consist of an overweight of small-cap companies. This means that if the small-cap portfolio increases by 1%, these portfolios are expected to increase by around 0.5%. This is in line with the negative coefficient previously observed for revenue in the logistic regression models. Interestingly, model 3, including patents and announcements, is not significantly correlated with the SMB portfolio. As seen in Table 6.1, the patent dummy is positively correlated with size variables. Thus, we would expect to see an increase in small-cap companies for this portfolio. However, the model has likely found other variables more influential when assembling the decision trees and calculating the probabilities. The "all variable" portfolio (model 4), on the other hand, is highly correlated with the SMB portfolio, suggesting that size variables or variables correlated with size have been prominent in the estimation of this model.

Thirdly, we examine the HML factor. We observe a significantly positive correlation for all the target portfolios with the HML portfolio, indicating a higher proportion of value than growth companies in the target portfolios. A value stock is a security trading at a lower price than what the company's performance otherwise may indicate, and thus a natural property for a potential takeover target. When considering the portfolios only using the benchmark variables (models 1 2), there are no obvious properties of the variables that suggest investing in value stocks. Additionally, the coefficients are around 0.3, which is in the gray area of whether to state that the portfolios include an overweight of value companies. This also holds for the portfolio from the model with patents and announcements (model 3). The "all variable" portfolio (model 4) surely has a higher proportion of value stocks or stocks that has acted similarly to that of value stocks over the test period. Features prominent for value stock companies such as low price-earnings (P/E) ratio and low market-to-book (P/B) ratio yield a negative coefficient when used as control variables in Table 6.1 (model 8), suggesting that value stocks are more likely takeover targets. However, the variables were not statistically significant. Nevertheless, the XGBoost algorithm may have picked up on combinations of these features in conjunction with other variables, increasing the takeover probability of a company. Moreover, the algorithm possibly considers other variables with properties of a value stock company, such as companies paying a high dividend yield.

Lastly, the portfolios are also tested against a momentum factor and a liquidity factor

(Appendix A7.1). However, these factors are not found to be statistically significant for any of the portfolios. Additionally, we test the portfolios by excluding actual targets to see the effect of these investments (Appendix A7.2). Interestingly, the results do not change in terms of abnormal returns for any of the models compared to the portfolios, including the targets, indicating that the announcement returns are either very low for the predicted targets or that the targets perform poorly in the period prior to acquisition announcements. Another reason could be in conjunction with the discussion on Precision-recall curves from section 6.3, suggesting that some takeovers are more obvious. According to the efficient market hypothesis, all publicly available information should be reflected in the price of a stock (Fama, 1970). Consequently, we likely do not obtain any abnormal return for these obvious takeovers, as this information is already reflected in the stock prices when we invest.

We examine a portfolio only consisting of only the actual targets (Table 6.10, model 5) to see if it is possible to obtain positive abnormal returns using a yearly rebalancing strategy. Here we observe a significant positive monthly abnormal return of 1.7%. Thus, the prediction of takeover targets is not in vain. One should disregard the risk factors to some extent in this model. These values come as a result of our investment strategy with targets and not necessarily the actual stocks within the portfolio, as we sell our positions right after the announcement is made and hold cash for the rest of the calendar year. This is reflected in the low adjusted-R^2.

Furthermore, we test the statistical differences in market risk and abnormal returns between the XGBoost model portfolios and the benchmark portfolio by implementing a long-short strategy. This is created by taking the returns of the algorithmic models' portfolios (long) and subtracting the benchmark portfolio returns (short), and regressing the long-short portfolio returns against the 3-factor model. The result of the coefficients will only yield the coefficients for both the intercept and the factors of the XGBoost model portfolios minus the respective coefficients for the benchmark model portfolio observed in Table 6.11. However, we run the regression to test for significance of the difference between the coefficients.

Table 6.11: XGBoost Target Portfolios Comparison to Benchmark Portfolio

	Dependent variable:		
	long_short_xgboost1	long_short_xgboost2	long_short_xgboost3
	(1)	(2)	(3)
market_excess	−0.345*	−0.179	0.076
	p = 0.073	p = 0.464	p = 0.772
SMB	−0.013	−0.388	0.466
	p = 0.954	p = 0.166	p = 0.123
HML	−0.100	−0.011	0.160
	p = 0.439	p = 0.947	p = 0.369
Constant	0.014**	0.012	0.008
	p = 0.027	p = 0.119	p = 0.343
Observations	84	84	84
R^2	0.051	0.029	0.031
Adjusted R^2	0.016	−0.008	−0.005
Residual Std. Error (df = 80)	0.052	0.067	0.072
F Statistic (df = 3; 80)	1.441	0.788	0.855
Note:			*p<0.1; **p<0.05; ***p<0.01

We observe that the portfolio created by the benchmark variables and the XGBoost algorithm significantly outperforms the benchmark model, with a monthly abnormal return of 1.4%. More interestingly is why we observe such a significant difference between the two portfolios, using the same variables when predicting targets. Firstly, we observe that the adjusted-R^2 (Table 6.11, model 1) is only 0.016, indicating that the factor portfolios explain the expected returns very poorly. We check whether the observed return can stem from the higher number of targets in the XGBoost portfolio. When comparing the portfolios by excluding the actual targets (Appendix A7.4, model 1), we still observe a monthly abnormal return at 1.5%, indicating that the excess returns cannot be attributed to the higher number of targets in the portfolio.

Furthermore, examining Appendix A7.5 (model 1), we observe that even though the momentum and liquidity variables are not significant independently, they still mitigate the monthly abnormal return to 0.8% and not statistically significant when included. The coefficients show positive signs for these variables, suggesting that the XGBoost model is more exposed to momentum- and liquidity risk than the benchmark model. The two remaining portfolios (Figure 6.11, model 2 & 3) are not found to obtain significantly positive abnormal returns.

In conclusion, neither of the target portfolios achieve a significant positive abnormal return when adjusted for the 3-factor model. When testing a portfolio consisting of only actual targets, we observe a significant monthly abnormal return of 1.7%. Thus, it is possible to achieve abnormal returns if one manages to increase the models' predictive skills. Targets that are harder to predict probably yield higher abnormal returns, which makes the task of achieving abnormal returns from takeover target portfolios even more difficult. A long-short portfolio using the XGBoost model (long) and the logistic regression model with the benchmark variables yields a monthly positive abnormal return of 1.4% when adjusted for the three-factor model. However, this abnormal return can not be attributed to the higher number of actual targets in the XGBoost portfolio. When including the momentum and liquidity factors, we no longer find any significance. Thus, we cannot obtain positive abnormal returns of a portfolio assembled using a gradient boosting algorithm compared to the logistic regression function.

7 Conclusion and Further Research

The aim of our thesis was to extend the existing takeover prediction literature on the Norwegian MA market through extension of the Khan and Myrholt (2018) model. We test the marginal contribution of new variables related to intellectual property and target firm announcements and whether an algorithmic prediction method outperforms the logistic regressions traditionally used in takeover prediction literature.

In analyzing intellectual property variables, we use patent data to determine a company's technological drift and as a proxy of its technological profile. We find that neither the technological drift nor the technological profile of a company has a significant effect on a company's probability of receiving a takeover bid when controlling for the variables in Khan and Myrholt (2018). Thus, they do not provide a significant marginal contribution to predictive power. We do, however, find that companies without a patent portfolio are significantly more likely to receive a takeover bid after controlling for the variables in Khan and Myrholt (2018).

The number of stock exchange announcements shows no significant effect on a company's probability of receiving a takeover bid after controlling for the variables in Khan and Myrholt (2018). Thus, the announcement quantity variables provide no significant marginal contribution to predictive power. To calculate the sentiment of the companies' stock exchange announcements, we use the positive- and negative word lists from Loughran and McDonald's Master dictionary. The three variables, average yearly sentiment, positivity ratio, and negativity ratio, suggest that companies with a more negative tone in their announcements are more likely to be the target of a takeover bid. However, we do not find significance beyond a 10% level after controlling for the variables in Khan and Myrholt (2018) and cannot conclude that the likelihood of a takeover bid increases with target-company stock exchange announcements.

Predicting our test sample using logistic regression models, we observe minor improvements in the AUC values for the Receiver Operating Characteristics (ROC) curve and Precision-Recall (PR) curve when extending the Khan and Myrholt (2018) model with patent data independently. However, the improvements are not significant. Extending the Khan and Myrholt (2018) model with announcement data independently and in combination with patent data result in minor deteriorations of the AUC-ROC and AUC-PR values. When introducing a gradient boosting prediction technique (XGBoost) to the predictions, we observe better AUC values for the Receiver Operating Characteristics (ROC) curves and the Precision-Recall curves compared to the corresponding logistic regression models. However, only the XGBoost model with the Khan and Myrholt (2018) variables shows an improvement significant at a 10% level. Thus, we cannot conclude that the predictive power is improved by applying a gradient boosting prediction technique.

Our test sample shows that it is possible to achieve abnormal returns significantly higher than the OSEAX after adjusting for the Fama-French 3-factors model, with a perfect prediction of takeover targets. However, none of our models show sufficient predictive skill to significantly outperform the "All share" index when adjusted for systematic risk factors. Comparing the portfolios constructed from the XGBoost models to the one from Khan and Myrholt's (2018) logistic regression with a long-short strategy, we find no significant positive abnormal returns from the XGBoost portfolios. This substantiates that the predictive power is not significantly improved by applying a gradient boosting prediction technique.

Our research shows that a portfolio consisting of only takeover targets yields monthly abnormal returns of 1.7% when adjusted for systematic risk factors when using a yearly rebalancing strategy. This finding motivates further research to improve the models for takeover prediction. We suggest further extending this study by incorporating quarterly data to increase the information fed to the prediction models. Additionally, by performing feature engineering, it should be possible to exclude variables changing over time which causes noise to prediction models. Including a year variable and constructing a model which weights the information of observations in later years more should also able the model to cope better with changes in market trends.

Moreover, there will always be new variables to be considered, which may increase the predictive power. Data quality is improving with higher standards in accounting practices, and more players working with storage of data and normalization of income statements will improve the information fed into the prediction models. As more data points are collected, one should also be able to use more advanced machine learning algorithms such as artificial neural nets in takeover target prediction. However, the amount and quality currently available for the Norwegian market are insufficient to utilize its full potential.

The findings on intellectual property were unexpected. We were not able to convert the marginal effect of the patent dummy variable to the predictive power of the models tested, suggesting that the observed effect is weaker in recent time or has reached a turning point. Interestingly, we found that the proportion of intangible assets positively correlated with the takeover probability, indicating the importance of intellectual property in target prediction. For further research, we recommend extending the research of intellectual property more extensively. One way to explore the use of patents in target prediction could be to use the SAO analysis method to cluster similar patents and assign a rank to each cluster based on the proportion of target companies possessing one or more patents within a given cluster. Thereafter, assign an overall acquisition probability score to each company based on its patent portfolio and the cluster scores to which each of their patents belongs. Another approach could be to examine the R&D capability of a firm more thoroughly with variables such as total number of patent citations, number of citations per patent, number of inventors, number of patents per inventor, number of citations per inventor, and a company's yearly spending on R&D.

The sentiment analysis conducted in our thesis is applied to stock exchange announcements. A way to further explore the applications of sentiment analyses in takeover target prediction is to focus on different text corpora. Most studies predicting acquisition participants in the U.S. focus on the MD&A-section of companies' annual reports or financial news articles. However, due to limitations to available data and computation power, we have not incorporated this in our thesis. There are several other natural language processing (NLP) techniques that could potentially produce better predictors. A way to further explore the use of NLP to predict takeover targets is to generate custom word lists containing the words shown to have the highest impact on a company's probability of receiving a takeover bid.

References

Aabø-Evensen, O. K. (2021). Norway: Mergers & acquisitions laws and regulations 2021. https://iclg.com/practice-areas/mergers-and-acquisitions-laws-and-regulations/norway. Retrieved May 11, 2021.

Alimov, A. and Officer, M. S. (2017). Intellectual property rights and cross-border mergers and acquisitions. *Journal of Corporate Finance*, 45:360–377.

Ambrose, B. W. and Megginson, W. L. (1992). The role of asset structure, ownership structure, and takeover defenses in determining acquisition likelihood. *Journal of Financial and Quantitative Analysis*, pages 575–589.

Betton, S., Eckbo, B. E., and Thorburn, K. S. (2008). Corporate takeovers. *Handbook of empirical corporate finance*, pages 291–429.

Bourne, M., Dircks, T., Kardatzke, J., Miller, D., Roberts, M., and White, C. (2019a). From ml to ma: Ten ma target predictions through a machine learning model. Technical report, Nicholas Center for Corporate Finance & Investment Banking, Wisconsin School of Business.

Bourne, M., Dircks, T., Kardatzke, J., Miller, D., Roberts, M., and White, C. (2019b). Process manual: From ml to ma: Ten ma target predictions through a machine learning model. Technical report, Nicholas Center for Corporate Finance & Investment Banking, Wisconsin School of Business.

Brar, G., Giamouridis, D., and Liodakis, M. (2009). Predicting european takeover targets. *European Financial Management*, 15(2):430–450.

Breiman, L. et al. (2001). Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3):199–231.

Central Bank of Norway (2020a). Exchange rates. https://www.norges-bank.no/en/topics/Statistics/exchange_rates/. Retrieved September 27, 2020.

Central Bank of Norway (2020b). Interest rates. https://www.norges-bank.no/en/topics/Statistics/Interest-rates/. Retrieved September 22, 2020.

Chen, H.-l. and Bassett, G. (2014). What does smb > 0 really mean? *Journal of Financial Research*, 37(4):543–552.

Danbolt, J., Siganos, A., and Tunyi, A. (2016). Abnormal returns from takeover prediction modelling: challenges and suggested investment strategies. *Journal of Business Finance & Accounting*, 43(1-2):66–97.

DeLong, E. R., DeLong, D. M., and Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, pages 837–845.

Dieudonne, S., Cretin, F., and Bouacha, S. (2014). M&a activity: where are we in the cycle? *Available at SSRN 2504412.*

DNB Markets (2020). Exchange rates. https://www.dnb.no/bedrift/markets/valuta-renter/valutakurser-og-renter/konverteringskurser-euro.html+&cd=2&hl=en&ct=clnk&gl=no. Retrieved September 27, 2020.

Eckbo, B. E. and Solibakke, P. (1991). Bedriftsoppkjøp og internasjonalisering: Norge og europa i 1980-årene. *Beta: Tidsskrift for bedriftsøkonomi*, 5:1–30.

Ellert, J. C. (1976). Mergers, antitrust law enforcement and stockholder returns. *The Journal of Finance*, 31(2):715–732.

EY (2017). Ey capital insights. http://www.ey.com/gl/en/services/transactions/ey-capital-insights-16-the-impact-of-falling-oil-prices-on-m-a?utm_source=ey-capital-insights&utm_medium=web&utm_campaign=eycom. Retrieved November 15, 2020.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417.

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

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of financial economics*, 116(1):1–22.

Frey, R. and Hussinger, K. (2006). The role of technology in m&as: a firm-level comparison of cross-border and domestic deals. Technical report, Discussion Paper Series 1.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.

Goergen, M. and Renneboog, L. (2004). Shareholder wealth effects of european domestic and cross-border takeover bids. *European Financial Management*, 10(1):9–45.

Graffin, S. D., Haleblian, J., and Kiley, J. T. (2016). Ready, aim, acquire: Impression offsetting and acquisitions. *Academy of Management Journal*, 59(1):232–252.

Heinze, G. and Dunkler, D. (2017). Five myths about variable selection. *Transplant International*, 30(1):6–10.

Heldal, J. and Fosen, J. (2006). Logistisk regresjon - kurskompendium i byråskolens kurs sm507. https://www.ssb.no/a/publikasjoner/pdf/notat_200654/notat_200654.pdf. Retrieved May 22, 2021.

Hellebust, H. and Sandvik, S. (2020). Norway m&a and ipo report 2019. https://wiersholm.no/en/newsletter/norway-ma-and-ipo-report-2019/.

Huang, L., Guo, Y., Zhao, Y., Wang, Y., and Porter, A. L. (2013). The role of text mining of patent in mergers and acquisitions. In *2013 Proceedings of PICMET'13: Technology Management in the IT-Driven Services (PICMET)*, pages 2271–2282. IEEE.

Hussinger, K. (2010). On the importance of technological relatedness: Smes versus large acquisition targets. *Technovation*, 30(1):57–64.

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). *An introduction to statistical learning*, volume 112. Springer.

Jensen, M. C. and Ruback, R. S. (1983). The market for corporate control: The scientific evidence. *Journal of Financial economics*, 11(1-4):5–50.

Katsafados, A. G., Androutsopoulos, I., Chalkidis, I., Fergadiotis, E., Leledakis, G. N., and Pyrgiotakis, E. G. (2021). Using textual analysis to identify merger participants: Evidence from the us banking industry. *Finance Research Letters*, page 101949.

Khan, N. and Myrholt, T. (2018). Predicting norwegian takeover targets: an empirical analysis of the norwegian m&a market. Master's thesis, NHH.

Kim, S., Park, I., and Yoon, B. (2020). Sao2vec: Development of an algorithm for embedding the subject–action–object (sao) structure using doc2vec. *Plos one*, 15(2).

Lall, R. (2016). How multiple imputation makes a difference. *Political Analysis*, 24(4):414–433.

Langetieg, T. C. (1978). An application of a three-factor performance index to measure stockholder gains from merger. *Journal of Financial Economics*, 6(4):365–383.

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. (2015). The use of word lists in textual analysis. *Journal of Behavioral Finance*, 16(1):1–11.

Loughran, T. and McDonald, B. (2020). Textual analysis in finance. *Annual Review of Financial Economics*, 12:357–375.

Ma, T., Zhang, Y., Huang, L., Shang, L., Wang, K., Yu, H., and Zhu, D. (2017). Text mining to gain technical intelligence for acquired target selection: A case study for china's computer numerical control machine tools industry. *Technological Forecasting and Social Change*, 116:162–180.

Mandelker, G. (1974). Risk and return: The case of merging firms. *Journal of financial economics*, 1(4):303–335.

Manski, C. F. and McFadden, D. (1981). Alternative estimators and sample designs for discrete choice analysis. *Structural analysis of discrete data with econometric applications*, 2:2–50.

Markusen, J. R. (1995). The boundaries of multinational enterprises and the theory of international trade. *Journal of Economic perspectives*, 9(2):169–189.

Martynova, M. and Renneboog, L. (2011). The performance of the european market for corporate control: Evidence from the fifth takeover wave. *European financial management*, 17(2):208–259.

Medby, K. O. and Nordgård, A. (2019). Maskinlæring for analyse av børsmeldinger og aksjekursprediksjon. Master's thesis, NTNU.

Ministry of Finance (1997). Lov om verdipapirhandel. https://lovdata.no/dokument/NLO/lov/1997-06-19-79. Retrieved October 20, 2020.

Ministry of Finance (2007). Lov om regulerte markeder (boersloven). https://lovdata.no/dokument/LTI/lov/2007-06-29-74. Retrieved February 20, 2021.

Mousavi, N. A. (2011). The evolving role of intellectual property in m&a transactions. *Intellectual Asset Management*, 48:28–34.

Nielsen, D. (2016). Tree boosting with xgboost-why does xgboost win" every" machine learning competition? Master's thesis, NTNU.

Ødegaard, B. A. (2021). Empirics of the oslo stock exchange. basic, descriptive, results 1980-2020. https://ba-odegaard.no/wps/empirics_ose_basics/empirics_ose_basics_2021_03.pdf. Version from March 2021, Retrieved May 22, 2021.

Office, N. I. P. (2020). Patents act. https://www.patentstyret.no/en/norwegian-patents-act. Retrieved October 10, 2020.

OSLO BØRS ASA (2007). Annual overview of published information. https://www.oslobors.no/Oslo-Boers/Regelverk/Boerssirkulaerer-arkiv/5-2007-OPPHEVET-AArlig-oversikt-over-offentliggjort-informasjon. Retrieved February 21, 2021.

Palepu, K. G. (1986). Predicting takeover targets: A methodological and empirical analysis. *Journal of accounting and economics*, 8(1):3–35.

Park, H., Yoon, J., and Kim, K. (2013). Identification and evaluation of corporations for merger and acquisition strategies using patent information and text mining. *Scientometrics*, 97(3):883–909.

Powell, R. G. (2001). Takeover prediction and portfolio performance: A note. *Journal of Business Finance & Accounting*, 28(7-8):993–1011.

Routledge, B. R., Sacchetto, S., and Smith, N. A. (2017). Predicting merger targets and acquirers from text. Technical report, Working Paper, Carnegie Mellon University.

Shmueli, G. et al. (2010). To explain or to predict? *Statistical science*, 25(3):289–310.

Statista (2020a). Average annual brent crude oil price from 1976 to 2020. https://www.statista.com/statistics/262860/uk-brent-crude-oil-price-changes-since-1976/. Retrieved September 22, 2020.

Statista (2020b). Share of the total volume of merger and acquisition (m&a) deals that were cross-border worldwide from 1985 to 2020. Retrieved May 11, 2021.

Statistics Norway (2020a). Consumer price index. https://www.ssb.no/en/priser-og-prisindekser/konsumpriser/statistikk/konsumprisindeksen. Retrieved September 22, 2020.

Statistics Norway (2020b). Household consumption. https://www.ssb.no/en/inntekt-og-forbruk. Retrieved September 22, 2020.

Statistics Norway (2020c). Tax statistics for companies. https://www.ssb.no/en/virksomheter-foretak-og-regnskap/statistikker/skattepl. Retrieved September 22, 2020.

The Financial Supervisory Authority of Norway (2020). Act on securities trading. https://www.finanstilsynet.no/globalassets/laws-and-regulations/laws/securities-trading-act.pdf. Version from March 2021, Retrieved October 15, 2020.

WIPO (2020). Guide to the international patent classification, version 2020. Retrieved May 11, 2021.

Appendix

A1 Patent Data

Figure A1.1: Excerpt Copied from WIPO's "Guide to the International Patent Classification" (2020)

II. LAYOUT OF CLASSIFICATION SYMBOLS
Section; Class; Subclass; Group; Complete classification symbol

SECTION

19. The Classification represents the whole body of knowledge which may be regarded as proper to the field of patents for invention, divided into eight sections. Sections are the highest level of hierarchy of the Classification.

(a) **Section Symbol** – Each section is designated by one of the capital letters A through H.

(b) **Section Title** – The section title is to be considered as a very broad indication of the contents of the section. The eight sections are entitled as follows:

A HUMAN NECESSITIES
B PERFORMING OPERATIONS; TRANSPORTING
C CHEMISTRY; METALLURGY
D TEXTILES; PAPER
E FIXED CONSTRUCTIONS
F MECHANICAL ENGINEERING; LIGHTING; HEATING; WEAPONS; BLASTING
G PHYSICS
H ELECTRICITY

(c) *[Deleted]*

(d) **Subsection** – Within sections, informative headings may form subsections, which are titles without classification symbols.

Example: Section A (HUMAN NECESSITIES) contains the following subsections:

AGRICULTURE
FOODSTUFFS; TOBACCO
PERSONAL OR DOMESTIC ARTICLES
HEALTH; LIFE SAVINGS; AMUSEMENT

CLASS

20. Each section is subdivided into classes which are the second hierarchical level of the Classification.

(a) **Class Symbol** – Each class symbol consists of the section symbol followed by a two-digit number.

Example: H01

(b) **Class Title** – The class title gives an indication of the content of the class.

Example: H01 BASIC ELECTRIC ELEMENTS

(c) **Class Index** – Some classes have an index which is merely an informative summary giving a broad survey of the content of the class.

Figure A1.2: IPC-class Raw Data after Initial Preprocessing

beskrivelse	status	dato	IPC	selskap	publikasjon	year	IPC_section	IPC_class
METHOD FOR AND EQUIPMENT FOR SUPPRESSING DISCOL...	Pending	2019-12-31	C22F1/047	NORSK HYDRO	US20200232080A1	2019	C	C22
HIGH VOLTAGE WET'MATE CONNECTION ASSEMBLY	Pending	2019-12-30	H01R13/523	AKER SOLUTIONS	NO20191536A1	2019	H	H01
IRON COMPOUND USED AS A CATALYST FOR HYDROSILYLA...	Pending	2019-12-30	C08F4/70	ELKEM	IN201917054401A	2019	C	C08
METHOD FOR PRODUCING SILICONE ELASTOMER MOULDS	Pending	2019-12-27	B29C33/38	ELKEM	IN201947054145A	2019	B	B29
DEVICE FOR STORING AND DELIVERING FERTILIZER MOBILE...	Pending	2019-12-27	A01C23/00	YARA INTERNATIONAL	CO2019014957A2	2019	A	A01
METHOD FOR OPERATING AN AGRICULTURAL SYSTEM AN...	Pending	2019-12-24	A01B63/22	KVERNELAND	IN201927053645A	2019	A	A01
ORGANIC'INORGANIC HYBRID POLYSILAZANE AND MANU...	Granted	2019-12-23	C08K3/36	DOF	KR102122032B1	2019	C	C08
ENHANCED STEERING CONTROL SYSTEM FOR PERSONAL ...	Pending	2019-12-23	B63H25/02	KONGSBERG AUTOMOTIVE	WO2020129033A1	2019	B	B63
HEAT EXCHANGER' USE OF ALUMINUM ALLOY AND ALUMI...	Pending	2019-12-23	C22C21/00	NORSK HYDRO	JP2020073721A	2019	C	C22
INSTALLING SUBSEA RISERS	Pending	2019-12-23	F16L1/12	SUBSEA 7	WO2020136378A1	2019	F	F16
SUBSEA WELLHEAD SYSTEMS AND METHODS	Pending	2019-12-20	E21B33/035	AKER SOLUTIONS	NO20191515A1	2019	E	E21
SUBSEA WELLHEAD SYSTEMS AND METHODS	Pending	2019-12-20	E21B33/035	AKER SOLUTIONS	NO20191516A1	2019	E	E21
METHOD FOR THE ADDITIVE MANUFACTURING OF A SILIC...	Pending	2019-12-20	B33Y70/00	ELKEM	WO2020127882A1	2019	B	B33
POWER CONVERTER AND METHOD OF CONTROLLING A P...	Pending	2019-12-20	H02M5/458	ELTEK	US20200204081A1	2019	H	H02
TREATMENT OF PRODUCED HYDROCARBONS	Pending	2019-12-20	C07C7/00	EQUINOR	WO2020130852A1	2019	C	C07
METHOD FOR SURFACE COATING IN AN AGRICULTURAL SP...	Pending	2019-12-20	B05B13/06	KVERNELAND	WO2020127893A1	2019	B	B05

Figure A1.3: IPC-class Variables after Variable Construction

conm	fyear	dependent_variable	A61	B01	B03	B04	B07	B21	B23	B29	B60	B63	B65	B66	B67	C01	C04
SPECTRUM	2018	1	0	2	0	0	0	0	0	0	0	0	0	0	0	0	0
BW LPG	2018	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
OCEAN YIELD	2018	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
NORSK HYDRO	2018	0	77	228	6	10	35	258	254	19	90	10	159	21	0	115	5
AQUA BIO TECHNOLOGY	2018	0	117	0	0	0	0	0	0	0	0	0	0	0	0	0	0
ELEMENT	2018	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
NORDIC MINING	2018	0	0	0	0	0	0	0	0	0	0	0	0	0	0	12	0
SCANA	2018	0	0	0	0	0	0	0	0	0	0	42	0	8	0	0	0
BORGESTAD	2018	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
ELKEM	2018	0	91	58	11	0	0	0	1	28	7	0	30	3	13	145	109
BORREGAARD	2018	0	10	0	0	0	0	0	0	0	0	0	4	0	0	0	7
YARA INTERNATIONAL	2018	0	8	353	0	1	2	0	0	1	0	0	43	0	0	279	65
ARENDALS FOSSEKOMPANI	2018	0	2	58	0	0	0	2	24	0	0	0	1	0	0	57	0
VOW	2018	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
MULTICONSULT	2018	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

A2 Stock Exchange Announcement Data

Table A2.1: Number of Announcements per Category

Announcement Category	Number of Observations
ANNEN INFORMASJONSPLIKTIG REGULATORISK INFORMASJON	16534
BGY - BØRSPAUSE	2
BØRSPAUSE / HANDELSPAUSE	3668
DERIVATMELDINGER	21
EKS.DATO	3150
ENDRINGER I RETTIGHETENE TIL AKSJER / VERDIPAPIRER	77
FLAGGING	7801
HALVÅRSRAPPORTER OG REVISJONSBERETNINGER	21953
IKKE-INFORMASJONSPLIKTIGE PRESSEMELDINGER	2242
INNSIDEINFORMASJON	63677
KAPITAL- OG STEMMERETTSENDRINGER	6171
MELDEPLIKTIG HANDEL FOR PRIMÆRINNSIDERE	28742
MELDING FRA ANDRE AKTØRER	106
MELDING FRA OSLO BØRS	122
NOTERING / OPPTAK AV VERDIPAPIRER	3130
PROSPEKT / OPPTAKSDOKUMENT	906
RAPPORTERING OM BETALING TIL MYNDIGHETER	3
RENTEREGULERING	849
SLUTTKURSER DERIVATER	1
SUSPENSJONER	279
SÆRLIG OBSERVASJON	1498
UTSTEDERS MELDEPLIKT VED HANDEL I EGNE AKSJER	2768
VALG AV HJEMSTAT	11
ÅRSOVERSIKT	1020
ÅRSRAPPORTER OG REVISJONSBERETNINGER	694

A2.1 Text Pre-Processing for Sentiment Analysis

To our knowledge, there are no Norwegian versions of either the LM- or HE dictionary. Although the English announcements were kept when available in the initial cleaning of the stock exchange announcement dataset, the dataset contains over 50 000 Norwegian announcements. To analyze the sentiment of the Norwegian announcements, both dictionaries were translated using Google's pre-trained machine learning translation models. Examining the translations, it became apparent that some of the Norwegian translations consisted of multiple words, e.g., "resigns" became "trekker seg". This would become an issue in the text analysis, as part of the processing involves "tokenizing" the text or splitting the text into single-word observations. To avoid this issue, we transformed every translation consisting of multiple words into a single word, e.g., "trekker seg" became "trekkerseg". This transformation also needed to be done for all the relevant combinations of words in the Norwegian announcements. To avoid searching through approximately 55 000 stock exchange announcements and manually identify every relevant combination of words, a custom function was written in R to automate the process. The function loops through every Norwegian translation, identifying the number of words, and generates the required regular expressions to identify and transform all the relevant combinations of words as described above.

In the pre-processing of the message texts from the announcements, punctuation, numbers, and stop words were removed. Examples of stop words are "it", "this" and, "that", which are words that can be removed without depriving the sentence of meaning. The "stopwords" package in R contains lists of stop words in English and Norwegian and was used to remove the stop words in the cleaning process. After cleaning the announcement texts, the announcements were tokenized into unigrams (single-word observations), and a Document Term Matrix (DTM) was constructed. In a DTM, each row represents a document (announcement), and each column represents a term (word). The values in the matrix denote the number of times a term is present in a document. Using the constructed DTM as input, the sentiment analysis was conducted using the "analyzeSentiment" function from the "SentimentAnalysis" package in R, with stemming set to FALSE. The function calculates the sentiment for every document according to a specified ruleset, in this case, which dictionaries to be used. The resulting data frame contained the sentiment score, positivity-, negativity-, and uncertainty ratio based on the LM Dictionary and the sentiment score, positivity- and negativity ratio based on the HE Dictionary for every single stock exchange announcement. The sentiment variables were created by calculating the mean of each of these variables for each company each year.

A3 Ownership Concentration Univariate Regression without Imputation

Table A3.1: Ownership Concentration Univariate Regression without Imputation

	Dependent variable:
	dependent_variable
Ownership Concentration	0.135
	p = 0.667
Constant	-2.605^{***}
	p = 0.000
Observations	2,623
Log Likelihood	-668.794
Akaike Inf. Crit.	1,341.588
Note:	*p<0.1; **p<0.05; ***p<0.01

A4 Patent and Stock Exchange Regressions Excluding Ownership Concentration

Table A4.1: Patent Analysis Regression without Ownership Concentration

	Dependent variable:							
	dependent_variable							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Patent Total	−0.001**	−0.0004						
	p = 0.038	p = 0.134						
Patent Dummy‘		−0.426***		−0.531***	−0.378***	−0.425***	−0.437***	−0.425***
		p = 0.004		p = 0.001	p = 0.010	p = 0.004	p = 0.004	p = 0.004
Technologic Drift			−0.775	0.181				
			p = 0.193	p = 0.774				
Growth					0.028*	0.028*	0.029*	0.027*
					p = 0.081	p = 0.083	p = 0.072	p = 0.096
Return on Equity					0.018	0.017	0.017	0.017
					p = 0.307	p = 0.330	p = 0.339	p = 0.434
Revenue					−0.00002*	−0.00002*	−0.00002**	−0.00002*
					p = 0.053	p = 0.059	p = 0.050	p = 0.057
Cash/Total Assets					−1.748***	−1.655***	−1.585***	−1.686***
					p = 0.0001	p = 0.0003	p = 0.001	p = 0.0002
Current Ratio					0.003**	0.003**	0.003**	0.003**
					p = 0.046	p = 0.028	p = 0.031	p = 0.026
Industry Disturbance					0.015	0.026	0.023	0.024
					p = 0.590	p = 0.361	p = 0.450	p = 0.407
Brent Crude Oil Price					0.009***	0.009***	0.009***	0.009***
					p = 0.002	p = 0.003	p = 0.003	p = 0.004
10 Year Government Bond Yield					0.112*	0.003	0.113*	0.121**
					p = 0.058	p = 0.982	p = 0.059	p = 0.040
(Godwill+Intangible)/Total Assets						0.785***	0.769***	0.768***
						p = 0.001	p = 0.002	p = 0.001
Year						−0.038		
						p = 0.353		
G-Sector 10							−0.134	
							p = 0.434	
G-Sector 35							−0.467	
							p = 0.233	
G-Sector 55							−0.973	
							p = 0.344	
Growth-Resource Mismatch								0.157
								p = 0.292
Price/Earnings								−0.000
								p = 0.989
Market-to-Book								−0.003
								p = 0.870
Debt Ratio								0.001
								p = 0.955
Constant	−2.562***	−2.388***	−2.582***	−2.388***	−3.248***	73.214	−3.337***	−3.469***
	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.000	p = 0.374	p = 0.000	p = 0.000
Observations	3,339	3,339	3,339	3,339	3,339	3,339	3,339	3,339
Log Likelihood	−822.370	−818.094	−825.832	−819.999	−798.705	−792.856	−791.675	−792.716
Akaike Inf. Crit.	1,648.740	1,642.188	1,655.665	1,645.998	1,617.409	1,609.712	1,611.351	1,615.433
Note:								*p<0.1; **p<0.05; ***p<0.01

Table A4.2: Number of Announcements Regression without Ownership Concentration

	Dependent variable:		
	dependent_variable		
	(1)	(2)	(3)
Total Announcements	−0.010***	−0.0004	
	p = 0.00002	p = 0.841	
Announcement Dummy		−1.794***	−1.866***
		p = 0.000	p = 0.000
INNSIDEINFORMASJON			0.003
			p = 0.393
‘MELDEPLIKTIG HANDEL FOR PRIMÆRINNSIDERE‘			−0.006
			p = 0.334
ÅRSOVERSIKT			0.210**
			p = 0.034
EKS.DATO			−0.070
			p = 0.167
‘BØRSPAUSE / HANDELSPAUSE‘			−0.036
			p = 0.470
Growth	0.017	0.015	0.013
	p = 0.321	p = 0.451	p = 0.485
Return on Equity	0.009	0.012	0.012
	p = 0.655	p = 0.500	p = 0.500
Revenue	−0.00001	−0.00001	−0.00001
	p = 0.300	p = 0.158	p = 0.174
Cash/Total Assets	−2.042***	−2.254***	−2.264***
	p = 0.00003	p = 0.00001	p = 0.00001
Current_Ratio	0.003**	0.004**	0.004***
	p = 0.037	p = 0.012	p = 0.008
Industry Disturbance	0.025	0.026	0.026
	p = 0.391	p = 0.388	p = 0.399
Brent Crude Oil Price	0.009***	0.012***	0.010***
	p = 0.006	p = 0.0003	p = 0.007
10 Year Government Bond Yield	0.195***	−0.020	−0.066
	p = 0.007	p = 0.809	p = 0.470
Constant	−3.169***	−1.679***	−1.313***
	p = 0.000	p = 0.0003	p = 0.009
Observations	2,744	2,744	2,744
Log Likelihood	−665.907	−641.295	−636.993
Akaike Inf. Crit.	1,351.814	1,304.590	1,303.985
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A4.3: Sentiment Regression without Ownership Concentration

	Dependent variable:							
	dependent_variable							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Yearly_SentimentHE	−15.802 p = 0.539							
Yearly_SentimentLM		−15.322* p = 0.057						−7.814 p = 0.368
Yearly_NegativeHE			−77.844 p = 0.310					
Yearly_NegativeLM				13.017 p = 0.115				
Yearly_PositiveHE					−26.916 p = 0.319			
Yearly_PositiveLM						−70.514* p = 0.091		
Yearly_RatioUncertaintyLM							−4.140 p = 0.863	
3-Year Market Returns								0.047 p = 0.886
Announcement Dummy	−1.770*** p = 0.000	−2.068*** p = 0.000	−1.750*** p = 0.000	−2.071*** p = 0.000	−1.710*** p = 0.000	−1.600*** p = 0.000	−1.806*** p = 0.000	−2.018*** p = 0.000
Growth	0.014 p = 0.454	0.015 p = 0.425	0.014 p = 0.469	0.016 p = 0.423	0.014 p = 0.463	0.013 p = 0.487	0.014 p = 0.455	0.017 p = 0.393
Return on Equity	0.012 p = 0.508	0.012 p = 0.516	0.014 p = 0.476	0.012 p = 0.516	0.012 p = 0.498	0.012 p = 0.490	0.012 p = 0.498	0.012 p = 0.487
Revenue	−0.00001 p = 0.119	−0.00001 p = 0.110	−0.00001 p = 0.117	−0.00001 p = 0.112	−0.00001 p = 0.116	−0.00001 p = 0.117	−0.00001 p = 0.121	−0.00001 p = 0.171
Cash/Total Assets	−2.283*** p = 0.00001	−2.318*** p = 0.00001	−2.337*** p = 0.00001	−2.290*** p = 0.00001	−2.324*** p = 0.00001	−2.368*** p = 0.00001	−2.269*** p = 0.00001	−2.281*** p = 0.00001
Current Ratio	0.004** p = 0.012	0.004*** p = 0.007	0.004** p = 0.012	0.004*** p = 0.007	0.004** p = 0.014	0.004** p = 0.014	0.004** p = 0.011	0.004*** p = 0.009
Industry Disturbance	0.026 p = 0.397	0.017 p = 0.579	0.027 p = 0.377	0.019 p = 0.547	0.026 p = 0.400	0.026 p = 0.393	0.027 p = 0.385	0.011 p = 0.719
Brent Crude Oil Price	0.012*** p = 0.0003	0.011*** p = 0.002	0.012*** p = 0.0003	0.011*** p = 0.001	0.012*** p = 0.0003	0.012*** p = 0.0003	0.012*** p = 0.0003	0.012*** p = 0.001
10 Year Government Bond Yield	−0.024 p = 0.766	0.015 p = 0.861	−0.028 p = 0.732	0.009 p = 0.910	−0.028 p = 0.731	−0.022 p = 0.789	−0.024 p = 0.772	−0.002 p = 0.979
Yearly_SentimentLM:3-Year Market Returns								30.457* p = 0.053
Constant	−1.650*** p = 0.0003	−1.723*** p = 0.0002	−1.629*** p = 0.0003	−1.718*** p = 0.0002	−1.626*** p = 0.0003	−1.651*** p = 0.0003	−1.660*** p = 0.0003	−1.671*** p = 0.0002
Observations	2,744	2,744	2,744	2,744	2,744	2,744	2,744	2,744
Log Likelihood	−641.121	−639.525	−640.753	−640.096	−640.792	−639.762	−641.300	−636.154
Akaike Inf. Crit.	1,304.243	1,301.051	1,303.507	1,302.192	1,303.583	1,301.525	1,304.601	1,298.309
Note:								*p<0.1; **p<0.05; ***p<0.01

A5 Sentiment of Stock Exchange Announcements

Figure A5.1: Yearly Average Negative Sentiment from all Stock Exchange Announcements

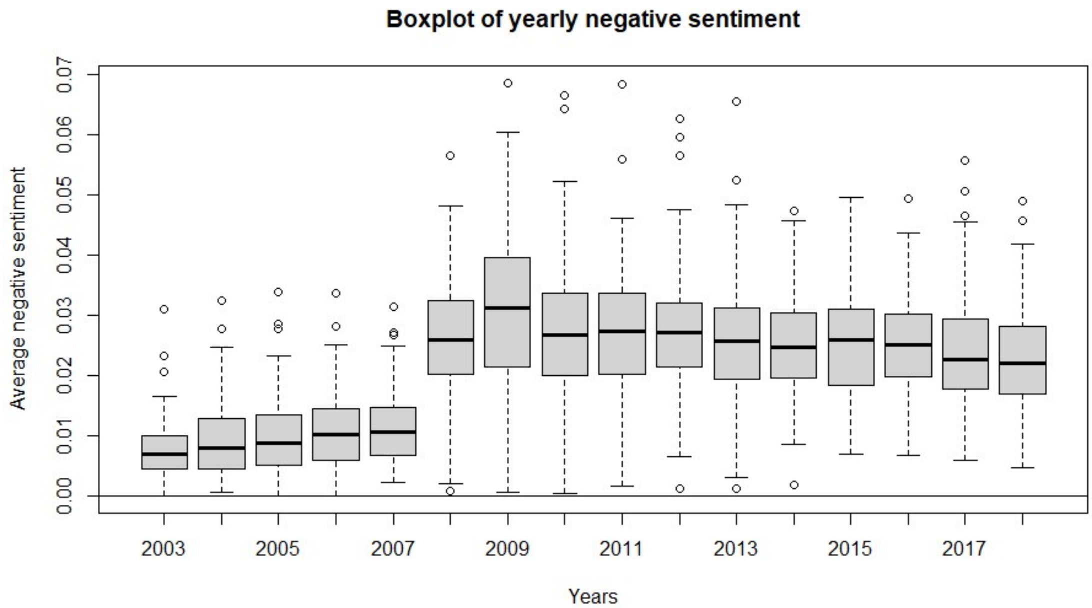

Figure A5.2: Yearly Average Positive Sentiment from all Stock Exchange Announcements

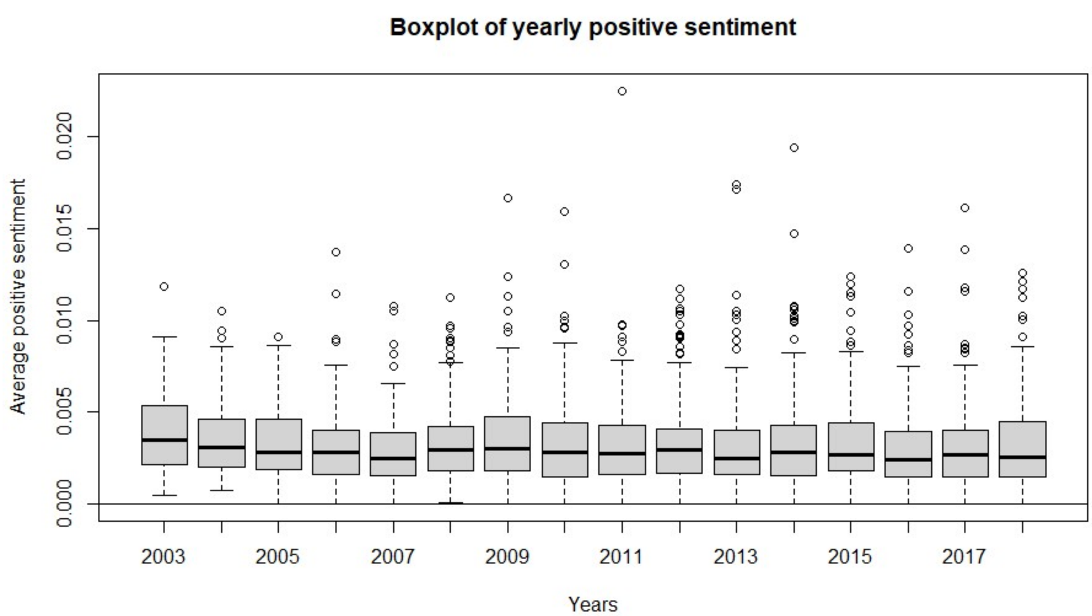

A6 ROC and Precision-Recall curves

Figure A6.1: ROC and Precision-Recall from Model with Benchmark Variables using XGBoost

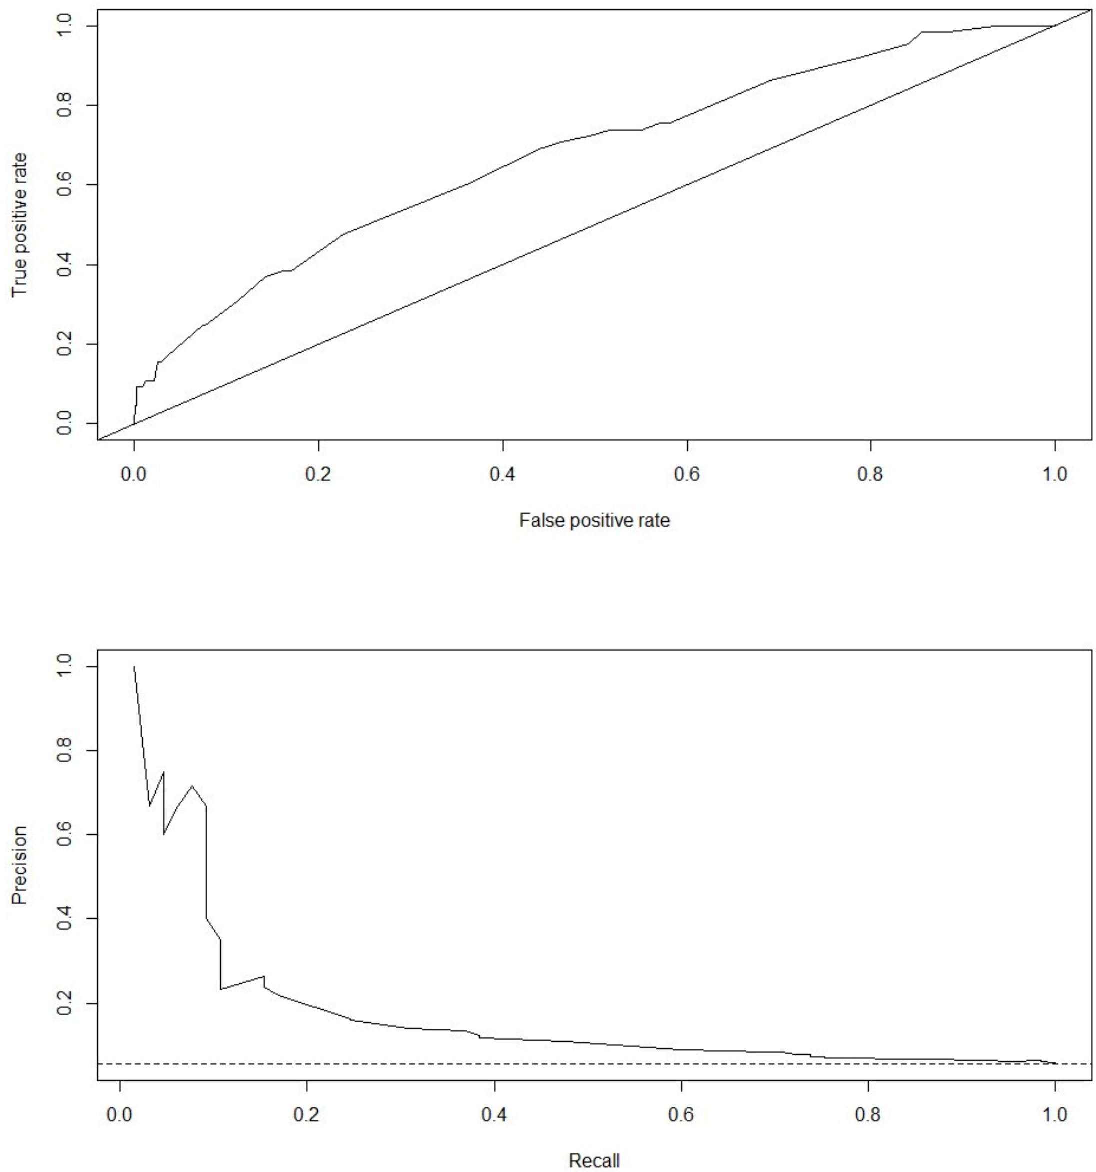

Figure A6.2: ROC and Precision-Recall from Model with Benchmark Variables plus Patents using Logistic Regression

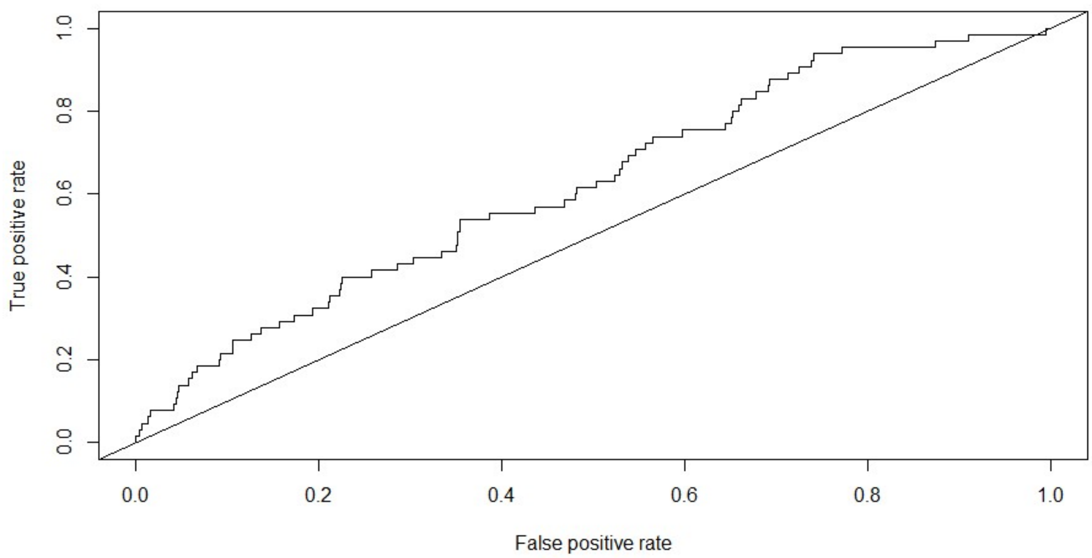

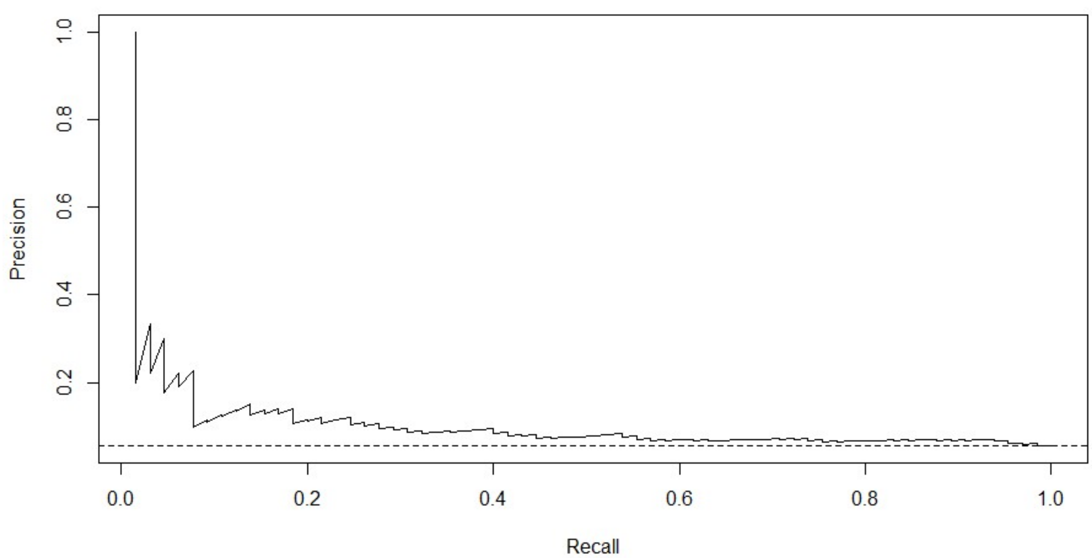

Figure A6.3: ROC and Precision-Recall from Model with Benchmark Variables plus Patents using XGBoost

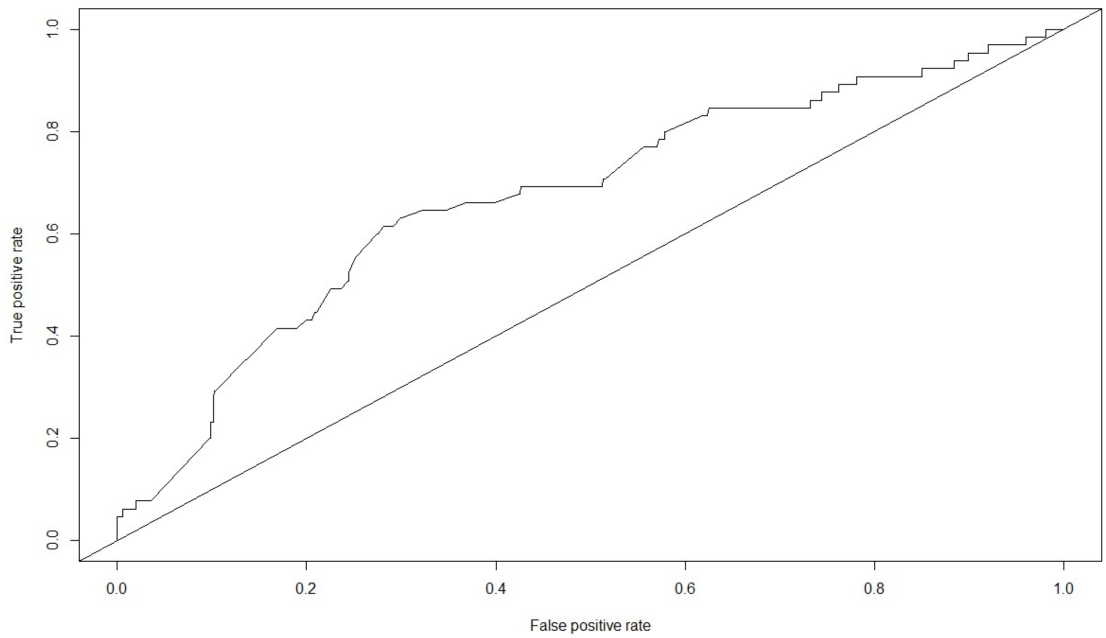

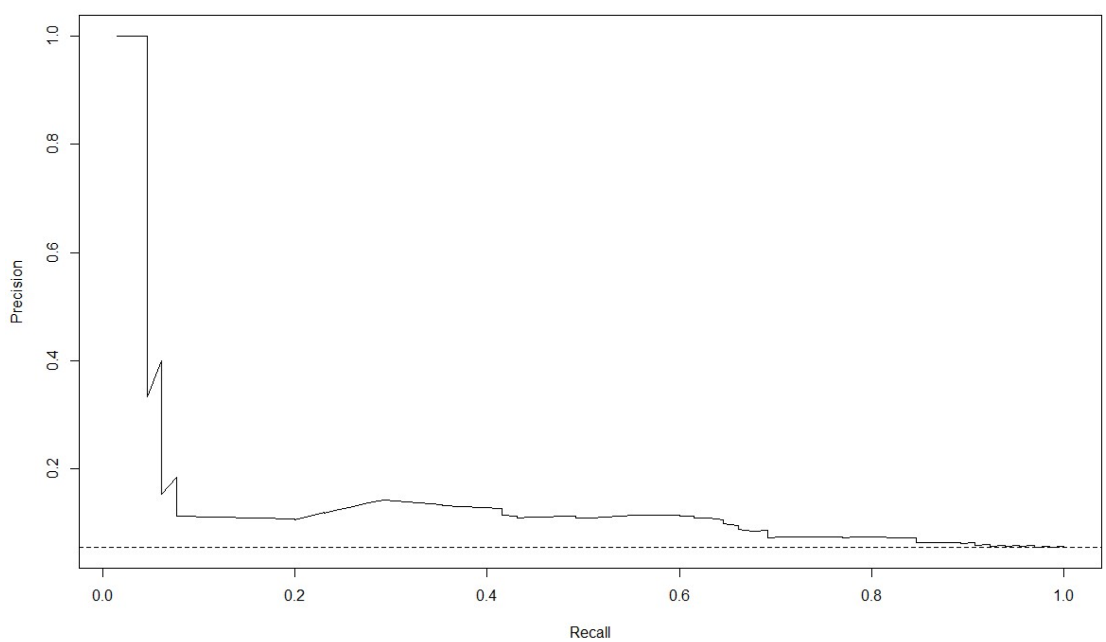

Figure A6.4: ROC and Precision-Recall from Model with Benchmark Variables plus Announcements using Logistic Regression

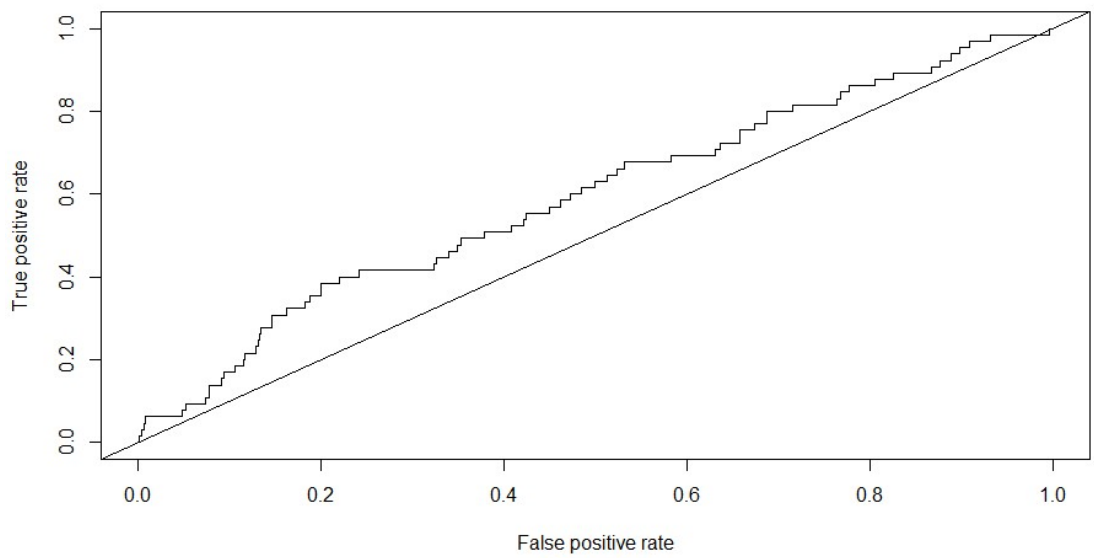

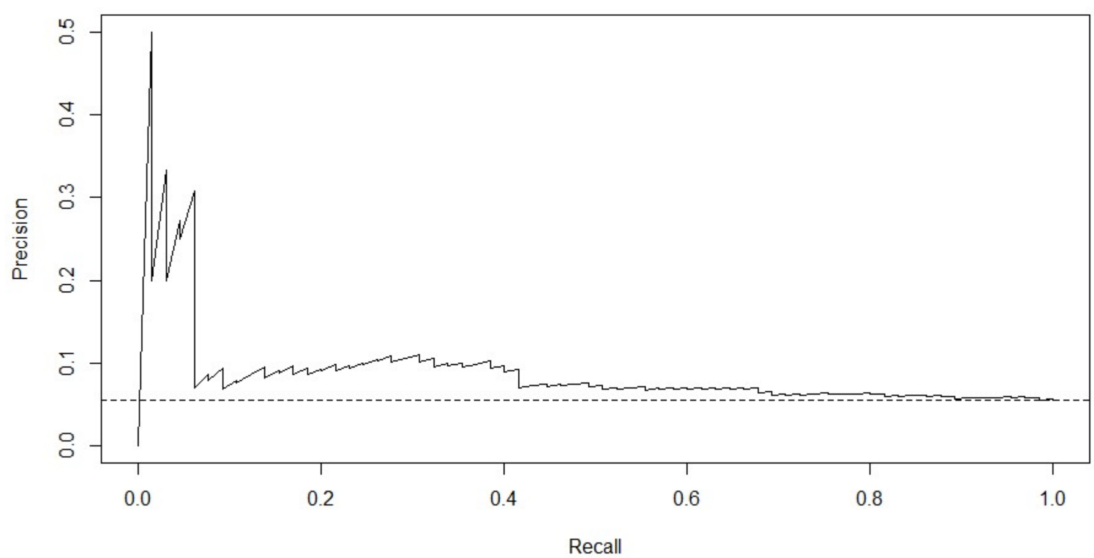

Figure A6.5: ROC and Precision-Recall from Model with Benchmark Variables plus Announcements using XGBoost

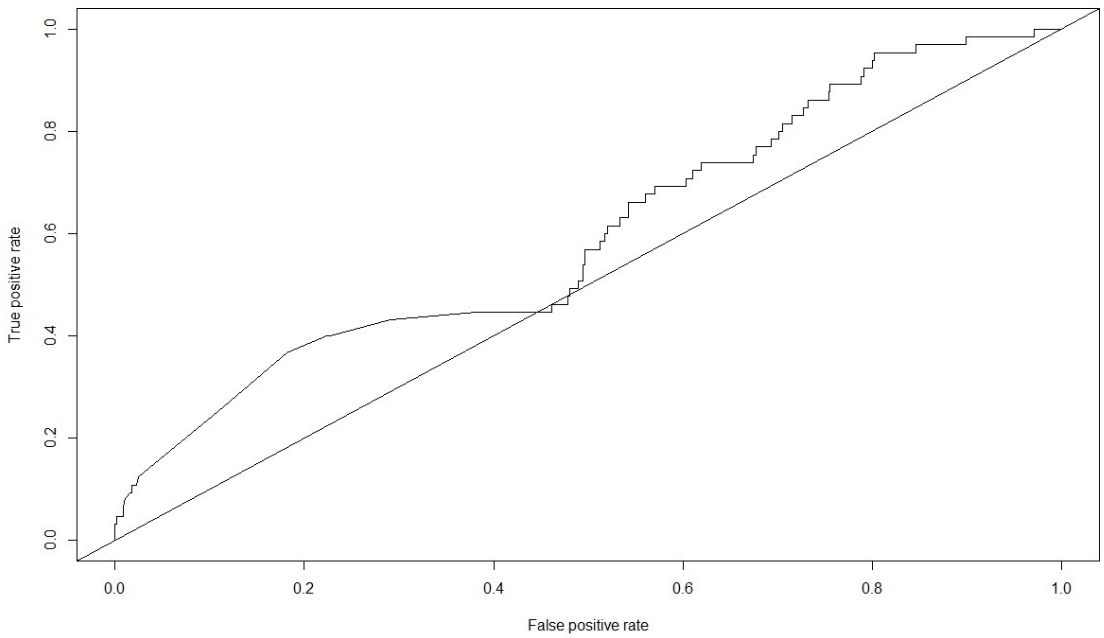

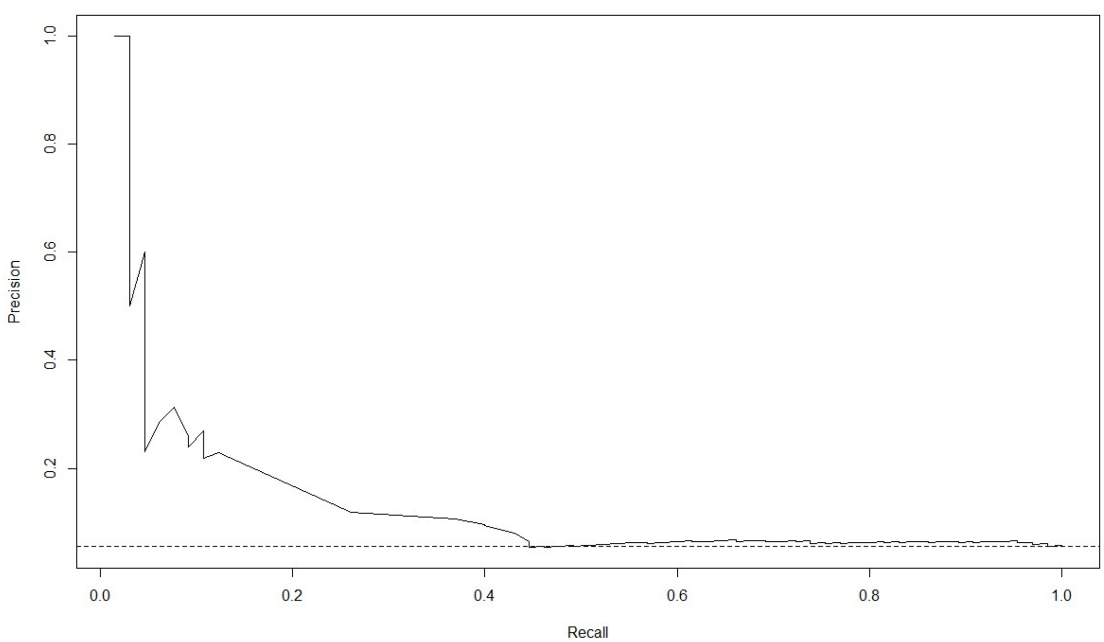

Figure A6.6: ROC and Precision-Recall from Model with Benchmark variables plus Patents and Announcements using Logistic Regression

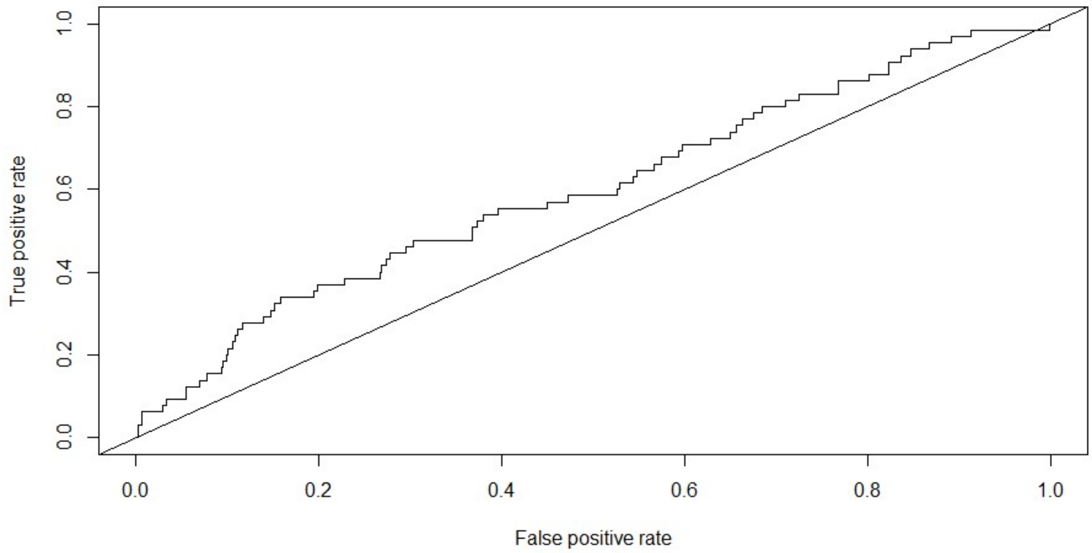

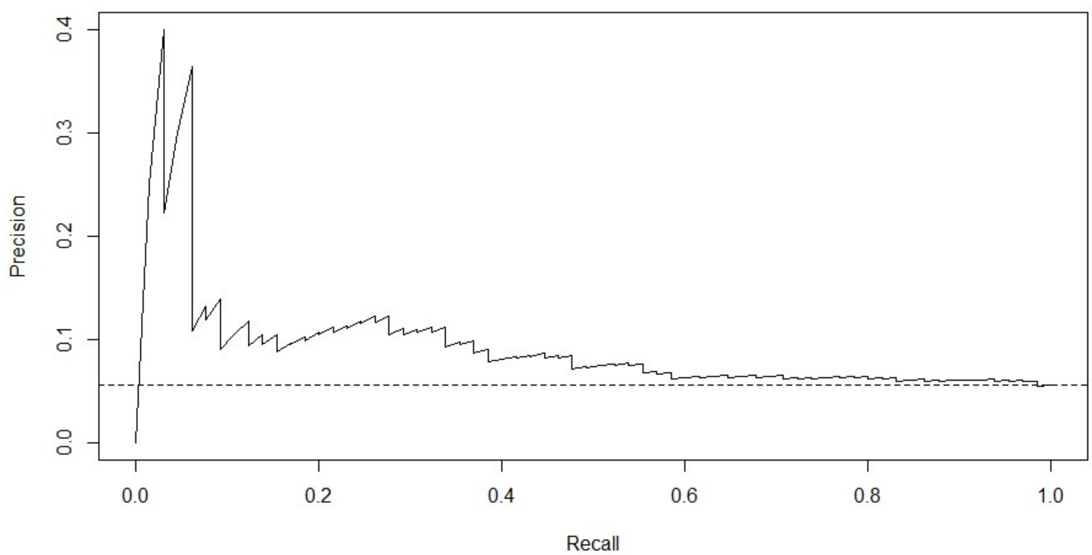

Figure A6.7: ROC and Precision-Recall from Model with Benchmark Variables plus Patents and Announcements using XGBoost

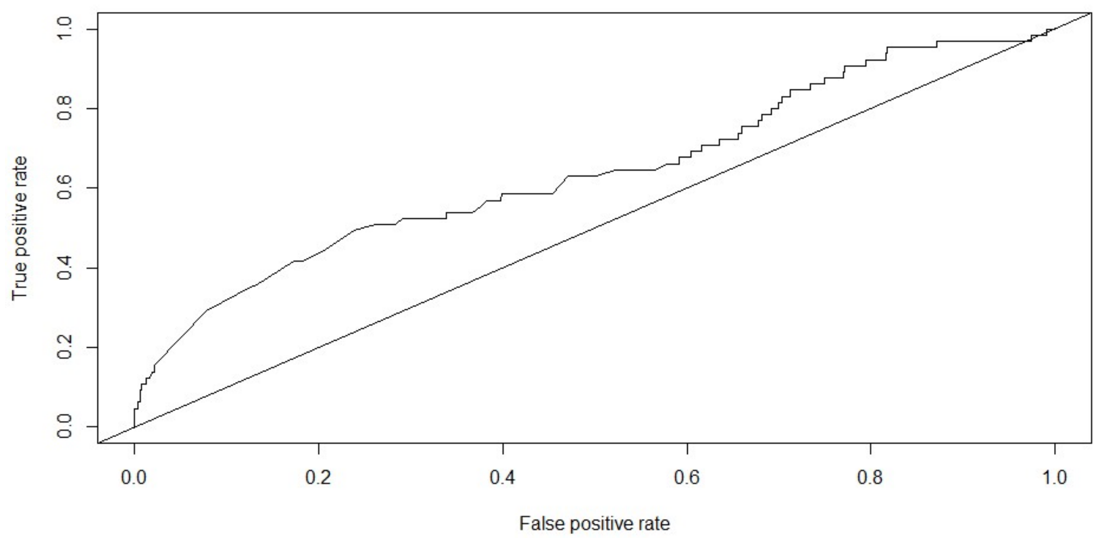

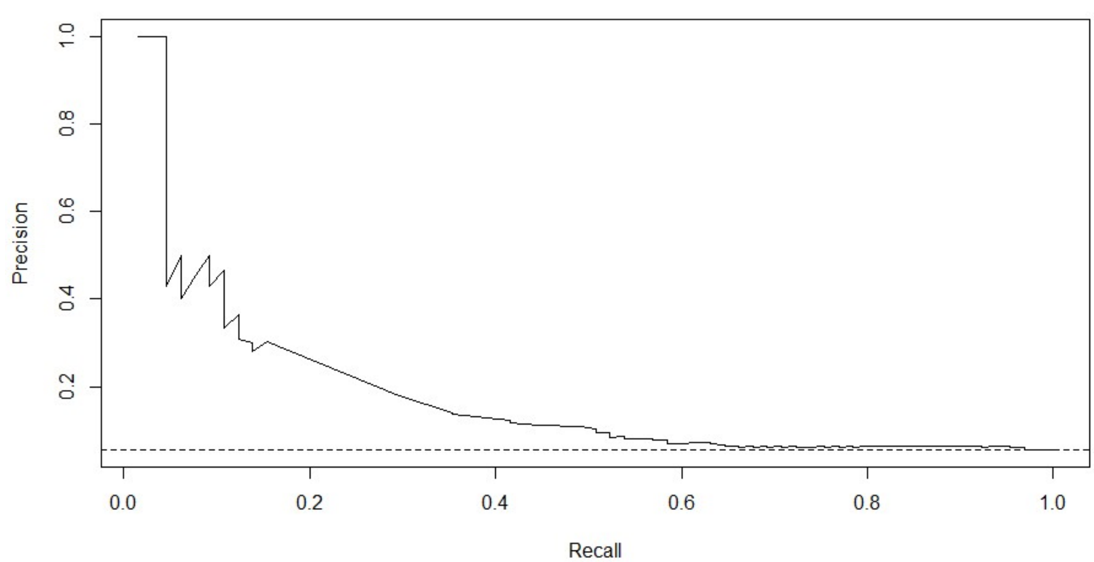

Figure A6.8: ROC and Precision-Recall from Model with All Variables using XGBoost

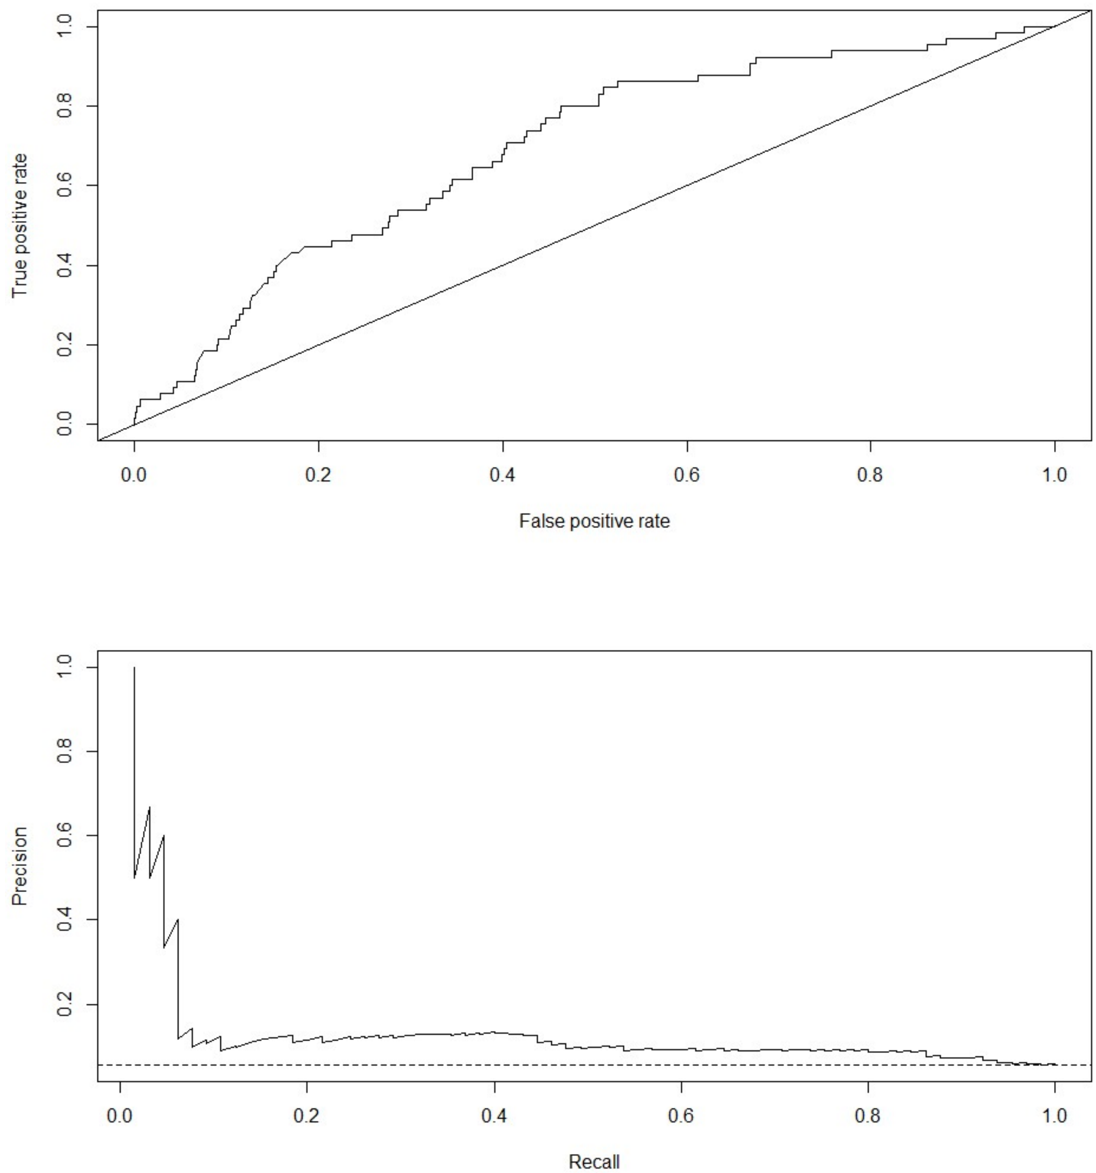

A7 Portfolio Performance Evaluation

Table A7.1: Target Portfolio Performance Evaluation including Momentum and Liquidity

	Dependent variable:			
	khanmyrholt_excess	xgboost1_excess	xgboost2_excess	xgboost3_excess
	(1)	(2)	(3)	(4)
market_excess	0.733***	0.580***	0.865***	0.859**
	p = 0.008	p = 0.007	p = 0.008	p = 0.026
SMB	0.510**	0.496***	0.023	1.016***
	p = 0.035	p = 0.009	p = 0.936	p = 0.004
HML	0.364***	0.244**	0.301*	0.527***
	p = 0.009	p = 0.023	p = 0.066	p = 0.008
PR1YR	−0.344*	−0.106	−0.162	−0.198
	p = 0.063	p = 0.456	p = 0.457	p = 0.447
LIQ	−0.064	0.097	0.355	−0.088
	p = 0.802	p = 0.627	p = 0.246	p = 0.810
Constant	−0.0004	0.008	0.007	0.004
	p = 0.955	p = 0.204	p = 0.452	p = 0.699
Observations	84	84	84	84
R^2	0.298	0.237	0.190	0.232
Adjusted R^2	0.253	0.188	0.138	0.182
Residual Std. Error (df = 78)	0.054	0.041	0.064	0.076
F Statistic (df = 5; 78)	6.637***	4.838***	3.661***	4.703***

Note: *p<0.1; **p<0.05; ***p<0.01

Table A7.2: Target Portfolio Performance Evaluation excluding Actual Targets

	Dependent variable:			
	khannmyrholt_excess	xgboost1_excess	xgboost2_excess	xgboost3_excess
	(1)	(2)	(3)	(4)
market_excess	0.940***	0.585***	0.787***	1.102***
	p = 0.00003	p = 0.0005	p = 0.003	p = 0.0003
SMB	0.596**	0.615***	0.483*	1.228***
	p = 0.015	p = 0.002	p = 0.096	p = 0.0004
HML	0.373***	0.286***	0.454***	0.595***
	p = 0.010	p = 0.010	p = 0.009	p = 0.004
Constant	−0.009	0.005	−0.001	−0.004
	p = 0.168	p = 0.311	p = 0.923	p = 0.696
Observations	84	84	84	84
R^2	0.257	0.228	0.174	0.257
Adjusted R^2	0.229	0.199	0.143	0.230
Residual Std. Error (df = 80)	0.057	0.044	0.069	0.079
F Statistic (df = 3; 80)	9.231***	7.865***	5.633***	9.248***
Note:				*p<0.1; **p<0.05; ***p<0.01

Table A7.3: Target Portfolio Comparison including Momentum and Liquidity

	Dependent variable:		
	long_short_xgboost1	long_short_xgboost2	long_short_xgboost3
	(1)	(2)	(3)
market_excess	−0.154	0.132	0.126
	p = 0.555	p = 0.692	p = 0.728
SMB	−0.014	−0.487	0.506
	p = 0.952	p = 0.104	p = 0.120
HML	−0.120	−0.063	0.163
	p = 0.369	p = 0.712	p = 0.378
PR1YR	0.238	0.182	0.146
	p = 0.186	p = 0.426	p = 0.557
LIQ	0.161	0.419	−0.024
	p = 0.520	p = 0.192	p = 0.946
Constant	0.008	0.007	0.005
	p = 0.285	p = 0.445	p = 0.654
Observations	84	84	84
R^2	0.073	0.052	0.036
Adjusted R^2	0.014	−0.009	−0.026
Residual Std. Error (df = 78)	0.052	0.067	0.072
F Statistic (df = 5; 78)	1.235	0.857	0.586
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A7.4: Target Portfolio Comparison excluding Actual Targets

	Dependent variable:		
	long_short_xgboost1	long_short_xgboost2	long_short_xgboost3
	(1)	(2)	(3)
market_excess	-0.355^{*}	-0.152	0.162
	p = 0.083	p = 0.547	p = 0.551
SMB	0.019	-0.113	0.632^{**}
	p = 0.936	p = 0.695	p = 0.044
HML	-0.087	0.081	0.222
	p = 0.526	p = 0.636	p = 0.230
Constant	0.015^{**}	0.009	0.006
	p = 0.028	p = 0.295	p = 0.515
Observations	84	84	84
R^2	0.048	0.010	0.053
Adjusted R^2	0.012	-0.027	0.017
Residual Std. Error (df = 80)	0.055	0.069	0.074
F Statistic (df = 3; 80)	1.344	0.280	1.491

Note: *p<0.1; **p<0.05; ***p<0.01

Table A7.5: Target Portfolio Comparison excluding Actual Targets, Including Momentum and Liquidity

	Dependent variable:		
	long_short_xgboost1	long_short_xgboost2	long_short_xgboost3
	(1)	(2)	(3)
market_excess	−0.139	0.096	0.179
	p = 0.616	p = 0.781	p = 0.632
SMB	0.002	−0.268	0.671**
	p = 0.995	p = 0.384	p = 0.048
HML	−0.112	0.024	0.228
	p = 0.427	p = 0.890	p = 0.235
PR1YR	0.236	−0.013	0.103
	p = 0.217	p = 0.957	p = 0.689
LIQ	0.206	0.458	−0.048
	p = 0.439	p = 0.168	p = 0.893
Constant	0.009	0.008	0.003
	p = 0.278	p = 0.421	p = 0.754
Observations	84	84	84
R^2	0.069	0.037	0.056
Adjusted R^2	0.009	−0.025	−0.005
Residual Std. Error (df = 78)	0.056	0.069	0.075
F Statistic (df = 5; 78)	1.152	0.603	0.921
Note:			*p<0.1; **p<0.05; ***p<0.01