



Predicting M&A targets using news sentiment and topic detection

Petr Hajek^{a,*}, Roberto Henriques^b

^a Science and Research Centre, Faculty of Economics and Administration, University of Pardubice, Studentska 84, Pardubice, Czech Republic

^b NOVA IMS, Universidade Nova de Lisboa, 1070-312 Lisboa, Portugal

ARTICLE INFO

Keywords:

M&A
Takeover
News
Sentiment
Topic detection
BERT

ABSTRACT

This paper uses news sentiment and topics to discuss the challenges and opportunities of predicting mergers and acquisition (M&A) targets. We explore the effect of investor sentiment on identifying M&As targets and how company-specific news articles can be used as a source of sentiment and topics to obtain richer information on various corporate events. We propose a framework incorporating news sentiment and topics into the M&A target prediction model, utilising state-of-the-art transformer-based sentiment analysis and topic modelling approaches. We evaluate the textual features' predictive power using a real-world dataset of US and UK target and non-target companies from 2020 to 2021, with several experiments conducted to reveal the contribution of sentiment and thematic focus of news to M&A target prediction. A profit-based objective function is proposed to overcome the inherent class imbalance problem in the dataset. Our findings suggest that news-based prediction models outperform traditional statistical and single machine learning methods, indicating the need for more robust and less prone to overfitting ensemble learning methods. Additionally, our study provides evidence for the positive effect of news-based negative sentiment on the likelihood of M&A. Our research has important implications for investors and analysts who seek to identify investment opportunities.

1. Introduction

Mergers and acquisitions (M&As) are a popular and effective way to consolidate and develop a company's business activities (Cumming et al., 2023). M&As are important corporate strategies with the potential for positive and negative outcomes for companies, their stakeholders, and the wider market. On one hand, M&As can increase market share, cost efficiency, synergy, and diversify business operations, as well as providing access to new markets and advanced technologies. On the other hand, they can result in reduced competition, jobs losses, market risks, regulatory hurdles, and high acquisition cost.

Following the global economic downturn triggered by the COVID-19 pandemic, M&A deal activities hit a new record high in 2021, with over 60 thousand deals exceeding \$5 trillion, fuelled by the changed global regulatory landscape (Chiaromonte et al., 2023). The average annual volume of global M&A deals between 2017 and 2022 was nearly \$4 trillion (Statista, 2023). Notably, the financial sector saw the largest increase in M&A deal value in 2021, rising by more than 120 % from the previous year, surpassing IT as the dominant sector in terms of M&A deal value.

M&As are major corporate events leading to considerable stock price

movements. More precisely, it is generally accepted that M&A announcements lead to significant abnormal returns for target firms but not for bidders (Wang and Lahr, 2017; Tunyi, 2021a). Therefore, numerous studies have focused on predicting M&A targets to build a profitable investment strategy (for their overview, see Tunyi, 2021b).

The area of M&A target prediction is attracting considerable interest because of its relevance for investors, managers, and regulators (Zhang et al., 2018; Tunyi, 2021b). From investors' perspective, predicting M&A targets is the foundation of successful investment strategies due to the considerable gains for M&A targets that occur when announcing deals. For managers, timely information on M&As allows managers to take measures to ensure the interests of stakeholders, including implementing an appropriate takeover defence strategy. From the regulators' perspective, M&A target prediction is essential to identify and investigate suspected insider trading.

However, predicting M&As is difficult as target companies tend to underperform in the run-up to a takeover, posing investors with timing risk (Danbolt et al., 2016). Earlier research reveals that potential M&A targets are systematically different from non-target companies concerning their financial characteristics. Specifically, preliminary empirical evidence highlights that M&A targets are comparatively smaller,

* Corresponding author.

E-mail addresses: petr.hajek@upce.cz (P. Hajek), roberto@novaims.unl.pt (R. Henriques).

younger, undervalued, poorly managed, and suffer from a mismatch between resources and growth potential (Cremers et al., 2009). Yet, prior studies recognise that additional investigation into the determinants of M&A targets is justified, given the poor predictive power of existing models (Tunyi, 2019). The low prediction accuracy is attributed to the observation that M&A prediction hypotheses often account for other corporate restructuring-related events, such as divestitures and bankruptcies (Powell and Yawson, 2007). Moreover, this prediction problem is characterised by a relatively high imbalance between the classes of target and non-target companies, which makes it challenging to achieve adequate prediction accuracy for both cases, that is, to correctly identify target companies while reaching a low number of misclassified non-target companies (Ouzounis et al., 2009).

The findings of previous studies are ambiguous regarding the effect of investor sentiment on the identification of M&A targets. On the one hand, earlier studies suggest that acquiring companies tend to be more likely to commit to targets when acquisition costs appear low and investor sentiment is high in the market (Tunyi, 2019). On the other hand, the sentiment of annual reports had reportedly an adverse effect on the likelihood of a company becoming an M&A target, indicating that companies with a higher proportion of negative sentiment are more likely to be M&A targets (Katsafados et al., 2021). This inconsistency may be due to the difference between the sentiment of the market (outsiders) and the sentiment communicated by managers, with the latter specific not only to the internal perspective on the company's performance but also thematically limited mainly to past financial results and future outlooks (Kearney and Liu, 2014). This study turns to company-specific news articles as a source of sentiment and topics to obtain richer information on various corporate events and avoid the low frequency of internal corporate documents. This will allow us to examine the effect of multiple events on detecting M&A targets. This study not only focuses on investor sentiment estimated from the text of news articles, as is common in similar studies in finance (Li et al., 2020) but also examines the effect of the thematic focus of these reports. The analysis of the thematic contents of the news allows us to investigate the impact of the type of information provided along with its incremental value to investors (Jiang et al., 2018; Huang et al., 2018). News articles provide information readily available to investors and facilitate their understanding of the important events covered. Indeed, the combination of sentiment and topic indicators proved beneficial in improving the prediction performance of financial distress models compared to accounting indicators alone (Jiang et al., 2022), which is one of the reasons that motivated this study.

A large volume of company-specific news articles is collected to extract sentiment features and topics that are economically meaningful, and state-of-the-art transformer-based sentiment analysis and topic modelling approaches are exploited in this study. More precisely, we adopt the BERT (Bidirectional Encoder Representations from Transformers) language model to effectively learn contextual word embeddings and fine-tune it to perform two distinct tasks: identify financial sentiment from news and extract a coherent representation of the topics covered. Taken together, this study proposes a framework that incorporates news sentiment and topics into the M&A target prediction model. The predictive power of the textual features is evaluated using a real-world dataset of US and UK target and non-target companies from 2020 to 2021. Several experiments are performed to reveal the contribution of sentiment and thematic focus of news to M&A target prediction. In addition, we propose a profit-based objective function to overcome the inherent class imbalance problem in the dataset. The results provide evidence of the predictive power of the news articles, even when compared to conventional accounting data. In addition, the sentiment and topic features extracted from news articles significantly improve M&A target prediction in terms of accuracy and the abnormal return achieved by investing in the identified companies.

In summary, the contributions of the current study are threefold. First, this study explores the potential of news articles' contents to

predict M&A targets. Although recent research has attempted to use additional, qualitative textual information (letters to shareholders and annual reports) to test its predictive power for identifying successful M&As (Parungao et al., 2022; Katsafados et al., 2021), there are, to the best of our knowledge, no studies that have used such information to detect M&A targets. Hence, this study expands the literature on using text analytics for M&A target prediction. This could be a catalyst for a re-evaluation of current financial forecasting and investment analysis models, encouraging a shift towards big data-driven approaches. Second, we propose BERT-based sentiment and topic textual features to reveal linguistic information provided in the news data. The studies by Katsafados et al. (2021) and Parungao et al. (2022) used only low-frequency data and small samples, respectively, while relying on hand-crafted dictionaries, which severely limits the accuracy and reliability of the qualitative information obtained (Frankel et al., 2022). So, while earlier research has used dictionary-based approaches to predict M&A participants, these approaches are difficult to be adapted to different contexts and sources of sentiment. Therefore, to overcome this limitation, this study explores the combination of advanced BERT-based contextual embedding models to identify financial sentiment and coherent topics in news articles. Thus, the paper shows how cutting-edge technology using transformer-based sentiment analysis and topic modelling can be applied to the prediction of financial markets, and in particular to the identification of M&A targets. This highlights how big data and NLP are becoming increasingly important in financial decision-making, which is in line with current research trends in technological forecasting (Dwivedi et al., 2023). Third, previous studies either failed to address the problem of imbalance between target and non-target companies in the data or attempted to find an equal number of matching non-target companies, thus ignoring potentially useful data on additional non-target companies. To address this issue, this study utilises cost-sensitive ensemble learning to provide the users with a profit-driven prediction model. This study also offers several important implications for stakeholders. In addition to identifying investment opportunities with potential abnormal returns for investors, the proposed model aims to help managers identify appropriate potential targets and improve regulators' ability to assess antitrust compliance by identifying potential targets on time.

The rest of this paper is organised as follows. Section 2 reviews the literature on M&A target prediction. Section 3 describes the dataset used. Section 4 presents the proposed framework for predicting M&A targets using news sentiment and topic detection. Section 5 outlines the experimental setup and shows the results. This is followed by Section 6, which discusses the results and their implications for stakeholders. Section 7 concludes with the future research agenda.

2. Related literature

Recent findings suggest that M&As typically lead to improvements in firm performance due to factors such as the realisation of synergies, risk diversification, improved profitability, increased market power, and the implementation of integrated management strategies (Hossain, 2021). On one hand, M&A activity is influenced by trade credit availability, exchange rates, monetary policy, and cross-border considerations. On the other hand, factors such as political stability, state-owned enterprise ownership, and regulatory conditions are found to be negatively associated with M&A deals.

Regarding the companies' performance around and after M&As, a large body of research literature provides substantial evidence that shareholders of acquiring companies tend to realize minimal to negative returns upon the announcement of a takeover (Renneboog and Vansteenkiste, 2019). Furthermore, when examining the subsequent performance of the merged company's stock price and operating performance over an extended period of time, typically two to three years post-transaction, numerous studies consistently find that bidder shareholders often experience limited or no positive returns following

takeover transactions (Moeller et al., 2004). It is at least partly due to the fact that expected synergies are overstated when the acquisition is announced. Several factors contribute to this overestimation, including behavioural biases of market participants and the presence of biased and optimistic statements in bidder press releases. The most common argument is that the market gradually adjusts to the information associated with takeover transactions (Renneboog and Vansteenkiste, 2019). A further issue is how to ensure post-deal value creation. These include how quickly to integrate, how to manage change, resource and knowledge sharing, communication, employee motivation and attrition, and how to integrate cultures.

In addition to traditional measures such as stock price performance, alternative approaches to assessing post-merger performance have gained prominence. In particular, Li (2013) showed that acquirers can create value in M&A deals by increasing the productivity of the target firm. Another dimension considered in assessing post-acquisition performance is market share. For example, Ghosh (2004) observed a significant increase in the acquiring company's market share three years after the acquisition.

Given the focus of this study, this section further reviews related work on M&A target prediction. However, note that another stream of research has focused on alternative prediction tasks in the area of M&As, particularly on predicting the success of M&As (Lee et al., 2020; Bi and Zhang, 2021; Dang et al., 2022).

As underlined in a recent review study (Tunyi, 2021b), early works in M&A target prediction focused primarily on capturing the common salient characteristics of M&A targets using financial ratios. However, it later became apparent that a successful investment strategy could not be developed based on such a prediction of the M&A target. Over the last ten years, the focus has shifted to the consequences of M&A predictability on stock prices. Notably, while early studies considered it impossible to reach abnormal returns using M&A predictions (Barnes, 1999), more recent studies provide evidence that when the predictions are combined with appropriate screening strategies, abnormal returns can be achieved (Danbolt et al., 2016; Tunyi, 2021b).

The early seminal studies provided theoretically grounded hypotheses guiding the identification of predictive features for effective discrimination between target and non-target companies, including growth resource mismatch and management inefficiency hypotheses (Palepu, 1986). Despite these methodological clues, early empirical studies suggested that although statistical models based on accounting data performed superior to random chance, they did not produce good enough results to lead to excess returns (Zanakis and Zopounidis, 1997; Barnes, 1999; Espahbodi and Espahbodi, 2003). This can be attributed to the fact that statistical models require, to be valid, assumptions to be met on the distribution of the data, which does not correspond to economic reality. A noticeable improvement in prediction performance was therefore achieved only later with the use of artificial intelligence methods.

Doumpos et al. (2004) showed that the UTADIS (UTILités Additives DIScriminantes) multi-criteria decision-making method and neural networks outperform conventional discriminant analysis and logistic regression regarding target detection accuracy. Ouzounis et al. (2009) investigated the performance of discriminant analysis, UTADIS, support vector machines, and neural networks, including the portfolio performance of the predicted targets. Encouraging results were documented for the UTADIS method, and further improvements were achieved by combining the predictions of all the models used, resulting in significantly higher returns than the buy-and-hold strategy. More recent evidence (Parungao et al., 2022) suggests that decision trees can also provide competitive performance in predicting takeover targets.

However, all the previously mentioned takeover prediction models suffer serious weaknesses. Perhaps the most serious disadvantage of these models is that class imbalance in the takeover data was not considered. For example, Tunyi (2019) identified nearly 35,000 observations for 1986–2016, but only 1400 observations involved mergers

and acquisitions. In the same vein, the ratio of target and non-target firms was found to be 1:4 in the study by Ouzounis et al. (2009). Such imbalance ratios led to poor performance for the target class. Most studies have addressed this problem by optimising the cut-off value of the classifier, which resulted in a substantial deterioration in performance for the non-target class (Ouzounis et al., 2009). Another body of studies has tended to overcome this problem by balancing the dataset achieved by selecting matched non-target firms (Pasiouras and Tanna, 2010). However, this leads to inadequate model robustness and limited performance on holdout data due to reduced training data. Therefore, in this study, we take a different approach and, inspired by advances in credit risk prediction (Papouškova and Hajek, 2019; Kou et al., 2021; Yang et al., 2022), instead of maximising model accuracy, we focus on maximising the model profit achieved by cost-sensitive ensemble learning. Our approach assigns more weight to the minority class of target firms and also considers the higher returns achieved by their correct early detection.

Regarding the features used, prior studies utilised financial ratios derived from financial statements. This is because unusual values of financial ratios may indicate inefficient management and undervaluation of the company. In their widely acclaimed work, Palepu (1986) compiled a comprehensive list of hypotheses concerning the probability of M&As and corresponding proxies, represented mainly by financial ratios. Specifically, the inefficient management hypothesis is typically characterised by profitability ratios, and market valuation ratios support the undervaluation hypothesis. In addition, companies with growth-resource mismatches and inefficient financial structures tend to be more susceptible to M&As, and managers usually prefer larger M&A deals. The results of later studies validate this theoretical framework (Ouzounis et al., 2009; Tunyi, 2021b). To further improve the prediction accuracy of M&A targets, later studies have extended the work by Palepu (1986) on M&A determinants. For example, Brar et al. (2009) argue that abnormal returns can be achieved by incorporating financial market data, such as trading volume and stock price momentum.

Furthermore, Akkus et al. (2016) provide evidence of geographical and market proximity's role in bank mergers. The financial features utilised in previous studies are summarised in Table 1, along with the prediction method and performance achieved. In addition to traditional measures of classification model performance, another widely pursued performance measure is associated with the profitability of the portfolio of predicted M&A targets. Danbolt et al. (2016) highlight three main reasons behind systematically underperforming portfolios: (1) a high number of misclassified non-target companies; (2) misclassification of financially distressed companies as target and distressed companies share some financial performance characteristics; and (3) poor timing of picking M&A targets.

While early studies believed financial ratios to be pivotal to the prediction of M&A targets, recent studies on successful M&As suggest that information extracted from company-related textual documents may be equally important (Ma et al., 2017; Katsafados et al., 2021; Parungao et al., 2022). Furthermore, the last decade has seen a remarkable increase in the performance of language models, which has contributed to the improvement of prediction models in many areas of finance, including bankruptcy prediction (Mai et al., 2019), financial fraud (Craja et al., 2020), and stock returns (Hajek, 2018; Yang et al., 2023). Over the last five years, the focus has shifted to contextualised deep learning-based language models representing words as low-dimensional vectors responsive to the context in which they occur. This study draws on these advances in NLP to integrate news-based sentiment and topic detection into the model of M&A target prediction.

3. Research methodology

3.1. Conceptual framework

The generic conceptual framework of the proposed system for

Table 1
Summary of previous studies on M&A target prediction.

Study	Method	Features	Performance
Zanakis and Zopounidis (1997)	LDA, LR	Profitability, managerial performance, solvency	Acc = 0.535
Slowinski et al. (1997)	Rough sets	Profitability, liquidity, leverage	Acc = 0.633
Barnes (1999)	LDA	Firm size, inefficient management, undervaluation, growth-resource mismatch, inefficient financial structure	Acc = 0.750
Espahbodi and Espahbodi (2003)	DT	Firm size, investment opportunity, intangible assets, free cash flow, leverage, dividend policy, growth-resource mismatch	Acc = 0.656
Doumpos et al. (2004)	UTADIS	Profitability, leverage, liquidity, activity, efficiency, growth	Acc = 0.796
Brar et al. (2009)	LR	Firm size, inefficient management, undervaluation, leverage, liquidity, firm age, market barrier	Prec = 0.450
Ouzounis et al. (2009)	LDA, UTADIS, SVM, NN	Firm size, undervaluation, growth-resource mismatch, free cash flow, inefficiency management, price-earnings, dividend policy	Acc = 0.670
Pasiouras and Tanna (2010)	LDA, LR	Bank size, capital strength, profitability, cost efficiency, liquidity, market share	AUC = 0.802
Tunyi (2019)	LR	Firm size, life cycle, inefficient management, undervaluation, growth-resource mismatch, free cash flow, firm age, block holder, price momentum, trading volume	AUC = 0.654
Meghouar and Ibrahim (2021)	LR	Firm size, firm performance, growth-resource mismatch, undervaluation, dividend policy, free cash flow, growth, ownership structure	Acc = 0.894
Parungao et al. (2022)	DT	Firms' letters to shareholders	Acc = 0.670
This study	profit-driven ensemble learning	News-based sentiment and topic detection + financial indicators (firm size, growth-resource mismatch, undervaluation, dividend policy, free cash flow, profitability, firm age, leverage, liquidity, price-earnings, ownership structure)	

Legend: Acc – accuracy, AUC – area under the receiver operating characteristic curve, DT – decision tree, LDA – linear discriminant analysis, LR – logistic regression, and UTADIS – UTilités Additives DIScriminantes.

predicting M&A targets is depicted in Fig. 1. First, data are collected on financial ratios and news articles, and then the system extracts topics covered and news sentiment as linguistic indicators. After the data are partitioned into training and testing data using sequential validation, relevant features were selected. Although we investigated several filter-based methods for feature selection in this step, the traditional chi-square feature selection method worked best in terms of accuracy. This procedure provides a set of relevant indicators for predicting M&A targets. As noted above, previous studies have been restricted to individual statistical or machine learning prediction methods, thus overlooking the advantages of ensemble learning. This is a surprise considering that recent studies have established that ensemble learning methods outperform individual prediction methods on related financial distress prediction tasks (Sun et al., 2020; Abedin et al., 2023). We posit that ensemble methods perform equally well in predicting targets since

they have similarly diverse company profiles as those in financial distress (Tunyi, 2021b). Indeed, combining individual classifiers in ensemble learning makes it suitable for such diverse data while maintaining robustness to overfitting (Chen et al., 2020). To address the problem of imbalanced classes and to provide the investor with a financially effective prediction system, profit-driven ensemble learning is employed by using a profit measure as the objective function, in which both class imbalance and the benefits (costs) of correctly (incorrectly) identified M&A targets are taken into consideration. The last step is to simulate the trading strategy to verify the abnormal return of trading based on the recommendations of the proposed system.

3.2. Financial features

As discussed above, previous literature has considered several categories of financial ratios for predicting M&A targets. Taken together, this body of literature has consistently produced sound theoretical evidence for the role of financial ratios in detecting M&A targets, suggesting that indicators representing inefficient management, growth-resource mismatch, undervaluation, and inefficient financial structure are the most important predictors of M&A targets. Indeed, to detect diverse motives for M&As, the set of financial features should be carefully constructed, and various aspects of a company's financial performance should be considered to achieve a highly accurate prediction system. Therefore, our selection of financial features (Table 2) was shaped by earlier research and, consistent with this line of research, the present study used the following categories of financial features: (1) inefficient management (Barnes, 1999); (2) growth (Doumpos et al., 2004); (3) growth-resource mismatch (Tunyi, 2019); (4) leverage (Slowinski et al., 1997); (5) firm age (Tunyi, 2019); (6) liquidity (Pasiouras and Tanna, 2010); (7) dividend policy (Meghouar and Ibrahim, 2021); (8) free cash flow (Ouzounis et al., 2009); (9) firm size (Meghouar and Ibrahim, 2021); (10) profitability (Doumpos et al., 2004); (11) ownership structure (Meghouar and Ibrahim, 2021); (12) undervaluation (Ouzounis et al., 2009); and (13) price-earnings (Ouzounis et al., 2009).

Financial features are useful for identifying and evaluating potential M&As and gauging a company's suitability as an M&A target. However, news articles can provide insights into market sentiment, rumors, and speculations about potential M&As that cannot be gauged from financial data. News articles may also reveal a company's strategic intentions, such as entering new markets or technologies, which could serve as precursors to M&A activities. To capture this news-based information, linguistic features can be used.

3.3. Linguistic features

Sentiment indicators are among the most commonly investigated linguistic features in the financial literature, typically derived using dictionary-based or machine learning-based methods (Kearney and Liu, 2014). While the dictionary-based methods rely on automatically or manually compiled word lists for each sentiment category, the machine learning-based methods require a large number of labelled documents. However, both approaches struggle with a fairly specific financial context, which makes general dictionaries, such as the General Inquirer or Diction, and machine learning models trained from non-financial corpora highly inaccurate for financial tasks (Loughran and McDonald, 2011; Huang et al., 2023). To address this problem, word lists have been compiled specifically for the financial domain (Loughran and McDonald, 2011), and large corpora of financial texts have been utilised for training the machine learning-based word representations adequately (Huang et al., 2023). In particular, the latter approaches based on neural word embeddings represent the most accurate language models increasingly used in finance, outperforming the domain-specific dictionaries (Azimi and Agrawal, 2021).

While early word embedding representations relied on static word

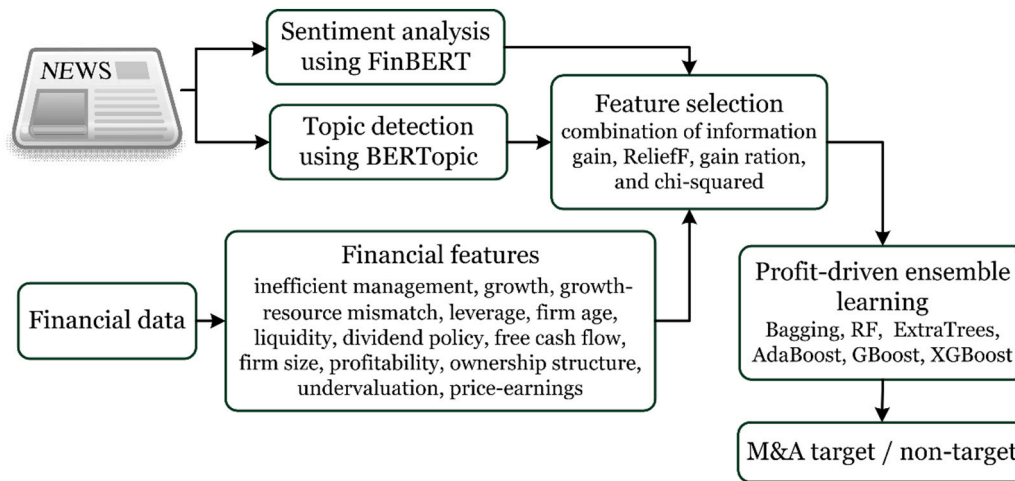


Fig. 1. Conceptual framework of the proposed system for predicting M&A targets.

Table 2
List of financial features used to predict M&A targets.

Category	Feature	Category	Feature
Inefficient management	ROCE	Firm size	TA
Growth	S_{growth}		S
	S_{growth_exp}		EV
Growth-resource mismatch	Mismatch ^a		TE
Leverage	BD/TA	Profitability	ROE
	TD/EV		GPM
	TD		NI
	IC		OM
	Rating		EBITDA
Firm age	Age		EV/EBITDA
Liquidity	CR	Ownership structure	InstHold
	CashR	Undervaluation	PBV
	Cash		Pvar
Dividend policy	RR		EPS_{growth_exp}
	Div_yield		PS
	PR		PEG
Free cash flow	FCF	Price-earnings	P/E
	FCF/S		

Legend: BD – book debt, CashR – cash ratio, CR – current ratio, Div_yield – dividend yield, EBITDA – earnings before interest, taxes, depreciation and amortisation, EPS – earnings per share, EV – enterprise value, exp. – expected (next 5 years), FCF – free cash flow, GPM – gross profit margin, IC – interest coverage, InstHold – shares held by mutual funds, NI – net income after interest, OM – operating margin, P/E – price earnings ratio, PEG – stock price to earnings to EPS growth, Pvar – stock price variation (3-year), PR – payout ratio, PS – stock price to sales, ROCE – return on capital employed, Rating – Standard & Poor’s credit rating (1 for Aaa, 2 for Aa, ..., 9 for C), RR – reinvestment rate, TA – total assets, TD – total debt, TE – total equity.

^a 1 for the following combinations: high S_{growth} -high BD/TA-low CR or low S_{growth} -low BD/TA-high CR (high (low) are defined as above (below) the sample average), 0 otherwise.

vectors, contextualised word embeddings allow the model to use different word vectors according to the context. The BERT model (Devlin et al., 2018) is one of the most popular contextualised models due to its capacity to capture both syntactic and semantic knowledge (Huang et al., 2023). This is done using a masked language model that mitigates the unidirectionality constraint of alternative transformer-based language models, such as ELMo and OpenAI GPT. The task of next sentence prediction represents the other pre-training objective of BERT. To produce the BERT model, a large pre-training corpus of general texts from BookCorpus and Wikipedia (3.3 million words) were used. The pre-trained BERT model can be further fine-tuned to meet the needs of various NLP tasks, including sentiment classification and topic detection. Since the original version of the BERT model was pre-trained on

general texts, due to the different writing styles and specific vocabulary, it is appropriate to perform both BERT pre-training and fine-tuning on financial texts (Huang et al., 2023). As this study focuses on financial news, we used the FinBERT (financial BERT) model,¹ which was pre-trained on the TRC2-financial corpus containing over 29 million words from news articles published by Reuters.

To fine-tune the FinBERT model for sentiment analysis, a sample of ~5000 sentences from financial news and firm press releases included in the Financial PhraseBank were utilised. Business experts annotated this corpus concerning the potential impact on the stock price. For the experimental settings, we followed Araci (2019) and used the following configuration: the maximum sequence length was 64, the number of epochs was 6, and the learning rate was 2e-5 with sequentially unfrozen layers. Then, the fine-tuned FinBERT model was used to label the news headlines collected for target and non-target M&A companies.

To detect topics in the content of the news articles, we employed BERTopic (Grootendorst, 2022), a neural topic model identifying coherent topic representations using the pre-trained BERT model. First, the pre-trained word embeddings are produced. Then, these embeddings are clustered, and finally, topics are detected using the class-based *tf-idf* (term frequency–inverse document frequency) method modelling the word weights in clusters. One advantage of BERTopic over the traditional topic detection methods is that it avoids the problem of vocabulary mismatch by using the semantic representation of terms (Jeon et al., 2023). Furthermore, the topic representation is improved by allowing noisy topic that covers texts unrelated to other topics. We opted for BERTopic also because this model outperformed existing topic modelling methods regarding topic coherence and topic diversity (Grootendorst, 2022). To perform topic detection, the BERTopic library² was employed. Text pre-processing included the removal of stop words (using the Nltk stopwords) and lemmatisation (WordNet lemmatiser). The UMAP (uniform manifold approximation and projection) method (with 5 components and 15 neighbours) was used to cluster the BERT embeddings because it is computationally effective and allows to preserve global and local features of the original high-dimensional dataset (Grootendorst, 2022).

3.4. Feature selection

Feature selection is an essential step in related tasks, including financial distress prediction (Kou et al., 2021; Hajek and Michalak,

¹ <https://huggingface.co/ProsusAI/finbert>.

² <https://github.com/MaartenGr/BERTopic>.

2013; Lu et al., 2023; Ding et al., 2023), because it helps identify features most predictive of the outcome, which in this case is the likelihood of an M&A target. The predictive model's performance can be optimised by selecting the most relevant features. As noted above, M&A target prediction typically involves analysing a large number of features. However, not all of these features may be equally important for predicting M&A targets. By performing feature selection, the goal is to identify the most relevant features highly correlated with the likelihood of M&A targets and discard the irrelevant or redundant ones.

In this study, we adopted a two-stage feature selection approach that has recently proven successful in predicting financial distress (Kou et al., 2021). To this end, in the first stage, feature relevance was evaluated using four filter methods, namely information gain, ReliefF, gain ratio, and chi-squared. Then, the relevance scores for the four feature ranking methods were normalized to intervals [0,1] and averaged, and features were ranked according to the overall relevance score. Thus, redundant and noisy features could be removed, reducing the search space of the wrapper method in the second stage. Specifically, in the second stage, the optimal feature subset was identified from those selected in the first stage based on the highest accuracy when using XGBoost as a classifier. That is, a sequential backward search was used by the wrapper method in the reduced search space to find the optimal feature subset. Details of the algorithms used in both stages can be found in Kou et al. (2021).

3.5. Profit-driven ensemble learning

Earlier research suggests that ensemble learning-based models surpass the prediction performance of single classifiers, including the profit-based performance criteria (Hajek and Abedin, 2020; Kou et al., 2021; Tsai et al., 2021). Indeed, ensemble methods can improve the accuracy of predictions by combining the predictions of multiple individual classifiers by reducing bias that may exist in single classifiers. Ensemble methods are also generally more robust to noise and outliers in the data than single classifiers and can help reduce overfitting. Ensemble methods, particularly those involving averaging predictions, can notably decrease the variance of predictions. Averaging predictions from multiple models trained on different subsets of data mitigates the impact of noise from individual models, thereby contributing to more stable and reliable predictions.

Among the ensemble methods, bagging and boosting are two popular ensemble learning techniques that differ in the training data creation approach. Bagging involves training multiple classifiers independently on random subsets of the training data and combining their predictions using a voting mechanism. Conversely, boosting consists of training multiple classifiers sequentially, where each subsequent classifier is trained to correct the errors made by the previous classifier. Bagging is primarily used to reduce variance in the predictions of individual classifiers by averaging their predictions, and this can help reduce overfitting and improve the generalization performance of the model. Boosting, is primarily used to reduce bias in the predictions of individual classifiers by emphasizing the hard-to-classify examples that were misclassified by previous classifiers, which can reduce the model's bias.

To derive a profit-driven ensemble learning model, the prediction model is trained to maximize the investor's expected profit (abnormal return) rather than simply maximising accuracy, which is inadequate for class imbalance problems such as M&A target prediction. Profit-driven M&A target prediction allowed us to address the class imbalance problem and consider the costs and benefits of different investment decisions. Inspired by the profit-based measures developed for related financial distress prediction tasks (Verbraken et al., 2014; Bahnsen et al., 2015), a confusion matrix with benefits and costs was considered, as shown in Table 3.

Considering the different abnormal returns achieved by different M&A target investments, the example-dependent expected profit $P(t)$ can then be defined as follows:

Table 3
Confusion matrix for predicting M&A targets.

Actual class	Predicted class	
	Target	Non-target
Target	$\pi_0 F_0(t)$ [benefit = b_0]	$\pi_0(1-F_0(t))$ [cost = 0]
Non-target	$\pi_1 F_1(t)$ [cost = c_1]	$\pi_1(1-F_1(t))$ [benefit = 0]
	↓ Trade	↓ No trade

Legend: π_0 and π_1 are the prior probabilities of target and non-target classes, respectively, $F_0(t)$ and $F_1(t)$ are the cumulative density functions for the target and non-target classes given the cut-off value t , b_0 is the average cumulative abnormal return (CAR*) for target companies, and c_1 is the loss (negative CAR) associated with an investment in a non-target company.

*CAR was calculated as the sum of AR over the observation period. In agreement with Yang et al. (2019), we also considered the pre-M&A period and used [-60, 3] window around the announcement date. The calculation of the AR was carried out according to the methodology of MacKinlay (1997) for the estimation period [-250, -60].

$$P(t) = \sum_{i=1}^{N_0} b_0^i - \sum_{j=1}^{N_1} c_1^j, \tag{1}$$

where i and j denote the i -th and j -th training sample predicted as target and non-target, respectively, N_0 and N_1 are the numbers of training samples classified as target and non-target. That is, the M&A target classifiers were trained using the example-dependent profit function in Eq. (1) that maximizes the expected abnormal returns of samples predicted as M&A targets.

4. Data

For our study, consistent with (Yang et al., 2019; Aramyan, 2022), we first identified the group of M&A target companies. Only M&As of US and UK publicly listed companies with a deal value of more than USD 100 million were included to eliminate marginal transactions. Considering transaction announcement dates from September 2020 to November 2021, 182 target companies were identified for which financial indicators were available, as presented in Table 2. Then, the set of non-target peer companies was produced by considering industry classification, financial metrics (e.g., revenue, earnings and market capitalization), and geographic location, which led to the identification of 3648 companies. That is, we obtained the final sample of 3830 companies, of which 182 were labelled as M&A targets. The values of the previous year's financial indicators were collected from the Value Line database (see Table 4 for the mean values of the features).

In addition, CAR was calculated for each sample to determine the example-dependent expected profit function. For the non-target companies included in our sample, the average CAR was negative (-2.99%), while the average CAR for the target companies was positive (10.12%).

To obtain the news corpus for the target and non-target M&A companies, news headlines were collected from the Reuters News service for the period from 30 to 5 days before the announcement date. In total, 298,134 news headlines were obtained from Reuters News during the sample period. From this news corpus we extracted linguistic variables, namely sentiment and topics, using the BERT model (results are presented in the following section). Finally, we merged the M&A financial data with the linguistic variables. The data were split using sequential validation at a ratio of 2:1 into training data (2553 samples from 15/09/2020 to 02/08/2021), followed by testing data (1277 samples from 03/08/2021 to 15/11/2021).

5. Experimental results

The experiments in this study were designed to demonstrate the contribution of news-based information to the accuracy of M&A target identification. To this end, we first present the results of a content

Table 4
Mean values of financial features for target and non-target companies.

Feature	Non-target	Target	Feature	Non-target	Target
ROCE	0.60	-0.05	TA	161.12×10^9	2.08×10^9
S _{growth}	0.02	0.07	S	93.27×10^9	1.46×10^9
S _{growth_exp}	0.15	0.09	EV	320.89×10^9	4.21×10^9
Mismatch	27.8 %	29.7 %	TE	92.59×10^9	1.08×10^9
BD/TA	0.42	0.44	ROE	-0.003	-0.06
TD/EV	0.32	0.33	GPM	39.32	43.02
TD	68.53×10^9	1.00×10^9	NI	10.86×10^9	29.89×10^6
IC	187.30	98.23	OM	-401.02	-30.88
Rating	4.87	5.03	EBITDA	14.08×10^9	328.07×10^6
Age	43.85	39.84	EV/EBITDA	6.09	58.12
CR	2.33	2.02	InstHold	0.54	0.64
CashR	0.20	0.19	PBV	26.43	12.05
Cash	17.56×10^9	216.50×10^6	Pvar	0.30	0.32
RR	4.94	2.61	EPS _{growth_exp}	0.12	0.11
Div_yield	0.016	0.018	PS	50.85	5.56
PR	0.88	1.18	PEG	8.73	9.40
FCF	5.77×10^9	119.33×10^6	P/E	90.51	78.83
FCF/S	-1.63	-0.03			

analysis of news articles related to the studied target and non-target M&A companies. We then compare the performance of state-of-the-art machine learning methods in detecting M&A targets. Finally, we show the economic impact of the proposed prediction model, its robustness, and compare the results with a baseline model based on financial features and alternative approaches to news text analysis.

First, the fine-tuned FinBERT model for sentiment analysis was employed to label the news headlines with positive, neutral or negative sentiment. The overall news-based sentiment was determined as the sum of positive news minus the number of negative news for the company divided by the total number of news over the sample period (Hajek, 2018). The descriptive statistics in Fig. 2 suggest that the news-based sentiment for target companies was more negative (on average 0.055) than non-target companies (0.065), indicating a negative relationship between positive news sentiment and the likelihood of M&A.

In the next step, BERTopic was used to extract 1069 topics. To derive a coherent and interpretable topic model, we only considered topics with a frequency greater than 1000. This resulted in 10 topics, which

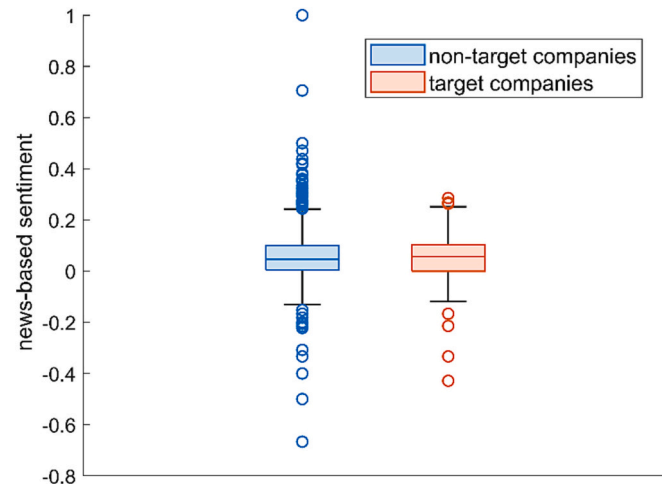


Fig. 2. News-based sentiment for target and non-target M&A companies.

descriptions are shown in Table 5, along with their most representative terms. Note that we gave each topic a name based on the representative terms. The results in Table 5 suggest that the news articles often discussed the overall performance of the target companies in the M&A transactions. In addition, analysts and the press evaluated the investment portfolio, focusing on target companies that exceed earnings expectations. Another topic discussed was the risk associated with holding ETFs that include shares of the M&A target company. Other themes are indicative of M&A targets in the telecommunications and real estate sectors. The focus of the news also reflected a growing emphasis on responsible and sustainable business practices.

The frequencies of the ten topics for the target and non-target companies are shown in Fig. 3. The results show that although the frequencies of topics in the news headlines of the two categories of companies were similar, the news for non-target companies were generally oriented towards business performance. In contrast, the reports of target companies were more likely to deal with financial performance, dividend policy, ESG ratings, and more often indicated corporate consolidation.

The outcome of the previous steps was a set of news-based linguistic features that were merged with the financial features, and feature selection was performed to identify features that were relevant for predicting M&A targets. Table 6 lists the optimal subset of features selected based on the relevance scores of the four filters and wrapper approach using XGBoost. As seen from Table 6, the set of financial features was substantially reduced, with some categories of features completely eliminated as irrelevant (e.g., firm age and dividend policy). For other categories, redundant variables were eliminated. Strikingly, all the linguistic features were retained, indicating their importance for distinction between target and non-target companies.

In further experiments, we used the feature subset and examined the classification performance of profit-driven ensemble learning models against the single classifiers used in previous studies. We used the conventional accuracy measure (Acc) as well as F1 score and AUC (area under the receiver operating characteristic curve) measures to evaluate the classification performance due to its robustness to class imbalance, which is precisely the case with M&A target prediction (Pasiouras and Tanna, 2010; Tunyi, 2019). To enhance the analysis of M&A target identification, we employed the true positive rate (TPR) and true negative rate (TNR). These metrics represent the percentage of correctly identified targets companies and misidentified non-target companies, respectively. In addition, the profit P measure was used to demonstrate the financial impact of making investments based on early detection of M&A deals.

The following methods from prior studies were used as benchmark profit-driven classifiers: (1) LDA (Pasiouras and Tanna, 2010), (2) LR (Meghour and Ibrahim, 2021), (3) DT (Parungao et al., 2022), (4) MLP (Ouzounis et al., 2009), and (5) SVM (Ouzounis et al., 2009).

The results in Table 7 suggest that, consistent with earlier research (Ouzounis et al., 2009), MLP and SVM performed well in AUC measure, indicating a balanced classification performance for both classes of companies. In contrast, the DT method achieved high accuracy, but, as

Table 5
Topics extracted using BERTopic and their representative terms.

Topic	Top 5 terms
Business performance	Global, industry, series, business, periodic
Portfolio scorecard	Scorecard, portfolio, watchdog, corporate, report
EPS beat	Beat, eps, fcf, 4q, guidance
Financial results	Capitalcube, results, ended, 3q, september
ETF risk report	Emptor, caveat, etf, risk, report
5G	5g, korea, huawei, china, mobile
Homebuilding	Jpmorgan, homebuilding, econ, homeserve, builder
Dividend quality	Quality, period, dividend, capitacube, ended
ESG rating	Msci, esg, rating, report, im report
M&A	Alliance, partnership, merger, marketline, acquisition

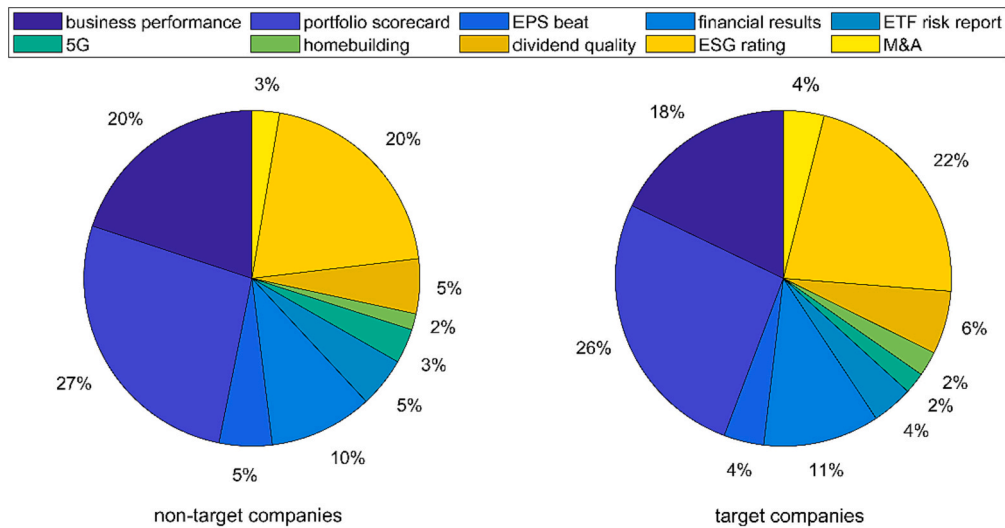


Fig. 3. Relative frequencies of news-based topics for target and non-target M&A companies.

Table 6
List of selected features.

Feature category	Feature	Feature category	Feature
Growth-resource mismatch	Mismatch	News-based sentiment	Overall sentiment
Leverage	TD	News-based topics	Business performance
Liquidity	Rating		Portfolio scorecard
Free cash flow	Cash		EPS beat
Firm size	FCF		Financial results
	S		ETF risk report
	EV		5G
Profitability	TA		Homebuilding
	NI		Dividend quality
	EBITDA		ESG rating
Ownership structure	InsttHold		M&A
Undervaluation	Pvar		
	PS		

Table 7
Classification performance of compared methods.

Method	Acc	AUC	F1 score	TPR	FPR	Profit P [%]
LDA	0.690	0.792	0.791	0.769	0.313	0.747
LR	0.889	0.833	0.920	0.769	0.108	-0.078
DT	0.934	0.755	0.946	0.564	0.055	3.823
MLP	0.895	0.875	0.924	0.821	0.103	0.867
SVM	0.873	0.860	0.911	0.846	0.127	0.229
ExtraTree	0.933	0.717	0.945	0.487	0.053	0.180
Bagging	0.934	0.929	0.946	0.513	0.053	4.355
RF	0.939	0.949	0.950	0.615	0.051	2.200
AdaBoost	0.923	0.921	0.942	0.846	0.074	0.286
GBoost	0.922	0.928	0.940	0.667	0.070	0.912
XGBoost	0.932	0.955	0.947	0.692	0.060	4.603

Note: the classification methods were trained with the parameter setting given in the Appendix 1, and the best results are in bold.

indicated by the relatively lower AUC value, only on the class of non-target companies.

Overall, the results of the single classifiers suggest that these models are not comprehensive enough to detect the diverse profiles of target and non-target companies (Table 7). Although SVM and MLP showed promising results by accurately identifying a high percentage of target companies (TPR), they did so at the expense of a notable increase in error rates for the majority of non-target companies (FPR). Consequently, this resulted in a significant decrease in profits. This limitation

was addressed mainly by using the profit-driven ensemble methods, namely bagging-based methods (Bagging, RF (random forest), and ExtraTrees (extremely randomized trees)) and boosting-based methods (AdaBoost (adaptive boosting), GBoost (gradient boosting), and XGBoost (extreme gradient boosting)).

The results in Table 7 for the bagging-based methods suggest that the additional randomization of ExtraTree was not effective and that this method was biased towards the majority class of non-target companies. In contrast, Bagging and RF performed well even in terms of AUC and F1 score, indicating the effectiveness of bootstrapping randomization. Finally, the boosting-based methods also worked good in terms of Acc and AUC. Notably, XGBoost excelled regarding the issue of class imbalanced data, which can be attributed to several built-in techniques to handle this problem (Hajek et al., 2023). In addition, the XGBoost method produced the highest economic benefits, indicating a good ratio of correctly classified target companies and incorrectly classified non-target companies. We used this method in subsequent experiments for this reason.

To demonstrate the marginal economic effect of the proposed model, the trend of profit P achieved by XGBoost as a function of the cut-off value t is presented in Fig. 4. It is worth noting that at $t = 0.0$, all 314 deals were executed with $P = 0.769\%$, whereas at $t = 1.0$, the number of deals was zero with $P = 0.000\%$. From Fig. 4 it can be seen that shifting

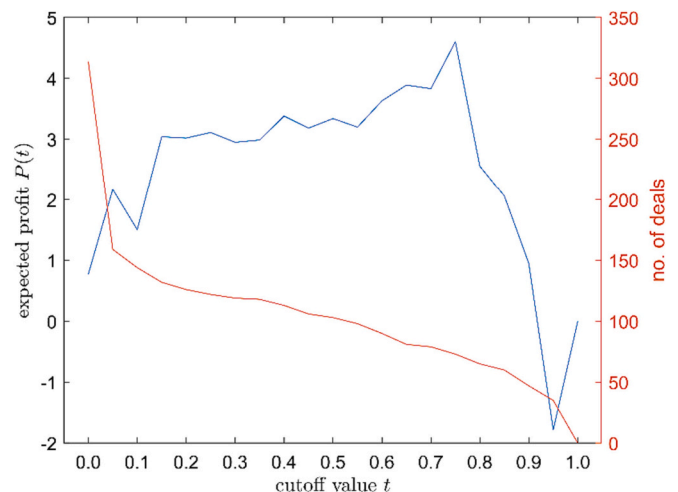


Fig. 4. Effect of shifting the cut-off value on the expected profit.

the cut-off value increases the percentage of correctly classified target companies while decreasing the percentage of correctly classified non-target companies. This is justified if the benefit b_0 is greater than the cost c_1 (see Table 3). However, shifting the cut-off value is only beneficial to a certain extent; at a value greater than 0.7, the number of misclassified non-target companies was already too high, reducing the overall profit. Therefore, we selected the cut-off value $t = 0.7$ to achieve maximum profit of $P = 4.603\%$ from 73 deals.

To demonstrate the effect of financial and linguistic features, we further considered three scenarios for the XGBoost method: (1) all selected features (financial + linguistic), (2) financial features only, and (3) linguistic features only. It is apparent from Fig. 5 that the linguistic features alone are highly accurate, in contrast to the financial features, which are less effective on all three performance measures. Fig. 5 is quite revealing in several ways. First, the results show the limited accuracy of detecting target companies using traditional financial indicators. Second, XGBoost's performance was improved by combining the financial and linguistic features, but surprisingly only slightly compared with the linguistic-only features.

To check for the robustness of the results, we further considered different split ratios of training and testing data. Again, sequential validation was applied, but this time for five different ratios ranging from 1:1 to 9:1, corresponding to different training and testing periods. More precisely, the end of the training periods ranged from 17/06/2021 (for the 1:1 ratio) to 26/10/2021 (for the 9:1 ratio). Generally speaking, the results in Fig. 6 indicate consistent classification performance, with the differences in profit attributable to the different CARs of the target companies over the different test periods.

For comparison purposes, we further used a dictionary-based approach to sentiment analysis, specifically that of Loughran and McDonald (2011), which are most commonly used word lists in the financial literature (Hajek and Henriques, 2017; Li et al., 2020; Katsafados et al., 2021). These finance-specific dictionaries have several advantages over general dictionaries. First, they reduce the noise caused in general dictionaries by misclassifying many negative words such as "liabilities" or "taxes". Furthermore, a negative meaning is not assigned to certain industry-specific words. In addition to commonly used positive and negative sentiment, this set of dictionaries includes uncertain, litigious and modal word lists, thus allowing for a richer representation of sentiment in the text of financial disclosures. Future abnormal returns have been shown to be associated with the higher frequency of all these word categories in annual reports (Loughran and McDonald, 2011). Overall, these dictionaries have shown greater predictive power in the stock market than general dictionaries, including Vader and SentiWordNet (Li et al., 2020). In agreement with Huang et al. (2023), a sentence was identified as negative if it contained at least one word from the negative dictionary, as positive if it contained a word from the

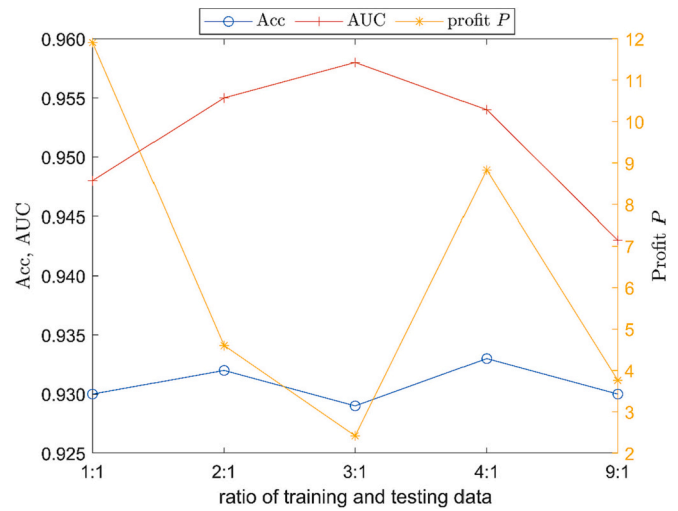


Fig. 6. Effect of split ratios on classification performance.

positive dictionary (and no negative words), otherwise the sentence was identified as neutral.

The results in Fig. 7 reveal that although higher classification accuracy can be obtained using the dictionary-based approach compared with the proposed BERT-based approach, this is achieved at the cost of less balanced performance on both classes, resulting in a lower detection capability of the target companies. In addition, the trading strategy based on the dictionary-based approach was less profitable than that for the proposed BERT-based model.

6. Discussion

Our study provides insight into the capacity to leverage the content of news articles to predict M&A targets. Unlike previous studies reliant on hand-crafted dictionaries (Katsafados et al., 2021; Parungao et al., 2022), we have adopted state-of-the-art sentiment and topic analysis methods based on large language models utilising transformer neural networks to explore the content of the news. We clearly demonstrated the contributions of the proposed approach in terms of both accuracy in predicting M&A targets and economic benefits. In contrast to the

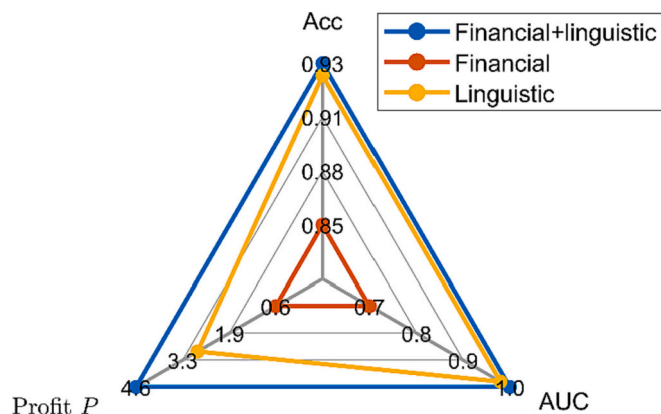


Fig. 5. Classification performance of XGBoost for financial features, linguistic features and the combination of financial and linguistic features.

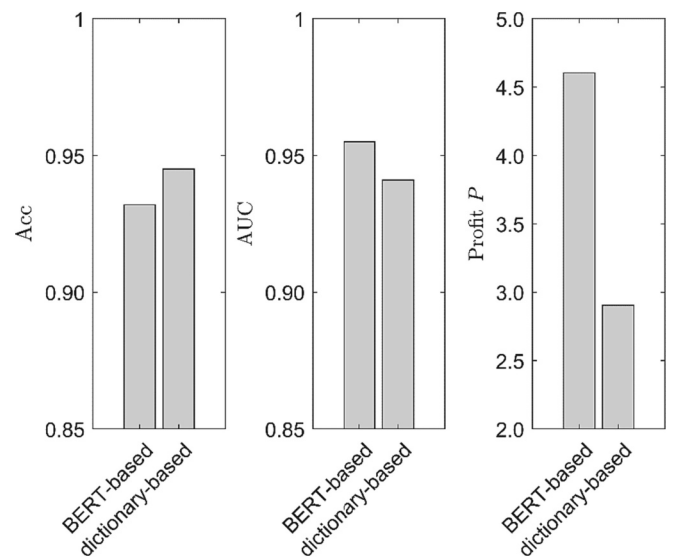


Fig. 7. Comparison of classification performance between the proposed BERT-based linguistic features and the dictionary-based approach to sentiment analysis.

dictionary-based approaches used in previous research, the use of the BERT-based model allowed us to adapt to the context of M&A transactions and the sources of sentiment.

We found that the linguistic features extracted from the news are essential for accurately predicting M&A targets, indicating that news articles, compared with not necessarily up-to-date financial indicators, carry timely information indicating M&As. Consistent with [Katsafados et al. \(2021\)](#), who find a positive relationship between the negative sentiment from the annual reports and the likelihood of a takeover, our results provide evidence for a positive effect of news-based negative sentiment and the likelihood of M&A.

It should also be underlined that traditional statistical and single machine learning methods failed to provide adequate prediction accuracy, which led us to employ more robust and less prone to overfitting ensemble learning methods. Thus, an M&A target prediction model was achieved in combination with the selection of relevant features. We demonstrated that the proposed cost-sensitive ensemble model can provide the investors with an accurate forecasting model. Unlike the machine learning models used in previous studies ([Ouzounis et al., 2009](#)), our model is robust enough to handle the uneven distribution of target and non-target companies in the data, making it a better fit with economic reality. Equally noteworthy is that the news-based prediction model also performed well in terms of the proposed profit measure, implying the financial significance of the improved classification performance of the model for identifying investment opportunities.

Overall, our findings based on feature selection and ensemble learning methods provide strong support for the key role of sentiment and topic analysis of news for the prediction of M&A targets, providing additional empirical support for the prospect theory ([Kahneman and Tversky, 2013](#)). There are several possible explanations for this result. First, news articles can provide up-to-date information on potential M&A targets, allowing investors and analysts to stay informed about market trends and identify potential M&A targets on time. While financial indicators provide historical data, linguistic features from news can provide forward-looking information. Second, news articles often contain information on market sentiment, such as positive or negative sentiment towards a company or industry, including financial performance, market trends, and strategic directions. Third, news articles can provide information on industry trends, such as emerging technologies and regulatory changes, which can help investors and analysts identify potential M&A targets that align with industry trends and have strong growth potential.

It is worth noting that both the COVID-19 pandemic and the conflict between Russia and Ukraine have introduced significant geopolitical uncertainty that have prompted a reassessment of M&A strategies across various industries as managers focused on liquidity rather than long-term growth strategy ([Bauer et al., 2022](#)). We believe that our news-based model can play an important role in identifying emerging trends in market sentiment and geopolitical events and use them as indicators of M&A activities, especially in dynamic and uncertain scenarios like the COVID-19 pandemic and the Russian-Ukraine war.

Although large language models (LLMs), such as BERT and GPT-4, provide significant benefits in analysing large amounts of data, their application carries certain risks and threats. First, LLMs may misinterpret intricate financial and legal texts, which could have severe consequences in the M&A process. Additionally, relying too heavily on LLMs for decision-making in M&As can be hazardous, as these models may not include all the qualitative factors and human judgments that are essential in these decisions. LLMs also lack human ethical considerations, which are frequently viral in business decisions, especially in complicated scenarios like M&As. Finally, the utilisation of AI to process M&A data must conform with data protection regulations to avoid potential legal and reputational hazards caused by any violation. Therefore, further research regarding the role of LLMs in the M&A process would be interesting.

7. Conclusion

In this study, we used BERT-based sentiment and topic features to identify M&A targets among US and UK publicly listed companies. One of the most significant findings from this study is that the news-based linguistic features greatly outperform traditional financial indicators in predictive power. To overcome the intrinsic problem of class imbalance in M&A target prediction, we proposed a profit-based measure utilising different prior probabilities and different abnormal returns obtained for target and non-target companies. It was also shown that profit-driven ensemble learning methods are superior to previous single classification methods.

The findings of this study have a number of important implications. Firstly, the news-based information can be used to support strategic decision-making. For instance, executives can utilise this information to assess market sentiment and potential reaction to the deal when evaluating a possible M&A. Secondly, it is crucial to implement policies for regular monitoring of news sentiment and relevant topics as part of the risk management strategy. This will facilitate the timely detection of potential risks or opportunities in the market. Thirdly, investors can use the news-based information to identify undervalued companies that could potentially be M&A targets. Fourthly, regulators could monitor unusual positive or negative news surrounding a company as part of their insider trading or market manipulation surveillance.

Several important limitations of the current study need to be considered. Most of these limitations arise from the study's focus on news articles as indicators of M&As. First, the Reuters news service limits the period to obtain these news. Therefore, further research should consider additional sources of news articles to extend the observation period, which would also allow for more M&A coverage. Second, news articles can be supplemented with other textual data sources, such as quarterly/annual reports or earnings conference calls to obtain more comprehensive linguistic information, including that from internal stakeholders. It would also be interesting to explore news-based sentiment and topics using alternative large language models due to their recent rapid development.

CRedit authorship contribution statement

Petr Hajek: Conceptualization, Data curation, Funding acquisition, Investigation, Methodology, Software, Supervision, Validation, Writing – original draft. **Roberto Henriques:** Formal analysis, Methodology, Visualization, Writing – review & editing.

Declaration of competing interest

None

Data availability

Data will be made available on request.

Acknowledgment

The authors acknowledge the financial support of the Czech Science Foundation [Grant No. 22-22586S].

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.techfore.2024.123270>.

References

- Abedin, M.Z., Guotai, C., Hajek, P., Zhang, T., 2023. Combining weighted SMOTE with ensemble learning for the class-imbalanced prediction of small business credit risk. *Complex Intell. Syst.* 9 (4), 3559–3579.
- Akkus, O., Cookson, J.A., Hortacsu, A., 2016. The determinants of bank mergers: a revealed preference analysis. *Manag. Sci.* 62 (8), 2241–2258.
- Araci, D., 2019. Finbert: financial sentiment analysis with pre-trained language models. In: *arXiv Preprint*. arXiv:1908.10063.
- Aramyan, H., 2022. Predicting M&A targets using ML: unlocking the potential of NLP based variables. Available at: Refinitiv. <https://developers.refinitiv.com/en/article-catalog/article/predicting-MnA-targets-using-ML-Unlocking-the-potential-of-NLP-variables>.
- Azimi, M., Agrawal, A., 2021. Is positive sentiment in corporate annual reports informative? Evidence from deep learning. *Rev. Asset Pricing Stud.* 11 (4), 762–805.
- Bahnsen, A.C., Aouada, D., Ottersten, B., 2015. Example-dependent cost-sensitive decision trees. *Expert Syst. Appl.* 42 (19), 6609–6619.
- Barnes, P., 1999. Predicting UK takeover targets: some methodological issues and an empirical study. *Rev. Quant. Financ. Acc.* 12 (3), 283–302.
- Bauer, F., Friesl, M., Dao, M.A., 2022. Run or hide: changes in acquisition behaviour during the COVID-19 pandemic. *J. Strateg. Manag.* 15 (1), 38–53.
- Bi, W., Zhang, Q., 2021. Forecasting mergers and acquisitions failure based on partial-sigmoid neural network and feature selection. *PLoS One* 16 (11), e0259575.
- Brar, G., Giamouridis, D., Liodakis, M., 2009. Predicting European takeover targets. *Eur. Financ. Manag.* 15 (2), 430–450.
- Chen, Z., Chen, W., Shi, Y., 2020. Ensemble learning with label proportions for bankruptcy prediction. *Expert Syst. Appl.* 146, 113155.
- Chiaramonte, L., Dreassi, A., Piserà, S., Khan, A., 2023. Mergers and acquisitions in the financial industry: a bibliometric review and future research directions. *Res. Int. Bus. Financ.* 64, 101837.
- Craja, P., Kim, A., Lessmann, S., 2020. Deep learning for detecting financial statement fraud. *Decis. Support. Syst.* 139, 113421.
- Cremers, K.M., Nair, V.B., John, K., 2009. Takeovers and the cross-section of returns. *Rev. Financ. Stud.* 22 (4), 1409–1445.
- Cumming, D., Jindal, V., Kumar, S., Pandey, N., 2023. Mergers and acquisitions research in finance and accounting: past, present, and future. *Eur. Financ. Manag.* 29 (5), 1464–1504.
- Danbolt, J., Siganos, A., Tunyi, A., 2016. Abnormal returns from takeover prediction modelling: challenges and suggested investment strategies. *J. Bus. Financ. Acc.* 43 (1–2), 66–97.
- Dang, M., Henry, D., Thai, H.A., Vo, X.V., Mazur, M., 2022. Does policy uncertainty predict the death of M&A deals? *Financ. Res. Lett.* 46, 102489.
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K., 2018. Bert: pre-training of deep bidirectional transformers for language understanding. In: *arXiv Preprint*. arXiv:1810.04805.
- Ding, S., Cui, T., Bellotti, A.G., Abedin, M.Z., Lucey, B., 2023. The role of feature importance in predicting corporate financial distress in pre and post COVID periods: evidence from China. *Int. Rev. Financ. Anal.* 90, 102851.
- Doumpos, M., Kosmidou, K., Pasiouras, F., 2004. Prediction of acquisition targets in the UK: a multi-criteria approach. *Oper. Res.* 4 (2), 191–211.
- Dwivedi, Y.K., Sharma, A., Rana, N.P., Giannakis, M., Goel, P., Dutot, V., 2023. Evolution of artificial intelligence research in technological forecasting and social change: research topics, trends, and future directions. *Technol. Forecast. Soc. Chang.* 192, 122579.
- Espahbodi, H., Espahbodi, P., 2003. Binary choice models and corporate takeover. *J. Bank. Financ.* 27 (4), 549–574.
- Frankel, R., Jennings, J., Lee, J., 2022. Disclosure sentiment: machine learning vs. dictionary methods. *Manag. Sci.* 68 (7), 5514–5532.
- Ghosh, A., 2004. Increasing market share as a rationale for corporate acquisitions. *J. Bus. Financ. Acc.* 31 (1–2), 209–247.
- Grootendorst, M., 2022. BERTopic: neural topic modeling with a class-based TF-IDF procedure. In: *arXiv Preprint*. arXiv:2203.05794.
- Hajek, P., 2018. Combining bag-of-words and sentiment features of annual reports to predict abnormal stock returns. *Neural Comput. & Applic.* 29 (7), 343–358.
- Hajek, P., Abedin, M.Z., 2020. A profit function-maximizing inventory backorder prediction system using big data analytics. *IEEE Access* 8, 58982–58994.
- Hajek, P., Henriques, R., 2017. Mining corporate annual reports for intelligent detection of financial statement fraud—a comparative study of machine learning methods. *Knowl.-Based Syst.* 128, 139–152.
- Hajek, P., Michalak, K., 2013. Feature selection in corporate credit rating prediction. *Knowl.-Based Syst.* 51, 72–84.
- Hajek, P., Abedin, M.Z., Sivarajah, U., 2023. Fraud detection in mobile payment systems using an XGBoost-based framework. *Inf. Syst. Front.* 25, 1985–2003.
- Hossain, M.S., 2021. Merger & Acquisitions (M&As) as an important strategic vehicle in business: thematic areas, research avenues & possible suggestions. *J. Econ. Bus.* 116, 106004.
- Huang, A.H., Lehavy, R., Zang, A.Y., Zheng, R., 2018. Analyst information discovery and interpretation roles: a topic modeling approach. *Manag. Sci.* 64 (6), 2833–2855.
- Huang, A.H., Wang, H., Yang, Y., 2023. FinBERT: a large language model for extracting information from financial text. *Contemp. Account. Res.* 40 (2), 806–841.
- Jeon, E., Yoon, N., Sohn, S.Y., 2023. Exploring new digital therapeutics technologies for psychiatric disorders using BERTopic and PatentSBERTa. *Technol. Forecast. Soc. Chang.* 186, 122130.
- Jiang, C., Wang, Z., Wang, R., Ding, Y., 2018. Loan default prediction by combining soft information extracted from descriptive text in online peer-to-peer lending. *Ann. Oper. Res.* 266 (1), 511–529.
- Jiang, C., Lyu, X., Yuan, Y., Wang, Z., Ding, Y., 2022. Mining semantic features in current reports for financial distress prediction: empirical evidence from unlisted public firms in China. *Int. J. Forecast.* 38 (3), 1086–1099.
- Kahneman, D., Tversky, A., 2013. Prospect theory: an analysis of decision under risk. In: *Handbook of the Fundamentals of Financial Decision Making, Part I*, pp. 99–127.
- Katsafados, A.G., Androutsopoulos, I., Chalkidis, I., Fergadiotis, E., Leledakis, G.N., Pyrgiotakis, E.G., 2021. Using textual analysis to identify merger participants: evidence from the US banking industry. *Financ. Res. Lett.* 42, 101949.
- Kearney, C., Liu, S., 2014. Textual sentiment in finance: a survey of methods and models. *Int. Rev. Financ. Anal.* 33, 171–185.
- Kou, G., Xu, Y., Peng, Y., Shen, F., Chen, Y., Chang, K., Kou, S., 2021. Bankruptcy prediction for SMEs using transactional data and two-stage multiobjective feature selection. *Decis. Support. Syst.* 140, 113429.
- Lee, K., Joo, S., Baik, H., Han, S., In, J., 2020. Unbalanced data, type II error, and nonlinearity in predicting M&A failure. *J. Bus. Res.* 109, 271–287.
- Li, X., 2013. Productivity, restructuring, and the gains from takeovers. *J. Financ. Econ.* 109 (1), 250–271.
- Li, X., Wu, P., Wang, W., 2020. Incorporating stock prices and news sentiments for stock market prediction: a case of Hong Kong. *Inf. Process. Manag.* 57 (5), 102212.
- Loughran, T., McDonald, B., 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *J. Financ.* 66 (1), 35–65.
- Lu, Y., Yang, L., Shi, B., Li, J., Abedin, M.Z., 2023. A novel framework of credit risk feature selection for SMEs during industry 4.0. *Ann. Oper. Res.* 1–28. <https://doi.org/10.1007/s10479-022-04849-3>.
- Ma, T., Zhang, Y., Huang, L., Shang, L., Wang, K., Yu, H., 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. *Technol. Forecast. Soc. Chang.* 116, 162–180.
- MacKinlay, A.C., 1997. Event studies in economics and finance. *J. Econ. Lit.* 35 (1), 13–39.
- Mai, F., Tian, S., Lee, C., Ma, L., 2019. Deep learning models for bankruptcy prediction using textual disclosures. *Eur. J. Oper. Res.* 274 (2), 743–758.
- Meghouar, H., Ibrahim, M., 2021. Financial characteristics of takeover targets: a French empirical evidence. *EuroMed J. Bus.* 16 (1), 69–85.
- Moeller, S.B., Schlingemann, F.P., Stulz, R.M., 2004. Firm size and the gains from acquisitions. *J. Financ. Econ.* 73 (2), 201–228.
- Ouzounis, G., Gaganis, C., Zopounidis, C., 2009. Prediction of acquisitions and portfolio returns. *Int. J. Bank. Account. Financ.* 1 (4), 381–406.
- Palepu, K.G., 1986. Predicting takeover targets: a methodological and empirical analysis. *J. Account. Econ.* 8 (1), 3–35.
- Papoukova, M., Hajek, P., 2019. Two-stage consumer credit risk modelling using heterogeneous ensemble learning. *Decis. Support. Syst.* 118, 33–45.
- Parungao, M.E., Galido, A., Suazo, M.L., Parungao, L.A., 2022. Exploring qualitative data as predictors for M&A: empirical analysis of target firms' letters to shareholders. *Cogent Bus. Manag.* 9 (1), 2084970.
- Pasiouras, F., Tanna, S., 2010. The prediction of bank acquisition targets with discriminant and logit analyses: methodological issues and empirical evidence. *Res. Int. Bus. Financ.* 24 (1), 39–61.
- Powell, R., Yawson, A., 2007. Are corporate restructuring events driven by common factors? Implications for takeover prediction. *J. Bus. Financ. Acc.* 34 (7–8), 1169–1192.
- Renneboog, L., Vansteenkiste, C., 2019. Failure and success in mergers and acquisitions. *Finance* 58, 650–699.
- Slowinski, R., Zopounidis, C., Dimitras, A.I., 1997. Prediction of company acquisition in Greece by means of the rough set approach. *Eur. J. Oper. Res.* 100 (1), 1–15.
- Statista, 2023. Mergers and acquisitions worldwide - statistics & facts. Available at: <https://www.statista.com/topics/1146/mergers-and-acquisitions/>.
- Sun, J., Li, H., Fujita, H., Fu, B., Ai, W., 2020. Class-imbalanced dynamic financial distress prediction based on Adaboost-SVM ensemble combined with SMOTE and time weighting. *Inf. Fusion* 54, 128–144.
- Tsai, C.F., Sue, K.L., Hu, Y.H., Chiu, A., 2021. Combining feature selection, instance selection, and ensemble classification techniques for improved financial distress prediction. *J. Bus. Res.* 130, 200–209.
- Tunyi, A., 2019. Firm size, market conditions and takeover likelihood. *Rev. Acc. Financ.* 18 (3), 483–507.
- Tunyi, A.A., 2021a. Revisiting acquirer returns: evidence from unanticipated deals. *Finance* 66, 101789.
- Tunyi, A., 2021b. Fifty years of research on takeover target prediction: a historical perspective. *Qual. Res. Financ. Mark.* 13 (4), 482–502.
- Verbraken, T., Bravo, C., Weber, R., Baesens, B., 2014. Development and application of consumer credit scoring models using profit-based classification measures. *Eur. J. Oper. Res.* 238 (2), 505–513.
- Wang, Y., Lahr, H., 2017. Takeover law to protect shareholders: increasing efficiency or merely redistributing gains? *Finance* 43, 288–315.
- Yang, B., Sun, J., Guo, J.M., Fu, J., 2019. Can financial media sentiment predict merger and acquisition performance? *Econ. Model.* 80, 121–129.
- Yang, F., Qiao, Y., Abedin, M.Z., Huang, C., 2022. Privacy-preserved credit data sharing integrating blockchain and federated learning for industrial 4.0. *IEEE Trans. Industr. Inform.* 18 (12), 8755–8764.

- Yang, C., Abedin, M.Z., Zhang, H., Weng, F., Hajek, P., 2023. An interpretable system for predicting the impact of COVID-19 government interventions on stock market sectors. *Ann. Oper. Res.* 1–28. <https://doi.org/10.1007/s10479-023-05311-8>.
- Zanakis, S.H., Zopounidis, C., 1997. Prediction of Greek company takeovers via multivariate analysis of financial ratios. *J. Oper. Res. Soc.* 48 (7), 678–687.
- Zhang, W., Wang, K., Li, L., Chen, Y., Wang, X., 2018. The impact of firms' mergers and acquisitions on their performance in emerging economies. *Technol. Forecast. Soc. Chang.* 135, 208–216.

Petr Hajek is a Professor of system engineering and informatics at the Faculty of Economics and Administration, University of Pardubice. He received a Ph.D. degree in system engineering and informatics from the University of Pardubice in 2006. He is currently the head of the Science and Research Centre with the Faculty of Economics and Administration. He is the author or co-author of four books and more than 100 articles published in

journals such as *Technological Forecasting and Social Change* and *Decision Support Systems*. His-current research interests include business decision-making and knowledge-based systems. He is also a member of the editorial board of several international leading journals.

Roberto Henriques is currently Associate Professor at NOVA IMS from NOVA University. He holds a PhD in Information Management (NOVA), a Master's in Geographic Information Systems and Science and a degree in Biophysics Engineering. His research interests include learning analytics, educational data mining, and decision support systems using artificial intelligence and data mining methods. His work has been published in several high-quality international journals and conferences. He has also been involved in several innovation projects, including machine learning, marketing analytics and CRM, patient management and clinical mining.