

LongFinBERT: A Language Model for Long Financial Documents*

Tri Minh Phan^a, Erik-Jan Senn^a

Version: August 22, 2024

Abstract

This paper introduces LongFinBERT, a modern language model specialized for processing long financial documents. Due to an adaptation in model architecture, LongFinBERT demonstrates substantially lower computational requirements for lengthy documents compared to other state-of-the-art language models. This characteristic enables processing of e.g. an entire annual accounting filing at once, which was previously computationally infeasible for LMs. We apply LongFinBERT to two empirical settings: Firstly, we aim to improve the detection of financial misreporting using text from 10-K filings from 1994 to 2018. Misreporting predictions that utilize text-based features from LongFinBERT outperform those based solely on accounting variables or other textual models, namely Latent Dirichlet Allocation, neural document embeddings, and FinBERT. Lastly, we find that market returns respond to year-over-year alterations of accounting disclosures, measured using LongFinBERT.

Keywords: 10-K filings, deep learning, language model, machine learning, financial misreporting detection

*We thank participants of the Conference of Computational and Financial Econometrics (CFE 2023), the GPEF-poster session at the University of St. Gallen, and Francesco Audrino, Simon Feistle, Matthias Fengler, Lyudmila Grigoryeva, and Jule Schuettler for their valuable feedback.

^a: Department of Economics, University of St.Gallen

Corresponding author: Phan Minh Tri, Department of Economics, University of St.Gallen, Switzerland, Bodanstrasse 6, 9000 St. Gallen, Email: triminh.phan@unisg.ch.

1 Introduction

Current accounting and financial research has rapidly adopted natural language processing (NLP) techniques to analyse financial documents (Loughran and McDonald, 2016; Gentzkow et al., 2019; Ash and Hansen, 2023; Bae et al., 2023; Bochkay et al., 2023). Most studies are based on classical NLP methods, e.g. word counting, despite recent advancements in NLP, notably deep-learning based language models (LMs).

LMs such as FinBERT (Huang et al., 2023) have strong linguistic capabilities, but are computationally costly, especially for long documents. Yet, many use cases in accounting and finance require to extract holistic information from long and complicated documents such as accounting and regulatory filings, company and central bank reports, earning call transcripts. This leads to a significant gap in the literature: long documents in accounting and finance have yet to be fully addressed by powerful tools like LMs.

To bridge the gap, this paper introduces LongFinBERT, an LM that efficiently handles lengthy financial documents. The main feature of LongFinBERT compared to previous LMs in the financial domain is its computational feasibility for long documents. As a consequence, while standard LMs are computationally restricted to analyzing sections of long documents separately, LongFinBERT can process the entire document at once, capturing potential dependencies across different sections.¹

The construction of LongFinBERT is the first focus of this paper. Besides incorporating *LongNet*, we allow LongFinBERT to inherit the architectural and linguistic foundations of FinBERT (Huang et al., 2023), a successful LM adapted to financial contexts. LongFinBERT is pre-trained using a large set of 10-K filings of US firms from January 1994 to December 2008, enabling it to grasp the language style used in financial and accounting contexts.² After constructing the model, we turn to answer the following questions: (i) *Is LongFinBERT more computationally efficient than other LMs?* (ii) *Can LongFinBERT extract more valuable information from financial texts than other NLP models?*

With respect to computational efficiency, Ding et al. (2023) provide the argument for why the computational complexity of the *LongNet* attention mechanism scales only linearly with the length of the document. However, in practice, also other model components

¹For example, when using LongFinBERT to analyse company risks from 10-K filings, two relevant sections discussing company risks (Item 1A - Risk Factors and Item 7 - Management’s Discussion and Analysis of Financial Condition and Results of Operations) are processed jointly by the model and their contents can interact, as opposed to processing each section separately with a standard LM.

²We differentiate *training* and *pre-training* of an LM that *pre-training* attempts to teach an LM fundamental linguistic knowledge, while *training* involves training an LM for a downstream task, such as risk management, sentiment analysis or fraud detection. *Pre-training* solely relies on textual data without supervised labels while *training* requires task-specific labels.

impact performance. We conduct a simulation study to compare the computational demands of LongFinBERT, FinBERT and three other LMs specialized in handling long documents. We find that LongFinBERT requires less time and memory usage than the competing LMs during the training and prediction phases. More importantly, LongFinBERT can process extremely long documents at once, while most competing LMs fail due to extensive memory usage exceeding the available GPU memory.³ This comparison showcases the low computational cost of LongFinBERT for extracting text features and indicates efficient fine-tuning for practical applications.

Then, we evaluate LongFinBERT’s ability to produce meaningful text-based features in the accounting and financial context, specifically to detect financial misreporting. Following the accounting literature, we use data of the SEC Accounting and Auditing Enforcement Releases (AAER) to identify misreporting firms (Beneish, 1999; Dechow et al., 2011; Cecchini et al., 2010; Brown et al., 2020), and match it with financial variables and 10-K filings. We then train a machine learning classifier to predict misreporting cases, using textual features extracted from 10-K filings and accounting variables proposed by Dechow et al. (2011). The results show that text-based features of LongFinBERT enhance the misreporting detection ability of accounting variables. In comparison to other NLP methods including FinBERT, LongFinBERTs text features yield the best misreporting predictions. In addition, we analyze the effects of a data leakage phenomenon related to long-lasting misreporting incidents, partly discussed by Walker (2021). We demonstrate that prediction performance is reduced when controlling for data leakage, raising concerns about findings in Bao et al. (2020) and Brown et al. (2020).

To further explore the capabilities of LongFinBERT in processing accounting disclosures, we revisit the research question explored by Brown and Tucker (2011), in which they examine the relationship between year-over-year modifications of management disclosures and the subsequent market reactions. Different from Brown and Tucker (2011) who employ a bag-of-words method, we compute the year-over-year cosine similarity of management disclosures using the corresponding document vectors derived from LongFinBERT. We uncover a trend toward the increased similarity of management disclosures over time, consistent with Brown and Tucker (2011)’s findings. However, in contrast to their conclusions, we find that the stock market is indeed responsive to the alternations of the Management’s Discussion and Analysis section of 10-K filings.

We primarily contribute to the burgeoning research on developing LMs for the accounting and finance domain. In spirit of Huang et al. (2023) that provide a flexible LM to process shorter texts for various financial and accounting tasks, this paper provides the computationally feasible and efficient counterpart for longer text documents. The ability to process entire documents at once makes LongFinBERT a valuable tool to analyze

³The simulation is conducted using a NVIDIA A100-SXM4 with 40 GB of GPU memory.

a wide range of questions relying on financial documents. We openly release the pre-trained `model` and `code` in hope to assist further research at the intersection of NLP with accounting and finance.

Our contribution extends to the field of misreporting detection, which has been under study since the 1990s (Dechow et al., 1996; Beneish, 1997, 1999). While recent research has attempted to leverage NLP techniques to exploit the abundance of accounting textual data (Brown et al., 2020), the increasing size and complexity of such documents have constrained these studies to traditional word-count-based methods. To our knowledge, we are the first to tackle this challenge using a state-of-the-art LM. The superior performance of LongFinBERT in this accounting-related task highlights the potential of text-based features extracted from accounting disclosures using modern NLP techniques in detecting misreporting instances.

The remaining paper is organized as follows: Section 2 reviews the current state of NLP models in accounting and financial research, Section 3 describes the construction and pre-training of LongFinBERT. A simulation study on LongFinBERT’s computational efficiency is provided in Section 4. Section 5 covers the application of LongFinBERT to financial misreporting detection by discussing the relevant literature, data, research design, and empirical results. Section 6 analyzes the year-over-year similarity of management disclosures using LongFinBERT and its associations with the stock market. Finally, Section 7 concludes the paper and discusses possible limitations and future applications of LongFinBERT.

2 NLP models in accounting and finance research

The main aim of natural language processing (NLP) is to enable computers to understand human language by converting text into meaningful numeric representations that can then be used for a variety of (financial) tasks; such as predicting a financial variable from text information (see Section 5). The numeric representation can be seen as a compressed version of the text that should still capture the essential information required for various (financial) applications.

A useful NLP-model should have semantic and contextual understanding of the text, and be computationally feasible. The most commonly used model (also in economics) can be categorized into three large classes: (i) word-count-based models; (ii) neural embedding models; and (iii) language models. Now follows a description those NLP model classes.

Word-count-based models

The most classical way to represent a document is via the occurrences of words in the document, referred to as *Bag-of-words* (Harris, 1954; Spaerck Jones, 1972). Particularly, the algorithm begins by defining a set of vocabulary. A document is then represented by a vector containing the frequency of each word from the pre-defined vocabulary appearing within that document. Later, Manning and Schutze (1999) suggest a transformation of the word frequency, named as *term frequency - inverse document frequency*, to obtain a more efficient document representation.

Due to its straightforward implementation and low computational cost, this method still maintains a foothold in current economic research. Jegadeesh and Wu (2013) represent 10-K filings by the transformed frequency of words chosen from the financial sentiment vocabulary of Loughran and McDonald (2011) to identify which words make a stronger impression on investors. Also based on the word-count method, Chen et al. (2019) examine the gain obtained by FinTech innovations of corporations in the financial service sector. Frankel et al. (2022) combine the word-count-based document representation and various machine learning (ML) methods, demonstrating that sentiment expressed in 10-K filings correlates with contemporaneous market reactions. Chen et al. (2022) propose a supervised approach to measure the tone of NASDAQ articles using their word-count representation. Deviating from document representation, a significant body of literature relies on the word-count method to measure several textual attributes and connect them to economic variables; see Tetlock (2007), Loughran and McDonald (2011), Brown and Tucker (2011), Price et al. (2012), Muslu et al. (2015), Jiang et al. (2019), Cho and Muslu (2021), among others.

Bag-of-word approaches are widely used despite two well-known downsides: First, because of the sole reliance on word frequency, the word-count method treats words as independent tokens. Therefore, the approach ignores word semantics and context (Mikolov et al., 2013). Second, word-count-based representation is typically high-dimensional due to the proliferation of words used to compile texts. For example, the vocabulary used in Loughran and McDonald (2011) and Jegadeesh and Wu (2013) contains more than 3,000 words; and Chen et al. (2022) derive vocabulary consisting of approximately 43,500 words. A high-dimensional text representation complicates downstream machine learning or econometric modeling e.g. by relying on regularization.

Neural embedding models

Neural embedding models like Word2Vec (Mikolov et al., 2013) or Doc2Vec (Le and Mikolov, 2014) are introduced to overcome the aforementioned issues of the word-count method. Word2Vec is a neural-network-based algorithm that examines the similarity between words based on their co-occurrences in a corpus. The output of Word2Vec is a set of dense and low-dimensional semantic vectors, often referred to as *word embeddings*, associated with every word within the vocabulary. Words sharing similar meanings or usage tend to have

closer embeddings, based on Word2Vec, in the vector space, thus capturing word semantics. Doc2Vec is a modification of Word2Vec applied for documents instead of words. Consequently, through Doc2Vec, documents are represented by dense, low-dimensional, and semantic vectors, called *document embeddings*, as opposed to the word-count method.

With these advances, neural embedding models have gained much attraction from economic researchers recently. Li et al. (2021) leverage Word2Vec to quantify corporate culture by analyzing earnings call transcripts of US firms. Fengler and Phan (2023) use Word2Vec to model topics in the Management’s Discussion and Analysis section of 10-K filings. Furthermore, Miranda-Belmonte et al. (2023) utilize various neural embedding models to construct an economic policy uncertainty index derived from US and Mexican newspapers. Using Doc2Vec, Rawte et al. (2020) measure the similarity of 10-K filings of the same firms between two consecutive years. They show that textual similarity strongly predicts key performance and risk measures of banks.

Neural embedding models can capture word semantics, but only capture limited context because embeddings are static after training. Huang et al. (2023) provide a canonical example of the word “bank” in which it has different meanings in the two following sentences: “I went to the bank to deposit a check” and “We walk along the river’s bank”. Nevertheless, Word2Vec embeds the word “bank” into a single vector, failing to capture its contextual nuances. Similarly, algorithms like Doc2Vec encode entire documents into fixed vectors, regardless of word order. This means that documents composed of the same words, even if in different arrangements, receive identical embeddings. For example, in our experiment using Doc2Vec, the two sentences (documents) “Sales increase but expenses decrease in this year” and “Sales decrease but expenses increase in this year” have identical embeddings despite their opposite meanings.⁴

Language models (LMs)

To address the issue of static embeddings of neural embedding algorithms, language models (LMs) are introduced. Early work on LMs modelled text as a time series by storing the context in the state of a recurrent neural network (Wang and Cho, 2016), (Soutner and Müller, 2013). However, these models were computationally expensive due to sequential estimation. The concept of *multi-head attention* introduced by Vaswani et al. (2017) quickly turned into the standard building block of every LM: The mechanism allows words to pay attention to other words in a document, thereby producing dynamically contextualized

⁴To conduct this small experiment, we use a Python package named *gensim* (Řehůřek and Sojka, 2010). With our settings, the cosine similarity between these two sentences is 0.9980, while the cosine similarity between the first sentence and itself is 0.9976. We want to note that document embeddings generated by this package are subject to some level of randomness. Consequently, a Doc2Vec model trained by *gensim* may produce different embeddings even with identical documents. To mitigate this variability, it is recommended to set a high value to the `epochs` argument in the `similarity_unseen_docs` method of the Doc2Vec model. In our case, we opt for a setting of 10,000 epochs.

word embeddings. While not formally belonging to the class of time series models, the self-attention idea models relationships between words well and is both expressive and computationally efficient. Numerous attention-based LMs have been developed and introduced, such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), DeBERTa (He et al., 2021). Smart training methodologies are employed that avoid the need for labeled training data: most LMs are (partly) trained using masked language modelling, where the model has to predict randomly hidden words in a document from the surrounding words.

Aiming to provide a powerful LM specialized in the financial context, Huang et al. (2023) introduce FinBERT, an LM fine-tuned using a large financial corpus of corporate filings, analyst reports, and earnings call transcripts. They show that FinBERT outperforms other text models in sentiment and ESG (environmental, social, and governance) classification. By using FinBERT, Cao et al. (2023) examine how firms adapt their communication behavior with investors and stakeholders knowing that their disclosures are analyzed by AI models. Besides that, several authors attempt to adapt BERT or generative pretrained transformer (GPT) models for specialized financial applications, namely SEC-BERT (Loukas et al., 2022), FinChain-BERT (Yang, Zhang, Sun, Pang, Jing, Wa and Lv, 2023), FinGPT (Yang, Liu and Wang, 2023) and BloombergGPT (Wu et al., 2023).

Recent studies regarding how LMs work reveal that their word embeddings consist of both syntactic and semantic knowledge (Rogers et al., 2021), remarkably enhancing their linguistic capability (Devlin et al., 2019). While LMs probably have the best capabilities in terms of capturing word semantics and context compared to the previously presented methods, this comes at a computational cost. Especially for long documents, the attention mechanism is a computational bottleneck: The multi-head attention mechanism operates by calculating a square matrix, also referred to as *attention matrix*, whose dimensionality is the document length (i.e., the number of tokens in a document).⁵ This calculation exhibits a quadratic computational complexity in terms of document length. Specifically, when a document is doubled in length, the cost to compute the attention matrix increases fourfold. For example, consider a document containing 55,000 tokens, which is the average length of 10-K filings (Cazier and Pfeiffer, 2016). To compute a BERT attention matrix with twelve attention heads using single-precision floating-point format (i.e., 32 bits or four bytes), one needs $4 \times 55,000 \times 55,000 \times 12 = 145,200,000,000$ bytes = 145 gigabytes of memory to process only one layer of BERT, which has twelve layers and other parameters. Such a memory demand exceeds the capacity of the best currently available individual Graphics Processing Units (GPUs), which are the main workhorses for training

⁵Text models that are not LMs typically treat words as the smallest units in processing texts. Therefore, they break documents into words (or phrases) when generating vocabulary. This process is called *tokenization*. LMs tokenize documents more granularly, further breaking words into *subwords* occasionally. For instance, BERT (Devlin et al., 2019) splits the word “volatility” into three subwords “vol”, “ati”, and “lity”, but does not split the word “cash”. We follow Huang et al. (2023) to define *tokens* as words and subwords.

LMs.

Overall, the computational burden remains one of the main factors delaying applications of LMs in economic research, where some long text documents are of particular interest (e.g. company or central bank reports). Many different attention mechanisms are proposed to deal with this computational challenge. One approach to handle long documents is dividing the text into finer units, such as paragraphs or sentences (Huang et al., 2023). However, this approach does not allow for interactions between individual text units in a document. Another approach is to sparsify the attention matrix by a combination of sliding-window and global attention mechanism (Zaheer et al., 2020; Beltagy et al., 2020) or via Nyström method to approximate the attention matrix (Xiong et al., 2021). The main idea of sparsification of the attention matrix is to not allow all tokens far away from each other in a document to interact, while still allowing interactions between tokens in a close neighborhood: this leads to a reduction in computational complexity at a hopefully low cost in terms of semantic and contextual understanding. Our model *LongFinBERT* introduced in Section 3 employs a dilated attention mechanism to sparsify the attention matrix, enabling efficient estimation of an LM for very long financial documents with reduced computational demand.

3 LongFinBERT - A BERT-like language model specialized for long financial documents

This section presents LongFinBERT, an efficient language model specialized for long financial documents. At the heart of LongFinBERT is the dilated attention mechanism developed by Ding et al. (2023), which is visualized in Figure 1. Specifically, the dilated attention mechanism aims to sparsify a standard attention matrix by applying multiple segments on a (long) document. Each segment has a different size and dilated rate. The segment size controls for local information captured by the segment. For example, a segment size of four tokens allows the model to focus on every four consecutive tokens. The segment dilated rate controls for the sparsity of the attention matrix. A dilated rate of four implies that the first token only pays attention to the fifth, ninth, and thirteenth tokens, etc. By increasing the dilated rate, one can allow the attention matrix to capture the interactions between tokens located at a far distance in a long document, thus guaranteeing the sparsity of the attention matrix. Finally, the attention matrices processed by multiple segments with increasing segment sizes and dilated rates are aggregated. Consequently, the aggregated attention matrix incorporates both local (by small segment sizes and dilated rates) and global (by large segment sizes) information while staying sparse (by large dilated rates).




Figure 1: This figure visualizes the attention matrix processed by LongNet Ding et al. (2023). Each small square indicates the interaction between two tokens. For example, the second small square in the first row presents the interaction between the first and second tokens in a document. Colored small squares indicate the considered interactions while white squares denote ignored ones, resulting in a highly sparse attention matrix. This contrasts with the vanilla attention matrix used by FinBERT (Huang et al., 2023), where all squares are colored, indicating a dense attention matrix.

Relying on this dilated attention mechanism, we build an encoder-only transformer architecture similar to BERT (Devlin et al., 2019). LongFinBERT is constructed by stacking up twelve encoders, as shown in Figure 2. Unlike FinBERT, which leverages the pre-trained weights (i.e., model parameters) from BERT, with the new attention mechanism of LongFinBERT there is no option to initialize our model with pre-trained weights. Intuitively, while FinBERT can capitalize on fundamental linguistic understanding from BERT’s extensive pre-training, LongFinBERT needs to learn these foundational elements anew. However, equipping an LM with fundamental linguistic skills requires massive data and computational costs (Devlin et al., 2019; Zaheer et al., 2020). To address this challenge, we enable LongFinBERT to adopt FinBERT’s tokenization and word embedding layer, the two deepest parts of FinBERT’s model parameters. This strategy allows LongFinBERT to leverage fundamental linguistic knowledge specific to the financial domain of FinBERT, thus relieving the demand for extensive data and computational power during the pre-training of LongFinBERT.⁶

⁶We differentiate *training* and *pre-training* of an LM. *Pre-training* attempts to teach an LM fundamental linguistic knowledge, while *training* involves training an LM for a specific task. In that sense, *pre-training* solely relies on textual data and does not require supervised labels. *training*, on the opposite, necessitates supervised labels for a specific task.




Figure 2: This figure illustrates the architecture of LongFinBERT, with the core component being LongNet, detailed in Figure 1. Initially, a document (e.g., 10-K filing) is converted to numerical representations by embedding layers. The encoded sequences are then passed through twelve encoders. Each encoder is constructed by the LongNet attention mechanism, followed by a linear projection and a normalization layer (i.e., LayerNorm). After the twelve encoders, the encoded document is processed by an optional pooling layer (i.e., Pooler) and a classification layer (i.e., Classifier).

Similar to BERT and FinBERT, we employ the masked language modeling task to pre-train LongFinBERT. Specifically, we randomly mask (hide) 15% of tokens in a document. Then, LongFinBERT learns to predict the masked tokens based on the remaining un-masked tokens. We pre-train LongFinBERT using a large set of 100,717 10-K filings of US firms from 1993 to 2008.⁷ The dataset from 2009 to 2013 is used for model validation and includes 44,342 10-K filings. We selected the dataset of 10-K filings for pre-training LongFinBERT because of three reasons. First, 10-K filings are usually very lengthy with approximately 55,000 words on average (Cazier and Pfeiffer, 2016), aligning well with the intended spirit of LongFinBERT in handling long financial documents. Second, 10-K filings are widely used by economic and financial studies involving text data (Loughran and McDonald, 2011; Li, 2010; Jegadeesh and Wu, 2013; Jiang et al., 2019; Brown et al., 2020; Huang et al., 2023). Third, 10-K filings are chosen over 10-Q filings as Griffin (2003) show that the stock markets are more responsive to the information of 10-K compared to 10-Q filings.

In total, our pre-training sample includes 5.85 billion tokens. Despite containing fewer

⁷10-K filings are long comprehensive annual financial reports submitted by U.S. public companies to the Securities and Exchange Commission (SEC), aimed at providing standardized public access to relevant information for shareholders and potential future investors. The filings contain detailed information about a company’s financial performance, risk factors, management discussion and analysis (MDA), corporate governance, executive compensation, and audited financial statements. They are typically released within 60-90 days after the end of a company’s fiscal year.

texts, our corpus is about 20% and 77% larger than the corpora used to pre-train FinBERT and BERT, respectively. This is because we consider the whole 10-K filings instead of focusing on only several sections as in [Huang et al. \(2023\)](#). We pre-train LongFinBERT with an AMD EPYC 7763 machine with 64 cores and an NVIDIA A100-PCIe-80GB GPU. We apply five different segments to LongFinBERT with the segment sizes of $\{16, 128, 512, 1024, 2048\}$ and dilated rates of $\{1, 16, 64, 256, 512\}$. LongFinBERT is trained with eight iterations, in which each iteration takes around 14 hours. Compared to FinBERT, which spent approximately 72 hours on one iteration, LongFinBERT exhibits a five-fold improvement in the pre-training time. The pre-trained model weights and training code are publicly available.⁸

4 Computational efficiency of LongFinBERT

In Section 3, LongFinBERT is introduced aiming to improve the computational efficiency of LMs. Particularly, it simplifies the computational complexity of the standard multi-head attention mechanism from quadratic to linear complexity. Therefore, LongFinBERT is expected to outperform other LMs in practice in terms of computational costs, in which time and memory are the main two aspects.

Accordingly, we design a simulation study in order to compare the computational performance of LongFinBERT with that of four other LMs. The competitors include: (i) Google Big Bird ([Zaheer et al., 2020](#)), Longformer ([Beltagy et al., 2020](#)), Nyströmformer ([Xiong et al., 2021](#)), and FinBERT ([Huang et al., 2023](#)). The first three LMs are selected due to their capability to handle long documents. FinBERT is chosen, however, because of its popularity in economic/financial/accounting studies ([Cao et al., 2023](#); [Park and Cho, 2023](#); [Schimanski et al., 2024](#)). It is worth noting that we refrain from comparisons in the computational performance of LongFinBERT with non-LM algorithms. This is because non-LM algorithms, such as Bag-of-words, Word2Vec, or Doc2Vec, are typically trained on Central Processing Units (CPUs) meanwhile Graphics Processing Units (GPUs) are more commonly used to train LMs. Therefore, a comparison of the computational costs between these two model classes would be generally unreasonable.

To set up the simulation study, we generate a dummy document with variable lengths ranging from 64 (2^6) tokens to 65,536 (2^{16}) tokens. Subsequently, we examine the computational time and memory of the five LMs, including LongFinBERT, to process this document in both two phases, training and inferring.⁹ In particular, during the train-

⁸The pre-trained model weights: <https://huggingface.co/minhtriphon/LongFinBERT-base>.
The training code: <https://github.com/minhtriphon/LongFinBERT-base/tree/main>.

⁹Different from the inferring phase, during the training phase, the gradients (i.e., derivatives of the loss




Figure 3: This figure reports the results of the simulation study comparing the computational performances of LongFinBERT (blue) with those of four other competing LMs, namely Google Big Bird (orange), Nyströmformer (green), FinBERT (red), and Longformer (violet). The results include the training time (top-left), training memory (top-right), inferring time (bottom-left), and inferring memory (bottom-right) of all five LMs. The document length (number of input tokens) is displayed in the x-axes using a logarithmic scale. The time (in seconds) and memory (in megabytes) are displayed in the y-axes.

ing phase, we train each LM with the synthetic document for one iteration per document length value. In the inference phase, we employ each LM to infer the vectorized representation of the dummy document. To mitigate the impact of noise, we run each experiment 500 times and report the median values of time and memory usage. All experiments are conducted using an NVIDIA A100-SXM4 with 40 gigabytes of GPU memory.

It should be noted that there are limits of the competing LMs regarding the maximum length of documents that can be passed through these LMs. For example, FinBERT can

function with respect to the model parameters evaluated at the current data point) are computed by the back-propagation algorithm (Linnainmaa, 1976) and stored in the memory of the used GPUs to implement the stochastic gradient descent optimization (Robbins and Monro, 1951). This operation demands high memory capacity from the GPUs in use as well as longer processing time.

only process documents with no longer than 512 tokens (Huang et al., 2023). Similarly, the maximum limit for Google Big Bird, LongFormer, and Nyströmformer is 4,096 tokens. We attribute these limitations to technical constraints inherent in the positional embedding layer of each LM.¹⁰ For example, when a document with 4,097 tokens is passed through Google Big Bird, the vector representation for the 4,097th position is unidentified because vector representations are provided for only 4,096 positions. Appearing as a technical problem, this limitation possibly causes unfair comparisons between the LMs, especially with FinBERT, when they can theoretically, handle documents longer than these limits. We, therefore, overcome this constraint by expanding the positional embedding layers in the LMs. This modification expectedly ensures an equitable comparison environment for all involved LMs.

Figure 3 showcases the results of the simulation study. In the training phase, LongFinBERT demonstrates the capability to handle a document with 2^{16} tokens. Meanwhile, the other LMs only process a document with only 2^{14} tokens at most with LongFormer, which is 4 times shorter than the longest document supported by LongFinBERT. In the inference phase, LongFinBERT remains the top performer in terms of both time and memory usage. In particular, in comparison with LongFormer, the most efficient among the competing LMs, which requires more than 1.25 seconds and 11,000 megabytes to vectorize a document with 2^{16} tokens, LongFinBERT only needs approximately 0.35 second and 8,000 megabytes to process a document with the same length. That constitutes an advantage of LongFinBERT by $3.6 \text{ (time)} \times 1.38 \text{ (memory)} \approx 5$ times over the best competing LM.

We observe that, however, LongFinBERT exhibits longer processing times compared to the other LMs when handling documents shorter than 512 tokens. Its performance improves when the document length surpasses 2^{10} tokens and becomes strongest with documents longer than 2^{12} tokens. This phenomenon is caused by the utilization of many segments with different sizes in LongFinBERT. Therefore, with short documents, small-size segments collide, leading to redundancy.¹¹ Notably, this phenomenon is primarily evident in time usage. In contrast, LongFinBERT consistently outperforms other LMs in terms of memory usage, regardless of the document length.

Overall, while LongFinBERT may require more time to process short documents, it demonstrates the capability to efficiently process long documents compared to other LMs, including those that are designed for lengthy texts.

¹⁰Positional embedding layers in LMs aim to encode different positions of input tokens by distinct numerical vectors. Consequently, a token (or word) appearing at two different locations in a text, at the end, will be represented by two different vectors. This mechanism allows LMs to capture better word contexts.

¹¹Consider the scenario where LongFinBERT uses two segments of sizes 64 and 512 respectively. When processing a document with 512 tokens, LongFinBERT utilizes both segments to compute the attention matrix. In contrast, the other LMs compute the attention matrix only once.

5 Performance of LongFinBERT on financial misreporting detection

So far, LongFinBERT has shown its superior computational performance in handling long texts. Now, we turn to investigating the economic value of LongFinBERTs textual features for detecting financial misreporting¹² using financial disclosures, which are long text documents often spanning 50 to hundreds of pages. The section introduces relevant literature on financial misreporting detection (5.1), describes the data (5.2) and research design (5.3), and presents the empirical results (5.4, 5.5).

5.1 Literature on financial misreporting detection

Financial misreporting can impose significant economic costs, such as damaging a firm’s reputation, leading to misinformed financing decisions by investors, and reducing overall trust in capital markets. Therefore, detecting misreporting is a crucial component of accounting research.

Many studies have highlighted the relationship between misreporting behavior and accounting variables, thus eliciting the usefulness of accounting data to detect misreporting. For instance, Dechow et al. (1996) observe that misreporting firms exhibit a high amount of accruals before the alleged accounting manipulation. Beneish (1999) propose that variables regarding firm revenues, expenses, and sales growth can help to identify earning manipulations. Building upon the empirical evidence, Dechow et al. (2011) construct a measure called *F-score*, a composite of several financial variables, that can detect misreporting firms. Using a similar wide set of financial variables, Cecchini et al. (2010) construct a non-parametric kernel estimator to predict misreporting.

Newer studies show that language-based information can enhance misreporting detection. Li (2008) discover a positive connection between the more frequent use of causation words, optimistic words, and future tense verbs in the Management’s Discussion and Analysis section of 10-K filings and managerial obfuscation behaviors. Through an analysis of conference call transcripts, Larcker and Zakolyukina (2012) demonstrate that attributes derived from emotion words outperform abnormal accrual measures in detecting misreporting. Brown et al. (2020) use Latent Dirichlet Allocation (LDA), a word-count-based topic modeling framework, to analyze 10-K filings. Their findings suggest that

¹²The term *misreporting* describes deliberately providing inaccurate, misleading, or false information about the financial condition or performance of a company, for a variety of reasons. We refrain from using the term *fraud* because legally, actions are classified as fraudulent only if market participants are adversely affected by relying on the false or misleading financial information.

thematic contents comprised in 10-K filings improves the misreporting detection ability compared to models purely based on financial variables.

So far, despite the immense linguistic capabilities of state-of-the-art LMs, the utilization of 10-K filings in misreporting detection heavily relies on word-counts-based methods, primarily due to the verbosity of the financial documents. Notably, 10-K filings have 55,000 words on average (Cazier and Pfeiffer, 2016) with the possibility of lengthening even further (Brown and Tucker, 2011; Dyer et al., 2017). This fact presents a challenge in employing LMs in detecting misreporting behavior using 10-K filings due to their quadratic complexity in both time and memory dimensions, thus necessitating more efficient approaches.

Addressing this challenge, applying LongFinBERT to 10-K filings in detecting misreporting is efficient in two ways. First, LongFinBERT inherits the superior linguistic capabilities of other LMs in the same model class such as FinBERT. Therefore, LongFinBERT is expected to capture more complicated linguistic structures in 10-K filings in comparison to non-deep-learning approaches. Second, in comparison to other LMs, LongFinBERT is more computationally efficient for long documents as shown in Section 4. Given these benefits, text features from LongFinBERT are promising to improve current text-based misreporting detection models in the literature.

5.2 Data and descriptive statistics

We rely on three different data sources in this application: (i) a text corpus, from which we extract text-based features; (ii) numerical financial and accounting data, used to construct Dechow et al. (2011)’s financial variables; and (iii) a dataset to identify misreporting incidents.

The text corpus

Similar to Brown et al. (2020), we analyse textual data from annual financial statements of US firms, 10-K filings, from 1994 to 2018. 10-K filings include contents such as detailed financial statements and the MD&A section that could contain irregularities, sudden shifts in accounting methods, or other discrepancies that might suggest possible financial misreporting.

We download 10-K filings from 1st January 1994 to 31st December 2018 in the Notre Dame Software Repository for Accounting and Finance (SRAF).¹³ The download yields in total 183,057 10-K documents filed by 34,972 US firms over the sample period. In accordance

¹³<https://sraf.nd.edu/sec-edgar-data/cleaned-10x-files/>

with the current norm in using LMs (Huang et al., 2023), we do not apply intensive text normalization steps except excluding the SEC header in each 10-K file. To apply non-LM approaches, inspired by Frankel et al. (2022), we execute several standard steps of text normalization, including word lowercasing, contraction fixing, number removal, punctuation removal, and lemmatization.

Financial variables

Following Dechow et al. (2011), using the Computstat and CRSP database, we construct 17 financial variables capturing firm size, audit quality, and involvement in complex business transactions, namely, merger and acquisition activities, and restructurings.¹⁴ These variables include: total assets ($\log(\text{Assets})$), accruals calculated by Richardson et al. (2005) (*RSST_Accruals*), changes in account receivables ($\Delta\text{Receivables}$), changes in inventory ($\Delta\text{Inventory}$), soft asset ratio ($\%\text{Soft_assets}$), changes in return on assets (ΔROA), the lag of firm stock excess log returns over one year prior to the filing date ($\text{Lag}(r)$), cash flow from financing activities (*TotFinancing*), the book-to-market ratio (B/M) and a set of dummy variables on: securities issuance (*Actual_issuance*), carrying future operating lease obligations (*Operating_leases*), merger and acquisition (*M&A*), being audited by the Big-4 auditors (KPMG, Deloitte, EY, and PwC) (*Big4Auditors*), being audited by mid-size auditors (i.e., BDO, Grant Thornton, McGladrey) (*MidSizeAuditors*), ex-ante financing needs (*ExanteFinancing*), nonzero restructuring charges (*Restructuring*).

After dropping all firm-year observations with missing values in any of the 17 financial variables, the dataset consists of 62,037 observations coming from 8,843 US firms over the sample period. The financial variables will serve as features for the baseline misreporting detection model.

Identifying misreporting behavior

We rely on the dataset SEC¹⁵ Accounting and Auditing Enforcement Releases (AAERs) to identify misreporting firms, and is widely used in misreporting detection research (Beneish, 1999; Dechow et al., 2011; Perols et al., 2017; Brown et al., 2020; Bertomeu et al., 2021; Yi et al., 2023).

The dataset was initially collected by Dechow et al. (2011) and recently updated by the Marshall Business School of the University of South California. The dataset comprises 4,278 SEC AAERs, revealing 1,816 manually detected cases of misreporting involving 1,768 distinct firms ranging from 1971 to 2019. Among these cases, two firms committed misreporting three times, 44 firms were found to misreport twice, and 1,722 firms were

¹⁴We refer readers to Dechow et al. (2011) and Appendix B in Brown et al. (2020) for detailed instructions on calculating these variables.

¹⁵U.S. Securities and Exchange Commission

implicated in misreporting once. On average, the identified period of misreporting for a single incident approximately 2.33 years and involves about 1.88 misreporting categories. The longest case involves Comverse Technology, Inc. (Ulticom) lasting from 1991 to 2005 in which they were alleged to have misreported the cost of goods sold, reserve accounts, expenses/shareholder equity, and liabilities.

In line with [Dechow et al. \(2011\)](#), misconduct events unrelated to accounting or auditing, such as bribery, disclosure-related and timing violations, or events that lack clear identification of the underlying cause of the violations are excluded. There are eleven remaining misconducts including the misreporting behaviors regarding: revenues (*REV*), other expense/shareholder equity accounts (*EXP_SE*), capitalized costs as assets (*ASSET*), account receivable (*REC*), inventory (*INV*), cost of goods sold (*COGS*), liabilities (*LIAB*), reserve accounts (*RES*), allowance for bad debt (*DEBT*), marketable securities (*MKT_SEC*), and payables (*PAY*). We identify a misreporting firm as a firm that committed at least one of the misconducts listed above.

One should note that misreporting data includes a risk of inaccurate information: Since the SEC likely focuses their investigations on firms with a strong indication of some type of misreporting, AAERs are likely to have a low type 2 error rate (labeling a firm as misreporting when it is not), but a relevant type 1 error rate (missing out on a misreporting firm). In addition, the SEC might have a selection bias for certain types of firms; and the substantially lower fraud rate after 2005-2008 hints at a changes in the SEC's focus and resource allocation in time.

The final dataset

Finally, we match the three datasets, i.e., 10-K filings, financial variables, and AAER misreporting cases, to obtain the final dataset used to train our misreporting detection models. The final dataset comprises 50,663 company-year observations of 7,342 distinct firms spanning the period from 1st January 1994 to 31st December 2018. There are 491 firm-year misreporting observations, approximately accounting for 1% of the total number of observations. In comparison with [Brown et al. \(2020\)](#), whose sample period ranges from 1994 to 2010, our dataset exhibits a misreporting rate of 1.21% in the same sub-sample period. This quantity is comparable to the corresponding proportion reported by [Brown et al. \(2020\)](#), which is 1.34%.

Descriptive statistics

Figure 4 shows the fraud rate and total number of firms in the sample over time. In general, consistent with other studies using the AAERs dataset, the misreporting rate is around 0.2-2%. The misreporting rate exhibit a sharp increase until 2001, followed by a gradual decline thereafter. This variation coincides with the enactment of the Sarbanes-




Figure 4: This figure displays the misreporting rate in percent and the total number of companies in the sample (sample size) for each year. The data sample spans from 1st January 1994 to 31st December 2018.

Oxley Act in 2002 which mandates several practices in financial and accounting reporting to protect investors from fraudulent behaviors of firms. From 1997 on, the sample size is stable around 2000 firms per year.

Table 1 displays the descriptive statistics of the 17 financial variables described in Section 5.2. Additionally, we also present Welch's *t*-tests of the mean difference of these variables between the no-misreporting and with-misreporting samples. Our statistics align with those of Brown et al. (2020) in terms of magnitude, relation direction, and statistical significance. Specifically, misreporting firms exhibit significantly higher total assets, accruals, and soft asset proportions. They demonstrate notably higher year-by-year changes in account receivables, inventory, and cash sales. Furthermore, misreporting corporations issue greater volumes of securities than their non-misreporting counterparts. Interestingly, our findings reveal that misreporting firms are more frequently audited by Big-4 auditing firms, consistent with Brown et al. (2020). With many variables displaying significant differences between non-misreporting and misreporting samples, these variables are potentially useful for identifying misconduct behaviors.

	No misreporting ($N = 50,172$)		Misreporting ($N = 491$)		Difference
	Mean	Std. Dev.	Mean	Std. Dev.	
$\log(\text{Assets})$	5.959	2.059	6.591	1.726	-0.633***
<i>RSST_Accurals</i>	0.025	0.202	0.071	0.247	-0.046***
$\Delta\text{Receivables}$	0.005	0.066	0.019	0.060	-0.013***
$\Delta\text{Inventory}$	0.004	0.047	0.011	0.063	-0.008***
$\%\text{Soft_assets}$	0.476	0.429	0.663	0.276	-0.186***
$\Delta\text{Cash_sales}$	0.171	1.591	0.242	0.609	-0.070**
ΔROA	-0.010	0.240	-0.009	0.154	-0.001
<i>Actual_issuance</i>	0.884	0.320	0.921	0.271	-0.036***
<i>Operating_leases</i>	0.777	0.417	0.876	0.330	-0.099***
<i>M&A</i>	0.151	0.358	0.088	0.283	0.063***
$\text{Lag}(r)$	-0.007	0.628	-0.023	0.653	0.016
<i>Big4Auditors</i>	0.727	0.446	0.778	0.416	-0.051***
<i>MidSizeAuditors</i>	0.101	0.301	0.088	0.283	0.014
<i>TotFinancing</i>	0.030	0.196	0.029	0.134	0.001
<i>ExanteFinancing</i>	0.087	0.281	0.029	0.167	0.058***
<i>Restructuring</i>	0.013	0.115	0.022	0.148	-0.009
<i>B/M</i>	0.695	0.705	0.632	0.524	0.064***

Table 1: This table reports means and standard deviations of 17 financial variables in two samples: one with misreporting and one without misreporting. The table also displays the difference in means between the samples and the corresponding two-sided Welch’s t -test testing for equality in means between samples. The definitions of the financial variables are given in Section 5.2. The data sample spans the period from 1st January 1994 to 31st December 2018. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

5.3 Research design

The aim of our empirical methodology is to compare how well LongFinBERT and other text models can extract relevant text information for misreporting detection.

Model structure

Figure 5 illustrates the model diagram for our misreporting detection problem. Specifically, each 10-K file is passed through a text model to extract meaningful text-based features in numerical form. It is worth noting that text models are adopted as feature extraction tools without further training or fine-tuning. The text-based features are then augmented with the 17 financial variables of the corresponding firm-year observation. This augmentation is achieved through vector concatenation (combining the vectors end-to-end to a single longer vector), resulting in a unified vector that integrates both textual and numerical financial information for each firm-year observation.¹⁶ Subsequently, this (combined) feature vector serves as input to train a pre-defined classifier tasked with

¹⁶We also report results on models that only include text or only include financial variables as benchmarks. For those models, no concatenation is required.




Figure 5: This figure illustrates the model diagram used to detect misreporting incidents. Text models include LongFinBERT, FinBERT, LDA, Doc2Vec, and TF-IDF representation. 17 financial variables are introduced in Section 5.2.

detecting misreporting. The classifier predicts the probability of misreporting incidents ($is_misreporting = 1$) as outputs.

Due to the high dimensionality of text-based features and the low frequency of misreporting cases, misreporting detection including text data is identified as a binary classification with rare events (few misreporting cases) in high-dimensional data. Consequently, we opt for gradient boosting (Freund et al., 1996) as the classifier.¹⁷ With the demonstration of strong performance in high-dimensional settings (Lusa et al., 2017), the algorithm is used by many studies in misreporting detection (Whiting et al., 2012; Bao et al., 2020; Bertomeu et al., 2021). However, Lusa et al. (2017) notice that boosting algorithms may struggle when encountering imbalanced data, particularly being prone to overfitting on training data. Accordingly, they suggest employing two regularization techniques to cope with imbalanced data: (i) shrinkage (Friedman, 2001); and (ii) subsampling (Friedman, 2002). We implement gradient boosting algorithms by adopting *XGBoost*, a Python library for gradient boosting developed by Chen and Guestrin (2016). *XGBoost* stands out as a preferred tool for developing boosting algorithms due to two crucial mechanisms: (i) embedded regularization within loss functions, and (ii) variable and observation sub-

¹⁷Another possibly feasible classifier might be regularized logit, however we expect better performance using gradient boosting.




Figure 6: This figure illustrates the expanding-window research design. The training of the classifier starts with data from 1994 to 1997. The data in 1998 are used to optimize the classifier hyperparameters (number of boosting rounds), thereby mitigating the risk of overfitting. The predictions are made for the year of 1999. The cycle repeats by expanding the training data by one year in each step. Consequently, the data for hyperparameter optimization and prediction are then shifted by one year to the future. The process continues until the prediction year reaches 2018.

sampling mechanism. According to [Lusa et al. \(2017\)](#), these two features are supposed to mitigate the risk of overfitting, thus expectedly improving the misreporting detection capability.

The presented setup ensures that measured differences in misstatement prediction performance are only related to the choice of text model and not other aspects of the pipeline, as each model uses the same information set (10-K and /or financial variables) and the same flexible classification algorithm (where the misstatement classifier is optimized separately for each NLP-model). The final models differ by the NLP-model used to extract text-based features, which include LongFinBERT, LDA, Doc2Vec or FinBERT.

Estimation and forecasting

Figure 6 describes our estimation and out-of-sample forecasting procedure. In contrast to the rolling-window approach employed by [Brown et al. \(2020\)](#), where a fixed number of five years is used in the training sets, we opt for the expanding-window procedure ([West and McCracken, 1998](#)). This choice is motivated by the high-dimensional nature of our problem, necessitating more data to obtain reliable estimations. We initiate the process using the data from 1994 to 1997 to train the classifier (gradient boosting). The data in 1998 is used to optimize regularization hyperparameters (i.e., validating the model), and data from 1999 to make predictions. The process continues by extending the training period by one year at a time. Concurrently, the sample periods for validation and prediction are shifted by one year with each extension. This procedure is executed separately for each combination of text model and feature set (text and/or financial variables). We provide

details regarding training hyperparameters in Appendix A.

Evaluation metrics

We follow the current stream of research in misreporting detection and use three metrics for binary classification problems to evaluate the performance of LongFinBERT (Beneish, 1997; Larcker and Zakolyukina, 2012; Perols et al., 2017; Bao et al., 2020; Brown et al., 2020): Area Under the Receiver Operating Characteristics Curve (AUC), the Normalized Discounted Cumulative Gain at the top $k\%$ (NDCG@ $k\%$), and the Expected Cost of Misclassification (ECM). All metrics are computed based on the out-of-sample predicted probabilities for misreporting from 1999 to 2018.

The Receiver Operating Characteristics (ROC) curve, plotted in a two-dimensional space, displays the true positive rate (on the y -axis) with the false positive rate (on the x -axis) of a prediction across all possible classification thresholds (from 0 to 1). The ROC-curve therefore shows all achievable combinations of false positives and false negatives (using different classification thresholds). One typically evaluates a model’s predictions by considering the Area Under the ROC Curve (AUC), where a high value indicates that the model separates the classes well. AUC values range from 0 to 1 in which a random prediction yields an AUC of 0.5. AUC does not require a choice of classification threshold, making it a useful comprehensive and robust evaluation tool.

Normalized Discounted Cumulative Gain at the top $k\%$ (NDCG@ $k\%$) assesses predictions focusing on the top $k\%$ of the predicted scores:

$$\text{NDCG@k} = \frac{\text{DCG@k}}{\text{ideal DCG@k}}$$

$$\text{with } \text{DCG@k} = \sum_{i=1}^k \frac{rel_i}{\log_2(i+1)},$$

where rel_i equals 1 if the i th observation in the ranking list is a true positive, and 0 otherwise; and observations are ordered starting from the highest predicted probability ($i = 1$). The k^{th} observation (cutoff) is the $(1-\alpha)$ -quantile of predicted probabilities, and the normalization using ideal DCG@k leads to a value range of 0 to 1. In words, NDCG@ $k\%$ compares the top $k\%$ of observations ranked by predicted probabilities with a perfectly ordered list where all correctly predicted misreporting incidents appear before any non-true-positive cases.

Normalized to range from 0 to 1, a higher NDCG@ $k\%$ value indicates that more true misreporting incidents appear within the top- $k\%$ set ranked based on predicted probabilities. A NDCG@ $k\%$ of 1 means that the ranking of predicted probabilities for the top $k\%$ is equivalent to the ranking of observed (true) probabilities (incidents); a NDCG@ $k\%$ of 0

implies predicted ranking is the inverse ranking of the observed (true) ranking. The interpretation of NDCG@ $k\%$ similar to *precision at top- $k\%$* , which is the share of true positives in the top $k\%$ of predicted probabilities. The metric takes a perspective of e.g. the SEC that could use the model as a screening tool to identify e.g. the $k\%$ most likely candidate cases before proceeding with a costly investigation.

The final evaluation metric is Expected Cost of Misclassification (ECM) introduced by Beneish (1997) and later adopted by Perols et al. (2017). Following Perols et al. (2017), we compute ECM is computed as follows:

$$ECM = C^{FN} \times P(is_misreporting = 1) \times \frac{n^{FN}}{n^P} + C^{FP} \times P(is_misreporting = 0) \times \frac{n^{FP}}{n^N} ;$$

where C^{FN} and C^{FP} are the (normalized) estimates of the economic cost of false negative (FN) and false positive (FP) classifications; $P(is_misreporting = 1)$ is the prior probability of misreporting, $P(is_misreporting = 0) = 1 - P(is_misreporting = 1)$; n^{FN} and n^{FP} are the number of false negative cases and false positive cases, determined by the predicted probabilities and a classification threshold; n^P and n^N are the number of misreporting and non-misreporting observations in the prediction set. ECM takes economic cost of misclassification directly into account, controlled by the ratio between C^{FN} and C^{FP} .

ECM is interpreted as the expected economic cost of a classification model, so a model with lower ECM is preferable to a model with higher ECM. A perfect classifier has 0 ECM, and there is a natural upper bound determined by the cost estimates and the prior probabilities. The straightforward practical interpretation of ECM requires more assumptions than the previously presented metrics. First, an estimate of the economic cost of false positives (C^{FP}) and false negatives (C^{FN}) is required. While C^{FN} is likely much larger than C^{FP} , the exact values are difficult to measure as they e.g. include intangible components such as reputational costs and opportunity cost for investors, and the cost can differ substantially for different types of misreporting. We follow the approach of Perols et al. (2017) by focusing on *relative* cost estimates. Based on an analysis of abnormal returns after announcement of GAAP violations, Beneish (1997) find that cost of overlooking a misreporting firm (false negative) is approximately 20 to 40 times more expensive than that of a false allegation against a non-misreporting firm (false positive) in terms of investor returns. Accordingly, we choose the cost ratio $C^{FN} : C^{FP}$, of 30 : 1 (Perols et al., 2017). Furthermore, ECM depends on the prior probability of misreporting, $P(is_misreporting = 1)$, which is usually chosen by the unconditional frequency of misreporting. Based on Bell and Carcello (2000), we select $P(is_misreporting = 1) = 0.006$. Finally, ECM also relies on the choice of classification threshold above which an incident is classified as misreporting. This threshold is typically optimized during the validation phases (Beneish, 1997; Feroz et al., 2000; Perols and Lougee, 2011; Perols et al., 2017). However, this introduces more computational burdens and may lead the models to overfit the validation data more

easily. Therefore, we compute ECM with the classification threshold of 0.6 because [Perols et al. \(2017\)](#) show that this threshold generates decisions that are least sensitive to misclassified labels.

With these metrics, models with small skewness, indicating a tendency towards left-skewness in the empirical distributions, suggest a higher probability of surpassing a certain value, thus being desirable. For ECM, the interpretation of these statistics turns upside-down as we seek to minimize ECM instead of maximizing it. A model producing a small standard deviation of the metric values is preferred because it serves as a signal of a stable predictive model. Accordingly, we emphasize the preferred quantities in Table 5 by making them bold for ease of identification. From a visual standpoint, for AUC, NDCG@1%, and NDCG@5%, models with predictions skewed towards the left of the x -axis are preferred, while the opposite holds for ECM.

Evaluation procedure

To compute our set of performance metrics, we use the following procedure: For each evaluation metric and model, we (i) train the model 500 times including hyperparameter selection on the metric using the same train sample, (ii) predict out-of-sample as described in Figure 6, and (iii) compute the metric using the predictions pooled over all prediction years. This approach yields an estimate of the distribution of metrics, which is particularly valuable because model estimation and predictions can be unstable for rare events such as misreporting. ([Perols et al., 2017](#)).¹⁸

Remarks and additional robustness

As detailed in Section 3, LongFinBERT is pre-trained on 10-K filings spanning from 1994 to 2008, while the misreporting detection task employs an expanding rolling window approach with predictions made from 1999 to 2018. Concerns may arise regarding the potential for using the overlapping data period (1999 to 2008) for both LongFinBERT’s pre-training and the prediction of misreporting detection, yielding overly optimistic results for LongFinBERT, especially before 2008. However, with Table 6 in Appendix C, we dispel this concern by comparing the average metrics over time obtained by two models, utilizing: (i) only the 17 financial variables and (ii) only the LongFinBERT text-based features. The comparison reveals that, with smaller AUC, NDCG@ k %, and higher ECM, the only-text models perform weaker than the only-financial-variable models before 2008. This effectively rules out the possibility of LongFinBERT producing overly optimistic results.¹⁹

¹⁸Different from [Brown et al. \(2020\)](#) who perform bootstrapping on the test set, we generate 500 predictions for each observation in the test set for each model using a fixed train set. The variation in our approach comes from the randomness in the estimation of the classifier.

¹⁹The practice of re-using pre-training data for a downstream task is, indeed, not unusual in economic research. Particularly, [Huang et al. \(2023\)](#) employ a large collection of analyst reports to pre-train FinBERT,

	AUC			ECM			NDCG@1%			NDCG@5%		
	Only Fin. Var.	Only Text	Text and Fin. Var.	Only Fin. Var.	Only Text	Text and Fin. Var.	Only Fin. Var.	Only Text	Text and Fin. Var.	Only Fin. Var.	Only Text	Text and Fin. Var.
<i>Mean</i>	0.571	0.566	0.568	0.186	0.181	0.180	0.024	0.032	0.032	0.068	0.095	0.090
<i>Std.</i>	0.033	0.027	0.024	0.003	0.001	0.002	0.014	0.018	0.019	0.024	0.025	0.027
<i>Min</i>	0.476	0.484	0.510	0.179	0.175	0.173	0.000	0.000	0.000	0.022	0.034	0.024
<i>1% – percentile</i>	0.488	0.497	0.516	-	-	-	0.003	0.003	0.003	0.031	0.042	0.037
<i>Median</i>	0.573	0.566	0.569	0.186	0.181	0.180	0.021	0.030	0.029	0.064	0.094	0.089
<i>99% – percentile</i>	-	-	-	0.192	0.184	0.185	-	-	-	-	-	-
<i>Max</i>	0.644	0.638	0.638	0.193	0.184	0.187	0.092	0.102	0.112	0.153	0.162	0.168
<i>Skewness</i>	-0.286	-0.081	0.018	0.019	0.206	-0.229	1.264	0.999	1.036	0.719	0.229	0.211

Table 2: This table reports the descriptive statistics of the four prediction evaluation metrics AUC, ECM, NDCG@1%, and NDCG@5% obtained from three different models with 500 runs. The three models are those which contain (i) only 17 financial variables (Only Fin. Var.); (ii) only text-based features extracted from LongFinBERT (Only Text); and (iii) a concatenation of the 17 financial variables and the text-based features from LongFinBERT (Text and Fin. Var.). The bold-faced numbers indicate the best quantities among the three models evaluated by each metric. Certain statistics are omitted (-) as they are either irrelevant or redundant for comparisons with others. The prediction sample spans the period from 1999 to 2018.

Another difference between Brown et al. (2020) and our work is that we prevent cases in which the same misreporting event appears in both training and validating/predicting sets, which causes data leakage and can lead to spurious results.²⁰ This issue is initially pinpointed by Walker (2021) as a critique against Bao et al. (2020) and is rooted in the fact that a misreporting incident can last for longer than a year (on average 2.33) as described in Section 5.2. However, the AAER report is only published some time after. Therefore, we address this data leakage issue by discarding all misreporting cases that appear in both training and validating/predicting sets. Cases that present in both validating and predicting sets but not in the training sets are retained. We demonstrate that this data handling can remarkably affect the result interpretation in Section 5.5, raising questions about the validity of certain findings presented in Brown et al. (2020).

5.4 Empirical results

5.4.1 Performance of LongFinBERT compared to the 17 financial variables

The descriptive statistics of the metric values obtained from the three models under examination are presented in Table 5. The three models are those that contain (i) only the 17 financial variables, (ii) only the text-based features extracted from LongFinBERT, and (iii) a concatenation of the 17 financial variables and the text-based LongFinBERT features. We follow the current literature and use the first model, using only the 17 financial variables, as the benchmark (Brown et al., 2020). The comparisons are made using four evaluation metrics. These kernel densities are made using the metric values obtained

subsequently re-using it for a sentiment classification task. They argue that this is a valid practice because the pre-training data contains only texts without sentiment labels. The same argument is applied to our case with misreporting labels. We refer readers to footnote 11 in Huang et al. (2023) for more details.

²⁰For brevity, this issue is termed *data leakage* hereafter.




Figure 7: This figure presents the kernel densities of the four metrics, i.e., AUC (top-left), ECM (top-right), NDCG@1% (bottom-left), and NDCG@5% (bottom-right), obtained by three models. The three models are those which contain (i) only 17 financial variables (blue), (ii) only text-based features extracted from LongFinBERT (orange), and (iii) a concatenation of the 17 financial variables and the text-based features from LongFinBERT (green), respectively. The vertical red dotted line in the AUC plot implies the AUC value of 0.5 of a random prediction.

from 500 repetitions per experimental run as described in Section 5.3. Besides numerical results presented in Table 5, we also provide visual comparisons of the three models using kernel densities in Figure 7.

Despite having the highest average AUC, the model trained solely on financial variables, i.e., the benchmark model, shows an AUC of 0.488 at the 1%-percentile and a minimum of 0.476. These figures suggest that this model fails to beat random guesses at the 1% significance level. Conversely, obtaining an AUC of 0.516 and 0.510 with the same statistics, the model combining both LongFinBERT-based and financial features consistently outperforms random predictions, which the baseline model could fail as described. The performance of the model using only LongFinBERT textual features falls between those

of the benchmark model and the combined model. Additionally, it is demonstrated that the LongFinBERT-based features contribute to stabilizing the model performances, evidenced by the smallest standard deviation of the AUC belonging to the combined model. The AUC plot in Figure 7 illustrates the aforementioned analysis. Particularly, the distribution of AUC obtained by the benchmark model showcases the fatter left tail in comparison to the other two densities, implying the possibility for the model to fall below random guesses.

When evaluated by ECM, LongFinBERT strongly demonstrates its superior performance over the benchmark by significantly improving the average ECM upon integration of its text-based features with the financial variables. Particularly, the average ECM enhances from 0.186 of the benchmark to 0.180 of the combined model. This improvement is quantified as twice the magnitude of one standard deviation, calculated from a sample of 500 ECM values derived from the benchmark model. Moreover, we observe favorable statistics across the results of models utilizing LongFinBERT features, whether integrated with financial variables or not. These findings substantiate the conclusion that LongFinBERT enhances significantly the misreporting detection power, as measured by ECM. The results are easily noticeable from the ECM plot depicted in Figure 7, where the ECM distributions of both the text-only and combined models are distinctly positioned to the left, setting them apart from the ECM distribution of the benchmark model.

The superior performance of LongFinBERT still holds when assessing the models using the NDCG@ $k\%$ metrics. The models using LongFinBERT features achieve better NDCG@1% and NDCG@5% on average compared to the benchmark models. The kernel densities of NDCG@1% and NDCG@5% in Figure 7 possess more mass on their right tails and are also more right-skewed. Recall that NDCG@ $k\%$ is maximized for the optimal models, these results suggest that models using LongFinBERT text-based features have a higher likelihood of producing high NDCG@ $k\%$ values with $k \in \{1, 5\}$, thus being more favorable.

Overall, we have shown that text-based features extracted from LongFinBERT enhance the efficacy of the existing benchmark set of financial variables proposed by Dechow et al. (2011) in detecting misreporting incidents. Notably, the text-based features from LongFinBERT exhibit comparable or superior performance compared to the 17 financial variables. The contribution made by LongFinBERT remains consistent and robust across different choices of evaluation metrics.

5.4.2 Performance of LongFinBERT compared to other text models

So far, LongFinBERT has been demonstrated to boost the misreporting detection capability across various evaluation metrics. The next question that naturally arises is: “How does LongFinBERT perform in comparison with other text models in misreporting detection?”. This section attempts to provide an answer to this question by comparing LongFinBERT with LDA, Doc2Vec, and FinBERT models in the misreporting detection task. These models are selected as representative examples of each model class discussed in Section 2, and also because of their popular applications in accounting studies. In general, we compare two types of models: (i) those trained with only text-based features derived by these models (Figure 8); and (ii) those integrating both text-based features and the 17 financial variables (Figure 9).

It is important to note that FinBERT struggles to process an entire 10-K file at once due to its quadratic complexity as discussed in Section 2, thereby requiring truncation. Accordingly, we explore two truncation methods when applying FinBERT to process 10-K filings: (i) retaining only the first 8,196 tokens; and (ii) retaining only the last 8,196 tokens of each 10-K document. Different from FinBERT, which offers readily available pre-trained model weights, LDA and Doc2Vec require training from scratch before extracting their text-based features of 10-K filings. Further details on their training procedures are provided in Appendix B.

In general, LongFinBERT (marked by the blue curves) produces text-based features that outperform those extracted from other text models, both when used alone and integrated with the 17 financial variables. Notably, among the models with solely text-based features, LongFinBERT is the only text model yielding predictions that significantly beat the random guesses at the 5% significance level. The superior performance of LongFinBERT compared to LDA and Doc2Vec aligns with our expectations due to the powerful linguistic capability of LMs over non-LM models.

However, FinBERT models with truncated documents exhibit relatively weak performance. We conjecture that this result stems from the limitation of FinBERT’s ability to process the whole 10-K filings, rather than from its linguistic capabilities. Consequently, FinBERT generates document embeddings that capture incomplete information. This result highlights the importance of processing the whole document to detect misreporting utilizing 10-K filings, thereby underscoring the advantage of LongFinBERT over the text models, especially with LMs, in processing extremely lengthy accounting disclosures.

In summary, together with Section 5.4.1, our analyses demonstrate the superior performance of LongFinBERT over both financial metrics and alternative text models in detecting misreporting. This finding holds across numerous evaluation metrics, affirming its




Figure 8: This figure presents the kernel densities of the four metrics, i.e., AUC (top-left), ECM (top-right), NDCG@1% (bottom-left), and NDCG@5% (bottom-right), obtained by five models trained with only the corresponding text-based features. The five models include LongFinBERT (blue), LDA (orange), Doc2Vec (green), FinBERT with the first 8,196 tokens (red), and FinBERT with the last 8,196 tokens (violet). The vertical red dotted line in the AUC plot implies the AUC value of 0.5 of a random prediction.

robust performance.

5.5 Discussion of the data leakage problem

In this section, we discuss the data leakage problem mentioned in Section 5.3. To recall, a misreporting incident likely persists for many years. This imposes a challenge in rolling window time series research designs, where the training data directly precedes the validation/prediction data. In such scenarios, a misreporting case may extend across both the training and prediction datasets. Consequently, models that are trained on the training data are also informed about the information in the prediction data, hence mak-




Figure 9: This figure presents the kernel densities of the four metrics, i.e., AUC (top-left), ECM (top-right), NDCG@1% (bottom-left), and NDCG@5% (bottom-right), obtained by five models trained with the combination of the 17 financial variables and the corresponding text-based features. The five models include LongFinBERT (blue), LDA (orange), Doc2Vec (green), FinBERT with the first 8,196 tokens (red), and FinBERT with the last 8,196 tokens (violet). The vertical red dotted line in the AUC plot implies the AUC value of 0.5 of a random prediction.

ing the predictions spuriously optimistic. We refer to this phenomenon as *data leakage* and the misreporting cases spanning both training and prediction data as *leaked misreporting cases*. We reckon that a discussion of the data leakage issue is crucial to obtaining a correct research design, especially in research using machine learning. Research designs overlooking this issue potentially generate interpretations under spuriously optimistic results, thereby driving the conclusions astray.

Panel A of Tabel 3 provides an overview of the distribution of the misreporting cases across the years during prediction. The interpretation of the table is as follows: We use the data from 1994 to 1997, which has 47 misreporting cases, to train a model. The model, in turn, is validated with the data in 1998, and used to detect misreporting with the data

Year (in prediction)	Number of misreporting cases					
	Panel A: No delay prediction			Panel B: Delay prediction		
	Training	Validating/ Predicting	Leaked cases	Training	Validating/ Predicting	Leaked cases
1999	47	57	19			
2000	65	63	24			
2001	85	68	26	47	68	11
2002	104	64	26	65	64	12
2003	127	54	28	85	54	11
2004	142	50	31	104	50	10
2005	153	40	25	127	40	15
2006	161	34	21	142	34	16
2007	168	24	12	153	24	6
2008	174	19	9	161	19	4
2009	180	18	9	168	18	3
2010	184	18	8	174	18	4
2011	189	17	8	180	17	4
2012	194	20	8	184	20	2
2013	198	15	4	189	15	0
2014	206	12	5	194	12	0
2015	209	11	5	198	11	1
2016	213	9	3	206	9	0
2017	215	11	4	209	11	0
2018	219	7	4	213	7	1

Table 3: This table reports the number of misreporting cases appearing in the training, and validating/predicting sets in two setups: (i) the validating/predicting set immediately follows the training set (i.e., Panel A: No delay prediction); and (ii) the validating/predicting set is delayed by two years after the training set (i.e., Panel B: Delay prediction). The leftmost column presents the years in the predicting sets. Columns *Training* (*Validating/Predicting*) present the number of misreporting cases appearing in the training (validating/predicting) sets in order to make predictions of misreporting in the corresponding years. The *Overlapping* columns present the number of misreporting cases appearing in both training and validating/prediction sets in a corresponding year, thereby being considered as *leaked cases*.

in 1999. The combined validating and predicting datasets contain 57 misreporting cases, with 19 cases overlapping between the training and validating/predicting sets. Overall, leaked cases in this setting account for up to 45.66% of the total misreporting cases in the prediction set. This substantial proportion of leaked cases potentially distorts the result interpretations under inappropriate controls.

We focus on AUC and NDCG@1% to investigate the effects of the data leakage issue on result interpretations because they are the main metrics utilized by Brown et al. (2020). We train boosting algorithms using the combination of the 17 financial variables and text-based features from LongFinBERT (our model) and LDA (employed by Brown et al. (2020)). First, we compare the metrics obtained by two setups: (i) leaked misreporting cases removed (i.e., data leakage controlled), and (ii) leaked misreporting cases retained (i.e., data leakage uncontrolled). The results of this analysis are presented in Figure 10. Our results indicate that retaining leaked cases is associated with better metric values, consistently across both AUC and NDCG@1%. Moreover, the removal of leaked cases




Figure 10: This figure illustrates the model performance between two experimental settings: (i) leaked misreporting cases removed (blue), and (ii) leaked misreporting cases retained (orange). We use the combination of the 17 financial variables and the text-based features extracted from LongFinBERT and LDA to train the boosting algorithm. The distributions are aggregated over 500 runs of each setup. The prediction data sample spans the period from 1999 to 2018.

appears to have a more pronounced impact on AUC compared to NDCG@1%.

One might argue that the shift in the metric kernel densities occurs due to a change in the distribution of misreporting incidents in the predicting set when leaked cases are discarded, rather than solely as a consequence of the data leakage phenomenon. In response to the argument, we establish the second experiment in which we control the data leakage issue by delaying the prediction by two years for each expanding window. Particularly, in Figure 6, we extend window 1 in the sense that the training horizon remains unchanged but we make validation and prediction in 2000 and 2001, correspondingly. Similar extensions are applied to the consecutive windows. Accordingly, we do not discard the leaked cases to focus on the sole effects of the delayed prediction on model performances.




Figure 11: This figure illustrates the model performance between two experimental settings: (i) with a two-year delay in prediction (blue), and (ii) without delay in prediction (orange). We use the combination of the 17 financial variables and the text-based features extracted from LongFinBERT and LDA to train the boosting algorithm. The distributions are aggregated over 500 runs of each setup. The prediction data sample spans the period from 1999 to 2018.

Panel B of Table 3 reports the distributions of leaked cases under the delay-prediction setting. Although the number of misreporting cases in the training set is shifted by two years, the count in the validating/predicting set remains unchanged. Due to the two-year delay in prediction, the number of leaked cases declines remarkably, now accounting for approximately 20.37% of total misreporting cases in the prediction set.

Figure 11 compares the performance of models under two settings: (i) delay misreporting detection by two years (i.e., data leakage controlled); and (ii) no delay misreporting detection (i.e., data leakage uncontrolled). Similar to the setting of leaked-case removal, delayed prediction weakens model performance. Specifically, when model validation and prediction are postponed for two years after the training dataset, the kernel densities of AUC and NDCG@1% shift to the left. In other words, the absence of delaying prediction

is connected with an increase in model performance across the two metrics utilized by [Brown et al. \(2020\)](#).

Overall, we demonstrate that not controlling for data leakage is associated with skeptically optimistic model performance evaluated by AUC and NDCG@1%. We do that by discarding leaked misreporting cases or by delaying prediction. While we retain any conclusions of misleading result interpretation under data leakage, our discussions raise concerns about obtaining a proper research design for research in misreporting detection.

6 10-K dynamics and market reaction

In this section, using LongFinBERT, we revisit the connection between the year-over-year similarity of 10-K filings and the associated market reaction, as explored by [Brown and Tucker \(2011\)](#). Specifically, [Brown and Tucker \(2011\)](#) employ a bag-of-words method to gauge the cosine similarity between the Management’s Discussion and Analysis (MD&A) sections in the 10-K filings of the same firms across two consecutive years. Consequently, they observe a negative relationship between the MD&A section’s similarity and stock market reaction, yet emphasizing that this effect has been eroded over time. With the introduction of LongFinBERT, we are now able to revise this empirical question under the lens of a more state-of-the-art textual analysis approach, thus potentially uncovering more insights than the traditional bag-of-words methods.

6.1 Document similarity

To this end, we first utilize LongFinBERT to derive the numerical representations of the 10-K filing and the corresponding MD&A section of firm i in year t , denoted by $w_{i,t}^{10K}$ and $w_{i,t}^{MDA}$, respectively. We then gauge the year-over-year similarity between two 10-K filings or MD&A sections by computing the cosine similarity between $w_{i,t}^k$ and $w_{i,t-1}^k$ with $k \in \{10K, MDA\}$. Technically,

$$Sim_{i,t}^k = \frac{\langle w_{i,t}^k, w_{i,t-1}^k \rangle}{\|w_{i,t}^k\| \cdot \|w_{i,t-1}^k\|},$$

where $\langle x, y \rangle$ is the inner product of two vectors x and y , and $\|x\|$ is the L2 norm of vector x .²¹ Cosine similarity ranges from -1 to 1. According to [Brown and Tucker \(2011\)](#),

²¹Given two d -dimensional vectors $x = (x_1, x_2, \dots, x_d)^\top$ and $y = (y_1, y_2, \dots, y_d)^\top$, the inner product of x and y is defined as $\langle x, y \rangle = x_1y_1 + x_2y_2 + \dots + x_dy_d$; the L2 norm of x is $\|x\| = \sqrt{x_1^2 + x_2^2 + \dots + x_d^2}$.




Figure 12: This figure illustrates the evolution of the aggregated year-over-year cosine similarity of 10-K filings (blue) and the corresponding MD&A section (orange) over time. These curves present the averages of firm-wise year-over-year similarities for the corresponding documents each year. Grey areas denote the 95% confidence intervals of the point estimates for the means of the year-over-year similarities. The data sample spans the period from January 2002 to December 2016.

$Sim_{i,t}^k$ captures the linguistic similarity between the accounting disclosures of a firm in two consecutive years and is therefore a good proxy for informational similarity in these documents.

Figure 12 presents the evolution of the aggregated year-over-year similarity of 10-K filings and the corresponding MD&A section over time. Consistent with [Brown and Tucker \(2011\)](#), we observe an increasing trend of the year-over-year similarity of both 10-K filings and the MD&A section over time. The MD&A section exhibits lower year-over-year similarity than the entire 10-K filings. This finding is in line with [Dyer et al. \(2017\)](#)'s results and our expectation that 10-K filings consist of more boilerplate and generic sections such as business description (Item 1) or company properties (Item 2). Therefore, 10-K filings are less prone to modifications compared to the MD&A section, where firm managers express their assessments and viewpoints regarding the firm's operations in the current fiscal year, thus exhibiting greater variability on an annual basis. To exemplify this point, we provide a comparison of Apple Inc.'s (CIK: 0000320193) 10-K filings for 2015 and 2016 in [Appendix D](#). Notably, the year-over-year similarity in the MD&A section decreases following the financial crisis, indicating that firm managers tend to substantially modify the language used in this section between 2009 and 2010. Recall that the disclosures in 2009 (2010) discuss the firm's operations in 2008 (2009), this result is consistent with the contemporaneous unstable business conditions.

6.2 Market reaction

We now examine the relationship between the magnitude of management disclosures' modifications and market reaction. To this end, we adopt the regression study designed by [Brown and Tucker \(2011\)](#) with small modifications. In particular, we use the absolute value of the cumulative abnormal returns around the 10-K filing dates to capture the market reaction to the filing. The cumulative abnormal return of firm i , $CAR_{i,t,[0 \rightarrow h]}$, is based on the Fama-French three-factor model and cumulated from the filing date (inclusively) up to h days after the filing date as in [Jegadeesh and Wu \(2013\)](#) and [Jiang et al. \(2019\)](#). We then regress $|CAR_{i,t,[0 \rightarrow h]}|$ on $Sim_{i,t}^k$ and a set of control variables as follows,

$$|CAR_{i,t,[0 \rightarrow h]}| = \alpha + \beta Sim_{i,t}^{10K} + \gamma Sim_{i,t}^{MDA} + \sum_{j=1}^6 \delta_j Controls_{i,j,t} + \epsilon_{i,t}, \quad (1)$$

where *Controls* contains 6 variables: the market capitalization ($\log(Size)$), the additional financial information measured by new items in the Compustat database (*New_items*), the dummy variable on filing delay (*Late_filing*), and the magnitude of market reaction before the filing date ($|CAR_{i,t,[-22 \rightarrow -1]}|$, $|CAR_{i,t,[-10 \rightarrow -1]}|$ and $|CAR_{i,t,[-5 \rightarrow -1]}|$). It is worth mentioning that we can include both $Sim_{i,t}^{10K}$ and $Sim_{i,t}^{MDA}$ in the regression model (correlation 0.197). A greater (smaller) value of $Sim_{i,t}^k$ represents fewer (more) modifications in the management disclosures, thus potentially associated with weaker (stronger) market reactions toward the disclosure releases. Accordingly, their regression coefficients β and γ are expected to be negative.

We use the sample of 10-K filings and MD&A sections from January 2002 and December 2016. Instead of extracting the MD&A sections from 10-K filings, we use a dataset compiled by [Ewens et al. \(2023\)](#), which can be downloaded from the Caltech Database.²² We use the Compustat, Audit Analytics Late Filers, and Wharton Research Data Service Beta Suite databases to compute the cumulative abnormal returns and the set of control variables. After matching the text-based and accounting/financial datasets, we obtain 33,049 firm-year observations from 4,770 unique companies spanning 14 years.

Table 4 reports the regression results of Equation (1) with market reaction over three-day and ten-day horizons across three different settings: (i) solely including Sim^{MDA} ; (ii) solely including Sim^{10K} ; and (iii) including both Sim^{MDA} and Sim^{10K} . The similarity measures are standardized to have mean 0 and standard deviation 1. As anticipated, the stock market reacts significantly to the year-over-year similarities of the MD&A section of 10-K filings. Moreover, the negative coefficients over the regressions imply that more modifications in the MD&A section (low Sim^{MDA}) are associated with a higher magni-

²² Available in: <https://data.caltech.edu/records/28m72-ysc35>

	CAR - 3 days			— CAR - 10 days		
	(1)	(2)	(3)	(4)	(5)	(6)
Sim^{MDA}	-0.005*		-0.005**	-0.004*		-0.005**
	(-1.697)		(-1.995)	(-1.767)		(-2.186)
Sim^{10K}		0.011	0.013		0.018*	0.020*
		(1.062)	(1.390)		(1.677)	(1.941)
$\log(Size)$	-0.010***	-0.010***	-0.010***	-0.015***	-0.015***	-0.015***
	(-8.004)	(-8.006)	(-8.007)	(-9.224)	(-9.213)	(-9.221)
New_items	0.004	0.004	0.004	-0.001	-0.001	-0.001
	(1.341)	(1.340)	(1.340)	(-0.233)	(-0.236)	(-0.235)
$Late_filing$	0.004**	0.004**	0.004**	0.006*	0.006*	0.006*
	(2.285)	(2.294)	(2.234)	(1.674)	(1.691)	(1.702)
$ CAR_{[-5,-1]} $	0.083***	0.083***	0.083***	0.095***	0.095***	0.095***
	(12.18)	(12.23)	(12.20)	(9.094)	(9.069)	(9.092)
$ CAR_{[-10,-1]} $	0.005***	0.005***	0.005***	0.018***	0.018***	0.018***
	(3.308)	(3.337)	(3.347)	(2.434)	(2.437)	(2.437)
$ CAR_{[-22,-1]} $	0.024***	0.024***	0.024***	0.035***	0.035***	0.035***
	(18.13)	(18.29)	(18.19)	(6.374)	(6.341)	(6.365)
$No. Obs$	33,049	33,049	33,049	33,049	33,049	33,049
R^2	0.031	0.031	0.031	0.029	0.029	0.029

Table 4: This table reports the regression results of Equation (1) with $h = 3, 10$ days across three different settings, i.e., including either Sim^{MDA} or Sim^{10K} or both, always standardized to have mean 0 and standard deviation 1. The dependent variable is the absolute value of the abnormal stock returns based on the Fama-French asset pricing model cumulated from the filing dates to h days later. Inspired by Brown and Tucker (2011), we control for firm size ($\log(Size)$), additional financial information (New_items), dummy on late filing ($Late_filing$), and past market reactions ($|CAR_{[-5,-1]}|$, $|CAR_{[-10,-1]}|$, and $|CAR_{[-22,-1]}|$). The intercept (α), firm-fixed effects, and time-fixed effects are also included in the regression equation. The coefficients, Newey-West heteroskedasticity- and autocorrelation-robust t -statistics (in parentheses), the number of observations, and R^2 are reported. The data sample spans the period from January 2002 to December 2016. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

tude of market reaction. Given more up-to-date data samples from 2002 to 2016, we find that the effect of MD&A modifications measured using LongFinBERT on the stock market remains pronounced in recent periods, contradictory to Brown and Tucker (2011). The relation between the year-over-year similarity of 10-K filings and market reaction, however, is relatively weak. Intuitively, while the modifications in the MD&A section attract more attention from investors, those of 10-K filings are usually overlooked. This finding aligns with the current literature documenting the importance of the MD&A section in 10-K filings (Tavcar, 1998; SEC, 2003; Li, 2010; Muslu et al., 2015).

In general, by using LongFinBERT, we demonstrate an increasing trend in the year-over-year similarity of the MD&A section in 10-K filings, aligning with Brown and Tucker (2011). However, while Brown and Tucker (2011) conclude that investors have become indifferent to the MD&A’s modifications, our analysis indicates the contrary. This vari-

ance could stem from the different methods employed; LongFinBERT is anticipated to capture more linguistic dependencies than bag-of-words methods, revealing deeper economic implications.

7 Conclusion

Despite the current advancements in natural language processing (NLP), the application of state-of-the-art language models (LMs) in economic research remains limited. This is mainly rooted in the substantial computational demands of LMs, especially, applied to financial documents, which are typically lengthy.

Aiming to bridge that gap, we introduce LongFinBERT, a language model (LM) specifically designed for lengthy financial documents. In particular, LongFinBERT adopts the novel dilated attention mechanism developed by [Ding et al. \(2023\)](#). This mechanism allows LongFinBERT to maintain interactions between tokens distant from each other while ensuring sparse attention matrices. Consequently, that allows LongFinBERT to efficiently tackle the computational hurdles of LMs applied to long documents while maintaining the power of LMs in capturing complex linguistic structures. We adapt LongFinBERT to the financial context by pre-training it with a large corpus of 10-K filings.

We demonstrate LongFinBERT’s superiority over other text models from two angles: computational efficiency and the ability to extract meaningful textual features. First, we design a simulation study to compare LongFinBERT with four other LMs in handling long documents based on their processing time and memory usage. The results show that LongFinBERT exhibits significantly lower demands in both time and memory than the competing LMs in processing extensively long documents. Notably, LongFinBERT showcases a five-fold improvement in computational demands over LongFormer, another LM specialized for long documents.

Second, to assess the ability of LongFinBERT to extract meaningful textual features for economic applications, we employ LongFinBERT for misreporting detection using 10-K filings, which are known for their extensive length. We find that the textual features extracted from LongFinBERT enhance the ability to detect misreporting of [Dechow et al. \(2011\)](#)’s financial variables. Furthermore, LongFinBERT outperforms other text models in detecting misreporting incidents. The results remain consistent over various evaluation metrics commonly used in the misreporting detection literature. Coupled with its computational superiority, LongFinBERT has showcased its efficiency in processing financial documents.

Third, we apply LongFinBERT to examine the year-over-year similarities of 10-K filings

and the corresponding Management’s Discussion and Analysis (MD&A) section. In line with the current literature, we reveal a trend towards increased similarity over time in both disclosures, with 10-K filings exhibiting greater similarity than the MD&A section. Complementary to the extant literature, we document a significant relationship between the market reaction and the modifications of the MD&A section but not 10-K filings. This result suggests that investors may respond more sensitively to alterations within the MD&A section, reaffirming the importance of the MD&A section in 10-K filings.

A potential limitation of LongFinBERT, however, resides in the interpretations of its extracted textual features, a challenge shared by many LMs. Nevertheless, we anticipate that LongFinBERT can be effectively utilized to extract interpretable topics from financial documents, which currently rely on the word-count-based Latent Dirichlet Allocation model. With the demonstrated linguistic capability, LongFinBERT is expected to produce more accurate topical information in economic and financial documents.

References

- Ash, E. and Hansen, S. (2023). Text algorithms in economics, *Annual Review of Economics* **15**: 659–688.
- Bae, J., Yu Hung, C. and van Lent, L. (2023). Mobilizing text as data, *European Accounting Review* **32**(5): 1085–1106.
- Bao, Y., Ke, B., Li, B., Yu, Y. J. and Zhang, J. (2020). Detecting accounting fraud in publicly traded us firms using a machine learning approach, *Journal of Accounting Research* **58**(1): 199–235.
- Bell, T. B. and Carcello, J. V. (2000). A decision aid for assessing the likelihood of fraudulent financial reporting, *Auditing: A Journal of Practice & Theory* **19**(1): 169–184.
- Beltagy, I., Peters, M. E. and Cohan, A. (2020). Longformer: The long-document transformer, *arXiv preprint arXiv:2004.05150*.
- Beneish, M. D. (1997). Detecting GAAP violation: Implications for assessing earnings management among firms with extreme financial performance, *Journal of Accounting and Public Policy* **16**(3): 271–309.
- Beneish, M. D. (1999). The detection of earnings manipulation, *Financial Analysts Journal* **55**(5): 24–36.
- Bertomeu, J., Cheynel, E., Floyd, E. and Pan, W. (2021). Using machine learning to detect misstatements, *Review of Accounting Studies* **26**: 468–519.
- Bochkay, K., Brown, S. V., Leone, A. J. and Tucker, J. W. (2023). Textual analysis in accounting: What’s next?, *Contemporary Accounting Research* **40**(2): 765–805.
- Brown, N. C., Crowley, R. M. and Elliott, W. B. (2020). What are you saying? Using *topic* to detect financial misreporting, *Journal of Accounting Research* **58**(1): 237–291.
- Brown, S. V. and Tucker, J. W. (2011). Large-sample evidence on firms’ year-over-year MD&A modifications, *Journal of Accounting Research* **49**(2): 309–346.
- Cao, S., Jiang, W., Yang, B. and Zhang, A. L. (2023). How to talk when a machine is listening: Corporate disclosure in the age of AI, *The Review of Financial Studies* **36**(9): 3603–3642.
- Cazier, R. A. and Pfeiffer, R. J. (2016). Why are 10-K filings so long?, *Accounting Horizons* **30**(1): 1–21.
- Cecchini, M., Aytug, H., Koehler, G. J. and Pathak, P. (2010). Detecting Management Fraud in Public Companies, *Management Science* **56**(7): 1146–1160.
- Chen, C. Y.-H., Fengler, M. R., Härdle, W. K. and Liu, Y. (2022). Media-expressed tone, option characteristics, and stock return predictability, *Journal of Economic Dynamics and Control* **134**: 104290.
- Chen, M. A., Wu, Q. and Yang, B. (2019). How valuable is FinTech innovation?, *The Review of Financial Studies* **32**(5): 2062–2106.
- Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system, *Proceedings of the 22nd ACM KDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794.
- Cho, H. and Muslu, V. (2021). How do firms change investments based on MD&A disclosures of peer firms?, *The Accounting Review* **96**(2): 177–204.

- Dechow, P. M., Ge, W., Larson, C. R. and Sloan, R. G. (2011). Predicting material accounting misstatements, *Contemporary Accounting Research* **28**(1): 17–82.
- Dechow, P. M., Sloan, R. G. and Sweeney, A. P. (1996). Causes and consequences of earnings manipulation: An analysis of firms subject to enforcement actions by the SEC, *Contemporary Accounting Research* **13**(1): 1–36.
- Devlin, J., Chang, M.-W., Lee, K. and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Long and Short Papers)*, Vol. 1, Association for Computational Linguistics, Minneapolis, Minnesota, pp. 4171–4186.
- Ding, J., Ma, S., Dong, L., Zhang, X., Huang, S., Wang, W. and Wei, F. (2023). LongNet: Scaling transformers to 1,000,000,000 tokens, *Proceedings of the 10th International Conference on Learning Representations*.
- Dyer, T., Lang, M. and Stice-Lawrence, L. (2017). The evolution of 10-K textual disclosure: Evidence from Latent Dirichlet Allocation, *Journal of Accounting and Economics* **64**(2-3): 221–245.
- Ewens, M., Peters, R. and Wang, S. (2023). Measuring intangible capital with market prices, *Working Paper*.
- Fengler, M. and Phan, T. M. (2023). *A Topic Model for 10-K Management Disclosures*, School of Economics and Political Science, Department of Economics.
- Feroz, E. H., Kwon, T. M., Pastena, V. S. and Park, K. (2000). The efficacy of red flags in predicting the sec's targets: an artificial neural networks approach, *Intelligent Systems in Accounting, Finance & Management* **9**(3): 145–157.
- Frankel, R., Jennings, J. and Lee, J. (2022). Disclosure sentiment: Machine learning vs. dictionary methods, *Management Science* **68**(7): 5514–5532.
- Freund, Y., Schapire, R. E. et al. (1996). Experiments with a new boosting algorithm, *the Thirteenth International Conference on Machine Learning*, Vol. 96, Citeseer, pp. 148–156.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine, *Annals of Statistics* pp. 1189–1232.
- Friedman, J. H. (2002). Stochastic gradient boosting, *Computational Statistics & Data Analysis* **38**(4): 367–378.
- Gentzkow, M., Kelly, B. and Taddy, M. (2019). Text as data, *Journal of Economic Literature* **57**(3): 535–574.
- Griffin, P. A. (2003). Got information? Investor response to form 10-K and form 10-Q EDGAR filings, *Review of Accounting Studies* **8**: 433–460.
- Harris, Z. S. (1954). Distributional structure, *Word* **10**(2-3): 146–162.
- He, P., Liu, X., Gao, J. and Chen, W. (2021). Deberta: Decoding-enhanced bert with disentangled attention, *International Conference on Learning Representations*.
- Huang, A. H., Wang, H. and Yang, Y. (2023). Finbert: A large language model for extracting information from financial text, *Contemporary Accounting Research* **40**(2): 806–841.
- Jegadeesh, N. and Wu, D. (2013). Word power: A new approach for content analysis, *Journal of Financial Economics* **110**(3): 712–729.

- Jiang, F., Lee, J., Martin, X. and Zhou, G. (2019). Manager sentiment and stock returns, *Journal of Financial Economics* **132**(1): 126–149.
- Larcker, D. F. and Zakolyukina, A. A. (2012). Detecting deceptive discussions in conference calls, *Journal of Accounting Research* **50**(2): 495–540.
- Le, Q. and Mikolov, T. (2014). Distributed representations of sentences and documents, *Proceedings of the 31st International Conference on Machine Learning*, Vol. 32 of *Proceedings of Machine Learning Research*, PMLR, Beijing, China, pp. 1188–1196.
- Li, F. (2008). Annual report readability, current earnings, and earnings persistence, *Journal of Accounting and Economics* **45**(2-3): 221–247.
- Li, F. (2010). The information content of forward-looking statements in corporate filings—a naïve Bayesian machine learning approach, *Journal of Accounting Research* **48**(5): 1049–1102.
- Li, K., Mai, F., Shen, R. and Yan, X. (2021). Measuring corporate culture using machine learning, *The Review of Financial Studies* **34**(7): 3265–3315.
- Linnainmaa, S. (1976). Taylor expansion of the accumulated rounding error, *BIT Numerical Mathematics* **16**(2): 146–160.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L. and Stoyanov, V. (2019). Roberta: A robustly optimized BERT pretraining approach, *arXiv preprint arXiv:1907.11692*.
- Loughran, T. and McDonald, B. (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-ks, *The Journal of Finance* **66**(1): 35–65.
- Loughran, T. and McDonald, B. (2016). Textual analysis in accounting and finance: A survey, *Journal of Accounting Research* **54**(4): 1187–1230.
- Loukas, L., Fergadiotis, M., Chalkidis, I., Spyropoulou, E., Malakasiotis, P., Androutsopoulos, I. and Paliouras, G. (2022). FiNER: Financial numeric entity recognition for XBRL tagging, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Dublin, Ireland, pp. 4419–4431.
- Lusa, L. et al. (2017). Gradient boosting for high-dimensional prediction of rare events, *Computational Statistics & Data Analysis* **113**: 19–37.
- Manning, C. and Schütze, H. (1999). *Foundations of statistical natural language processing*, MIT press.
- Mikolov, T., Chen, K., Corrado, G. and Dean, J. (2013). Efficient estimation of word representations in vector space, *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Miranda-Belmonte, H. U., Muñoz-Sánchez, V. and Corona, F. (2023). Word embeddings for topic modeling: An application to the estimation of the economic policy uncertainty index, *Expert Systems with Applications* **211**: 118499.
- Muslu, V., Radhakrishnan, S., Subramanyam, K. and Lim, D. (2015). Forward-looking MD&A disclosures and the information environment, *Management Science* **61**(5): 931–948.
- Park, J. and Cho, S. (2023). Incorporation of company-related factual knowledge into pre-trained language models for stock-related spam tweet filtering, *Expert Systems with Applications* **234**: 121021.

- Perols, J. L., Bowen, R. M., Zimmermann, C. and Samba, B. (2017). Finding needles in a haystack: Using data analytics to improve fraud prediction, *The Accounting Review* **92**(2): 221–245.
- Perols, J. L. and Lougee, B. A. (2011). The relation between earnings management and financial statement fraud, *Advances in Accounting* **27**(1): 39–53.
- Price, S. M., Doran, J. S., Peterson, D. R. and Bliss, B. A. (2012). Earnings conference calls and stock returns: The incremental informativeness of textual tone, *Journal of Banking & Finance* **36**(4): 992–1011.
- Rawte, V., Gupta, A. and Zaki, M. J. (2020). A comparative analysis of temporal long text similarity: Application to financial documents, *Workshop on Mining Data for Financial Applications*, Springer, pp. 77–91.
- Řehůřek, R. and Sojka, P. (2010). Software Framework for Topic Modelling with Large Corpora, *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, ELRA, Valletta, Malta, pp. 45–50. <http://is.muni.cz/publication/884893/en>.
- Richardson, S. A., Sloan, R. G., Soliman, M. T. and Tuna, I. (2005). Accrual reliability, earnings persistence and stock prices, *Journal of Accounting and Economics* **39**(3): 437–485.
- Robbins, H. and Monro, S. (1951). A stochastic approximation method, *The Annals of Mathematical Statistics* pp. 400–407.
- Rogers, A., Kovaleva, O. and Rumshisky, A. (2021). A primer in BERTology: What we know about how BERT works, *Transactions of the Association for Computational Linguistics* **8**: 842–866.
- Schimanski, T., Reding, A., Reding, N., Bingler, J., Kraus, M. and Leippold, M. (2024). Bridging the gap in ESG measurement: Using NLP to quantify environmental, social, and governance communication, *Finance Research Letters* **61**: 104979.
- SEC (2003). Interpretation: Commission guidance regarding management’s discussion and analysis of financial condition and results of operations, *Securities Act Release* (33-8350): 34–48960.
- Soutner, D. and Müller, L. (2013). Application of LSTM Neural Networks in Language Modelling, in D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, I. Habernal and V. Matoušek (eds), *Text, Speech, and Dialogue*, Vol. 8082, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 105–112.
- Spaerck Jones, K. (1972). A statistical interpretation of term specificity and its application in retrieval, *Journal of Documentation* **28**(1): 11–21.
- Tavcar, L. R. (1998). Make the md&a more readable, *The CPA Journal* **68**(1): 10.
- Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market, *The Journal of Finance* **62**(3): 1139–1168.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u. and Polosukhin, I. (2017). Attention is all you need, *Advances in Neural Information Processing Systems*, Vol. 30, Curran Associates, Inc.
- Walker, S. (2021). Critique of an article on machine learning in the detection of accounting fraud, *Econ Journal Watch* **18**(2): 61.

- Wang, T. and Cho, K. (2016). Larger-Context Language Modelling with Recurrent Neural Network, *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* pp. 1319–1329.
- West, K. and McCracken, M. (1998). Regression-based tests of predictive ability, *International Economic Review* **39**(4): 817–40.
- Whiting, D. G., Hansen, J. V., McDonald, J. B., Albrecht, C. and Albrecht, W. S. (2012). Machine learning methods for detecting patterns of management fraud, *Computational Intelligence* **28**(4): 505–527.
- Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D. and Mann, G. (2023). BloombergGPT: A Large Language Model for Finance.
- Xiong, Y., Zeng, Z., Chakraborty, R., Tan, M., Fung, G., Li, Y. and Singh, V. (2021). Nyströmformer: A nyström-based algorithm for approximating self-attention, *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35, pp. 14138–14148.
- Yang, H., Liu, X.-Y. and Wang, C. D. (2023). FinGPT: Open-source financial large language models, *FinLLM Symposium at IJCAI 2023*.
- Yang, X., Zhang, C., Sun, Y., Pang, K., Jing, L., Wa, S. and Lv, C. (2023). FinChain-BERT: A high-accuracy automatic fraud detection model based on NLP methods for financial scenarios, *Information* **14**(9): 499.
- Yi, Z., Cao, X., Pu, X., Wu, Y., Chen, Z., Khan, A. T., Francis, A. and Li, S. (2023). Fraud detection in capital markets: A novel machine learning approach, *Expert Systems with Applications* **231**: 120760.
- Zaheer, M., Guruganesh, G., Dubey, K. A., Ainslie, J., Alberti, C., Ontanon, S., Pham, P., Ravula, A., Wang, Q., Yang, L. et al. (2020). Big bird: Transformers for longer sequences, *Advances in Neural Information Processing Systems* **33**: 17283–17297.

A Hyperparameters of the gradient boosting classifiers

Hyperparameter	Value	Description
<code>n_estimators</code>	250	The number of rounds for boosting
<code>booster</code>	<code>gbtree</code>	The base booster. In our case, <code>gbtree</code> is a decision tree
<code>max_depth</code>	2	The maximum depth of a tree
<code>objective</code>	<code>binary:logistic</code>	The loss function. In our case, the binary cross-entropy loss
<code>colsample_bytree</code>	0.5	The subsample ratio of columns when constructing each tree
<code>reg_alpha</code>	0.1	The L1 regularization term on tree weights (a higher value makes the model more conservative)
<code>subsample</code>	0.9	The subsample ratio of the training observations
<code>learning_rate</code>	0.05	The step size shrinkage used in update to prevents overfitting (a higher value makes the model more conservative)
<code>tree_method</code>	<code>gpu_hist</code>	The tree construction algorithm used in XGBoost
<code>min_child_weight</code>	10	The minimum sum of instance weight (Hessian) needed in a child
<code>max_bin</code>	64	The maximum number of discrete bins to bucket continuous features (only use if <code>tree_method = hist</code> or <code>gpu_hist</code>)
<code>early_stopping_round</code>	25	This parameter will stop training if the validation metric is not improving after the last early stopping round

Table 5: This table presents the hyperparameters used to train the gradient boosting classifiers described in Section 5.3.

B Training details of the LDA and Doc2Vec models

Different from LMs, which typically require intense computational costs, non-LM models such as LDA and Doc2Vec can be efficiently trained without GPUs. We, therefore, adopt the rolling window approach to train the LDA and Doc2Vec models to capture the linguistic variations of 10-K filings, following Brown et al. (2020). We utilize the Python library *gensim* (Řehůřek and Sojka, 2010) to obtain efficient training of both models.

For training the LDA model, we implement an expanding window strategy similar to the setup for training the classifiers in Figure 6. The initial LDA model is trained on the period from 1994 to 1997. Subsequently, the model is employed to infer latent topics within the 10-K filings in 1998. These two steps constitute the first window. Transitioning to the second window, we update the LDA model trained in the first window with the 1998 filings. Subsequently, we utilize this updated model to infer topics from the 1999 10-K documents. This iterative process continues until we reach the year 2018 in terms of topic inference. Suggesting by Brown et al. (2020), we fix the number of topics to 31. Other hyperparameters are set by default following *gensim*.

Unlike LDA which *gensim* offers a continual training mechanism,²³ Doc2Vec requires

²³This feature enables a fully trained model to be revived and continue training with new, unseen data.

training from scratch upon the arrival of unseen data. Consequently, the expanding-window strategy potentially renders training Doc2Vec computationally challenging, especially when the training window spans several years. Therefore, we apply the rolling-window approach to train Doc2Vec in which we maintain a two-year training window. This method significantly reduces computational time while effectively capturing linguistic changes over time, leveraging a substantial training dataset of approximately 270 million words on average. As for the training hyperparameters, we set the context window to ten adjacent words, discard words appearing in less than 15 filings, and assign document vectors with 300 dimensions.

C Performance of misreporting detection models over time

Year	AUC		ECM		NDCG@1%		NDCG@5%	
	Fin. Var.	Text	Fin. Var.	Text	Fin. Var.	Text	Fin. Var.	Text
1999	0.511	0.505	0.180	0.191	0.013	0.009	0.029	0.028
2000	0.584	0.541	0.184	0.180	0.042	0.026	0.079	0.051
2001	0.611	0.579	0.180	0.180	0.022	0.023	0.045	0.048
2002	0.599	0.586	0.184	0.180	0.025	0.010	0.072	0.036
2003	0.616	0.585	0.188	0.193	0.013	0.012	0.044	0.041
2004	0.635	0.609	0.180	0.180	0.038	0.009	0.071	0.026
2005	0.576	0.507	0.182	0.181	0.014	0.001	0.038	0.008
2006	0.537	0.459	0.193	0.180	0.007	0.001	0.022	0.003
2007	0.634	0.552	0.180	0.180	0.016	0.003	0.050	0.013
2008	0.602	0.592	0.185	0.180	0.005	0.015	0.020	0.043
2009	0.586	0.456	0.180	0.180	0.005	0.001	0.019	0.009
2010	0.531	0.374	0.180	0.180	0.008	0.002	0.027	0.005
2011	0.601	0.503	0.184	0.180	0.016	0.005	0.044	0.016
2012	0.639	0.59	0.185	0.179	0.017	0.023	0.043	0.050
2013	0.694	0.647	0.188	0.176	0.009	0.017	0.031	0.056
2014	0.709	0.543	0.196	0.180	0.020	0.000	0.057	0.002
2015	0.740	0.539	0.171	0.180	0.031	0.005	0.100	0.015
2016	0.603	0.631	0.19	0.180	0.018	0.003	0.038	0.012
2017	0.557	0.544	0.181	0.180	0.006	0.000	0.020	0.001
2018	0.192	0.456	0.206	0.180	0.000	0.000	0.000	0.001

Table 6: This table reports the means of four metrics obtained by two models using (i) only the 17 financial variables (Fin. Var.), and (ii) only LongFinBERT text-based features (Text). Each entry in this table is the average metric value in each year over 500 runs. The prediction sample spans the period from 1999 to 2018.

D Excerpts from Apple’s 2015 and 2016 10-K Filings

	2015	2016	Similarity
	Company Background	Company Background	
<i>Item 1 - Business</i>	The Company designs, manufactures and markets mobile communication and media devices, personal computers and portable digital music players, and sells a variety of related software, services, accessories, networking solutions and third-party digital content and applications. The Company's products and services include iPhone®, iPad®, Mac®, iPod®, Apple Watch®, Apple TV®, a portfolio of consumer and professional software applications, iOS, macOS™, watchOS® and tvOS™ operating systems, iCloud®, Apple Pay® and a variety of accessory, service and support offerings. The Company sells and delivers digital content and applications through the iTunes Store®, App Store®, Mac App Store, TV App Store, iBooks Store™ and Apple Music® (collectively "Internet Services"). The Company sells its products worldwide through its retail stores, online stores and direct sales force, as well as through third-party cellular network carriers, wholesalers, retailers and value-added resellers. In addition, the Company sells a variety of third-party Apple compatible products, including application software and various accessories through its retail and online stores. The Company sells to consumers, small and mid-sized businesses and education, enterprise and government customers. The Company's fiscal year is the 52 or 53-week period that ends on the last Saturday of September. The Company is a California corporation established in 1977.	The Company designs, manufactures and markets mobile communication and media devices, personal computers and portable digital music players, and sells a variety of related software, services, accessories, networking solutions and third-party digital content and applications. The Company's products and services include iPhone®, iPad®, Mac®, iPod®, Apple Watch®, Apple TV®, a portfolio of consumer and professional software applications, iOS, macOS™, watchOS® and tvOS™ operating systems, iCloud®, Apple Pay® and a variety of accessory, service and support offerings. The Company sells and delivers digital content and applications through the iTunes Store®, App Store®, Mac App Store, TV App Store, iBooks Store™ and Apple Music® (collectively "Internet Services"). The Company sells its products worldwide through its retail stores, online stores and direct sales force, as well as through third-party cellular network carriers, wholesalers, retailers and value-added resellers. In addition, the Company sells a variety of third-party Apple compatible products, including application software and various accessories through its retail and online stores. The Company sells to consumers, small and mid-sized businesses and education, enterprise and government customers. The Company's fiscal year is the 52 or 53-week period that ends on the last Saturday of September. The Company is a California corporation established in 1977.	1.000 (whole Item: 0.828)
<i>Item 2 - Properties</i>	The Company's headquarters are located in Cupertino, California. As of September 26, 2015, the Company owned or leased 25.6 million square feet of building space, primarily in the U.S. The Company also owned or leased building space in various locations, including throughout Europe, China, Singapore and Japan. Of the total owned or leased building space 18.5 million square feet was leased building space, which includes approximately 5.3 million square feet related to retail store space. Additionally, the Company owns a total of 1,757 acres of land in various locations.	The Company's headquarters are located in Cupertino, California. As of September 24, 2016, the Company owned 7.1 million square feet and leased 22.3 million square feet of building space, primarily in the U.S. Additionally, the Company owned a total of 2,583 acres of land primarily in the U.S.	0.708 (whole Item: 0.715)
	<i>Fiscal 2015 highlights</i>	<i>Fiscal 2016 highlights</i>	
<i>Item 7 - MD&A</i>	Net sales rose 28% or \$50.9 billion during 2015 compared to 2014, driven by a 52% year-over-year increase in iPhone® net sales. iPhone net sales and unit sales in 2015 increased in all of the Company's reportable operating segments. The Company also experienced year-over-year net sales increases in Mac®, Services and Other Products. Apple Watch®, which launched during the third quarter of 2015, accounted for more than 100% of the year-over-year growth in net sales of Other Products. Net sales growth during 2015 was partially offset by the effect of weakness in most foreign currencies relative to the U.S. dollar and lower iPad® net sales. Total net sales increased in each of the Company's reportable operating segments, with particularly strong growth in Greater China where year-over-year net sales increased 84%.	Net sales declined 8% or \$18.1 billion during 2016 compared to 2015, primarily driven by a year-over-year decrease in iPhone net sales and the effect of weakness in most foreign currencies relative to the U.S. dollar, partially offset by an increase in Services. In April 2016, the Company announced a significant increase to its capital return program by raising the expected total size of the program from \$200 billion to \$250 billion through March 2018. This included increasing its share repurchase authorization from \$140 billion to \$175 billion and raising its quarterly dividend from \$0.52 to \$0.57 per share beginning in May 2016. During 2016, the Company spent \$29.0 billion to repurchase shares of its common stock and paid dividends and dividend equivalents of \$12.2 billion. Additionally, the Company issued \$23.9 billion of U.S. dollar-denominated term debt and A\$1.4 billion of Australian dollar-denominated term debt during 2016.	0.611 (whole Item: 0.782)

Table 7: This table displays parts of the 10-K filings and the corresponding MD&A sections of Apple Inc. (CIK: 0000320193) in 2015 and 2016.