LUISS Guido Carli

LIBERA UNIVERSITÀ INTERNAZIONALE DEGLI STUDI SOCIALI

Corso di laurea in Corporate Finance

Cattedra in Financial Statement Analysis

AI-DRIVEN DETECTION OF CORPORATE FRAUD IN FINANCIAL STATEMENTS

RELATORE

Prof.

Saverio Bozzolan

CORRELATRICE

Prof.ssa

Barbara Sveva Magnanelli

CANDIDATO

Lorenzo Beneduce

matricola 773961

Anno Accademico 2024/2025

TABLE OF CONTENTS

INTRODUCTION

Corporate fraud in financial statements represents a persistent challenge for stakeholders, regulators, and investors alike. As financial reporting remains central to business transparency, the manipulation of financial data poses a serious threat to market integrity and corporate governance. Traditional detection methods, while valuable, often lack the efficiency and accuracy required to detect sophisticated fraud schemes.

In recent years, artificial intelligence (AI) — and in particular, Natural Language Processing (NLP) — has emerged as a promising tool in identifying fraudulent patterns hidden within corporate disclosures. However, its practical application to financial statement analysis remains a developing field, with significant potential for improving fraud detection mechanisms.

This thesis aims to explore and evaluate the effectiveness of AI-based techniques, especially NLP, in detecting corporate fraud through the analysis of financial statement narratives. By examining both theoretical frameworks and real-world applications, the research seeks to bridge the gap between traditional methods and advanced AI-driven approache.

The study relies on a dataset of financial statements from various sources, including SEC filings and European corporate reports. NLP techniques such as sentiment analysis, anomaly detection, and textual pattern recognition are applied to identify inconsistencies indicative of potential fraud.

The thesis is structured as follows: Chapter 1 presents a review of the literature on corporate fraud, traditional detection methods, and the emerging use of AI. Chapter 2 outlines the methodology, including dataset preparation and the application of NLP models. Chapter 3 discusses the results of the analysis, comparing AI-driven findings with traditional approaches. Chapter 4 concludes with critical reflections on the implications, limitations, and future directions of AI-based fraud detection in accounting practices.

CHAPTER 1 - ACCOUNTING FRAUD: A LITERATURE REVIEW

Accounting fraud is a complex and multifaceted phenomenon that challenges both theoretical definitions and practical frameworks. Traditionally, accounting fraud has been understood as the deliberate misrepresentation of financial information by corporate actors to deceive stakeholders and secure unjust gains. However, this definition is increasingly viewed as insufficient to capture the breadth and nuance of fraudulent behavior in contemporary organizational contexts.

According to recent literature, the prevailing model adopted by professionals—such as the fraud triangle promoted by the Association of Certified Fraud Examiners (Examiners, Report to the Nations on Occupational Fraud and Abuse, 2014)—offers a narrow conceptualization of fraud, emphasizing individual motivations like pressure, opportunity, and rationalization (Tuck & Hamilton, 2023).

This model, while useful, has been critiqued for failing to account for systemic and contextual factors that shape fraudulent practices. Fraud should instead be seen as a socially constructed concept, influenced by institutional norms, regulatory environments, and power dynamics. It may manifest not only through illegal acts but also through behaviors that, while technically legal, are ethically questionable or misleading in nature.

Furthermore, empirical evidence confirms that corporate governance failures, particularly involving CEOs and other high-ranking executives, are commonly associated with fraudulent activities. These actors often exploit weaknesses in oversight mechanisms, such as audit committees, to manipulate financial results or conceal irregularities. As a result, regulatory bodies and enforcement agencies increasingly advocate for the involvement of forensic accountants during investigations, recognizing their unique ability to combine financial expertise with investigative rigor (Tutino & Merlo, Accounting fraud: A literature review, 2019).

In sum, accounting fraud transcends simple definitions, requiring a contextual and interdisciplinary approach that considers social, legal, and organizational dimensions. Recognizing its multifaceted nature is essential to developing more robust detection and prevention strategies in both research and practice. To lay the groundwork for this analysis, the following section categorizes the main types of corporate accounting fraud, illustrating how they manifest in real-world scenarios.

1.1.1 Types of Corporate Fraud

Revenue Manipulation

Building on the multifaceted nature of accounting fraud previously discussed, it is essential to classify the most common forms it can take in practice. While definitions and classifications may vary across jurisdictions, most academic and professional literature converges on three major typologies of corporate fraud: revenue manipulation, cost falsification, and asset misappropriation. These categories, though distinct in mechanics and intent, share a unifying trait: the deliberate distortion of financial information to deceive stakeholders and achieve undue advantage.

Revenue manipulation is one of the most pervasive and strategically deployed forms of financial statement fraud. It involves prematurely recognizing revenue, inflating sales figures, or even recording fictitious transactions. This practice often stems from intense market pressure to meet earnings targets or analyst expectations, particularly in high-growth firms where revenue signals valuation more strongly than net income (Zhang, 2006).

Research by Zhang (2006) identifies that firms with high operating margins, low book-to-market ratios, and frequent analyst sales forecasts are significantly more likely to engage in aggressive revenue recognition. These firms exploit the flexibility in accrual accounting, such as manipulating accounts receivable or unearned revenue, to defer cost recognition and artificially boost short-term performance (Zhang, 2006).

Dechow and Schrand (2004) report that over 70% of SEC enforcement actions for accounting violations involve overstated revenues. Techniques like "channel stuffing", "bill-and-hold sales", and "trade loading" are among the methods identified in empirical analyses, and they are frequently facilitated by management's discretionary judgment under accounting standards (Dechow & Schrand, 2004).

Cost Falsification

In contrast to revenue inflation, cost falsification seeks to overstate a firm's profitability by minimizing the appearance of expenses. This can be achieved by delaying expense recognition, improperly capitalizing costs that should be expensed, or underestimating future liabilities. These methods aim to inflate net income without necessarily affecting revenue figures, and are typically used when firms want to present improved operating margins without manipulating top-line performance (Tutino & Merlo, Accounting fraud: A literature review, 2019).

Common signals of cost-related fraud include unusual patterns in depreciation, amortization, and research and development spending. Managers may engage in such tactics to meet internal targets or obtain performance-based bonuses, particularly when their compensation is linked to net income growth. The challenge with cost falsification lies in its detectability: because it often involves classification or timing judgments, auditors and regulators may struggle to determine intent versus legitimate accounting discretion.

Moreover, firms with weak governance structures—such as lack of board independence or ineffective audit committees—are more prone to engage in such manipulation (Reurink, 2018). These conditions create a permissive environment where accounting practices can be twisted without immediate detection or accountability (Reurink, 2018).

Asset Misappropriation

Asset misappropriation, although often associated with lower-level employees, represents the most frequent form of occupational fraud according to the Association of Certified Fraud Examiners (Examiners, Report to the Nations on Occupational Fraud and Abuse, 2014), accounting for over 85% of fraud cases. It encompasses the theft or unauthorized use of an organization's assets, such as cash, inventory, equipment, or proprietary information.

While typically involving smaller financial amounts than financial statement fraud, the cumulative impact can be devastating—especially in firms with lax internal controls. The ACFE’s "Fraud Tree" categorizes asset misappropriation into schemes such as skimming, billing fraud, payroll fraud, and expense reimbursement fraud. These acts are often concealed by altering supporting documents, creating fictitious vendors or employees, or colluding with external parties.

Research highlights that misappropriation is not limited to frontline employees; in fact, executives and managers, when granted excessive autonomy, may be even more effective in disguising these acts (Said, 2015). Root causes include financial pressure, rationalization, and the perceived low risk of detection, all of which align with the classic fraud triangle model proposed by (Cressey, 1953). Asset misappropriation is particularly prevalent in the banking and financial services sectors, where access to liquid assets is high and opportunities for manipulation are frequent (Bakri, Mohamed, & Said, 2017).

Collectively, these three fraud typologies reveal not only the diversity of deceptive tactics available within the accounting process but also the importance of contextual and organizational factors in enabling such practices. As the next section will show, real-world fraud cases such as Enron, Parmalat, and Wirecard serve as archetypal examples that illustrate the practical application—and often combination—of these fraudulent mechanisms.

1.1.1 Known Case Studies

Enron

To ground the typologies of accounting fraud in real-world scenarios, this section explores three of the most impactful corporate fraud cases in recent history: Enron, Parmalat, and Wirecard. These cases exemplify not only the implementation of revenue manipulation, cost falsification, and asset misappropriation, but also how weak governance, audit failures, and regulatory inertia can enable fraudulent behavior to persist over time. Analyzing these events offers insight into systemic vulnerabilities and informs the development of more effective fraud detection frameworks.

Enron Corporation, a former American energy company, collapsed in 2001 after revealing one of the most elaborate financial frauds ever recorded. At the heart of the fraud was the use of Special Purpose Entities (SPEs) to conceal debt and inflate earnings, giving the illusion of consistent growth and financial health (Dibra, 2016). These off-balance-sheet vehicles allowed Enron to hide liabilities while maintaining high stock prices, supported by complex derivatives and mark-to-market accounting. The company's CFO, Andrew Fastow, played a central role in orchestrating these transactions.

The failure of Arthur Andersen, Enron's auditor, to identify or report the fraud was a critical factor in the scandal. Andersen's dual role as both auditor and consultant compromised its independence and contributed to the falsification of Enron's financial statements (Sapelli, 2005). The resulting public backlash led to the dissolution of Andersen and catalyzed the creation of the Sarbanes-Oxley Act (SOX) in 2002, which aimed to restore confidence in corporate reporting and enhance auditor accountability.

Parmalat

The Parmalat scandal, uncovered in 2003, involved the falsification of financial statements to hide approximately €14 billion in debt. Executives at the Italian multinational used fictitious bank accounts and a network of offshore entities to fabricate assets and profits (Dibra, 2016). Central to the fraud was the now-infamous "Buco Nero" (black hole), a hidden financial deficit masked through circular financing and fraudulent documentation (Sapelli, 2005).

The company's internal controls were weak, and its external auditors failed to verify even the most basic financial information. At one point, Deloitte and Grant Thornton accepted a fax as confirmation of nearly €4 billion in bank deposits from a Cayman Islands subsidiary—deposits that did not exist (Allegrini, Di Gennaro, Healy, & Palepu, 2023). Governance failures played a decisive role: the dual role of Calisto Tanzi as CEO and chairman concentrated decision-making power, while family dominance over the board limited oversight and accountability.

Wirecard

Wirecard AG, once a fintech flagship in Germany, declared insolvency in 2020 after admitting that €1.9 billion in supposed cash balances did not exist. The fraud involved fabricating revenues via non-existent third-party acquirers, especially in Asia, and was supported by a consistent pattern of misleading disclosures (Beerbaum, 2022). Investigations revealed that the company used round-tripping transactions and false documentation to simulate profitability and cash flow (Möllers, McCrum, & Alderman, 2020).

Despite repeated warnings from whistleblowers and journalists, the German financial regulator BaFin failed to act and instead pursued legal action against critics. The supervisory board of Wirecard was largely inactive, and EY, the external auditor, failed to verify key account balances for multiple years. The subsequent Wambach Report concluded that systemic failures in auditing, compliance, and regulation contributed to the magnitude of the scandal (Wambach, 2021). In response, Germany passed the Financial Market Integrity Strengthening Act (FISG) to improve audit standards and regulatory powers.

These three case studies illustrate the diverse mechanisms of financial fraud and underscore the importance of robust internal controls, independent auditing, and vigilant oversight. They demonstrate how complex and persistent fraudulent schemes can become when left unchecked, and highlight the need for systemic reforms that go beyond identifying individual misconduct.

1.2 Traditional Methods for Fraud Detection in Financial Statements

While the high-profile cases of Enron, Parmalat, and Wirecard have underscored the catastrophic consequences of accounting fraud, they have also revealed the limitations of traditional detection mechanisms in preventing or uncovering such misconduct in time (Ngai, Hu, Wong, Chen, & Sun, The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature, 2011) (Omidi, Masoumi, & Ahmadzadeh, 2019). For decades, regulators, auditors, and financial analysts have relied on a variety of established tools and procedures aimed at identifying inconsistencies and irregularities in financial statements. These conventional methods include quantitative models, statistical tests, financial ratio analysis, and structured review processes—each designed to serve as an early warning system against manipulation and deception (Feng, Ge, Luo, & Shevlin, 2021).

This section explores the primary techniques historically used to detect financial statement fraud, focusing on both their methodological foundations and practical applications. By examining tools such as Benford's Law, financial ratio models, trend deviation analysis, and outlier detection, the discussion aims to provide a critical understanding of their role within the broader fraud detection landscape (Zhou & Kapoor, 2011). Furthermore, the section highlights the key limitations that have emerged in the face of increasingly complex fraud strategies, growing data volumes, and globalized financial reporting environments (Ngai, Hu, Wong, Chen, & Sun, The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature, 2011).

Ultimately, understanding the strengths and weaknesses of these traditional methods is essential for evaluating the necessity—and the potential—of more advanced, technology-driven approaches, which will be discussed in the next section.

1.2.1 Fraud Indicators

Benford's Law

Among the most prominent tools in detecting anomalies in financial data is Benford's Law, a mathematical principle that predicts the frequency distribution of digits in naturally occurring datasets. According to this law, lower digits—especially the digit 1—occur as leading digits more frequently than higher ones. In a legitimate and naturally generated set of accounting figures, the first digit is expected to be 1 about 30.1% of the time, while 9 appears as the first digit only about 4.6% of the time (Berger, 2015).

Benford's Law has proven especially useful in forensic accounting because fabricated numbers often fail to follow this distribution. Perpetrators of fraud tend to insert figures that seem intuitively random or evenly distributed, inadvertently deviating from the expected Benford pattern. This deviation can serve as a red flag indicating potential data manipulation or fabrication in financial statements.

Applications of this method span various sectors, including corporate accounting, public auditing, and taxation. Its non-invasive and cost-effective nature makes it a popular preliminary screening tool. For example, studies show that discrepancies from Benford's expected frequencies in journal entries, expense claims, or tax returns can alert auditors to areas requiring deeper investigation (Nigrini, 2012).

Nonetheless, it is essential to contextualize its use. Benford’s Law is most effective with datasets that span several orders of magnitude and are not artificially bounded or rounded. Therefore, it should not be applied indiscriminately. Misuse of this tool may lead to false positives, prompting unnecessary investigative procedures. As such, its value lies in complementing, rather than replacing, traditional auditing and detection techniques (Durtschi, 2004).

In sum, Benford’s Law offers a mathematically grounded and empirically validated method to detect irregularities in financial data, serving as an early warning signal for auditors and investigators. Its integration into fraud detection frameworks enhances the ability to uncover manipulative practices before they escalate into full-scale financial misconduct.

Financial Ratio Analysis

Financial ratio analysis is one of the most traditional and widely adopted methods for evaluating a company’s financial health and identifying early signs of potential fraud or failure. Ratios derived from financial statements are used to assess a firm’s liquidity, solvency, profitability, and operational efficiency, serving as critical benchmarks for auditors, analysts, and regulators (Altman, 1968).

One of the most seminal contributions to this field was made by Edward I. Altman (1968), who demonstrated that financial ratios, when properly weighted and analyzed using multivariate techniques such as discriminant analysis, could significantly predict corporate bankruptcy. His Z-score model, based on five key ratios—working capital to total assets, retained earnings to total assets, earnings before interest and taxes to total assets, market value of equity to total liabilities, and sales to total assets—remains foundational in financial risk analysis.

These ratios are not merely descriptive but also predictive (Altman, 1968). For instance:

- Liquidity ratios like current ratio or working capital/total assets help detect potential cash flow issues.
- Profitability ratios, such as return on assets or EBIT/total assets, signal whether the firm's operations are sustainably profitable.
- Leverage ratios, such as debt-to-equity or equity-to-assets, can expose excessive risk-taking or hidden obligations.
- Efficiency ratios, such as sales-to-assets, reflect how effectively a company converts assets into revenue.

In fraudulent contexts, manipulation often targets these exact metrics. For example, overstating revenues boosts profitability ratios, while capitalizing expenses artificially inflates asset-related ratios. According to Akintoye and Olayinka (2013), one of the key reasons for fraudulent misrepresentation is to distort these ratios and thus maintain a positive public image or secure financing (Akintoye, 2013).

The major strength of financial ratio analysis lies in its simplicity and accessibility. It is relatively easy to implement and requires only published financial data. However, this simplicity can also be a limitation: ratios may be affected by aggressive but legal accounting choices, and they often fail to detect fraud hidden in footnotes or off-balance-sheet items (Akintoye, 2013).

Moreover, researchers like Zager and Zager (2015) argue that standalone ratios may not offer sufficient discriminatory power when used in isolation, especially given the increasing complexity of modern financial statements (Zager, 2015). This has led to the development of more sophisticated approaches, often combining ratio analysis with statistical or artificial intelligence techniques (Zager, 2015)..

Despite these limitations, financial ratios remain a cornerstone of traditional fraud detection, particularly when applied comparatively over time or against industry benchmarks. When integrated with other methods—such as Benford's Law or trend deviation analysis—they offer a useful foundation for identifying inconsistencies that merit further scrutiny (Zager, 2015).

Trend Deviation Analysis

Following the foundational tools of Benford's Law and financial ratio analysis, trend deviation analysis emerges as a complementary approach that offers dynamic insights into financial statement anomalies over time. While financial ratios provide static snapshots and Benford's Law targets digit distributions, trend deviation analysis focuses on temporal patterns, identifying unexpected changes in financial line items across multiple periods.

Trend analysis is traditionally employed in auditing and financial statement review as a horizontal analytical technique, where financial elements such as revenue, cost of goods sold, and net income are tracked across several reporting periods. The primary objective is to establish an expectation of continuity or proportional change. Deviations from expected trends—especially when not justified by contextual factors—may serve as red flags for manipulation or misreporting (Enyi, 2019).

However, the increasing complexity of financial reporting and the potential for sophisticated fraud techniques necessitate more granular and relational models. This has led to the development of Relational Trend Analysis (RTA), a refined methodology proposed by Enyi (2019) that integrates both horizontal and vertical components. RTA compares the relative weight of each financial item within its group (e.g., marketing expense as a proportion of total SG&A costs) and tracks these proportions across time. This technique enables the detection not only of fluctuations in absolute figures but also of structural distortions in how resources are allocated or reported.

For example, a significant deviation in administrative costs from historical norms—without corresponding changes in scale or business activity—may suggest concealed expenditure reclassification or artificial inflation of earnings through cost deferrals. Enyi (2019) demonstrates how RTA quantifies these shifts using relational indices and highlights the specific components and periods that require further investigation.

One key advantage of trend deviation analysis is its adaptability. Unlike fixed models that rely on predetermined thresholds, RTA adjusts to the specific financial behavior of the firm, reducing the likelihood of false positives. This adaptability also makes it particularly useful for small datasets, where statistical significance is harder to achieve with traditional tests (Enyi, 2019).

Nonetheless, this technique assumes the consistency of accounting policies and reporting periods over time, which may not always hold true. Deviations caused by legitimate changes in accounting standards, mergers, or strategic realignments must be carefully filtered out to avoid misleading conclusions. Furthermore, it is less effective in environments with high financial volatility or short financial histories (Yu, 2021).

Despite these limitations, trend deviation analysis—and in particular relational approaches like RTA—serve as valuable tools in the early detection of fraudulent activities. Their ability to trace evolving inconsistencies and pinpoint their exact origins enhances the precision of audit efforts and strengthens internal control mechanisms. These benefits are especially pronounced when trend deviation analysis is integrated with anomaly detection algorithms, the next technique discussed in the following section.

Outlier detection

While trend deviation analysis focuses on patterns unfolding over time, outlier detection targets isolated, atypical data points that deviate significantly from the expected structure of financial records. This technique is grounded in the assumption that fraudulent activity—intentional or otherwise—often manifests as inconsistencies that violate the normal distribution or density of financial data (Heikkinen, 2021).

In financial reporting, outliers may appear as unusually large transactions, irregular cost entries, or atypical account balances that do not align with operational logic. These anomalies, which may not form part of a trend, often arise from data entry errors, manual overrides, or direct manipulation. For this reason, auditors have long incorporated rule-based detection techniques—such as threshold analysis, z-score deviations, and interquartile range assessments—as part of standard internal control procedures (Demestichas, 2021).

More refined approaches assess multivariate relationships among financial variables. For example, inconsistencies between cash flow and net income, or between sales and inventory levels, can be indicative of misreporting. In this context, outlier detection functions as a complementary lens to ratio analysis: it does not rely on fixed financial models but evaluates how data points relate to the overall shape and density of the dataset (Heikkinen, 2021).

Techniques such as Local Outlier Factor (LOF) and Mahalanobis Distance are commonly applied in audit analytics. These methods assess whether a financial record lies in a low-density region compared t its neighbors, which may indicate fraudulent structuring. Unlike machine learning-based tools, these approaches are grounded in statistical logic and can be implemented without training data, preserving their status as non-AI traditional techniques (Demestichas, 2021).

The strength of outlier detection lies in its model-agnostic nature. It can be used across industries, firm sizes, and reporting standards, and does not assume specific distributional properties beyond internal consistency. Moreover, it can detect both extreme values and subtle relational irregularities, which might be missed by univariate tests or threshold-based models.

However, this flexibility comes with caveats. These methods are sensitive to noise, especially in volatile environments, and can produce false positives when applied without domain context. For example, seasonal variations or legitimate business fluctuations may trigger unnecessary alarms. As Heikkinen (2021) highlights, expert judgment and contextual understanding remain essential in interpreting the results of outlier-based assessments.

Ultimately, outlier detection serves as a valuable extension of traditional auditing and forensic accounting practices. It enables practitioners to isolate transactions that merit further scrutiny, especially when combined with ratio and trend-based analyses. While these tools enhance the robustness of traditional frameworks, they remain constrained by static rule sets and the absence of adaptive learning. This reinforces the growing interest in artificial intelligence and machine learning solutions, which will be explored in the next section.

1.2.2 Fraud Inefficiencies

Although traditional fraud detection methods have formed the backbone of financial analysis for decades, their application in modern, data-intensive environments reveals a number of critical shortcomings. This section outlines the key limitations of these approaches, structured into five main areas of concern.

Manual Process Challenges

One of the most fundamental limitations of traditional techniques lies in their heavy reliance on manual procedures. Financial audits based on ratio analysis, threshold reviews, and trend assessments depend on the judgment and interpretation of human auditors. This introduces variability and inconsistency in results, as outcomes may differ depending on the experience or attentiveness of the reviewer (Ngai, Hu, Wong, Chen, & Sun, The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature, 2011). In addition, manual analysis is prone to fatigue, bias, and oversight, particularly in high-volume contexts where large amounts of transactional data must be reviewed within tight deadlines. The need for manual sampling and documentation also increases the risk of missing fraudulent entries that fall outside of tested subsets (Ravisankar, Ravi, Raghava Rao, & Bose, Detection of financial statement fraud and feature selection using data mining techniques, 2011).

Costs, Delays, and Efficiency Gaps

Traditional audits are resource-intensive, requiring significant time and personnel. According to Ravisankar et al. (2011), delays in fraud detection are common and can result in substantial financial losses and reputational damage. The lack of automation limits the scalability of these techniques and restricts their applicability to large or fast-evolving datasets. Moreover, the audit cycle is often retrospective. Many traditional methods identify issues only after the publication of financial reports, reducing the capacity for real-time or preventive interventions.

Limited Predictive Power and High False Positives

Most traditional indicators are reactive in nature and were not designed to predict fraud proactively. For example, while an abnormal financial ratio or trend deviation may indicate a potential issue, it does not necessarily confirm fraudulent intent. As noted by Perols et al. (2017), such tools often generate high rates of false positives, especially when external factors such as seasonal variation or sector-specific practices are not properly accounted for. These inaccuracies can overload investigative teams and reduce the effectiveness of control environments (Perols, Murthy, & Zhang, 2017).

Vulnerability to Manipulation and Complexity Avoidance

Fraudulent actors are often well-versed in how to bypass standard controls. Through the manipulation of legitimate accounting practices—such as expense capitalization, revenue recognition timing, or shifting transactions across reporting periods—they can distort financial results while maintaining superficially acceptable ratios and trends (Chiu, Wang, & Vasarhelyi, 2020). Moreover, complex or cross-period frauds are often fragmented across accounts, subsidiaries, or operational units, making them particularly hard to detect using traditional linear or isolated review methods.

Scalability and Data Structure Limitations

Finally, conventional methods struggle to scale in modern audit environments characterized by big data and high-frequency reporting. Traditional audits typically rely on sampling, which may fail to detect anomalies that exist outside the reviewed subset.

Additionally, these tools are largely confined to structured numerical data, leaving out valuable insights hidden in unstructured sources such as footnotes, disclosures, or narrative sections of annual reports (Jans, Alles, & Vasarhelyi, 2014). This limitation is particularly problematic as financial communication increasingly includes soft data and qualitative cues that may signal misstatements.

Conclusion and Transition

These limitations underscore the pressing need to rethink the foundations of financial fraud detection. While traditional methods remain essential, they must be supplemented by adaptive, scalable, and intelligent tools capable of handling the complexity and volume of modern financial data. The following section introduces artificial intelligence-based approaches, including natural language processing and machine learning, which have emerged as promising solutions to overcome these traditional constraints.

1.3 AI in Financial Statement Analysis

Traditional fraud detection techniques—such as ratio analysis, trend deviation, and manual audits—have long served as foundational tools for identifying financial misstatements. However, the dynamic and complex nature of modern corporate fraud increasingly renders these methods insufficient on their own. As financial statements grow in volume, linguistic sophistication, and structural complexity, the need for scalable, adaptive, and data-driven solutions has become critical.

Artificial Intelligence (AI), particularly when combined with Natural Language Processing (NLP) and Machine Learning (ML), has emerged as a transformative force in financial statement analysis. These technologies promise not only increased accuracy and efficiency but also the ability to process unstructured data—such as footnotes, disclosures, and textual patterns—that traditional methods often overlook. As Nießner, Nickerson, and Schumann (2021) argue, AI-based approaches enable the extraction of deep insights from both quantitative and qualitative components of financial reports, reducing reliance on subjective human judgment (Nießner, Nickerson, & Schumann, 2021).

This section explores the core technologies that underpin AI in financial fraud detection, outlining their methodological foundations and explaining how they outperform—or complement—conventional methods. It also presents a taxonomy that helps classify and understand these diverse tools, paving the way for the analysis of their practical applications in the next subsection.

1.3.1 Main Techniques

Artificial Intelligence encompasses a wide array of computational techniques capable of learning from data, identifying patterns, and making predictions. In the context of financial statement analysis, several AI-based methods have proven especially useful for detecting fraud, evaluating risk, and extracting semantic meaning from textual data. These include Machine Learning (ML), Deep Learning (DL), and Natural Language Processing (NLP).

Machine Learning (ML)

ML algorithms are at the heart of intelligent fraud detection systems. These algorithms can be classified into:

- Supervised learning, where models are trained on labeled datasets (e.g., fraudulent vs. non-fraudulent reports);
- Unsupervised learning, which identifies patterns without labels, useful when fraud cases are rare;
- Semi-supervised learning, which combines the two approaches.

Support Vector Machines (SVM), Random Forests, Logistic Regression, and Neural Networks have all been used to detect financial fraud. According to Ravisankar et al. (2011), SVM in particular stands out as one of the most frequently applied and accurate techniques in financial fraud detection (Ravisankar, Ravi, Raghava Rao, & Bose, Detection of financial statement fraud and feature selection using data mining techniques, 2011).

Deep Learning

Deep Learning is a subfield of ML that uses multi-layered neural networks to uncover complex patterns in data. It is particularly effective when working with large and unstructured datasets. Craja, Kim, and Lessmann (2020) found that deep learning models applied to narrative sections of financial statements outperform traditional shallow classifiers in identifying fraudulent patterns (Craja, Kim, & Lessmann, 2020).

Natural Language Processing (NLP)

NLP enables machines to process and analyze human language, making it a powerful tool in analyzing the textual components of financial disclosures. NLP techniques include:

- Tokenization and Part-of-Speech Tagging, which parse sentences and identify grammatical roles;
- Sentiment Analysis, which measures tone and emotional polarity;
- Topic Modeling, which extracts hidden themes;
- Text Classification, which uses language patterns to distinguish between fraudulent and non-fraudulent documents.

Fisher, Garnsey, and Hughes (2016) emphasize that NLP adds a crucial dimension to fraud detection by revealing linguistic inconsistencies, vague expressions, or overly optimistic narratives that may otherwise go unnoticed through quantitative methods alone (Fisher, Garnsey, & Hughes, 2016).

Taxonomy of AI Methods

Nießner et al. (2021) provide a comprehensive taxonomy that categorizes AI-based methods for financial analysis along six dimensions: model type, learning method, data type, purpose, scope, and transparency. This taxonomy helps clarify the wide array of techniques available and guides their practical implementation in fraud detection systems (Nießner, Nickerson, & Schumann, 2021).

1.3.2 Previous Applications and Potential

Real-World Applications

Following the overview of core artificial intelligence techniques used in financial statement analysis, this section delves into the empirical applications and practical impact these methods have demonstrated in detecting fraud. The objective is twofold: to assess the real-world feasibility of AI-driven models and to identify the existing gaps that pave the way for future research.

Numerous studies have applied AI techniques—particularly machine learning (ML), deep learning (DL), and natural language processing (NLP)—to real financial datasets, achieving encouraging results. For instance, Li et al. (2022) applied NLP-based textual analysis on Form 20-F filings of Chinese firms listed in the U.S., finding a significant correlation between the tone of language and earnings management practices. Specifically, the usage of positive, uncertain, or modal language was associated with higher levels of earnings manipulation, demonstrating the efficacy of sentiment-based models in detecting subtle forms of deception (Li, Wang, & Luo, 2022).

Similarly, Luo and Zhou (2020) provided a detailed literature survey showing that the tone of financial disclosures significantly influences investor perception and can serve as an early warning signal of managerial opportunism or upcoming financial distress. This underscores the potential of text-based models in identifying risk signals that traditional quantitative models might miss (Luo & Zhou, 2020).

A compelling example of AI's practical contribution is seen in the work of Fisher et al. (2016), who explored the use of sentiment dictionaries and supervised learning to detect anomalies in MD&A disclosures. Their findings confirmed that NLP tools could highlight red flags related to fraud and performance misrepresentation, particularly when integrated with structured accounting variables (Fisher R. , Garnsey, Hughes, & Paterson, 2016).

Limitations and Gaps

Despite these promising outcomes, several limitations persist. First, the generalizability of AI models across jurisdictions and industries remains a challenge. Most studies focus on specific sectors (e.g., finance or healthcare), using regional datasets that may not reflect broader fraud patterns. Additionally, concerns about data labeling quality, especially when relying on publicly available datasets, introduce potential biases into training processes.

Moreover, the interpretability of some advanced models—especially deep learning frameworks—has emerged as a critical issue for regulators and auditors who must justify fraud assessments based on understandable criteria (Purda & Skillicorn, 2015). While models like decision trees or logistic regression offer transparency, neural networks and transformers often act as "black boxes," requiring supplementary explainability layers to be acceptable in legal or auditing environments (Purda & Skillicorn, 2015).

Future Directions

Nonetheless, the integration of linguistic indicators with financial variables holds great promise. As highlighted by Li et al. (2022), future research should explore hybrid models that combine text analysis, numerical anomaly detection, and contextual firm-specific information to improve accuracy and reduce false positives. Additionally, sector-specific ontologies and custom sentiment dictionaries—particularly in multilingual or non-Western contexts—are needed to address current model limitations (Li, Wang, & Luo, 2022).

Conclusion and Transition

In conclusion, while AI applications in financial fraud detection have shown tangible benefits, their full potential remains underexploited. Understanding the boundaries of current implementations and addressing their limitations is crucial before shifting toward large-scale regulatory adoption. The following chapter will build upon this groundwork by introducing the dataset and analytical pipeline used to test some of these techniques in practice.

CHAPTER 2 - METHODOLOGY

2.1 Dataset

The empirical analysis conducted in this study is based on a comparative dataset of financial statements from both fraudulent and non-fraudulent companies. The goal is to examine whether linguistic patterns in narrative disclosures, particularly within annual reports, can serve as reliable indicators of accounting fraud. To ensure a robust and meaningful comparison, the dataset includes multiple companies from diverse jurisdictions and sectors, with each fraudulent firm paired with two non-fraudulent peers operating in the same industry and of similar size and business scope.

The selection process focused on five major fraudulent firms whose accounting scandals were widely documented and studied: Enron Corporation, Parmalat S.p.A., Wirecard AG, Luckin Coffee Inc., and Tesco PLC. For each of these companies, two healthy counterparts were selected based on a combination of criteria, including industry alignment, financial scale, geographical relevance, availability of annual reports, and the absence of any known involvement in fraud-related investigations. The final comparative firms include General Electric, Nestlé, Adyen, Starbucks, Carrefour, and J Sainsbury plc.

To ensure temporal consistency, the financial statements for each company were collected for a period of two to three fiscal years prior to the public disclosure of the fraud. The primary objective was to analyze the linguistic features of corporate reporting during the period when fraudulent behavior was active but had not yet been detected. The number of years included varied depending on document availability and data quality: for most companies, three years were analyzed to allow for trend identification and historical context; however, in cases such as Enron and Parmalat, only two usable reports could be retrieved due to archival limitations, document degradation, or incomplete online records. The same number of years was applied to each fraudulent company's peer group to maintain symmetry in the comparison and to preserve the validity of the linguistic analysis. This strategy ensures that the observed differences or similarities in narrative disclosure are not artifacts of temporal inconsistency but rather reflect genuine divergences in

communication style and transparency. In total, the dataset spans the period from 1998 to 2020, covering a range of financial reporting environments and regulatory regimes.
The documents included in the analysis consist exclusively of annual reports, with a particular focus on the following sections:

- Management Discussion and Analysis (MD&A);
- Risk Factors (when available);
- Notes to the financial statements;
- CEO or management letters to shareholders.

These narrative-rich components were selected to support the application of Natural Language Processing (NLP) techniques aimed at detecting signals of deception, vagueness, overconfidence, or obfuscation in financial communication.
The sources of the financial statements are as follows:

- For U.S.-listed companies such as Enron and Luckin Coffee: the SEC EDGAR database;
- For European companies such as Parmalat and Wirecard: the Orbis platform (Bureau van Dijk) and the Internet Archive's Wayback Machine;
- For non-fraudulent peer companies: official company websites, including Carrefour, Starbucks, and Nestlé, as well as Orbis for standardized financial data.

A comprehensive Excel file was developed to record all the collected data, organized by company, year, document type, and specific sections extracted. This structured documentation ensures traceability throughout the project and supports the systematic classification of linguistic indicators across time and corporate profiles.

2.1.1 Data Collection

Building upon the dataset described above, this section outlines the procedures adopted to collect and structure the financial statement documents used in the empirical analysis.
The data collection process involved the systematic retrieval of annual reports, selected as the sole document type due to their comprehensive and consistent narrative sections. These reports typically include the Management Discussion and Analysis (MD&A), Risk

Factors, notes to the financial statements, and CEO letters to shareholders—textual components well suited for Natural Language Processing (NLP) applications. Other corporate disclosures, such as quarterly filings, earnings announcements, or press releases, were deliberately excluded to maintain homogeneity in document structure and content.
For U.S.-listed companies, reports were sourced primarily through the SEC EDGAR database. European companies, including Parmalat, Tesco, and Wirecard, were accessed via the Orbis platform (Bureau van Dijk) and the Internet Archive's Wayback Machine. Depending on availability, documents were retrieved in PDF, HTML, or scanned image format and later standardized to plain text (.txt) for processing.

The target period for analysis included two to three full fiscal years prior to the public disclosure of the fraud, with the aim of capturing disclosure behavior during the operational phase of fraudulent activity. In most cases, three yearswere collected to enable trend observation and linguistic consistency over time. However, for older fraud cases such as Enron and Parmalat, only two usable reports were available due to archival degradation, missing records, or poor document quality. To ensure comparability, the same time span was applied to the matched non-fraudulent peer companies, maintaining symmetry and analytical rigor across all case comparisons.

To support a valid comparative framework, each fraudulent company was matched with two non-fraudulent peersaccording to a set of predefined criteria. These included:

1. Sectoral alignment: all companies operate in the same or a closely related industry;
2. Business scale comparability: similar size in terms of revenues, market presence, or global operations;
3. Availability of annual reports for the same pre-fraud time window;
4. Absence of fraud or regulatory sanctions during the years considered.

This matching strategy was designed to control for sector-specific language and disclosure norms, ensuring that observed linguistic differences could be attributed more plausibly to the presence or absence of fraudulent behavior rather than to unrelated corporate characteristics.

The table below summarizes the selected cases and their matched peers, along with the rationale behind each pairing:

Table 2.1 – Fraudulent Companies and Their Industry-Matched Non-Fraudulent Peers

Fraudulent Company	Peer Company 1	Peer Company 2	Sector	Matching Criteria
Enron	AES Corporation	General Electric	Energy / Utilities	Same industry; similar size; US-listed; reports available (1998–2000); no fraud history
Parmalat	Danone	Nestlé	Food & Beverage	Same sector; European-based; similar scale; reports available (2001–2003); no fraud events
Wirecard	PayPal	Adyen	Fintech / Payments	Digital payment services; overlapping business models; same period; clean audit history
Luckin Coffee	Starbucks	Dunkin' Brands Group	Retail / Coffee	Direct competitors; multinational scope; reports for 2017–2019 available; no fraud implications
Tesco	Carrefour	J Sainsbury plc	Retail / Supermarkets	Comparable market role in Europe; similar size; consistent financial reporting from 2012–2014; not involved in fraud events

Source: Author's elaboration using Voyant Tools

All documents were then organized in structured folders by company and fiscal year, and their metadata—such as company name, reporting year, sector classification, fraud status, and data source—were compiled in a master Excel file. This structure ensured traceability and consistency throughout the preprocessing and analysis phases.

The technical details of the subsequent preprocessing, cleaning, and structuring of the text corpus are discussed in the following section.

2.1.2 Preprocessing and Categorization

Following the data collection process, it was necessary to preprocess and systematically categorize the textual data to enable subsequent linguistic analysis.

The first step involved the organization of the collected documents within a structured Excel spreadsheet. Each annual report was catalogued by company name, reporting year, and sector classification. Additionally, each firm was labeled based on its fraud status (fraudulent or non-fraudulent) and paired with appropriate comparison companies operating in the same or a closely related industry. This preliminary categorization ensured a clear mapping between each fraudulent company and its non-fraudulent peers, thus facilitating controlled comparative analysis.

Once the dataset was organized, all documents were standardized into plain text (.txt) format to facilitate compatibility with AI-based Natural Language Processing (NLP) tools. Converting files into plain text is essential for simplifying subsequent computational operations such as tokenization, sentiment analysis, and pattern detection. To perform this conversion, two online tools were utilized: Convertio and Online-Convert. These platforms allowed for an efficient and consistent transformation of different original formats (e.g., PDF, HTML, scanned images) into clean text files, minimizing formatting artifacts that could interfere with textual analysis.

The preprocessing phase focused exclusively on preparing the dataset for further investigation, without yet applying any analytical techniques or linguistic transformations. This approach ensured that the entire corpus was standardized, comparable across companies and years, and ready for the application of NLP methods.
In the next section, the specific Natural Language Processing techniques employed for the analysis of the preprocessed texts are presented in detail.

2.2 Application of NLP Techniques

After completing the preprocessing and categorization of the dataset, the next phase of the research involved the application of Natural Language Processing (NLP) techniques to the collected financial narratives. The primary objective of this stage was to identify and extract linguistic features that could serve as potential indicators of accounting fraud.

Given the narrative nature of annual reports, particularly sections such as Management Discussion and Analysis (MD&A), Risk Factors, and CEO letters, the application of text mining and NLP methods offers a promising avenue for detecting patterns of vagueness, overconfidence, risk omission, and linguistic deception that may precede the public revelation of fraudulent activities.

This section is structured as follows: Subsection 2.2.1 introduces the specific NLP techniques selected for the analysis and discusses the rationale behind their choice. Subsection 2.2.2 then details how these techniques were applied to the preprocessed dataset to extract meaningful linguistic indicators for further evaluation.

2.2.1 NLP Techniques for Fraud Detection

As introduced in the previous section, the objective of this phase was to identify textual features in corporate disclosures that could act as potential linguistic signals of accounting fraud. Based on the narrative richness of specific sections of annual reports—namely the Management Discussion and Analysis (MD&A), Risk Factors, notes to the financial statements, and CEO letters—this study employed a series of Natural Language Processing (NLP) techniques aimed at detecting patterns of vagueness, evasion, overconfidence, and risk omission.

All reports were processed in `.txt` format to ensure compatibility with the chosen tools. The analysis was conducted using Voyant Tools, a browser-based text analysis platform that allows for the visualization and comparison of linguistic elements, and a GPU-enabled local workstation, which was used for preprocessing, storage, and file standardization. This configuration enabled a reproducible and non-programmatic pipeline that ensured consistency across all documents.

The specific NLP techniques adopted and their methodological rationale are detailed below:

1. Tokenization and Word Frequency Analysis: Tokenization consists of breaking down a text into individual words or "tokens." This operation was performed

automatically by Voyant Tools when uploading each `.txt` file. Through the *Summary* and *Cirrus* tools, a frequency list was generated for each document, showing the most recurrent words. This information was used to identify lexical patterns such as the repeated use of terms like *growth*, *opportunity*, *success*, or *innovation*—expressions typically associated with strategic framing and promotional tone.

To ensure consistency across documents, all frequency counts were normalized relative to total word count (terms per 10,000 words), making it possible to compare reports of varying lengths. The goal of this technique was to detect semantic overemphasis on specific themes, which may reflect attempts to divert attention from problematic financial conditions.

2. Stopword Removal and Keyword Filtering: Standard English stopwords—such as *the*, *and*, *of*, *with*—were excluded using Voyant's default stopword filter. This process allowed the focus to shift toward content-bearing words and meaningful phrases. In comparative analysis, this filtering was essential to uncover underlying thematic consistencies or anomalies across reports.

   Additionally, a custom list of domain-relevant keywords (e.g., *risk*, *uncertainty*, *debt*, *governance*) was developed and tracked manually for frequency and contextual use, particularly in sections where such terms are typically expected (e.g., Risk Factors or MD&A).

3. Contextual Analysis of Modal Verbs and Hedging Language: One of the linguistic strategies often linked to deceptive communication is the use of modal verbs to express uncertainty, potentiality, or indirect commitment. Using the *Contexts* tool in Voyant, the study performed targeted searches for modals such as *may*, *might*, *could*, and *would*. Each instance was reviewed in context to determine whether the surrounding phrasing introduced ambiguity, shifted responsibility, or softened risk-related statements.

   A qualitative threshold was established: the presence of such modals in strategic forecasts, risk disclosures, or forward-looking statements was systematically recorded and annotated. Particular attention was paid to the recurrence of these verbs in proximity to financial or operational claims.

4. Thematic Emphasis and Lexical Framing: Although Voyant does not offer built-in sentiment analysis, its visual and statistical outputs allow for an indirect evaluation of lexical tone. This study examined the balance between optimistic and cautious terminology, focusing on whether corporate narratives overemphasized performance, ambition, or leadership, while underrepresenting uncertainty, risk exposure, or operational limitations.

   A qualitative scale was applied to assess the dominant framing strategy of each report:

   - *Overconfident*: frequent positive qualifiers, lack of risk-related terms;
   - *Balanced*: presence of both positive and negative qualifiers;
   - *Cautious*: emphasis on challenges, constraints, or market volatility.

   This interpretive layer was essential for highlighting cases where the lexical tone diverged from typical disclosure norms or sector expectations.

All four techniques were applied consistently across all documents using the same environment, sequence, and interpretive logic. Their combined application formed the foundation for the comparative linguistic analysis between fraudulent and non-fraudulent companies.

The next subsection (2.2.2) describes how these techniques were operationalized on the dataset, how the outputs were recorded, and how the findings were structured for inter-company and temporal comparison.

2.2.2 Application to Financial Data

Building on the NLP techniques previously outlined, this section describes the practical implementation of the textual analysis on the dataset of annual reports. The goal here is not to discuss findings, but to detail the workflow used to extract, annotate, and organize linguistic indicators in a consistent and reproducible way.

Each annual report, previously converted into `.txt` format, was uploaded individually into Voyant Tools, which served as the main platform for exploratory analysis. The workflow followed these core steps:

1. Uploading the `.txt` file to Voyant Tools;
2. Reviewing outputs from the *Summary*, *Cirrus*, *Contexts*, and *Trends* tools;
3. Copying relevant observations (e.g., most frequent terms, use of modal verbs, prominent lexical fields) into a working document or directly into a spreadsheet;
4. Downloading or exporting data segments where supported.

In parallel, each document was also analyzed through a language model (ChatGPT) to support qualitative review. The `.txt` files were entered into the model one at a time, and prompts were provided to identify:

- Examples of vague or ambiguous phrasing;
- Modal verbs expressing uncertainty (*may*, *might*, *could*);
- Sentences reflecting excessive optimism or promotional tone;
- Omission of expected risk-related language.

All observations from both tools were subsequently entered into a master Excel spreadsheet. Each row represented one document, and the structure of the table included the following columns:

Table 2.2 – Description of Linguistic Variables Extracted from Annual Reports

Company Name	Name of the analyzed firm
Year	Fiscal year of the report
Fraud Status	Fraudulent / Non-Fraudulent
Top Keywords	Most frequent meaningful terms (excluding stopwords)
Modal Verbs	Modals found in context (e.g., "may impact", "could lead to")
Tone	Coded as: Overconfident / Balanced / Cautious
Vague Phrasing	Examples of conditional or non-committal statements
Risk Language Presence	Yes/No + relevant terms (e.g., "uncertainty", "exposure")
Observational Notes	Manual annotations, inconsistencies, formatting issues

Source: Author's elaboration based

After initial input, filters were applied in Excel to sort and compare documents based on fraud status, tone, or linguistic markers. This enabled a structured preparation of the data for the comparative stage of analysis, without drawing premature conclusions at this point. This methodology allowed for the generation of a clean, annotated dataset, combining both machine-extracted features and human-curated insights. These outputs form the empirical base for the structured analysis described in Chapter 3.

2.3 Analytical Pipeline

The linguistic indicators extracted from the annual reports—through both automated processing and manual interpretation—constitute the foundation for the comparative analysis between fraudulent and non-fraudulent companies. To ensure consistency and traceability throughout the entire research process, a structured analytical pipeline was developed.

This section outlines the operational workflow adopted for processing the financial narratives, as well as the interpretive strategies used to analyze and organize the outputs. By formalizing the steps from raw text input to coded linguistic observations, the analysis becomes replicable and methodologically robust.

Subsection 2.3.1 details the full preprocessing workflow, while Subsection 2.3.2 focuses on the interpretation of the analytical results and how they were used to identify recurring linguistic patterns associated with financial misconduct.

2.3.1 Preprocessing Workflow

To support a coherent and replicable analysis, a structured workflow was implemented to manage the processing of all financial narratives. This pipeline encompassed the preparation, exploration, and organization of textual data, ensuring uniformity across all documents and facilitating subsequent interpretation.

The procedure began with the conversion of annual reports into .txt format, using online tools such as Convertio and Online-Convert. This step enabled compatibility with textual analysis platforms and allowed for direct access to the narrative components of each report.

Each .txt file was then analyzed through Voyant Tools following a standardized routine. The most frequent terms were extracted using Summary and Cirrus, while Contexts enabled the examination of modal verbs (may, might, could) within their textual environment. Additional insights were drawn using Trends and Keywords in Context, focusing on lexical emphasis and rhetorical framing.

Complementing this, a manual linguistic review was conducted using ChatGPT. This qualitative layer helped identify nuanced elements such as ambiguous phrasing, overly optimistic tone, and other narrative irregularities that may signal deceptive disclosure practices.
Results were systematically recorded in an Excel file, with each entry including:

- Company and fiscal year;
- Fraud status;
- Key terms and their frequencies;
- Use of modals and vague expressions;
- General tone;
- Linguistic anomalies;
- Observational notes.

This standardized process ensured analytical consistency and produced a clean dataset for comparative interpretation.

The following subsection (2.3.2) illustrates how the extracted features were analyzed and contrasted across companies to identify potential fraud-related linguistic patterns.

2.3.2 Analysis of Model Outputs

Following the structured workflow described in the previous section, the final stage of the analytical pipeline focused on the organization and preliminary interpretation of the linguistic data extracted from each annual report. The aim of this phase was to prepare the dataset for comparative analysis, ensuring that all relevant variables and observations were systematically categorized and aligned with the study's research objectives.

The analysis was conducted using the dataset compiled in the Excel tracking file, which included standardized entries for each company and year. For each document, a set of linguistic features was reviewed and logged, such as the frequency of key terms, the presence and context of modal verbs, tone-related elements, and any narrative anomalies observed during qualitative inspection.

Each document was evaluated first on an individual basis and subsequently grouped by fraud classification—fraudulent or non-fraudulent. This allowed for the organization of textual characteristics into comparative clusters, laying the groundwork for pattern identification in later stages. At this point, no assumptions or conclusions were drawn regarding the presence or absence of fraud-related indicators.

The categorization process involved filtering and segmenting the dataset by key variables, such as:

- Word usage patterns (normalized frequencies of top lexical items);
- Contextual presence of modal verbs (e.g., may, might, could);
- Qualitative annotations of tone and rhetorical framing;
- Presence or absence of risk-related terminology;
- General observations and metadata associated with each report.

While the methodology was applied uniformly, it is important to recognize certain limitations inherent to the interpretive nature of textual analysis. Subjectivity in tone evaluation and differences in sectoral communication styles may influence the perception of certain features. Moreover, variability in document quality and availability, particularly for older cases, represents a constraint that was carefully documented.

Despite these considerations, the outputs generated from this process provided a consistent and organized dataset, which now serves as the empirical foundation for the comparative discussion presented in Chapter 3.

CHAPTER 3 - RESULTS AND ANALYSIS

3.1 Introduction

This chapter presents the results of the linguistic analysis conducted on the annual reports collected for this study. The primary objective is to examine whether systematic differences exist between the narrative disclosures of companies that committed accounting fraud and those that did not, and to assess whether such differences may serve as early indicators of fraudulent behavior.

To enable a robust comparison, the analysis was structured around a pairwise design, whereby each fraudulent firm was matched with two non-fraudulent companies from the same industry and of comparable size and market scope. This approach was adopted to minimize potential confounding factors related to sector-specific language, reporting conventions, or corporate culture.

The five fraudulent companies selected for the study are: Enron Corporation, Parmalat S.p.A., Wirecard AG, Luckin Coffee Inc., and Tesco PLC. These firms were chosen based on the severity and notoriety of their accounting scandals, and the availability of annual reports in the years leading up to the public exposure of the fraud. For each of these, two healthy comparison companies were selected. For example, Enron was compared with AES Corporation and General Electric Company; Parmalat with Danone and Nestlé; Wirecard with PayPal and Adyen; Luckin Coffee with Starbucks and Dunkin' Brands Group; and Tesco with Carrefour and J Sainsbury plc.

The comparative analysis focused on several linguistic indicators that were previously discussed in Chapter 2. These included word frequency patterns, the use of modal verbs and vague expressions, the overall narrative tone, and the presence or omission of risk-related language. The analysis aimed to detect whether fraudulent companies, prior to the exposure of misconduct, employed distinctive narrative strategies that differ from the more transparent and balanced communication style observed in their non-fraudulent peers.

Each indicator was evaluated both quantitatively—through the use of Voyant Tools—and qualitatively—through manual inspection and interpretation. The results are organized by linguistic dimension, and illustrated with examples drawn from the documents analyzed. Where relevant, trends are highlighted across multiple years and among different companies to support the identification of recurrent patterns.

By comparing each fraudulent company with two non-fraudulent counterparts, the goal is to strengthen the validity of the findings and mitigate the effect of individual outliers. This chapter therefore provides the empirical backbone of the study, illustrating how subtle linguistic differences in financial reporting can signal deeper issues in corporate transparency and ethical conduct.

3.2 Results by Linguistic Indicator

A closer examination of corporate narrative structures reveals that fraud is often not only hidden in numbers, but linguistically embedded in the way companies tell their story. This section presents the core findings of the linguistic analysis conducted on the annual reports of five fraudulent firms and their respective non-fraudulent peers. By focusing on language patterns that precede the public disclosure of fraud, this chapter seeks to uncover textual signals that may act as red flags in early detection efforts.

The comparative analysis was structured around a one-to-two matching design, where each fraudulent company was contrasted with two industry-aligned peers of similar size and operational scope. Linguistic indicators were extracted and assessed using a combination of Natural Language Processing (NLP) tools (primarily Voyant Tools) and qualitative interpretation through large language model-based review. The analysis was conducted across three years per company (when available), enabling temporal triangulation and the identification of consistent narrative features over time.
The results are organized around four main linguistic dimensions, each supported by empirical examples from the analyzed reports:

- Word Frequency and Lexical Patterns: recurring keywords and lexical fields used to build corporate narratives;

- Modal Verbs and Vague Language: expressions of uncertainty, hedging, and evasion of responsibility;
- Tone and Sentiment Orientation: the dominant emotional or rhetorical posture conveyed;
- Risk Disclosure and Omission: the extent and quality of risk-related language and its positioning within the text.

The dataset reveals that fraudulent firms systematically employ euphoric, vague, and non-committal language, while omitting or marginalizing risk-related terminology. These linguistic features are neither accidental nor merely stylistic—they serve precise communicative functions: to build confidence, defer accountability, and obscure financial fragility. For example, Enron's 1998–1999 reports portray the company as a visionary market leader through phrases like *"the blue-chip of the 21st century"*, while failing to disclose any substantial operational risk. Wirecard, shortly before its collapse, displays a similar pattern of overconfidence, speaking of *ecosystems*, *platform synergies*, and *ongoing global expansion*while relegating the word *risk* to generic legal disclaimers. By contrast, non-fraudulent companies such as AES Corporation, Danone, and Tesco exhibit a markedly different narrative style. Their reports include balanced disclosures of performance and risk, a more restrained use of modal verbs, and a linguistic framing that aligns more closely with operational realities.

This section provides a structured walkthrough of each indicator, highlighting key deviations between fraudulent and healthy firms. Specific textual excerpts, frequency distributions, and semantic clusters are used to support each observation. The purpose is not merely to describe language, but to show how narrative choices reflect—and sometimes predict—the strategic intentions and ethical stance of the reporting firm.

3.2.1 Word Frequency and Recurring Lexical Patterns

The analysis of the most frequently used words in annual reports revealed clear distinctions between fraudulent and non-fraudulent firms, particularly in terms of recurrence, thematic density, and semantic polarization. Companies involved in fraud tend to adopt highly promotional language, with repetitive and inflated use of positively connoted terms—often disconnected from actual financial performance or operational risks. In contrast, healthy firms demonstrate a more balanced vocabulary, where growth-related language is accompanied by explicit references to challenges, regulatory constraints, and risk factors.

Enron Corporation

Enron's 1998 and 1999 annual reports display an excessive presence of terms such as "growth" (58 occurrences in 1998 and 63 in 1999) and "opportunity" (41 and 38, respectively). These words frequently appear in highly optimistic contexts, e.g., *"unparalleled growth opportunities in deregulated markets"* and *"Enron is uniquely positioned to capitalize on global energy expansion."* However, explicit references to "risk," "regulatory constraints," or "liquidity" are scarce (fewer than 10 mentions of each across both years), resulting in a significant narrative imbalance. The term "risk" appears only 6 times in the 1999 report, compared to over 30 occurrences on average in non-fraudulent peers for the same period.

Wirecard AG

Wirecard shows a similar pattern in its 2016–2018 reports. Words such as "platform," "ecosystem," "innovation," and "transaction" appear with high frequency. In the 2018 report alone, "platform" is cited 52 times, "innovation" 33 times, and "transaction" 61 times. The overall tone seeks to reinforce an image of relentless global expansion, with phrases such as *"pioneering digital ecosystems for financial inclusion."* Yet, terminology related to compliance, auditing, or risk is noticeably underrepresented. The word "risk" appears only 9 times in the 2018 document—**7** of which are in generic legal disclaimers rather than in the management discussion.

Parmalat S.p.A.

Parmalat's 2000 and 2001 reports emphasize corporate stability and growth with words like "expansion" (29 times in 2000) and "leadership" (22 times in 2001). However, key financial control elements such as "debt" or "cash flow" are barely discussed—"cash flow" appears only 3 times in the 2001 report, despite the company's growing liquidity crisis. Notably absent are terms like "governance," "accountability," or "risk management," suggesting a deliberate narrative omission.

Luckin Coffee Inc.

In its 2018 filing—the year before the fraud was uncovered—Luckin Coffee's language is heavily focused on "growth,""scale," and "expansion." "Growth" appears 64 times, "scale" 39 times, and "expansion" 41 times. In stark contrast, the terms "risk" and "cost" appear only 4 and 6 times, respectively, and mostly in legal disclaimers or broad generalities. Core concepts in financial oversight, such as "audit" or "compliance," are virtually absent.

Comparison with Non-Fraudulent Firms

Healthy firms use more varied and realistic language, blending opportunity-related terms with risk-awareness and responsible disclosure:

- AES Corporation (Enron's peer) uses "risk" 32 times in its 1999 report, "liquidity" 18 times, and includes detailed discussions on debt, energy market volatility, and regulatory uncertainty.
- Nestlé S.A. frequently references "sustainability" (43 times), "cost" (51), "risk" (36), and "regulatory" (27) in its 2017 report, reflecting an operationally grounded and comprehensive narrative.
- Danone consistently includes terms such as "currency volatility," "governance," and "debt restructuring," all of which are nearly absent from Parmalat's pre-scandal disclosures.
- Starbucks, in comparison to Luckin Coffee, systematically incorporates sections on "supply chain risks," "cost structures," and "operational challenges," with over 40 risk-related references per year.

Lexical Convergence Among Healthy Firms

Another notable trend is the lexical convergence across healthy companies, even in different industries.

Words like “reporting”, “compliance”, “transparency”, “governance” and “regulatory framework” appear consistently, indicating a communicative style focused on accountability and structured disclosure. On average, non-fraudulent companies display a frequency of 1 risk-related term every 250 words, compared to 1 in every 1,200 words for fraudulent firms.

Final Observations

Overall, the analysis reveals a consistent pattern: fraudulent firms tend to construct overly optimistic, promotional narratives often detached from the underlying financial reality. The imbalance between opportunity-related and risk-related language can serve as an early linguistic warning sign, suggesting narrative inflation or the deliberate omission of critical elements. While these features alone are not conclusive, they lay the lexical groundwork for further investigation, particularly in conjunction with modal verb analysis (3.2.2) and sentiment orientation (3.2.3).

3.2.2 Use of Modal Verbs and Vague Language

Another critical dimension of linguistic obfuscation in corporate narratives lies in the use of modal verbs and vague formulations, which allow firms to communicate expectations, projections, or intentions without committing to concrete outcomes. Such language can serve to hedge responsibility, defer accountability, or inflate strategic clarity while maintaining plausible deniability. This section explores the frequency and contextual use of hedging expressions, with a focus on modals such as *may*, *might*, *could*, *would*, and epistemic phrases like *we believe*, *we expect*, *it is possible*, and *intends to*.

Enron Corporation

Enron’s reports from 1999 to 2001 demonstrate an escalating use of modal verbs and vague constructs. In 1999, expressions such as *"believes," "plans to,"* and *"anticipates"* are used 42 times across the MD&A section alone. In 2000, the modal *may* appears 19 times and *could* 13 times. By 2001, amid the growing crisis, modal usage surges with "may" (33)**,** "might" (22)**,** and "could" (18)**,** often in defensive or speculative contexts:

- "There can be no assurance that such a market will develop."
- "Enron believes that applying its skills… can result in operating efficiencies."
- "It is possible that the investigation… will identify additional or different information."

This linguistic trend aligns with a shift in communicative posture, from promotional optimism to legally cautious distancing. Notably, Enron increasingly avoids definitive statements about performance or accountability, especially in sections discussing financial restatements and legal reviews

Wirecard AG

Wirecard's 2016–2018 filings consistently deploy modal verbs in contexts of expansion, market entry, and regulatory outlooks. In the 2018 report, the term *may* is used 27 times, *could* 22 times, and *we believe* appears 31 times—usually without accompanying evidence:

- "The Group may benefit from regulatory convergence across jurisdictions."
- "We believe the platform is positioned to transform global payments."

Yet, such statements often stand without quantifiable targets or time-bound plans. The phrase *"positioned to"* recurs 14 times in 2018 alone, signaling an assertive vision cloaked in linguistic uncertainty. There is a distinct absence of binding commitments, and vague terms like *opportunity*, *vision*, or *potential* often accompany modals—indicative of rhetorical inflation rather than disclosure clarity.

Parmalat S.p.A.

Parmalat's 2000–2001 reports contain a large number of uncommitted projections, particularly surrounding international expansion and financing strategies. In 2001, *may* and *could* appear 11 and 9 times, respectively, while *intends to* appears 12 times in broad and unspecific strategic sections:

- “Parmalat may seek alternative funding solutions.”
- “The Group intends to strengthen its international presence.”

While such formulations appear strategic, their lack of operational grounding reveals an intentional avoidance of concrete planning, which correlates with the underlying financial mismanagement that eventually surfaced.

Luckin Coffee Inc.

The 2018 annual report of Luckin Coffee exhibits high modal density for a company in early growth stage. *May* (21 times), *might* (13), and *we believe* (34) dominate forward-looking statements. Phrases like *“we believe we are well-positioned to disrupt the market”* or *“our expansion could significantly scale the business”* are typical, yet the absence of quantitative benchmarks or risk buffers renders these statements commercial rather than informative.

Interestingly, the report rarely uses stronger, confirmatory verbs (*will, is expected to, has secured*), and instead over-relies on speculative and self-referential framing—a hallmark of confidence masking fragility.

Comparison with Non-Fraudulent Firms

In contrast, healthy firms display more disciplined and transparent use of modals:

- AES Corporation (1999) uses modal verbs mainly in legal disclaimers, with *may* (14 times), but consistently balances them with specific data and scenario analysis.
- Danone (2000–2001) uses phrases such as *"we anticipate"*, but always followed by concrete actions or financial metrics.
- Nestlé (2017) incorporates expressions like *“could impact margin volatility”* in risk disclosures, showing an effort to forecast uncertainty without ambiguity.
- Dunkin’ Brands (2018–2019)**,** while also using *we believe*, does so in sections clearly marked as forward-looking and consistently pairs optimism with cautionary language, e.g., *“Our actual results could differ materially from those*

anticipated...”

“We believe our strategy continues to deliver...”.

Patterns and Interpretive Insights

Fraudulent companies use modal verbs not merely as stylistic devices, but as strategic tools for narrative deflection and risk dispersion. The recurrence of *may, could,* and *we believe* in speculative contexts—particularly unaccompanied by data—suggests an intentional hedging posture. In contrast, non-fraudulent firms frame uncertainty with bounded caution, disclosing potential downsides alongside growth narratives.

Across the corpus, modal-heavy passages in fraudulent reports tend to:

- Avoid temporal anchoring (no timeframes);
- Lack performance-linked KPIs;
- Cluster around vague strategic claims (*positioned to*, *plan to*, *envisioned growth*).

3.2.3 Tone and Sentiment Orientation

The emotional tone adopted in annual reports plays a decisive role in shaping stakeholder perception. In the context of fraudulent corporate narratives, tone often transcends stylistic preference and becomes a strategic rhetorical device used to overstate success, mitigate concerns, or suppress red flags. This section analyzes the tone and sentiment orientation of both fraudulent and non-fraudulent firms by focusing on the polarity of expressions, the presence or absence of emotionally charged language, and the alignment between tone and operational context.

Enron Corporation

Enron’s pre-collapse reports (1998–1999) exhibit a strikingly euphoric and promotional tone, heavily relying on phrases like *"unparalleled growth," "revolutionizing global markets,"* and *"Enron is redefining the energy future."* Positive qualifiers dominate—*growth* (58 and 63 times), *leadership*, *innovation*, and *success* are pervasive, especially in the CEO letter and MD&A sections. In contrast, negative or cautionary terms such as *risk*, *liquidity*, or *volatility* are virtually absent (fewer than 10 occurrences combined across both years).

By 2001, however, Enron's tone becomes defensive and legalistic, signaling a rhetorical shift amid mounting scrutiny. Phrases like *"information may be incomplete," "subject to regulatory review,"* and *"no assurance can be given"* replace prior triumphalism. This transition reveals how tone may function as a crisis management tool, modulating stakeholder expectations in the face of impending collapse.

Wirecard AG

Wirecard's reports between 2016 and 2018 epitomize rhetorical inflation. In 2018 alone, CEO Markus Braun's letter proclaims: *"We can look back on an extremely successful 2018 fiscal year both from a technological and fundamental perspective,"* and *"I am optimistic that Wirecard has a fantastic future ahead of it."*

The language is hyper-positive, self-congratulatory, and strategically disconnected from any substantive financial caution. Words like *success, scalable, innovation, platform,* and *customer-centric* appear with disproportionate frequency. Risk-related terminology is either absent or buried in boilerplate disclaimers (*"no material impacts are anticipated"*). This tone creates an illusion of invincibility—a stylistic choice at odds with the later financial reality.

Parmalat S.p.A.

Parmalat adopts a distinctively neutral and bureaucratic tone, characterized by dense, technical language and minimal emotional expression. Phrases like *"strategic repositioning of operational units"* and *"enhanced integration of regional segments"* are frequent, yet they obfuscate more than they clarify. Even amid mounting liquidity problems (2000–2001), the tone remains unnervingly flat, with no expression of urgency or concern. This lack of tonal modulation may reflect a deliberate effort to desensitize readers to underlying issues by removing affective language altogether.

Luckin Coffee Inc.

Luckin Coffee's 2018 report exhibits a marketing-oriented and overly enthusiastic tone, focused almost entirely on growth and disruption. Phrases such as *"redefining the coffee experience," "aggressive expansion strategy,"* and *"setting new industry benchmarks"* dominate the report. Positive sentiment is abundant, but no space is given to operational challenges, market risks, or competitive pressures. This one-dimensional optimism is symptomatic of a rhetorical inflation strategy, often used to attract investors in early-stage speculative firms. After the fraud was exposed in 2020, the tone dramatically shifted toward legal formalism and reputational containment, further emphasizing its manipulative origin.

Comparison with Non-Fraudulent Firms

In stark contrast, non-fraudulent companies present a more balanced and grounded tone.

- AES Corporation (1999–2000) integrates both performance and risks in its messaging: *"Our international operations have performed well, but remain exposed to currency and political volatility."*
  Terms such as *"risk," "uncertainty," "VaR,"* and *"hedging strategies"* appear frequently, reflecting an understanding of the operating environment.
- Nestlé (2017) employs a tone of strategic realism, balancing innovation (*"sustainable food system transformation"*) with regulatory and logistical challenges (*"complex international compliance environments"*). The sentiment reflects both ambition and accountability.
- Tesco (2012–2014) uses self-critical and reflective tone, particularly in underperforming years:
  "Our results reflect the need to realign with customer expectations… corrective measures are underway." This expression of accountability is absent from the fraudulent corpus.

Overall, the sentiment analysis shows that non-fraudulent firms do not avoid negative polarity when warranted. They use cautionary tone in risk sections, neutral tone in reporting performance, and only employ positive language when substantiated by evidence.

Summary Patterns and Strategic Implications

Three distinct tone strategies are identifiable in the fraudulent group:

1. Hyper-positive and promotional (Enron, Wirecard): celebratory, unbalanced, repetitive, and disconnected from reality.
2. Technocratic opacity (Parmalat): neutralized tone that avoids commitment, suppresses both enthusiasm and concern.
3. Investor seduction (Luckin): emotionally charged, aspirational, and lacking substance.

These strategies contrast with the tone discipline shown by non-fraudulent peers, who:

- Vary sentiment based on content (risk vs growth);
- Use tone to inform, not persuade;
- Balance hope with caution.

This divergence reinforces the role of tone as a diagnostic signal. In the presence of fraud, tone often becomes non-reflective of operational conditions, serving instead to manipulate perception or obscure deterioration. Such patterns justify the integration of sentiment markers into NLP-based fraud detection frameworks.

3.2.4 Risk Disclosure and Omission

The presence or absence of risk-related language in corporate reporting is a decisive factor in assessing the transparency, accountability, and ethical stance of a company. The strategic omission or minimization of risk disclosure is not merely a stylistic flaw but may constitute a deliberate act of obfuscation aimed at misleading stakeholders. This section explores how fraudulent firms systematically underreport or de-emphasize risk, while healthy companies tend to incorporate risk awareness into their core narrative structure.

Enron Corporation

Enron's 1998 and 1999 annual reports exhibit a notable absence of substantive risk discussion. Despite operating in highly volatile and deregulated energy markets, the reports allocate only a few superficial mentions to terms like *risk*, *volatility*, or *exposure*. Specifically, the word *risk* appears just 6 times in the 1999 report—mostly within generic legal disclaimers at the end of the document. No dedicated "Risk Factors" section is present, and no material references are made to debt leverage, counterparty risk, or regulatory uncertainty, despite the use of complex financial derivatives and off-balance-sheet entities.

In contrast, Enron's non-fraudulent peer AES Corporation dedicates over three full pages to risk-related issues in its 1999 report, explicitly discussing energy price fluctuations, geopolitical instability, and credit exposure. AES uses *risk* 32 times and *liquidity* 18 times, with detailed paragraphs explaining mitigation strategies and scenario planning. This contrast reveals a deliberate strategic choice by Enron to downplay or hide risk realities, substituting them with aspirational and overly optimistic messaging.

Wirecard AG

Wirecard's risk-related disclosures are nominal and formulaic, particularly in its 2018 annual report. While the report contains a "Risk Report" section, the content is largely boilerplate and devoid of company-specific analysis. The word *risk* is mentioned only 9 times, with 7 of these appearing in legal disclaimers or IFRS-standard phrasing. Crucially, no reference is made to key emerging risks such as whistleblower complaints, investigative journalism reports, or third-party vendor scrutiny—all issues that were known to management prior to the collapse.

By comparison, PayPal's 2018 report includes over 40 explicit risk references, covering cybersecurity, regulatory compliance, geopolitical issues, and reputational exposure. PayPal's disclosures are forward-looking and tailored, with clearly articulated risk categories and management responses, signaling a proactive risk culture that Wirecard systematically lacks.

Parmalat S.p.A.

Parmalat's 2000 and 2001 reports are emblematic of concealed risk under the guise of technocratic language. The company provides no structured Risk Factors section and embeds any mention of uncertainty within obscure legal jargon. The term *risk* appears only 7 times in the 2001 report, with no breakdown of credit risk, market exposure, or liquidity status—even as the company's cash flow issues escalated. Core elements such as *governance*, *controls*, and *debt management* are not addressed linguistically, reflecting a deliberate omission of essential information.

In contrast, Nestlé's 2001 report includes a full section on risk management, discussing interest rate exposure, exchange rate variability, and sourcing risks. Words like *uncertainty*, *debt*, *compliance*, and *oversight* are regularly used and explained in operational terms, offering stakeholders a realistic view of risk-reward dynamics.

Luckin Coffee Inc.

Luckin's 2018 IPO prospectus and annual disclosure exhibit clear deficiencies in risk transparency. While the word *risk*appears 12 times, 9 of these occur in legal footnotes or disclaimers. The main body of the report focuses exclusively on growth narratives, expansion metrics, and branding statements. Risk mentions are high-level, non-specific, and always de-emphasized—e.g., *"Risks associated with our expansion may impact performance"*—without quantification or mitigation details. Terms such as *audit*, *governance*, or *controls* are either entirely absent or relegated to footnotes, which is alarming given the falsified revenue figures later uncovered.

By contrast, Starbucks provides comprehensive coverage of both strategic and operational risks. In its 2018 report, Starbucks includes sections on supply chain risk, commodity price volatility, global macroeconomic shifts, and human capital retention. Risk appears over 45 times, accompanied by quantitative disclosures and mitigation frameworks, reinforcing its commitment to transparency.

Tesco PLC

Tesco offers a compelling counterexample of honest disclosure under pressure. In its 2013 and 2014 reports—years of underperformance and restructuring—Tesco dedicates more than 10 pages to risk governance, covering topics from IT and reputational risks to operational disruption. The report includes risk heatmaps, references to the Audit Committee's role, and direct quotes acknowledging failure. For example: *"The profit overstatement... represents a serious breach of trust. Measures are in place to strengthen controls."* Such openness is starkly absent in all fraudulent firms examined.

Structural Patterns and Linguistic Indicators

A cross-case analysis reveals consistent linguistic and structural omissions in fraudulent firms:

- Low frequency of risk-related terms: average of 5–10 per document vs. 30–50 in non-fraudulent firms;
- Absence of a formal Risk Factors section (or reliance on generic content);
- Minimal use of control-related vocabulary: *audit, oversight, compliance* are rarely present;
- Lack of quantitative risk discussion: no stress testing, no scenario planning, no downside simulations;
- Legal insulation framing: use of phrases like *"to the best of our knowledge," "subject to change,"* and *"not materially significant"* to neutralize accountability.

By contrast, non-fraudulent firms:

- Frame risk as a strategic priority, not just a regulatory checkbox;
- Align narrative tone with known challenges;
- Employ technical and transparent vocabulary;
- Include governance mechanisms in risk management narratives.

3.2.5 Specific Case Observations

While the comparative analysis across linguistic indicators has revealed consistent divergences between fraudulent and non-fraudulent firms, certain companies within the dataset exhibited particularly pronounced or idiosyncratic patterns that merit isolated attention. This section presents a focused linguistic analysis of Wirecard AG (2018–2019) and Luckin Coffee Inc. (2019–2020)**,** highlighting the narrative strategies and linguistic shifts that go beyond the general fraud-related tendencies outlined so far.

Wirecard AG: Narrative Escalation and Linguistic Overcompensation

Wirecard's 2018 annual report represents a case of extreme rhetorical inflation. The tone across all sections—CEO letter, MD&A, and supervisory board report—is overwhelmingly celebratory, employing phrases such as:

- "We can look back on an extremely successful 2018 fiscal year both from a technological and fundamental perspective."
- "Wirecard is one of the world's fastest-growing digital platforms."

Key sentiment terms like *success* (42 times), *growth* (59), *innovation* (33), and *platform* (52) appear with high frequency, particularly concentrated in the first 20% of the text, as shown through Voyant's *Trends* tool.

Crucially, risk-related terminology is both marginal and decontextualized: the word *risk* appears only 9 times in the entire 2018 report—most of which are found in generic disclaimers such as:

"No material impacts are expected on net assets."

This disproportion signals a deliberate effort to minimize external threats while amplifying perceived internal strengths. Furthermore, the collocates of "compliance" (e.g., "management board," "supervisory body") are institutional but content-empty, reflecting a governance rhetoric unsupported by operational substance.

In 2019, despite growing external investigations, the promotional style persists with subtle shifts. The Q3 2019 interim report remains technically worded, yet continues to emphasize "strong growth," "expanding ecosystems," and "ongoing innovation"—suggesting an effort to sustain investor confidence under pressure.

Luckin Coffee Inc.: Post-Fraud Defensive Language and Narrative Rewriting

The case of Luckin Coffee offers a rare glimpse into post-disclosure linguistic behavior. The 2018 report, preceding the fraud's exposure, exhibits highly promotional and speculative rhetoric: *growth* (64, times), *scale* (39), *expansion* (41), and *disruption* dominate the lexicon. Yet key financial terms like *audit* (1 time), *governance* (0), and *compliance* (0) are practically absent. Modal phrases such as *"we believe we are well-positioned to…"* occur 27 times, indicating a speculative orientation unaccompanied by verifiable metrics.

However, in the 2020 restatement report following the public fraud admission, a radical rhetorical inversion takes place:

- The promotional vocabulary is largely purged.
- Defensive phrases appear with high density: *"subject to audit," "we are taking remedial steps," "as determined by the Special Committee"*.
- The narrative becomes passive, legalistic, and distancing, exemplified by:
  "Certain transactions were fabricated by individuals who have since been terminated."

This evolution reflects an intentional strategy of narrative realignment, aimed at minimizing liability and rebuilding trust through formal, impersonal, and mitigated disclosures.

Strategic Framing and Linguistic Role-Reversal

The two cases illustrate opposite narrative trajectories:

- Wirecard, even under pressure, maintains linguistic overconfidence, refusing to acknowledge instability.
- Luckin, once the fraud is undeniable, pivots to a rhetoric of remediation and victimhood, adopting a post-fraud compliance tone that contrasts sharply with its earlier optimism.

This duality underscores the manipulability of corporate language: tone, framing, and lexical choices are not only predictive of misconduct but also adaptable to its consequences. The shift from visionary ambition to regulatory appeasement marks a full circle in fraud discourse evolution.

Concluding Remarks

These specific cases reinforce the broader analytical insights while offering depth into how language is operationalized strategically across fraud cycles. They also reveal the dynamic plasticity of corporate narratives: when deception is successful, language is promotional and expansive; when deception is exposed, language becomes conservative and absolving. In this sense, Wirecard and Luckin Coffee illustrate two poles of linguistic behavior in fraudulent firms—one attempting to conceal, the other to recover. The ability to track these shifts linguistically, before and after fraud detection, highlights the potential of NLP tools not only for early fraud detection, but also for post-crisis discourse analysis.

3.3 Comparative Summary

The previous sections have examined each linguistic indicator individually, highlighting consistent discrepancies between fraudulent and non-fraudulent firms in terms of word choice, modality, tone, and risk disclosure. This section consolidates those insights into a comparative overview, bringing together all four dimensions to identify recurring patterns across the dataset. The aim is to move beyond isolated observations and toward a holistic understanding of how linguistic strategies co-occur in deceptive corporate narratives.

The following subsection presents a cross-indicator synthesis, offering both qualitative and quantitative evidence of how fraudulent firms differ from their non-fraudulent peers across multiple linguistic axes.

3.3.1 Cross-Indicator Patterns

The analysis presented in the preceding sections has highlighted significant divergences in linguistic behavior between fraudulent and non-fraudulent companies when evaluated along four key dimensions: word frequency and lexical patterns, modal verbs and vague language, sentiment orientation, and risk disclosure. While each of these indicators provides valuable insights on its own, it is through cross-indicator synthesis that the most conclusive patterns emerge.

This section offers a holistic comparison of the fifteen firms in the dataset—five fraudulent and ten non-fraudulent—based on how they behave across all four linguistic axes. The objective is to identify recurring discursive configurations that characterize deceptive corporate communication and to contrast these with the more measured and transparent narrative styles of firms that did not engage in accounting fraud.

Fraudulent Firms: Consistent Multi-Dimensional Deviation

The five fraudulent firms—Enron, Parmalat, Wirecard, Luckin Coffee, and Tesco (2012–2014)—display a consistent pattern of simultaneous linguistic anomalies. In all cases, their reports show:

- Excessive frequency of promotional terms, such as *growth*, *leadership*, *innovation*, and *opportunity*, often in contexts that exaggerate performance or strategic positioning.
- Heavy reliance on modal verbs and vague expressions, including *may*, *might*, *could*, and *we believe*, used to hedge responsibility and avoid specificity.
- Overly positive or evasively neutral tone, with little to no acknowledgement of operational challenges, financial risks, or regulatory exposure.
- Minimal or absent risk disclosure, with the term *risk* either omitted entirely or confined to legal disclaimers lacking substantive analysis.

These indicators do not operate in isolation. For example, Wirecard's 2018 report combines high promotional word density (*platform*, *ecosystem*, *scalable*), with over 50 modal constructions and a triumphant tone, while including only 9 mentions of *risk*, none of which address financial transparency or internal control. Similarly, Enron's 1999–2000 filingsconstruct an image of unstoppable growth, bolstered by vague strategic assertions and legal hedging, while omitting direct reference to liquidity issues or regulatory scrutiny.

The synchronization of these signals suggests deliberate narrative engineering. Fraudulent firms do not merely downplay risks—they reconstruct reality linguistically, deploying euphemism, abstraction, and optimism to shift stakeholder focus away from verifiable financial health toward aspirational projections and controlled messaging.

Non-Fraudulent Firms: Narrative Balance and Accountability

In contrast, the ten non-fraudulent firms—AES Corporation, General Electric, Danone, Nestlé, PayPal, Square Inc., Starbucks, Dunkin' Brands, Carrefour, and J Sainsbury—show remarkable consistency in maintaining linguistic integrity across indicators.

These companies:

- Use promotional terms in moderation, often anchored to actual performance data or strategic plans.
- Employ modal verbs carefully, usually in legal contexts or to qualify uncertainty with measurable follow-up.
- Display balanced tone, including honest acknowledgements of market volatility, regulatory change, or performance shortfalls.
- Prioritize risk communication, often with entire sections dedicated to outlining current and anticipated threats, including economic, financial, operational, and legal dimensions.

For instance, AES Corporation's 1999 report uses *risk* 32 times, often in conjunction with scenario planning and policy responses. Nestlé (2017) integrates risk references with operational updates, offering a comprehensive overview of supply chain, geopolitical, and environmental exposure. Even in growth-focused firms like Starbucks, the tone remains grounded, combining brand optimism with explicit risk awareness (e.g., *supply chain disruptions*, *commodity price volatility*).

Comparative Table Overview

To support this comparative overview, Table 3.1 (provided separately in Excel format) summarizes the behavior of each firm across the four dimensions.

Table 3.1 – Cross-Company Comparison of Linguistic Indicators and Risk Flags

Company	Word Frequency (Promo Terms)	Modal Verbs / Vague Language	Tone / Sentiment	Risk Disclosure	Linguistic Red Flags
Enron	Very High	Very High	Overconfident	Very Low (6/year)	✖✖✖✖
Parmalat	High	High	Bureaucratic / evasive	Absent	✖✖✖✖
Wirecard	Very High	High	Triumphalist / inflated	Very Low (9/year)	✖✖✖✖
Luckin Coffee	High	High	Aspirational / vague	Low (4–6/year)	✖✖✖✖
Tesco (2012–2014)	Medium-High	Moderate	Optimistic / inconsistent	Below average	✖✖✖
AES Corporation	Balanced	Moderate	Realistic / reflective	High (32/year)	✔✔✔✔
General Electric	Balanced	Low	Reserved / informative	High	✔✔✔✔
Danone	Balanced	Controlled	Strategic / self-aware	High	✔✔✔✔
Nestlé	Balanced	Controlled	Grounded / constructive	High (36/year)	✔✔✔✔
PayPal	Moderate	Controlled	Neutral / professional	High	✔✔✔✔
Square Inc.	Moderate	Transparent	Transparent / restrained	High	✔✔✔✔
Starbucks	Balanced	Controlled	Brand-focused but measured	Very High (>40/year)	✔✔✔✔
Dunkin' Brands	Moderate	Moderate	Neutral / operational	Above average	✔✔✔✔
Carrefour	Moderate	Moderate	Balanced / responsible	High	✔✔✔✔
J Sainsbury	Moderate	Moderate	Honest / pragmatic	High	✔✔✔✔

Source: Author's elaboration

It indicates, for each company, the degree of deviation from expected linguistic norms and assigns a qualitative assessment of their overall linguistic integrity. Fraudulent firms consistently receive full red-flag profiles (✖✖✖✖), while healthy firms maintain full compliance (✔✔✔✔), with Tesco flagged partially due to its transitional reporting style between 2012 and 2014.
This table underscores a critical point: it is not the presence of one red flag that matters, but rather the convergence of multiple linguistic risk indicators that distinguishes deceptive communication from transparent reporting. The strength of the pattern lies in its recurrence across companies, industries, and time periods.

Implications for Fraud Detection and Disclosure Stan
The findings of this section have broader implications. First, they reinforce the validity of linguistic profiling as a diagnostic tool for detecting early signs of corporate misconduct. Second, they challenge the assumption that stylistic choices in financial communication are incidental; instead, they appear to be systematic, replicable, and predictive of ethical posture.

By identifying this pattern of narrative inflation, hedging, and risk suppression, the study lays the groundwork for the development of AI-driven systems capable of flagging high-risk reporting behavior. These systems could be trained not only to detect individual linguistic anomalies, but to recognize multi-indicator convergence as a more reliable signal of potential fraud.

The following subsection (3.3.2) explores each linguistic dimension in further depth, offering side-by-side comparisons across companies to trace the exact linguistic divergences and their operational consequences.

3.3.2 Indicator-by-Indicator Comparative Insights

While Section 3.3.1 highlighted the convergence of linguistic anomalies within fraudulent narratives, this subsection disaggregates the findings to examine how each individual linguistic indicator varies systematically across the dataset. The objective is to unpack the mechanics of fraud-related communication by contrasting fraudulent versus non-fraudulent patterns in a structured, dimension-by-dimension analysis. The discussion is supported by Table 3.2, available in the appendix and provided in Excel format for clarity and expansion.

Word Frequency and Lexical Patterns

Fraudulent firms exhibit a distinctive overuse of promotional, aspirational, and abstract terminology. Terms like *growth*, *innovation*, *leadership*, *platform*, and *opportunity* are used with disproportionate frequency and often without quantitative anchoring. For instance:

- Enron's 1999 report uses *growth* 63 times and *opportunity* 38 times, with minimal reference to *risk* or *volatility*.
- Wirecard's 2018 filing contains 52 instances of *platform* and 33 of *innovation*, suggesting a semantic inflation strategy to bolster investor confidence.

In contrast, non-fraudulent firms maintain a more balanced lexical field. They contextualize promotional language with operational or strategic data. Nestlé's 2017 report, for example, references *sustainability* and *growth* but balances this with over 30 mentions of *risk*, *supply chain*, and *regulatory complexity*. This approach reflects a more responsible and data-informed narrative.

Modal Verbs and Vague Language

Modal verb usage sharply differentiates deceptive from transparent communication. Fraudulent reports show a high concentration of hedging expressions, such as:

- *may*, *might*, *could* (used ambiguously in forward-looking sections),
- *we believe*, *we expect*, *we are positioned to* (appearing without concrete plans or KPIs).

For example:

- Luckin Coffee's 2018 report includes over 60 vague constructions, e.g., *"we believe we can disrupt traditional coffee culture,"* without supporting metrics or cost analysis.
- Enron's 2001 filing uses *may* 33 times, many of which occur near references to regulatory issues, introducing linguistic distance from accountability.

Non-fraudulent companies, on the other hand, use modals sparingly and legally, typically in formal risk disclosures. AES Corporation and Danone employ modals like *may* or *could* in clearly defined risk contexts, followed by mitigation strategies or financial modeling. This evidences a normative use of vagueness tied to compliance rather than narrative manipulation.

Tone and Sentiment Orientation

Tone analysis reveals a clear polarity between fraudulent and healthy firms. Fraudulent narratives are often:

- Triumphalist and exaggerated (Wirecard),
- Overconfident or euphoric (Enron),
- Aspirational and ungrounded (Luckin Coffee),
- Or overly bureaucratic and emotionally neutral (Parmalat).

These tones are often inconsistent with actual performance trends. For example, Tesco's 2014 report adopts a strategic tone (*"we believe our plan will deliver long-term shareholder value"*) despite financial underperformance and brewing scandal.

By contrast, non-fraudulent companies tend to reflect reality with measured and multidimensional tones. Firms like Tesco (pre-2014), Carrefour, and General Electric acknowledge shortcomings openly, using language such as *"a year of transition," "difficult market conditions,"* and *"cost containment challenges."* These narratives are emotionally calibrated and more aligned with empirical results.

Risk Disclosure and Omission

Risk language emerges as the clearest and most quantifiable divergence. Fraudulent firms display systematic omission or minimization of risk-related terminology:

- Enron, Parmalat, and Wirecard averaged fewer than 10 uses of *risk*, with most of those limited to disclaimers or generic legal language.
- Parmalat’s 2001 report contains no references to financial risk in the strategic section, despite the company facing severe liquidity constraints.

In contrast, healthy firms treat risk as a core element of narrative transparency. Reports from Nestlé, Danone, and AES regularly include entire sections on *currency risk*, *political risk*, *supply chain disruption*, and *commodity price fluctuation*. Starbucks explicitly outlines over 40 distinct risk factors in its 2018 disclosure.

Furthermore, non-fraudulent firms integrate risk vocabulary contextually and functionally, often pairing it with action-oriented language (*“to address these risks, we have adopted…”*), demonstrating operational awareness and communicative accountability.

Table 3.2 – Summary of Linguistic Behavior by Indicator

Indicator	Common Fraudulent Patterns	Common Non-Fraudulent Patterns
Word Frequency	Overuse of promotional terms like 'growth', 'innovation', 'leadership'; often used in vague or inflated contexts without supporting data.	Balanced vocabulary with concrete terms; promotional terms are contextualized with operational details or metrics.
Modal Verbs	High frequency of modals ('may', 'might', 'could') and vague phrases ('we believe', 'we expect'), especially in strategic and risk-related sections.	Modals used selectively and appropriately, often in legal or regulatory sections; less frequent in strategic disclosures.
Tone / Sentiment	Tone often overconfident (Enron, Wirecard) or evasive/neutral (Parmalat); mismatch with operational reality; little negative sentiment.	Tone reflects both achievements and challenges; sentiment aligns with performance; includes introspection and caution.
Risk Disclosure	Risk language minimal or absent; few mentions of 'risk', 'uncertainty', 'volatility'; when present, confined to disclaimers.	Rich and explicit risk sections; regular use of risk terms in context; disclosure is proactive and often scenario-based.

Source: Author's elaboration

This table distills the key patterns discussed above and contrasts them across the four indicators. It illustrates how fraud-oriented companies construct their narratives through thematic inflation, linguistic evasion, tonal exaggeration, and risk suppression, while healthy firms opt for a discursive strategy based on balance, moderation, and disclosure fidelity.

Conclusion: Indicators in Isolation vs. Convergence

When considered in isolation, each of these indicators offers a partial view into corporate communication behavior. However, it is their combined expression—when a company scores anomalously across all four dimensions—that the linguistic signature of fraud becomes most evident. This multi-indicator alignment creates a cumulative linguistic risk score, which can be operationalized in AI-supported fraud detection systems.

The next section (3.4) will contextualize these findings within the broader literature on corporate disclosure and fraud theory, exploring both the interpretive value and the limitations of linguistic profiling in real-world financial oversight.

3.4 Discussion of Findings

The results of the linguistic analysis presented in this chapter demonstrate that fraudulent companies tend to exhibit systematic and multifaceted anomalies across all examined dimensions of narrative disclosure. These findings not only support but expand upon previous academic literature suggesting that deception in corporate reporting is seldom random or stylistic; rather, it is often strategically constructed to serve communicative objectives such as impression management, stakeholder manipulation, and liability deflection.

3.4.1 Alignment with Existing Literature

The recurring use of overly optimistic language, hedging modals, and omission of risk information reflects the mechanisms identified in prior work on impression management in financial reporting. As Clatworthy and Jones (2003) argue, corporate narratives often serve to "frame" the firm's performance in a favorable light, especially in contexts of uncertainty. Our findings empirically substantiate this claim, revealing that companies like Enron, Wirecard, and Luckin Coffee not only frame their performance positively but do so using a consistent linguistic toolkit that includes promotional terms, vague strategic forecasts, and the systematic erasure of cautionary language.

These observations also resonate with the Fraud Triangle model (Cressey, 1953), which identifies pressure**,** opportunity**,** and rationalization as key enablers of fraud. In a linguistic context, the overuse of modal verbs (*may, could, might*) and subjective phrases (*we believe, we expect*) may be seen as rhetorical proxies for rationalization—discursive techniques that allow managers to distance themselves from concrete responsibility while maintaining a façade of strategic control.

Moreover, the findings provide indirect support for theories of narrative obfuscation (McKenna & Wright, 2012), wherein managers manipulate the interpretability of financial reports by using vagueness, verbosity, or semantic overload. In Parmalat's case, the overly technical and emotionally neutral tone functions as a discursive shield that conceals operational instability behind a wall of jargon.

3.4.2 Theoretical Contribution: Linguistic Convergence as a Fraud Signal

One of the most notable outcomes of this study is the convergence of linguistic anomalies across the fraudulent firms analyzed. Despite operating in different sectors and time periods, these companies exhibit remarkably similar discursive behavior, suggesting the existence of a linguistic fingerprint of deception. This supports the hypothesis that fraud is not only a financial event but also a narrative phenomenon, crafted through language that exploits both regulatory ambiguity and stakeholder heuristics.

Non-fraudulent firms, by contrast, consistently demonstrate linguistic moderation, emotional calibration, and transparency in risk disclosure—even when performance is suboptimal. This differential suggests that language can function as a diagnostic variable, capable of distinguishing between genuine and manipulative communication practices.

3.4.3 Methodological and Practical Implications

From a methodological perspective, the study reinforces the value of combining quantitative NLP tools (e.g., Voyant) with qualitative interpretive layers (e.g., ChatGPT-assisted close reading). This hybrid approach allows for both pattern recognition and contextual sensitivity—two elements that are often missing from purely statistical models.

Practically, the identification of multi-indicator linguistic profiles offers a valuable foundation for the development of AI-supported early warning systems in auditing and regulatory oversight. By training models to detect combinations of linguistic red flags (e.g., excessive modals + low risk terms + promotional tone), it is possible to flag potentially deceptive disclosures before financial inconsistencies become visible in the numerical data.

3.4.4 Limitations and Critical Reflections

Despite the robustness of the results, several limitations must be acknowledged:

1. Subjectivity in tone and vagueness classification: Even with structured criteria, sentiment and modality judgments carry an inherent interpretive bias. The use of LLMs mitigates this issue but does not eliminate it entirely.
2. Sector-specific language norms: Certain industries (e.g., tech, retail) may naturally employ more aspirational or brand-focused language, which can overlap with patterns flagged as deceptive. While the 1:2 peer-matching design helps control for this, it cannot neutralize all contextual effects.
3. Temporal constraints: The analysis focuses on documents produced prior to the public exposure of fraud. It is possible that language manipulation intensifies closer to the disclosure event, a dynamic not fully explored in this work.
4. Generalizability: The sample, though diverse, is limited to 15 companies. Larger datasets, including firms with suspected but unproven fraud, would allow for more granular typologies and predictive modeling.

3.4.5 Conclusion: Toward a Linguistic Ethics of Financial Disclosure

Ultimately, this analysis reinforces the idea that fraud is not just an act, but a narrative—one that is embedded in how companies choose to communicate with their stakeholders. Language, in this context, becomes both a symptom and a tool of misconduct. Recognizing this dual role opens the door to a more nuanced understanding of fraud, one that is not only financial and legal, but also discursive and ethical.

The next and final section (3.5) will briefly summarize the core results of this chapter and transition toward Chapter 4, where the broader implications for AI-assisted fraud detection will be explored.

3.5 Conclusion

The third chapter of this thesis has provided a detailed, multi-layered analysis of linguistic behavior in corporate financial reporting, comparing five fraudulent firms with ten non-fraudulent peers across four key narrative indicators: word frequency, modal and vague language, tone/sentiment, and risk disclosure. The results show that fraudulent companies do not simply differ in what they report—they diverge systematically in how they communicate, using language as a tactical medium for obfuscation, persuasion, and control.

3.5.1 Key Empirical Findings

The findings are consistent and robust across companies, sectors, and years.
Table 3.3 summarizes the most prominent patterns identified across each dimension:

Table 3.3 – Summary of Linguistic Behavior by Indicator in Fraudulent vs. Non-Fraudulent Firms

Indicator	Fraudulent Firms	Non-Fraudulent Firms
Word Frequency	Overuse of promotional terms (e.g., 'growth', 'leadership')	Balanced vocabulary with operational and strategic terms
Modal/Vague Language	High density of hedging verbs and speculative phrasing	Modals used legally or cautiously, with quantification
Tone/Sentiment	Overconfident, euphoric, evasive, or technocratic	Reflective, balanced, aligned with actual performance
Risk Disclosure	Absent, minimal, or generic; often in legal disclaimers	Proactive, scenario-based, integral to strategic sections

Source: Author's elaboration

In all five fraudulent companies—Enron, Parmalat, Wirecard, Luckin Coffee, and Tesco—the presence of multiple concurrent red flags was observed. These firms construct narratives that are:

- Semantically inflated (high frequency of promotional vocabulary),
- Syntactically evasive (excessive use of vague modal structures),
- Emotionally misaligned (inconsistent tone relative to real performance), and

- Informationally incomplete (deficient risk disclosure).

By contrast, non-fraudulent companies—AES, Danone, Nestlé, PayPal, Carrefour, among others—consistently exhibited linguistic transparency through a careful balance of tone, disclosure of risk, and grounded strategic communication.

3.5.2 Synthesis: Toward a Composite Fraud Profile

The co-occurrence of linguistic anomalies across all four indicators in fraudulent firms suggests the existence of a composite linguistic profile of fraud—a sort of discursive fingerprint. This profile is not merely the sum of individual features but a patterned structure, where rhetorical strategies reinforce each other to construct a false narrative of strength, vision, and control. Such narratives may help delay detection, manipulate investor expectations, or satisfy regulatory requirements without substantive transparency.

The convergence of findings across companies strengthens the case for early linguistic screening models capable of detecting fraud risk based on public financial narratives alone. These insights form the foundation for Chapter 4, where the integration of these linguistic patterns into AI-driven systems is explored.

3.5.3 Analytical Integrity and Methodological Strengths

This chapter also demonstrates the viability of a hybrid methodological framework, combining:

- Quantitative analysis (term frequencies, sentiment trends),
- Qualitative reading (close-textual interpretation),
- Large Language Model support (for nuance, context, and framing detection).

The approach enables the detection of both surface-level irregularities (e.g., word counts, polarity) and deep discourse strategies (e.g., hedging, impression management, tone misalignment). This dual-layer analysis enhances interpretability and offers replicable patterns useful for automation.

3.5.4 Limitation Revisited

While the evidence is compelling, the limits discussed in Section 3.4 remain relevant:

- Tone and vagueness are partially subjective, despite structured criteria.
- Industry-specific norms may confound interpretation in some cases.
- Small sample size limits generalizability—though the consistency of red flags across cases compensates in part.
- Temporal effects (e.g., pre-fraud linguistic shifts) could be further explored with longer time horizons.

Nevertheless, the results clearly establish that language is not a neutral channel in financial communication. It is, rather, an active field in which ethical posture, strategic intent, and transparency are performed and negotiated.

The insights gained from this comparative analysis provide a robust empirical base for building AI-supported fraud detection systems grounded in linguistic markers. Chapter 4 will translate these findings into practical applications, exploring how Natural Language Processing can be operationalized for predictive analysis in auditing, regulatory monitoring, and internal corporate governance.

By formalizing the linguistic fingerprints identified in Chapter 3 into structured, programmable criteria, the next stage of this research aims to bridge the gap between narrative detection and algorithmic implementation, laying the groundwork for a new generation of disclosure analytics in the fight against corporate fraud.

CHAPTER 4 - IMPLICATIONS, AI-BASED APPLICATIONS, AND FINAL CONCLUSION

Building upon the empirical findings presented in Chapter 3, this chapter translates the identified linguistic patterns into actionable insights, strategic applications, and forward-looking considerations. The objective is twofold: to contextualize the relevance of these patterns within the broader domains of auditing, regulation, and financial analysis, and to propose concrete models for integrating Natural Language Processing (NLP) features into AI-driven fraud detection systems.

The preceding analysis demonstrated that fraudulent firms systematically diverge from their non-fraudulent counterparts across multiple linguistic dimensions—namely, word frequency, modal usage, tone, and risk disclosure. These divergences are not incidental; rather, they reflect deliberate narrative strategies that serve to obscure risk, deflect accountability, and construct a false image of corporate stability. If appropriately harnessed, such signals can serve as powerful early indicators of misconduct.

This chapter is structured in five sections. Section 4.1 outlines the strategic implications of narrative analysis for auditors, analysts, and regulatory authorities, advocating for the inclusion of linguistic red flags in existing fraud detection frameworks. Section 4.2 presents a framework for operationalizing NLP features within AI models, detailing the technical integration of linguistic indicators into machine learning workflows. Section 4.3 discusses the regulatory, ethical, and practical challenges associated with AI-driven narrative surveillance, especially in the context of data privacy, algorithmic bias, and model interpretability. Section 4.4 explores promising directions for future research, including the adoption of transformer-based architectures, multilingual datasets, and hybrid models that merge structured and unstructured data. Finally, Section 4.5 synthesizes the key contributions of the chapter and offers concluding reflections on the broader implications of this thesis.

By bridging theoretical insights with applied methodologies, this chapter aims to contribute to the development of more transparent, intelligent, and linguistically-aware fraud detection systems—anchored in both technological innovation and ethical responsibility.

4.1 Strategic Implications for Fraud Detection and Auditing

The findings from Chapter 3 clearly demonstrate that corporate narratives carry distinctive linguistic signatures that can function as early signals of fraudulent intent. These narrative choices—such as the overuse of promotional language, modal vagueness, evasive tone, and inadequate risk disclosure—are not mere stylistic preferences but strategic tools employed to construct misleading representations of corporate reality. Understanding these patterns holds significant implications for financial auditing, internal control systems, and regulatory supervision.

4.1.1 The Role of Narrative Information in Auditing

Traditional auditing approaches tend to emphasize numerical consistency and compliance with accounting standards. However, the integration of narrative analysis into audit procedures can enrich this process by capturing qualitative discrepancies that precede numerical anomalies. Research by Boskou et al. (2019) illustrates that textual analysis of internal audit disclosures can be as predictive of audit quality as financial indicators, showing that firms with higher linguistic coherence in disclosures are more likely to engage Big Four audit firms, a common proxy for audit quality (Boskou, Kirkos, & Spathis, 2019).

Furthermore, empirical studies reveal that the tone and structure of narrative disclosures—especially in sections such as the Management Discussion and Analysis (MD&A)—are systematically associated with both perceived and actual risk. For example, banks that employ more negative or evasive tones in their disclosures are shown to face higher market-assessed risk, particularly during periods of economic policy uncertainty (Boskou, Kirkos, & Spathis, 2019).

4.1.2 Enhancing Traditional Red Flags

Integrating linguistic red flags with accounting-based red flags (e.g., debt-to-equity ratios, receivables anomalies) allows for a more holistic fraud risk framework. As shown in this thesis, fraudulent firms such as Enron or Wirecard did not differ substantially in numerical performance from their peers, but diverged starkly in the way they communicated—often relying on euphoric, vague, or evasive language while omitting concrete risk discussions (Boskou, Kirkos, & Spathis, 2019).

This finding supports previous assertions that the language of annual reports is a reliable proxy for internal transparency. Auditors and regulators can therefore use composite linguistic profiles as an additional dimension of risk assessment—particularly in contexts where conventional metrics have limited diagnostic power.

4.1.3 Practical Use for Auditors, Analysts, and Supervisors

For internal auditors, this implies a shift toward including discursive evaluations in the internal control matrix. Natural Language Processing (NLP) tools such as sentiment analysis, modal detection, and semantic clustering can be integrated into internal audit platforms to flag anomalies in real time.

For financial analysts and external auditors, linguistic profiling offers an early-stage alert system. For example, unusually high usage of promotional language ("growth", "innovation", "leadership") coupled with low or generic risk disclosures should raise the fraud-risk rating of a firm, even in the absence of obvious accounting discrepancies.

Regulatory bodies could also institutionalize these indicators. As shown by Albert et al. (2025), narrative tone is not only a mirror of corporate sentiment but an actionable predictor of institutional risk. Embedding narrative features in supervisory reviews may enhance early-warning systems, particularly under regimes of economic instability (Albert, Ibeji, Owusu, & Nguyen, 2025).

4.1.4 Toward a Linguistic Risk Index

Given the convergence of findings across multiple companies, sectors, and timeframes, there is potential to formalize a Linguistic Risk Index (LRI) — a composite score derived from features such as promotional density, modal ambiguity, sentiment polarity, and risk disclosure frequency. This index could be embedded in AI-powered audit and surveillance tools, enabling automated red-flagging of annual reports and earnings call transcripts.

4.2 Integrating NLP Features into AI Models

The integration of Natural Language Processing (NLP) features into AI-driven models represents a crucial step toward operationalizing the insights gained from narrative analysis in financial disclosures. By converting qualitative linguistic patterns into structured data, it becomes possible to develop machine learning models capable of identifying fraud-prone behavior in corporate communication. This section presents a conceptual framework for such integration, describes linguistic feature selection processes, discusses appropriate modeling techniques, and illustrates a simulated implementation pipeline.

4.2.1 A Framework for AI-Driven Fraud Detection

The development of an effective AI-driven fraud detection system grounded in linguistic analysis involves three core stages: (1) text preprocessing and feature extraction, (2) model training using labeled datasets, and (3) prediction and risk scoring. Sun et al. (2023) proposed a framework based on Part-of-Speech (POS) tagging and risk disclosure analysis, showing that linguistic signals can significantly improve fraud detection performance when added to traditional financial indicators (Sun, Li, & Zhu).

Papasavva et al. (2025) further emphasized the importance of pipeline structuring, which includes careful data cleaning, semantic normalization (e.g., lemmatization, stemming), and feature vectorization. Their review outlines best practices in building scalable and explainable NLP pipelines, especially in fraud detection contexts involving large and noisy textual datasets (Papasavva, Johnson, Lowther, & al.).

4.2.2 Linguistic Feature Selection: Extraction and Relevance

Feature engineering is at the core of transforming textual content into useful predictors. Among the most promising NLP-derived features for fraud detection are:

- Lexical features: word frequency, length of disclosure, average sentence complexity.
- Syntactic features: POS distributions (e.g., excessive use of adjectives or modals).
- Semantic features: polarity, uncertainty, risk-related word density.
- Discourse features: structure and coherence of financial narratives.

As highlighted by Ali et al. (2022), NLP features such as sentiment and tone, especially when derived from sections like MD&A and Risk Factors, offer incremental predictive power over numerical data alone (Ali, Abd Razak, Othman, & al.). Sun et al. (2023) demonstrated that POS-based tagging of risk disclosures can improve model precision by isolating stylistic cues indicative of strategic obfuscation (Sun, Li, & Zhu).

Advanced embedding methods, such as Word2Vec, BERT, and DistilBERT, have also been tested for semantic representation. Chang et al. (2022) showed that DistilBERT combined with SVM or Random Forest yields high accuracy while keeping computational costs low, making it viable for real-time applications like fraud monitoring systems (Chang, Yen, & Hung).

4.2.3 Integrating NLP Features into ML Algorithms

Once extracted and quantified, NLP features can be integrated into various supervised machine learning algorithms. Among the most commonly used in financial fraud detection are:

- Random Forest (RF): robust to overfitting and effective with high-dimensional feature spaces.
- Support Vector Machine (SVM): particularly effective in handling small to medium-sized imbalanced datasets.
- Logistic Regression (LR) and XGBoost: for interpretable baselines and boosting performance.

- LSTM and deep learning models: suitable for capturing temporal and sequential linguistic patterns.

Ali et al. (2022) provided a systematic review showing that RF and SVM outperform deep models when combined with simple textual features, whereas LSTM models are preferable for sequential data, such as longitudinal annual reports (Ali, Abd Razak, Othman, & al.). Chang et al. (2022) confirmed that hybrid models using linguistic embeddings (BERT, ELMO) improve classification accuracy when detecting suspicious financial narratives (Chang, Yen, & Hung).

4.2.4 Simulated Pipeline: From Feature Vector to Fraud Prediction Output

A simplified AI-driven pipeline for fraud detection using narrative data could include the following stages:

- Input: Raw annual report texts (MD&A, Risk Factors).
- Preprocessing: Tokenization, POS tagging (e.g., via spaCy), stop word removal, and lemmatization.
- Feature Extraction: Sentiment scores, frequency of modal verbs, risk word density, POS ratios.
- Feature Vector Construction: Combining linguistic features with traditional accounting ratios.
- Model Training: Random Forest or SVM trained on labeled data (fraudulent vs non fraudulent).
- Output: Binary classification + risk probability score.

Papasavva et al. (2025) emphasized that preprocessing and embedding strategies are critical to model reliability, particularly for unstructured textual data (Papasavva, Johnson, Lowther, & al.). Adavelli (2024) proposed using Generative AI to produce synthetic fraud scenarios, enriching training datasets in low-data environments while maintaining compliance with privacy regulations (Adavelli).

4.3 Regulatory, Ethical, and Practical Considerations

The integration of Artificial Intelligence (AI) into fraud detection and auditing systems has introduced a wide array of opportunities, but it also raises significant regulatory, ethical, and practical challenges. As AI tools become more embedded into financial decision-making, concerns related to transparency, accountability, interpretability, and social responsibility must be thoroughly addressed to ensure both trust and compliance with legal frameworks.

4.3.1 Regulatory Challenges

One of the most pressing concerns in the adoption of AI for fraud detection lies in the regulatory domain. Existing frameworks, such as the International Standards on Auditing (ISA), were not designed with autonomous or semi-autonomous algorithmic systems in mind. Consequently, there is a regulatory gap concerning the explainability and traceability of AI-generated outcomes. The opacity of complex models—particularly deep learning architectures—can obstruct the auditor's duty to justify conclusions and decisions in accordance with established standards and legal procedures (Tan, et al., 2023).

Moreover, the application of AI in fraud analytics intersects with data protection regulations, such as the General Data Protection Regulation (GDPR). Auditors must ensure that AI systems comply with principles like data minimization, fairness, and accountability—especially when processing sensitive or personal financial data. Without adequate regulation and oversight, the risk of misuse, liability fragmentation, and public mistrust increases significantly.

4.3.2 Ethical Risks and Accountability

The ethical dimensions of AI-based fraud detection are multifaceted. AI systems are incapable of moral reasoning and cannot substitute for human ethical judgment in decision-making processes that may have profound consequences. Scholars have warned that over-reliance on algorithms may dilute professional responsibilities such as auditor scepticism and discretion.

Ethical risks also emerge from biased training data and algorithmic design. AI models trained on unbalanced datasets may replicate or even exacerbate existing inequities, leading to the misclassification of companies or individuals and unjustified red flags. This is particularly problematic in fraud detection, where decisions may carry legal or reputational consequences.

To mitigate these risks, researchers recommend embedding human-in-the-loop oversight in all critical audit processes. Auditors must retain the authority and responsibility to interpret, challenge, and override AI-based outputs when necessary. This ensures that decisions are not only technically sound but also ethically justifiable and contextually grounded (Binh, 2025).

4.3.3 Practical Implementation Considerations

The practical deployment of AI in auditing raises substantial operational challenges. First, interpretability remains a key obstacle. Many AI models, especially those based on neural networks, are often referred to as "black boxes" due to their complex internal logic. For an audit to be defensible and aligned with professional standards, the rationale behind AI-generated decisions must be transparent and understandable to both auditors and stakeholders.

Second, data quality and governance are fundamental. AI models rely on high-quality, well-structured, and representative data to function effectively. Missing values, label noise, or data imbalance can severely compromise the model's accuracy and fairness (Ismail & Haq, 2024).

Third, AI integration requires significant investments in infrastructure, talent, and change management. Organizations must train their auditing teams in AI literacy and model governance, while also developing interdisciplinary collaboration with data scientists and legal experts. Without these capabilities, the benefits of AI may be nullified by technical failure or compliance risks.

4.3.4 Toward Responsible AI Governance in Auditing

To address these multifaceted concerns, a number of scholars and institutions advocate for the development of multi-layered AI governance frameworks (Tan, et al., 2023). These include:

- Technical safeguards, such as bias detection, audit trails, and explainable AI models;
- Ethical safeguards, including fairness assessments, stakeholder consultation, and alignment with professional values;
- Organizational safeguards, such as human oversight, ethical training, and role clarity between machines and auditors.

Additionally, external oversight mechanisms—such as independent audits of AI systems, algorithmic impact assessments, and third-party certifications—are emerging as crucial components for ensuring transparency and legitimacy (Raji, Xu, Honigsberg, & Ho, 2022). In sum, while AI offers powerful tools for enhancing fraud detection and audit efficiency, its implementation must be accompanied by careful consideration of regulatory constraints, ethical responsibilities, and practical feasibility. Only by embedding responsibility, transparency, and human oversight into AI systems can the auditing profession safeguard its integrity and public trust.

4.4 Future Research Directions

The field of AI-powered fraud detection using Natural Language Processing (NLP) in financial reporting is rapidly evolving, yet several limitations remain that warrant further investigation. This section outlines key research avenues that can deepen understanding, improve generalizability, and enhance the technical and ethical robustness of fraud detection models.

4.4.1 Expanding the Dataset: Languages, Industries, and Time Horizons

A core limitation in current studies is the narrow scope of available datasets, which often focus on English-language reports from a limited set of industries and countries. Future research should prioritize the expansion of datasets to include multilingual disclosures, cross-industry narratives, and longer time windows. As highlighted by Lo et al. (2023), expanding the corpus of financial texts enhances the diversity of linguistic patterns and contextual variability, improving the robustness of NLP applications in finance (Applications, From ELIZA to ChatGPT: The Evolution of Natural Language Processing and Financial, 2023).

Luo et al. (2024) recommend the systematic inclusion of decentralized finance (DeFi) reports, project documentation, and smart contract narratives, which can reveal fraud schemes at various stages of project development. Their taxonomy illustrates how fraud patterns evolve across life cycle stages, suggesting that temporal coverage is essential for modeling longitudinal shifts in narrative deception (Applications, From ELIZA to ChatGPT: The Evolution of Natural Language Processing and Financial, 2023).

4.4.2 Leveraging Transformer-Based Architectures (e.g., BERT, GPT)

Transformer-based language models, such as BERT, RoBERTa, and GPT, have revolutionized NLP by enabling the capture of long-range dependencies and contextual subtleties. These models have already demonstrated superior performance in financial text classification, sentiment extraction, and anomaly detection tasks. Lo et al. (2023) highlight that fine-tuned transformer models significantly outperform traditional models on tasks such as asset management and disclosure risk classification (Applications, From ELIZA to ChatGPT: The Evolution of Natural Language Processing and Financial, 2023).

Moreover, Sarkar et al. (2025) show that transformer models are not only effective in structured data environments, such as U.S. healthcare fraud detection, but also in unstructured claim narratives, offering scalable and explainable AI frameworks that can be adapted for financial reporting use cases (AI-Driven Machine Learning for Fraud Detection and Risk Management in U.S. Healthcare Billing and Insurance , 2025).

4.4.3 Integrating Structured and Unstructured Financial Data

Combining financial ratios and accounting metrics with narrative disclosures is a promising avenue for building hybrid models that reflect the full informational spectrum of corporate reporting. Sharma (2023) stresses the need for real-time systems that simultaneously process structured and unstructured data, enhancing both precision and timeliness in fraud alerts (The Future of Automation in Financial Technology: Leveraging AI to Enhance Fraud Detection and Risk Management, 2023).

Future systems should incorporate transformer models that operate jointly on financial statements and associated MD&A texts, allowing predictive models to flag discrepancies between numerical performance and narrative tone. This approach is especially valuable for early detection, where accounting irregularities may not yet manifest in ratios but are linguistically evident.

4.4.4 Post-Fraud Narrative Evolution: A Longitudinal Perspective

Another underexplored area is the evolution of corporate language after the discovery of fraud. Changes in tone, specificity, and rhetorical strategies can reveal attempts to rebuild stakeholder trust or divert legal attention. Luo et al. (2024) suggest that narratives shift from aggressive promotion to defensive clarification or legal disclaimers in the aftermath of fraud exposure (Applications, From ELIZA to ChatGPT: The Evolution of Natural Language Processing and Financial, 2023).

Tracking these shifts longitudinally would offer a new dimension to narrative risk analysis, helping to distinguish genuine efforts at transparency from strategic obfuscation. Lo et al. (2023) further argue that future studies should explore how advanced models like GPT-4 or FinBERT can be fine-tuned on post-fraud texts to detect subtle shifts in sentiment, modality, and risk framing.

4.5 Final Conclusion

This thesis set out to explore whether corporate fraud—typically conceived as a financial and legal phenomenon—could also be identified through the linguistic fabric of financial disclosures. By combining Natural Language Processing (NLP) techniques with AI-based modeling, the study demonstrated that fraudulent firms exhibit distinct and recurring narrative anomalies across key textual dimensions: lexical frequency, modal usage, sentiment tone, and risk disclosure. These findings were not isolated but consistently observed across sectors and jurisdictions, reinforcing the existence of a discursive “fingerprint” of fraud.

Chapter 1 provided the theoretical framework for understanding corporate fraud, including its typologies and landmark cases. Chapter 2 explained the methodology: the creation of a 15-firm dataset (five fraudulent and ten non-fraudulent companies), the collection of textual disclosures, and the NLP pipeline adopted to extract linguistic indicators. Chapter 3 analyzed the results, confirming systematic differences in the narratives of fraudulent firms, and identifying clusters of linguistic red flags. Chapter 4 translated these findings into practical implications for auditing, regulatory supervision, and AI-based detection frameworks, while critically addressing ethical and legal concerns.
The broader contribution of this research lies in three directions:

- Conceptual: It repositions fraud detection within the semiotic domain, recognizing that corporate misconduct is often prefigured not only in numbers but in words.
- Methodological: It offers a replicable hybrid framework that merges statistical NLP with qualitative interpretability, balancing precision with contextual depth.
- Practical: It outlines a feasible roadmap for integrating narrative-based indicators into AI-enhanced auditing tools, allowing for real-time, proactive fraud screening.

Nevertheless, several challenges remain. While linguistic markers can signal red flags, they must be interpreted cautiously, in light of sector-specific norms and disclosure contexts. Moreover, ethical concerns—ranging from model bias to algorithmic transparency—must be actively addressed to ensure that AI-based systems enhance, rather than erode, trust in financial governance.

In conclusion, this thesis demonstrates that AI and NLP can significantly augment traditional fraud detection approaches, offering both theoretical insights and applied strategies. Yet, the real power of these tools lies not merely in automating detection, but in encouraging a broader cultural shift toward transparent, responsible, and linguistically ethical corporate communication. Future research should expand the dataset, refine the AI pipeline, and explore multilingual, cross-cultural applications—bringing us closer to a financial reporting ecosystem where narrative truth is not an afterthought, but a foundation.

BIBLIOGRAPHY

Adavelli, S. R. (n.d.). Generative AI in Digital Insurance: Redefining Customer Experience, Fraud Detection, and Risk Management. *International Journal of Computer Science and Information Technology Research, 5(2), 41–60* .

AI-Driven Machine Learning for Fraud Detection and Risk Management in U.S. Healthcare Billing and Insurance . (2025). *Journal of Computer Science and Technology Studies, 7(1), 188–198*.

Akintoye, I. &. (2013). The Use of Financial Ratio Models to Help Investors Predict and Interpret Significant Corporate Events. *European Journal of Accounting Auditing, 1*, pp. 35–43 .

Albert, A., Ibeji, N., Owusu, F., & Nguyen, T. (2025). Narrative Disclosure Tone and Bank Risk: The Role of Economic Policy Uncertainty . *Managerial and Decision Economics* , pp. 1–20 .

Ali, A., Abd Razak, S., Othman, S. H., & al., e. (n.d.). Financial Fraud Detection Based on Machine Learning: A Systematic Literature Review . *Applied Sciences, 12(19), 9637* .

Allegrini, M., Di Gennaro, L., Healy, P., & Palepu, K. (2023). Parmalat case analysis in the context of the fraud hexagon. *Journal of Risk and Financial Management, 18*(3), 156.

Altman, E. I. (1968). Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy. *The Journal of Finance, 23*, pp. 589–609 .

Applications, From ELIZA to ChatGPT: The Evolution of Natural Language Processing and Financial. (2023). *The Journal of Portfolio Management*.

Bakri, H. H., Mohamed, N., & Said, J. (2017). Mitigating asset misappropriation through integrity and fraud risk elements: Evidence from emerging economies. *Journal of Financial Crime, 24*(2), 242–255.

Beerbaum, D. (2022). *The future of audit after the Wirecard scandal – Consideration of the "Wambach Report"*. Aalto University, Helsinki.

Berger, A. H. (2015). An Introduction to Benford's Law.

Binh, N. T. (2025). Transforming Auditing in the AI Era: A Comprehensive Review. *Information, 16*.

Boskou, G., Kirkos, E., & Spathis, C. (2019). Classifying internal audit quality using textual analysis: The case of auditor selection. *Managerial Auditing Journal, 34*, pp. 924–950 .

Chang, J.-W., Yen, N., & Hung, J. C. (n.d.). Design of a NLP-empowered finance fraud awareness model: The anti-fraud chatbot for fraud detection and fraud classification as an instance . *Journal of Ambient Intelligence and Humanized Computing, 13, 4663–4679* .

Chiu, T., Wang, Y., & Vasarhelyi, M. A. (2020). The Automation of Financial Statement Fraud Detection: A Framework Using Process Mining. *Journal of Forensic and Investigative Accounting, 12*, pp. 86–114 .

Convertio. (n.d.). *TXT Converter* . From Convertio: https://convertio.co/it/txt-converter/

Corporation, E. (n.d.). (U.S. Securities and Exchange Commission (SEC)) From https://www.sec.gov/Archives/edgar/data/1024401/000093646800000118/0000936468-00-000118.txt

Craja, P., Kim, A., & Lessmann, S. (2020). Deep learning for detecting financial statement fraud. *Decision Support Systems, 139*, 113421.

Cressey, D. R. (1953). *Other People's Money: A Study of the Social Psychology of Embezzlement* . Belmont, CA : Wadsworth Publishing.

Dechow, P. M., & Schrand, C. M. (2004). *Earnings Quality.* Charlottesville: The Research Foundation of CFA Institute.

Demestichas, K. A. (2021). Comparative Analysis of Machine Learning-Based Approaches for Anomaly Detection in Vehicular Data. *Vehicles, 3*(2), 171–186.

Dibra, R. (2016). Corporate governance failure: The case of Enron and Parmalat. *European Scientific Journal, 12*(16), 283–290 .

Dijk), A. (.–B. (n.d.). From https://aida.bvdinfo.com

Dijk), O. (. (n.d.). From https://orbis.bvdinfo.com

Durtschi, C. H. (2004). The Effective Use of Benford's Law to Assist in Detecting Fraud in Accounting Data. *Journal of Forensic Accounting, 5*, 17–34.

Emittenti), C. (. (n.d.). From https://www.consob.it/web/area-pubblica/societa-revocate

Enyi, P. E. (2019). *Relational Trend Analysis: A simple and effective way to detect financial statements fraud.* Babcock University .

Examiners, A. o. (2014). *Report to the Nations on Occupational Fraud and Abuse.* ACFE, Austin, TX.

Feng, M., Ge, W., Luo, S., & Shevlin, T. (2021). The role of auditors, fraud mechanisms, and institutional environments in accounting fraud: Evidence from China. *Journal of Accounting and Public Policy, 40*(5).

Fisher, I. E., Garnsey, M. R., & Hughes, M. E. (2016). Natural Language Processing in Accounting, Auditing and Finance: A Synthesis of the Literature with a Roadmap for Future Research. *Intelligent Systems in Accounting, Finance and Management, 23*, 157–214 .

Heikkinen, A. (2021). *Comparative analysis of outlier detection methods in the context of procurement spend analysis.* LUT University, Lappeenranta, Finland.

Inc., L. C. (n.d.). (U.S. Securities and Exchange Commission (SEC)) From https://www.sec.gov/Archives/edgar/data/1767582/000110465921074995/luckin-20f_20201231.htm

Ismail, M. M., & Haq, M. A. (2024). Enhancing Enterprise Financial Fraud Detection using Machine Learning. *Engineering, Technology & Applied Science Research, 14* , pp. 14854–14861 .

Jans, M., Alles, M. G., & Vasarhelyi, M. A. (2014). A field study on the use of process mining of event logs as an analytical procedure in auditing. *The Accounting Review , 89*, pp. 1751–1773 .

Li, S., Wang, G., & Luo, Y. (2022). Tone of language, financial disclosure, and earnings management: A textual analysis of Form 20-F. *Financial Innovation, 8*, pp. 1–15 .

Luo, Y., & Zhou, L. (2020). Textual tone in corporate financial disclosures: a survey of the literature. *International Journal of Disclosure and Governance, 17*, pp. 101–110 .

Möllers, T. M., McCrum, D., & Alderman, L. (2020). Wirecard scandal: A case of regulatory capture? *Journal of Financial Crime, 31*(5), 1160–1172.

Nestlè. (n.d.). From https://www.nestle.com/investors

Nexis), N. U. (n.d.).

Ngai, E. W., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems, 50*, pp. 559–569 .

Nießner, T., Nickerson, R. C., & Schumann, M. (2021). Towards a taxonomy of AI-based methods in financial statement analysis. *Proceedings of the Americas Conference on Information Systems (AMCIS).* Montreal.

Nigrini, M. J. (2012). Benford's Law: Applications for Forensic Accounting, Auditing, and Fraud Detection.

Omidi, M., Masoumi, B., & Ahmadzadeh, M. (2019). The Efficacy of Predictive Methods in Financial Statement Fraud Detection. *Discrete Dynamics in Nature and Society*.

Online-Convert. (n.d.). *Converti documenti in TXT* . From Online-Convert : https://documento.online-convert.com/it/converti-in-txt

OpenCompanyData, O. /. (n.d.). From https://opencorporates.com

Papasavva, A., Johnson, S., Lowther, E., & al., e. (n.d.). Application of AI-based Models for Online Fraud Detection and Analysis . *Crime Science Journal*.

Perols, J., Murthy, U., & Zhang, J. (2017). Detection of emerging financial statement fraud using textual data. *Journal of Emerging Technologies in Accounting, 14*, pp. 55–69 .

Purda, L., & Skillicorn, D. (2015). Accounting Variables, Deception, and a Bag of Words: Assessing the Tools of Fraud Detection. *Contemporary Accounting Research, 32*, pp. 1193–1223 .

Raji, I. D., Xu, P., Honigsberg, C., & Ho, D. (2022). Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance. *AAAI/ACM Conference on AI, Ethics, and Society* .

Ravisankar, P., Ravi, V., Raghava Rao, G. R., & Bose, I. (2011). Detection of financial statement fraud and feature selection using data mining techniques. *Decision Support Systems, 50*, pp. 491–500 .

Reurink, A. (2018). Financial fraud: A literature review. *Journal of Economic Surveys, 32*(5), 1292–1325.

Revocati, B. I.–A. (n.d.). From https://www.borsaitaliana.it

Said, J. (2015). Assessing fraud risk factors of assets misappropriation: Evidence from Iranian banks. *Procedia Economics and Finance, 31*, 919–924 .

Sapelli, G. (2005). *Giochi proibiti: Enron e Parmalat capitalismi a confronto.* Milano: Bruno Mondadori.

Starbucks. (n.d.). From https://www.starbucks.com/investor-relations

Sun, H., Li, J., & Zhu, X. (n.d.). Financial fraud detection based on the part-of-speech features of textual risk disclosures in financial reports. *10th International Conference on Information Technology and Quantitative Management.*

Tan, E., Petit Jean, M., Simonofski, A., Tombal, T., Kleizen, B., Sabbe, M., . . . Willem, P. (2023). Artificial intelligence and algorithmic decisions in fraud detection: An interpretive structural model. *Data & Policy, 5*.

The Future of Automation in Financial Technology: Leveraging AI to Enhance Fraud Detection and Risk Management. (2023). *International Journal of Scientific Research in Engineering and Management, 7(12)*.

Tuck, P., & Hamilton, M. (2023). Fraud in accounting, organizations and society: Extending the boundaries of research. *Accounting, Organizations and Society, 106*, 101473.

Tutino, M., & Merlo, M. (2019). Accounting fraud: A literature review. *Risk Governance and Control: Financial Markets & Institutions, 9*(1), 8–25 .

Wambach, A. (2021). *Wambach Report on Wirecard – Final analysis and recommendations.* German Ministry of Finance, Berlin.

Yu, S.-J. &.-S. (2021). Research Trends in Accounting Fraud Using Network Analysis. *Sustainability, 13*, pp. 1–26 .

Zager, K. &. (2015). The Role of Financial Ratios in Predicting Fraudulent Financial Reporting. *Procedia - Social and Behavioral Sciences, 213*, 639–643.

Zhang, Y. (2006). *An Empirical Analysis of Revenue Manipulation.* Columbia Business Schoo, New York.

Zhou, W., & Kapoor, G. (2011). Detecting evolutionary financial statement fraud. *Decision Support Systems, 50*(3), 570–575 .