Corporate Communication vs Operational Behavior in the US Banking Sector

Nikolaos Kliropoulos
Athens University of Economics and Business, AUEB
Department of International and European Economic Studies, DIEES

Advisor: Vasia Panousi

June 30, 2026

Abstract

This paper examines the congruence between the language and sentiment expressed in the annual financial filings in a sample of 32 large US banks versus the information captured in quantitative non-interest income (NII) data from those filings. Overall, the findings indicate a difference in the pre-2008 versus the post-2013 expansions. In the pre-2008 period, NII was increasing, while both sentiment and language reflected the enthusiasm from the increased profitability generated by underlying securitization activities. In the post-2013 period, however, while NII increases even faster than before, the language and sentiment instead reflect the prudence suggested by post-crisis oversight, even though this oversight may not be strictly enforced anymore. In other words, the language and sentiment in the financial sector may not be reflecting the reality of banks' continued and ever-increasing risk taking. Concerningly, this decoupling is present for Traditional as well as for Shadow Banks.

Email: nikoskliro2@gmail.com. I would like to thank my thesis advisor, Vasia Panousi. I hereby certify that the work carried out and presented in this submitted thesis is exclusively my own. Any information and material included that has been drawn from other sources has been properly cited in this thesis. Furthermore, I am aware that if it is established that the above declaration is not true, my thesis will be annulled.

1 Introduction

This paper examines the congruence between corporate communication and operational behavior in a sample of 32 large US financial institutions. The main question is whether the US banking sector indeed shifted its operations toward prudence and long-term stability following the 2008 financial crisis, or whether banks merely adopted a rhetoric of compliance while maintaining high-risk short-term-profit models.

To that end, the thesis utilizes a hand-collected panel of 10-K filings of 32 major US financial institutions over the 32-year period from 1993 to 2025. The sample includes 15 Traditional Banks (TB) and 17 Shadow Banks (SB). Building upon recent macroeconomic literature that utilizes textual analysis to quantify economic behavior and institutional deliberation, such as Hansen, McMahon, and Prat (2018), the paper tracks the dynamic usage of terminology representing short-term profit versus terminology representing long-term stability. It also utilizes FinBert, the most advanced large language model (LLM) for the determination of financial sentiment. The linguistic analysis results are then compared to extracted bank-level quantitative data on non-interest income (NII). Specifically, NII generated from volatile operations, such as market trading, complex securitization, and investment banking, can be a proxy for actual operational risk appetite. Juxtaposing the linguistic findings against this financial metric allows us to quantify the divergence between stated stability objectives and empirical risk-taking behavior.

A comparison across the linguistic and quantitative methods yields the following main results. First, in the pre-2008 period, bank sentiment was positive, bank language was related to short-term profitability, and bank activities were increasingly NII-related.

Second, the intermediary period from 2008 to 2012 represents a transitional phase characterized by the immediate crisis fallout and the subsequent drafting and implementation of sweeping legislation like the Dodd-Frank Act. During this transition, bank language shifts sharply to prudential and long-term goals, bank sentiment switches to negative, and NII bank activities decline. Entering the post-2013 period, as these new regulatory frameworks became fully institutionalized, bank sentiment is overall slightly positive, bank language remains focused on oversight terminology, but NII continues to increase. In fact, it increases at rates higher than those of the pre-2008 period.

Third, the above findings indicate a difference in the pre-2008 versus the post-2013 expansions. In the pre-2008 period, NII was increasing, while both sentiment and language reflected the enthusiasm from the increased profitability generated by underlying securitization activities. In the post-2013 period, however, while NII increases even faster than before, the language reflects regulations that are no longer strictly enforced and the

sentiment keeps pace with the language, but not with the increasing profitability that is probably being generated from the increased NII activities. In other words, the language and sentiment in more recent years, due to anchoring around regulation legislation, may not be reflecting the reality of banks' continued and ever-increasing risk taking. Concerningly, this decoupling is present for Traditional Banks as well as for Shadow Banks.

Finally, this decoupling of language-sentiment and quantitative data across the two periods highlights the need for increased pressure by central banks for quantitative data collection from the financial sector.

2 Literature Review

Researchers increasingly use Natural Language Processing (NLP) to analyze corporate disclosures. Loughran and McDonald (2011) established the need for domain-specific lexicons to measure financial sentiment accurately, as general dictionaries often misclassify accounting terms. Building on this, Cohen, Malloy, and Nguyen (2020) found that changes in the language of 10-K filings often precede periods of elevated operational risk, suggesting managers alter their reporting style to obscure vulnerabilities.

The application of textual analysis extends beyond corporate filings to broader institutional communication to uncover hidden economic realities. Hansen, McMahon, and Prat (2018) show that shifts in central bank language reflect underlying economic conditions and shape market expectations. Similarly, Schmeling and Wagner (2019) find that the tone of central bank communication affects asset prices and signals future macroeconomic risks. Expanding on this, Jegadeesh and Wu (2015) demonstrate how dissecting Federal Open Market Committee (FOMC) meeting minutes can objectively extract the tone and uncertainty level of specific economic topics, transmitting superior information to the market. More recently, Raviv, Hilscher, and Nabors (2024) applied machine learning to central bank communications, proving that extracted sentiment from the Fed and ECB accurately predicts future policy interest rates. Together, these studies demonstrate that formal financial texts contain highly quantifiable data regarding actual market conditions and institutional intent.

Beyond sentiment, the literature increasingly employs diverse NLP techniques to extract unobservable risks and structural topics from dense, formal texts. For instance, Davis, Hansen, and Seminario-Amez (2020) demonstrate how textual analysis of corporate disclosures can measure firm-level risk exposures to specific macroeconomic shocks, capturing realities that standard quantitative metrics may lag in reflecting. Furthermore, researchers use techniques like Latent Dirichlet Allocation (LDA) to navigate complex

compliance language; O'Neill et al. (2017) apply topic modeling to legislative texts to identify salient legal topics, illustrating how unsupervised machine learning can extract core themes from highly formalized regulatory documents. Even in survey research, as shown by Ferrario and Stantcheva (2022), open-ended text analysis is recognized as a superior method for eliciting agents' true first-order concerns. These methodological advancements validate the approach of using textual patterns to uncover the true operational focus of financial institutions.

To accurately capture context in these dense texts, recent research relies on machine learning models rather than static dictionaries. Araci (2019) and Huang, Wang, and Roache (2020) demonstrated the effectiveness of FinBERT, a language model pre-trained specifically on financial corpora. Their work shows that FinBERT outperforms traditional methods by evaluating the context of financial terminology, allowing it to differentiate between standard accounting reporting and actual market sentiment. This justifies the use of FinBERT as the core sentiment analysis tool in this thesis.

Finally, to compare these linguistic findings against quantitative realities, the literature identifies non-interest income (NII) as a standard metric for financial risk. Stiroh (2004) established that an increased reliance on NII raises revenue volatility, as trading and fee-based activities are inherently riskier than traditional credit origination. DeYoung and Torna (2013) further showed that heavy reliance on investment banking and complex securitization revenues increased the likelihood of institutional failure during the 2008 crisis. These studies validate the use of NII as the empirical proxy for bank risk-taking in this paper.

3 Data Description

The primary data source used in this analysis comprises the 10-K reports submitted by 32 major financial institutions based in the United States. The temporal scope spans 32 years from 1993 to 2025. According to Securities and Exchange Commission regulations, the 10-K report provides a standardized overview of institutional financial condition, operational activities, and market risks. The textual analysis specifically targets sections containing management perspectives, namely Item 1A, Risk Factors, and Item 7, Management Discussion and Analysis. Unlike the strictly audited numerical tables found elsewhere in the 10-K, these narrative sections afford executives discretionary leeway to explain market conditions, rationalize past performance, and outline future projections. This qualitative flexibility makes them the most accurate corpus within the filing for extracting authentic corporate sentiment and strategic intent.

Data collection utilized automated web crawling via Python scripts, yielding a comprehensive longitudinal panel of 1,056 distinct 10-K documents. Prior to linguistic analysis, the documents underwent standardized preprocessing to isolate the core managerial narrative. First, the algorithm removed common English stopwords, including frequent conjunctions and articles that provide no document differentiation. Second, the algorithm removed domain-based stopwords. These are specialized, non-informative terms occurring frequently within financial datasets but carrying negligible analytical value. For the current dataset, these removed terms include fiscal, annual, report, SEC, shareholders, statements, notes, financial, item, part, ended, bank, banking, institution, company, assets, liabilities, loans, deposits, borrowings, interest, income, expense, credit, capital, and balance.

To explicitly structure the comparative analysis, the 32 institutions are segmented into two distinct classifications defined by their operational and regulatory environments: Traditional Banks and Shadow Banks. Traditional Banks are institutions that accept customer deposits, engage heavily in standard credit lending, and operate under strict regulatory frameworks overseen by the Federal Reserve and Federal Deposit Insurance Corporation. Shadow Banks are alternative financial institutions that provide credit and capital services without accepting retail deposits, thereby operating under minimal regulatory constraints. The left column of Table 3 in Appendix C, lists 15 Traditional Banks which include JPMorgan Chase, Bank of America, Citigroup, Wells Fargo, U.S. Bancorp, PNC Financial, Truist, Capital One, KeyCorp, Fifth Third, Regions Financial, Huntington, M&T Bank, Citizens Financial, and WesBanco. The right column lists the 17 Shadow Banks, which include Goldman Sachs, Morgan Stanley, Charles Schwab, American Express, State Street, Northern Trust, BlackRock, KKR & Co, Blackstone, Apollo Global, Carlyle Group, MetLife, AIG, Prudential, Travelers, S&P Global, and Moody's.

In addition to the textual corpus, the research collected quantitative financial data to serve as an objective empirical baseline. The primary goal of this financial analysis is to validate whether the self-reported corporate rhetoric aligns with actual operational behavior. The selected metric for this empirical validation is Non-Interest Income. To construct this dataset, an automated retrieval process extracted historical Non-Interest Income figures directly from the standardized financial tables within the Item 8, Financial Statements, section of the 10-K filings for all 32 institutions spanning 1993 to 2025. The research then aggregated these extracted data, calculating the total combined sum of Non-Interest Income across all institutions for each individual year.

4 Methodology

The methodology combines quantitative financial modeling and qualitative natural language processing across four distinct analytical approaches to identify divergences between corporate rhetoric and market behaviors.

The first method employs an aggregate analysis of simple word frequencies. This technique tracks the dynamic evolution of the top 50 most frequently used words across the entire corpus of 10-K filings. Monitoring these raw frequency counts establishes a baseline understanding of how broad thematic vocabularies within the banking sector shifted between the pre-crisis and the post-crisis eras.

The second method employs a lexicon-based analysis to quantify corporate communication strategies by calculating term frequencies across two custom dictionaries. These lexicons were constructed by expanding the baseline 2011 Loughran and McDonald financial dictionary with high-frequency terms derived empirically from the aggregate word frequency analysis of the 10-K corpus. This augmentation integrated two specific structural categories. The first category includes operational terms such as *securities*, *trading*, *capital*, and *derivatives*, which are indicative of the originate-to-distribute (OTD) banking model, a framework wherein institutions originate loans and sell them as packaged securities to investors. The second category incorporates post-2008 regulatory terminology reflecting the updated legal architecture.

The short-term profit dictionary captures terminology associated with immediate financial returns and active market operations. Conversely, the long-term stability dictionary comprises terminology related to structural safety, regulatory compliance, capital adequacy, and operational resilience, incorporating terms such as *risk*, *stress*, and *losses*. Within formal Securities and Exchange Commission disclosures, an elevated frequency of these terms does not necessarily denote increased operational risk-taking; rather, it indicates compliance-driven reporting mandated by post-crisis legislation requiring detailed disclosures of stress-testing and loss-mitigation strategies.

The third method applies advanced sentiment analysis using the FinBert model. FinBert is a neural network large language model (LLM) trained on financial texts so as to understand the precise context of industry vocabulary. Standard sentiment algorithms routinely misinterpret financial terminology. For example, a standard algorithm may flag the word *liability* as a negative emotion rather than a neutral accounting metric. Instead, the specialized training of the FinBert model prevents these contextual errors. The sentiment analysis here specifically targets the Risk Factors section and the Management Discussion and Analysis section of the 10-K filings. These targeted sections contain the

official future projections and market-condition assessments made by company executives.

The fourth method provides an empirical validation of the linguistic results through quantitative financial analysis. The research tracks the variable of *non-interest income* as a proxy for high-risk-taking activities. Non-interest income consists of revenue generated from volatile operations, like market trading, packaging loans into securities and charging investment banking fees. This differs fundamentally from net interest income, which is the Traditional highly regulated revenue derived from the margin between interest paid on deposits and interest earned on loans. Comparing the linguistic results with the quantitative data on non-interest income identifies the congruence or divergence between self-reported risk appetite and actual operational risk-taking.

5 Linguistic Analysis Results

This section presents the results of the first two Natural Language Processing methodologies. The first methodology is an aggregate word frequency analysis, and the second is a targeted lexicon-based analysis.

The aggregate frequency approach identifies broad shifts in industry vocabulary. The data for this analysis are derived by extracting the top 50 most frequently used words across the entire corpus of 10-K filings. Figure 1 in Appendix D maps how the frequency of these specific words shifted across the pre-crisis and post-crisis periods. The vertical axis lists the words, and the horizontal axis displays their normalized frequency percentage within the documents. Data points derived from pre-2008 filings appear as blue dots, while post-crisis data points appear as orange dots.[1] For example, the term *liabilities* appeared with a 20 percent frequency before 2008, as indicated by the blue dot, while its usage increased to 50 percent after the crisis, as indicated by the orange dot. As an additional data point, the term *disclosure* is positioned near the bottom of the vertical axis; its pre-2008 usage, as indicated by the blue dot, registers at roughly 10 percent, while its post-crisis usage, as indicated by the orange dot, shifts to the far right, approaching 40 percent.

Analysis of the pre-2008 data, indicated by the blue dots, demonstrates that the industry vocabulary centered on expansion and market activity. Transactional terms such as *exchange*, along with language focused on rapid growth and short-term profit generation, such as *leverage*,*securities*, *trading*, *capital*, and *yield*, constituted the majority of corporate filings. This reflects the prevailing originate-to-distribute (OTD) banking model of

[1]The grey line connecting points is a visual aid used to link the two time periods for each specific word, and it does not represent a statistical confidence interval.

the early 2000s. The OTD model is a financial strategy where a bank or lender originates loans, such as mortgages or auto loans, and immediately sells them to third-party investors rather than holding them to maturity. These loans are typically packaged together and converted into tradable securities, a process known as securitization, which are then sold to asset managers, pension funds, or other institutional investors. In this context, pre-2008 financial reports functioned as prospectuses, highlighting immediate financial opportunities to investors.

In contrast, the post-2008 data, represented by the orange dots, indicate a shift toward risk management, prudence, and regulatory compliance. A prime example is the word *disclosure*, which started appearing so frequently as to form a central thematic construct. The data also identify a renewed emphasis on Traditional banking terms such as *balance* and *credit*, alongside attention to regulatory and hedging terms, such as *tax* and *insurance*, respectively.

The conclusion drawn from this aggregate analysis is that a structural shift in vocabulary occurred. After the 2008 collapse, institutions reduced expansionist rhetoric in favor of narratives emphasizing long-term stability and compliance. This linguistic transition from transactional to compliance-focused language potentially provides a quantitative measure of the regulatory pressures introduced by the Dodd-Frank Act of 2010.

The targeted lexicon approach evaluates rhetoric across distinct institution types to examine specific communication strategies. To execute this methodology, two distinct dictionaries are defined. First, the Short-Term Profit Lexicon represents aggressive market expansion, utilizing words such as *yield*, *margin*, and *gain*. Second, the Long-Term Stability Lexicon represents regulatory prudence, utilizing words such as *hedging* and *prudent*. The complete list of words comprising both lexicons is available in Table 1 and Table 2 of Appendix B. Table 1 structurally presents the 30 distinct terms of the Short-Term Profit Lexicon, arranged in a four-column grid. Table 2 follows the identical structural format, cataloging the 30 distinct terms of the Long-Term Stability Lexicon. To quantify corporate strategy, this methodology measures the relative frequency of targeted lexicons rather than their absolute frequency against total document word count. Normalizing lexicon occurrences against total document length introduces noise, as 10-K filings exhibit high variance in length due to standard accounting boilerplate and shifting disclosure mandates. Instead, this analysis isolates the specific semantic domain of operational strategy. For each document, the raw frequencies of all terms within the Short-Term Profit Lexicon and the Long-Term Stability Lexicon are aggregated. The final metric is expressed as the proportional share of each lexicon relative to their combined total. This normalization isolates the rhetorical trade-off between competing corporate priorities. Consequently, the metric evaluates the structural composition of the institution's strategic narrative,

effectively neutralizing variations in aggregate document length.

Applying this targeted lexicon methodology to the Traditional Banks identifies a linguistic shift before and after 2008. Before the crisis, the reports of Traditional Banks used Short-Term Profit terms, such as yield, margin, and gain. Afterwards, the focus shifted toward Long-Term Stability, with words like risk, compliance, and liquidity. In quantifiable terms calculated via the methodology outlined above, where percentages represent the proportion of each lexicon against their combined total, the relative use of Short-Term Profit words by Traditional Banks decreased from 64.0 percent pre-crisis to 59.1 percent post-crisis. This represents a 4.9 percentage point rhetorical pivot away from profit terminology, as detailed in Figure 2 (a) of Appendix D. By contrast, their relative use of Long-Term Stability words increased from 36.0 percent to 40.9 percent. This demonstrates a compliance hierarchy where heavily regulated depository institutions altered their public messaging to project safety.

The methodology is next applied to Shadow Banks, which operate under lighter regulations. By measuring these relative proportions, the data show Shadow Banks reduced their Short-Term Profit language from a pre-crisis baseline of 66.8 percent to a post-crisis level of 61.7 percent. This represents a nearly identical 5.1 percentage point rhetorical pivot away from profit terminology. Concurrently, their relative Long-Term Stability terminology increased from 33.2 percent to 38.3 percent, as detailed in Figure 2 (b) of Appendix D. The post-crisis Short-Term Profit rhetoric of Shadow Banks of 61.7 percent is comparable to the pre-crisis rhetoric of Traditional Banks at 64.0 percent, reflecting the fact that Shadow Banks maintained their profit-seeking rhetoric to a higher degree, likely due to lower regulation.

The conclusion drawn from this targeted lexicon analysis indicates that systemic risk language did not dissipate from the financial sector; rather, it persisted. To contextualize these findings, it is necessary to observe the macroeconomic factors driving them. Before the crisis, Short-Term Profit language mirrored the originate-to-distribute banking model. Traditional Banks operated as securitization intermediaries, generating fee revenue by originating loans, packaging them into complex securities, and distributing them to remove localized asset risk. As the lexicon data demonstrate, post-crisis regulations constrained Traditional Banks from maintaining this transactional model, shifting their language toward stability. However, the sustained Short-Term Profit rhetoric within the Shadow Bank dataset confirms that these high-risk operations continued to thrive in the lightly regulated sector, allowing Shadow entities to openly maintain the operational framework that Traditional institutions rhetorically abandoned.

6 Sentiment Analysis Results

This section presents the results of the third Natural Language Processing methodology. The targeted lexicon analysis in Section 5 quantified the exact terminology utilized by financial institutions. Building upon those results, Method 3 utilizes the FinBert model to measure the emotional sentiment embedded within the corporate discourse.

To execute this methodology, the FinBert algorithm was applied directly to the pre-processed corpus extracted from Item 1A and Item 7 of the 10-K filings. The algorithm evaluates the text and assigns a probability score, classifying the phrasing as positive, negative, or neutral based on contextual market-impact parameters rather than rigid dictionaries. Translating FinBert's training on expert-annotated financial data to this specific analytical context, positive sentiment is categorized as discussions regarding growth prospects, financial strength, and effective risk management. Conversely, negative sentiment is captured through discussions of unmitigated risk, financial losses, and adverse market conditions. Finally, neutral sentiment captures the standardized compliance terminology and routine accounting metrics common in regulatory disclosures. Methodologically, FinBERT categorizes projections of operational expansion and yield generation as positive, whereas assessments detailing liquidity threats resulting from market volatility register as negative. Standard regulatory clauses affirming institutional adherence to federal capital adequacy requirements are strictly scored as neutral. The individual scores for each bank were then averaged to generate annual sentiment scores for the financial industry.

Figure 3 in Appendix D plots the time series for positive and negative sentiment across the industry. The horizontal axis presents the years from 2000 to 2025, and the vertical axis shows the average percentage probability score of each sentiment category. The green line tracks positive sentiment, and the red line tracks negative sentiment. An analysis of the data reveals that during the pre-2008 period, positive sentiment peaked at approximately 11 percent while negative sentiment remained subdued at roughly 3 percent. The onset of the 2008 crisis triggered a sharp inversion where negative sentiment surged to approximately 9 percent and positive sentiment concurrently declined to 4 percent. Following this period of volatility, positive sentiment recovered between 2014 and 2020 to consistently exceed negative sentiment and eventually plateaued at a maximum of roughly 6 percent.

In the post-2020 period, positive sentiment declines and reaches the level of negative sentiment. The two lines then stay at similar percentages, but in more recent times a shift starts to appear away from negative and into more positive sentiment.

Overall, the linguistic analysis indicates shifts consistent with phases of the business cycle. In the pre-2008 period, banks expressed positive sentiment associated with the pursuit of short-term profitability. During the crisis and recession, the banks expressed negative sentiment, likely due to declining profit prospects, along with a switch to language reflecting the regulatory requirements of the Dodd-Frank and Basel III legislations. In the post-2013 period, the temporary bump in positive sentiment is thwarted by the pandemic, while the banks continue to talk about the new regulatory requirements, as opposed to risk-taking.

7 Quantitative Analysis Results

This section presents the results of the fourth analytical methodology, which utilizes quantitative data from the annual financial reports.

The designated metric for this analysis is Non-Interest Income (NII), defined as revenue generated from fee-based market activities rather than traditional interest-based credit origination. The structural composition of this revenue varies significantly across institutional classifications. For traditional depository institutions, NII typically derives from retail-level transactions, including mortgage servicing rights, overdraft fees, and wire transfers. Conversely, for Shadow Banks and systemically important financial institutions in the post-crisis era, NII is predominantly generated through institutional-level market activities, including private credit origination, wealth management, proprietary trading, and leveraged loan syndication. Because these latter activities exhibit structurally higher volatility, Non-Interest Income serves as a proxy for measuring operational risk. The data were automatically extracted from the Item 8 or Financial Statements section of the 10-K filings for the 32 financial institutions in the sample, spanning the years 1993 to 2025. Tracking absolute aggregate nominal dollar volume rather than asset-scaled ratios is a deliberate methodological choice; it captures the true systemic scale and absolute growth of risk-associated revenues within the financial architecture, which proportional metrics can obscure during phases of aggressive balance sheet expansion.

Figure 4 in Appendix D plots the longitudinal Non-Interest Income for specific Traditional and Shadow Banks from 1993 to 2025. These individual trajectories are benchmarked against the average of all these banks, depicted by the solid black line. The individual trajectories are highlighted using distinct colors, with JPMorgan Chase depicted in light blue, Bank of America in orange, Citigroup in green, Wells Fargo in red, Goldman Sachs in light brown, and Morgan Stanley in purple. The data demonstrate that this average begins near 5 billion dollars in the early 1990s, expands to nearly 40 billion dollars pre-crisis, contracts during the recession, and subsequently resumes an upward trajectory to

surpass 50 billion dollars by 2025.

While the average provides a baseline, evaluating the individual trajectories reveals much sharper operational extremes. For example, JPMorgan Chase expanded its Non-Interest Income from approximately 35 billion dollars prior to the 2008 crisis to nearly 90 billion dollars by 2025. Similarly, Morgan Stanley doubled its volatile revenue generation from roughly 30 billion dollars pre-crisis to 60 billion dollars by the end of the sample period. Conversely, Citigroup demonstrates a starkly different trajectory, peaking at over 40 billion dollars pre-crisis before contracting to approximately 25 billion dollars by 2025.

However, relying solely on averages or individual case studies obscures the aggregate magnitude of systemic risk across the broader sector. To reveal the true scale of this expansion, Figure 5 aggregates the absolute total Non-Interest Income for the two distinct sectors.

Figure 5 in Appendix D plots the aggregate total NII for the 15 Traditional Banks, as depicted by the blue line, versus the 17 Shadow Banks, as depicted by the red line. The horizontal axis spans from 1993 to 2025, and the vertical axis measures total NII in billions of dollars. Both lines exhibit an overall upward trajectory with distinct cyclical declines. For example, during the 2008–2009 period, the Traditional Banks NII shows a decline from approximately 150 billion to 100 billion dollars. The Shadow Banks' NII begins at a higher baseline pre-2008, experiences a dip in 2009, remains slightly below the Traditional Banks' line between 2010 and 2020, and ultimately crosses back above the Traditional Banks' line in the post-2020 period. This aggregate view confirms that total risk-taking is ballooning across the entire industry.

8 Comparison

A comparison across the linguistic and quantitative methods yields the following main results.

First, in the pre-2008 period, bank sentiment was positive, bank language was related to short-term profitability, and bank activities were increasingly NII-related. Second, during the financial crisis and subsequent recession, bank language shifts to prudential and long-term goals, bank sentiment switches to negative, and NII bank activities decline. In the post-2013 period, bank sentiment is overall slightly positive, bank language remains focused on oversight terminology, but NII continues to increase. In fact, it increases at rates higher than those of the pre-2008 period.

Third, the findings reveal a structural divergence between the pre-2008 and post-2013

expansionary periods. In the pre-2008 period, the growth in NII associated with positive sentiment and profit-oriented terminology, accurately reflecting the financial gains generated by underlying securitization activities. Conversely, in the post-2013 period, aggregate Non-Interest Income expanded at rates exceeding pre-crisis levels, yet corporate language shifted to emphasize regulatory compliance. Institutional sentiment aligned with this regulatory framework rather than the corresponding increase in operational profitability. As a result, contemporary corporate discourse systematically masks the underlying reality of sustained operational risk-taking. This empirical decoupling between stated institutional stability and actual risk exposure is evident across both Traditional and Shadow Banks.

Finally, this decoupling of language-sentiment and quantitative data across the two periods highlights the need for increased pressure by central banks for quantitative data collection from the financial sector.

9 Conclusions

This thesis investigated the empirical validity of the rhetorical shift toward long-term stability within the United States banking sector following the prior financial crisis. By systematically juxtaposing Natural Language Processing analytics, specifically aggregate frequency tracking, custom lexicon analysis, and neural network sentiment modeling, against quantitative financial evaluation over a multi-decade scope, this research establishes a definitive statistical divergence between corporate communication and operational execution. The findings confirm that post-crisis regulatory frameworks did not eliminate systemic risk, but rather incentivized a structural compliance hierarchy that altered corporate messaging while high-risk financial activities continued to grow.

The linguistic and sentiment analyses demonstrate the construction of a standardized regulatory defense narrative among Traditional depository institutions. Following the implementation of post-crisis legislation, Traditional Banks significantly reduced their utilization of short-term profit terminology while concurrently increasing long-term stability rhetoric to emphasize compliance, disclosure, and balance sheet strength. Furthermore, aggregate industry sentiment stabilized from a pre-crisis peak into a static, neutral equilibrium, indicating that contemporary corporate filings function primarily as legal compliance instruments rather than authentic market assessments. Conversely, Shadow Banks sustained elevated short-term profit rhetoric in the post-crisis era, statistically mirroring the pre-crisis baseline of Traditional Banks.

This rhetorical division relates directly with an empirical decoupling of words from ac-

tions, quantified through the measurement of Non-Interest Income. A quantitative financial assessment provides a divergent perspective from the industry's stated stability objectives. The aggregate financial data demonstrate that the banking sector maintains a structural reliance on non-interest revenue streams, with total Non-Interest Income exceeding pre-crisis baselines. The analysis indicates that this operational expansion is a broad, as institutions across the sample have systematically increased their non-interest revenue. Furthermore, sector-level disaggregation shows that both Traditional and Shadow Banks exhibit parallel growth trajectories in these activities. Rather than reducing their exposure, Traditional Banks grew their fee-based revenues in tandem with the lightly regulated Shadow Banks. This trajectory culminated in an empirical parity, demonstrating that both sectors generate immense levels of total risk-associated, non-interest revenue despite their divergent public posturing.

Ultimately, this study concludes that contemporary corporate assertions of institutional stability actively obscure the underlying operational realities of the United States financial architecture. Post-crisis legislative frameworks failed to eliminate systemic economic threats; rather, they inadvertently facilitated a systemic decoupling, where vulnerabilities continue to grow across the financial system masked by compliance rhetoric. Because the banking sector continues to rely heavily on volatile, fee-driven operational strategies comparable to those that precipitated the prior economic collapse, the foundational threats to macroeconomic stability remain intact. The current regulatory paradigm has merely reclassified and obscured, rather than resolved, systemic risk

A Appendix A Data Compilation and Processing Outline

This study relies on a multifaceted dataset incorporating corporate self-reporting narratives and objective financial performance metrics. The data acquisition and preparation process involved four distinct phases:

- **Phase 1 Qualitative Data Acquisition from SEC Filings**
  - **Primary Source** United States Securities and Exchange Commission EDGAR database.
  - **Document Type** Audited Annual 10-K Filings.
  - **Temporal Scope** 1993 to 2025 capturing pre-crisis deregulation, the 2008 collapse, Dodd-Frank implementation, and the modern high-rate environment.
  - **Sample Size** 32 major United States financial institutions segmented into 15 Traditional Banks and 17 Shadow Banks.
  - **Targeted Disclosures** Programmatic extraction focused strictly on sections containing forward-looking statements and management psychology:
    - *Item 1A Risk Factors*
    - *Item 7 Management Discussion and Analysis of Financial Condition and Results of Operations*

- **Phase 2 Textual Preprocessing and Corpus Generation**
  - **Extraction** Automated programmatic downloading and HTML parsing utilizing Python.
  - **Cleaning pipeline** Strict removal of boilerplate legal disclaimers, standard SEC formatting, numeric tables, and raw HTML tags to isolate core narrative text.
  - **Linguistic Normalization** Implementation of tokenization, lemmatization, and domain-specific stopword filtering to eliminate non-value-added vocabulary.
  - **Data Storage** Processed text serialized via Python pickle module to ensure efficient high-speed retrieval during machine learning deployment.

- **Phase 3 Quantitative Financial Data Aggregation**

  - **Primary Empirical Metric** Total Non-Interest Income.
  - **Data Sourcing** Extracted from *Item 8 Financial Statements and Supplementary Data* within the 10-Ks and cross-referenced with macroeconomic aggregate data from standard financial databases to ensure accuracy.
  - **Analytical Function** Non-Interest Income served as the empirical proxy for risk, measuring reliance on volatile fee-based market activities versus Traditional credit lending.

- **Phase 4 Analytical Modeling**
  - **Lexicon Analysis** Algorithmic frequency mapping of custom Short-Term Profit dictionaries versus Long-Term Stability dictionaries.
  - **Sentiment Tracking** Deployment of the FinBert model. Because standard sentiment analyzers routinely misclassify financial jargon, the domain-specific pre-trained architecture of FinBert was essential for accurately scoring the positive, negative, and neutral valence of corporate rhetoric year over year.

B Appendix B Lexicon Dictionaries

Table 1: Short-Term Profit Lexicon

yield	margin	gain	fee
commission	spread	premium	trading
arbitrage	speculation	deal	transaction
proprietary	prop	leverage	quarterly
short-term	current	fiscal	interim
bonus	compensation	stock	options
shareholder	return	dividend	securities
equity	revenue	growth	acquisition
profit	income	sale	payment
amount	outstanding		

Table 2: Long-Term Stability Lexicon

risk	hedging	buffer	reserve
collateral	solvency	liquidity	stress
test	compliance	regulation	basel
dodd-frank	mandate	audit	oversight
governance	capital	sustainability	long-term
resilience	systemic	prudent	stability
retention	transparency	ethics	accountability
conduct	culture	allowance	losses
regulatory	liability	deposits	obligations
impairment	credit	loss	derivatives

C Appendix C Data Tables

Table 3: Bank Classifications and Focus Areas

Group 1 Traditional Banks, 15 institutions *Focus: Deposits, Lending, Mortgages*	**Group 2 Shadow Banks, 17 institutions** *Focus: Trading, Wealth Management, Private Equity*
JPMorgan Chase	Goldman Sachs
Bank of America	Morgan Stanley
Citigroup	Charles Schwab
Wells Fargo	American Express
U.S. Bancorp	State Street
PNC Financial	Northern Trust
Truist	BlackRock
Capital One	Blackstone
KeyCorp	KKR & Co
Fifth Third	Apollo Global
Regions Financial	Carlyle Group
Huntington	MetLife
M&T Bank	AIG
Citizens Financial	Prudential
WesBanco	Travelers
	S&P Global
	Moody's

Note: This table classifies the 32 institutions into two groups based on their primary business operations and regulatory status.

D Appendix D Data Figures

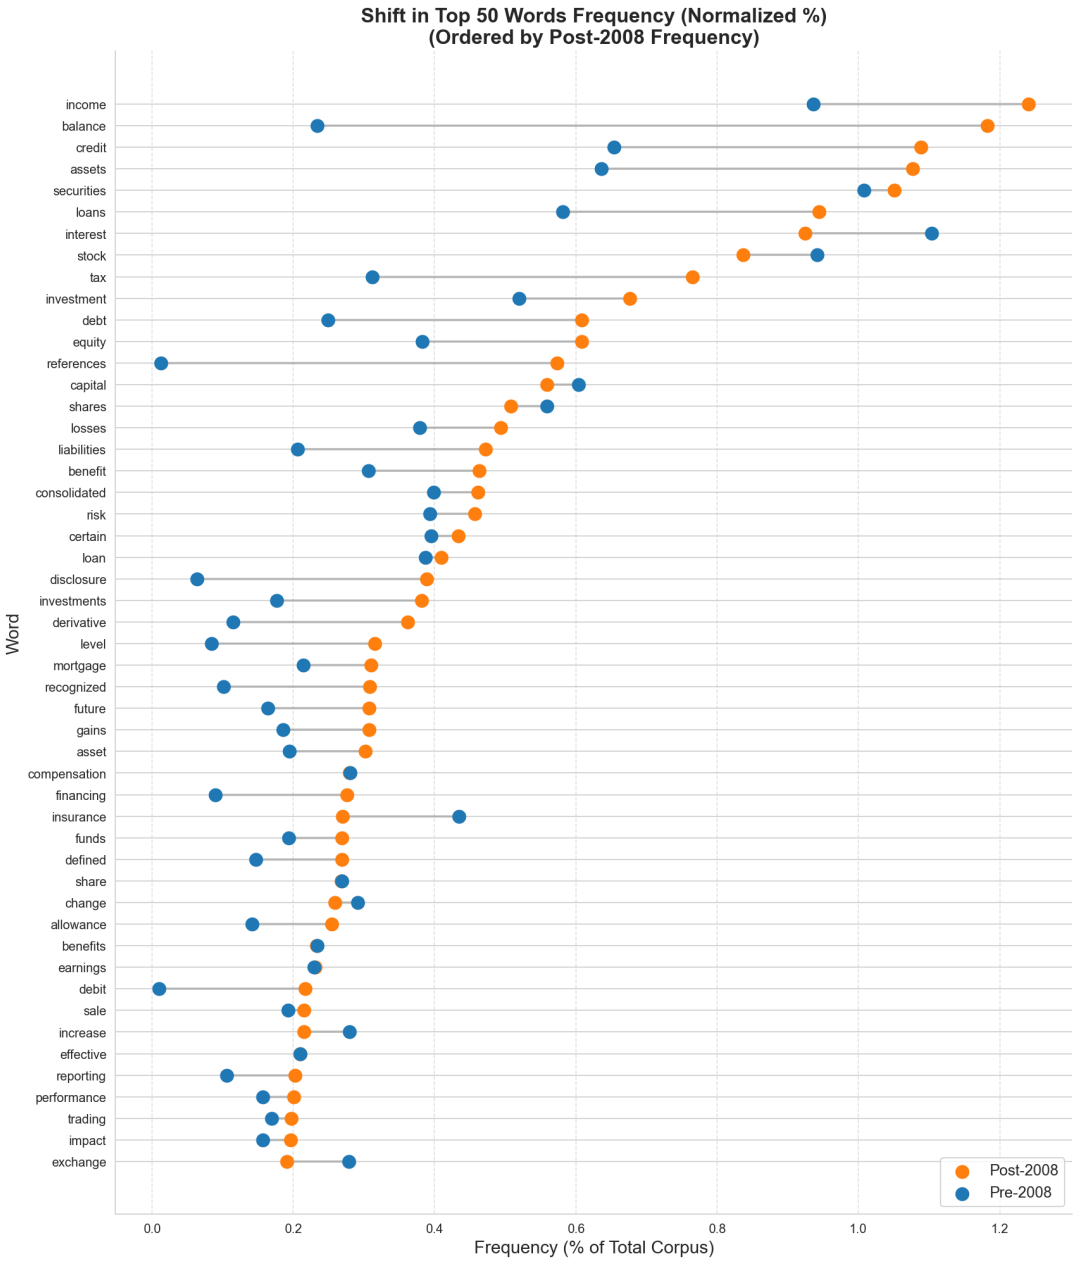
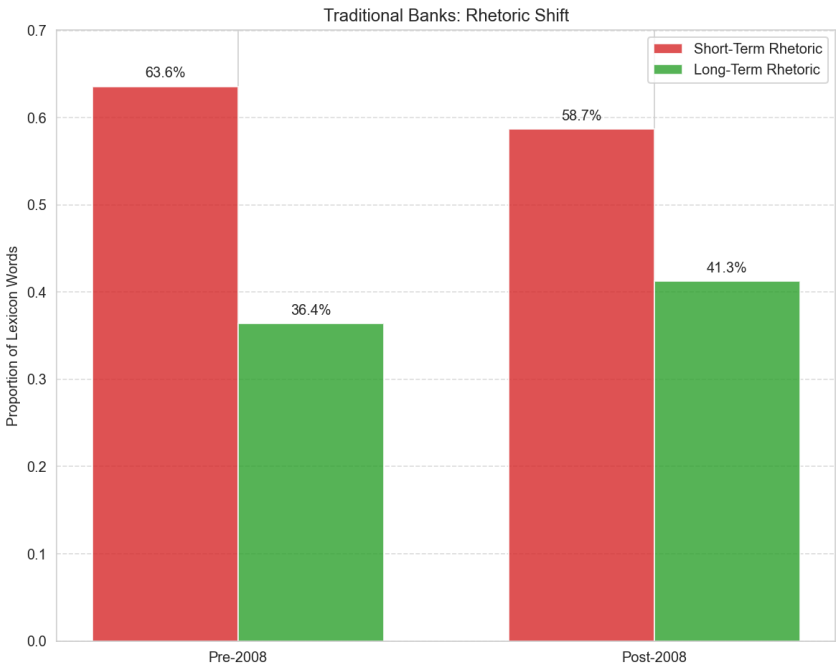


Figure 1: Shift in Top 50 Words Frequency Normalized Percentage

Note: This figure illustrates the proportional shift in the normalized frequency of the top 50 terms utilized in 10-K filings across the 32 sampled institutions. The vertical axis details the specific terminology, and the horizontal axis quantifies their normalized occurrence percentage. Observations from the pre-2008 period, indicated by blue markers, denote a higher frequency of expansionist and transactional terms such as yield and leverage. In contrast, the post-2008 period, indicated by orange markers, reveals a structural linguistic pivot toward compliance, long-term stability, and risk-management terminology. Notably, the term liabilities shifts from a 20 percent frequency pre-crisis to 50 percent post-crisis, while disclosure approaches a 40 percent occurrence rate, highlighting the sector's rhetorical anchoring around new regulatory frameworks.



(a) Traditional Banks Shift



(b) Shadow Banks Shift

Figure 2: Rhetoric Shift Traditional versus Shadow Banks

Note: This comparative figure tracks the rhetorical transition between the custom Short-Term Profit and Long-Term Stability lexicons before and after 2008. The reported percentages reflect the relative proportion of each lexicon against their combined total, rather than a fraction of aggregate document length. The first subfigure demonstrates that Traditional Banks altered their strategic discourse post-crisis, decreasing the relative concentration of short-term profit terminology from 64.0% to 59.1%. The second subfigure indicates that Shadow Banks executed a rhetorical pivot of identical magnitude. However, owing to a higher initial baseline, their post-crisis lexicon focus on short-term profit of 61.7% remains empirically comparable to the pre-crisis baseline of Traditional Banks at 64.0%.

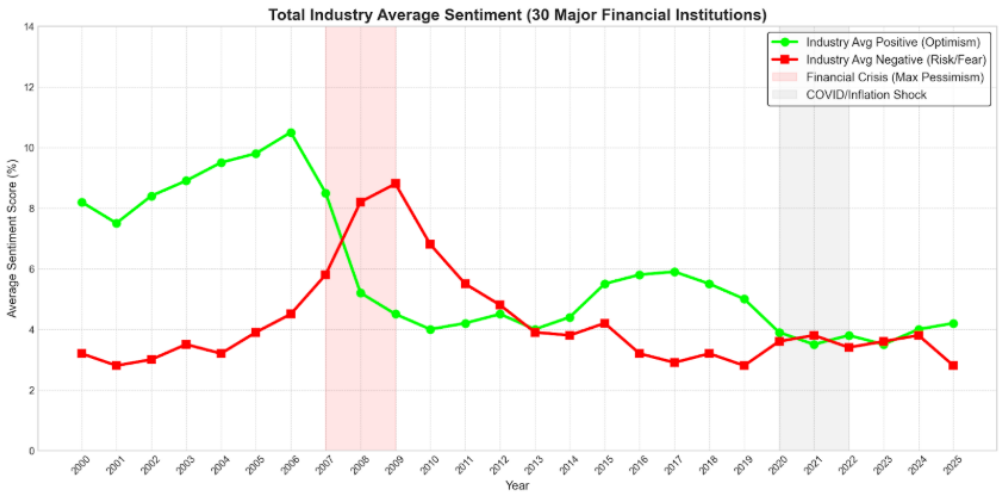


Figure 3: Total Industry Average Sentiment 2000 to 2025

Note: This longitudinal analysis tracks aggregate industry sentiment derived via the FinBert large language model from 2000 to 2025. The vertical axis represents the average probability score for positive, negative, and neutral sentiment classifications. The green line tracks positive sentiment while the red line tracks negative sentiment. The data capture a definitive inflection point during the 2008 financial crisis, where negative sentiment spikes to approximately 9% and sharply intersects the declining positive sentiment. Following a period of stabilized positive sentiment between 2014 and 2020, the post-2020 data indicate a temporary convergence of positive and negative outlooks due to pandemic pressures. This is followed by a gradual decoupling as positive sentiment recovers in the final years of the sample despite ongoing systemic risk-taking.

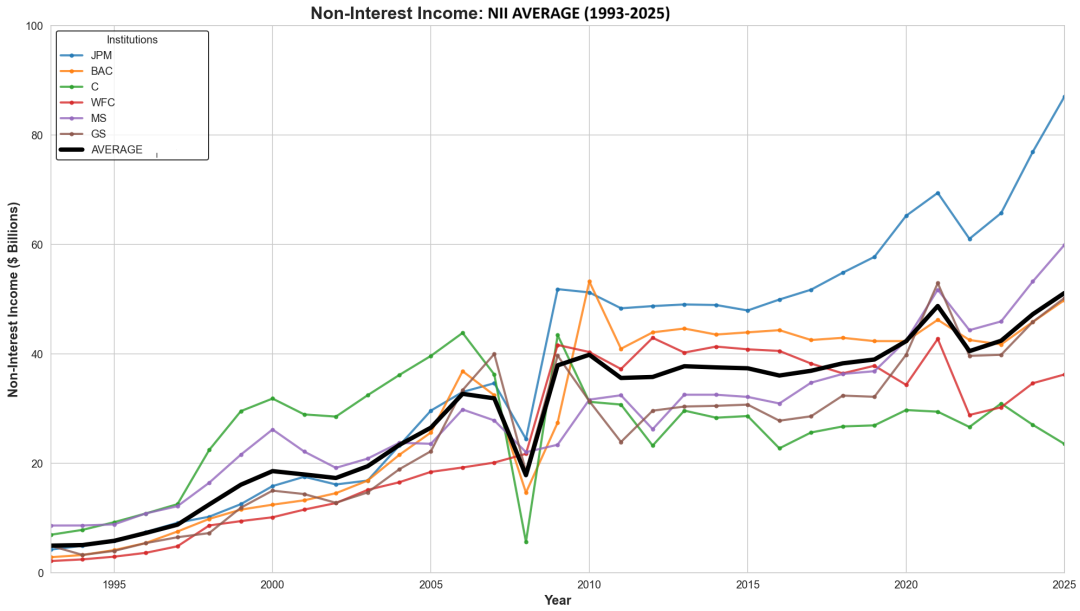


Figure 4: Non-Interest Income: Individual Institutions vs. Industry Average (1993 to 2025)

Note: This time-series tracks Non-Interest Income as an empirical proxy for fee-based operational activities. The data benchmark individual institutional performance against the industry average (solid black line). This average exhibits substantial pre-crisis expansion peaking near $40 billion, a pronounced contraction during the 2008 crisis, and a sustained post-crisis recovery exceeding $50 billion. Top-tier institutions, such as JPMorgan Chase, increased non-interest revenue to over $80 billion. This confirms a continued structural reliance on non-interest revenue streams, divergent from the stability-focused corporate rhetoric.

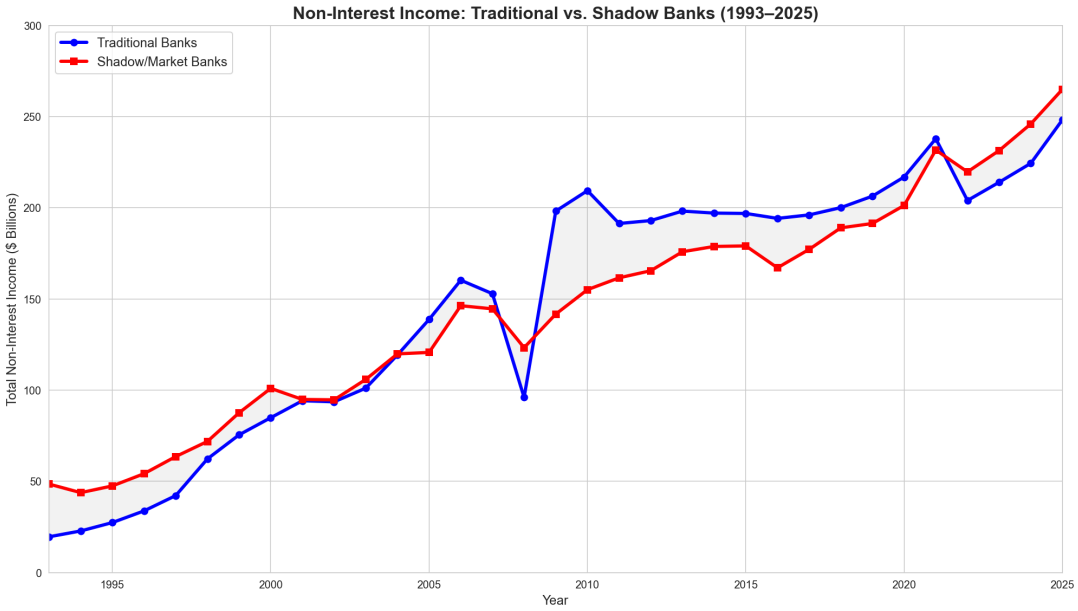


Figure 5: Non-Interest Income: Traditional versus Shadow Banks (1993 to 2025)

Note: This chart depicts the comparative historical trajectory of aggregate Non-Interest Income for 15 Traditional Banks, shown by the blue line, versus 17 Shadow Banks, shown by the red line, from 1993 to 2025. While Traditional Banks experienced significant income contractions during both the 2008 and 2020 economic downturns, Shadow Banks maintained higher pre-crisis baselines and exhibited robust post-crisis recoveries. Crucially, the post-2020 period displays an empirical parity. This visually confirms that both the Traditional and lightly regulated Shadow Banking sectors demonstrate parallel expansion, reaching comparable and unprecedented levels of total risk-associated, non-interest revenue generation.

References

[1] Araci, D. (2019). *FinBERT: Financial Sentiment Analysis with Pre-trained Language Models*. arXiv preprint arXiv:1908.10063.

[2] Cohen, L., Malloy, C., & Nguyen, Q. (2020). Lazy Prices. *The Journal of Finance*, 75(3), 1371–1415.

[3] Davis, S. J., Hansen, S., & Seminario-Amez, C. (2020). Firm-Level Risk Exposures and Stock Returns in the Wake of COVID-19. *National Bureau of Economic Research*, Working Paper 27867.

[4] DeYoung, R., & Torna, G. (2013). Nontraditional Banking Activities and Bank Failures During the Financial Crisis. *Journal of Financial Intermediation*, 22(3), 397–421.

[5] Ferrario, B., & Stantcheva, S. (2022). Eliciting People's First-Order Concerns: Text Analysis of Open-Ended Survey Questions. *AEA Papers and Proceedings*, 112, 163–169.

[6] Hansen, S., McMahon, M., & Prat, A. (2018). Transparency and Deliberation within the FOMC: A Computational Linguistics Approach. *The Quarterly Journal of Economics*, 133(2), 801–870.

[7] Huang, A., Wang, H., & Roache, M. (2020). *FinBERT: A Large Language Model for Extracting Information from Financial Text*. Working Paper.

[8] Jegadeesh, N., & Wu, D. (2015). Deciphering FOMC Minutes with a Dictionary-Based Machine Learning Approach. *Available at SSRN*.

[9] Loughran, T., & McDonald, B. (2011). When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *The Journal of Finance*, 66(1), 35–65.

[10] O'Neill, J., et al. (2017). *Topic Modeling for Legislative Texts*. Working Paper.

[11] Raviv, E., Hilscher, J., & Nabors, R. (2024). *Extracting Central Bank Sentiment Using Machine Learning*. Working Paper.

[12] Schmeling, M., & Wagner, C. (2019). Does Central Bank Tone Move Asset Prices? *Journal of Financial and Quantitative Analysis*, 54(2), 535–565.

[13] Stiroh, K. J. (2004). Diversification in Banking: Is Noninterest Income the Answer? *Journal of Money, Credit and Banking*, 36(5), 853–882.