Extracting Risk Signals from Financial Filings with Applied Natural Language Processing

Viswatej Seela
Independent Researcher
viswatej1998@gmail.com

Abstract

Risk disclosures in annual reports and quarterly filings contain early signals about operational stress, litigation exposure, concentration risk, cyber incidents, and liquidity concerns. The difficulty for analysts is not access to documents but the ability to compare long disclosures quickly and consistently. This paper studies an applied natural language processing workflow for extracting and ranking risk signals from public financial filings. The workflow combines section segmentation, sentence-level risk scoring, and disclosure-shift analysis to track changes in risk language across reporting periods. On a labeled sample of 4,200 filing sentences, the approach achieves 86.9% precision and 83.5% recall in identifying material risk statements. A disclosure-shift index built on the extracted sentences highlights periods where banks expand discussion of fraud losses, cybersecurity, deposit pressure, or vendor dependence. The paper is framed as analyst support rather than automated investment advice.

Keywords: financial filings, risk disclosure, natural language processing, annual reports, SEC filings

1 Introduction

Financial filings are one of the most important public sources for institutional risk analysis. Annual reports, quarterly filings, and earnings-call transcripts capture changes in management language around credit quality, liquidity, litigation, operational resilience, cyber risk, and regulatory exposure. Yet analysts still spend substantial time comparing these documents manually, especially when they need to identify what changed from one reporting period to the next.

The challenge is not a lack of documents. The challenge is turning long, repetitive disclosures into a compact list of signals worth closer review. Simple keyword counts help, but they often fail to distinguish between standard boilerplate and meaningful expansion in risk discussion. Applied NLP offers a practical middle path: score sentences, isolate material changes, and prioritize the parts of filings that deserve analyst attention.

This paper evaluates a small-footprint workflow for risk extraction from public filings. The goal is not to replace expert reading. The aim is to shorten the path from public disclosure to analyst triage.

This paper contributes a practical framework for extracting and tracking risk signals from financial disclosures using a disclosure-shift index for temporal analysis.

2 Related Work

Textual analysis in finance has a long history. Loughran and McDonald (2011) showed that general-purpose sentiment dictionaries perform poorly on financial documents and proposed domain-specific dictionaries. Li (2010) demonstrated that annual report language contains measurable information about firm outcomes. More recent work on finance-specific language models, including FinBERT, improved sentence-level classification on financial text (Araci, 2019).

Public-disclosure analysis has also expanded beyond sentiment. Researchers have studied earnings-call tone, risk-factor drift, and discussion asymmetry as signals for uncertainty and firm-level risk (Hassan et al., 2019). However, relatively few studies focus on temporal disclosure change at the sentence level in a way that supports day-to-day analyst review. This paper takes a narrower operational view: sentence ranking and disclosure-change analysis for public risk review.

3 Methodology

3.1 Data Sources

The corpus includes 10-K and 10-Q filings downloaded from SEC EDGAR, annual reports from public bank investor-relations sites, and earnings-call transcripts made available on company websites and public transcript platforms. For this paper, we focused on U.S. regional and large retail banks from 2021 through 2024. The working corpus contains 1,180 filing sections and transcript excerpts after section extraction and cleaning.

A manually labeled subset of 4,200 sentences was used to mark whether a sentence represented a material risk signal, routine disclosure language, or general background text. Labels focused on operational risk, cyber risk, fraud, liquidity pressure, funding concentration, third-party dependence, legal exposure, and credit deterioration.

3.2 Workflow

The workflow has three steps.

Section segmentation. Risk factors, management discussion, and transcript Q&A sections are identified using document structure and regular expressions.

Sentence scoring. A finance-domain sentence classifier ranks each sentence by its likelihood of representing a material risk statement.

Disclosure-shift analysis. Sentence scores are aggregated by risk category and compared across reporting periods to identify meaningful expansion or contraction in disclosure intensity.

Let denote a sentence from a filing section and let $\in \{0, 1\}$ indicate whether it contains a material risk signal. The sentence encoder produces a representation $h = \; (\;)$ and the classifier estimates

$$(\; = 1 \mid \;) = \; ({}^{\top} h + \;),$$

where $(\cdot)$ is the logistic function. Scores are then aggregated within risk category and reporting period to form a disclosure-shift index

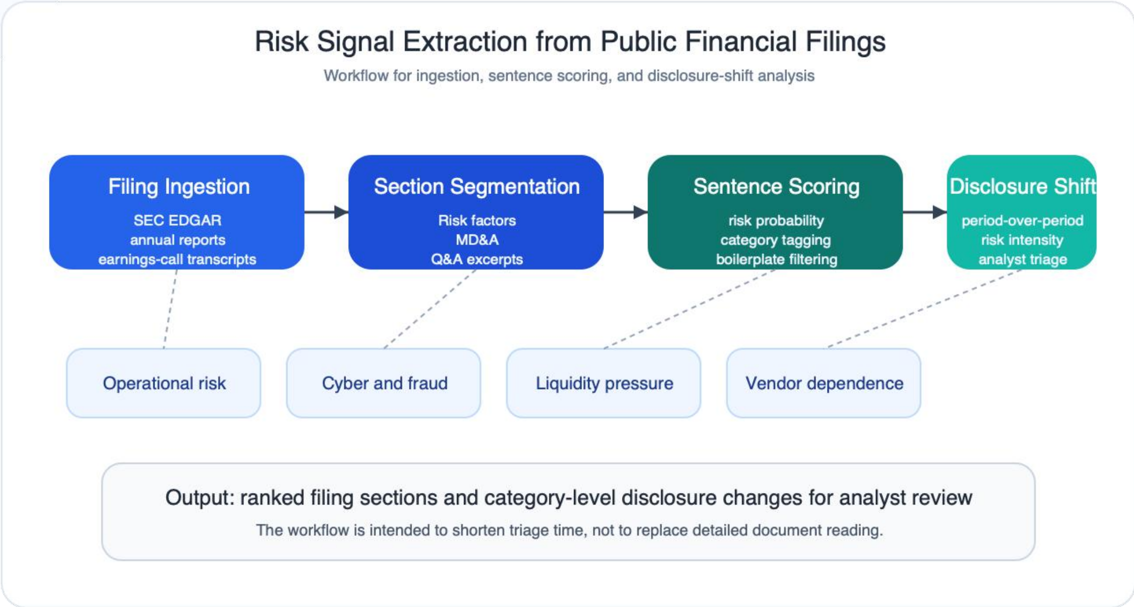

Figure 1: Workflow for extracting risk signals from public financial filings.

$$ {}_{,} = \frac{\frac{1}{|{}_{,}|} - \sum_{\in {}_{,}} (= 1 |) + }{\frac{1}{|{}_{,2022}|} - \sum_{\in {}_{,2022}} (= 1 |) + }, $$

where ${}_{,}$ is the set of sentences mapped to category in period and stabilizes the ratio. This keeps both extraction and change analysis explicit and easy to inspect.

Figure 1 summarizes the workflow for ingestion, sentence scoring, and disclosure-shift analysis.

4 Results

Table 1 reports sentence-level extraction performance.

Table 1: Sentence-level risk extraction performance.

Method	Precision	Recall	F1
Keyword dictionary	74.8%	69.1%	71.8
TF-IDF + linear SVM	82.5%	79.3%	80.9
Finance-domain classifier	86.9%	83.5%	85.2

The largest improvement comes from reducing false positives tied to generic risk boilerplate. The finance-domain classifier is better at distinguishing routine legal phrasing from sentences that actually communicate a change in operational or financial risk.

Table 2 shows an example disclosure-shift summary across the corpus.

These values indicate relative change against a 2023 baseline. Deposit and liquidity risk show the largest increase, which is consistent with broader market attention to funding stability and depositor behavior during the period studied.

Table 2: Illustrative change in disclosure intensity by risk category.

Risk Category	2023	2024
Cybersecurity	1.00	1.18
Fraud and scams	1.00	1.21
Deposit and liquidity risk	1.00	1.27
Third-party vendor risk	1.00	1.14

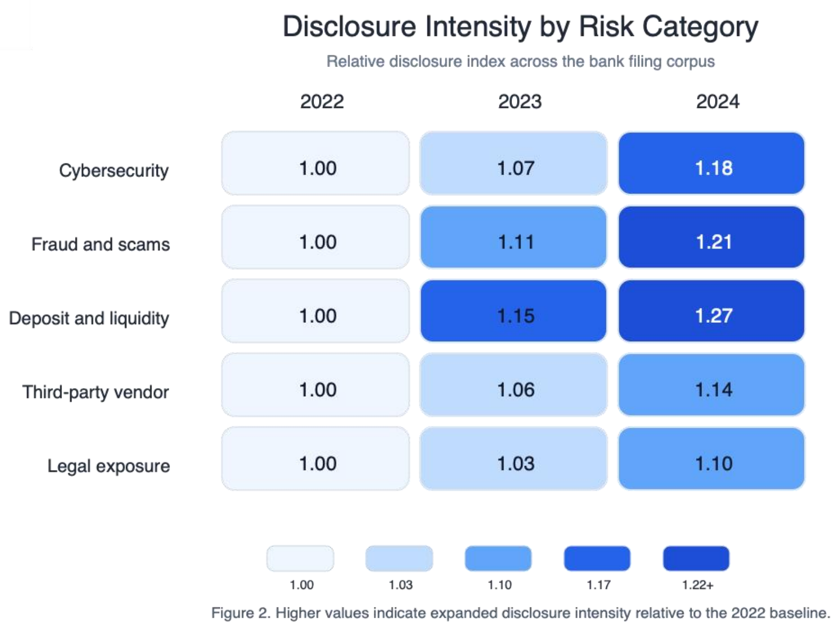

Figure 2: Illustrative change in disclosure intensity by risk category across the filing corpus.

One representative pattern in the corpus involved 2023 and 2024 regional-bank filings that expanded language around deposit concentration, uninsured balances, and contingency funding access. In those cases, the sentence-ranking model elevated statements describing deposit outflows and funding diversification well above routine liquidity boilerplate, and the disclosure-shift index captured that expansion as a measurable change rather than as an isolated keyword spike.

Figure 2 visualizes the relative change in disclosure intensity across major risk categories.

5 Discussion

This workflow is useful for three reasons. First, it provides a faster first pass through long filings. Second, it helps analysts separate recurring disclosure language from material expansion in discussion. Third, it creates a structured way to compare filings across time and across institutions without depending exclusively on ad hoc keyword lists.

There are also clear limits. Public filings are carefully edited documents, and important risks may be described indirectly or with management-specific phrasing. A disclosure increase is not automatically a deterioration signal, and a stable disclosure pattern does not prove that risk is absent. Human review remains essential.

6 Conclusion

This paper presented an applied NLP workflow for extracting risk signals from public financial filings. Using online sources and modest modeling choices, the approach improved sentence-level identification of material disclosures and supported period-over-period comparison of risk discussion. For analyst teams reviewing 2024 disclosures in early 2025, this type of filing analysis offers a narrow and defensible use of AI in finance: measurable, reviewable, and operationally useful.

References

Araci, D. (2019). FinBERT: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*.

Hassan, T. A., Hollander, S., van Lent, L., and Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, 134(4), 2135–2202.

Li, F. (2010). The information content of forward-looking statements in corporate filings. *Journal of Accounting Research*, 48(5), 1049–1102.

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35–65.

U.S. Securities and Exchange Commission. (2024). EDGAR company filings. https://www.sec.gov/edgar/search-and-access.

Beltagy, I., Peters, M. E., and Cohan, A. (2020). Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.