Thematic Analysis of Banking Regulatory Changes Following the Collapse of Silicon Valley Bank

Nicholas Wong
Massachusetts Institute of Technology
nicwjh@mit.edu

Priscilla Clark
Massachusetts Institute of Technology
poclark@mit.edu

Abstract

We propose an application of natural language processing (NLP) methods for thematic analysis of regulatory trends surrounding the collapse of Silicon Valley Bank (SVB). Our project investigates the hypothesis that regulatory focus shifted toward the governance of small banks and liquidity in the aftermath of the collapse. We leverage NLP techniques, including TF-IDF and BERTopic, to identify recurring patterns and themes in regulatory documents in the time periods 18 months pre- and post-SVB. By applying these methodologies, we uncover nuanced insights into how policy priorities evolved during this critical period. Our work aims to bridge the gap between understanding the root causes of the SVB collapse and evaluating the regulatory response, offering valuable implications for policymakers, regulators, and financial institutions navigating similar crises. Across our three deployed methods, we observe a confirmation of our initial hypothesis.

1 Introduction

The 2023 collapse of Silicon Valley Bank (SVB) marked a significant disruption in the financial landscape, exposing vulnerabilities in governance as well as highlighting liquidity and capital adequacy concerns. The event underscored the interplay between regulatory oversight and institutional decision-making. Our study aims to examine how regulatory focus shifted in response to the SVB collapse, providing insights to financial institutions and timing of new rules related to the root causes of the crisis.

We posit that regulatory focus on the governance of small and mid-sized banks, liquidity, and capital requirements intensified in the 18-month period post-SVB. To test this hypothesis, we employ natural language processing (NLP) techniques to perform a thematic analysis of regulatory documents published in the 18-month period pre- and post-SVB. We utilize naive keyword analysis, TF-IDF weighting, and BERTopic-based topic modeling to uncover shifts in regulatory focus and provide empirical evidence for a pivot in regulatory discourse. Experimental results showed significant statistical evidence confirming the hypothesis.

2 Related work

Numerous prior works have studied the root causes of the SVB collapse. A recurring theme across these papers is a link to rising interest rates which impacted the value of SVB's holdings of government bonds and mortgage-backed securities. In [3], Jiang states that interest rate hikes by the Federal Reserve led to a sharp drop in the price of securities on the asset side because of heightened interest rate sensitivity from their long duration. Simultaneously, SVB had a large proportion of short-term liabilities in the form of deposits that created a severe asset-liability mismatch that was exacerbated by these rate hikes [1, 7].

Most works also point to SVB's liquidity crisis as a root cause of the SVB collapse. In [1], Chen mentions the liquidity crisis as part of the broader collapse and focuses on how the Federal Reserve's interest rate hikes led to liquidity pressures which were exacerbated by sudden withdrawals by startups. Saif and Faour in [5] similarly mention the liquidity crisis as a result of SVB's inability to meet withdrawal demands from depositors, linking it to both the Federal Reserve's interest rate hikes and SVB's investment in long-term government bonds. Jiang in particular highlights the liquidity crisis as being the root cause of the SVB collapse, where the asset-liability mismatch and rising interest rates caused a liquidity crisis that ultimately led to SVB's collapse[3].

Another common feature amongst existing literature points to poor risk management on the part of SVB and a lack of regulatory oversight as significant causes of the collapse. SVB's investment decisions, many of which were highly interest rate sensitive, did not account for potential interest rate hikes and led to high exposure on the interest rate front [3]. Wang in [7] points to poor risk management as a primary cause of the SVB collapse, particularly as it relates to SVB's decision to invest heavily in long-term securities without sufficient diversification or hedging. Regarding regulatory oversight, Chen in [1] highlights regulatory easing during the Trump administration and delayed responses from regulators as significant contributing factors.

In this work, we contribute to the field of existing literature by focusing on the regulatory response to the SVB collapse. Utilizing select NLP techniques, we analyze shifts in regulatory focus, particularly in areas concerning liquidity and the regulatory oversight of small banks. This approach complements existing analyses of the collapse's cause and provides a novel perspective on the evolution of regulatory priorities in the wake of financial crises. At this project's conclusion, we aim to offer insights into the interplay between market vulnerabilities and regulatory responses.

3 Data

3.1 Source

The data source for the analysis included publicly available proposed and finalized rules extracted from Regulations.gov [4]. The following agencies were included: Consumer Financial Protection Bureau (CFPB), Federal Deposit Insurance Corporation (FDIC), Federal Reserve System (FRS), Financial Stability Oversight Council (FSOC), Department of the Treasury, Securities and Exchange Commission (SEC), and Office of the Comptroller of the Currency (OCC). The data source includes lengthy unstructured text in PDF/HTML files and structured metadata.

3.2 Extraction and Preprocessing

The Regulations.gov API is used to programmatically access metadata related to federal agency documents. This API provides essential metadata such as document titles, publication dates, docket numbers, and direct links to any attachment text. Using the metadata provided by the API, the pipeline identifies and downloads the relevant attachment HTML text.

API search parameters were used to filter for the required agencies and the time period in question: 18 months before and after March 10, 2023. A specific time window of 36 months, spanning 18 months before and 18 months after the collapse, is defined to analyze regulatory trends around the event. The dates marking this window are calculated as September 10, 2021, for the starting point and September 10, 2024, for the endpoint.

Once the time window is established, the dataset is further filtered to retain only those documents posted within this range. The preprocessing includes calculating the total number of documents within the specified period and grouping them by agency to determine which regulatory bodies were most active during this time. This approach ensures a focused dataset that highlights regulatory actions during a period of heightened financial scrutiny, setting the stage for further analysis to evaluate shifts in regulatory priorities related to systemic risks, liquidity, and governance in the aftermath of the SVB collapse.

3.3 Exploratory Data Analysis

Exploratory Data Analysis (EDA) aimed to examine shifts in the regulatory focus of financial agencies before and after the collapse of Silicon Valley Bank (SVB) in March 2023. The analysis primarily

leveraged document metadata and content to extract insights. Key steps included filtering relevant documents of the "Rule" or "Proposed Rule" type and organizing them into pre- and post-SVB datasets based on their posting dates. A total of 368 documents were analyzed, with 170 from the pre-SVB period and 198 from the post-SVB period.

The EDA uncovered several trends. First, a count of documents per agency revealed variations in the focus and contributions of different regulatory bodies. A pie chart visualization was used to highlight these differences as seen in Figure 4 in the appendix.

Additionally, temporal analysis identified shifts in the number of regulatory documents posted monthly, segmented by document type, with a distinct inflection point aligning with the SVB collapse. Document word counts were also analyzed, revealing their distribution and the prevalence of extremely long texts, necessitating filtering for outliers. The corresponding figures 5 and 6 are available in the appendix.

This EDA forms the foundation for further advanced modeling, such as topic analysis using TF-IDF and BERTopic, to uncover semantic shifts in regulatory priorities. The EDA findings corroborate initial hypotheses regarding an intensified regulatory focus on liquidity, systemic risk management, and small to mid-sized bank governance post-SVB.

4 Methods

We run 3 methods to assess regulatory changes pre- and post-SVB: a naive method utilizing simple keyword counts, a baseline TF-IDF analysis, and a cutting-edge BERTopic topic modeling approach. Results from these methods are then compared with our initial hypothesis.

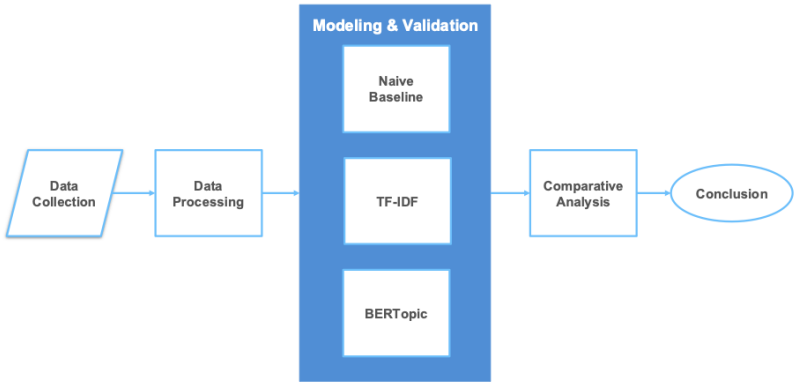

Figure 1: Proposed Project Architecture

4.1 Naive Baseline

The naive baseline model was developed to provide an initial, straightforward analysis of how the regulatory focus shifted before and after the Silicon Valley Bank (SVB) collapse. By analyzing the frequency of predefined keywords in regulatory documents, the model identifies themes that gained or lost attention following the collapse. It establishes a foundational comparison by quantifying the presence of specific terms, such as "liquidity," "capital," and "community bank," in documents from the pre- and post-collapse periods. The naive approach offers a simple and interpretable framework for understanding the regulatory response without relying on complex natural language processing techniques. Keywords for the naive model were determined based on team member subject-matter expertise and a review of event news articles:

Selected Keywords
Liquidity
Capital
Capital Requirement
Small Bank
Mid-Sized Bank
Community Bank
Interest Rate Risk
Concentration

Table 1: Keywords for Naive Analysis on SVB Collapse

4.2 TF-IDF

TF-IDF is a statistical measure used to evaluate the importance of a term within a document relative to the entire corpus of documents. It combines term frequency, which captures how often a term appears in a single document, and inverse document frequency, which downweighs terms that are common across the corpus. In doing so, TF-IDF highlights terms that are frequent and distinctive, providing a weighted representation of the document's content. We employ TF-IDF to identify the 20 most important keyword pre- and post-SVB to analyze thematic shifts.

4.3 BERTopic

BERTopic is a state-of-the-art topic modeling technique that leverages transformer-based embeddings and clustering algorithms to identify themes within text data. It is the most advanced technique we employ. We review three core aspects of BERTopic that are most relevant to this project (sentence embeddings, dimensionality reduction, topic representation). We refer the reader to [2] for a more comprehensive review.

The first step in BERTopic involves transforming text into dense vector representations using transformer-based models, such as Sentence-BERT (SBERT). Given a document $D = \{w_1, w_2, \ldots, w_n\}$, where w_i represents the words in the document, the embedding model generates a vector:

$$\mathbf{E}_D = f_{\text{transformer}}(D) \in \mathbb{R}^d$$

Here, $f_{\text{transformer}}$ is the embedding function, and d is the dimension of the embedding space. By using transformer-based sentence embeddings to capture semantic relationships in regulatory text, BERTopic is able to provide a nuanced understanding of the data.

To make high-dimensional embeddings computationally manageable, BERTopic applies Uniform Manifold Approximation and Projection (UMAP). UMAP reduces the embeddings while preserving their local structure. The similarity between points $\mathbf{z}_i$ and $\mathbf{z}_j$ in the low-dimensional space is defined as:

$$Q_{ij} = \frac{1}{1 + \alpha \|\mathbf{z}_i - \mathbf{z}_j\|^2}$$

where α is a hyperparameter, and $\|\mathbf{z}_i - \mathbf{z}_j\|$ is the Euclidean distance between the points in the reduced space. Its dimensionality reduction with UMAP ensures efficient handling of high-dimensional embeddings while preserving local structures, crucial for identifying distinct regulatory themes.

After clustering the reduced embeddings, BERTopic identifies the most relevant keywords for each topic using Term Frequency-Inverse Document Frequency (TF-IDF). The importance of a word w in a topic t is computed as:

$$\text{TF-IDF}(w, t) = \text{TF}(w, t) \cdot \log \frac{N}{\text{DF}(w)}$$

where $\text{TF}(w, t)$ is the frequency of word w in topic t, $\text{DF}(w)$ is the number of documents containing w, and N is the total number of documents. BERTopic's topic representation with TF-IDF extracts interpretable keywords, making it easier to uncover and compare trends across pre- and post-SVB periods.

Compared to traditional topic modeling techniques, BERTopic is able to leverage transformer-based embeddings to capture contextual and semantic relationships between words - crucial for analyzing

complex regulatory language. BERTopic's use of UMAP allows for greater flexibility in clustering and identifying nuanced topics, even in unbalanced datasets, which makes it particularly suited for analyzing pre- and post-SVB regulatory trends.

5 Experiments

5.1 Naive Baseline

The results of the naive baseline analysis highlight significant changes in regulatory focus following the collapse of Silicon Valley Bank (SVB). The frequency of documents mentioning predefined keywords in pre- and post-SVB periods provides insight into evolving regulatory priorities. Several keywords exhibited statistically significant increases, indicating heightened attention to specific themes in the aftermath of the crisis.

The process begins by scanning regulatory documents from pre- and post-SVB periods for occurrences of the selected keywords and their variations (e.g., singular, plural, or gerund forms). Only documents containing a keyword at least three times are counted, ensuring the focus remains on substantive mentions rather than incidental appearances. The keyword counts are then normalized by the total number of documents in each period, and a Poisson regression analysis is applied to assess whether the observed changes in keyword frequency are statistically significant. This approach identifies shifts in regulatory emphasis while accounting for differences in the size of the document corpus between the two periods.

Key terms like "liquidity," "capital," and "capital requirement" saw substantial increases in mentions post-SVB. For instance, "liquidity" appeared in 13 documents pre-SVB and 40 documents post-SVB, with a post-rate of 20.2% compared to 7.6% pre-SVB, yielding a significant rate ratio of 2.64 ($p = 0.001$). Similarly, "capital" increased from 31 to 60 documents, showing a significant rate ratio of 1.66 ($p = 0.019$). The largest relative increase was observed for "capital requirement," which rose from 1 pre-SVB document to 10 post-SVB documents, with a rate ratio of 8.59 ($p = 0.007$). These results reflect an intensified focus on liquidity management and capital adequacy, addressing core vulnerabilities exposed during the SVB collapse.

Emerging themes also gained prominence, with "community bank" mentions rising from 7 to 22 documents and "concentration" increasing from 6 to 20 documents. These shifts, with rate ratios of 2.70 and 2.86, respectively, underscore growing regulatory attention to systemic risks in smaller financial institutions and concentrated client bases, key factors in the SVB collapse. However, terms like "small bank" and "mid-sized bank" showed limited changes, with the former slightly declining and the latter remaining statistically insignificant. This suggests a broader regulatory pivot toward systemic risk management rather than institution-specific issues.

Overall, the results demonstrate how the SVB collapse reshaped regulatory discourse, emphasizing liquidity, capital resilience, and systemic vulnerabilities to prevent future crises. The naive baseline serves as a foundational analysis, capturing the key shifts in regulatory focus and paving the way for more nuanced exploration using advanced modeling techniques.

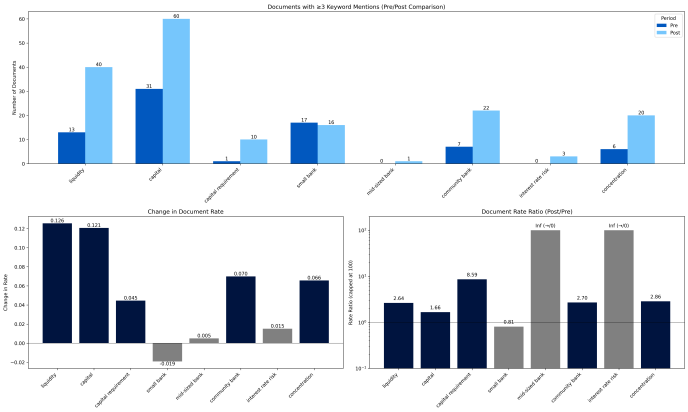

Figure 2: Naive Baseline Results

5.2 TF-IDF

We implement TF-IDF by vectorizing the text using *TfidfVectorizer*[6] to create a TF-IDF matrix encoding the weighted importance of terms. TF-IDF scores are averaged across pre- and post-SVB corpora.

Term	Change Type	Pre-SVB Score	Post-SVB Score
rov	New term	0.0000	0.0193
appraisal	Increased	0.0018	0.0145
avms	New term	0.0000	0.0129
egrpra	Increased	0.0002	0.0126
advisers	Increased	0.0004	0.0124
agenda	Increased	0.0002	0.0123
ncua	Increased	0.0029	0.0122
idi	Increased	0.0046	0.0117
fdta	New term	0.0000	0.0117
valuation	Increased	0.0002	0.0117
fincen	Increased	0.0026	0.0112
avm	New term	0.0000	0.0099
nprm	Increased	0.0005	0.0099
aml	New term	0.0000	0.0098
edgar	New term	0.0000	0.0095
quality control	Increased	0.0001	0.0088
real estate	Increased	0.0058	0.0087
issuers	Increased	0.0052	0.0084
trading	Increased	0.0006	0.0084
guidelines	Increased	0.0011	0.0081

Table 2: Top 20 Keywords Pre- and Post-SVB Comparisons

TF-IDF highlights significant changes in regulatory language pre- and post-SVB. Terms relating to valuation methodologies, regulatory frameworks, and institutional oversight gained prominence post-SVB. Notably, reconsideration of value (*rov*), automatic valuation models (*avms*), and financial data transparency act (*fdta*), which signal a pivot toward accurate valuation and transparency in financial assessments. Significant weight increases in terms such as *appraisal* and *valuation* further reflect the shift in regulatory focus toward assessing financial health.

5.3 BERTopic

We implement BERTopic for topic modeling the periods pre- and post-SVB utilizing standard and pruned approaches. Following these experiments, we validate the performance of BERTopic with topic coherence.

5.3.1 Standard Topic Modeling

To begin with, we implement a standard BERTopic model on the two distinct corpora of data pre- and post-SVB. In this experiment, we apply BERTopic without pruning constraints, allowing topics to emerge organically and for an unbiased assessment of semantic shifts.

Pre-SVB topics emphasize broader development and community (Topic 2: Pre-SVB), while post-SVB topics have a considerably narrowed focus on systemic risks and institutional mechanisms, such as FDIC resolution frameworks (Topic 2: Post-SVB). Post-SVB, we also notice an explicit focus on risk management (Topic 1: Post-SVB) and systematic regulatory oversight that is not as prevalent in the topic pre-SVB. Moreover, post-SVB, we see topics related to valuation guidance (Topic 5: Post-SVB) and civil penalties with inflation considerations (Topic 7: Post-SVB), suggesting a reactive response by regulators toward financial stability and penalties in the aftermath of SVB.

Topic	Top 5 Keywords
0	bureau, financial, consumer, sec, index
1	board, reserve, payment, bank, regulation
2	bank, agencies, community, assessment, development
3	cbp, subheading, broker, customs, brokers
4	assessment, fdic, deposit, rate, ratio
5	amount, threshold, account, cpi, exempt
6	hhs, enrollment, coverage, waiver, departments
7	banking, incident, organization, service, organizations
8	threshold, amount, january, cpi, increase

Table 3: Pre-SVB - Standard

Post-SVB, we also see health-policy (HHS) related topics (Topic 6: Pre-SVB) completely vanish, signaling a shift away from non-financial regulatory themes. We also see increased liquidity and credit concerns post-SVB through the emergence of topics related to debt eligibility (Topic 8: Post-SVB) and mortgage control (Topic 6: Post-SVB).

Topic	Top 5 Keywords
0	see, bureau, commission, sec, financial
1	risk, bank, agencies, banking, sec
2	fdic, resolution, cidi, deposit, deposits
3	cbp, merchandise, investigation, sec, date
4	regulations, occ, egrpra, comments, agencies
5	rov, guidance, appraisal, valuation, cfr
6	mortgage, avms, control, avm, quality
7	inflation, act, civil, adjustment, penalty
8	covered, ltd, ihc, eligible, debt

Table 4: Post-SVB - Standard

The intertopic distance maps illustrate the thematic structure of regulatory discourse pre- and post-SVB, generated using BERTopic.

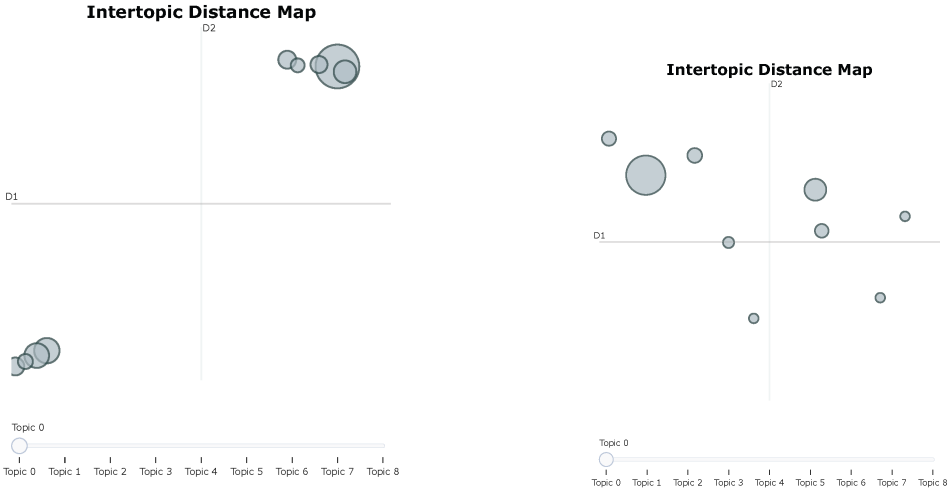

(a) Pre-SVB Intertopic Distance Map

(b) Post-SVB Intertopic Distance Map

Figure 3: Comparison of Intertopic Distance Maps Pre- and Post-SVB Collapse

Pre-SVB, the topics are tightly clustered with greater overlap and reflect a more cohesive regulatory focus. Post-SVB, the topics are more dispersed, indicating a shift toward diverse and specialized regulatory priorities (such as liquidity and credit concerns) in response to the crisis.

5.3.2 Pruned Topic Modeling

Standard cross-validation methods for hyperparameter tuning are not viable given the unsupervised nature of BERTopic. We implement a heuristic approach for BERTopic pruning where we change key parameters for the *HDBSCAN* core clustering mechanism of BERTopic. *min_cluster_size* is increased from 5 to 20 to reduce topic fragmentation and *cluster_selection_epsilon* is kept low at 0.01 to enforce stricter clustering boundaries. In the BERTopic model, *min_topic_size* is set to 20 in the pruned model to ensure that only topics with a sufficient amount of associated documents are retained, improving signal-to-noise.

Topic	Top 5 Keywords
0	bank, agencies, bureau, sec, assessment
1	board, reserve, bank, payment, regulation

Table 5: Pre-SVB - Pruned

Topic	Top 5 Keywords
0	see, sec, bureau, commission, financial
1	risk, bank, agencies, sec, banking

Table 6: Post-SVB - Pruned

We see an increased emphasis on risk post-SVB which was not as dominant in the results pre-SVB. This aligns with the results from the previous unpruned approach where we see the anticipated pivot toward systemic risk management post-SVB. Compared to the results from our unpruned approach, pruning consolidates peripheral topics, revealing dominant macro-level regulatory concerns pre- and post-SVB. Topic 0 post-SVB suggests increased emphasis on public accountability, centralized oversight, and reporting mechanisms through keywords such as *commission*, *sec*, and *financial*.

5.3.3 Validation

We utilize topic coherence as our evaluation metric for BERTopic to measure the semantic alignment and interpretability of identified topics. Nuanced themes have to be precisely captured for regulatory topics and coherence ensures that the identified topics are meaningful and interpretable. Metrics like perplexity or likelihood that prioritize mathematical fit over semantic quality are less aligned with our goal of understanding shifts in regulatory focus.

Scenario	Standard	Pruned
Pre-SVB	0.1454	0.0398
Post-SVB	0.1213	0.0107

Table 7: BERTopic Topic Coherence Scores

The unpruned model generates consistently better coherence scores pre- and post-SVB, suggesting that the standard approach generates more semantically meaningful topics. The pruned approach, better for identifying dominant trends and summarizing macro-level themes, sacrifices semantic alignment to emphasize thematic consolidation. Both approaches show higher coherence scores pre- as compared to post-SVB, suggesting greater semantic clarity in regulatory discourse before the SVB collapse. Overall, coherence scores across both methods and text corpora are relatively low, reflecting the limitations of BERTopic for topic modeling complex regulatory text without customization. Fine-tuning methods or alternative approaches should be utilized for more granular thematic analyses.

6 Conclusion

Results substantiate the hypothesis that regulatory focus on the governance of small and mid-sized banks, liquidity, and capital requirements intensified in the 18 months following the SVB collapse. The naive model revealed a post-SVB increase in documents mentioning "liquidity", "capital", "capital requirement", "community bank", and "concentration", indicating heightened regulatory

attention in these areas. TF-IDF analysis complemented these findings by identifying a surge in the importance of terms related to accurate valuation and assessment of financial institutions post-SVB. BERTopic further highlighted this shift, showing a transition from broad development themes pre-SVB to a narrowed emphasis on systemic risks, risk management, and liquidity concerns post-SVB. Collectively, these results demonstrate a cohesive regulatory pivot toward strengthening oversight and mitigating risks associated with small and mid-sized banks. Ensembling models of varying complexities in this fashion allows for improved robustness and confidence.

However, several important limitations and confounding factors should be considered when interpreting these results. First, the 18-month timeframe, while capturing immediate regulatory responses, may be insufficient to establish long-term regulatory trends. Regulatory changes often evolve over multiple years, and a longer study period would better capture the full scope of policy shifts. Additionally, the implementation timeline for regulatory changes varies significantly - while some changes can be implemented quickly through guidance and interpretive rules, others require lengthy notice-and-comment periods and gradual phase-ins. Our analysis captures regulatory intent through document analysis but may not fully reflect the timing and impact of actual policy implementation.

Several confounding variables could influence our findings. The post-SVB period coincided with broader economic factors including inflation concerns, cryptocurrency market volatility, and ongoing post-pandemic recovery efforts. These concurrent events may have independently influenced regulatory focus. Furthermore, our document-centric analysis approach may be affected by variations in regulatory document quality, consistency, and comprehensiveness across different agencies and time periods.

The study's domestic U.S. focus represents another limitation. International regulatory responses, particularly from jurisdictions with similar banking structures, could provide valuable comparative insights. For instance, European and Asian regulatory frameworks may have evolved differently in response to SVB's collapse, offering alternative approaches to similar challenges. Future research should consider cross-border regulatory dynamics and potential policy spillover effects.

Potential extensions include an exploration of broader and more nuanced trends in financial regulations. Comparative analyses with similar crises, such as the 2008 collapse of Lehman Brothers, could reveal consistent patterns in regulatory shifts. This could also be supplemented with a comparative analysis with a relatively uneventful time period that can serve as a null baseline model. Additionally, sentiment analysis of regulatory texts to contrast polarity of these sentiments may offer further insights into regulators' attitude toward topics such as risk, small banks, and liquidity. Future studies should also examine regulatory implementation timelines more explicitly, perhaps through case studies of specific policy changes and their execution. International comparative analyses could illuminate how different jurisdictions balance similar regulatory challenges. Cross-industry analysis may offer insights into the interplay between regulatory policy and market behavior, while deeper examination of document quality and consistency issues could help refine methodological approaches. Lastly, extending the analysis to compare regulatory shifts across industries may offer insights into the interplay between regulatory policy and market behavior, particularly in how different sectors respond to and implement regulatory changes.

References

[1] Haonan Chen. "An Study on the Causes and Consequences of the SVB Collapse". In: *SHS Web of Conferences* 188 (2024). Ed. by Azlina Md Yassin, p. 01018. ISSN: 2261-2424. DOI: 10.1051/shsconf/202418801018. URL: http://dx.doi.org/10.1051/shsconf/202418801018.

[2] Maarten Grootendorst. *BERTopic latest documentation — bertopic.readthedocs.io*. https://bertopic.readthedocs.io/. [Accessed 20-11-2024].

[3] Shengye Jiang. "Reasons and Enlightenment for The Bankruptcy of Silicon Valley Bank". In: *Frontiers in Business, Economics and Management* 13.2 (Feb. 2024), pp. 160–162. ISSN: 2766-824X. DOI: 10.54097/kvh6bm11. URL: http://dx.doi.org/10.54097/kvh6bm11.

[4] Regulations.gov. *Regulations.gov*. Accessed: 2024-10-09. n.d. URL: https://www.regulations.gov/.

[5] Abdullah Saif. S. S. Al-Sowaidi and Ahmad M. W. Faour. "Causes and Consequences of the Silicon Valley Bank Collapse: Examining the Interplay Between Management Missteps and

the Federal Reserve's Floundering Decisions". In: *Journal of World Economic Research* 12.1 (2023), pp. 38–46. DOI: 10.11648/j.jwer.20231201.15. eprint: https://article.sciencepublishinggroup.com/pdf/10.11648.j.jwer.20231201.15. URL: https://doi.org/10.11648/j.jwer.20231201.15.

[6] *TfidfVectorizer — scikit-learn.org*. https://scikit-learn.org/1.5/modules/generated/sklearn.feature_extraction.text.TfidfVectorizer.html. [Accessed 04-12-2024].

[7] Xuanqi Wang. "Silicon Valley Bank Collapse: Causes amp; Consequences". In: *Highlights in Business, Economics and Management* 32 (May 2024), pp. 53–57. ISSN: 2957-952X. DOI: 10.54097/s3scbn33. URL: http://dx.doi.org/10.54097/s3scbn33.

Appendix

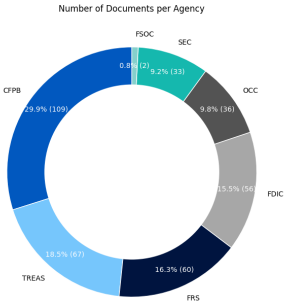

Figure 4: Included Regulatory Agencies

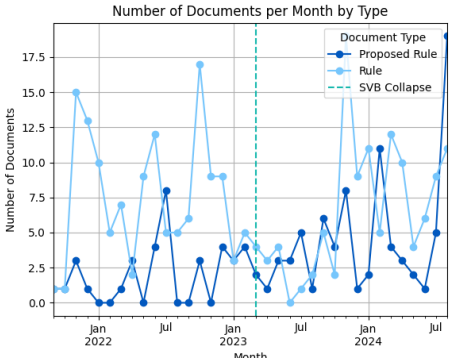

Figure 5: Number of Documents per Month

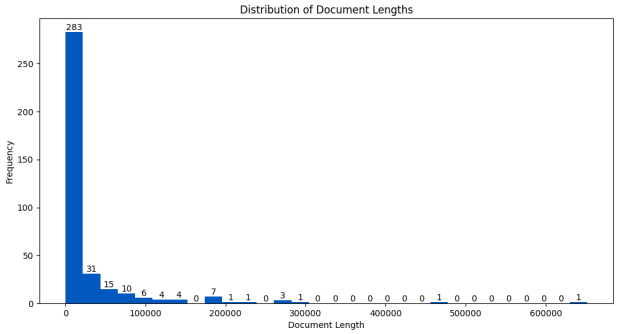

Figure 6: Distribution of Document Length