Online Appendix

A Data Sources

Figure A.1: **Data Construction**
This figure demonstrates how to link the FDIC data to the SEC EDGAR system. This task requires a bridge between the FDIC certificate number and the central index key (CIK), the key ID number used by the SEC to identify public companies. We first find the corresponding identification numbers used by regulators to identify commercial banks and then link the commercial banks to their parent holding companies. The CIK number is available in the CRSP-COMPUSTAT dataset. The Federal Reserve Bank (FRB) of New York provides the corresponding CRSP's permanent company identifier (PERMCO) for each bank holding company (BHC). Finally, merging with the CRSP-COMPUSTAT, the process identifies the corresponding CIK for each BHC in the sample, including the failed ones. From bottom to top, the "FDIC certif" denotes the certificate id used to identify banks at the FDIC level, which is item RSSD9050 in the Call Report. Parent ID corresponds to the higher holder of the commercial bank, which corresponds to item RSSD9364 in the Call Report. RSSD9001 is the unique ID given to the bank holding company in the FR Y-9C report (consolidated statement).

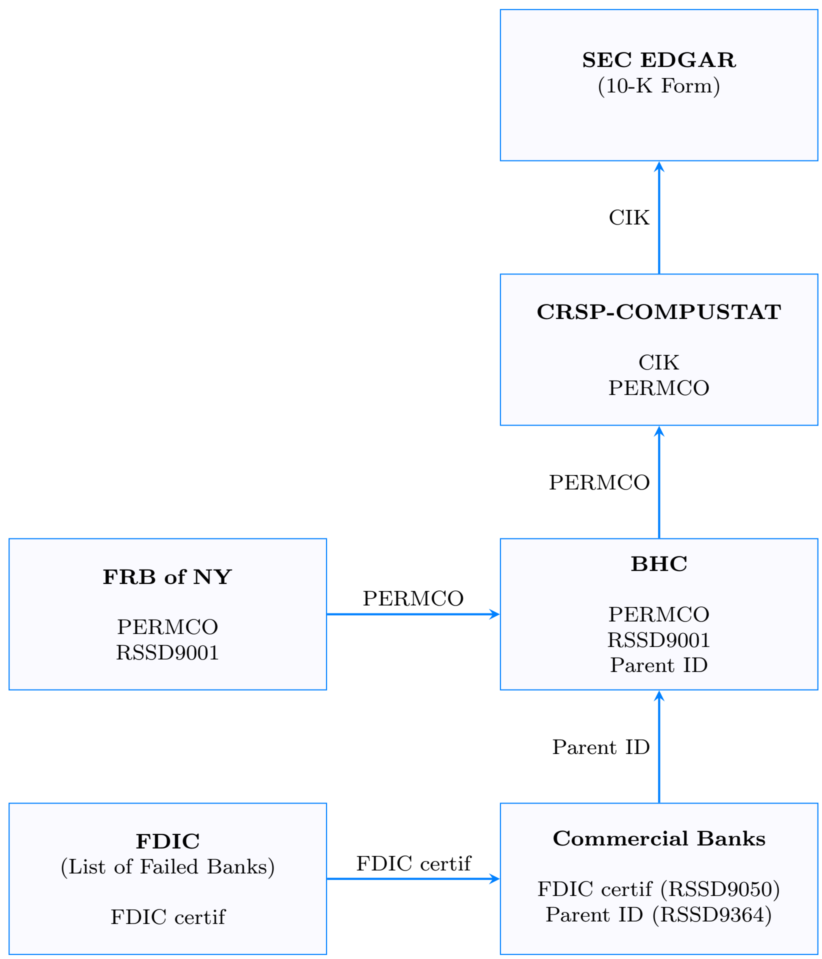

B List of Failed Banks

Table B.1: **List of Failed Banks**
This table reports all public failed banks identified in our sample. Failure is defined according to the FDIC. CIK is the Central Index Key used to identify banks in the SEC EDGAR system. Last Date denotes the last quarterly FRY-9C available at the bank holding company level. The Assets column reports the bank's total assets in millions of dollars. Delisting Code is the delisting code from the CRSP database - the 200 code relates to mergers, and the 500 states whether the company was dropped from the sample. Finally, the Acquired column returns one if the bank was acquired and zero otherwise.

	Name	CIK	Last Date	Assets	DLSTCD	Acquired	Resolution
1	WACHOVIA CORP	36995	2008-09-30	764378	231	1	PI
2	U S BC	101542	1997-06-30	33989		1	PI
3	COLONIAL BANCGROUP	92339	2009-03-31	26440	561	0	PA
4	W HOLD CO	1084887	2009-09-30	13490	585	0	PA
5	UCBH HOLD	1061580	2009-03-31	13419		0	PA
6	R&G FNCL CORP	1016933	2004-12-31	11931		0	PA
7	ALABAMA NBC	926966	2007-09-30	7967		1	PA
8	IMPERIAL BC	49899	2000-12-31	7549		1	PI
9	MAGNA GROUP	36094	1998-03-31	7251	231	1	PA
10	CORUS BSHRS	51939	2009-06-30	7064		0	PA
11	IMPERIAL CAP BC	1000234	2008-12-31	4440		0	PA
12	AMCORE FNCL	714756	2009-12-31	3777		0	PA
13	FRONTIER FC	716457	2009-12-31	3595	574	0	PA
14	IRWIN FC	52617	2009-06-30	3367		0	PA
15	MIDWEST BANC HOLD	1051379	2010-03-31	3182		0	PA
16	EUROBANCSHARES	1164554	2009-09-30	2807	580	0	PA
17	SECURITY BK CORP	925464	2009-03-31	2786		0	PA
18	FIRST ST BC	897861	2010-06-30	2627		0	PA
19	INTEGRA BK CORP	764241	2011-03-31	2180	560	0	PA
20	FIRST RGNL BC	356708	2009-09-30	2175		0	PA
21	VINEYARD NAT BC	840256	2008-09-30	2099		0	PA
22	SILVER ST BC	1362719	2008-06-30	1961		0	PI
23	CAPITAL CORP OF THE WEST	1004740	2008-09-30	1873		0	PA
24	HAMILTON BC	894172	2001-09-30	1411		0	PI
25	BANC CORP	1065298	2005-09-30	1376		0	PA
26	COMMONWEALTH BSHRS	835012	2010-12-31	1097		0	PA
27	NEXITY FC	1084727	2008-12-31	1062		0	PA
28	COLUMBIA BC	1010002	2009-09-30	1058		0	PA
29	PAB BSHRS	705200	2010-09-30	1044	584	0	PA
30	CRESCENT BKG CO	883476	2010-03-31	1014		0	PA
31	COOPERATIVE BSHRS	923529	2009-03-31	967		0	PA
32	PRINCETON NAT BC	707855	2012-06-30	954		0	PA
33	DEARBORN BC	895541	2010-12-31	916		0	PA
34	OMNI FNCL SVC	1362599	2007-09-30	887		0	PO
35	MERCANTILE BANCORP	1289701	2011-12-31	834		0	PA
36	WGNB CORP	1115568	2009-09-30	833		0	PA
37	TEAM FNCL	1082484	2008-09-30	766		0	PA
38	CAPITALSOUTH BC	1338977	2008-09-30	696		0	PA
39	BEACH FIRST NAT BSHRS	949228	2009-12-31	614		0	PA
40	SUN AMER BC	1042521	2009-09-30	555		0	PA
41	COMMUNITY VALLEY BC	1170833	2009-09-30	554		0	PA
42	COWLITZ BC	894267	2010-03-31	530	584	0	PA
43	COMMUNITY CENTRAL BK CORP	1014133	2010-09-30	514		0	PA
44	HABERSHAM BC	754597	2009-12-31	456		0	PA
45	FIRST ST FNCL CORP	1282582	2009-03-31	441		0	PA
46	PACIFIC ST BC	1169424	2005-12-31	310		0	PA
47	SDNB FC	702147	1996-09-30	192	233	0	PA

C Variables Definition

Table C.1: **Variables Definition**
The Code column denotes the FR 9-YC's relevant items used to compute the financial variables.

	Variable	Definition	Code
1	size	Size is the natural log of total assets	log(BHCK2170)
2	cap_ratio	Capital ratio is total book equity over total assets	$\frac{\text{BHCK3210}}{\text{BHCK2170}}$
3	roa	Return on assets (ROA), net income divided by total assets	$\frac{\text{BHCK4340}}{\text{BHCK2170}}$
4	non_perf_loan	Non-performing loans to total loans	$\frac{\text{BHCK5525 - BHCK3506 + BHCK5526 - BHCK3507 + BHCK1616}}{\text{BHCK2122}}$
5	alll	Allowance for loan and lease losses to total assets	$\frac{\text{BHCK3123}}{\text{BHCK2170}}$
6	llp	Loan loss provisions to total assets	$\frac{\text{BHCK4230}}{\text{BHCK2170}}$
7	deposit_ratio	Total deposits to total assets	$\frac{\text{BHDM6631 + BHDM6636 + BHFN6631 + BHFN6636}}{\text{BHCK2170}}$
8	ci_loan_ratio	Commercial and industrial loans to total loans	$\frac{\text{BHCK1763 + BHCK1764}}{\text{BHCK2122}}$
9	re_ratio	Real estate loans to total loans	$\frac{\text{BHCK1410}}{\text{BHCK2122}}$
10	loan_assets_ratio	Total loans to total assets	$\frac{\text{BHCK2122}}{\text{BHCK2170}}$
11	cost_income_ratio	Cost-to-income ratio	$\frac{\text{BHCK4093}}{\text{BHCK4074 + BHCK4079}}$
12	non_int_income	Non-interest income share	$\frac{\text{BHCK4079}}{\text{BHCK4074 + BHCK4079}}$

D Additional Results

Table D.1: **Summary Statistics for Textual Data**
This table provides summary statistics for the Loughran and McDonald (2011) dictionaries. Each panel corresponds to a different group of banks. Non-failed banks in Panel (a) are the subset of all banks that did not drop from the sample since becoming public. The failed banks in Panel (b) denote the group that were received by the FDIC as failed. The target banks in Panel (c) correspond to the banks that were acquired without flagged as failed.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Panel (a) Non-Failed Banks							
N_Tone	3,012	−1.073	0.423	−3.125	−1.347	−0.791	0.698
N_Negative	3,012	1.640	0.380	0.159	1.362	1.887	3.513
N_Positive	3,012	0.567	0.135	0.215	0.482	0.634	1.507
N_Uncertainty	3,012	1.257	0.247	0.127	1.092	1.429	2.127
N_Litigious	3,012	1.266	0.679	0.356	0.793	1.558	4.113
N_ModalWeak	3,012	0.441	0.120	0.095	0.356	0.526	0.865
N_Modal_Moderate	3,012	0.311	0.078	0.000	0.256	0.366	0.580
N_Modal_Strong	3,012	0.250	0.082	0.073	0.199	0.287	0.892
N_Constraining	3,012	0.863	0.158	0.365	0.754	0.974	1.448
Panel (b) Failed Banks							
N_Tone	371	−0.999	0.395	−2.053	−1.284	−0.713	0.092
N_Negative	371	1.574	0.353	0.780	1.296	1.833	2.463
N_Positive	371	0.574	0.150	0.241	0.482	0.645	1.541
N_Uncertainty	371	1.224	0.237	0.636	1.058	1.371	2.146
N_Litigious	371	1.336	0.766	0.501	0.824	1.610	4.219
N_ModalWeak	371	0.421	0.113	0.165	0.322	0.504	0.733
N_Modal_Moderate	371	0.296	0.076	0.123	0.244	0.346	0.613
N_Modal_Strong	371	0.260	0.069	0.122	0.219	0.291	0.774
N_Constraining	371	0.839	0.163	0.464	0.736	0.943	1.345
Panel (c) Targets							
N_Tone	1,840	−0.899	0.392	−2.999	−1.128	−0.644	0.422
N_Negative	1,840	1.491	0.344	0.383	1.248	1.687	3.066
N_Positive	1,840	0.593	0.158	−0.391	0.489	0.669	1.663
N_Uncertainty	1,840	1.127	0.224	0.543	0.971	1.272	2.040
N_Litigious	1,840	1.427	0.770	0.416	0.837	1.859	4.381
N_ModalWeak	1,840	0.392	0.111	0.120	0.314	0.462	0.806
N_Modal_Moderate	1,840	0.284	0.071	0.094	0.233	0.333	0.546
N_Modal_Strong	1,840	0.254	0.091	0.090	0.198	0.291	1.117
N_Constraining	1,840	0.808	0.166	0.406	0.684	0.917	1.522

E FinBERT Framework and Application

Background

The FinBERT large language model (LLM) (Huang et al., 2023) builds on the Bidirectional Encoder Representations from Transformers (BERT), a groundbreaking natural language processing (NLP) model developed by Google (Devlin, 2018) that transformed the way machines understand human language. BERT's key innovation lies in its ability to analyze words in their bidirectional context, where it simultaneously considers the words that come before and after a given word. This bidirectional approach allows BERT to capture the more complex meaning of words given a specific context. This makes the model highly effective for tasks like question answering, text classification, and sentiment analysis.

Building on BERT, the FinBERT model leverages the foundational architecture of BERT but adapts it specifically to the financial domain. Unlike general-purpose NLP models, which are trained on diverse and broad datasets like Wikipedia or news articles, FinBERT is fine-tuned on financial texts, including earnings call transcripts, sell-side analyst reports, and 10K/Q reports. Such domain-specific tuning enables FinBERT to understand better financial language's unique terminology, syntax, and context.

One of FinBERT's key innovations is its ability to extract sentiment signals from financial texts with high precision. It outputs probabilities for three sentiment classes: positive, neutral, and negative. This capability is particularly valuable in financial applications, where sentiment can drive market reactions and inform decision-making processes. For example, the FinBERT model can evaluate the tone of a CEO's commentary during an earnings call or assess the sentiment embedded in a company's 10-K report.

Compared to traditional sentiment analysis methods, such as bag-of-words approaches that rely on predefined dictionaries, FinBERT excels in capturing the nuanced and contextual nature of financial texts. It can handle subtle variations in tone, complex sentence structures, and domain-specific language, providing more accurate and actionable insights. According to Huang et al. (2023), FinBERT "substantially outperforms the Loughran and McDonald dictionary and other machine learning algorithms, including naive Bayes, support vector machine, random forest, convolutional neural network, and long short-term memory, in sentiment classification." For this reason, we choose the FinBert model to represent state-of-the-art natural language processing.

Workflow and Implementation

In the discussion below, we provide an overview of the model's workflow. In particular, Figure E.2 illustrates the workflow of the FinBERT model, where the process begins with "Input Embeddings." Each token in the input sequence is represented as the sum of three embeddings: token embeddings, segment embeddings (indicating different text segments), and position embeddings (reflecting token order). These embeddings capture the linguistic and positional context of each token.

The token embeddings are fed into a "Transformer Block," which consists of two key components: an attention mechanism and a "Feed-Forward Network." The attention mechanism captures relationships between tokens in the sequence, allowing the model to weigh their importance in context. The Feed-Forward Network further refines the output to generate "Contextualized Embeddings" for each token. The contextualized representation of the special [CLS] token is passed to a classification layer, consisting of a linear layer followed by a softmax activation. This layer generates sentiment scores (positive, neutral, or negative), enabling the model to analyze the financial text effectively. Regarding implementation, we employ the FinBERT model developed by Huang et al. (2023).[25]

To run the FinBERT model, we first segment each 10-K report into smaller text chunks. Since FinBERT can process up to 512 tokens per inference, we divide each document into segments of approximately 500 tokens. For our sample, this procedure yields about 100 chunks per 10-K report. We then feed each chunk into FinBERT, which generates both a sentiment label (positive, neutral, or negative) and corresponding scores. These scores represent the model's probabilistic assessment of each sentiment category and range between 0 and 1, indicating the confidence level of the assigned sentiment.

After processing all chunks within a 10-K report, we aggregate the chunk-level results to form document-level sentiment features. Specifically, we average the sentiment scores from all chunks to obtain positive, neutral, and negative sentiment measures for the entire report. We further calculate a Net sentiment measure as the difference between the aggregated positive and negative sentiment scores. This approach allows us to distill a large volume of textual information into a manageable set of sentiment features for subsequent bankruptcy prediction models. Specifically, for each document, the

[25] The implementation is publicly available on Hugging Face `https://huggingface.co/yiyanghkust/finbert-tone`.

following features are computed:

$$\text{Positive} = \frac{\sum_i^{N_p} POS_i}{\sum_i^{N_{POS}} POS_i + \sum_i^{N_{NEG}} NEG_i + \sum_i^{N_{NEU}} NEU_i}, \tag{E.1}$$

$$\text{Negative} = \frac{\sum_i^{N_{NEG}} NEG_i}{\sum_i^{N_{POS}} POS_i + \sum_i^{N_{NEG}} NEG_i + \sum_i^{N_{NEU}} NEU_i}, \tag{E.2}$$

$$\text{Neutral} = \frac{\sum_i^{N_{NEU}} NEU_i}{\sum_i^{N_{POS}} POS_i + \sum_i^{N_{NEG}} NEG_i + \sum_i^{N_{NEU}} NEU_i}, \tag{E.3}$$

$$\tag{E.4}$$

where N_{POS}, N_{NEG}, and N_{NEU} denote the number of chunks classified as positive, negative, and neutral, respectively. POS_i, NEG_i, and NEU_i are the sentiment scores for the i-th positive, negative, and neutral chunks, respectively. Finally, given the scores of each, we compute the document net tone (i.e., the net_bert variable) as

$$\text{net\_bert} = \frac{\text{Positive} - \text{Negative}}{\text{Neutral}}. \tag{E.5}$$

FinBERT-Related Figures

Figure E.1: **Partial Dependence Plots for Failed versus Non-Failed in 2008 using FinBERT** This figure demonstrates the effect of three features on bank failures. The illustration is conducted with respect to the leave-one-out approach from Table E.2 for 2008. The panels below demonstrate the partial dependency plots (PDPs; Greenwell (2017)). PDPs render the mapping function and provide a low-dimensional illustration of the relationship between the outcome variable (event of default) and the predictors of interest. The first three panels are univariate and investigate the relationship between one feature and the probability of default (PD). The plot is consistent with Figure 6 of the manuscript, whereas the main difference is that the analysis below relies on the FinBERT net tone to represent the textual data.

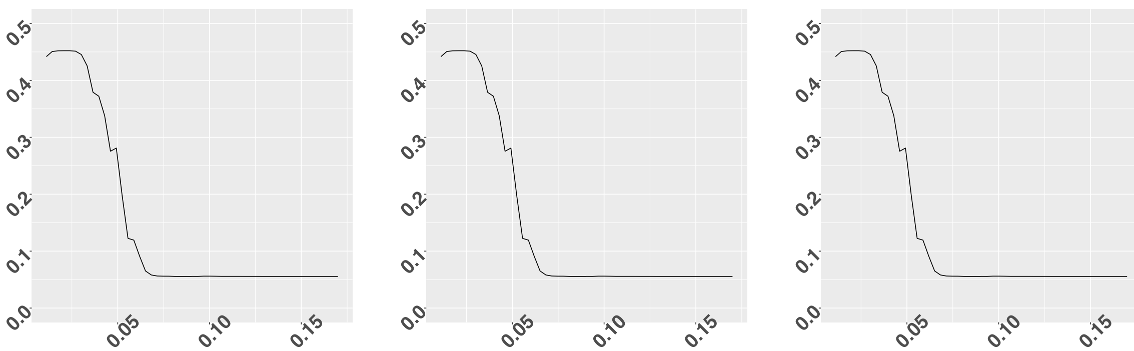

(a) Capital Ratio (b) Non-Performing Loans (c) FinBERT Net Tone

Figure E.2: **FinBERT Structure**
This diagram illustrates the architecture of FinBERT, a transformer-based model fine-tuned for financial sentiment analysis. The model begins with "Input Embeddings," which include token embeddings, segment embeddings, and positional embeddings. These embeddings pass through multiple "Transformer Blocks" comprising multi-head self-attention mechanisms and feed-forward networks, allowing the model to generate "Contextualized Embeddings" per Token. The [CLS] token's contextualized embedding is fed into the "Classification Head," which employs a linear layer followed by a softmax function to output Sentiment Scores (positive, neutral, or negative). FinBERT's architecture (Huang et al., 2023), adapted from BERT Devlin (2018), is designed to capture nuanced relationships in financial text, making it particularly effective for tasks like financial sentiment analysis, event detection, and risk assessment in unstructured data.

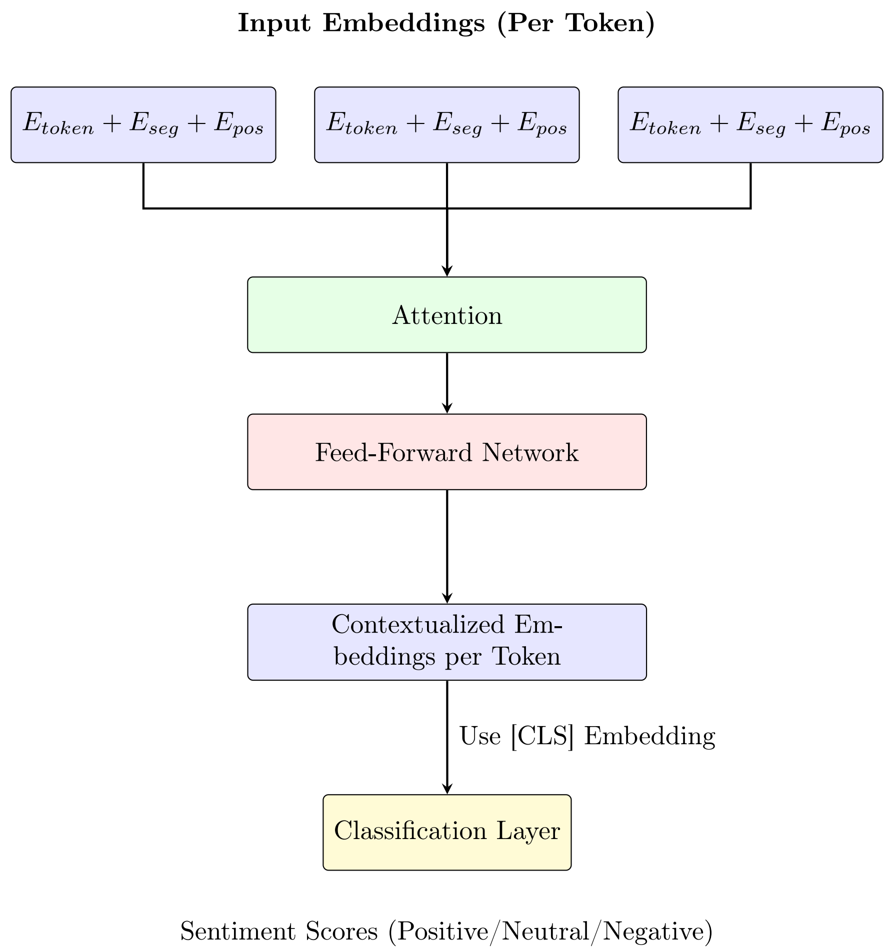

FinBERT-Related Tables

Table E.1: **FinBERT: Summary Statistics and Correlation Matrix**
The table below reports summary statistics and the correlation matrix for the FinBERT sentiment scores. In particular, the positive_bert and negative_bert correspond to the confidence levels that the 10-K report is associated with positive and negative sentiment, respectively. The net tone (net_bert) is the difference between the two adjusted by the confidence level that the report's sentiment is neutral. Panel (a) covers the summary statistics of the final annual panel used in our machine learning analysis. Note that the confidence scores for positive and negative sentiment range between zero and one, reflecting a probability value, whereas the values of net_bert are unbounded. Panel (b) reports the pairwise correlation coefficients between the FinBERT scores and the sentiment values from Loughran and McDonald (2011), which are denoted by N_.

Panel (a) Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Max
net_bert	5,380	0.022	0.070	−0.268	0.474
negative_bert	5,380	0.030	0.035	0.000	0.218
positive_bert	5,380	0.050	0.042	0.000	0.322

Panel (b) Correlation Matrix Statistics

	N_Tone	N_Negative	N_Positive	net_bert	negative_bert	positive_bert
N_Tone	1	-0.938	0.483	0.356	0.569	-0.561
N_Negative	-0.938	1	-0.149	-0.237	-0.505	0.599
N_Positive	0.483	-0.149	1	0.415	0.350	-0.087
net_bert	0.569	-0.505	0.350	0.830	1	-0.694
negative_bert	-0.561	0.599	-0.087	-0.179	-0.694	1
positive_bert	0.356	-0.237	0.415	1	0.830	-0.179

Table E.2: **Descriptive Analysis using Leave-One-Out Approach using FinBERT**
This table reports the leave-one-out (LOO) results for the years ranging between 2000 and 2009 (inclusive). For each year y, the LOO approach trains random forest (rf) using $N_y - 1$ observations, leaving one for testing. This procedure is repeated until N_y fitted values are attained. See Section 4.2 for further information. Accuracy, Precision, Recall, and F_1 denote the performance metrics from Section 4.4. In Panel (a), the LOO classifies between failed and non-failed banks. The former group is identified as the banks that the FDIC received for resolution. The latter group is the surviving banks after controlling for mergers & acquisitions and delisting (see Table 1). Panel (b) is similar to Panel (a), whereas the control group comprises the target banks acquired through acquisition. The column #Fail denotes the number of observations out of N_y corresponding to the failed banks. Each panel is split into two: "Financial Data Only" and "Financial Data and FinBERT." The latter repeats the same experiment, however, with the inclusion of net tone computed using the FinBERT model (Huang et al., 2023). The table is consistent with Table 3 of the manuscript, whereas the main difference is that the analysis below relies on the FinBERT net tone to represent the textual data.

Financial Data Only						**Financial Data and FinBERT**					
Year	#Fail	Accuracy	Precision	Recall	F_1	Year	#Fail	Accuracy	Precision	Recall	F_1
Panel (a) Failed Vs. Non-Failed											
2000	24	0.85	0.17	0.04	0.07	2000	24	0.86	0.00	0.00	
2001	26	0.87	0.40	0.08	0.13	2001	26	0.87	0.50	0.08	0.13
2002	28	0.86	0.33	0.04	0.06	2002	28	0.86	0.25	0.04	0.06
2003	29	0.86	0.38	0.10	0.16	2003	29	0.87	0.40	0.07	0.12
2004	29	0.88	0.57	0.14	0.22	2004	29	0.89	0.80	0.14	0.24
2005	29	0.87	0.45	0.17	0.25	2005	29	0.87	0.40	0.21	0.27
2006	32	0.90	0.75	0.38	0.50	2006	32	0.90	0.71	0.38	0.49
2007	34	0.94	0.85	0.68	0.75	2007	34	0.94	0.82	0.68	0.74
2008	28	0.96	0.83	0.71	0.77	2008	28	0.96	0.83	0.71	0.77
2009	11	0.99	0.82	0.82	0.82	2009	11	0.99	0.82	0.82	0.82
Panel (b) Failed Vs. Targets											
2000	24	0.86		0.00		2000	24	0.81	0.10	0.04	0.06
2001	26	0.82	0.20	0.04	0.06	2001	26	0.84		0.00	
2002	28	0.81	0.33	0.07	0.12	2002	28	0.80	0.33	0.11	0.16
2003	29	0.83	0.62	0.34	0.44	2003	29	0.84	0.69	0.31	0.43
2004	29	0.79	0.50	0.17	0.26	2004	29	0.80	0.62	0.17	0.27
2005	29	0.78	0.56	0.31	0.40	2005	29	0.79	0.64	0.31	0.42
2006	32	0.72	0.57	0.38	0.45	2006	32	0.75	0.69	0.34	0.46
2007	34	0.88	0.88	0.82	0.85	2007	34	0.89	0.90	0.82	0.86
2008	28	0.88	0.85	0.82	0.84	2008	28	0.88	0.85	0.82	0.84
2009	11	0.89	0.69	0.82	0.75	2009	11	0.89	0.69	0.82	0.75

Table E.3: **Predictive Analysis using Capital Ratio Forecast with FinBERT**
This table reports the predictive results based on the capital ratio forecast described in Section 4.3. The predictive model is trained at the end of each year (between 1999 and 2008) based on the data from the last four years. The forecast denotes the next year's capital ratio. Banks are deemed as failed over the next year if their forecast ranks at the bottom 20% percentile. Accuracy, Precision, Recall, and F_1 denote the performance metrics from Section 4.4. In Panel (a), the treated group is the failed banks that the FDIC received for resolution. The control group in Panel (a) covers all other banks in the sample. In Panel (b), we adjust the treated and control groups such that the former includes distressed targets (see Section 3.1). The column #Fail denotes the number of observations out of N_y corresponding to the failed banks. Each panel is split into two: "Financial Data Only" and "Financial Data and FinBERT Data." The latter repeats the same experiment, however, with the inclusion of net tone computed using the FinBERT model (Huang et al., 2023). The table is consistent with Table 5 of the manuscript, whereas the main difference is that the analysis below relies on the FinBERT net tone to represent the textual data.

Financial Data Only						**Financial Data and FinBERT**					
Year	#Fail	Accuracy	Precision	Recall	F_1	Year	#Fail	Accuracy	Precision	Recall	F_1
Panel (a) Identifying Failures											
2000	24	0.79	0.14	0.42	0.21	2000	24	0.80	0.15	0.46	0.23
2001	26	0.77	0.11	0.31	0.16	2001	26	0.78	0.12	0.35	0.18
2002	28	0.75	0.07	0.18	0.10	2002	28	0.75	0.07	0.18	0.10
2003	29	0.74	0.05	0.14	0.08	2003	29	0.75	0.07	0.17	0.10
2004	29	0.76	0.08	0.21	0.12	2004	29	0.76	0.10	0.24	0.14
2005	29	0.77	0.13	0.31	0.18	2005	29	0.77	0.13	0.31	0.18
2006	32	0.76	0.15	0.31	0.20	2006	32	0.76	0.15	0.31	0.20
2007	34	0.87	0.44	0.82	0.57	2007	34	0.87	0.44	0.82	0.57
2008	28	0.87	0.37	0.89	0.53	2008	28	0.87	0.38	0.93	0.54
2009	11	0.84	0.17	1.00	0.29	2009	11	0.84	0.17	1.00	0.29
Panel (b): Identifying Failures and Acquisition through Distress											
2000	34	0.77	0.15	0.32	0.21	2000	34	0.77	0.17	0.35	0.23
2001	38	0.75	0.12	0.24	0.16	2001	38	0.75	0.14	0.26	0.18
2002	38	0.73	0.08	0.16	0.11	2002	38	0.73	0.08	0.16	0.11
2003	44	0.72	0.08	0.14	0.10	2003	44	0.72	0.08	0.14	0.10
2004	44	0.74	0.14	0.23	0.17	2004	44	0.74	0.15	0.25	0.19
2005	45	0.76	0.22	0.33	0.26	2005	45	0.76	0.22	0.33	0.26
2006	49	0.74	0.22	0.31	0.26	2006	49	0.74	0.21	0.29	0.24
2007	49	0.85	0.50	0.65	0.57	2007	49	0.85	0.50	0.65	0.57
2008	41	0.86	0.46	0.76	0.57	2008	41	0.87	0.47	0.78	0.59
2009	23	0.88	0.36	1.00	0.53	2009	23	0.88	0.36	1.00	0.53

Table E.4: **Interpretability for Descriptive Analysis using FinBERT**
This table reports the sensitivity analysis of the leave-one-out (LOO) approach. The numbers below describe the importance of each variable with respect to the Recall value. The importance scores are computed with respect to the Permutation Feature Importance according to Biecek (2018). A higher sensitivity score implies higher importance for the corresponding variable. Importance is determined in terms of contribution to the Recall performance reported in Table E.2. The features are ranked with their respective score according to 2008. The table is consistent with Table 4 of the manuscript, whereas the main difference is that the analysis below relies on the FinBERT net tone to represent the textual data.

	Variable	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
	Panel (a): Failed vs. Non-Failed										
1	cap_ratio	30.42	7.31	11.43	4.83	15.52	0.69	10.31	29.41	61.43	96.36
2	non_perf_loan	8.33	0.38	0.00	3.10	7.24	7.93	34.06	24.41	17.50	24.55
3	cost_income_ratio	2.50	16.92	12.86	17.24	12.41	0.00	0.62	2.94	9.29	0.00
4	non_int_income	15.83	0.00	18.21	5.52	19.31	31.38	19.37	1.76	2.50	0.00
5	alll	0.00	3.85	6.79	4.83	9.66	6.21	5.94	0.29	0.00	0.00
6	ci_loan_ratio	0.00	0.00	0.00	8.97	5.17	5.52	5.31	0.00	0.00	0.00
7	deposit_ratio	3.75	10.77	6.79	4.48	8.97	6.55	8.13	4.12	0.00	0.00
8	llp	2.92	0.00	11.07	8.62	8.97	18.62	0.00	7.94	0.00	0.00
9	loan_assets_ratio	9.17	18.08	21.07	12.76	3.79	10.69	18.44	0.00	0.00	0.00
10	net_bert	6.67	3.85	7.86	4.14	12.07	2.76	0.94	2.06	0.00	0.00
	Panel (b): Failed vs. Targets										
1	non_perf_loan	3.33	4.62	14.29	0.00	2.41	6.55	3.44	3.24	23.93	11.82
2	cap_ratio	40.83	4.62	11.07	0.00	4.48	3.10	11.56	8.24	9.29	57.27
3	size	10.83	2.69	10.00	8.97	12.41	10.00	2.81	0.00	3.21	0.00
4	alll	0.00	8.08	7.14	3.10	6.21	3.45	0.00	0.00	0.00	0.00
5	ci_loan_ratio	0.00	3.46	2.50	9.31	7.93	6.90	5.00	0.00	0.00	0.00
6	cost_income_ratio	12.50	15.00	22.14	26.21	17.59	12.07	0.00	0.00	0.00	0.00
7	deposit_ratio	3.33	3.46	5.00	4.83	1.72	5.17	0.31	0.00	0.00	0.00
8	llp	20.42	10.77	7.86	17.59	14.48	10.69	4.06	4.12	0.00	0.00
9	loan_assets_ratio	33.33	10.38	19.29	11.38	6.90	3.10	1.56	0.00	0.00	0.00
10	net_bert	2.92	8.85	16.07	9.66	11.72	0.00	0.00	0.00	0.00	0.91

Table E.5: **Sensitivity Analysis for Predictive Model using FinBERT**
This table reports the sensitivity analysis of the predictive model from Table E.3. The numbers below describe the importance of each variable with respect to next year's capital ratio forecast. The importance scores are computed with respect to the Permutation Feature Importance according to Biecek (2018). A higher sensitivity score implies higher importance for the corresponding variable. Importance is determined in terms of contribution to the following year's capital ratio and, hence, the financial soundness of the bank over the next year (Berger and Bouwman (2013)). The features are ranked with their respective score for 2008. The table is consistent with Table 6 of the manuscript, whereas the main difference is that the analysis below relies on the FinBERT net tone to represent the textual data.

	variable	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
1	cap_ratio	99.67	99.84	100.18	99.46	100.12	100.04	100.36	88.38	79.07	97.52
2	roa	1.48	2.05	0.25	1.68	0.39	0.33	0.42	6.62	4.43	0.71
3	non_perf_loan	0.20	0.02	0.02	0.00	0.28	0.08	0.41	5.77	2.84	0.10
4	cost_income_ratio	0.48	0.06	0.34	0.15	0.64	0.12	0.89	1.36	1.41	0.35
5	size	0.29	0.22	0.07	0.07	0.63	0.27	1.56	1.01	1.37	0.67
6	re_ratio	0.17	0.07	0.22	0.55	0.17	0.05	0.24	1.62	0.90	0.22
7	llp	0.38	0.08	0.03	0.16	0.04	0.61	0.68	2.73	0.65	0.05
8	loan_assets_ratio	0.24	0.08	0.00	0.08	0.17	0.23	0.20	0.19	0.54	1.24
9	net_bert	0.25	0.13	0.06	0.16	0.21	0.04	1.14	0.01	0.45	0.04
10	ci_loan_ratio	0.14	0.06	0.06	0.16	0.37	0.09	0.04	0.09	0.29	0.14