Social Media as a Bank Run Catalyst*

J. Anthony Cookson
Colorado – Boulder

Corbin Fox
James Madison University

Javier Gil-Bazo
Universitat Pompeu Fabra

Juan F. Imbet
Université Paris Dauphine

Christoph Schiller
Arizona State University

April 18, 2023

Abstract

Social media fueled a bank run on Silicon Valley Bank (SVB), and the effects were felt broadly in the U.S. banking industry. We employ comprehensive Twitter data to show that preexisting exposure to social media predicts bank stock market losses in the run period even after controlling for bank characteristics related to run risk (i.e., mark-to-market losses and uninsured deposits). Moreover, we show that social media amplifies these bank run risk factors. During the run period, we find the intensity of Twitter conversation about a bank predicts stock market losses at the hourly frequency. This effect is stronger for banks with bank run risk factors. At even higher frequency, tweets in the run period with negative sentiment translate into immediate stock market losses. These high frequency effects are stronger when tweets are authored by members of the Twitter startup community (who are likely depositors) and contain keywords related to contagion. These results are consistent with depositors using Twitter to communicate in real time during the bank run.

Keywords: Bank Runs; Social Media; Social Finance; FinTech
JEL Codes: F30; F36; G38; Q50

*We thank Asaf Bernstein, Jamie Brown, Erica Jiang, Will Mullins and Christophe Spaenjers for helpful comments. All remaining errors are our own.

Author contact: Cookson (tony.cookson@colorado.edu), Fox (fox2ca@jmu.edu), Gil-Bazo (javier.gil-bazo@upf.edu), Imbet (juan.imbet@dauphine.psl.eu), and Schiller (christoph.schiller@asu.edu).

1 Introduction

This paper investigates social media's role in the bank run that led to the failure of Silicon Valley Bank (SVB) on March 10, 2023. Swaths of clients with deposits exceeding the FDIC insurance threshold requested immediate withdrawal of their funds. The failure of SVB was preceded by a large spike of public communication on Twitter by apparent depositors in SVB who used the forum to discuss the trouble the bank was facing, and more importantly, their intentions to withdraw their deposits from SVB. The openness and speed of this coordination around a bank run is unprecedented. Indeed, given the importance of communication in leading to bank runs in equilibrium (Diamond and Dybvig, 1983; Angeletos and Werning, 2006), observers were quick to suggest that social media may have fueled the run on SVB.

In this paper, we present evidence that social media did, indeed, contribute to the run on SVB. More importantly, our analysis suggests that other banks face similar risks. We collect a comprehensive sample of tweets about all publicly traded banking stocks, and analyze their content, dynamics, and the social transmission of bank run ideas from investor social networks to connected networks of depositors. One core insight is that SVB faced a novel channel of bank run risk that is unique to the social media era. SVB depositors active on social media played a central role in the bank run. These depositors were concentrated and highly networked through the venture capital industry and founder networks on Twitter, amplifying other bank run risks. More importantly, SVB is not the only bank to face this novel risk channel: Open communication by depositors via social media increased the bank run risk for other banks that were *ex ante* exposed to such discussions in social media.

To fix ideas about why social media can amplify bank run risk, we build on the model intuition in Jiang, Matvos, Piskorski, and Seru (2023b) who focus on the run risk originating from uninsured deposits. They argue formally that bank run risk depends on depositors' beliefs about the fraction of uninsured depositors s who request to withdraw their deposits. A simple, reduced form way to illustrate this is to examine the formula for the insured deposit coverage ratio (ins_cov_ratio), which depends on the share of uninsured deposits that run, s:

$$ins\_cov\_ratio_i = \frac{Assets_i - s \times uninsured\_dep_i - insured\_dep_i}{insured\_dep_i} \quad (1)$$

Jiang et al. (2023b) simulate scenarios where all or half of uninsured deposits run (i.e., $s = 1$ and $s = 0.5$) to gain an understanding of system-wide bank run risk. Of course, s varies across banks (i.e., s_i). We argue that greater exposure to social media increases s_i, which amplifies bank run risk through a novel channel.

This is precisely what we find using a comprehensive sample of bank stock returns, collected from FirstRate Data, and tweets about U.S. banks. We begin by estimating a cross-sectional specification that explains the collapse in bank stock prices from March 1 through March 14 in relation to the extent of Twitter conversation before this "run period", mark-to-market bank losses, and the percentage of uninsured bank deposits. We find that banks in the top tercile of preexisting Twitter conversation have an average of 6 percentage points larger stock market loss during the run period. This main effect is greater than the decline in bank stock prices for a standard deviation increase in uninsured deposits, and it is also robust to controlling for uninsured deposits. Consistent with Jiang et al. (2023b), we find that mark-to-market losses on their own do not explain bank exposure to run risk. Rather, the interaction between mark-to-market losses and percent uninsured deposits – which we dub 'run exposure' – matters for the decline in bank stock prices during the run.

More importantly, we find that pre-run exposure to Twitter conversation *interacts* significantly with both uninsured deposits and the combination of mark-to-market losses and uninsured deposits. Both the interaction between uninsured deposits and Twitter conversation intensity and the triple interaction between mark-to-market losses, uninsured deposits, and Twitter conversation intensity significantly predict larger bank stock losses during the run period. Moreover, the main effects of uninsured deposits are insignificant and small after including the interaction with Twitter conversation intensity.

Next, we investigate the content of the conversation surrounding the run. Using a dictionary of terms associated with depositors withdrawing their deposits (e.g., tweets that mention "withdraw"), we find that most of the asset declines associated with pre-market social media exposure are driven by banks with more intensive "run behavior" discussions during the run period. Indeed, from March 8 through March 13, users posted 6,528 'run' tweets about SVB, which is roughly five times the number of next most run discussed ticker (i.e., First Republic Bank, FRC, which also had a notable run discussion). However, the mechanism linking pre-run Twitter exposure to run-based tweets and bank run price declines is not limited to SVB, but rather is broadly present in our sample of publicly traded banking firms. We see a similar pattern with tweets that mention banking contagion: SVB had 9,662 such tweets during the run period, and controlling for this systematically moderates the main effect of pre-run Twitter conversation intensity.

Going further, we exploit the high frequency nature of the Twitter conversation to provide evidence that social media exposure led to bank run risk rather than simply reflecting it. To do this, we combine the high frequency of Twitter posts and the fact that what matters for bank run risk is conversation *among depositors*. The Twitter data allows us to separate and track conversations about SVB's ticker S*"I"*VB separately from more general conversations about SVB (i.e., SVB, Silicon Valley Bank, etc.) during the bank run (March 8th through March 10th). The use of the ticker (SIVB) reliably distinguishes investor-contributed tweets from general conversation. This decomposition is informative: Viewing these series separately (see Figure 1), investors began discussing SVB with tweets that reference the ticker SIVB. Distinctly later, these conversations were followed by a much larger volume of general tweets about the bank. This pattern is consistent with depositors communicating on Twitter about withdrawing their deposits from the bank. This case study suggests that investor conversations spilled over into depositor conversations that fueled the run on SVB.

SVB's depositors were also highly concentrated in the start-up community, which held large uninsured deposits. Due to the highly networked nature of this community, encouraged by venture capital firms, these depositors were not only concentrated in their deposits in SVB, but they also exhibit a high degree of communication on Twitter. To speak to the role of the startup community in the bank run, we classify users as part of the startup community by constructing a dictionary of startup and founder words in user profile descriptions (e.g., "entrepreneur" or "founder"). Consistent with this role in the bank run, we see startup community tweets spike distinctly after the initial increase in tweet volume, and they are more likely to mention SVB than SIVB (see Figure 2). In the cross-section of banks, we then repeat the main analysis, based on pre-period exposure to the startup community distinct of general exposure to it. The findings strengthen when we focus on exposure to startup community tweets. Together with our main results, these findings point to the importance of highly concentrated and highly networked communication for bank run risk, beyond SVB.

A potential concern with these cross-sectional tests is that the timing of the tweets in relation to the bank stock returns is imprecise, particularly for tweets posted during the run period. To address this issue, we estimate similar regressions using *hourly* bank stock returns. First, we find a negative effect of 'run exposure' (i.e., the interaction of MTM asset losses with % uninsured deposits) on bank returns. Crucially, this effect emerges only starting on March 09, 2023 using hourly returns. Next, we estimate a triple-difference specification with an indicator for after the onset of the run (i.e., post-March 9, 9am), the number of tweets in the past four hours, and banks' run-exposure. By lagging the number of tweets, these tests better distinguish between Twitter's effect on returns and the reverse effect, helping us us shed light on the social contagion effect of social media on the bank run risk in our sample. We find a large, significant effect of Twitter conversation intensity in the preceding 4 hours on bank returns, which is higher for banks with greater run-exposure after the onset of the SVB run. Comparing only within the group of 'high' run-exposure banks, we find that banks in the top tercile of Twitter activity lost an additional 15% average cumulative return compared to banks with low Twitter activity (returns of -30% vs. -15%) between March 06 and March 14. This effect holds even after excluding SVB from the sample.

As a final leg of our empirical analysis, we employ several tests that exploit the high frequency timing of tweets in and out of the run period, building on the approach in Bianchi, Cram, and Kung (2021) and Bianchi, Gómez-Cram, Kind, and Kung (2023). These tests, which rely on high frequency bank stock price information from FirstRate, are useful to rule out confounding events that occur outside of a narrow window from 5 minutes before the tweet to 5 minutes afterwards. We estimate a significant negative impact of neutral sentiment tweets posted during the run period that is small and insignificant prior to the run period. However, the sentiment of the tweets is a significant predictor: more negative sentiment tweets correspond to more negative high-frequency returns. These sentiment effects are stronger if the tweets are authored by someone in the Twitter startup community, and they are stronger when the tweets contain words in the "contagion" dictionary. Taken together with our other empirical tests, these findings provide strong evidence that social media helped fuel the bank run, and indeed, played this significant role by amplifying balance sheet risks.

Our main contribution is to provide novel evidence of rapid communication via social media as a bank run risk. Diamond and Dybvig (1983)'s insight is that coordination is foundational to bank run models. Building on this insight, scholars developed models of coordination via communication, which have enjoyed broad application in our understanding of bank runs and financial crises (Peck and Shell, 2003; Angeletos and Werning, 2006; Brunnermeier, 2009; Acemoglu, Ozdaglar, and Tahbaz-Salehi, 2015). Since the Global Financial Crisis and the rise of shadow banks, research turned to investigate the risks posed by these new forms of banking (Jiang, Matvos, Piskorski, and Seru, 2020; Buchak, Matvos, Piskorski, and Seru, 2022, 2018), and the importance of run-prone uninsured depositors in generating possible bank run equilibria (Egan, Hortaçsu, and Matvos, 2017).[1] Relative to this literature, which focuses on bank balance sheets and risk taking, our contribution is to highlight the role of social media as a communication technology that amplifies existing risks

[1]As part of a series of papers on this March 2023 episode of banking distress, Jiang, Matvos, Piskorski, and Seru (2023a) emphasizes the role of limited hedging of bank asset exposures, while Jiang, Matvos, Piskorski, and Seru (2023c) highlights aspects of banks' involvements with commercial real estate loans.

identified in the literature. The implication that social media matters for banking stability is potentially troubling because social platforms can spread inaccurate information, which could serve as a sunspot that leads to bank runs.[2]

Our research also relates to the banking literature that focuses specifically on financial contagion via social networks during times of banking distress (Calomiris and Mason, 1997; Iyer and Puri, 2012). Like our work, this literature focuses on communication as a contagion mechanism in bank runs. However, the emphasis in this literature has been on how communication spreads through traditional social networks such as immigrant networks or word of mouth (e.g., Kelly and Gráda, 2000). Our focus is distinct because social *media* has at least two features that make it more powerful as a coordination mechanism than offline social networks. One important distinction is the speed of communication on social media platforms versus personal connections. Another aspect that interacts with the rapid speed of communication is the fact that information posted to Twitter is visible publicly, which transmits information well beyond close personal connections. Both of these aspects of social media lead to more rapid and widespread coordination. If social media continues to be a forum for depositors to share information, this would prove to be particularly powerful mechanism that can amplify bank run risk.

Our emphasis on social media as a communication and coordination medium relates to a broader literature in economics that has studied communication and contagion in networks (Chwe, 2000; Van Bommel, 2003; Golub and Jackson, 2010; Elliott, Golub, and Jackson, 2014; Falato, Hortacsu, Li, and Shin, 2021). With the advent of myriad technologies to communicate, there has been research into the real effects of this communication, particularly in politics (Parmelee and Bichard, 2011; Müller and Schwarz, 2021). In this respect, our contribution advances what is known about communication technologies such as radio (Strömberg, 2004)

[2] Related to the communication channel we identify, another strand of the bank run literature has emphasized the role of trust in the financial system (Iyer and Puri, 2012; Traweek and Wardlaw, 2022). Indeed, trust is instrumental in encouraging participation in banking and financial markets more broadly (Gurun, Stoffman, and Yonker, 2018; Brown, Cookson, and Heimer, 2019; Stein and Yannelis, 2020; Dupont, 2022). This literature ought to find our analysis relevant, particularly as mistrust of information on social media and otherwise increases.

and television (Kearney and Levine, 2019), which have been shown to have important effects due to their impacts on communication. In this broader literature, the closest paper is Ziebarth (2013), which studies the effect of radio as a communication technology on bank distress during the Great Depression. In contrast to radio, which provides sparse information and one-way communication, social media aggregates information from many sources, making it potentially a much stronger communication and coordination device in the context of communication about bank runs.

We also contribute to the social economics literature (Levy, 2021; Chetty et al., 2022a,b). In particular, we complement recent work on social finance, which has emphasized social transmission of ideas (Akçay and Hirshleifer, 2021; Hirshleifer, 2020; Kuchler and Stroebel, 2021).[3] Bank runs are by their nature a social phenomenon, and social media is a natural place to study the formation and transmission of ideas. Much of the literature in social finance has focused on investment outcomes (Cookson, Engelberg, and Mullins, 2022; Han, Hirshleifer, and Walden, 2022; Pedersen, 2022) only a few studies have looked into real outcomes – e.g., innovation and merger outcomes (Hirshleifer and Plotkin, 2021; Cookson, Niessner, and Schiller, 2022) Although some research on social media has investigated political outcomes and ideologies (Müller and Schwarz, 2022), we are the first to provide direct evidence of a social transmission channel via social media for bank runs. As the financial sector is core to the macroeconomy, our finding that social media meaningfully amplifies bank run risk is a fundamental contribution to the literature's understanding the real effects of social media.

[3]In a narrow sense, we relate to growing literature on financial social media (Chen, De, Hu, and Hwang, 2014; Giannini, Irvine, and Shu, 2019; Cookson and Niessner, 2020; Cookson, Lu, Mullins, and Niessner, 2022; Gil-Bazo and Imbet, 2022). This literature has focused on how social media relates to asset markets (Renault, 2017; Cookson, Engelberg, and Mullins, 2020; Bianchi et al., 2021, 2023), whereas our focus is on the broader conversation that emerged about depositor issues.

2 Data and Measurement

2.1 Twitter Sample

We begin by constructing a comprehensive picture of Twitter discussions about bank stocks in the last several years. Twitter grants academics access to their API to collect historical tweet-level and user-level data. Using Twitter's API, we collected all tweets from 1 January 2020 until 13 March 2023 containing at least one cashtag ($ followed by the company's ticker) for the universe of publicly traded bank firms (all depository institutions with the three-digit SIC code 602, 603, or 609). The practice of using cashtags on Twitter began in 2012 and has since been widely adopted as a way to reference publicly traded companies. Other than requiring an instance of a banks cashtag, we only consider original tweets (i.e. no retweets) that are written in English. These sample filters leave us with a sample of 5,399,740 tweets about these financial institutions, with more than two years of pre-period to observe normal levels of conversation about each banking firm. Tweet-level variables include the text of the tweet, a timestamp, an author identifier, number of retweets, and much more.

In addition to querying tweet level data, we separately query author-level data for each of the unique Twitter users that authored one of the tweets collected. This resulted in collecting author data for 544,888 unique Twitter users. Author-level variables provide information about the authors account including number of followers, date of account creation, and a user provided description.

2.1.1 Processing and scoring of tweets

We score the sentiment of each tweet with Valence Aware Dictionary and sEntiment Reasoner (VADER), a sentiment classifier designed specifically for social media. Hutto and Gilbert (2014) propose VADER as a sentiment classifier for short messages that contain social text, like tweets, and document that it out performs naive dictionary-based methods for scoring tweets. Specifically, we use the VADER algorithm to read through the text of each tweet and compute a raw sentiment score based on the words contained in the tweet. The VADER algorithm is an aggregate of a negative component, a neutral component, and a positive component that load on negative, neutral and positive sentiment tokens, respectively. To compute sentiment, the VADER algorithm aggregates across these sub-components to form a sentiment score for each tweet, which is a number between -1 (very negative) and 1 (very positive). For much of the analysis, we aggregate the number of messages (of different types) to the bank-day level, but for the high-frequency analysis, we specifically use the timestamp.

For the high frequency tweet level tests, we employ the negative and positive components of the VADER scoring of the tweet as separate variables in the analysis. We do this because it is natural to observe an asymmetry in the impact of negative content versus positive content in the context of communication about a bank run, which is mostly coordination around negative information. Moreover, in the context of sentiment analysis in accounting and finance, it is common to include positive and negative keywords separately (Loughran and McDonald, 2011, 2020).

2.1.2 Content dictionaries

We define four content dictionaries of terms used in tweets about banks during this period: "balance sheet", "cryptocurrency", "run behavior", and "contagion." We construct these dictionaries by adapting an iterative method in the spirit of Gentzkow and Shapiro (2011), which is similar to what was applied to a tweet sample by Cookson et al. (2020).

For each topic, we identify a small set of seed words. For example, the balance sheet dictionary includes 'hold-to-maturity' or 'HTM'; For run behavior, 'run' and 'get out'; For Contagion, 'systemic' or 'spillover'; For cryptocurrency, 'crypto' (the full set of seed words is italicized in the contextual dictionaries in Table 1). Using these seed words, we identify the subsample of tweets containing those words, which we use to compute the frequency distribution of words used by topic. We identify the top 40 most *salient* words by topic in comparison to the overall distribution of words using the saliency measure from Goldsmith-Pinkham, Hirtle, and Lucca (2016), and then, we consolidate words into meaningful bigrams that emerged from the salience analysis (e.g., “liquidity” “management” → “liquidity management”), and eliminate terms that appear in multiple dictionaries at the same time (e.g., ’bank’ appeared in all dictionaries). The resulting contextual dictionaries, summarized in the first four columns of Table 1, flag tweets that cover specific topics, which is useful for identifying mechanisms.

In addition to these dictionaries that allow us to flag tweets based on the tweet’s content, we also construct a user-level dictionary by text mining the user description field for the users in our sample. Our aim is to reliably capture whether the user is part of the Twitter “startup community” that was heavily focused in SVB’s uninsured depositor conversation. To do this, we employ a list of terms around the idea of “founder,” “entrepreneur,” and more generally, the VC-backed startup industry. This dictionary is reported in the final column of Table 1. Tweets flagged as “startup community” tweets are defined as those that are authored by a user with at least one of these terms in their user profile description.

2.2 Bank Balance Sheet Information

To construct measures of run exposure for banks in our sample, we closely follow Jiang et al. (2023b) and obtain asset maturity and repricing data for all FDIC-insured banks from Call Reports provided on the Federal Financial Institutions Examination Council (FFIEC) website. These bank balance sheet data are available at the quarterly frequency. We first collect the value of residential MBS (RMBS) and non-residential MBS, loans secured by family residential properties, and all other loans and leases across maturities and repricing breakdowns for Q1 of 2022. To impute changes in the value of loans and securities on the bank balance sheets since 2022:Q1, we use U.S. Treasury Bond ETFs from iShares and S&P Treasury Indices across different maturities m to match the maturity and repricing breakdowns of the bank assets. As shown in Appendix Figure A.1, Treasury Bond Indices and

ETFs with the longest maturity were most strongly affected by recent interest rate increases and incurred the highest losses. Consequently, banks with a high share of long-duration assets incur the largest mark-to-market asset decline.

We estimate changes in the value of banks' assets between 2022:Q1 and 2023:Q1 due to mark-to-market assets as follows:

$$\begin{aligned}\Delta Assets\ MTM = & \sum_{m} (RMBS_m + Mortgages_m) \times \Delta Treasury\ Price_m \times Multiplier \\ & + \sum_{m} \left(Treasuries,\ other\ securities,\ loans_m \times \Delta Treasury\ Price_m \right), \end{aligned} \tag{2}$$

where m indexes the maturity and repricing breakdowns (i.e., 0-3 months, 3-12 months, 1-3 years, 3-5 years, 5-10 years, and 15+ years). $\Delta Treasury\ Price_m$ is the market price change (in %) of Treasury Bonds corresponding to maturity m from 2022:Q1 to 2023:Q1. To account for repayment risk in RMBS and mortgages, we follow Jiang et al. (2023b) and construct the real-estate $Multiplier$ as the ratio of the change in the iShares MBS ETF over the change in the S&P Treasury Bond Index between 2022 and 2023 (Jiang et al. (2023b) provides more detail on the variable construction). We impute a bank's implied mark-to-market (MTM) asset value in 2023:Q1 as the 2022:Q1 value plus $\Delta Assets\ MTM$ from Equation (2) and define the variable "% Asset Decline MTM" as the % change in asset value from 2022:Q1 to 2023:Q1-MTM.

Appendix Figure A.2 shows the distribution of $\Delta Assets\ MTM$ across all FDIC-insured banks (Fig. A.2a) and across publicly listed bank holding companies (BHCs) (Fig. A.2b), which are aggregated across banks if a bank holding company (BHC) has multiple FDIC-insured banks. The range and mean of the distribution are comparable to the numbers documented in Jiang et al. (2023b) for both the full sample of banks and the publicly listed BHCs. Similar to Jiang et al. (2023b), we find that Silicon Valley Bank (SIVB) was not an outlier in terms of % Asset Decline MTM, with a mark-to-market loss in asset value of about 12% due to interest rate increases, which is close to the sample mean. Appendix Figure A.3

displays the distribution % Asset Decline MTM across bank size (i.e., asset value) percentiles. We find that medium to large size institutions (60th to 90th percentile) exhibit the highest 'run exposure' due to asset value declines.

In addition, we obtain information on the total value of bank deposits below (data field RCONF049) and above (data field RCONF051) the FDIC insurance threshold of 250,000 USD from Call Report Schedule RC-O for Q4 of 2022. We use these data to calculate the % of uninsured deposits (i.e., RCONF051/(RCONF049+RCONF051)), which is then multiplied by % Asset Decline MTM to construct our measure of 'run exposure.' In robustness tests, we alternatively use the 'estimate of uninsured deposits' (data field RCON5597) to total deposits (data field RCONF236), and we find similar results.

2.3 Stock Price Data

The stock price data we use throughout the paper are collected from FirstRate Data, a provider of intraday stock trade data. From this data set, we obtain intraday stock price and volume data for 428 banks in intervals of one, five and thirty minutes. For each time interval, the data set includes the price of the first and last trade, the highest and lowest prices, and the trade volume. Prices are adjusted for both splits and dividends.

We use bank stock price data to compute the change in stock prices from March 1st to March 14th as well as hourly stock returns

We also use these data to perform a high-frequency analysis that allows us to identify the impact of bank-related tweets on the stock return by looking at price changes in a narrow time-window around the moment when a tweet is posted during trading hours. More specifically, following the methodology of Bianchi, Cram, and Kung (2021), we use the timestamp (in Eastern Time) of every tweet in our sample to determine two 10-minute windows: one from 15 to 5 minutes *before* the tweet and another one from 5 to 15 minutes *after* the tweet. Then, for each bank, we identify the price of the last trade in the [-15,-5] minute window and the price of the first trade in the [+5,+15] minute window and define log returns over the 5-minute window around each tweet as the difference in the logs of those prices. We exclude the observation from the analysis if either the last price in the window before the tweet or the first price in the window after the tweet have zero associated volume in our dataset. The timeline of the analysis is depicted below:

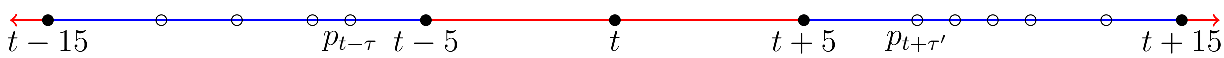

where $p_{i,t-\tau}$ is the last stock price observed in the [-15,-5] minute window, $p_{i,t+\tau}$ is the first stock price observed in the [5,15] minute window, and $p_{it} = \log(P_{it})$ and $\Delta p_{it} = p_{i,t+\tau} - p_{i,t-\tau}$ is the log return. The advantage of this high-frequency approach is that it is unlikely that other value-relevant information about the stock price becomes public during the short interval around the tweet. By estimating the log return using a starting price that was observed *before* the tweet was posted and an ending price after, we further ensure that any findings are unlikely to be driven reverse causality, i.e., the Twitter activity reacting to stock price changes.

2.4 Summary Statistics and Empirical Approach

In our empirical tests, the main period of analysis spans from January 1, 2023 through March 14, 2023. We have comprehensive Twitter data through March 13, and we have high-frequency stock return data from FirstRate through the end of March 14. Although we could feasibly go back to the beginning of 2020 with our tests, we choose to start the pre-run period in January 2023 to maintain comparability of the population of users engaged with Twitter, which helps the measurement in this pre-run period to be more relevant to what happened once the run on SVB began.

Specifically, we define the pre-run period to be January 1 through February 15, which ends well before any discussions begain regarding bank runs in traditional banks. To validate this choice, we compute the word frequency distribution for this 46-day pre-run period, and contrast it with the 5 day run-period (March 8 through March 13) through word clouds in Figure 3. The pre-run period word cloud contains no prominent mentions of bank run terms, such as depositors or withdrawing money, but the run-period is heavily populated with these terms, as well as references to SVB. Reflected in the size of the words in the word cloud, the run period also has a much greater concentration around a few salient terms, whereas the pre-run period is not as concentrated (though it includes many references to crypto markets).

Our empirical tests exploit the wide variation in Twitter conversation intensity during the pre-period. In Panel (a) of Table 2, we present cross-bank distribution of total number of tweets in the pre-run period and run period, rescaled to be the number of tweets about a bank *per 30 days* so that the numbers are comparable to one another, despite the run period having only 5 days in it. The distribution of tweets in the pre-run period has a wide dispersion, leading to substantial variation in the extent to which Twitter users comment about different banks: For example, the 10th percentile bank has only 33 tweets per 30 days while the 90th percentile bank has 511 tweets per 30 days written about it during the pre-run period. This distribution is highly skewed, with a mean of 536 tweets per 30 days, but a median of only 88.7 tweets per 30 days. To capture this wide variation while not giving undue influence to observations in the extreme tail, we divide the distribution into terciles. Panel (b) of Table 2 presents the mean and median number of tweets in the pre-run period and the run period. The top tercile's median number of tweets in the pre-run period is 344 in comparison to a median of 22.5 tweets in the bottom tercile.

The summary statistics in Panels (a) and (b) of Table 2 provide more information about the nature of the run's risks. Specifically, per 30 days, the run period has roughly 4 times the average number of tweets than the pre-period (mean of 2,278 versus 536). However, most of this increase is concentrated among the largest quantiles. The 66th percentile bank in the run period has a comparable tweet rate to the 66th percentile bank in the pre-run period (120 versus 136 per 30 days). However, the 90th percentile and above have much greater Twitter activity during the run period than during the pre-run period.

Once we obtain our analysis sample, we perform a validation exercise in which we identify the top 5 banks by the number of run period tweets in the "run behavior" dictionary. Panel (c) in Table 2 presents counts of run tweets, contagion tweets, pre-run tweets, and crypto tweets pre-run for these top 5 banks, as well as the 90th percentile of these . Consistent with causal empiricism, these tweets identify SIVB as the bank ticker with the most run-based conversation, with 6,528 "run" tweets. First Republic Bank (FRC), which also faced notable troubles, was next in this list with 1,249 run-based tweets. Moreover, all of these top-5 "run" mention banks have an abnormally high number of contagion tweets. In addition, all 5 banks are well above the 90th percentile for number of pre-run tweets (784), and they all have abnormally high cryptocurrency mentions (consistent with the visual evidence in Figure 3).

These summary statistics highlight that the banks that faced the greatest distress during this period were also the ones that had the most Twitter attention in the pre-period. In the next section, we systematically test this idea while relating the potential influence of social media to factors that have been identified to lead to bank distress, such as the percentage of uninsured deposits and losses from marked-to-market assets.

3 Results

In this section, we present our main results that link social media conversations to bank runs in early 2023. Our empirical strategy is to focus on bank stock returns for publicly traded banks as a proxy for the severity of the run period. We present three sets of results: (1) tests at the hourly frequency that link social media activity in a recent time window (e.g., 4 hours, 8 hours, 12 hours) to hourly returns, (2) purely cross-sectional tests that explain the total bank stock loss in the run period (Mar 1 through Mar 14) to *preexisting* social media exposure, measured distinctly before the run period, and (3) high frequency tests in the spirit of Bianchi et al. (2023) that examine the contemporaneous impact of social media activity and sentiment on bank stocks in 5-minute windows around the tweet.

3.1 Event Analysis around the Onset of the SVB Run

We employ bank stock returns as a proxy for the severity of the bank run risk that emerged during this period. This approach is second best to studying high frequency deposit flows, but such data on deposits are not easily available.

We validate the use of returns by estimating how returns for banks with high "run exposure" versus banks with low "run exposure" change around the onset of the run period. To understand the timing of the onset of the run period, consider Figure 1, which shows that the first substantial uptick of tweets about SVB (or SIVB) occurred in the morning of March 9. Thus, we take the onset of the bank-run period to be March 9. Aside from helping to identify the timing of the run on SVB, this figure also provides evidence on the social transmission of ideas: the run discussion originated in investment tweets (i.e., those that use the stock ticker SIVB), and then later, spilled over into discussions of the acronym (SVB), and later the name of the bank (Silicon Valley Bank).

To measure "run exposure," we follow Jiang et al. (2023b)'s procedure and construct the percentage loss due to mark-to-market assets from price changes in bond ETFs as detailed in Section 2. In addition, we compute the percentage of uninsured deposits relative to total deposits from the 2022:Q4 FDIC Call Reports. Appealing to arguments in Jiang et al. (2023b), the interaction between these two banking characteristics is our measure of "run exposure." The intuition for this measure is simple: a higher % drop in assets due implies that the bank holds fewer to cover depositor claims. However, if deposits are fully insured depositors have little incentive to withdraw. Hence, the combination of high mark-to-market asset value decline with a low deposit insurance ratio leaves a bank potentially vulnerable to a run, as explained in Jiang et al. (2023b).

Table 3a provides evidence consistent with this intuition. We contrast the top tercile in "run exposure", with the bottom two terciles, and separately compute the average hourly stock returns for March 1-8 (before the run on SVB) and March 9-14 (the run period). In line with the idea that high 'run exposure' banks were more vulnerable after the onset of the SVB episode, the hourly returns on bank stocks in the top tercile of run-exposed stocks are about 0.2 percentage points lower, with 95% confidence intervals that do not overlap.

We refine this unconditional test by estimating the following triple difference model using a bank-hour panel and hourly return data from FirstRate, which allows us to control for firm and time fixed effects:

$$\begin{aligned} r_{i,t} = a &+ b_1 \times 1(\geq \text{Mar } 09)_t + b_2 \times \% \text{ Asset} \downarrow \text{MTM}_i + b_3 \times \% \text{ Uninsured}_i \\ &+ b_4 \times 1(\geq \text{Mar } 09)_t \times \% \text{ Asset} \downarrow \text{MTM}_i + b_5 \times 1(\geq \text{Mar } 09)_t \times \% \text{ Uninsured}_i \\ &+ b_6 \times \% \text{ Asset} \downarrow \text{MTM}_i \times \% \text{ Uninsured}_{i,t-1} \\ &+ b_7 \times 1(\geq \text{Mar } 09)_t \times \% \text{ Asset} \downarrow \text{MTM}_i \times \% \text{ Uninsured}_i + \delta_i + \gamma_t + \epsilon_{i,t}, \end{aligned} \tag{3}$$

where $r_{i,t}$ is the hourly return for bank stock i during trading hour t (in %), $\%$ Asset $\downarrow$ MTM$_i$ is the percentage loss of bank i's assets due to mark-to-market securities, $\%$ Uninsured$_i$ is the percentage of deposits below the FDIC insurance threshold relative to total deposits, and $1(\geq \text{Mar } 09)_t$ is an indicator that equals 1 for all hours t after the onset of the SVB run at the beginning of March 9th (9am). The triple interaction model includes all lower order terms, as well as bank (δ_i) and day-by-hour (γ_t) fixed effects in some specifications. Standard errors are clustered at the bank level. The coefficient of interest is b_7, which reflects the degree to which banks with high run exposure – captured by the interaction $\%$ Decline MTM$_i \times \%$ Uninsured$_i$ – experience greater stock market losses in the run period (i.e., when $1(\geq \text{Mar } 09)_t$ is equal to one) than before the run period. To facilitate interpretation in the regression tables, we standardize continuous RHS variables to have zero mean and standard deviation of one. This standardization does not change the statistical tests, but scales the tabulated magnitude to be for a one-standard deviation increase. We indicate that a variable is standardized in a regression table by writing '(z)' next to the variable name.

Table 3b presents the results, which robustly show that $\beta_7 < 0$. In column 1, we estimate a triple interaction of -0.0412, which implies that a bank with high run exposure (one SD above the mean for both '% Asset $\downarrow$ MTM (z)' and '%Uninsured (z)') is expected to lose an additional 4.12 basis points of return per hour more during the run period than before it. The two-way interactions of '1($\geq$ Mar 09)' with '% Asset $\downarrow$ MTM (z)' and '%Uninsured (z)' provide a natural benchmark for the economic magnitude of this estimate. Both coefficients on '1($\geq$ Mar 09)' $\times$ '% Asset $\downarrow$ MTM (z)' and '1($\geq$ Mar 09)' $\times$ '%Uninsured (z)' are strongly negatively related to average returns (between 6.2 and 10.7 basis points per hour, respectively). Aside from these estimates being highly significant, they are also stable upon the inclusion of bank and day-by-hour fixed effects, even though these fixed effects explain substantial variation (R^2 goes from 0.0098 to 0.2576 upon their inclusion).

To provide evidence that this difference in returns, indeed, emerges precisely at the onset of the social media discussion about SVB, we plot the cumulative returns from March 06 through March 14 for banks with *ex ante* high versus low run exposure (i.e., '% Asset $\downarrow$ MTM (z)' $\times$ '%Uninsured (z)') in Figure 4. In the pre-run period up until the end of March 09, there are virtually no differences in the returns for high versus low "run exposure" banks. However, following the onset of the run, banks with both a high asset value decline due to MTM *and* low deposit insurance ratios experience sharper declines in returns. Not only are these results consistent with the message in Jiang et al. (2023b) that bank run risk is tied to uninsured deposits, but this evidence shows that our identification of the timing of the onset of the run via Twitter activity is valid and that our use of returns to understand the extent of bank run risk is reasonable.

3.2 The Role of Preexisting Exposure to Twitter Activity

Building on the finding that return movements of bank stocks during the run period reflect exposure to bank run risks, we next conduct a series of tests that examine how *preexisting* exposure to Twitter interacts with mark-to-market bank losses and the percentage of uninsured deposits.

To do this, we measure of exposure to social media conversation by counting the number of tweets that mention a bank's cashtag before the run on SVB. We define this pre-run period as January 1, 2023 through February 15, 2023, which occurs well before any run-specific discussion began on Twitter. Consistent with this choice of pre-period, Figure 3 presents word clouds that emphasize the most frequently used words in this pre-run period versus the run-period. The run period's conversation is concentrated in discussion about the implications of the SVB run, whereas the pre-run period has a wider range of topics, including many cryptocurrency terms and hashtags. Our goal is to construct a measure of intensity of Twitter conversations in the absence of bank run discussions, and thus, attribute any estimated effects to the bank's *exposure to social media conversation* rather than any information contained in social media about how run prone is one bank versus another.

We estimate a cross-sectional specification that interacts balance sheet characteristics that make bank i more or less run prone with preexisting exposure to Twitter:

$$Stock\_Loss_i = \beta_1 Loss_i \times Uninsured_i \times SocialExp_i + lower\ order\ terms + \epsilon_i \tag{4}$$

where $Stock\_Loss_i$ is the percentage of bank stock market value that was lost from March 1st until March 14th. We include tercile of tweet activity indicators in our main tests, instead of the continuous variable $SocialExp_i$, because the distribution of tweets is right skewed. The main coefficients of interest are those on the top tercile indicator (on its own and interacted with $Loss_i \times Uninsured_i$). The coefficients on the lower order terms in the specification are also of interest. For example, being largely reliant on uninsured deposits can be risky irrespective of mark-to-market losses. Thus, we are interested in the main effect on $Uninsured$, but also its interactions with $Loss$ and $SocialExp$. These coefficients provide a natural benchmark for magnitudes as well. In fact, the relative importance of these terms in comparison to one another speaks to how important exposure to social media is versus well-studied balance sheet characteristics (e.g., Egan et al., 2017; Jiang et al., 2020, 2023b).

Table 4 presents the main results. Column 1 presents a specification without the interaction with social media exposure. Broadly, this specification confirms the emphasis of Jiang et al. (2023b) on the role of uninsured deposits. Interpreting the main effect, a standard deviation increase in the percentage of uninsured deposits is associated with a 4.1 percentage point decline in the bank's stock during the SVB bank run. The relation of marked-to-market bank losses to bank stock losses is not statistically significant, nor is the interaction with uninsured deposits. However, the magnitudes of these estimated coefficients are economically meaningful and in the direction hypothesized: A standard deviation increase in marked-to-market losses is associated with 0.8 percentage points more stock market losses during the run, and for a bank with a high fraction of uninsured deposits (one standard deviation above the mean), this magnitude more than doubles due to the interaction coefficient estimate of 0.943.

In column 2 of Table 4, we present evidence on the link between high (or medium) preexisting social media exposure and the percentage bank stock loss during the SVB run. This social media channel predicts bank stock market losses during the run and the estimated magnitude is similar to the coefficient on percent uninsured deposits. A bank with top tercile preexisting social media exposure has, on average, 6.66 percentage points greater loss during the run period than a bank in the bottom tercile of social media exposure. By contrast to top tercile social media exposure, we see no significant difference between banks in the middle tercile of social media exposure and banks in the bottom tercile of exposure.

Communication via social media ought to affect bank run risk through amplifying existing risks. To test this, in columns 3 through 5, we estimate specifications that interact $SocialExp$ and bank run characteristics. Consistent with social media communication amplifying these bank balance sheet risks, we see that the inclusion of these interactive terms leads the main effects on % Loss and % Uninsured to become smaller and insignificant. Moreover, we eatimate a significant positive interaction between *top tercile social media exposure* and % Uninsured. These results emphasize the importance of communication via social media: % Uninsured deposits matters for bank run risk, but it relates to bank run risk only insofar as it coincides with significant social media exposure, which facilitates coordination among uninsured depositors. Moreover, we estimate a positive and highly significant triple interaction, which implies that a bank with high "run exposure" (one sd above the mean for both MTM bank losses and uninsured deposits) can expect 3.014 percentage points lower return during the run period. These findings highlight the importance of social media communication as an *ex ante* risk factor that leads to greater bank run risk via heightened communication.

3.2.1 Tweets by Startup Community Users

An important feature of the SVB bank run is that SVB's uninsured depositors were highly concentrated in one industry: VC-backed startups who were encouraged to deposit large sums of money with the bank. In addition to the usual risks of being exposed to a single industry's deposits, this particular startup community is also highly connected through social media. Thus, it is natural that the social media communication risks faced by SVB were unusually high.

We examine the Twitter startup community's influence on the run by flagging users who employ a set of startup terms (e.g., "startup," or "founder") in their user description field. This field allows us to follow the discussion of startup community users over time and to evaluate how cross-sectional exposure to this highly connected group of users influences the observed bank run risks.

Using preexisting Twitter conversation by startup community users instead of total tweets in the pre-run period, we re-estimate (4). Table 5 reports the estimated coefficients using preexisting startup community tweet activity. The results yield a slightly weaker triple interaction coefficient than using overall tweets. However, the two-way interaction between top tercile startup tweets and % Uninsured remains highly significant, and the main effect terms on % Uninsured, % Loss, and their interaction are all small and non-significant in the fully interacted model. Thus, even using a subset of the Twitter activity in the pre-run period, we conclude that exposure plays an important role in amplifying bank run risk due to banking characteristics known to be related to bank runs, particularly percent of uninsured deposits.

3.2.2 Relation to Tweets about Running and Contagion During the Run

Next, we investigate how preexisting exposures of banks to Twitter activity could lead to heightened attention and discussion of bank run topics during the run period. For this analysis, we estimate a cross-sectional test with the bank stock market loss as the dependent variable, but enriching the specification in equation (4) by controlling at the bank level for Twitter activity *during* the run period on run-specific topics. Specifically, we draw upon the content dictionaries we developed to identify tweets that mention running or contagion. We also investigate the role of startup community tweets during the run period.

Table 6 presents the results from estimating these empirical specifications. Column 1 presents the regression without any controls for run-period social media activity. Consistent with our main specification, top tercile exposure to social media predicts a 6.66 percentage point loss during the bank run. However, upon controlling individually for Contagion Tweets, Run Tweets, or Startup Tweets during the run period, individually (columns 2 through 4) or all together (column 5), the coefficient on preexisting social media exposure diminishes in both economic and statistical significance. This finding suggests that preexisting exposure to social media affects bank run risk because it leads to more run and contagion tweets, as well as more Twitter activity by members of the startup community on Twitter. These tweets served to coordinate and communicate among depositors during the run.

One concern in these cross-sectional tests in that the tweets during the run period – at this level of aggregation – need not precede the return reactions. Though this evidence links preexisting exposure to run period tweet activity, a cross-sectional test is ill equipped to attribute run period activity to return reactions rather than the other way around. To provide finer-grained insight into the timing of run period tweets, we next turn to examine the impact of tweet activity in a bank-hourly panel setting.

3.3 Hourly Bank Stock Returns and Twitter Activity during the Run

In this section, we employ a high frequency specification similar to the hourly return specification in Equation (3). Specifically, we estimate the following specification on our data set of returns and Twitter activity at the hourly frequency:

$$\begin{aligned} r_{i,t} = a &+ b_1 \times 1(\geq \text{Mar } 09)_t + b_2 \times \text{Run Exposure}_i + b_3 \times \text{N Tweets}_{i,t-1} \\ &+ b_4 \times 1(\geq \text{Mar } 09)_t \times \text{Run Exposure}_i + b_5 \times 1(\geq \text{Mar } 09)_t \times \text{N Tweets}_{i,t-1} \\ &+ b_6 \times \text{Run Exposure}_i \times \text{N Tweets}_{i,t-1} \\ &+ b_7 \times 1(\geq \text{Mar } 09)_t \times \text{Run Exposure}_i \times \text{N Tweets}_{i,t-1} + \delta_i + \gamma_t + \epsilon_{i,t}, \end{aligned} \tag{5}$$

where $r_{i,t}$ is the hour t return for bank i measured in percentage points, the indicator '$1(\geq \text{Mar } 09)_t$' captures hours after the onset of the run on March 09th. 'Run Exposure$_i$' is defined as the interaction between mark-to-market bank losses and percentage uninsured deposits of bank i, 'N Tweets$_{i,t-1}$' is a count of the number of tweets about bank i over the past 4 *hours*, lagged by one hour. This specification contrasts hourly return reactions related to different intensities of Twitter activity for highly run-exposed banks during the run period. By observing this high frequency (4 hour) lag of Twitter conversation intensity ('N Tweets$_{i,t-1}$') about bank i, we alleviate concerns that stock returns may drive twitter activity rather than the reverse. Our main sample of bank-hour observations is drawn from March 6th until March 14th (March 11 and 12 were weekend days, i.e., non-trading days).

Table 7a presents the results from estimating this specification. The number of Tweets in the past four hours exhibits a significant triple interaction with Run Exposure during the run period, particularly after controlling for firm fixed effects and in the specification with day-by-hour fixed effects. That is, we can attribute significantly lower hourly returns of run-exposed stocks to a higher *recent* Twitter conversation intensity about the bank. Our estimates are consistent with a social contagion interpretation, suggesting that bank returns are following Twitter conversation for run-exposed firms during the "run" period.

We perform two robustness exercises on this hourly interaction result. First, in Table 7b, we report the results if we drop the ticker for Silicon Valley Bank (SIVB) from the sample. The magnitude of the triple interaction remains highly significant, while being slightly smaller. Second, we consider a shorter time window around the onset of the bank run period: one day before (March 8) through the first full day (March 9). This sample of days leaves us with approximately one third the sample, but we find larger coefficient estimates for b_7 across all specifications, indicating that much of our result is driven by the period around the onset of the SVB run.

To examine the dynamics of this finding throughout the event period and rule out pre-trends, we plot the cumulative returns of banks with high and low Twitter activity in Figure 5. The sample underlying this figure includes only firms with 'high' run-exposure (i.e., above the median), to isolate the effect of Twitter activity in the pre- and post run period. Consistent with the emergence of the run driving the interaction between run exposure and social media activity, we see differences in cumulative returns diverge meaningfully for the first time at the onset of the run period (March 9th). Moreover, the results are not merely a reflection of SIVB's returns because we see a very similar pattern if we drop SIVB from this analysis (see Figure 5b). These results show that social media played a role, not only in the bank run on SVB, but the banking instability that affected a broader set of banks.

3.4 High frequency tests of Market Impacts of Tweets

In this last leg of our empirical analysis, we analyze price changes in the stocks of banks over very short time intervals around tweets that make reference to banks. More specifically, we study changes in log prices of bank stocks in 10-minute windows around the time of the bank-related tweets in our sample constructed as detailed in Section 2.3. The high frequency analysis allows for clean identification of the impact of *contemporaneous* conversations on social media on stock prices. The key identifying assumption is that no other events take place in such narrow time windows that could confound the observed effects. Our analysis follows Bianchi et al. (2021) who use a similar approach to study the effect of legislators' tweets on the stock prices of targeted firms. A similar methodology has also been used by Bianchi et al. (2023) to study the impact of presidential tweets on Fed funds futures.

We define the log return of bank i's stock over the 10-minute window around a tweet taking place at time t and containing the bank's cashtag as $\Delta p_{it} = p_{i,t+\tau} - p_{i,t-\tau}$, where $p_{i,t+\tau}$ and $p_{i,t-\tau}$ denote the log price of the last trade in the $[t-15', t-5']$-window and the log price of the first trade in the $[t+5', t+15']$-window, respectively. An advantage of this design is that we observe the starting price before the tweet is posted and the ending price after the tweet is posted when constructing returns. This alleviates concerns that stock prices drive Twitter activity, rather than the reverse. In our main high-frequency specification, we then regress log returns of bank i on the tone of the tweet:

$$\Delta p_{i,t} = a + b \times \text{VADER Pos}_{i,t} + c \times \text{VADER Neg}_{i,t} + \gamma_i + \epsilon_{i,t}, \tag{6}$$

where $\Delta p_{i,t}$ is winsorized at the 1% level and expressed in basis points, VADER $\text{Pos}_{i,t}$ and VADER $\text{Neg}_{i,t}$ are the positive and negative components of the sentiment score assigned by VADER to the tweet, and γ_i denotes bank fixed effects. To ease comparability comparability, both negative and positive components have been standardized to have a mean of zero and a standard deviation of one in the full sample. In all regressions, we cluster standard errors at the bank-day level.

Table 8 reports the results from estimating equation (6). As shown in column 1, we estimate a strong negative and immediate price response to the negative component of the VADER sentiment score (VADER $\text{Neg}_{i,t}$). The estimate indicates a 1.60 basis point lower return for a one-standard deviation increase in the negative sentiment component of the tweet's VADER score. Importantly, we do not find a similar effect for positive tweet sentiment, the coefficient estimate is economically small and indistinguishable from zero despite the very large sample size ($N = 1,521,078$) and relatively small standard error ($SE = 0.16$).

Next, we investigate the role of the startup community as a driver of this effect. To this end, we augment the specification in Equation (6) by interacting both VADER $\text{Pos}_{i,t}$ and VADER $\text{Neg}_{i,t}$ with the dummy variable '*Startup Flag*', which takes the value of one if the tweet author is in our 'Startup community' dictionary, depicted in Table 1. Silicon Valley Bank's deposits were highly concentrated among members of this startup community, and to the point of our paper, members of the startup community (VC investors, founders, etc.) are highly networked and active on Twitter. Consistent with the idea that members of the startup community were particularly important for driving the Twitter conversation and its bank run risks around the collapse of SVB, we estimate a significantly negative interaction effect of -2.13 for VADER $\text{Neg}_{i,t} \times 1(\text{Startup Community})$. The magnitude of this estimate implies that the negative high frequency effect of negative tweet sentiment on bank stock returns more than doubles when the tweet came from a member of the startup community.

In the remaining tests in Table 8, we restrict the sample period to after the onset of the SVB run on March 09. By comparing estimates obtained from the full sample period (Columns 1 and 2) to the estimates from this post-run sample, we can infer how the onset of the SVB run interacts with tweet sentiment and Twitter user identity (i.e., startup community) and tweet content. Consistent with the hourly-frequency estimates in Table 7, we find a stronger effect of negative tweets on returns and a much stronger effect of tweets by members of the startup community in the post-run period: the coefficient estimate for VADER $\text{Neg}_{i,t} \times 1(\text{Startup Community})$ jumps to -21.82 (significant at the 1% level). In line with this finding, we also document a stronger effect of negative tweet sentiment on returns when the Tweet is part of our 'contagion' dictionary (see Section 2.1.2). Last, similar to Table 7, we find much more pronounced unconditional negative returns in the post-run period, i.e., the constant term in columns (3) through (6) is negative, large in magnitude, and highly significant. The negative constant term shows that even neutral tweets have negative impact during the run period.

Interestingly, when we consider interactions for whether the tweet contained terms in the "run behavior" dictionary or tweets about banks with high run exposure − i.e., those with a large fraction of mark-to-market losses *and* a large percentage of uninsured deposits − we do not observe a significant interaction between tweet sentiment and these characteristics. In light of our findings in the hourly tests, there are at least two possibilities for these non-results. First, run behavior and high exposure banks could be characteristics for which the price reaction is not immediate because it takes time for the market to impound this information into prices. Since the time horizon extends only up to 15 minutes after the tweet, this may not allow enough time for the full effect to emerge. A second possibility is that these characteristics and their impact on price could be relatively independent of the sentiment expressed in such tweets. For example, once a Twitter user utters the word "run," it may not be as relevant to the price reaction that the tweet used moderately negative tone or extremely negative tone. In this case, the volume of "run behavior" tweets is likely to be much more informative. We do not distinguish these hypotheses, as the required tests to do so are beyond the scope of this analysis.

Table 9 presents tests analogous to Table 8, excluding Silicon Valley Bank (SIVB) and First Republic Bank (FRC) from the sample. This data filter reduces our sample size noticeably: the post-run sample size declines from 43,597 in Table 9 to 19,673 in Table 8, which is unsurprising given that a sizable volume of tweets posted during the run episode included the tickers of SIVB and FRC. While the coefficient estimates are smaller in magnitude compared to Table 8, the main takeaways remain unchanged. We continue to find that the negative tweet sentiment component has a significant negative effect on 10-min bank stock returns, and this effect is much stronger when the tweet was posted by a member of the startup community.

Overall, the tests in this section provide strong support for Twitter's role in amplifying bank distress during the run period. In addition to the tests in previous sections, the use of VADER negative sentiment reinforces the view that the content of the social media conversation had much to do with the banking distress that surrounded the collapse of SVB. Moreover, the relation of bank stock returns strengthens in its connection to both this negative sentiment and the role of tweets by startup community members. These findings reinforce evidence for our core mechanism that bank distress following SVB was amplified by social media conversations.

4 Conclusion

This paper presents comprehensive evidence that exposure to social media conversation about bank stocks amplifies classical bank run risks. Our empirical tests show that banks with a large *preexisting* exposure to social media performed much worse during the recent SVB bank run, particularly if they have large mark-to-market losses and a large percentage of uninsured deposits. At the hourly frequency, we show that negative returns emerge after periods of intense Twitter conversation, but this effect only emerges after the run on SVB begins. When we perform a high-frequency analysis as in Bianchi et al. (2023), the immediate price reactions reinforce many of these core lessons. The risk to banks increases markedly when bank firms are in a Twitter conversation.

The effect that we uncover − though a modern phenomenon linked to rapid communication over social media − is one that is classically rooted in bank run models, going back to Diamond and Dybvig (1983). These models have always posited that communication and coordination pose a risk to banks, especially when many of the deposits in the bank are uninsured. Increasingly, in today's society, social media provides a means for individuals to coordinate and communicate beyond what older technologies allow. The amplification of bank run risk via Twitter conversations is a unique opportunity to observe communication and coordination that shapes a critically important economic outcome − distress in banks.

Given the increasingly pervasive nature of social communication on and off Twitter, we do not expect this risk to go away, but rather, it is likely to influence other outcomes, as well.

References

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2015). Systemic risk and stability in financial networks. *American Economic Review 105*(2), 564–608.

Akçay, E. and D. Hirshleifer (2021). Social finance as cultural evolution, transmission bias, and market dynamics. *Proceedings of the National Academy of Sciences 118*(26), e2015568118.

Angeletos, G.-M. and I. Werning (2006). Crises and prices: Information aggregation, multiplicity, and volatility. *American Economic Review 96*(5), 1720–1736.

Bianchi, F., R. G. Cram, and H. Kung (2021). Using social media to identify the effects of congressional viewpoints on asset prices. NBER Working Paper.

Bianchi, F., R. Gómez-Cram, T. Kind, and H. Kung (2023). Threats to central bank independence: High-frequency identification with twitter. *Journal of Monetary Economics 135*, 37–54.

Brown, J. R., J. A. Cookson, and R. Z. Heimer (2019). Growing up without finance. *Journal of Financial Economics 134*(3), 591–616.

Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007–2008. *Journal of Economic Perspectives 23*(1), 77–100.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics 130*(3), 453–483.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2022). Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy. *Journal of Political Economy forthcoming*.

Calomiris, C. W. and J. R. Mason (1997). Contagion and bank failures during the great depression: The june 1932 chicago banking panic. *THE AMERICAN ECONOMIC REVIEW*.

Chen, H., P. De, Y. J. Hu, and B.-H. Hwang (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *The Review of Financial Studies 27*(5), 1367–1403.

Chetty, R. et al. (2022a). Social capital i: measurement and associations with economic mobility. *Nature 608*(7921), 108–121.

Chetty, R. et al. (2022b). Social capital ii: determinants of economic connectedness. *Nature 608*(7921), 122–134.

Chwe, M. S.-Y. (2000). Communication and coordination in social networks. *The Review of Economic Studies 67*(1), 1–16.

Cookson, J. A., J. E. Engelberg, and W. Mullins (2020). Does partisanship shape investor beliefs? evidence from the covid-19 pandemic. *The Review of Asset Pricing Studies 10*(4), 863–893.

Cookson, J. A., J. E. Engelberg, and W. Mullins (2022, 08). Echo Chambers. *The Review of Financial Studies*. hhac058.

Cookson, J. A., R. Lu, W. Mullins, and M. Niessner (2022). The social signal. Available at SSRN.

Cookson, J. A. and M. Niessner (2020). Why don't we agree? Evidence from a social network of investors. *The Journal of Finance 75*(1), 173–228.

Cookson, J. A., M. Niessner, and C. Schiller (2022). Can social media inform corporate decisions? evidence from merger withdrawals. Working Paper.

Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy 91*(3), 401–419.

Dupont, Q. (2022). Households' financial fallout from institutional scandals: Evidence from the u.s. catholic clergy abuse crisis. SSRN Working Paper 3875648.

Egan, M., A. Hortaçsu, and G. Matvos (2017). Deposit competition and financial fragility: Evidence from the us banking sector. *American Economic Review 107*(1), 169–216.

Elliott, M., B. Golub, and M. O. Jackson (2014). Financial networks and contagion. *American Economic Review 104*(10), 3115–3153.

Falato, A., A. Hortacsu, D. Li, and C. Shin (2021). Fire-sale spillovers in debt markets. *The Journal of Finance 76*(6), 3055–3102.

Gentzkow, M. and J. M. Shapiro (2011). Ideological segregation online and offline. *The Quarterly Journal of Economics 126*(4), 1799–1839.

Giannini, R., P. Irvine, and T. Shu (2019). The convergence and divergence of investors' opinions around earnings news: Evidence from a social network. *Journal of Financial Markets 42*, 94–120.

Gil-Bazo, J. and J. F. Imbet (2022). Tweeting for money: Social media and mutual fund flows. *Université Paris-Dauphine Research Paper* (3719169).

Goldsmith-Pinkham, P. S., B. Hirtle, and D. O. Lucca (2016). Parsing the content of bank supervision. FRB of NY Staff Report.

Golub, B. and M. O. Jackson (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics 2*(1), 112–149.

Gurun, U. G., N. Stoffman, and S. E. Yonker (2018). Trust busting: The effect of fraud on investor behavior. *The Review of Financial Studies 31*(4), 1341–1376.

Han, B., D. Hirshleifer, and J. Walden (2022). Social transmission bias and investor behavior. *Journal of Financial and Quantitative Analysis 57*(1), 390–412.

Hirshleifer, D. (2020). Presidential address: Social transmission bias in economics and finance. *The Journal of Finance 75*(4), 1779–1831.

Hirshleifer, D. and J. B. Plotkin (2021). Moonshots, investment booms, and selection bias in the transmission of cultural traits. *Proceedings of the National Academy of Sciences 118*(26), e2015571118.

Hutto, C. and E. Gilbert (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, Volume 8, pp. 216–225.

Iyer, R. and M. Puri (2012). Understanding bank runs: The importance of depositor-bank relationships and networks. *American Economic Review 102*(4), 1414–1445.

Jiang, E., G. Matvos, T. Piskorski, and A. Seru (2020). Banking without deposits: Evidence from shadow bank call reports. NBER Working Paper.

Jiang, E. X., G. Matvos, T. Piskorski, and A. Seru (2023a). Limited hedging and gambling for resurrection by u.s. banks during the 2022 monetary tightening? Working Paper, SSRN Number 4410201.

Jiang, E. X., G. Matvos, T. Piskorski, and A. Seru (2023b). Monetary tightening and us bank fragility in 2023: Mark-to-market losses and uninsured depositor runs? NBER Working Paper.

Jiang, E. X., G. Matvos, T. Piskorski, and A. Seru (2023c). U.s. bank fragility to credit risk in 2023: Monetary tightening and commercial real estate distress. Working Paper, SSRN Number 4413799.

Kearney, M. S. and P. B. Levine (2019). Early childhood education by television: Lessons from sesame street. *American Economic Journal: Applied Economics 11*(1), 318–350.

Kelly, M. and C. Ó. Gráda (2000). Market contagion: Evidence from the panics of 1854 and 1857. *American Economic Review 90*(5), 1110–1124.

Kuchler, T. and J. Stroebel (2021). Social finance. *Annual Review of Financial Economics 13*, 37–55.

Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review 111*(3), 831–870.

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

Loughran, T. and B. McDonald (2020). Textual analysis in finance. *Annual Review of Financial Economics 12*, 357–375.

Müller, K. and C. Schwarz (2021). Fanning the flames of hate: Social media and hate crime. *Journal of the European Economic Association 19*(4), 2131–2167.

Müller, K. and C. Schwarz (2022). From hashtag to hate crime: Twitter and anti-minority sentiment. *American Economic Journal, Applied Economics*.

Parmelee, J. H. and S. L. Bichard (2011). *Politics and the Twitter revolution: How tweets influence the relationship between political leaders and the public*. Lexington books.

Peck, J. and K. Shell (2003). Equilibrium bank runs. *Journal of Political Economy 111*(1), 103–123.

Pedersen, L. H. (2022). Game on: Social networks and markets. *Journal of Financial Economics 146*(3), 1097–1119.

Renault, T. (2017). Intraday online investor sentiment and return patterns in the us stock market. *Journal of Banking & Finance 84*, 25–40.

Stein, L. C. and C. Yannelis (2020). Financial inclusion, human capital, and wealth accumulation: Evidence from the freedman's savings bank. *The Review of Financial Studies 33*(11), 5333–5377.

Strömberg, D. (2004). Radio's impact on public spending. *The Quarterly Journal of Economics 119*(1), 189–221.

Traweek, V. and M. Wardlaw (2022). Societal inequality and financial market participation: Evidence from the freedman's savings bank.

Van Bommel, J. (2003). Rumors. *The Journal of Finance 58*(4), 1499–1520.

Ziebarth, N. L. (2013). The radio and bank distress in the great depression. *Manuscript, University of Iowa.*

Figure 1: Silicon Valley Bank Stock Price and Twitter Mentions

Investor tweets (ticker = SIVB) lead general discussion tweets (SVB and Silicon Valley Bank) within the run period.

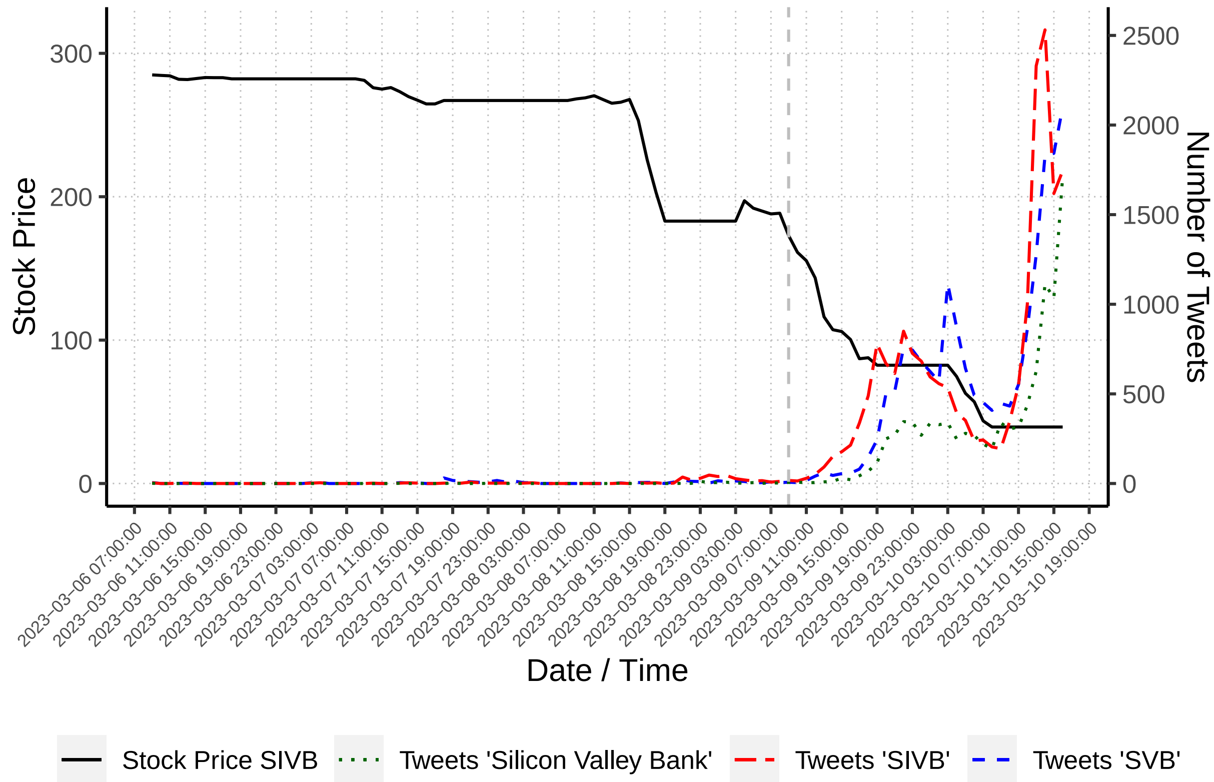

Notes: This figure shows the stock price of Silicon Valley Bank (SIVB) (black solid line), as well as the number of tweets mentioning "SIVB" (blue dashed line), "SVB" (red long-dashed line), and "Silicon Valley Bank" (green dotted line), over the period from March 06, 2023 at 9am (market open) to March 10, 2023 at 4pm (market close). The stock price is indicated on the left y-axis and the number of Tweets is indicated on the right y-axis. The grey vertical dashed line indicates March 09, 2023 at 9am.

Figure 2: Startup Community Tweets about Silicon Valley Bank

Twitter Startup Community users post mostly general discussion tweets, which start distinctly after the initial wave of tweets.

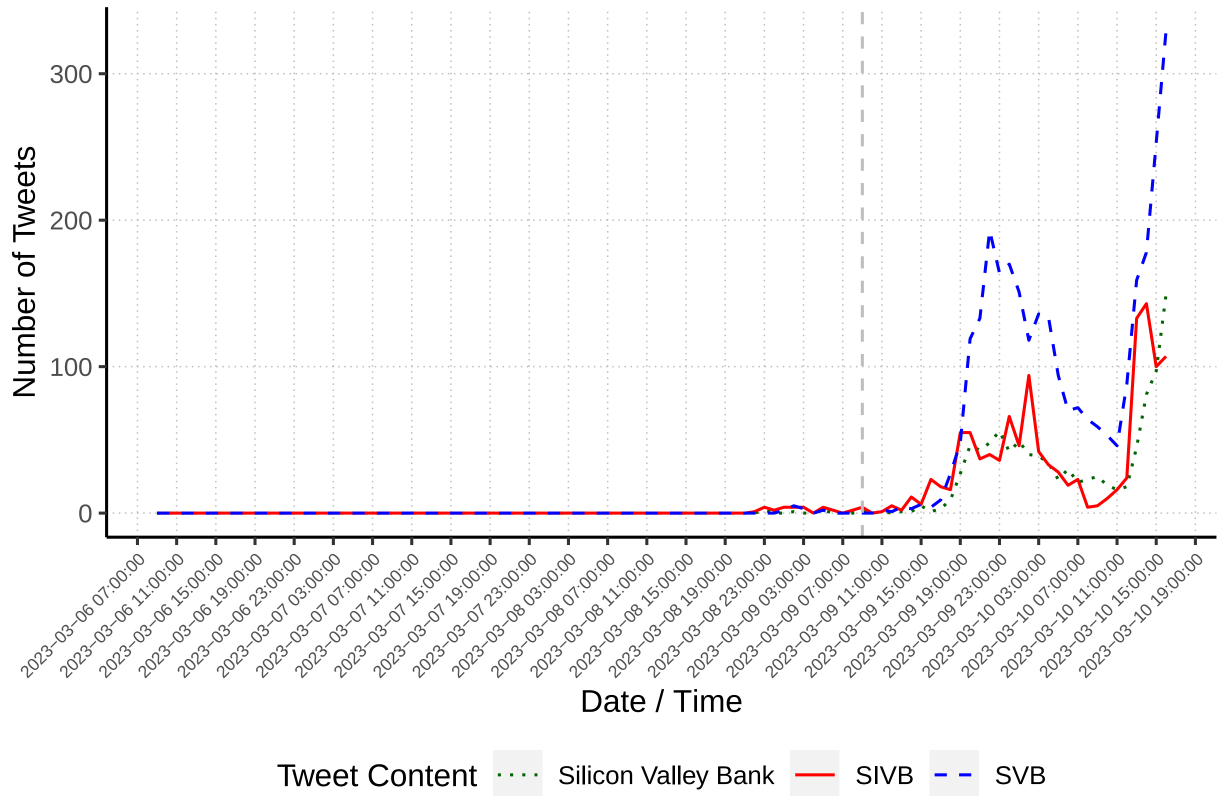

Notes: This figure shows the number of tweets mentioning “SIVB” (blue solid line), “SVB” (red dashed line), and “Silicon Valley Bank” (green dotted line), over the period from March 06, 2023 at 9am (market open) to March 10, 2023 at 4pm (market close). The sample includes only tweets sent by users who are part of the startup community as defined by corresponding dictionary of terms in the Twitter user biography. The grey vertical dashed line indicates March 09, 2023 at 9am.

Figure 3: Content of Tweets about Bank Stocks – pre-run period and run period

Tweets in the pre-run period are unrelated to bank run risks while tweets in the post-run period are.

(a) Pre-run period (Jan 1 through Feb 15)

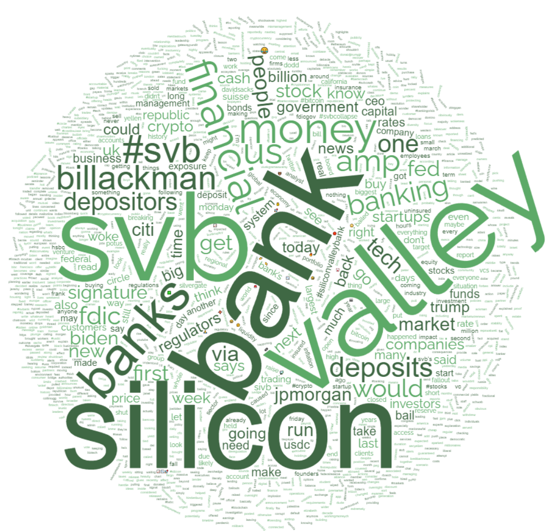

(b) Run period (Mar 8 through Mar 13)

Notes: This figure presents word clouds, depicting the most commonly used words in the pre-run period from January 1 through February 15 (Panel a) and the run period from March 8 through March 13 (Panel b). The size of the words in the word cloud reflects their relative frequency in comparison to other words.

Figure 4: Run Exposure and Cumulative Returns

Banks with more uninsured deposits and mark-to-market losses experience greater stock losses in the run period.

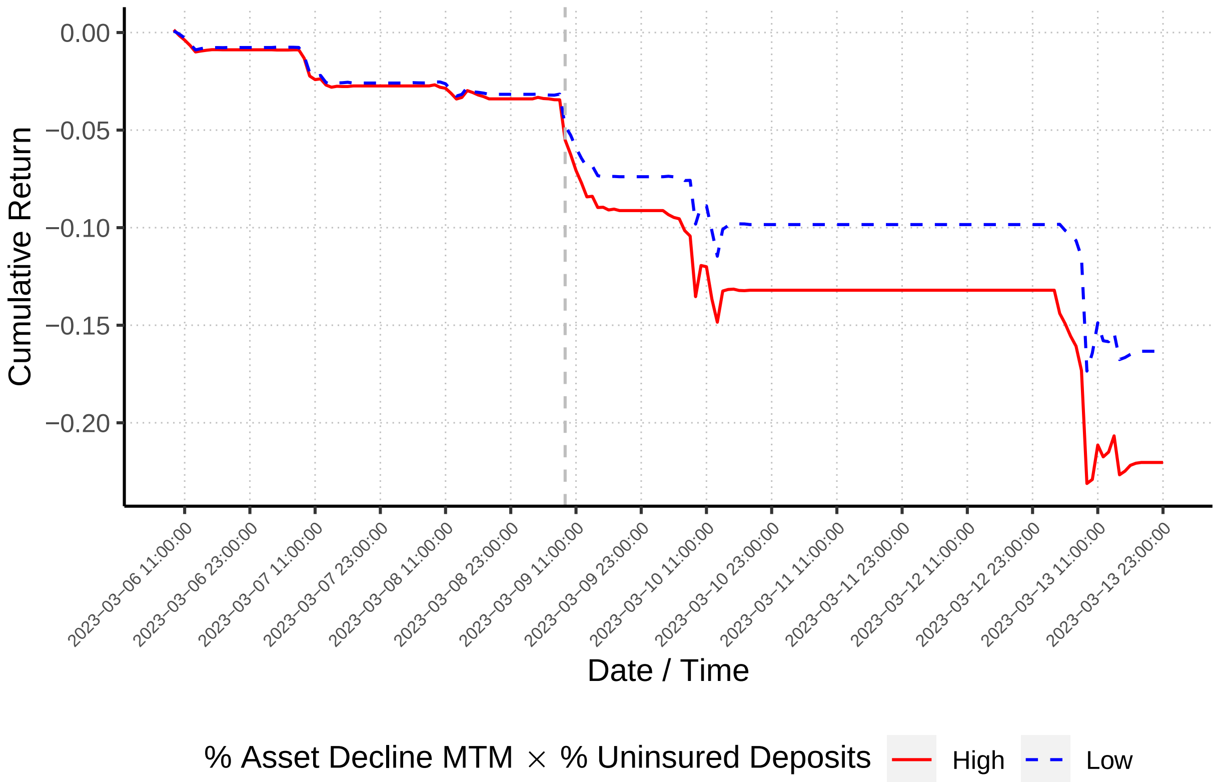

Notes: This figure shows the cumulative returns of a group of US publicly listed bank holding companies (BHCs) from March 06 until March 14, 2023. The sample is split by "run exposure" defined as the % Asset Decline due to mark-to-market × the % of uninsured deposits. 'High' represents the top tercile (red solid line), 'Low' represents the bottom two terciles (blue dashed line). The grey vertical dashed line indicates March 09, 2023 at 9am.

Figure 5: Social Media and the Bank Run – Only "High" Risk Exposure BHCs

Comparing only banks at risk, more Twitter conversation corresponds to greater stock losses after the onset of the run.

(a) Including SIVB

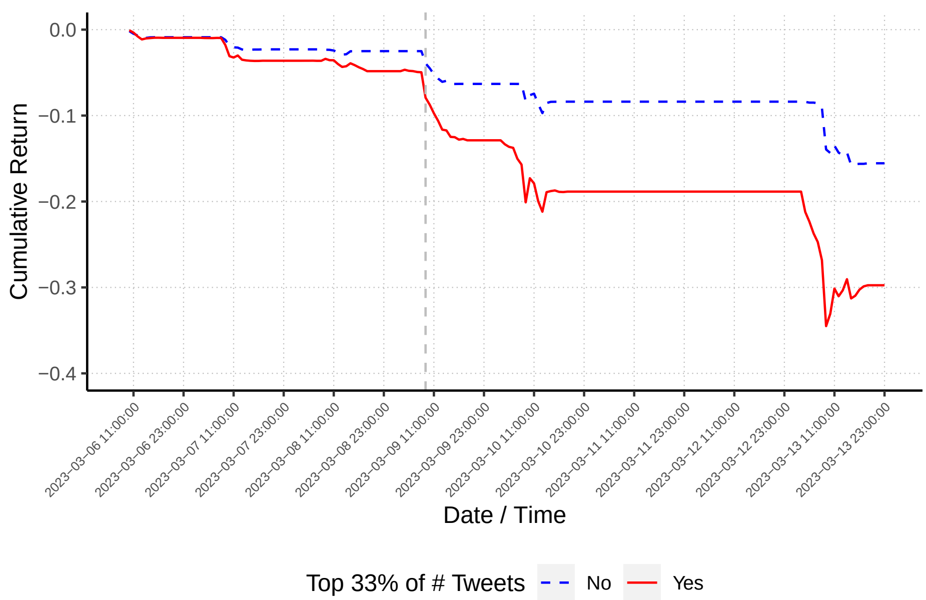

(b) Excluding SIVB

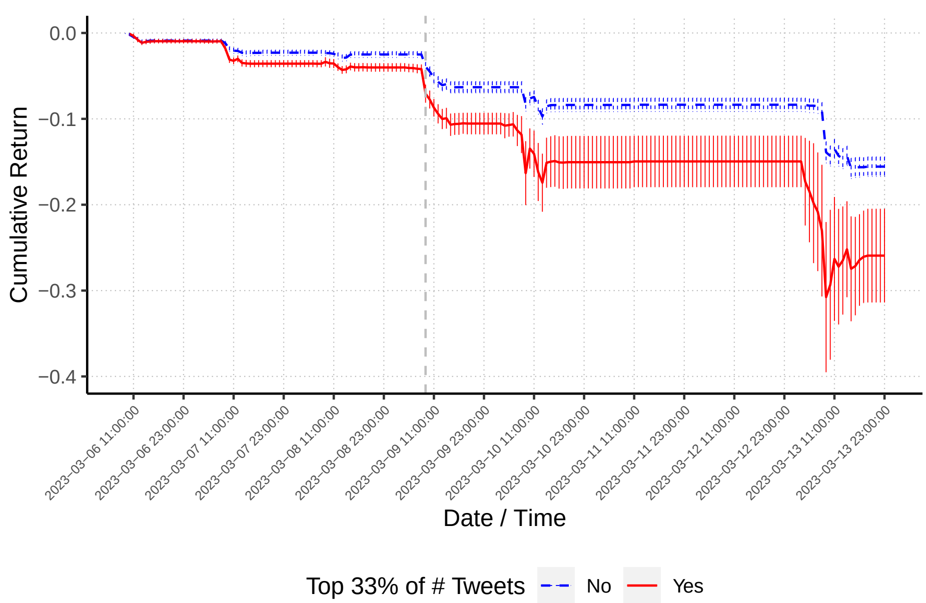

Notes: This figure shows the cumulative returns of a group of US publicly listed bank holding companies (BHCs) from March 06 until March 14, 2023. The sample includes only BHCs with bank run risk-exposure (i.e., % Asset Decline MTM × % Uninsured Deposits) in the top 50%. We then split the sample by the number of Tweets about the BHC in the sample period. 'Yes' represents the top tercile (red solid line), 'No' represents the bottom two terciles (blue dashed line). The grey vertical dashed line indicates March 09, 2023 at 9am. Figure 5b includes 95% confidence intervals around the mean cumulative return.

Table 1: Contextual Dictionaries for Classifying Types of Tweets

This table presents the keyword dictionaries for the major content dictionaries that we employ to classify tweets in to "Run," "Contagion,", "Balance Sheet,", "Cryptocurrency," and "Startup Community" authors. For each of our content dictionaries we provide the seed words used, in the iterative creation process, in italics.

Balance Sheet	*Run Behavior*	*Contagion*	*Crypto*	*Startup Community*
duration	*run*	*systemic*	*crypto*	VC
cover & cash	*withdraw*	*spillover*	USDC	entrepreneur
mortgage backed securities	deposit money	fed	Circle	start up
mismatch	access accounts	regulator	Bitcoin	startup
long maturity	pull & out	#contagion	stablecoin	founder
maturity mismatch	*get & out*	backstop	tech	_venture
marked to market		whole system	FTX	venture capital
mark to market		spreading	peg	
portfolio management		sparks	BlockFi	
liquidity		broader effects	Ripple	
insured deposits		financial system	depeg	
MBS		meltdown	#crypto	
hold to maturity		*contagion*	#blockchain	
HTM			BTC	
portfolio of loans			*silverlake*	
liquidity management				
uninsured deposits				
balance sheet				

Table 2: Summary Statistics on Tweet Activity

This table presents summary statistics on the number of tweets at the bank level. Panel (a) presents the number of tweets about a bank's cashtag per 30 days for the pre-run period (Jan 1 through Feb 15, 46 days) and the run period (Mar 8 through Mar 13, 5 days). Panel (b) presents the average and median number of tweets in and out of the run period by tercile. Panel (c) presents the counts of tweets by bank for "Run" behavior during the run period, "Contagion" tweets during the run period, the total number of tweets in the pre-run period ("Tweets Pre-Run"), and tweets that co-mention cryptocurrency discussion in the pre-run period ("Crypto Pre-Run"). For this table, the pre-run period is Jan 1 through Feb 15, 2023. The run period is defined as Mar 8 through Mar 13. As a comparison to these specific examples, we also report the 90th percentile of tweet counts across all banks in our sample.

(a) Distribution of tweets per 30 days across banks − Run and Pre-Run Periods

	Mean	Std	Min	10%	33%	50%	66%	90%	Max
Pre-run period total tweets	536.74	2207.19	1.96	33.26	60.65	88.70	136.30	511.30	23478.91
Run period	2278.98	23178.72	6.00	12.00	42.00	66.00	120.00	732.00	415746.00

(b) Average and Median Number of Tweets by Tercile, Run and Pre-Run Periods

	Bottom Tercile		Middle Tercile		Top Tercile	
	Mean	Median	Mean	Median	Mean	Median
Pre-run period total tweets	22.24	22.5	94.58	89	2,061.47	344
Run period	9.35	8	40.72	39	4,622.62	165

(c) Top Five Banks by "Run Behavior" Dictionary

	Run	Contagion	Tweets Pre-Run	Crypto Pre-Run
SIVB	6,528	9,662	1,163	20
FRC	1,249	1,368	1,257	343
SI	343	342	20,774	356
SBNY	260	106	2,403	106
JPM	206	245	30,063	275
90th Percentile	3	2	784	3

Table 3: Hourly Bank Stock Losses and Bank Run Exposure

This table presents sample splits and ordinary least squares (OLS) estimates for the effect of mark-to-market losses and percentage of uninsured deposits on bank stock losses at the hourly frequency during the run period. Panel 3a shows the average hourly returns for bank holding companies in our sample, split by "Run Exposure" (i.e., mark-to-market losses × percentage of uninsured deposits), for the period from March 01st to 08th and for March 9th to 14th, as well as the mean and confidence interval of the difference. The dependent variable in Panel 3b is the hourly return (in %) for a bank stock in our sample. '% Uninsured (z)' is the percentage of deposits at the bank that exceed the FDIC threshold of $250,000, drawn from 2022:Q4's FDIC Call Reports. '% Asset ↓ MTM (z)' is the percentage of mark-to-market bank asset losses between 2022:Q1 to 2023:Q1, following the construction of Jiang et al. (2023b). All variables labeled with '(z)' are scaled to have mean of zero and standard deviation of one. The indicator variable '$1(\geq$ Mar 9)' equals one after the onset of the SIVB bank run on March 9th, at 9am, and zero otherwise. Firm- and Day-by-Hour fixed effects are included as indicated. Robust standard errors that are clustered at the bank level are reported in parentheses. *, ** and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

(a) Sample Split by Run Exposure and Run Occurrence

	Mean Hourly Stock Returns (in %)			
% ↓ MTM × % Uninsured	[Mar 01–08]	[Mar 09–14]	Mean Diff.	90% CI Diff.
Bottom 67%	-0.0929	-0.5344	0.4414	[0.3576; 0.5253]
Top 33%	-0.0959	-0.7100	0.6141	[0.5302; 0.6980]

(b) Regression with Hourly Stock Returns

	Hourly Stock Return (%)		
	(1)	(2)	(3)
(Intercept)	-0.1404***		
	(0.0079)		
$1(\geq$ Mar 09)	-0.4287***	-0.4339***	
	(0.0217)	(0.0253)	
% Asset ↓ MTM (z)	0.0054		
	(0.0099)		
% Uninsured (z)	-0.0391**		
	(0.0162)		
$1(\geq$ Mar 09) × % Asset ↓ MTM (z)	-0.0616**	-0.0631**	-0.0725***
	(0.0240)	(0.0259)	(0.0271)
$1(\geq$ Mar 09) × % Uninsured (z)	-0.1073**	-0.1401*	-0.1324*
	(0.0503)	(0.0798)	(0.0743)
% Asset ↓ MTM (z) × % Uninsured (z)	0.0044		
	(0.0065)		
$1(\geq$ Mar 09) × % Asset ↓ MTM (z) × % Uninsured (z)	-0.0412**	-0.0352**	-0.0401**
	(0.0169)	(0.0178)	(0.0195)
Observations	13,026	13,026	13,026
R^2	0.0098	0.0211	0.2576
Within R^2		0.0091	0.0010
Firm FE		✓	✓
Day-by-Hour FE			✓

Table 4: Bank Stock Losses and Social Media Exposure

This table presents ordinary least squares estimates for the effect of social media exposure, mark-to-market losses, and percentage of uninsured deposits on bank stock losses during the run period. The dependent variable is the percentage of bank stock value that is lost by March 14. Social media exposure is measured as the number of tweets in the pre-run period from January 1 through February 15 that contain the bank's cashtag. All specifications employ terciles of this social media exposure variable to mitigate the influence of outlier observations. % Uninsured is the percentage of deposits at the bank that exceed the FDIC threshold of $250,000, drawn from 2022:Q4's FDIC Call Reports. % Loss is the percentage of mark-to-market bank asset losses, construction following Jiang et al. (2023b). Robust standard errors and reported in parentheses. *, ** and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	Dependent variable:				
	% of Stock Value Lost During Run				
	(1)	(2)	(3)	(4)	(5)
% Uninsured (z)	4.117***		1.223		1.288
	(1.025)		(0.895)		(0.893)
% Loss (z)	0.804			−0.069	−0.487
	(0.873)			(0.362)	(0.733)
% Uninsured (z):% Loss (z)	0.943				−0.980
	(0.735)				(0.782)
Mid SocialExp (T2)		0.579	0.074	0.575	0.276
		(0.798)	(0.870)	(0.834)	(0.861)
... × % Uninsured (z)			1.527		1.588
			(1.143)		(1.150)
... × % Loss (z)				0.461	1.425
				(0.689)	(0.966)
... × % Uninsured (z):% Loss (z)					0.990
					(1.005)
High SocialExp (T3)		6.660***	5.209***	6.464***	6.302***
		(1.490)	(1.306)	(1.542)	(1.497)
... × % Uninsured (z)			3.278*		4.157**
			(1.831)		(2.016)
... × % Loss (z)				−0.866	2.170
				(1.201)	(1.990)
... × % Uninsured (z):% Loss (z)					3.014**
					(1.277)
Constant	16.368***	13.453***	13.893***	13.477***	13.735***
	(0.618)	(0.538)	(0.686)	(0.587)	(0.665)
Observations	280	280	280	280	280
R^2	0.158	0.093	0.219	0.097	0.258
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 5: Bank Stock Losses and Startup Community Tweets during the Run

This table presents ordinary least squares estimates for the effect of social media exposure, mark-to-market losses, and percentage of uninsured deposits on bank stock losses during the run period. The specification is the same as in Table 4, except that it employs terciles of the fraction of startup community tweets about a cashtag during the run period, instead of aggregate tweet activity in the pre-period. This specification tests whether exposure to VC-backed startup depositors on Twitter matters for bank run risk. Robust standard errors and reported in parentheses. *, ** and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	Dependent variable:				
	% of Stock Value Lost During Run				
	(1)	(2)	(3)	(4)	(5)
% Uninsured (z)	4.117***		0.691		0.654
	(1.025)		(0.555)		(0.615)
% Loss (z)	0.804			−0.051	0.753
	(0.873)			(0.303)	(0.518)
% Uninsured (z):% Loss (z)	0.943				0.666
	(0.735)				(0.411)
Mid Startup Tweets (T2)		0.801	0.602	0.689	0.556
		(0.747)	(0.745)	(0.744)	(0.773)
... × % Uninsured (z)			1.853**		2.026**
			(0.869)		(1.012)
... × % Loss (z)				−1.047**	−0.993
				(0.528)	(0.795)
... × % Uninsured (z):% Loss (z)					−0.060
					(0.779)
High Startup Tweets (T3)		8.200***	6.322***	8.112***	6.706***
		(1.561)	(1.269)	(1.569)	(1.341)
... × % Uninsured (z)			3.633**		4.544**
			(1.668)		(1.831)
... × % Loss (z)				−0.949	−0.989
				(1.391)	(1.422)
... × % Uninsured (z):% Loss (z)					2.083
					(1.412)
Constant	16.368***	13.116***	13.346***	13.124***	13.464***
	(0.618)	(0.470)	(0.526)	(0.479)	(0.526)
Observations	280	280	280	280	280
R^2	0.158	0.139	0.244	0.145	0.280
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 6: Bank Stock Losses and Social Media Exposure, Controlling for Tweet Activity Posted During Run Period

This table presents ordinary least squares estimates for the effect of social media exposure on bank stock losses during the run period, controlling for tweet activity on different topics and by different authors during the run period. The dependent variable is the percentage of bank stock value that is lost by March 14. Social media exposure is measured as the number of tweets in the pre-run period from January 1 through February 15 that contain the bank's cashtag. All specifications employ terciles of this social media exposure variable to mitigate the influence of outlier observations. Total "Run" and "Contagion" tweets during the run are calculated by counting the number of tweets with at least one word in the "Run" and "Contagion" dictionaries during the run period of March 8 through March 13, and then divided into three groups: zero (the omitted group), below-median and above-median. Startup community tweets are counted during the same period, but based on a count of the number of tweets about the bank stock authored by users with at least one startup community term in their Twitter user description. Like the social media exposure variable, this variable is divided into terciles. Robust standard errors and reported in parentheses. *, ** and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	Dependent variable:				
	% of Stock Value Lost During Run				
	(1)	(2)	(3)	(4)	(5)
Mid SocialExp (T2)	0.579	0.562	0.456	−0.298	0.133
	(0.798)	(1.280)	(1.271)	(1.459)	(1.275)
High SocialExp (T3)	6.660***	2.589*	1.950	2.533	0.210
	(1.490)	(1.368)	(1.396)	(1.741)	(1.543)
Below-median "Contagion" Tweets in Run		10.814***			5.050
		(2.914)			(3.117)
Above-median "Contagion" Tweets in Run		17.712***			9.301***
		(1.983)			(2.651)
Below-median "Run" Tweets in Run			3.282		1.128
			(2.212)		(2.288)
Above-median "Run" Tweets in Run			19.177***		11.282***
			(1.992)		(2.762)
Mid Startup Tweets in Run (T2)				0.589	0.507
				(1.407)	(1.229)
High Startup Tweets in Run (T3)				6.536***	2.705*
				(1.624)	(1.514)
Constant	13.453***	13.204***	13.453***	12.927***	13.034***
	(0.538)	(0.992)	(0.983)	(1.180)	(1.030)
Observations	280	280	280	280	280
R^2	0.093	0.310	0.322	0.151	0.363

Table 7: Social Media and the Bank Run – Hourly Frequency

This table presents OLS estimates for the effect of social media activity on the relationship between hourly stock returns and run exposure around the onset of the SVB bank run on March 09, 2023. The dependent variable in all Panels is the hourly return (in %) for a bank stock in our sample. Similar to Table 3, "Run Exposure (z)" is defined as % Asset ↓ MTM' (i.e., the percentage of mark-to-market bank asset losses as in Jiang et al., 2023b) × '% Uninsured' (i.e., the percentage of deposits below the FDIC insurance threshold). The indicator variable '1(≥ Mar 9)' equals one after the onset of the SIVB bank run on March 9th, at 9am, and zero otherwise. Firm- and Day-by-Hour fixed effects are included as indicated. '# Tweets (4h) (z) (t-1)' is the number of Tweets posted about the bank in the 4 hours prior to the current period t, i.e., in hours $t-5$ through $t-1$. The sample is organized at the firm-by-day-by-hour level. All variables labeled with '(z)' are standardized to have mean zero and standard deviation of one. Panel 7a presents the baseline estimations using observations between March 06th and March 14th. Panel 7b uses a similar sample period but excludes Silicon Valley Bank (SIVB). Panel 7c uses a sample from March 07 (9am) to March 09 (4pm). Firm- and Day-by-Hour fixed effects are included as indicated. Robust standard errors that are clustered at the bank level are reported in parentheses. *, ** and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

$$
\begin{aligned}
r_{i,t} = a &+ b_1 \times 1(\geq \text{Mar } 09)_t + b_2 \times \text{Run Exposure}_i + b_3 \times \text{N Tweets}_{i,t-1} \\
&+ b_4 \times 1(\geq \text{Mar } 09)_t \times \text{Run Exposure}_i + b_5 \times 1(\geq \text{Mar } 09)_t \times \text{N Tweets}_{i,t-1} \\
&+ b_6 \times \text{Run Exposure}_i \times \text{N Tweets}_{i,t-1} \\
&+ b_7 \times 1(\geq \text{Mar } 09)_t \times \text{Run Exposure}_i \times \text{N Tweets}_{i,t-1} + \delta_i + \gamma_t + \epsilon_{i,t}
\end{aligned}
$$

(a) Run Exposure and Tweets over the last 4h

	Hourly Stock Return (%)		
	(1)	(2)	(3)
(Intercept)	-0.1437***		
	(0.0087)		
1(≥ Mar 09)	-0.4462***	-0.4712***	
	(0.0226)	(0.0281)	
Run Exposure (z)	-0.0002		
	(0.0131)		
# Tweets (4h) (z) (t-1)	-0.0435	0.1233	-0.3499
	(0.1189)	(0.2322)	(0.2643)
1(≥ Mar 09) × Run Exposure (z)	-0.0960***	-0.1374***	-0.1321***
	(0.0321)	(0.0378)	(0.0346)
1(≥ Mar 09) × # Tweets (4h) (z) (t-1)	-0.3022	-0.4407	-0.1424
	(0.3453)	(0.3139)	(0.3604)
Run Exposure (z) × # Tweets (4h) (z) (t-1)	0.2839	1.175***	1.103***
	(0.1951)	(0.3947)	(0.3650)
1(≥ Mar 09) × Run Exposure (z) × # Tweets (4h) (z) (t-1)	-0.1908	-1.058***	-0.9453***
	(0.2093)	(0.3443)	(0.3264)
Observations	12,915	12,915	12,915
R^2	0.0138	0.0263	0.2630
Within R^2		0.0135	0.0085
Firm FE		✓	✓
Day-by-Hour FE			✓
SE Cluster	Firm	Firm	Firm

... continued

(b) Excluding Silicon Valley Bank (SIVB)

	Hourly Stock Return (%)		
	(1)	(2)	(3)
(Intercept)	-0.1404***		
	(0.0080)		
1(≥ Mar 09)	-0.4354***	-0.4513***	
	(0.0182)	(0.0227)	
Run Exposure (z)	0.0042		
	(0.0122)		
# Tweets (4h) (z) (t-1)	-0.1064	0.2045	-0.4067**
	(0.1154)	(0.2234)	(0.1664)
1(≥ Mar 09) × Run Exposure (z)	-0.0867***	-0.1027***	-0.0991***
	(0.0316)	(0.0341)	(0.0303)
1(≥ Mar 09) × # Tweets (4h) (z) (t-1)	0.1904**	-0.1001	0.2701**
	(0.0862)	(0.2105)	(0.1319)
Run Exposure (z) × # Tweets (4h) (z) (t-1)	0.2100	0.7163**	0.6261**
	(0.2074)	(0.2982)	(0.2842)
1(≥ Mar 09) × Run Exposure (z) × # Tweets (4h) (z) (t-1)	-0.2063	-0.6944**	-0.5487**
	(0.1997)	(0.2887)	(0.2713)
Observations	12,865	12,865	12,865
R^2	0.0098	0.0162	0.2563
Within R^2		0.0106	0.0028
Firm FE		✓	✓
Day-by-Hour FE			✓
SE Cluster	Firm	Firm	Firm

(c) Shorter Time Window: March 08, 9am – March 09, 4pm

	Hourly Stock Return (%)		
	(1)	(2)	(3)
(Intercept)	-0.0734***		
	(0.0223)		
1(≥ Mar 09)	-0.5873***	-0.6261***	
	(0.0287)	(0.0460)	
Run Exposure (z)	-0.0073		
	(0.0280)		
# Tweets (4h) (z) (t-1)	-0.1953	0.5753	0.5607
	(0.1717)	(0.4701)	(0.5914)
1(≥ Mar 09) × Run Exposure (z)	-0.0834***	-0.1306***	-0.1293***
	(0.0319)	(0.0495)	(0.0472)
1(≥ Mar 09) × # Tweets (4h) (z) (t-1)	-0.7844***	-0.9962**	-1.133**
	(0.1919)	(0.4897)	(0.5570)
Run Exposure (z) × # Tweets (4h) (z) (t-1)	0.5214*	1.702***	1.645***
	(0.2722)	(0.5486)	(0.4804)
1(≥ Mar 09) × Run Exposure (z) × # Tweets (4h) (z) (t-1)	-0.9332***	-1.492***	-1.480***
	(0.2717)	(0.5342)	(0.4344)
Observations	4,109	4,109	4,109
R^2	0.0766	0.1415	0.2728
Within R^2		0.0754	0.0245
Firm FE		✓	✓
Day-by-Hour FE			✓
SE Cluster	Firm	Firm	Firm

Table 8: High Frequency Return Responses to Tweets

This table presents OLS estimates for the impact of the sentiment and content on the tweets in our sample on banks' stock price changes. In all regressions, the dependent variable $\Delta p_{i,t}$ is the log change in Bank i's price between the last trade of the $[-15, -5]$ minute window and the first trade of the $[+5, +15]$ minute window around the tweet happening at time t. $\Delta p_{i,t}$ is winsorized at the 1% level and expressed in basis points. VADER Pos (z) and VADER Neg (z) are the positive and negative components of the VADER sentiment respectively standardized using the sample used in each regression. 'Startup Flag' indicates that a given tweet was posted by a member of the startup community, as constructed in Section 2.1.2. 'Contagion Tweet' and 'Run Tweet' indicate if the tweet contains at least one token in the contagion and run behavior dictionaries respectively. 'High Exposure Bank' indicates if the bank belongs to the top 50 banks with highest bank run exposure (MTM losses × uninsured deposits). We exclude observations for which the traded volume associated with the last price in the window before the tweet or the first price in the window after the tweet is zero. In all regressions, we include fund fixed effects. Standard errors are reported in parenthesis and doubled clustered at the bank-day level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

$$\Delta p_{i,t} = a + b \times \text{VADER Pos (z)}_{it} + c \times \text{VADER Neg (z)}_{it} + \gamma_i + \epsilon_{i,t}$$

	(1) $\Delta p_{i,t}$	(2) $\Delta p_{i,t}$	(3) $\Delta p_{i,t}$	(4) $\Delta p_{i,t}$	(5) $\Delta p_{i,t}$	(6) $\Delta p_{i,t}$
VADER Pos(z)	-0.06	-0.02	-1.59	-1.46	-1.54	-0.79
	(0.16)	(0.16)	(1.43)	(1.44)	(1.57)	(0.92)
VADER Neg(z)	-1.60***	-1.56***	-2.72	-2.62	-3.21	-4.61***
	(0.27)	(0.28)	(2.20)	(2.38)	(1.97)	(1.41)
Startup Flag		3.49***	4.92			
		(1.29)	(10.86)			
VADER Pos(z) × Startup Flag		-1.49*	9.85			
		(0.82)	(8.89)			
VADER Neg(z) × Startup Flag		-2.13**	-21.82***			
		(0.93)	(7.29)			
Contagion Tweet				41.71		
				(36.77)		
VADER Pos(z) × Contagion Tweet				21.68		
				(23.73)		
VADER Neg(z) × Contagion Tweet				-28.18**		
				(14.32)		
Run Tweet					-2.68	
					(8.12)	
VADER Pos(z) × Run Tweet					5.32	
					(7.63)	
VADER Neg(z) × Run Tweet					-0.52	
					(9.69)	
VADER Pos(z) × High Exposure Bank						-0.79
						(2.41)
VADER Neg(z) × High Exposure Bank						1.93
						(3.23)
Constant	-0.78	-0.85	-26.17***	-26.06***	-25.90***	-26.19***
	(0.78)	(0.76)	(4.79)	(4.88)	(4.83)	(4.63)
Observations	1521078	1521078	43597	43597	43597	43597
R^2 (%)	1.01	1.02	2.47	2.47	2.46	2.45
Bank FE	✓	✓	✓	✓	✓	✓
Sample	All	All	≥ Mar09	≥ Mar09	≥ Mar09	≥ Mar09

Table 9: High Frequency Bank Responses to Tweets – excluding SIVB and FRC

This table presents OLS estimates for the impact of the sentiment and content on the tweets in our sample on banks' stock price changes excluding Silicon Valley Bank (SIVB) and First Republic Bank (FRC) from the sample. In all regressions, the dependent variable $\Delta p_{i,t}$ is the log change in Bank i's price between the last trade of the $[-15, -5]$ minute window and the first trade of the $[+5, +15]$ minute window around the tweet happening at time t. $\Delta p_{i,t}$ is winsorized at the 1% level and expressed in basis points. VADER Pos (z) and VADER Neg (z) are the positive and negative components of the VADER sentiment respectively standardized using the sample used in each regression. 'Startup Flag' indicates that a given tweet was posted by a member of the startup community, as constructed in Section 2.1.2. 'Contagion Tweet' and 'Run Tweet' indicate if the tweet contains at least one token in the contagion and run behavior dictionaries respectively. 'High Exposure Bank' indicates if the bank belongs to the top 50 banks with highest bank run exposure (MTM losses × uninsured deposits). We exclude observations for which the traded volume associated with the last price in the window before the tweet or the first price in the window after the tweet is zero. In all regressions, we include fund fixed effects. Standard errors are reported in parenthesis and doubled clustered at the bank-day level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

$$\Delta p_{i,t} = a + b \times \text{VADER Pos (z)}_{it} + c \times \text{VADER Neg (z)}_{it} + \gamma_i + \epsilon_{i,t}$$

	(1) $\Delta p_{i,t}$	(2) $\Delta p_{i,t}$	(3) $\Delta p_{i,t}$	(4) $\Delta p_{i,t}$	(5) $\Delta p_{i,t}$	(6) $\Delta p_{i,t}$
VADER Pos(z)	-0.00	0.04	-0.98	-0.52	-0.78	-0.72
	(0.16)	(0.16)	(0.94)	(0.91)	(1.27)	(0.85)
VADER Neg(z)	-1.38***	-1.34***	-3.63***	-4.03***	-4.26***	-3.95***
	(0.26)	(0.27)	(1.08)	(1.23)	(1.14)	(1.21)
Startup Flag		3.21**	0.06			
		(1.25)	(11.72)			
VADER Pos(z) × Startup Flag		-1.79**	18.60			
		(0.74)	(13.62)			
VADER Neg(z) × Startup Flag		-1.82**	-17.72***			
		(0.82)	(5.94)			
Contagion Tweet				-6.47		
				(34.30)		
VADER Pos(z) × Contagion Tweet				-10.09		
				(21.78)		
VADER Neg(z) × Contagion Tweet				-2.11		
				(16.77)		
Run Tweet					8.98	
					(6.13)	
VADER Pos(z) × Run Tweet					5.82	
					(20.67)	
VADER Neg(z) × Run Tweet					-0.86	
					(9.00)	
VADER Pos(z) × High Exposure Bank						2.51
						(6.27)
VADER Neg(z) × High Exposure Bank						-3.13
						(3.16)
Constant	-0.09	-0.15	-6.45	-6.53	-6.99	-6.57
	(0.77)	(0.76)	(4.64)	(4.52)	(4.61)	(4.52)
Observations	1484132	1484132	19673	19673	19673	19673
R^2 (%)	.49	.5	.91	.87	.88	.87
Bank FE	✓	✓	✓	✓	✓	✓
Sample	All	All	≥ Mar09	≥ Mar09	≥ Mar09	≥ Mar09

Appendix

Figure A.1: Prices of Treasury ETFs across maturities

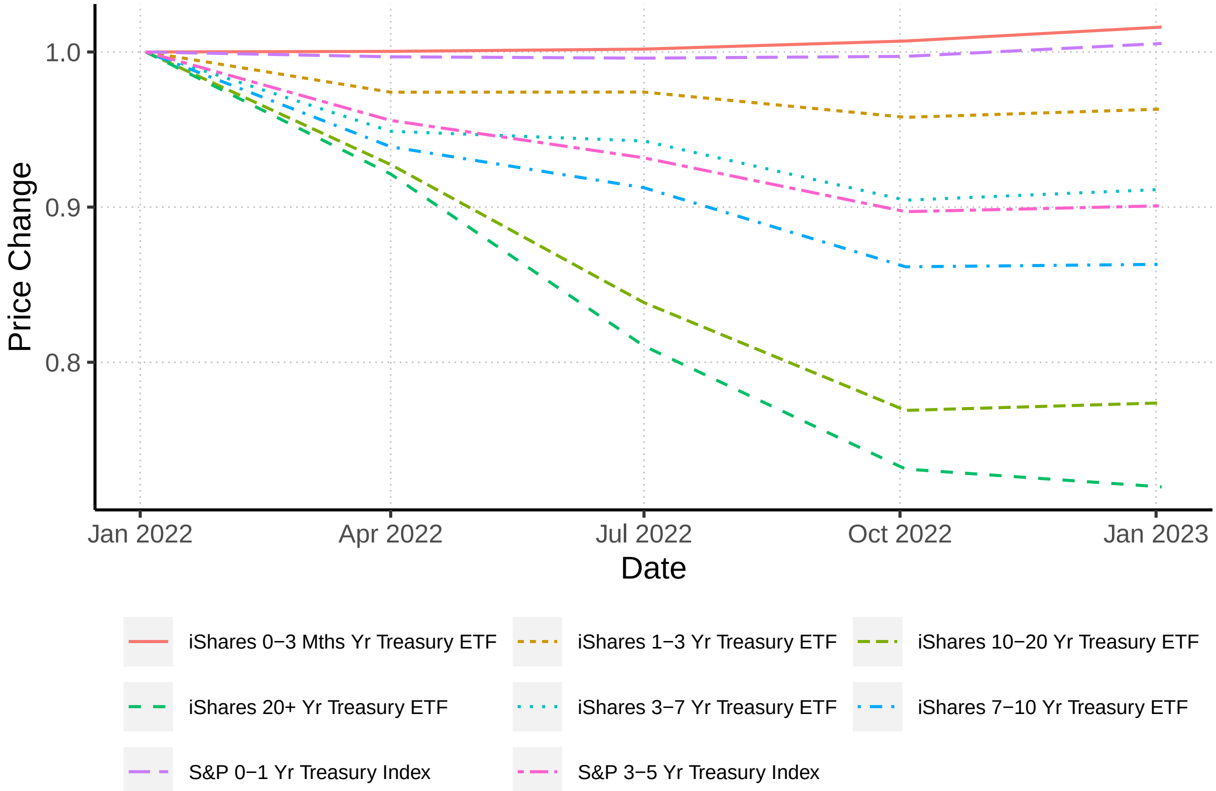

Notes: This figure shows changes in the prices of iShares Treasury Bond ETFs and S&P Treasury Indices across various maturities at the quarterly frequency, for the period from 2022:Q1 to 2023:Q1. ETF Prices and Treasury Indices are scaled by starting values in 2022:Q1.

Figure A.2: Distribution of mark-to-market asset declines across banks

(a) All banks

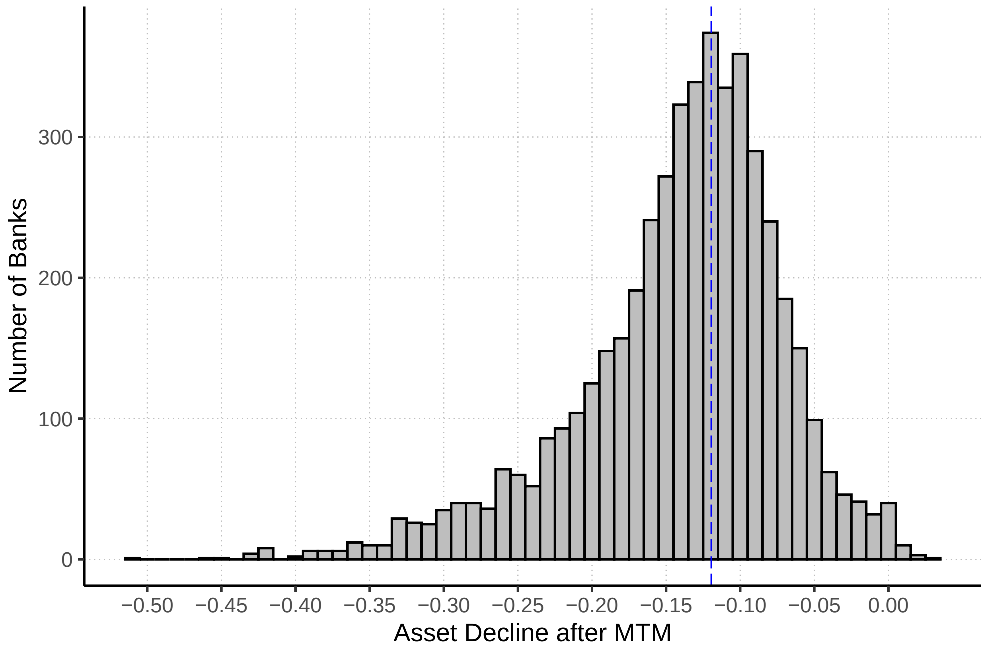

(b) Publicly listed bank holding companies (BHCs)

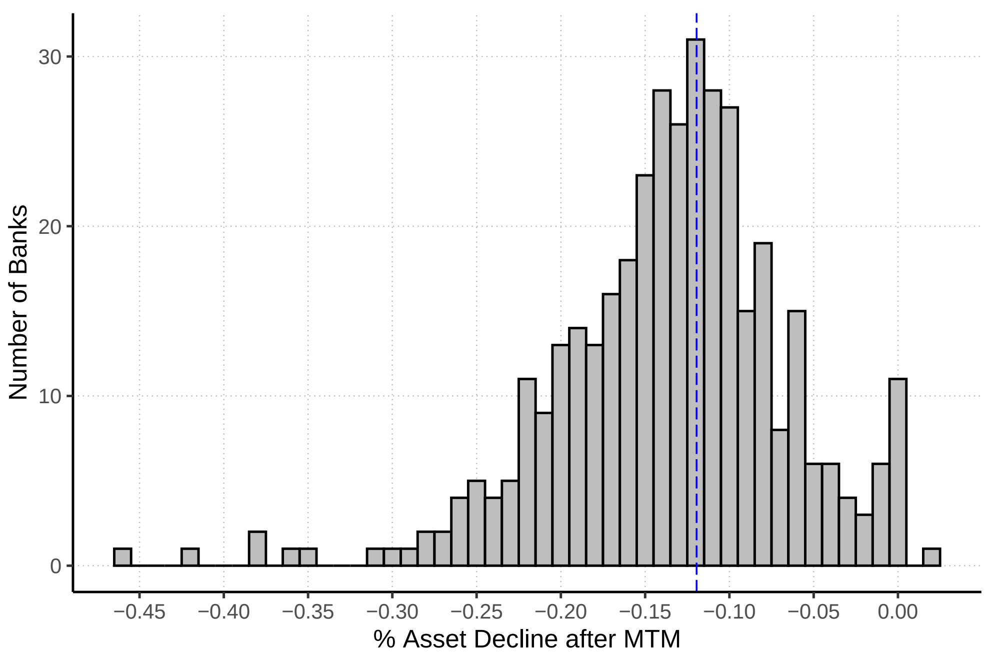

Notes: This figure shows the distribution of mark-to-market implied asset value changes between Q1 of 2022 and Q1 of 2023, constructed following Jiang et al. (2023b). The blue, dashed, vertical line indicates Silicon Valley Bank (SVB). Figure A.2a includes all FDIC-insured banks with call reports available from the FFIEC. Figure A.2b includes all public bank holding companies with available stock return data after aggregating at the bank holding company level.

Figure A.3: Mark-to-market asset declines and bank asset size

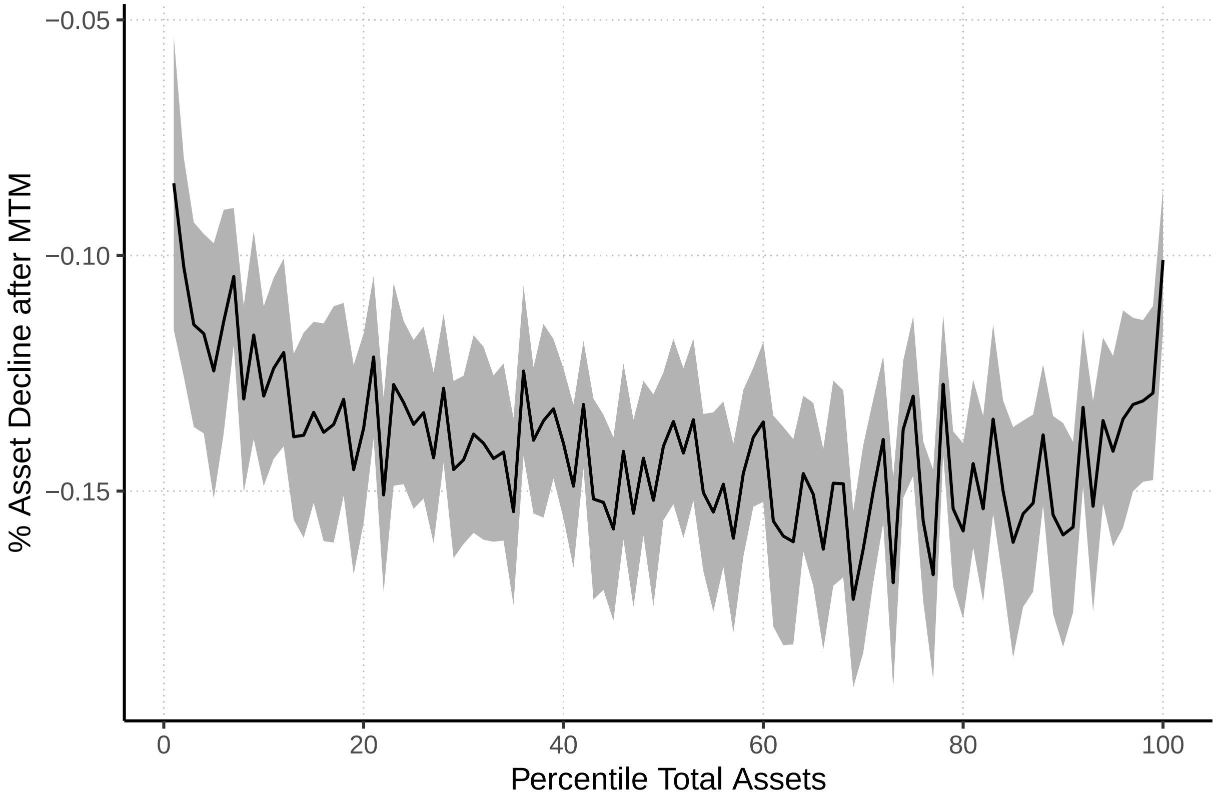

Notes: This figure shows the distribution of average mark-to-market implied bank asset changes between 2022:Q1 and 2023:Q1 (constructed following Jiang et al., 2023b) across percentiles of bank size measured as total assets in 2022:Q1. The smallest banks are in the percentile on the very left of the figure, the largest banks are in the percentile on the very right of the figure. The grey-shaded are indicates two standard deviations around the mean.