Bad Bank, Bad Luck?
Evidence from 1 Million Firm-Bank Relationships

Yannick Schindler
LSE

Peter John Lambert
LSE

February 11, 2025*

[Click here for the latest version]

Abstract

This paper studies the effects of bank failure on firm performance. We collect 36 million loan records to build a novel dataset on the credit relationships of 1.8 million US firms, predominantly composed of small and medium-sized enterprises (SMEs). We then implement a staggered treatment difference-in-differences estimation strategy using 179 bank failures from 1990 to 2023 to estimate the impact of bank failure on firm-level survival and employment growth. Although the US regulatory framework resolves failed banks through forced acquisitions by healthier institutions—a process designed to minimize disruption—we find substantial negative effects. Firms with a credit relationship to a bank that subsequently fails are 6.7 percentage points (44.3%) more likely to fail themselves within five years, while surviving firms exhibit 25% lower employment growth compared to those banking with non-failed institutions. These impacts persist for more than 10 years, are evident during both crisis and non-crisis periods, and are strongest among very small firms—a firm size segment that we are the first to study in this context. Our estimated effects are further supported by two natural experiments. Surprisingly, we also observe that a small subset of bank failures had positive effects on firm outcomes, suggesting that, in some cases, bank failure can be fortuitous for affected firms. Overall, our findings suggest that bank failures exert a substantially larger influence on the real economy than previously recognized, possibly requiring a re-evaluation of current regulatory approaches to managing such events.

Keywords: bank failure, bank-firm relationships, credit shocks, banking, financial intermediation, firm performance.
JEL Codes: G01, G21, G33, E44, R11

*The authors are indebted to Tim Besley, Wouter Den Haan, Ben Moll, John Van Reenen, and Andrew Scott for their PhD supervision and mentorship. We would also like to express our gratitude to Robin Burgess, Wei Cui, Gabriel Chodorow-Reich, Arnaud Dyèvre, Mark Gertler, Simon Gilchrist, Ethan Ilzetzki, Camille Landais, Hugo Reichardt, Ricardo Reis, and Alwyn Young for their insightful feedback and advice. We also thank Luca Barbato, Suraj Rajesh, and Juan Bautista Sosa along with other participants of the AEAI 2024 Summer Fellowship. We are grateful to seminar participatins in the macroeconomics and industrial organization groups at LSE as well as to Jesse Edgerton (Two Sigma) and Kirk White (US Census Bureau). The authors also gratefully acknowledge financial support from the Hayek Programme and the Programme on Innovation and Diffusion (POID) at LSE.

1 Introduction

A fundamental question in macroeconomics is how shocks to the financial sector propagate to the real economy. Periods of heightened financial distress have historically coincided with severe economic downturns—ranging from the Great Depression to the Global Financial Crisis (GFC)—raising an important question: do disruptions in the financial sector merely reflect underlying weakness in the broader economy, or do they actively contribute to economic decline? The recent high-profile collapses of Silicon Valley Bank and First Republic Bank in the United States (US) have brought renewed attention to understanding the extent to which financial sector distress transmits to the real economy.

This paper provides new evidence on how financial shocks affect the real economy by examining the impact of bank failures on small businesses. Small businesses serve as a crucial transmission channel between bank distress and the real economy, as they face unique financial constraints that make them particularly vulnerable to banking disruptions. Unlike larger firms, small businesses cannot access market-based financing through equity or bond markets, making them heavily dependent on bank credit for funding (Cole and Wolken 1996; Berger and Udell 1998; Robb and Robinson 2014). This dependence is compounded by small businesses typically relying on a single bank for credit (Petersen and Rajan 1994) and the substantial costs they face when attempting to switch their banking relationship from one bank to another (Petersen and Rajan 1994; Berger and Udell 1995; Cole 1998; Elsas and Krahnen 1998; Drexler and Schoar 2014).

The prevalence and economic significance of small businesses makes them a critical component for understanding how financial shocks affect the broader economy. In the US, firms with fewer than 100 employees account for over 98% of employer firms and more than one-third of private sector employment.[1] These businesses are key drivers of job creation, innovation, and economic growth in both developed and developing economies (Acs and Audretsch 1988; Cohen and Klepper 1996; Baumol 2004; Neumark, Wall, and Zhang 2011; Haltiwanger, Jarmin, and Miranda 2013; Ayyagari, Demirguc-Kunt, and Maksimovic 2011). Yet, despite the central role that small businesses play in the economy, our understanding of how firms weather financial disruptions—in particular the failure of banks—has been limited to larger firms (see e.g. Slovin, Sushka, and Polonchek (1993), Brewer, et al. (2003), Minamihashi (2011), Giannetti and Simonov (2013)).

The real effects of bank failures on small business borrowers are not obvious ex ante. One view, articulated by Fama (1980), argues that bank failures should have limited effects on firms: in a

[1] See the US Census 2021 SUSB Annual Data Tables by Establishment Industry for detailed breakdowns.

competitive banking system with well-functioning markets, firms should be able to seamlessly switch to other banks. Moreover, the failed bank resolution process—generally executed by the Federal Deposit Insurance Company (FDIC) forcing an acquisition of the struggling bank by a better-capitalized bank—aims to seamlessly transfer borrowers' loans to healthier institutions. In practice, several factors can make these transitions highly disruptive for small businesses. For example, bank acquisitions trigger branch closures (Nguyen 2019; Vij 2020), which disrupt firm-bank relationships and potentially destroy the "soft information"—such as detailed knowledge of the business owner's character, local market conditions, and community relationships—that loan officers accumulate over time and that are critical for small businesses and banks to overcome information asymmetries.[2] Whether the benefits of being matched with a healthier lender through the FDIC's resolution process outweigh the costs of disrupted lending relationships is ultimately an empirical question—one that is crucial both for understanding how financial sector distress propagates to the real economy and for evaluating the effectiveness of current bank resolution practices.

Studying how bank failures affect small businesses has traditionally been difficult because of a lack of loan-level data for smaller US firms. To overcome this challenge, we collect and process new data spanning 36 million loan documents from 1990 to 2023. We also develop and implement a novel approach to link these documents to firm-level microdata. This approach combines two key components: First, we employ modern text embedding models—neural networks trained on vast text corpora—to transform unique business names into high-dimensional vectors that capture semantic relationships. This technology allows us to match businesses even when their trading names differ from their legal registered names (i.e. matching "BoA" to "Bank of America") and handles variations in spelling and corporate suffixes (e.g. "Corp." vs. "Corporation"). Second, we geocode over 73 million unique addresses to augment the name matching with geographic information. We then use machine learning techniques to optimally combine name similarity and location proximity to predict match probability. The resulting dataset is, to the best of our knowledge, the most comprehensive dataset of firm-bank relationships for the small business segment in the US.[3] By analyzing these data and studying 179 bank failures, we examine the impact of bank failure on small business survival and employment growth. The large number of bank failures in our sample allows us to explore how specific aspects of a bank's

[2]For empirical evidence documenting the important role that long-term relationships between firms and banks play in the extension of credit to smaller firms, see e.g. Berger and Udell (1995), Petersen and Rajan (1994), and Petersen and Rajan (1995). The importance of these long-term relationships in propagating aggregate shocks is explored in den Haan, Ramey, and Watson (2003).

[3]In Appendix A.3 we provide a more detailed comparison of our dataset to other datasets commonly relied on in the literature that studies the firm-level effects of financial shocks. These datasets include *Dealscan*, FR Y-14, and the Shared National Credit (SNC) data.

collapse—e.g. regional or industry concentration—can worsen the effects of bank failure on firms.

Our headline finding is that firms with pre-existing relationships to banks that later failed typically experience large and lasting negative consequences. Compared to businesses that source credit from non-failing banks, a business that gets credit from a bank that fails is 3 percentage points more likely to cease operating within one year of the bank failure. This represents a 46.3% increase in the likelihood of business failure. These negative effects persist and grow over time such that five years after a bank failure, treated firms are 6.7 percentage points more likely to have ceased operations. Even a decade after the bank failure, treated firms show a 7.9 percentage point higher failure rate. The persisting and long-lasting impact of bank failures is a novel finding of this research, made possible by our analysis of three decades of bank failures.

Importantly, the magnitude and persistence of these effects suggest that bank failures do not merely accelerate the exit of already vulnerable firms. If that were the case, we would expect to see a sharp initial increase in failures among treated firms, followed by a convergence in survival rates between treated and untreated firms over time. Instead, the persisting gap in survival rates over 5 and 10 years indicates a scarring, long-lasting negative impact on firm viability. This pattern implies that bank failures create enduring challenges for small businesses, affecting even those firms that initially weathered the shock, and potentially altering their long-term growth trajectories and survival prospects.

Another notable finding in this paper is that the magnitude and time profile of our estimated effects are similar both for bank failures during and after the 2008 US financial crisis, suggesting that the mechanisms through which bank failures harm small businesses operate similarly in both crisis and non-crisis periods. This provides external validity for a number of studies which leverage the 2008 US financial crisis to study the firm-level effects of financial shocks (Almeida, et al. 2012; Chodorow-Reich 2014; Paravisini, et al. 2015; Huber 2018; Bentolila, Jansen, and Jiménez 2018; Berton, et al. 2018). Understanding these channels outside of crisis times is crucial for designing effective bank resolution policies, particularly given that bank failures often occur during non-crisis periods, such as the high-profile bank failures in the US during 2023.[4]

Another key finding in this paper is that we find a very clear size gradient among our already relatively small firms. We find that firms in the bottom tercile of by employment (median: 2 employees) see a two-fold higher negative impact on survival post-failure, compared to firms

[4] Silicon Valley Bank, First Republic, and Signature Bank were the 14th, 16th, and 29th largest banks in the US at the time of failure Federal Reserve (2022b). Their combined assets totaled $532 billion, more than the combined $526 billion of assets (adjusted for inflation) held by the 25 banks that failed in 2008, at the height of the Global Financial Crisis Russell and Zhang (2023).

in the top tercile (median: 12 employees). This highlights the important role that firm size—a predictor of larger credit constraints due to more severe information frictions—plays in shaping the firm-level effects of a bank shock.

While bank failures have significant negative effects on firm survival on average, this masks substantial heterogeneity across different bank failures. When we estimate the effect of each bank failure separately, we find that the impact on 5-year firm survival rates ranges from a 23 percentage point reduction to an 11 percentage point improvement in firm survival probabilities. Rather than clustering into distinct "good" and "bad" failures, the effects follow a nearly linear gradient, suggesting that multiple characteristics of bank failures combine to determine the net impact on firm outcomes. When comparing the magnitude of effects across the distribution, we find a striking pattern: not only are the majority of bank failures damaging to firm survival, but the negative effects among the most damaging failures (up to 23 percentage points) are substantially larger than the positive effects observed among the minority of least damaging ones (up to 11 percentage points). This suggests that while good resolution processes can help firms, the potential for damage from poor resolutions is considerably larger. Our finding of substantial variation in how bank failures affect borrowers allows us to identify which characteristics of failures and resolution processes are most damaging to firms. These results have important implications for bank regulation and resolution policy: they suggest that regulators can potentially minimize the real economic damage of bank failures by targeting specific features of the resolution process.

We also find that the negative consequences of bank failures extend beyond firm survival, significantly affecting employment growth for the firms that remain operational after their bank fails. From the year before the bank failure to the third year after the bank failure, firms that had a banking relationship with the failed bank experience 2.5 percentage point lower employment growth compared to firms banking with non-failed banks. This translates to a substantial 22.5% reduction in growth relative to unaffected firms. The employment effects are not only severe but also persistent. Five years post-failure, the gap in employment growth is 3.7 percentage points, representing a 224.8% lower growth rate. Even a decade after the bank failure, affected firms still lag behind, with employment growth 6 percentage points (or 29.1%) lower than their counterparts. Importantly, we find no evidence of accelerated hiring among affected firms within the 10-year period of our analysis, suggesting that the initial shock to employment growth rates is not offset by subsequent catch-up growth. These findings underscore that bank failures inflict lasting damage on firm-level employment growth, extending well beyond the immediate aftermath of the failure and affect even those businesses resilient enough to survive the initial shock.

To identify the impact of bank failures on small business outcomes, we employ a staggered local projection difference-in-differences (DiD) estimation strategy (Dube, et al. 2024), comparing the outcome trajectories of firms affected by bank failures with those of similar firms that were not. We take several steps to address identification concerns. First, to account for region-specific or industry-specific shocks that might simultaneously affect bank failures and firm outcomes, we include a rich set of county-by-year-by-industry fixed effects. Second, to address concerns that firms might anticipate bank failures and switch lenders, or that poor firm performance might drive bank failures, we conduct extensive analyses of pre-failure firm behavior and bank health. These analyses draw on regulatory post-mortem reports, bank Call Report data, and detailed firm switching patterns, all of which suggest that corporate loan losses were not a significant driver of bank failures and that firms did not systematically switch banks before failure. Third, to address potential sorting of healthier firms away from—and weaker firms into—soon-to-fail banks, we adopt a conservative definition of treated firms that includes only those with relationships established at least three years before failure. By maintaining a firm's treated status even if it switched to a healthy bank before failure, while excluding firms that started new relationships with the soon-to-fail bank in its final years, this approach eliminates bias from strategic firm sorting and produces a lower-bound estimate. Finally, we complement our main analysis by exploiting two natural experiments where bank failures were plausibly exogenous to borrower health: the forced closure of Park National Bank and Citizens National Bank via the FDIC's cross-guarantee powers, and the abrupt failure of Colonial Bank following the discovery of a large-scale mortgage lending fraud scheme.

Related literature: The effects of financial shocks on the real economy have been studied at various levels of aggregation, from country-level analyses (Bernanke 1983; Cardarelli, Elekdag, and Lall 2011) to examinations of regional outcomes (Peek and Rosengren 2000; Ashcraft 2005; Kandrac 2014; Greenstone, Mas, and Nguyen 2020). Our analysis complements this broader literature by providing direct evidence on the firm-level mechanisms through which bank failures propagate to real economic activity. Within the firm-level literature specifically, our paper makes four distinct contributions.

First, our focus on the smallest firms in the economy—including non-employer firms—alongside our analysis of smaller banks marks a significant departure from the existing literature, which primarily examines publicly-traded firms or large private firms borrowing from major banks. The median firm in our data has just three employees and the median bank has $367 million in assets—approximately one-seventh the size of the average FDIC-insured bank and well below the $10 billion threshold commonly used to define "small banks" in US regulatory frameworks.[5] In contrast, important contributions to the literature on the real firm-level effects of financial shocks—Chodorow-Reich (2014) and Huber (2018)—study not only much larger firms (with median employment of 620 and 132 respectively) but also focus on large banks. Chodorow-Reich (2014) examines lending through the syndicated loan market dominated by large banks, while Huber (2018) studies a shock to Commerzbank, one of Germany's largest banks. Other influential studies of financial shocks—such as Gan (2007), Almeida, et al. (2012), Amiti and Weinstein (2011), and Paravisini, et al. (2015)—likewise focus on larger firms. While Khwaja and Mian (2008) and Schnabl (2012) include smaller firms in their analysis, their datasets restrict them to examining how bank shocks affect firm borrowing, whereas our data allow us to examine real firm-level outcomes such as employment and survival. Similarly, while Berton, et al. (2018) study small firms borrowing from small banks, they examine credit supply shocks rather than bank failures, which represent disruptions to both the availability of credit as well as the firm-bank relationship. Bentolila, Jansen, and Jiménez (2018) analyze how bank bailouts affect firm-level employment during the Spanish financial crisis, but also focus on relatively large banks—their sample of bailed-out banks are 15 times larger than the average Spanish bank. Other evidence specifically on the firm-level effects of bank failure or bailouts comes from studies of large, publicly-traded firms (Slovin, Sushka, and Polonchek 1993; Brewer, et al. 2003; Minamihashi 2011; Giannetti and Simonov 2013).

Secondly, we provide evidence on how financial shocks affect firms outside of crisis periods, in contrast to the existing literature which draws primarily from major financial crises such as the Global Financial Crisis (Almeida, et al. 2012; Chodorow-Reich 2014; Paravisini, et al. 2015; Huber 2018; Bentolila, Jansen, and Jiménez 2018; Berton, et al. 2018), Japan's 1990s asset price collapse (Brewer, et al. 2003; Gan 2007; Amiti and Weinstein 2011), or the 1988-1991 Norwegian banking crisis (Ongena, Smith, and Michalsen 2003). Understanding how financial shocks affect firms outside of crises is crucial, as their transmission may differ when many firms face simultaneous distress and limited alternative funding sources. While we find that the effects are similarly severe in both crisis and non-crisis periods, evidence from non-crisis periods remains essential for designing effective bank resolution policies, particularly given that bank failures occur regularly outside of crises.

Thirdly, we complement the existing literature by analyzing how the effects vary across 179 bank failures, in contrast to previous studies that focus on shocks emanating from single institutions such as Continental Illinois Bank (Slovin, Sushka, and Polonchek 1993), Lehman Brothers

[5]Average bank size for FDIC-insured banks obtained from the FDIC's Quarterly Banking Profile historical data. Average calculated from 2004 through 2023.

(Chodorow-Reich 2014), or Commerzbank (Huber 2018).The large number of bank failures covered by our data allows us to identify specific features of bank failures that amplify or attenuate their impact on borrowers.

Fourthly, the long time span and comprehensive nature of our data enable us to chart the full dynamic path of bank failure effects. While data limitations have typically restricted previous studies to estimating average effects over a fixed post-treatment window (Bentolila, Jansen, and Jiménez 2018; Huber 2018; Chodorow-Reich 2014), we are able to estimate separate effects for each year following bank failure for up to a decade. This year-by-year estimation reveals how the impact of bank failures evolves over time, rather than providing just an average post-treatment effect.

The evidence in this paper also contributes to the literature on banking relationships. A substantial literature establishes that relationships help overcome information asymmetries and moral hazard problems (Jaffee and Stiglitz 1990), generating benefits for small firms through better loan terms (Berger and Udell 1995), improved credit access (Petersen and Rajan 1994; Cole 1998), and liquidity insurance (Elsas and Krahnen 1998; Agarwal and Hauswald 2010; Drexler and Schoar 2014). These benefits are particularly pronounced for small firms borrowing from small banks, which Berger, et al. (2005) shows rely more heavily on 'soft' information gathered through personal contact. While this literature extensively documents the value of banking relationships, we provide comprehensive evidence on the costs firms bear when these relationships are severed through bank failure.

The remainder of this paper is organized as follows. Section 2 discusses the institutional background of US bank failures. Section 3 describes our data and record linkage process. Section 4 presents descriptive statistics. Section 5 outlines our empirical strategy and discusses identification. Section 6 presents our main results on how bank failures affect firm survival and growth, and Section 7 concludes with implications for bank regulation and resolution policy.

2 Background on bank failure and small business lending

Understanding the institutional framework of bank failures and small business lending in the United States is crucial for our empirical analysis. This section outlines key features of bank failure regulation and resolution, describes the mechanisms through which bank failures affect small business borrowers and discusses how these institutional features inform our identification strategy. The section concludes by connecting these institutional details to our empirical approach in Section 5.

Prevalence and patterns of bank failure: Bank failures are remarkably common in the United States—nearly three thousand FDIC-insured banks have failed since 1980, with peaks of 530 and 157 annual failures during the Savings & Loans Crisis and Global Financial Crisis respectively (FDIC 2024; Hoggarth, Reis, and Saporta 2002). During the 2007-2011 financial crisis period, approximately 5% of banks, weighted by total asset, failed.[6] While bank failures tend to cluster during periods of financial crisis and economic downturns, significant failures occur in normal times too: in 2023, Silicon Valley Bank, First Republic, and Signature Bank failed, together accounting for over $500 billion in assets (Federal Reserve 2022a).[7]

Bank failures disproportionately affect smaller banks, which are the primary source of credit for small businesses (Berger, Kashyap, and Scalise 1995; Berger, et al. 1998; Peek and Rosengren 1998). During the Global Financial Crisis, the average failed bank was one-third the size of the average FDIC-insured bank, and banks under $10 billion in assets account for 42.5% of small business loans despite holding only 28% of total bank assets (Aibangbee 2022). This relationship is particularly important as small businesses heavily depend on bank credit for funding, being unable to access market-based financing through equity or bond markets (Cole and Wolken 1996; Berger and Udell 1998; Robb and Robinson 2014).

The US bank resolution process: The bank failure and resolution process begins when regulators determine a bank is critically undercapitalized or insolvent.[8] The FDIC then initiates a comprehensive evaluation process (Gnanarajah 2023), appointing an independent financial advisor to assess the bank's assets and establish a reservation value for the resolution process (Johnston-Ross, Ma, and Puri 2021).

The FDIC primarily employs three resolution methods: (1) Purchase and Assumption (P&A) agreements, (2) Deposit Payoffs, or (3) creating a Deposit Insurance National Bank. The P&A method, used in approximately 95% of cases during the financial crisis, involves a first-price sealed bid auction where qualified bidders must meet strict eligibility requirements including well-capitalized status and satisfactory regulatory ratings (Johnston-Ross, Ma, and Puri 2021). The FDIC evaluates these complex bids using a proprietary least-cost test model, selecting the bid that minimizes the cost to the Deposit Insurance Fund.

The impact of bank failures on small businesses is theoretically ambiguous. One view, articulated by Fama (1980), suggests that bank failures should have limited effects in competitive banking

[6] The asset-weighted failure percentage is calculated by summing the assets of all failed banks in this period obtained from the FDIC failed banks assistance data and dividing by total assets for all commercial banks in 2009 obtained FRED series `TLAACBW027SBOG`.

[7] These bank assets made up approximately 2% of all assets at all banks insured by the FDIC in 2023.

[8] A bank is considered critically undercapitalized when its ratio of tangible equity to total assets falls to or below 2%.

markets where firms can easily switch lenders. However, small firms are known to face high costs when switching lenders (Elsas and Krahnen 1998) and typically maintain single-bank relationships (Petersen and Rajan 1994; Berger and Udell 1995; Cole 1998; Drexler and Schoar 2014), making it difficult for small firms to switch their credit demand to another bank following a bank failure.

While P&A transactions—used in over 98% of bank failures included in our sample—are designed to ensure a smooth transition of banking relationships, several features of these forced acquisitions can create significant disruptions for small business borrowers. While the bank acquiring the failed bank's assets and liabilities must honor existing loan terms, it has full discretion over future lending decisions. This can be particularly problematic for small businesses, which often rely on credit lines subject to frequent renewals. The acquiring bank may be less inclined to renew loans upon maturity or extend additional credit, even if the original bank might have done so. This reluctance can stem from the acquiring institution's different assessment of the borrower's financial health or ability to pay, or simply because the new bank may not specialize in lending to that particular type of business (Paravisini, Rappoport, and Schnabl 2023).

Furthermore, acquiring banks frequently close or consolidate branches of failed banks, especially in markets where they already operate (Nguyen (2019), Vij (2020)). These closures can disrupt the relationship lending that small businesses depend on, as the detailed knowledge about borrowers accumulated by loan officers—critical for overcoming information asymmetries in small business lending—is often lost in the process.

The role of deposits vs. loans: While it may seem that loss of deposits is the main concern when a bank fails, we argue that access to deposits is actually minimally disrupted during a bank failure. Bank deposits up to $250,000 (or $100,000 before 20008) are insured by the FDIC and the FDIC has historically paid insured depositors within days following the failure of a bank when opting for the Deposit Payoff resolution method. Furthermore, of the 574 banks resolved by the FDIC since 2000, 543 have been resolved via a P&A agreement, with deposits being transferred to a DINB created by the FDIC while the bidding process for the failed bank's assets and liabilities is carried out.[9] During this time, and also once the failed bank's deposits have been transferred to the acquiring bank, depositors retain full and complete access to their deposits. This applies to both insured and uninsured depositors. No insured depositor has experienced a loss of funds during the FDIC's history and even losses to uninsured deposits are rare. The average loss

[9] A breakdown of bank failures by resolution method can be obtained from the FDIC Bank Failures & Assistance dataset (FDIC 2024).

experienced by uninsured depositors across all bank failures since 1992 has only been 6% and only 3% since 2008 (Federal Deposit Insurance Corporation 2023).

Implications for identification strategy: These institutional features of bank failure and small business lending have important implications for our empirical strategy, which we discuss in detail in Section 5.2. The first and most fundamental challenge is that bank failures and firm outcomes could be driven by common local economic shocks. For example, during the subprime crisis—which coincides with many of the bank failures in our sample— there was substantial geographic variation in economic conditions, including employment, house prices, and foreclosure rates (Gertler and Gilchrist (2018)). Such region-specific shocks—like a local real estate price collapse—could simultaneously increase both bank failure probability and firm exit rates, violating the parallel trends assumption necessary for causal interpretation. To address this, we include granular county-by-year-by-industry fixed effects, comparing only firms within the same local market and industry.

The second challenge is the potential for reverse causality or anticipation effects. However, evidence suggests that bank failures during our sample period were primarily driven by losses on mortgage-backed securities (Erel, Nadauld, and Stulz (2014), Antoniades (2015)) and high exposure to commercial real estate loans (Furlong and Knight (2010), Cole and White (2012)), rather than losses on business loans. We also find no evidence that firms systematically switch banks before failure, which aligns with Correia, Luck, and Verner (2024)'s finding that even depositors do not respond to increases in bank failure risk. Nevertheless, to further address potential sorting concerns, we employ an alternative treatment definition that maintains firms' treatment status even if they switch banks, while excluding firms that initiate new relationships with soon-to-fail banks. This approach eliminates any potential bias from strategic sorting of firms across banks, though it likely understates the true effect of bank failure since it includes in our treatment group firms that may have avoided the worst effects by switching to healthy banks. Finally, we complement our analysis with two natural experiments where bank failures were driven by factors plausibly exogenous to borrower health: the forced closure of banks through the FDIC's "cross-guarantee" powers and a bank failure triggered by the sudden discovery of large-scale fraud in the bank's mortgage lending divison.

3 Data

This section introduces a novel dataset that links firms to their lenders, offering an unprecedented view of the lending relationships that underpin the US economy, particularly for small and medium enterprises (SMEs). We construct this dataset by combining data from Uniform Commercial Code (UCC) filings and comprehensive firm-level information from Dun & Bradstreet.

The following subsections detail the construction of this dataset. Subsection 1 discusses the role of UCC filings in securing loans and their usefulness in identifying lending relationships. Subsection 2 describes the firm-level data used to obtain comprehensive information on businesses. Subsection 3 outlines our methodology for linking the UCC filings to the firm-level data to create the final dataset. This new data product addresses a critical gap in the study of SME lending relationships and offers new opportunities to explore the interplay between the financial sector and the real economy.

3.1 UCC financing statements

The micro-level data on banking relationships we use are derived from Uniform Commercial Code (UCC) financing statements filed with the Secretary of State office in each US state. The UCC is a set of laws that regulates commercial transactions including sales, leases, and rentals, within the US. First drafted in the 1940s, the aim of the UCC is to reduce complexity and costs for firms operating across multiple states by harmonizing commercial laws across the entire US.

UCC filings play a pivotal role in the extension of credit to small businesses as they allow lenders to secure interests in collateral offered by the business, which can include inventory, equipment, or accounts receivable.[10] By filing a UCC financing statement, lenders assert their rights over the business' collateral, ensuring priority over other creditors in case of default. This mechanism significantly mitigates lenders' risk. UCC filings also serves a second purpose: they inform other potential creditors of existing claims on the business' assets thereby ensuring that the business is unable to re-collateralize its assets.

Though UCC regulations are standardized nationwide, each state operates an independent registry for UCC filings. The incentive for lenders to make UCC filings is significant as in their absence lenders would be considered unsecured creditors, and therefore face lower recovery rates in the event of a borrower's default. The low cost of filing, typically ranging from \$15 to \$25, coupled with the convenience of electronic submission, further encourages this practice.

[10] Notably, UCC filings do not encompass real estate transactions or personal property governed by title laws such as airplanes and automobiles. Instead, these filings are used to acquire security interests in assets like equipment, inventories, bank accounts, and receivables.

We are not the first to use UCC filings in academic research.[11] However, to the best of our knowledge, we are the first to have assembled a UCC dataset of this size and scope and to use these data to study the effects of bank failures on firm performance.

Some mass marketing companies in the US collect and sell UCC data. However, these data provided by marketing companies have several limitations. First of all, these data are limited in how far back in time they go. Marketing companies that collect UCC data do so to create mailing lists that can be used for mailing campaigns (i.e. brochures, etc.), cold calling, or other marketing purposes. As a consequence, these data collection efforts tend to be focused on very recently filed filings, where the the business contact details are more likely to still be valid. Some marketing companies do provide historical data but the furthest back we have seen these data go is 2006 which makes assessing pre-trends for financial shocks occurring the financial crisis difficult. Second, the UCC data sold by marketing companies are far from being complete. The leading data provider for marketing UCC data offers data from 36 million UCC records for all US states going back to 2006. However, our data over the same time period but for only five states—representing about 1/3 of US GDP—comprise over 26 million records suggesting that UCC records sold by marketing companies reflect less than half of the universe of UCC filing. Furthermore, these companies do not provide details about their data collection efforts, raising concerns that the data they sell may be selected on dimensions unbeknownst to the researcher.

We hand-collect over 36 million UCC filings from the Secretary of State offices in California, Florida, New York, Texas, and Colorado. Though these are only five states, taken together they account for 39% of US GDP. Collection efforts varied from state to state: for most states we scraped the public-facing UCC database that is traditionally used by banks and other creditors to check for existing liens on potential borrowers.[12] For states for which web scraping was not possible, we negotiated directly with the Secretary of State office to obtain the filings. The final corpus of filings we assemble, across all five states, spans the period from 1993 to 2023, with the majority of filings coming from the period 2000 to 2023.

It is difficult to estimate precisely the fraction of loans that result in a UCC filing. However, small business loans are predominantly collateralized and since the benefits of securing a loan with a UCC filing are significant, we believe that the fraction of loans that result in a UCC filing is

[11] See e.g. Edgerton (2012) uses a sample of California UCC filings filed between 2005 and 2011 to study how lender distress during the GFC affected business invesmtent. Murfin and Pratt (2019), Ma, Murfin, and Pratt (2022), and Gopal and Schnabl (2022) use data from a commercial marketing data provider which relies on multiple data sources, including UCC filings.

[12] To respect fair use policies, we limited our scraping efforts to outside of business hours. Given this limitation and the very large number of documents that needed to be downloaded, our web scraper ran for multiple months in 2023.

substantial. Berger and Udell (1998) analyses data on small businesses with fewer than 500 employees in the 1993 National Survey of Small Business Finances (NSSBF) and finds that 92% of the bank debt obtained by these firms is collateralized. Chodorow-Reich, et al. (2020) documents that firms with less than $50 million in assets virtually never obtain unsecured credit, a finding that applies both to credit lines as well as to term loans.[13] Luck and Santos (2023) documents that less than 4% of loans made to small businesses in the US are unsecured.

Furthermore, data from the Federal Reserve Small Business Credit Survey (SBCS) shows that within the group of employer firms that report having outstanding debt, only 17% do not use collateral to secure their debt. Nearly half of small employer firms report using business assets to secure business debt and one-third report using personal assets to secure business debt (Federal Reserve 2020).[14] Requiring collateral for small business loans is particularly common among smaller banks (banks with less than $10 billion in assets). In the FDIC Small Business Lending Survey, 82% of small banks indicate that they required some form of collateral for their main small business loan product. For larger banks, roughly half require collateral for their main small business loan product (FDIC 2018). The most cited type of collateral required by both small and large banks is business assets and equipment.[15]

While a very large fraction of small business lending is collateralized, it is possible that some of this collateral is real estate, which is a form of collateral not governed by UCC filings but rather by mortgages which are recorded as liens on deeds in property registries. Luck and Santos (2023) use the Federal Reserve's FR Y-14Q data to show that about one-fifth of small business loans are secured by real property. Gopal and Schnabl (2022) rely on this finding, in conjunction with the finding in Luck and Santos (2023) that only 5% of small business borrowing is done via credit cards, to conclude that UCC filings cover approximately 73% of all small business lending in the US.

Furthermore, we conduct a data validation exercise—described in more detail below—which shows that our data track very closely with small business lending data reported in the Community Reinvestment Act (CRA) dataset published by the Federal Financial Institutions Examination Council (FFIEC). All banks with assets above $1 billion are mandated to report their small business lending data to the FFIEC, which publish the data aggregated at the Census tract level. This regulatory dataset is regarded as highly accurate and comprehensive due to its mandatory nature

[13] See Table 6 in Chodorow-Reich, et al. (2020) for a detailed breakdown.

[14] The high incidence of secured debt in small business lending appears to be a longstanding practice. The National Federation of Independent Business (1983) report that 78% of the total volume of small business loans were secured by collateral in the early 1980s.

[15] See Figure 5.4 in FDIC (2018) for a detailed breakdown of required collateral types.

and standardized reporting requirements. Importantly, Greenstone, Mas, and Nguyen (2020) show that the CRA data encompass approximately 86% of all loans under $1 million originated in the US, which highlights the CRA dataset's broad, near-complete representation of small business lending activity in the US.

3.2 Firm and bank-level data

Our business-level data is provided by Dunn & Bradstreet (D&B), a commercial data and analytics provider. D&B aims to collect information on the universe of US businesses by drawing from an extensive range sources, including state secretaries of state, Yellow Pages, court records, credit inquiries, licensing data, and direct telephone contact with businesses.[16] The quality of D&B data has been reviewed by several studies (Neumark, Wall, and Zhang 2011; Haltiwanger, Jarmin, and Miranda 2013; Barnatchez, Crane, and Decker 2017; Crane and Decker 2020) who find that while these data are not without limitations, they provide a comprehensive and accurate picture of the US business landscape, both in the cross section and across time. For example, Neumark, Wall, and Zhang (2011) compare employment data in D&B against employment data in the Bureau of Labor Statistics' (BLS) Current Emploment Survey (CES) at the county-by-industry level and find a correlation of 0.99. The same comparison against the BLS's Quarterly Census of Employment and Wages (QCEW) dataset generates a correlation of 0.95. Barnatchez, Crane, and Decker (2017) demonstrate that once very small businesses (<10 employees), very large businesses (>1,000 employees), and educational entities are excluded, the D&B dataset agrees well with the US Census Bureaus' County Business Patterns (CBP) data and the BLS's QCEW data. Nevertheless, we are careful to follow the advice and rules-of-thumb provided by Neumark, Wall, and Zhang (2011) and Crane and Decker (2020) while working with these data.

We obtain detailed information on bank balance sheets from the US extract of Orbis Bank Focus (formerly Bankscope), maintained by Moody's Analytics and Bureau van Dijk. Bank Focus is a global database providing financial and ownership information on banks worldwide. It compiles data from regulatory filings, annual reports, stock exchange disclosures, and direct institutional correspondence. The data include bank balance sheet information as well as information on bank structures and ownership.

[16] In the D&B's dataset, a "business" is broadly defined to include private for-profit entities, nonprofit organizations, and government agencies, encompassing a wider scope than many official business statistics such as the US Census Bureau's County Business Patterns (CBP) and Nonemployer Statistics (NES) and the Bureau of Labor Statistics (BLS) Quarterly Census of Employment and Wages (QCEW).

We obtain data on bank failures directly from the FDIC which provide data on failures of over three thousand FDIC-insured banks since 1934. The FDIC data include information on the date of failure, the bank's location, the bank's assets at time of failure, and the FDIC's resolution strategy.[17]

3.3 Data linkage

As there is no unique identifier that links UCC filings to establishments, we use a combination of name and location matching to link UCC filings to our establishment-level data. We first create text embedding for the 46 million unique business name strings present in our UCC filing corpus and the D&B dataset using the "text-embedding-3-large" embedding model from OpenAI. We make sure to account for both the registered business name as well as the "trading style" or "doing-business-as" name of the businesses—for example, "Starbucks Coffee" is the trading style name corresponding to the registered business name "Starbucks Corporation".

The embedding model we use transforms each business name into a high-dimensional vector of numbers. OpenAI has trained this embedding model on vast corpora of text to capture semantic relationships, allowing it to represent similar concepts with vectors that are close in the embedding space.[18] Using these embeddings allows us to compare business names not just based on so-called "fuzzy" matches, but on their overall semantic similarity. For example, "Kentucky Fried Chicken" will have a similar embedding not just to strings like "Kentucky Fried Chicken Corporation" but also to strings that are completely different letter-wise but which refer to the same entity, such as "KFC". This capability is crucial for our matching process, as it allows us to overcome common issues—such as abbreviations, acronyms, and alternative phrasings—that have traditionally plagued large-scale data linkage efforts. To the best of our knowledge, we are among the first to use these modern embedding techniques to link data in the applied economics literature.

To operationalize our data linking strategy, we need to compute semantic similarity for each candidate pair of business names across our three datasets (UCC, DnB, and Lightcast). However, with approximately 38 million embeddings for DnB and 12 million embeddings in UCC, a naive

[17] The data also include bank failures of institutions insured by the Federal Savings and Loan Insurance Corporation (FSLIC), the Resolution Trust Corporation (RTC), the Savings Association Insurance Fund (SAIF), and the Bank Insurance Fund (BIF), which were predecessor or related insurance entities that have since been merged into or replaced by the FDIC.

[18] The embedding technology employed in this study is fundamentally the same as that which underpins modern Large Language Models (LLMs) such as GPT-3, GPT-4, and BERT. These models rely on a technique called "self-supervised learning" to create dense vector representations of words and phrases that capture nuanced semantic relationships.

pairwise calculation for just these two datasets would require computing cosine similarity for 456 trillion possible pairs—which is computationally very costly. To overcome this challenge, we leverage the Faiss (Facebook AI Similarity Search) index, a technology developed by Meta AI Research. Faiss is designed for efficient similarity search and clustering of dense vectors, allowing us to perform nearest neighbor search in high-dimensional spaces with billions of vectors. This approach dramatically reduces the computational complexity of our matching process. Given the scale of our data and the high-dimensional nature of our embeddings, we utilized specialized GPU hardware to execute this step efficiently.

The next step of our data linkage pipeline involves geocoding the 73.2 million unique address strings contained in our data. We achieve this using the application programming interface (API) provided by the US Census Bureau which takes as an input an address string and returns structured geographical information including latitude and longitude. We then take the addresses that the US Census API was not able to resolve and pass them to an API provided by Amazon Web Services (AWS), which powers Amazon's package delivery operation. We take care to geocode both the physical address as well as the mailing address for each establishment in our data. This is important because the UCC stipulates that a UCC filing must include the mailing address of the borrower, which may differ from the physical address of the business.[19]

With the name embeddings and geographic coordinates in hand, we are able to create a large candidate set of 82 million possible matches where a possible match is defined to have a name similarity score of at least 0.85 and a geographic distance below 10 kilometers. We stratify this sample into cosine similarity and geographic distance deciles for a total of 100 strata. We then randomly draw ten candidate matches from each of these strata and manually label them as being a match or not a match. With the labelled data, we train a LASSO logistic regression that predicts the probability of a match based on features of the candidate match. These features include the geographic distance, the cosine similarity of the business names, as well as the total number of addresses associated with a given business name (e.g. "Burger King" may have thousands addresses). This latter feature is informative about the size of the business as well as the likelihood that the business is a chain or franchise. We also include the natural log as well as squared and all possible interaction terms of the cosine similarity and geographic distance and address count variables for a total of 54 possible predictors that the LASSO model can select from. The LASSO approach performs both predictor selection and regularization, effectively shrinking the coefficients of less important predictors to zero.The aim of this approach is to balance the trade-off between model simplicity and predictive accuracy.

[19]§9-516 of the UCC stipulates that a UCC filing must include the mailing address of the debtor to be legally effective.

As a final step, we utilize an active learning strategy with uncertainty-based sampling, iteratively selecting and manually labeling the most uncertain predictions to efficiently improve the model's performance. The result is a matching accuracy model that predicts the probability of a match between a UCC filing and an establishment in our establishment-level data. We use this model to predict the probability of a match for the 82 million candidate matches. Table 1 summarizes the final matches.

The linking procedure we use to link UCC filings to bank-level data follows a similar approach as the one outlined here for linking UCC filings to our firm-level data. The UCC-to-bank linkage is described in more detail in Appendix A.1.

4 Descriptive statistics

Our final panel consists of 1,889,782 unique firms from 1990 to 2023. To be included in the panel, a firm must have filed one or more UCC filings that we were able to successfully match to the firm and to a bank. As we are interested in information about the existence of a banking relationship, we include not just "UCC-1 Initial Financing Statements"—which reflect the creation of a new credit arrangement—but also other types of UCC filings such as "UCC-3 Continuation Statements" and "UCC-3 Amendment Statements". UCC-3 Continuation Statements extend the legal effectiveness of an existing financing statement for an additional number of years, indicating the continuation of a previously established credit relationship. UCC-3 Amendment Statements modify existing filings and can provide updated information about ongoing banking relationships, such as changes in the debtor's information or modifications to the collateral. To ensure that a UCC filing reflect credit arrangements, we only consider. Other filing types such as judgment liens and IRS tax liens are excluded as these reflect court judgments against a debtor or tax delinquencies, respectively.

To reduce the potential for violations of the "hidden variation in treatment" assumption in SUTVA, we exclude firms that form relationships with more than one bank that ultimately fails. Doing so also helps more clearly pin down the effect of a particular bank failure on firms in a later exercise where we analyze the firm-level effects of bank failure across different failed banks. We also drop firms whose first contact with a failed bank occurs after the bank failure, as these cases likely represent situations where both the failed bank and the acquiring bank are listed as creditors in the UCC financing statements and where our matching process erroneously linked a firm to the failed bank instead of the acquiring bank. These data cleaning procedures result in the removal of less than 9% of the original observations.

Table 2 presents summary statistics for our sample. The average firm in our dataset has 11.8 employees. The employment distribution is highly skewed, with a median of 3 employees and a 90th percentile of 20 employees. The median firm age at the time of its first UCC filing is 4 years, suggesting that many firms in our sample are relatively young when they first secure bank financing financing. This fact aligns with the findings in Robb and Robinson (2014) who show that startup firms in the US predominantly rely on banks for external debt during their first few years of operation.

For the subsample of firms banking with failed banks—the treated sample—we observe that the average duration of the banking relationship prior to bank failure is 7.6 years, with a median of 6 years. This statistic agrees with Berger and Udell (1998) who analyses the 1993 National Survey of Small Business Finance (NSSBF) to show that the average banking relationship for US firms with fewer than 500 employees is 7.77 years.

Approximately half of the firms in our sample are concentrated in the business services, health services, eating and drinking places, construction, real estate, and wholesale trade industries as categorized by two-digit Standard Industrial Classification (SIC) code. Figure 1 illustrates the geographic distribution of our sample, showing that while we have data on businesses from virtually all US counties, the majority come from California, Colorado, Texas, Florida, and New York, the states in which the UCC filings in our dataset were filed.

It's important to note some limitations of our data. Firstly, we do not directly observe when a firm ceases operations; instead, we infer this from a firm's exit from the D&B dataset. The D&B data collection and data quality management processes are exhaustive (see Section 3.2) and should accurately reflect cessation of business operations with the exit of a firm from the D&B panel. However, to further test the correlation between dataset exit and business operation cessation we assess how many of our UCC filings match to firms after they exit from the D&B panel. We find that less than 5% of our UCC filings match to a firm more than a year after it has exited from the D&B dataset, suggesting that exit from the dataset does correlate very strongly with the operational "death" of that firm.[20]

Table 6 provides information for the top 25 bank failures, as determined by how many firms have a credit relationship with the bank in our data. The distribution of credit relationships across failed banks is highly skewed with the top 25 banks representing 75.92% of the failed bank credit relationships in our dataset and the remaining 154 banks making up the balance.

[20] In this analysis we only consider "UCC-1 Initial Financing Statements" as other UCC filings such as credit amendment liens, judgment liens or IRS tax liens may well be filed against a business after it has *de facto* stopped operating. See Appendix A.4 for more details on this particular analysis.

5 Empirical framework

In this section, we articulate our empirical strategy for estimating the causal effects of bank failures on firm performance. Our two key outcomes of interest are firm survival (i.e., whether the firm is still operating in a given year) and employment headcount. Our panel data allow us to track firms over time; however, in our survival specification, the panel is balanced on the right, meaning firms never exit the panel, though their survival status permanently changes if and when they exit. This section first introduces our baseline specification, then discusses identification, and closes with a subsection on extensions and robustness exercises.

5.1 Baseline design

Since we study multiple bank failures occurring at different points in time, our empirical setting characterized by staggered treatment timing. The data track firm-level outcomes $y_{i,t}$ across $i = 1, ..., N$ firms over $t = 1, ..., T$ periods, where each period is a calendar year. Whether a firm has been exposed to bank failure is captured by the indicator $D_{i,t} \in \{0, 1\}$. Treatment is assumed to be absorbing, meaning that once a firm's bank fails, this status persists—formally, $D_{i,s} \leq D_{i,t}$ for $s < t$. Each firm i has an initial year of bank failure exposure p_i, with $p_i = \infty$ for "never treated" firms whose bank never fails.

We partition firms into cohorts $g \in 0, 1, ..., G$ based on the calendar year t in which they first experience bank failure. Cohort $g = 0$ comprises firms whose banks never fail during our sample period. We denote the time period of bank failure for group g as p_g. Using the potential outcomes framework of **?**, let $y_{i,t}(0)$ represents the outcome of firm i at time t if its bank never fails. Similarly, $y_{i,t}(p_i)$ represents the outcome for firm i if its bank fails at time $p = p_i$.

The group-specific and dynamic average treatment effect on the treated (ATT) for cohort g at horizon h (i.e., the effect h periods after bank failure) is:

$$\tau_h^g = \mathbb{E}\left[Y_{i,p_g+h}(p_g) - Y_{i,p_g+h}(0) \,\middle|\, p_i = p_g \right] \tag{1}$$

Estimating these treatment effects in a staggered treatment setting introduces complexity to our estimation strategy, as it creates multiple cohorts of treated units (firms whose banks failed at different dates), each with potentially different treatment effects.[21] To address this issue, we use

[21] Estimating a dynamic difference-in-differences (DiD) regression in this setting using a standard two-way fixed effects (TWFE) estimator is known to generate biased estimates because TWFE implicitly makes "forbidden comparisons" between newly treated units and previously treated units, which may still be experiencing treatment

a local projection difference-in-differences (LP-DiD) approach proposed by Dube, et al. (2024) as our main specification. Our baseline regression takes the following form:

$$y_{i,t+h} - y_{i,t-1} = \beta_h \Delta D_{i,t} + \delta_t^h + \gamma X_{i,t} + e_{i,t}^h, \quad \forall h \in \{0,\dots,h\} \tag{2}$$

where δ_t^h is a time fixed effect, X is vector of additional controls, and $e_{i,t}^h$ is the error term, which we cluster at the county and bank-level in our baseline design.

For each horizon h, we restrict the sample to observations that are either:

$$\begin{cases} \text{newly treated by bank failure} & \Delta D_{it} = 1, \\ \text{or `clean control'} & D_{i,t+h} = 0 \end{cases}$$

Under two identification assumptions discussed further below (conditional no anticipation and conditional parallel trends), the β_h coefficient from equation (2) consistently estimates a convex combination of all group-specific effects τ_g (i.e. a weighted average of the effects across all bank failure cohorts). Concretely:

$$\mathbb{E}(\hat{\beta}_h) = \sum_{g \neq o} \omega_{g,h} \tau_h^g \tag{3}$$

where the weight for each group-specific effect is:

$$\omega_{g,h} = \frac{N_{CCS_{g,h}}[n_{gh}(n_{c,g,h})]}{\sum_{g \neq o} N_{CCS_{g,h}}[n_{g,h}(n_{c,g,h})]} \tag{4}$$

where $N_{CCS_{g,h}}$ is the count of observations in the 'clean control' sample for group g at horizon h, such that $n_{g,h} = N_g / N_{CCS_{g,h}}$ is the fraction of treated units in the $CCS_{g,h}$ subsample. Finally, $n_{c,g,h} = N_{c,g,h} / N_{CCS_{g,h}}$ is the share of control units in the $CCS_{g,h}$ subsample[22].

effects from their own treatment start (de Chaisemartin and DHaultfuille 2020; Goodman-Bacon 2021; Callaway and SantAnna 2021; Sun and Abraham 2021; Athey and Imbens 2022; Borusyak, Jaravel, and Spiess 2024). For surveys of recent advances in the dynamic DiD literature, see deăChaisemartin and DHaultfuille (2023) and Roth, et al. (2023).

[22] We also consider alternatives to the 'variance-weighted ATT', namely the 'equally-weighted' ATT, along with a host of alternative robustness exercises discussed below.

5.2 Identification

Conditional parallel trends: One identification concern in estimating the effect of bank failure on firm performance is that there may be confounding shocks such as region or industry-specific events that affect both bank failure and firm outcomes. For example, during the US subprime crisis, which coincided with the majority of the bank failures in our data, not all regions within the US were equally affected. During the crisis, there was substantial geographic variation in employment, house price changes, household net worth, and foreclosure rates[23]. It is possible therefore that region-specific or industry-specific shocks that simultaneously raise the probability of bank failure and firm closure could lead to a violation of the conditional parallel trends identification assumption:

$$E\left[Y_{i,t}(0) - Y_{i,1}(0) \,\middle|\, p_i = p; X_{i,t}\right] = E\left[Y_{i,t}(0) - Y_{i,1}(0) \,\middle|\, X_{i,t}\right] \tag{5}$$

for all $t \in \{2, ..., T\}$ and all $p \in \{1, ..., T, \infty\}$. This condition requires that, in the counterfactual scenario where treated firms did not experience bank failures, their outcomes would have evolved in parallel with those of control firms, conditional on covariates.

This issue is particularly concerning for smaller, regional banks whose solvency is closely tied to the economic conditions of the region in which they operate. To address this, we include a triple interaction of county, year, and industry fixed effects in our regression, where an industry is defined by a firm's two-digit SIC code. With this set of fixed effects, we interpret our β^h coefficients as the estimated difference in outcomes between firms that bank with failed banks and firms that do not bank with failed banks, within the same county, year, and industry, while holding all time-invariant characteristics of the firms constant.

Conditional no anticipation: Another identification threat comes from the possibility that treated firms begin to perform poorly before their bank's failure. Such a pattern could emerge, for example, if poor firm performance is the cause of bank the failure or if firms are able to anticipate bank failure and systematically either switch away from, or towards, the failed bank (i.e. selection prior to treatment). Such anticipation effects would violate the no anticipation assumption:

$$\mathbb{E}[Y_{i,t}(p) - Y_{i,t}(0)|X_i] = 0, \text{ for all } p \text{ and } t \text{ such that } t < p \tag{6}$$

[23]See Kochhar and Gonzalez-Barrera (2009), Mian, Rao, and Sufi (2013), and Gertler and Gilchrist (2018) for a discussion of geographic variation in economic outcomes during the subprime crisis.

This condition requires that firms' outcomes are not affected by future bank failures before they occur, ruling out both anticipatory responses by firms and reverse causality from firm performance to bank failures.

There are multiple reasons that it is unlikely that poor firm performance lead to the bank failures studied in our sample. First, the existing literature documents that banks primarily failed during our sample period due to steep losses on mortgage-backed securities (MBS)[24], high loan portfolio concentration in commercial real estate loans, or high reliance on short-term debt[25]

Furthermore we collect 120 regulatory post-mortem reports (material loss reviews) conducted by the FDIC's Office of Inspector General. These reports, required when a bank failure results in a loss of $50 million or more to the FDICs deposit insurance fund, investigate the reasons for the bank's failure and assess the FDICs supervision. Using an LLM, we extracted key sections detailing the primary reasons for the bank's failure[26]. Our review of these sections indicates that not a single one of the 120 material loss reviews cites corporate loans as a reason for the bank's failure.

We also consult the FFIEC Call Report data to examine the primary sources of financial losses to failed banks prior to failure. A bank Call Report (Consolidated Report of Condition and Income) is a mandatory quarterly filing for all regulated US financial institutions that contains comprehensive information on the institution's financial condition. Using these data, we find that loan charge-offs for ADC and other real estate loans increase sharply before failure, while C&I loan charge-offs rise only marginally[27].

Another identification concern is the possibility that firms may anticipate bank failure and switch their banking relationships in advance. This would create a selection bias if financially weaker firms are more likely to bank with failed institutions. To address this, we conduct an event study analysis of firm switching behavior. Our analysis shows no statistically significant increase in the likelihood of switching banks in the years preceding bank failure[28]. We also utilise alternative assignment mechanisms to treatment, including that firms must have formed a relationship with the failed bank 2 or more years prior to the event. This analysis allows for firms who switch out prior to failure to remain in the treated group, yet shows little change relative to our baseline estimates.

[24] Erel, Nadauld, and Stulz (2014) and Antoniades (2015) provide an analysis of MBS-related losses and their impact on banks.

[25] Furlong and Knight (2010), Cullen (2011), and Cole and White (2012) explore the reliance on commercial real estate loans and short-term debt among failing banks.

[26] For full details on the method of review and section references, see Appendix X.

[27] Appendix X provides a detailed summary of charge-off patterns based on Call Report data.

[28] See Appendix C for detailed analysis of switching behavior.

In support of this, Correia, Luck, and Verner (2024) and Martin, Puri, and Ufier (2018) find no evidence that depositors respond to increases in bank failure risk. Small businesses, which typically maintain a relationship with only a single bank for both deposits and loans, also appear to not switch banking relationships before bank failure.

5.3 Extensions and robustness

To validate our main findings, we conduct several extensions and probe the robustness of our results to alternative specifications and assumptions.

Matching approach: Our baseline approach is to utilize a matched sample of our full data that is balanced evenly between control and treatment firms across these dimensions. However, we also verify that our results are robust to using the full sample of our data and employing granular controls to ensure the variation being exploited is 'within' county × industry × firm-size quartiles × calendar year.

Alternative treatment definitions: While there are reasons to believe that small businesses are generally unresponsive to increases in the probability that their bank fails—as discussed in detail in the previous section—concern may remain that prior to bank failure healthy firms switch their banking relationship away from the soon-to-be-failed bank and that unhealthy firms may switch their relationship towards the soon-to-be-failed bank. Both of these types of sorting would artificially worsen the quality of treated firms and bias our estimates away from zero.

To address these concerns, we check whether our estimtes are robust to changing our definition of treated firms. We now define a treated firm as one that established a credit relationship with a bank at least two years prior to its failure. Importantly, we maintain this treatment status even if the firm subsequently switches to another bank. Additionally, we exclude from our treatment group any firms that initiated a credit relationship with the failed bank within the two-year window preceding its failure. We choose a two-year window in line with findings from Correia, Luck, and Verner (2024) that show that for the vast majority of banks, bank failure is not predictable based on bank financials three years prior to failure.[29] This refined definition mitigates the potential bias from firms anticipating failure and preemptively switching banks.

This new treatment definition is likely to produce more conservative estimates of the effect of bank failure on firm outcomes. By including firms that switched away from the failed bank in our treatment group, we are potentially incorporating firms that were less affected by the failure. The

[29] See Figure 6 Panel C in Correia, Luck, and Verner (2024) which shows that even very unhealthy financials of banks do not substantially increase the risk of bank failure at a three-year time horizon.

same argument applies to our exclusion of firms that switch their banking relationships from a healthy bank to the ultimately-failed bank in the three years prior to the failure. Consequently, our estimates are likely to represent a lower bound of the true effect. This approach trades some potential upward bias for a more robust and defensible estimate, albeit one that may understate the full impact of bank failure on the most affected firms.

Pre-treatment dynamics: To address concerns about potential endogenous treatment timing, we augment our baseline specification with multiple lags of our outcome variables:

$$y_{i,t+h} - y_{i,t-1} = \beta_h \Delta D_{i,t} + \delta_t^h + \sum_{k=1}^{K} \gamma_k y_{i,t-k} + e_{i,t}^h \tag{7}$$

where K is the total number of lags. This specification helps control for any pre-existing trends in firm-level outcomes that might be correlated with bank failure timing or selection into particular banking relationships. For example, a large number of firms exiting could induce a bank failure. While we don't see any evidence of this kind of reverse causality, we nonetheless implement this adjustment to confirm the robustness of our baseline approach.

Alternative estimators: Following recent methodological advances in the difference-in-differences literature, we consider alternative estimators suitable for a staggered-treatment difference-in-difference approach. These include the two-stage estimator as developed by Gardner (2021), and designs based on Borusyak, Jaravel, and Spiess (2024), Callaway and SantAnna (2021).[30]

Alternative weighting and standard error clustering: Our baseline estimation produces ATTs that are variance-weighed across treatment cohorts, but we verify robustness to alternative weighting schemes including equal weighting across firms and weights based on firm size. The statistical inference in our main results clusters standard errors at the bank level, but we verify robustness to alternative clustering approaches including at the state and industry levels.

Placebo tests: To further validate our findings, we conduct a series of placebo tests. These include reassigning bank failures randomly across healthy banks, shifting actual failure dates to earlier periods, and constructing placebo treatments using healthy banks with similar characteristics to failed institutions.

[30] Due to the size of our data and the large number of cohorts/treatment years, at least one of the benchmark approaches to staggered difference-in-difference is infeasible, namely the Sun and Abraham (2021) estimator - which utilizes high dimensional cohort controls and interactions within a single estimating equation. Another reason we prefer local-projection difference-in-difference introduced by Dube, et al. (2024) is the computational ease of distributing these high-dimensional cohort features across multiple estimating equations.

6 Results

6.1 Firm survival

We first discuss the impact of bank failure on firm survival. Since the D&B microdata does not explicitly record the dates when firms cease to operate, we use a firm's disappearance from the D&B panel as a proxy for operational cessation. This approach appears valid: our linkage of UCC filings to D&B records reveals that only 16% of filings match with firms outside their D&B panel presence years, suggesting that entry and exit from the D&B dataset largely aligns with firm "births" and "deaths". For more details on this exercise see Figure .

We construct a firm-year panel such that firms enter our panel in the year they first appear in the D&B dataset, but we retain all firms in our panel until 2022, regardless of when they exit the D&B dataset. To denote a firm's operational status, we use an indicator variable defined as: $y_{i,t} = \mathbb{1}\{t \leq \tau_i\}$ where τ_i is the last year firm i is observed in the D&B data.[31]. In our main regression specification, we define a treated firm as one that formed a credit relationship with a bank that subsequently failed and did not establish relationships with any other banks between the formation of this relationship and the bank's failure—we change this restriction for subsequent estimations and discuss the results below. To address potential violations of the no hidden variation in treatment assumption of SUTVA, we further restrict our treated sample to exclude firms that formed relationships with multiple failed banks prior to their failure, which results in dropping less than 1% of the sample. A treated firm enters treatment in the year of the bank's failure and remains treated indefinitely—i.e. treatment is fully absorbing.

A UCC filing may arise when a firm obtains credit from a non-bank. Since our treated firms are by definition firms that have obtained credit from a bank, we restrict our control group in our headline specification to only include firms that have obtained credit from a bank at least once. This sample selection also allows us to control for bank size, which we proxy with total book value of the bank's loans. Other controls include county-by-year-by-industry fixed effects—where industry is defined at the two-digit SIC code—as well firm age and firm age squared. Figure 7 plots the estimated treatment coefficients from our headline regression specification with a time horizon ranging from 5 years prior to 10 years after the bank failure year. The failure event marks the beginning of a clear differential effect on the survival rates of treated firms relative to their untreated counterparts. One year after the bank failure, treated firms are 3 percentage points less likely to still be operating. The five-year and ten-year survival rates of treated firms are 6.7 an

[31] Similar empirical designs where a LP-DiD is used to study a binary outcome variable of this nature are used, for example, by Goda and Soltas (2023).

7.9 percentage points lower. Figure 6 translates these percentage point effects into percentage increases relative to the survival and death rates of never treated firms. One year after a bank failure, treated firms are 46.3% more likely to have ceased operating. The 5-year and 10-year death rates for treated firms are 44.3% and 34.4% percent higher compared to untreated firms.

The DiD estimates for the 5 year period preceding bank failures show no statistically significant differences from zero. This indicates that, conditional on our controls, there is no observable divergence in survival rates between treated and untreated firms prior to the bank failure year at the annual level. These pre-trend results are relevant to the question of whether firm performance prior to bank failures might have influenced the likelihood of bank failure. The flat estimates we observe align with our earlier discussion of the primary drivers of bank failures during this period, which were largely unrelated to commercial lending portfolios. We complement these DiD estimates with an event study analysis detailed in Appendix B.

6.2 Employment

Following our analysis of firm survival, we now turn to the impact of bank failures on firm-level employment. We employ the same headline regression specification as used for survival analysis, with employment growth as the dependent variable and run on the subset of firms that survive to each time horizon analyzed.

Figure 7 presents the estimated effects of bank failure on firm-level employment growth over time. The results reveal a clear negative impact on employment growth for treated firms compared to untreated firms. This effect becomes evident immediately following the bank failure and persists over the long term.

In the first year following the bank failure, we observe a statistically significant decrease in the employment growth rate of approximately 0.8 percentage points for treated firms relative to untreated firms. This negative effect persists over longer horizons, with the 5-year employment growth rate showing a decrease of about 3.7 percentage points. The effect continues to widen for longer horizons, with the 10-year employment growth rate showing a decrease of 6 percentage points for treated firms relative to untreated firms. The persistence of these effects over a decade suggests that the disruption of banking relationships has lasting consequences for small business employment growth.

The pre-trend analysis, covering the 5-year period preceding bank failures, shows no statistically significant differences between treated and untreated firms. This aligns with our earlier discussion on firm survival and further supports the notion that bank failures during this period were largely unrelated to the performance of their commercial lending portfolios.

Figure 8 shows the estimated impact on employment growth scaled by the employment growth experienced by never treated firms across our entire dataset. The 1-year employment growth rate for treated firms is 13% lower compared to untreated firms. This gap persists over longer horizons, with the 5-year employment growth rate being 24.8% lower and the 10-year employment growth rate being 29.14% lower for treated firms. The apparent decrease in the relative effect over time can be attributed to two factors. As depicted in Figure 8, divergence in employment growth begins to accelerate five years after the bank failure. We believe this pattern to be driven by the fact that the small businesses in our sample (median employment: 3) only sporadically change their headcount, meaning that differences in hiring patterns do not emerge immediately after the bank failure but only some years afterwards. The results presented in Figure 8 demonstrate the long-term scarring effects of bank failures on small business employment growth, with the impact becoming more pronounced over time and indicating that these events can hinder firms' ability to hire and grow for years after the initial shock.

6.3 Robustness

Alternative treatment definition: As discussed in Section 5.3, we address potential bias from strategic bank switching by redefining treated firms as those with credit relationships established at least two years before bank failure, regardless of subsequent switching behavior. We exclude firms that initiated relationships with the failed bank within two years of failure. This definition, based on the limited predictability of bank failures beyond a three-year horizon, produces conservative estimates by including firms that may have successfully switched away from failing banks while excluding those that may have switched to them. Figure 5 shows that the estimated coefficients using this alternative treatment definition are not statistically different from the estimates obtained from our headline regression.

Alternative estimators: We employ several alternative estimators suitable for staggered treatments with heterogeneous effects (discussed in more detail in 5.3). Figure 4 shows that estimates of our headline specification using estimators from Borusyak, Jaravel, and Spiess (2024), Callaway and SantAnna (2021). and Dube, et al. (2024) produce similar results, while the Gardner (2021) two-stage estimates suggest even larger effects. In contrast, standard TWFE estimates are substantially attenuated—a not-so-surprising result given known limitations of TWFE estimators in settings with staggered treatments and effect heterogeneity.[32]

Two natural experiments: To strengthen our identification, we examine two cases of bank failure that provide more exogenous variation in the timing of the bank's failure: Park National Bank (part of FBOP Corporation) and Colonial Bank. These cases offer instances where the bank's failure was unexpected by firms and plausibly exogenous to the health of the banks' business borrowers.

Park National Bank's failure in October 2009 was a direct result of the FDIC's use of its "cross-guarantee" powers under the Financial Institutions Reform, Recovery, and Enforcement Act of 1989 (FIRREA). Congress granted these powers to the FDIC to address moral hazard concerns within bank holding groups. Previously, if one bank in a holding company failed and incurred losses to the FDIC's insurance fund, the other affiliated banks were not held liable. This lack of accountability created moral hazard issues, as banks could shift bad assets to a single institution, exploiting the FDIC insurance fund when that bank ultimately failed. The cross-guarantee provisions were designed to prevent such practices by making all banks within a holding group jointly liable for losses, thereby incentivizing better risk management across the entire organization.

Colonial Bank's failure in 2009 resulted from a complex chain of events that were largely unrelated to the financial health of its business borrowers. Initially, the bank incurred losses in its acquisition, development, and construction (ADC) loan portfolio. In response to these losses, the bank's management applied for funds from the Troubled Asset Relief Program (TARP), which subjected the bank to increased regulatory scrutiny. This heightened examination uncovered a significant fraud scheme involving Colonial Bank and its mortgage originator counterparty, Taylor, Bean & Whitaker (TBW). The fraud entailed TBW susbtantially inflating the quantity and quality of mortgages it sold to Colonial, along with unauthorized credit limit breaches that were facilitated by complicit employees within Colonial Bank.

Following the discovery of the fraud, the FBI conducted a raid on both Colonial Bank and TBW, and Colonial Bank subsequently lost access to $875 million held in an escrow account with TBW. This sudden loss of liquidity, prompted the FDIC to close Colonial Bank just 11 days after the raid. The rapid sequence of these events indicates that the bank's failure was primarily driven by internal fraud and mismanagement rather than the financial conditions of its business borrower.

[32] See (de Chaisemartin and DHaultfuille 2020; Goodman-Bacon 2021; Callaway and SantAnna 2021; Sun and Abraham 2021; Athey and Imbens 2022; Borusyak, Jaravel, and Spiess 2024; Dube, et al. 2024) for more detailed discussions of TWFE limitations in staggered treatment settings.

Due to data limitations, specifically the lack of sufficient observations around the time of Park National Bank's failure, we are unable to estimate the full dynamic LP-DiD regression specification as in our main analysis for this case. Instead, we estimate a pooled effect by collapsing the pre-treatment and post-treatment periods into single periods, effectively capturing an average treatment effect of the bank failure on firm survival. The results, presented in Table 5, show that firms connected to Park National experienced a statistically significant decrease in survival probability of approximately 1.8 percentage points following the bank's failure. For Colonial Bank, the estimated effect is larger, with a decrease of 6.9 percentage points. When considering all bank failures collectively, the average effect is a decrease of 4.6 percentage points in firm survival probability. These findings align with our earlier results, underscoring the significant negative impact of bank failures on firm survival and highlighting the importance of stable banking relationships for small businesses.

6.4 Heterogeneity across bank failures

The results discussed so far speak to the average effects of bank failure across the 179 bank failures included in our analysis. However, failed banks vary both in terms of their geographic and industry concentration as well as their size and the size of the bank that acquires the failed bank post failure. To assess how the treatment effects discussed so far vary across bank failure events, we conduct a bank-by-bank analysis. Specifically, we repeat our regression analysis separately for each of the 25 largest banks in our dataset, where bank size is measured by the number of firms that have a banking relationship with each institution.

Table 6 provides a summary of these 25 banks including their failure date and the size of these banks' balance sheets. All but one of these banks was resolved via the P&A resolution method—discussed in more detail in Section 2—meaning that the failed bank's assets and liabilities were transferred to another financial institution, which assumed the banking operations and customer relationships of the failed bank.

Figure 9 plots the results from this analysis showing the effect on 5-year firm-level survival rates for each bank studied, ordered by the size of the effect. The heterogeneity observed in our analysis indicates that while the average effect of bank failures on firm survival is negative, the magnitudeand even the directionof this effect varies significantly across different bank failures. For most banks, we find substantial negative impacts on the five-year survival rates of associated firms, with some effects as large as a 15 percentage point decrease. However, in a few cases, firms linked to failed banks exhibit higher survival rates compared to the control group. This variation suggests that the consequences of bank failures are not uniform and are influenced by specific characteristics of the failed banks, the acquiring institutions, and the firms themselves.

Several factors may explain this heterogeneity. Some failure may occur during periods of heightened financial distress, during which firms find it harder to switch to a new banking partner. Furthermore, well-capitalized acquirers or those offering improved credit terms could potentially offset the negative effects of the failed bank relationship, leading to better outcomes for some affected firms. And lastly, the effects of bank failure may be exascerbated in instances where the failure leads to a large number of local branch closures and the destruction of the "soft" information on borrowers acquired by local branches. We test these diferent channels in the subsections below.

Firm size heterogeneity: The literature on banking relationships emphasizes their significance for small firms, as these relationships help overcome informational asymmetries that are particularly pronounced for very small, informationally opaque enterprises. Through relationship building, loan officers acquire "soft" information about prospective borrowers, enhancing screening and monitoring capabilities. Several studies support this view, demonstrating that longer and stronger banking relationships yield benefits to smaller firms (Petersen and Rajan 1994; Cole 1998; Elsas and Krahnen 1998; Drexler and Schoar 2014).

To further investigate size effects within our sample of already small firms, and to test the hypothesis that a firm-bank relationship disruption is more harmful as the size of the firm decreases, we run our headline regression specification separately for different size terciles. Figure 10 illustrates that the bottom tercile of firms (median firm size of 2) experiences more severe effects from bank failures compared to the top tercile (median firm size of 12). This pattern is evident across multiple years following the event, with the coefficient estimates for the bottom tercile consistently more negative than those for the top tercile. These results underscore the heightened vulnerability of the very smallest firms to disruptions in their banking relationships, even within a sample of small enterprises.

Crisis vs. non-crisis period bank failures: To investigate whether the effects of bank failures on firm survival differ during periods of widespread financial distress compared to more stable economic times, we separately estimate our baseline model for bank failures occurring during the Global Financial Crisis (GFC) period (2007-2011) and the post-GFC period (2012-2022). To justify our choice of 2012 as the cut-off for the non-crisis period, we examine several financial stress indicators that capture different aspects of market conditions and risk perceptions.[33] By

[33]The indicators used are as follows: (1) VIX index (FRED code: VIXCLS) measures market expectations of near-term volatility conveyed by S&P 500 stock index option prices. (2) St. Louis Fed Financial Stress Index (FRED code:

2012, these indicators had largely returned to pre-crisis levels or showed significant reduction in financial stress. Figure 11 presents the results of this analysis. The estimates show broadly similar patterns for both periods, with negative effects on firm survival emerging immediately after bank failure and persisting over time. The similarity in patterns across both periods underscores the importance of effective bank resolution strategies and support mechanisms for affected firms, not only during financial crises but also in more stable economic environments.

Bank branch networks: The resolution of failed banks through Purchase and Assumption (P&A) transactions can be viewed as a forced acquisition of the failed bank by a healthy institution. The literature on bank acquisitions and mergers suggests that these transactions create substantial disruptions to small business credit supply at regional levels (Nguyen 2019; Vij 2020). In particular, Nguyen (2019) provides quasi-experimental evidence showing that counties in which both the target and the acquiring bank had branches experience a 158% increase in the probability of a bank branch closure in the two years following the acquisition. This finding is further supported by Vij (2020), who shows that after acquiring a failed bank, the acquiring bank is more likely to shut down a branch previously belonging to the failed bank if it already had a presence in the same market.

The closure of bank branches following failures and acquisitions is closely linked to declines in small business lending. Nguyen (2019) finds that branch closings lead to a persistent decline in local small business lending, with annual originations falling by nearly 10% and remaining depressed for up to six years. The negative effects of branch closures on small business lending are further documented in Berger, et al. (1998), Di Patti and Gobbi (2007), Vij (2020), and Amberg and Becker (2024). These studies highlight the importance of physical branch proximity and the relationship-based nature of small business lending (Petersen and Rajan 1994; Berger and Udell 1995; Drexler and Schoar 2014).

Motivated by this literature, we explore the effect of branch availability on firm-level outcomes following bank failures. To test the hypothesis that firms in markets with fewer bank branches may face greater difficulties in establishing new banking relationships and securing credit after their primary bank fails, we examine the impact of bank failures on firm survival rates across different quartiles of branch availability.

STLFSI2) measures the degree of financial stress in the markets using 18 weekly data series. (3) TED Spread (FRED code: TEDRATE) is the difference between the 3-month LIBOR and the 3-month Treasury bill rate, indicating perceived credit risk in the general economy. (4) ICE BofA US High Yield Index Option-Adjusted Spread (FRED code: BAMLH0A0HYM2) measures the spread between yields on high-yield bonds and Treasury bonds, reflecting risk premiums in the corporate bond market.

Figure 12 presents the results, revealing that firms in the bottom quartile of branch availability—as measured by the total number of branches within five miles of the firm—experience more severe negative effects compared to those in the top quartile. The coefficient estimates for the bottom quartile are consistently more negative suggesting that firms with fewer nearby banking options are more vulnerable to the disruption caused by their bank's failure.

These findings align with the evidence on the persistent effects of branch closures on small business credit availability (Nguyen 2019) and underscore the spatial dimension of banking relationships. The results also highlight potential disparities in access to credit across different local banking markets, with firms in areas with less dense branch networks being more susceptible to the adverse consequences of bank failures.

7 Conclusion

This paper provides new evidence on how financial shocks propagate to the real economy by studying how bank failures affect small businesses. Using a novel dataset linking over 36 million loan documents to detailed firm-level information, we track how 179 bank failures between 1990 and 2023 affect small business survival and employment growth. Our focus on small firmsthe median firm in our sample has just three employeesborrowing from small banks provides insight into a crucial but understudied transmission channel through which financial sector distress affects the real economy.

We find that bank failures have large and persistent negative effects on small businesses. Firms that borrow from failed banks are 33% more likely to cease operations within one year compared to similar firms borrowing from non-failed banks, with this gap in survival rates growing to 37% after five years. The effects extend beyond survival: firms that remain operational experience substantially lower employment growth for up to a decade after their bank fails, with no evidence of catch-up growth. These persistent effects suggest that bank failures do not merely accelerate the exit of already vulnerable firms but rather create lasting disruptions that alter firms' long-term growth trajectories.

The effects of bank failures are particularly severe for the smallest firms in our sample, with the bottom third by employment experiencing effects twice as large as the top third. This size gradient highlights how firm characteristics that correlate with opacity and difficulty in switching lenderssuch as sizeshape the transmission of financial shocks. While bank failures have significant negative effects on average, we document substantial heterogeneity across different failures, with some resolution processes proving far more disruptive than others. This variation provides insight into which features of bank failures and resolution processes drive the worst outcomes for borrowers.

Our findings challenge the view that bank failures should have limited real effects in competitive banking markets where firms can seamlessly switch lenders. While the US bank resolution process appears designed to create frictionless transitions—with the FDIC arranging for healthier banks to acquire failed institutions—the reality for small businesses is far more disruptive. Branch closures following acquisitions can destroy the soft information that banks accumulate about borrowers over time, making it difficult for small, informationally opaque firms to access credit even after their loans are transferred to healthier institutions.

These findings have important implications both for our understanding of how financial shocks affect the real economy and for bank regulation and resolution policy. Our results provide direct evidence of a key mechanism through which financial sector disruptions create lasting damage to real economic activity: the destruction of valuable lending relationships between small banks and small businesses. The evidence that these effects persist for a decadeeven outside of crisis periodssuggests that the spillovers from financial sector distress to the real economy can be more severe and longer-lasting than previously documented. From a policy perspective, our findings indicate that regulators should carefully weigh the costs of different resolution strategies. The heterogeneity we document across bank failures suggests that regulators can potentially minimize the real economic damage of bank failures by targeting specific features of the resolution process.

References

ACS, ZOLTAN J., AND DAVID B. AUDRETSCH (1988): "Innovation in Large and Small Firms: An Empirical Analysis," *The American Economic Review*, 78(4), 678–690, Publisher: American Economic Association.

AGARWAL, SUMIT, AND ROBERT HAUSWALD (2010): "Distance and Private Information in Lending," *Review of Financial Studies*, 23(7), 2757–2788.

AIBANGBEE, YANY (2022): "How Has the Size Distribution of Banks Evolved Over the Last 30 Years?," .

ALMEIDA, HEITOR, MURILLO CAMPELLO, BRUNO LARANJEIRA, AND SCOTT WEISBENNER (2012): "Corporate Debt Maturity and the Real Effects of the 2007 Credit Crisis," *Critical Finance Review*.

AMBERG, NIKLAS, AND BO BECKER (2024): "Banking Without Branches," *Sveriges Riksbank working paper series*.

AMITI, M., AND D. E. WEINSTEIN (2011): "Exports and Financial Shocks," *The Quarterly Journal of Economics*, 126(4), 1841–1877.

ANTONIADES, ADONIS (2015): "Commercial Bank Failures During The Great Recession: The Real (Estate) Story," *Bank for International Settlements working paper*, 530.

ASHCRAFT, ADAM B. (2005): "Are Banks Really Special? New Evidence from the FDIC-Induced Failure of Healthy Banks," *American Economic Review*.

ATHEY, SUSAN, AND GUIDO W. IMBENS (2022): "Design-based analysis in Difference-In-Differences settings with staggered adoption," *Journal of Econometrics*, 226(1), 62–79.

AYYAGARI, MEGHANA, ASLI DEMIRGUC-KUNT, AND VOJISLAV MAKSIMOVIC (2011): "Small vs. Young Firms across the World: Contribution to Employment, Job Creation, and Growth," Discussion paper, The World Bank.

BARNATCHEZ, KEITH, LELAND D. CRANE, AND RYAN A. DECKER (2017): "An Assessment of the National Establishment Time Series (NETS) Database," *Finance and Economics Discussion Series*, 2017(110).

BAUMOL, WILLIAM J. (2004): "Entrepreneurial Enterprises, Large Established Firms and Other Components of the Free-Market Growth Machine," *Small Business Economics*, 23(1), 9–21, Publisher: Springer.

BENTOLILA, SAMUEL, MARCEL JANSEN, AND GABRIEL JIMÉNEZ (2018): "When Credit Dries Up: Job Losses in the Great Recession," *Journal of the European Economic Association*, 16(3), 650–695.

BERGER, ALLEN N., ANIL K. KASHYAP, AND JOSEPH M. SCALISE (1995): "The transformation of the U.S. banking industry: what a long, strange trip its been," *Brookings Papers on Economic Activity*, 2, 55–218.

BERGER, ALLEN N, NATHAN H MILLER, MITCHELL A PETERSEN, RAGHURAM G RAJAN, AND JEREMY C STEIN (2005): "Does function follow organizational form? Evidence from the lending practices of large and small banks," *Journal of Financial Economics*.

BERGER, ALLEN N, ANTHONY SAUNDERS, JOSEPH M SCALISE, AND GREGORY F UDELL (1998): "The eects of bank mergers and acquisitions on small business lending," *Journal of Financial Economics*.

BERGER, ALLEN N., AND GREGORY F. UDELL (1995): "Relationship Lending and Lines of Credit in Small Firm Finance," *The Journal of Business*, 68(3), 351–381, Publisher: University of Chicago Press.

BERGER, ALLEN N, AND GREGORY F UDELL (1998): "The economics of small business finance: The roles of private equity and debt markets in the őnancial growth cycle," .

BERNANKE, BEN S. (1983): "Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression," *The American Economic Review*, 73(3), 257–276, Publisher: American Economic Association.

BERTON, FABIO, SAURO MOCETTI, ANDREA F. PRESBITERO, AND MATTEO RICHIARDI (2018): "Banks, Firms, and Jobs," *The Review of Financial Studies*, 31(6), 2113–2156.

BORUSYAK, KIRILL, XAVIER JARAVEL, AND JANN SPIESS (2024): "Revisiting Event Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*.

BREWER, ELIJAH, HESNA GENAY, WILLIAM CURT HUNTER, AND GEORGE G. KAUFMAN (2003): "The value of banking relationships during a financial crisis: Evidence from failures of Japanese banks," *Journal of the Japanese and International Economies*, 17(3), 233–262.

CALLAWAY, BRANTLY, AND PEDRO H.C. SANTANNA (2021): “Difference-in-Differences with multiple time periods,” *Journal of Econometrics*, 225(2), 200–230.

CARDARELLI, ROBERTO, SELIM ELEKDAG, AND SUBIR LALL (2011): “Financial stress and economic contractions,” *Journal of Financial Stability*, 7(2), 78–97.

CHODOROW-REICH, GABRIEL (2014): “The Employment Effects of Credit Market Disruptions: Firm-Level Evidence From the 20089 Financial Crisis,” *Quarterly Journal of Economics*, 129(1).

CHODOROW-REICH, GABRIEL, OLIVIER DARMOUNI, STEPHAN LUCK, AND MATTHEW PLOSSER (2020): “Bank Liquidity Provision Across the Firm Size Distribution,” *NBER working paper*, (27945).

COHEN, WESLEY M., AND STEVEN KLEPPER (1996): “Firm Size and the Nature of Innovation within Industries: The Case of Process and Product R&D,” *The Review of Economics and Statistics*, 78(2), 232–243, Publisher: The MIT Press.

COLE, REBEL A. (1998): “The importance of relationships to the availability of credit,” *Journal of Banking & Finance*, 22(6-8), 959–977.

COLE, REBEL A., AND LAWRENCE J. WHITE (2012): “Déjà Vu All Over Again: The Causes of U.S. Commercial Bank Failures This Time Around,” *Journal of Financial Services Research*, 42(1-2), 5–29.

COLE, REBEL A., AND JOHN D. WOLKEN (1996): “Bank and nonbank competition for small business credit: evidence from the 1987 and 1993 National Surveys of Small Business Finances,” *Federal Reserve Bulletin*, 82, 983–995.

CORREIA, SERGIO, STEPHAN LUCK, AND EMIL VERNER (2024): “Failing Banks,” *NBER working paper*.

CRANE, LELAND D, AND RYAN A DECKER (2020): “Research with Private Sector Business Microdata: The Case of NETS/D&B,” *Federal Reserve Bank Working Paper Series*.

CULLEN, ALEX (2011): “Why do Banks Fail?,” *working paper*.

DE CHAISEMARTIN, CLÉMENT, AND XAVIER DHAULTFUILLE (2020): “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 110(9), 2964–2996.

DEN HAAN, WOUTER J., GAREY RAMEY, AND JOEL WATSON (2003): “Liquidity flows and fragility of business enterprises,” *Journal of Monetary Economics*.

DEĂCHAISEMARTIN, CLÉMENT, AND XAVIER DHAULTFUILLE (2023): "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey," *The Econometrics Journal*, 26(3), C1–C30.

DI PATTI, EMILIA BONACCORSI, AND GIORGIO GOBBI (2007): "Winners or Losers? The Effects of Banking Consolidation on Corporate Borrowers," *The Journal of Finance*, 62(2), 669–695.

DREXLER, ALEJANDRO, AND ANTOINETTE SCHOAR (2014): "Do Relationships Matter? Evidence from Loan Officer Turnover," *Management Science*, 60(11), 2722–2736, Publisher: INFORMS.

DUBE, ARINDRAJIT, DANIELE GIRARDI, ÒSCAR JORDÀ, AND ALAN TAYLOR (2024): "A Local Projections Approach to Difference-in-Differences," *NBER working paper*, p. w31184.

EDGERTON, JESSE (2012): "Credit Supply and Business Investment During the Great Recession: Evidence from Public Records of Equipment Financing," *working paper*.

ELSAS, RALF, AND JAN PIETER KRAHNEN (1998): "Is relationship lending special? Evidence from credit-őle data in Germany q," .

EREL, ISIL, TAYLOR NADAULD, AND RENÉ M. STULZ (2014): "Why Did Holdings of Highly Rated Securitization Tranches Differ So Much across Banks?," *Review of Financial Studies*, 27(2), 404–453.

FAMA, EUGENE F. (1980): "Banking in the theory of finance," *Journal of Monetary Economics*, 6(1), 39–57.

FDIC (2018): "Small Business Lending Survey," Discussion paper.

——— (2024): "Bank Failures & Assistance Data," .

FEDERAL DEPOSIT INSURANCE CORPORATION (2023): "Options For Deposit Insurance Reform," Discussion paper.

FEDERAL RESERVE (2020): "2020 Report on Employer Firms: Findings from the 2019 Small Business Credit Survey," Small Business Credit Survey (SBCS).

——— (2022a): "Availability of Credit to Small Businesses," Discussion paper, Board of Governors of the Federal Reserve System.

——— (2022b): "Insured U.S.-Chartered Commercial Banks That Have Consolidated Assets of $300 Million or More, Ranked by Consolidated Assets," Discussion paper.

FURLONG, FRED, AND ZENA KNIGHT (2010): "Loss Provisions and Bank Charge-offs in the Financial Crisis: Lesson Learned," *Federal Reserve Board - San Francisco - Economic Letter*.

GAN, JIE (2007): "The Real Effects of Asset Market Bubbles: Loan- and Firm-Level Evidence of a Lending Channel," *Review of Financial Studies*, 20(6), 1941–1973.

GARDNER, JOHN (2021): "Two-stage dierences in dierences," .

GERTLER, MARK, AND SIMON GILCHRIST (2018): "What Happened: Financial Factors in the Great Recession," *Journal of Economic Perspectives*, 32(3), 3–30.

GIANNETTI, MARIASSUNTA, AND ANDREI SIMONOV (2013): "On the Real Effects of Bank Bailouts: Micro Evidence from Japan," *American Economic Journal: Macroeconomics*.

GNANARAJAH, RAJ (2023): "Bank Failures and the FDIC," Discussion paper.

GODA, GOPI SHAH, AND EVAN J. SOLTAS (2023): "The impacts of Covid-19 absences on workers," *Journal of Public Economics*, 222, 104889.

GOODMAN-BACON, ANDREW (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225(2), 254–277.

GOPAL, MANASA, AND PHILIPP SCHNABL (2022): "The Rise of Finance Companies and FinTech Lenders in Small Business Lending," *The Review of Financial Studies*, 35(11), 4859–4901.

GREENSTONE, MICHAEL, ALEXANDRE MAS, AND HOAI-LUU NGUYEN (2020): "Do Credit Market Shocks Affect the Real Economy? Quasi-experimental Evidence from the Great Recession and Normal Economic Times," *American Economic Journal: Economic Policy*, 12(1), 200–225, Publisher: American Economic Association.

GREENWALD, DANIEL L., JOHN KRAINER, BOARD OF GOVERNORS OF THE FEDERAL RESERVE, AND PASCAL PAUL (2023): "The Credit Line Channel," *Federal Reserve Bank of San Francisco, Working Paper Series*.

HALTIWANGER, JOHN, RON S. JARMIN, AND JAVIER MIRANDA (2013): "Who Creates Jobs? Small versus Large versus Young," *Review of Economics and Statistics*, 95(2), 347–361.

HOGGARTH, GLENN, RICARDO REIS, AND VICTORIA SAPORTA (2002): "Costs of banking system instability: Some empirical evidence q," .

HUBER, KILIAN (2018): "Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties," *American Economic Review*, 108(3), 868–898.

JAFFEE, DWIGHT, AND JOSEPH STIGLITZ (1990): “Chapter 16 Credit rationing,” in *Handbook of Monetary Economics*, vol. 2, pp. 837–888. Elsevier.

JOHNSTON-ROSS, EMILY, SONG MA, AND MANJU PURI (2021): “Private Equity and Financial Stability: Evidence from Failed Bank Resolution in the Crisis,” *FDIC Working Paper Series*.

KANDRAC, JOHN (2014): “Bank Failure, Relationship Lending, and Local Economic Performance,” .

KHWAJA, ASIM IJAZ, AND ATIF MIAN (2008): “Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market,” *American Economic Review*, 98(4), 1413–1442.

KOCHHAR, RAKESH, AND ANA GONZALEZ-BARRERA (2009): “Foreclosures in the U.S. in 2008,” Discussion paper.

LUCK, STEPHAN, AND JOÃO A. C. SANTOS (2023): “The Valuation of Collateral in Bank Lending,” *Journal of Financial and Quantitative Analysis*, pp. 1–30.

MA, SONG, JUSTIN MURFIN, AND RYAN PRATT (2022): “Young Firms, Old Capital,” *Journal of Financial Economics*, 146(1), 331–356.

MARTIN, CHRISTOPHER, MANJU PURI, AND ALEXANDER UFIER (2018): “Deposit Inflows and Outflows in Failing Banks: The Role of Deposit Insurance,” .

MIAN, ATIF, KAMALESH RAO, AND AMIR SUFI (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *The Quarterly Journal of Economics*, 128(4), 1687–1726, Publisher: Oxford University Press.

MIAN, ATIF, AND JOÃO A.C. SANTOS (2018): “Liquidity risk and maturity management over the credit cycle,” *Journal of Financial Economics*, 127(2), 264–284.

MINAMIHASHI, NAOAKI (2011): “Credit Crunch Caused by Bank Failures and Self-Selection Behavior in Lending Markets,” *Journal of Money, Credit and Banking*, 43(1), 133–161.

MURFIN, JUSTIN, AND RYAN PRATT (2019): “Who Finances Durable Goods and Why It Matters: Captive Finance and the Coase Conjecture,” *The Journal of Finance*, 74(2), 755–793, Publisher: [American Finance Association, Wiley].

NEUMARK, DAVID, BRANDON WALL, AND JUNFU ZHANG (2011): “Do Small Businesses Create More Jobs? New Evidence for the United States from the National Establishment Time Series,” *The Review of Economics and Statistics*, 93(1), 16–29, Publisher: The MIT Press.

NGUYEN, HOAI-LUU Q. (2019): "Are Credit Markets Still Local? Evidence from Bank Branch Closings," *American Economic Journal: Applied Economics*, 11(1).

ONGENA, STEVEN, DAVID C SMITH, AND DAG MICHALSEN (2003): "Firms and their distressed banks: lessons from the Norwegian banking crisis," *Journal of Financial Economics*, 67(1), 81–112.

PARAVISINI, DANIEL, VERONICA RAPPOPORT, AND PHILIPP SCHNABL (2023): "Specialization in Bank Lending: Evidence from Exporting Firms," *The Journal of Finance*, 78(4), 2049–2085.

PARAVISINI, DANIEL, VERONICA RAPPOPORT, PHILIPP SCHNABL, AND DANIEL WOLFENZON (2015): "Dissecting the Effect of Credit Supply on Trade: Evidence from Matched Credit-Export Data," *The Review of Economic Studies*, 82(1 (290)), 333–359, Publisher: [Oxford University Press, The Review of Economic Studies, Ltd.].

PEEK, JOE, AND ERIC S. ROSENGREN (1998): "Bank consolidation and small business lending: It's not just bank size that matters," *Journal of Banking & Finance*, 22(6-8), 799–819.

——— (2000): "Collateral Damage: Effects of the Japanese Bank Crisis on Real Activity in the United States," *The American Economic Review*, 90(1), 30–45, Publisher: American Economic Association.

PETERSEN, MITCHELL A., AND RAGHURAM G. RAJAN (1994): "The Benefits of Lending Relationships: Evidence from Small Business Data," *The Journal of Finance*, 49(1), 3–37, Publisher: [American Finance Association, Wiley].

——— (1995): "The Effect of Credit Market Competition on Lending Relationships," *The Quarterly Journal of Economics*, 110(2), 407–443, Publisher: Oxford University Press.

ROBB, ALICIA M., AND DAVID T. ROBINSON (2014): "The Capital Structure Decisions of New Firms," *The Review of Financial Studies*, 27(1), 153–179, Publisher: [Oxford University Press, The Society for Financial Studies].

ROTH, JONATHAN, PEDRO H.C. SANTANNA, ALYSSA BILINSKI, AND JOHN POE (2023): "Whats trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 235(2), 2218–2244.

RUSSELL, KARL, AND CHRISTINE ZHANG (2023): "3 Failed Banks This Year Were Bigger Than 25 That Crumbled in 2008," *The New York Times*.

SCHNABL, PHILIPP (2012): “The International Transmission of Bank Liquidity Shocks: Evidence from an Emerging Market,” *The Journal of Finance*, 67(3), 897–932.

SLOVIN, MYRON B., MARIE E. SUSHKA, AND JOHN A. POLONCHEK (1993): “The Value of Bank Durability: Borrowers as Bank Stakeholders,” *The Journal of Finance*, 48(1), 247–266, Publisher: [American Finance Association, Wiley].

SUN, LIYANG, AND SARAH ABRAHAM (2021): “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225(2), 175–199.

THE NATIONAL FEDERATION OF INDEPENDENT BUSINESS (1983): “Credit, Banks and Small Business,” Discussion paper, National Federation of Independent Business, San Mateo, CA.

VIJ, SIDDHARTH (2020): “Acquiring Failed Banks,” *working paper*.

Figure 1: Total Loan Filings by US County (1990-2023)

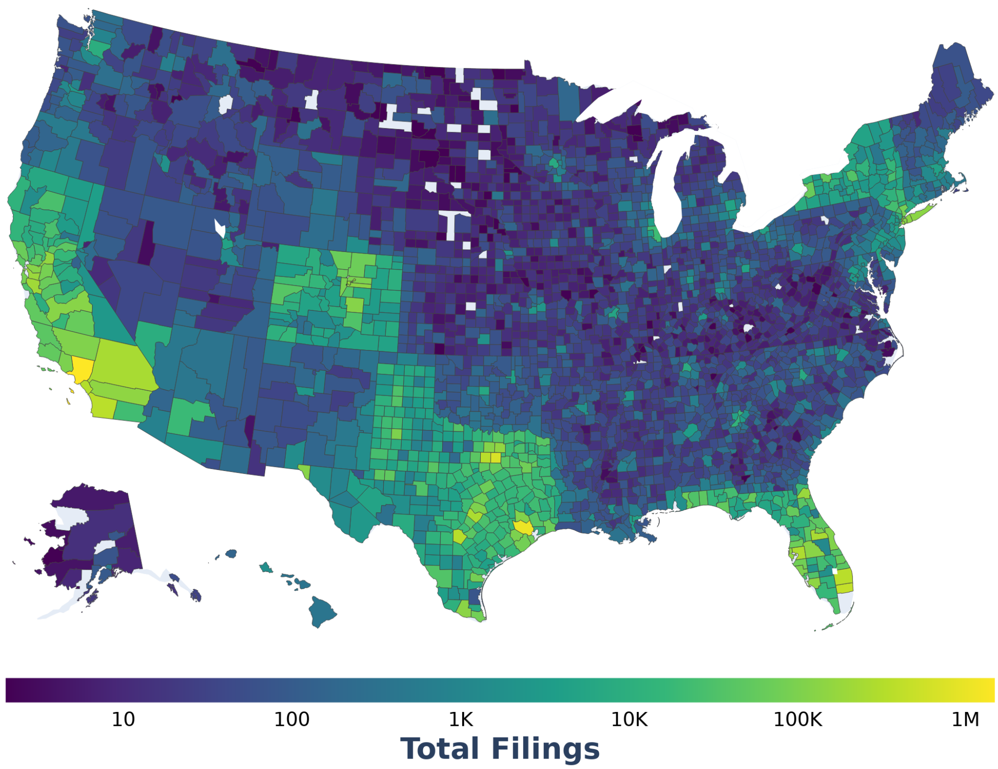

Note: We source our data from California, Texas, Florida, New York, and Colorado. This figure shows the geographic distribution of the UCC financing statements in our dataset across US counties from 1990 to 2023. Despite sourcing from these 5 States, our coverage extends nationally owing to many other firm-bank relationships having loan documentation in our focal states. Darker colors indicate a lower number of UCC filings. Allocation to counties is according to the debtor's address listed in the filing.

Figure 2: Effect of Bank Failure on Firm-Level Survival Rates

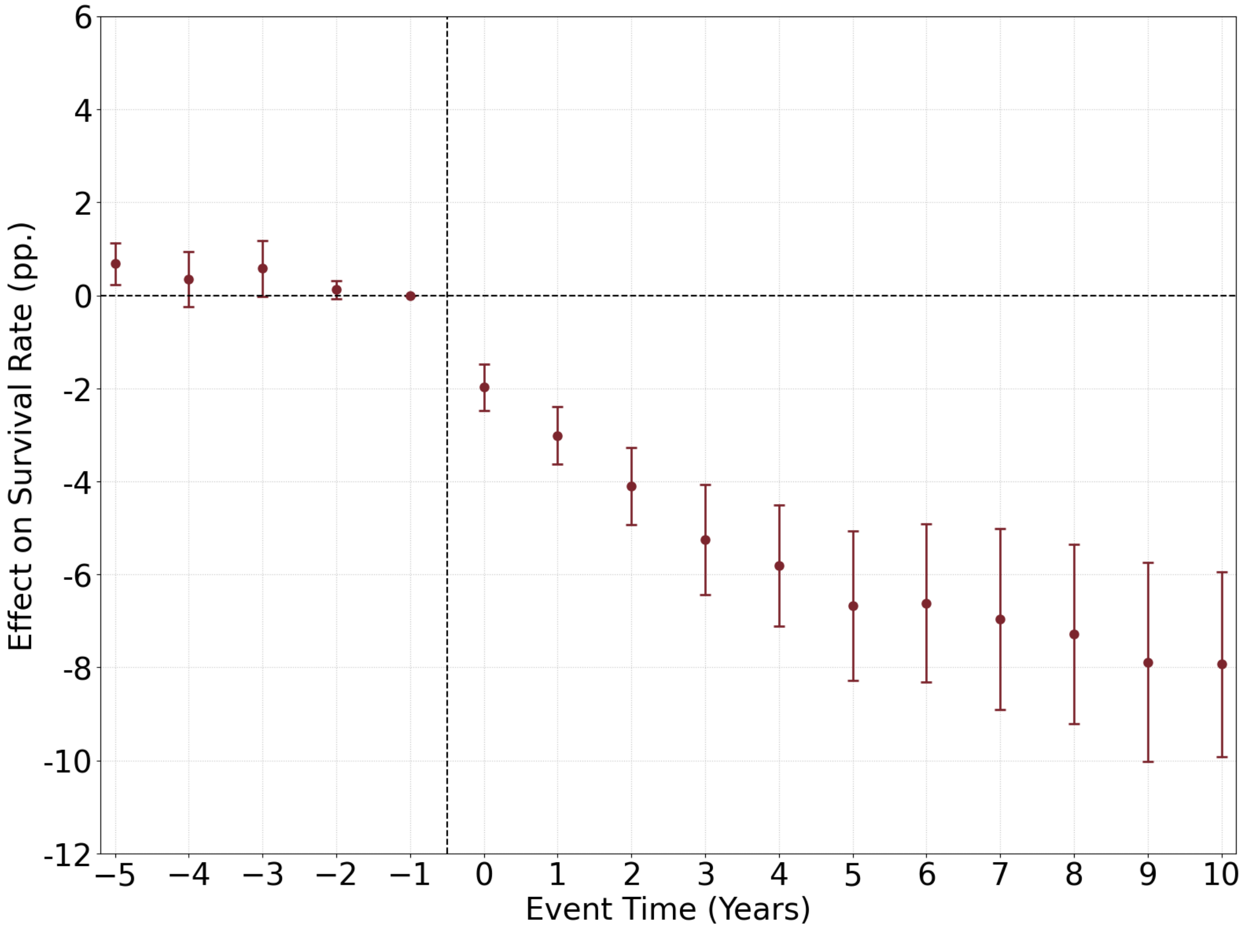

Note: This figure shows estimated effects of bank failure on firm-level survival rates using a matched local projection difference-in-difference (LP-DiD) estimator. Units are matched to balance across firm-size, industry, and county. The x-axis depicts event time (years) relative to bank failure. The y-axis displays the estimated effects. Estimates come from our baseline specification (see Equation 2). We control for year-by-county-by-industry-by-firm-size fixed effects as well as for firm age, firm age squared, and a lag of first difference in firm survival. Industry is defined at the two-digit SIC level. The sample comprises 483,403 observations over 19,762 firms (10,105 treated). Standard errors are clustered by both bank-by-county and year levels, and bars depict the 95% confidence interval.

Figure 3: Effect of Bank Failure on Firm-Level Survival Rates (alternative specifications)

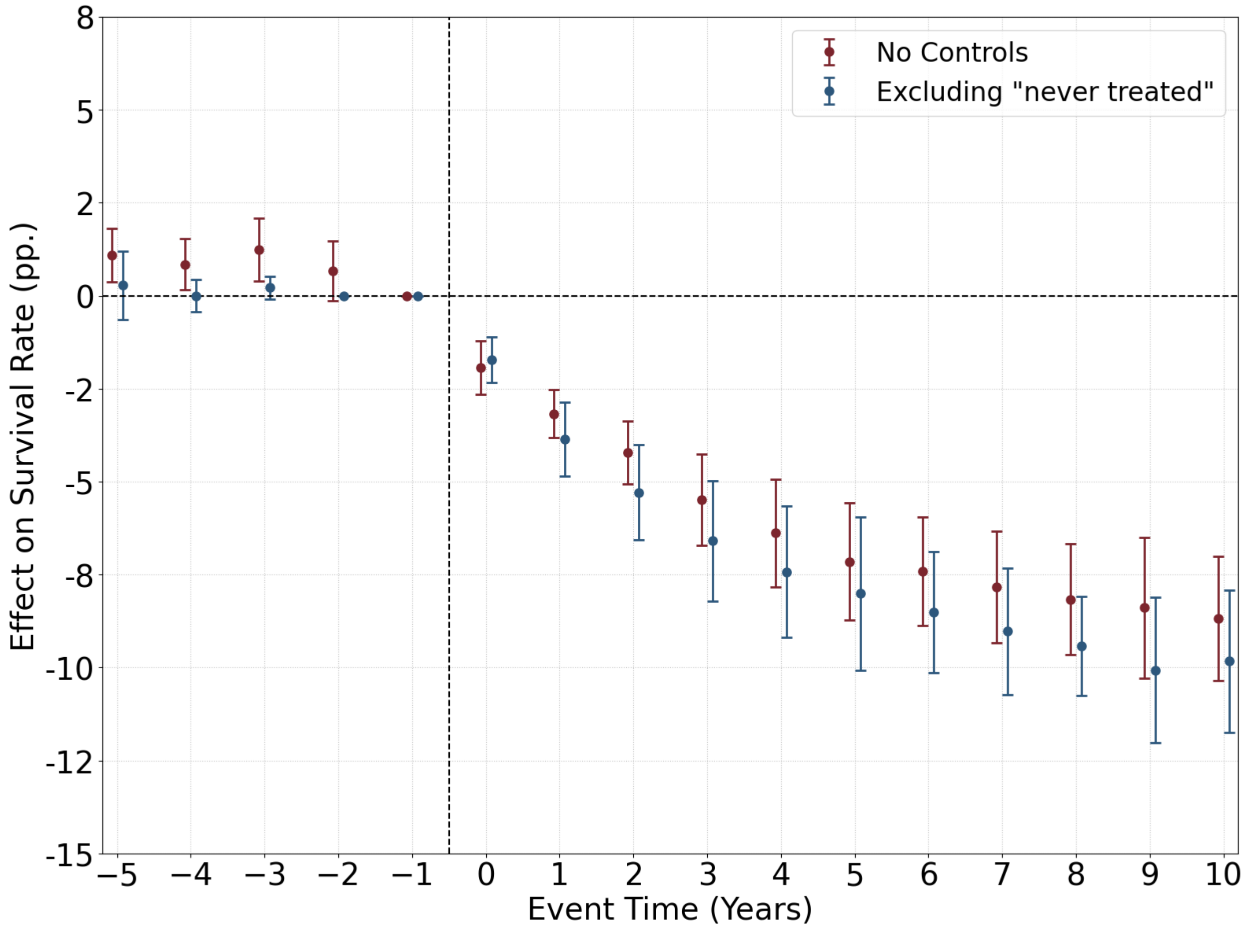

Note: This figure shows estimated effects of bank failure on firm-level survival rates using a matched local projection difference-in-difference (LP-DiD) estimator. Units are matched to balance across firm-size, industry, and county. The x-axis depicts event time (years) relative to bank failure. The y-axis displays the estimated effects. Estimates come from our baseline specification (see Equation 2). Burgundy estimates are obtained from an estimation of Equation 2 without any controls. Navy estimates are obtained from an estimation of Equation 2 with a full set of controls and excluding any firms who never bank with a failed bank (i.e. “never treated” firms). Controls include year-by-county-by-industry-by-firm-size fixed effects as well as firm age, firm age squared, and a lag of first difference in firm survival. Industry is defined at the two-digit SIC level. The estimation without controls is run on a sample that comprises 483,403 observations over 19,762 firms (10,105 treated). The estimation excluding “never treated” firms is run on a sample that comprises 423,525 observations over 19238 firms (all treated). Standard errors are clustered by both bank-by-county and year levels, and bars depict the 95% confidence interval.

Figure 4: Effect of Bank Failure on Firm-Level Survival Rates (alternative estimators)

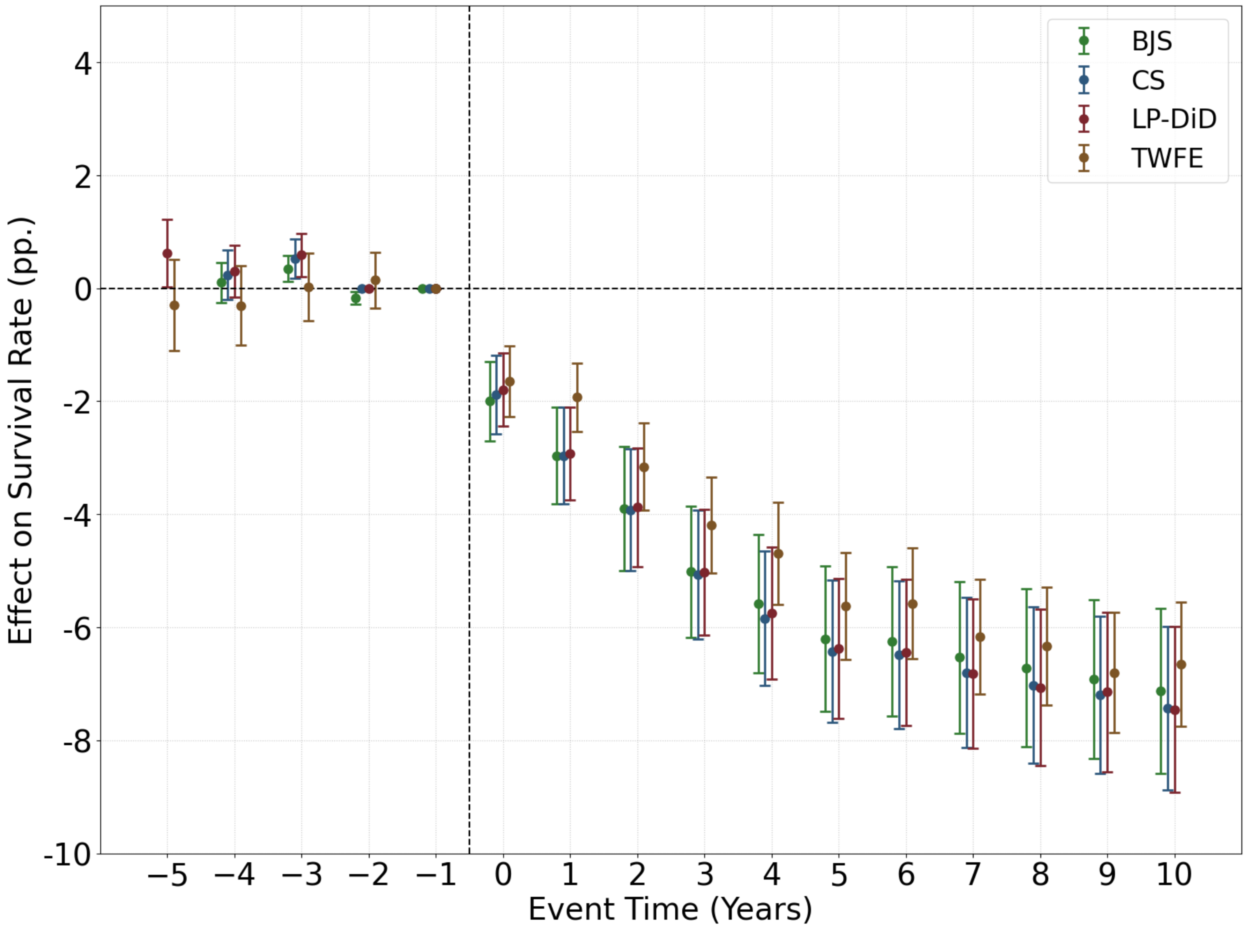

Note: This figure shows estimated effects of bank failure on firm-level survival rates using a matched using four different estimators: our baseline local projection estimator (LP-DiD) due to Dube, et al. (2024), the staggered treatment estimators proposed by Borusyak, Jaravel, and Spiess (2024) and Callaway and SantAnna (2021) abbreviated as BJS and CS respectively, as a well as a standard two-way fixed effects (TWFE) estimators. Units are matched to balance across firm-size, industry, and county. The x-axis depicts event time (years) relative to bank failure. The y-axis displays the estimated effects. Estimates come from our baseline specification (see Equation 2), controlling for firm age, firm age squared, and a lag of first difference in firm survival. Standard errors are clustered by both bank-by-county and year levels, and bars depict the 95% confidence interval.

Figure 5: Effect of Bank Failure on Firm-Level Survival Rates (including switching firms)

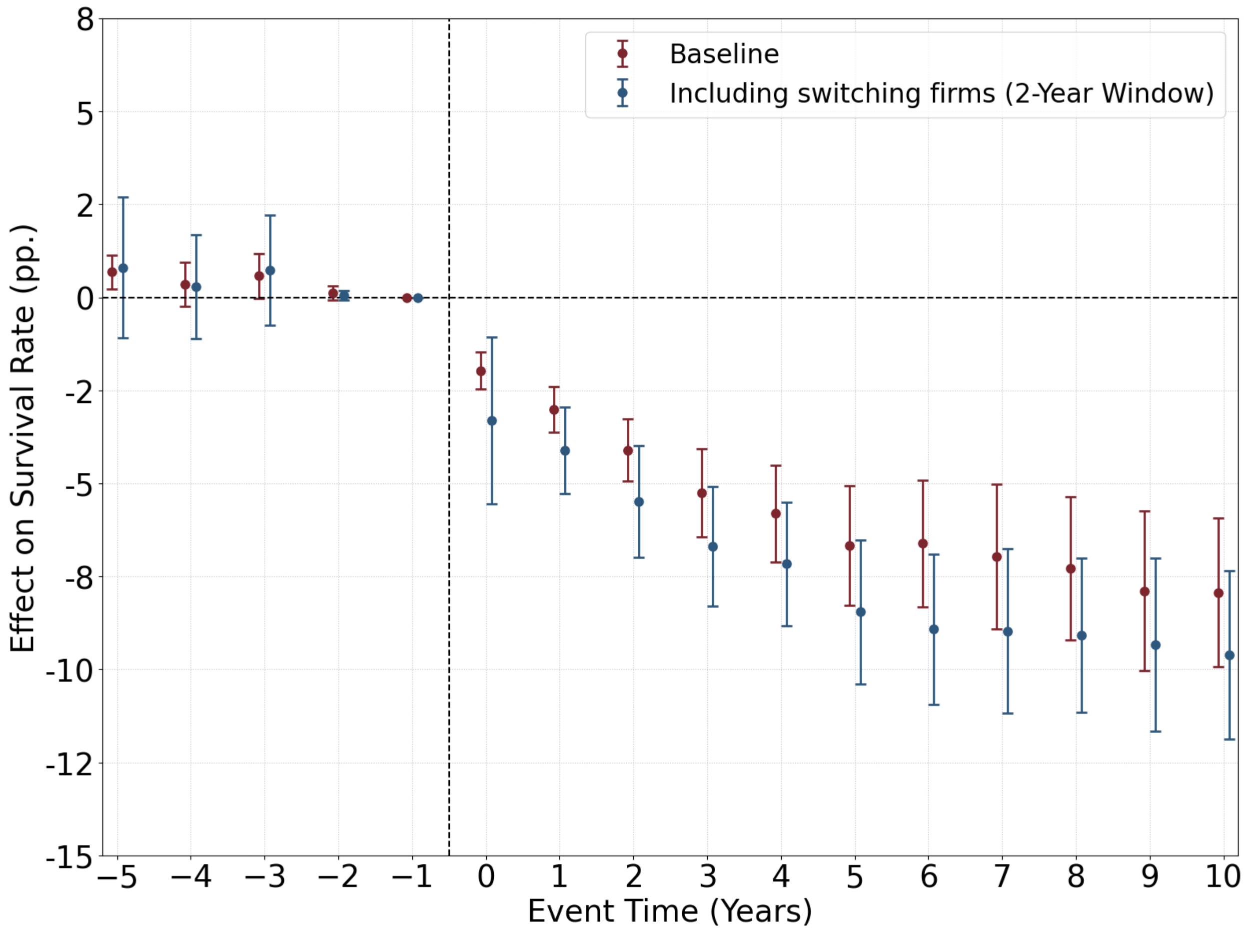

Note: This figure shows estimated effects of bank failure on firm-level survival rates using a matched local projection difference-in-difference (LP-DiD) estimator. Units are matched to balance across firm-size, industry, and county. The x-axis depicts event time (years) relative to bank failure. The y-axis displays the estimated effects. Estimates come from our baseline specification (see Equation 2). Burgundy estimates are obtained from an estimation of Equation 2 using our baseline specification (see note for Figure 3). Navy estimates are obtained from the same regression but changing the definition of treatment such that firms continue to be considered treated even if they switch to a healthy bank in the two years prior to bank failure. Similarly, firms that switch from a healthy bank to a soon-to-fail bank within two years of the bank failure are considered to be untreated firms. Standard errors are clustered by both bank-by-county and year levels, and bars depict the 95% confidence interval.

Figure 6: Implied Firm Death Rates Following Bank Failure

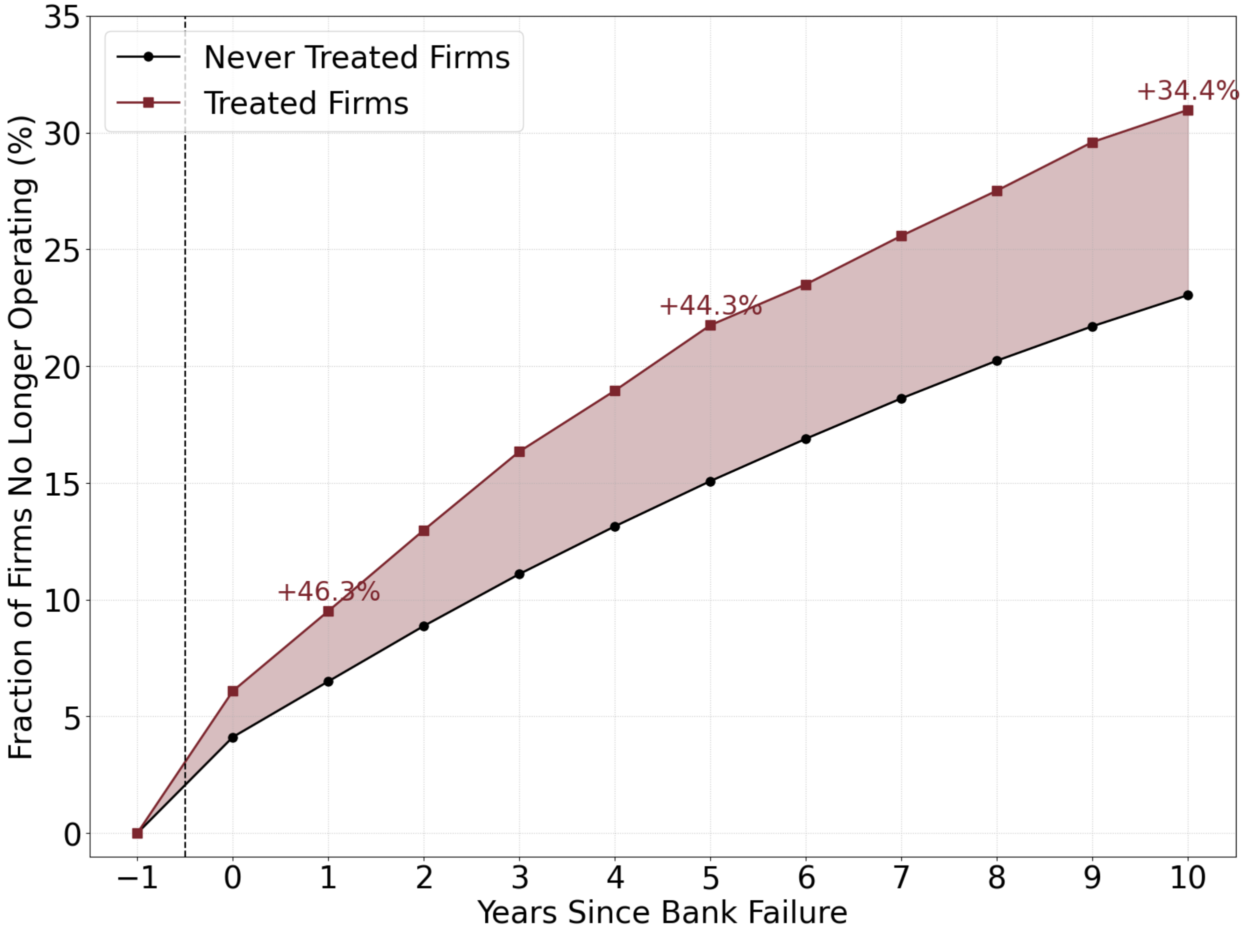

Note: This figure compares the death rates of firms affected by bank failures (treated firms) with those of firms which never bank with a failed bank (untreated firms) over a 10-year period following the bank failure event. Untreated firm death rates are calculated as the average death rates to each time horizon for firms not exposed to bank failures across our entire dataset. Treated firm death rates are derived by applying the estimated treatment effects to these baseline untreated rates. Percentage annotations indicate the relative increase in death rates for treated firms compared to untreated firms at years 1, 5, and 10 after the bank failure. The vertical dashed line at -0.5 years marks the pre-treatment period.

Figure 7: Effect of Bank Failure on Firm-Level Employment Growth

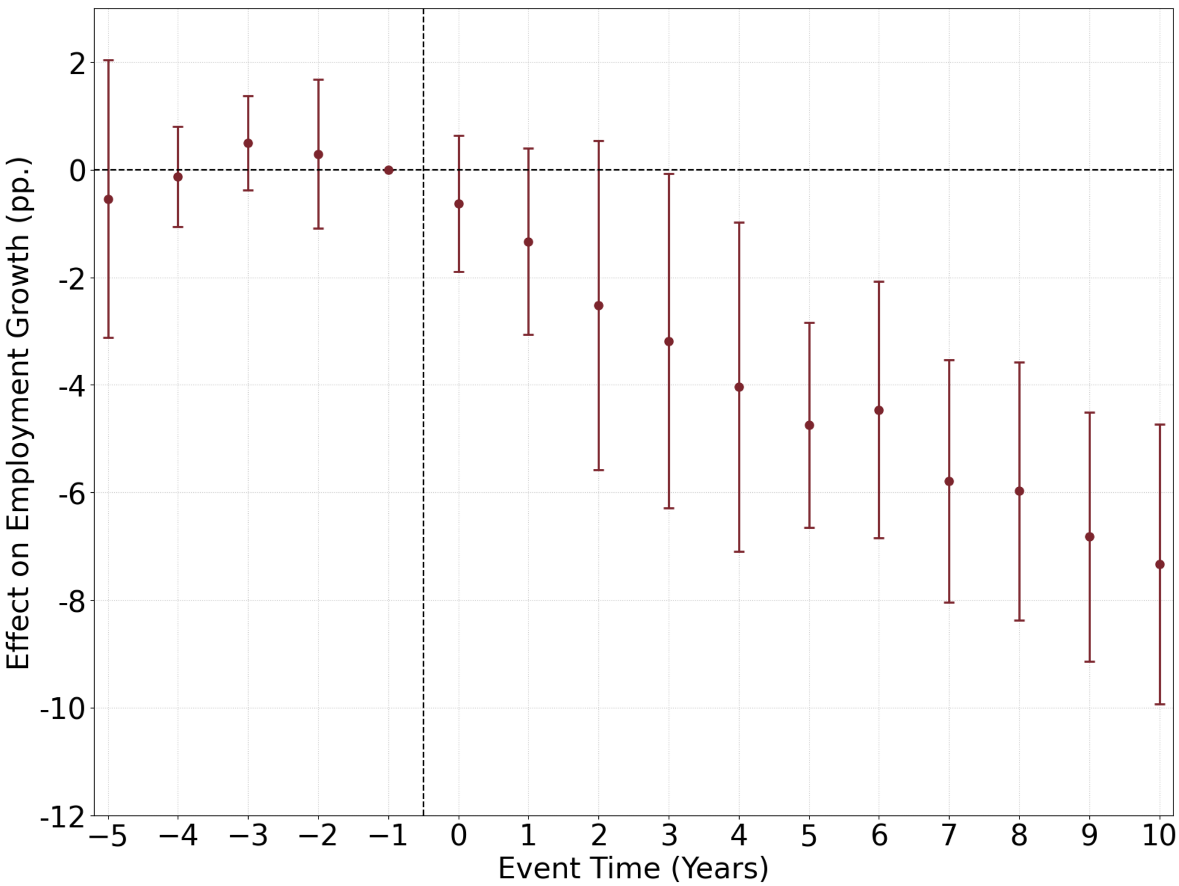

Note: This figure presents the estimated effects of bank failure on changes in firm-level log employment for different time horizons. The x-axis shows the event time in years relative to the bank failure, with 0 representing the year of failure. The y-axis displays the coefficient estimates, representing the percentage point difference in survival probability between treated firms (those exposed to bank failure) and never treated firms. Each point estimate is accompanied by its 95% confidence interval. The estimates are derived from a local projection DiD estimation of Equation 2 with county-by-year-by- industry fixed effects, controlling for bank sizeas proxied by book value of the banks loansfirm age and age squared. Standard errors are clustered at the bank and year levels

Figure 8: Comparison of Average Employment Growth for Treated and Untreated Firms

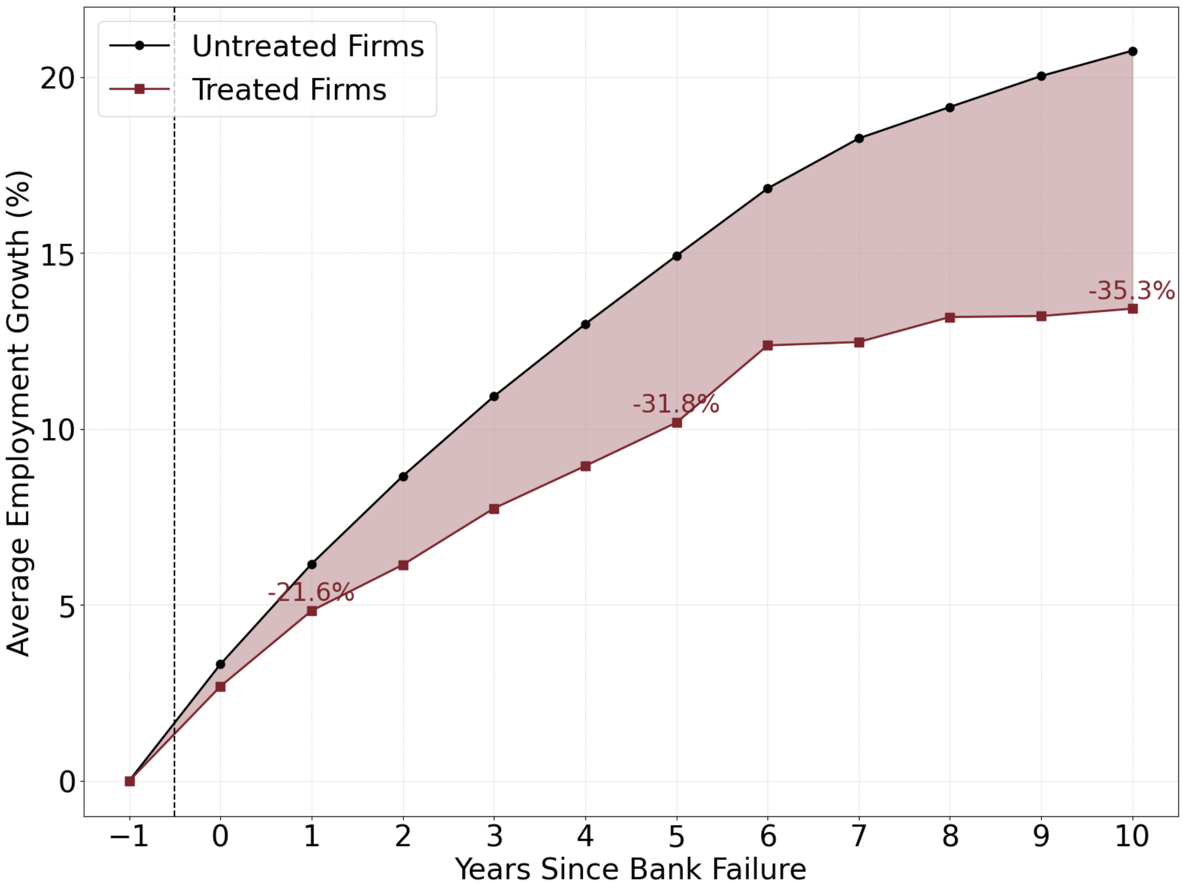

Note: This figure compares the firm-level employment growth firms affected by bank failures (treated firms) with those of firms which never bank with a failed bank (untreated firms) over a 10-year period following the bank failure event. Untreated firm employment growth rates are calculated as the average employment growth rates to each time horizon for firms not exposed to bank failures across our entire dataset. Treated firm employment growth rates are derived by applying the estimated treatment effects to these baseline untreated employment growth rates. Percentage annotations indicate the relative decrease in firm-level employment growth rates for treated firms compared to untreated firms. The vertical dashed line marks the bank failure event time.

Figure 9: Effect of Bank Failure on Firm-Level Survival Rates By Individual Failed Bank

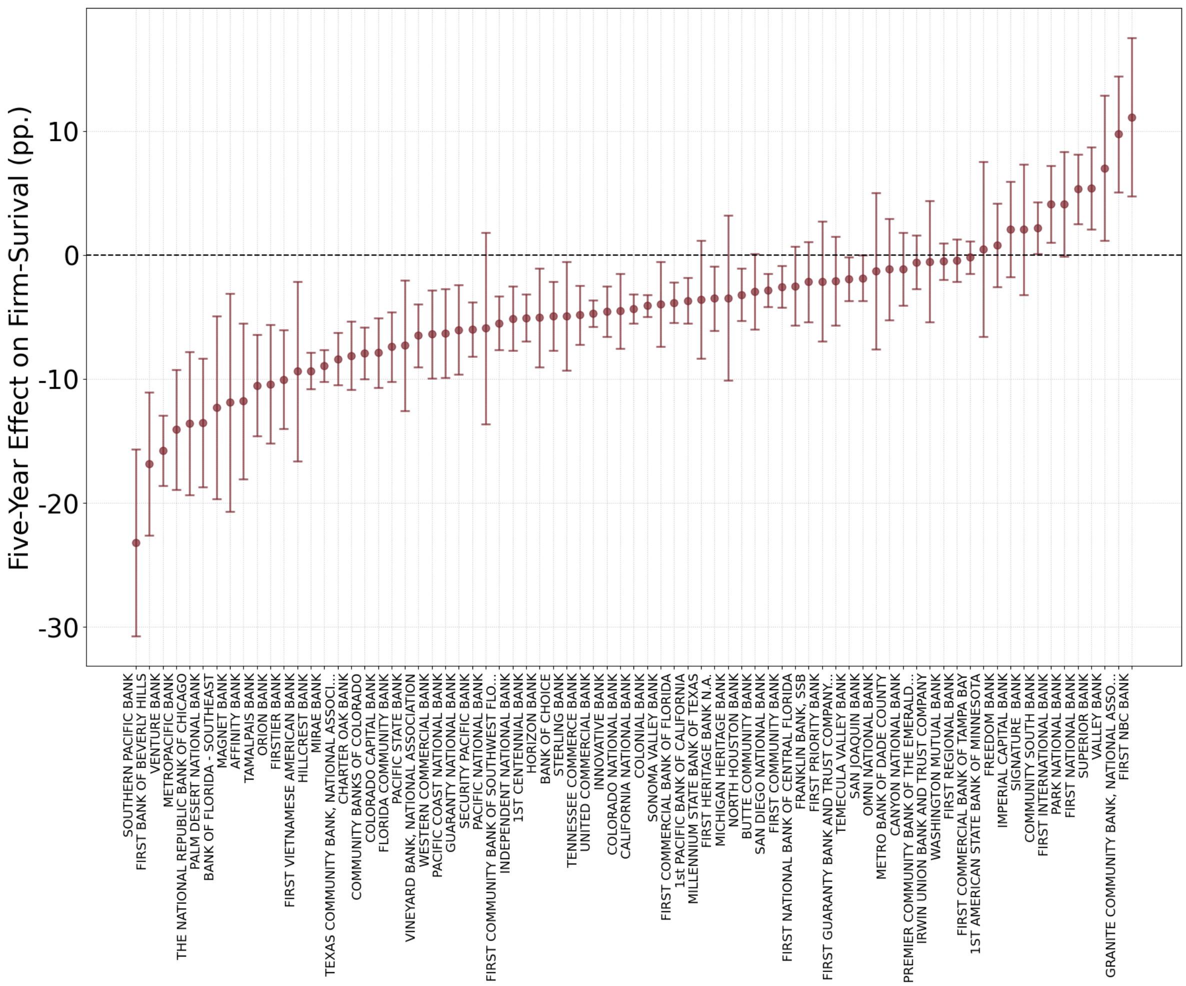

Note: This figure presents the estimated effects of bank failure on firm-level survival rates, disaggregated by individual failed banks. The x-axis represents the event time in years relative to the bank failure, with 0 representing the year of failure. The y-axis shows the coefficient estimates, indicating the percentage point difference in survival probability between treated firms (those exposed to the specific bank's failure) and never treated firms. Each line represents a separate regression for one of the 25 largest failed banks in our sample, as measured by the number of firms banking with these institutions. The estimates are derived from our baseline local projection DiD estimation with county-by-year-by-industry fixed effects, controlling for firm age, age squared, and bank size (proxied by the book value of the bank's loans). Standard errors are clustered at the firm and year levels.

Figure 10: Effect of Bank Failure on Firm-Level Survival Rates By Firm-Size Tercile

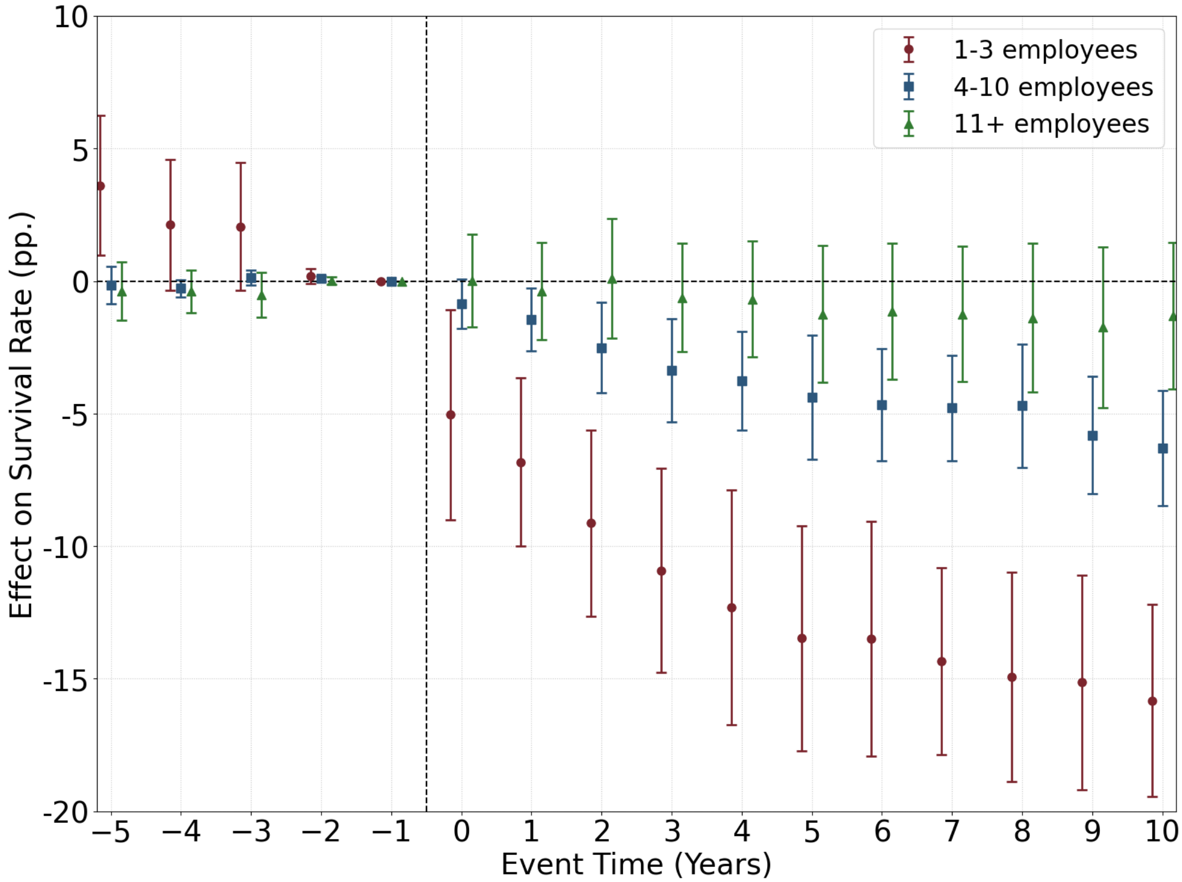

Note: This figure presents the estimated effects of bank failure on firm survival rates, disaggregated by firm size terciles. The x-axis represents years relative to bank failure, and the y-axis shows the coefficient estimates for the probability of firm survival. The bottom tercile (smallest firms) has a median size of 2 employees, while the top tercile has a median size of 12 employees. The more pronounced negative effects for the bottom tercile, particularly in the years following bank failure, highlight the greater vulnerability of the smallest firms to banking relationship disruptions. These results are consistent with the literature emphasizing the importance of banking relationships for small, informationally opaque firms in overcoming informational asymmetries. The estimates are derived from our headline regression specification, run separately for each size tercile, and include county-by-year-by-industry fixed effects and controls for firm age, age squared, and bank size.

Figure 11: Effect of Bank Failure on Firm-Level Survival Rates During and After US Financial Crisis

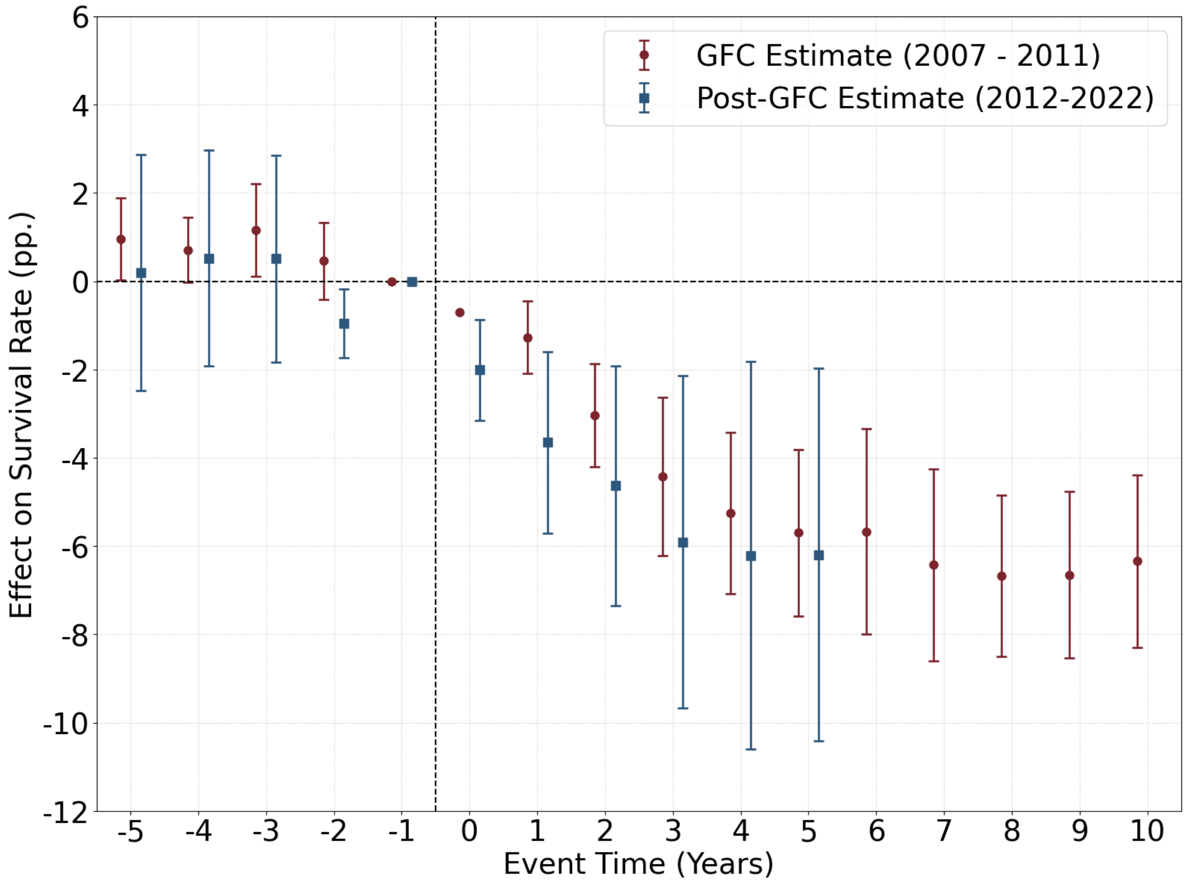

Note: This figure presents the estimated effects of bank failure on firm-level survival rates, comparing the Global Financial Crisis (GFC) period (2007-2011) to the post-GFC period (2012-2022). The x-axis shows the event time in years relative to the bank failure, with 0 representing the year of failure. The y-axis displays the coefficient estimates, representing the percentage point difference in survival probability between treated firms (those exposed to bank failure) and never treated firms. Each point estimate is accompanied by its 95% confidence interval. The estimates are derived from our baseline local projection DiD estimation with county-by-year-by-industry fixed effects, controlling for bank size, firm age and age squared. Standard errors are clustered at the bank and year levels. The vertical dashed line at event time 0 marks the bank failure year.

Figure 12: Effect of Bank Failure on Firm-Level Survival Rates by Branch Availability

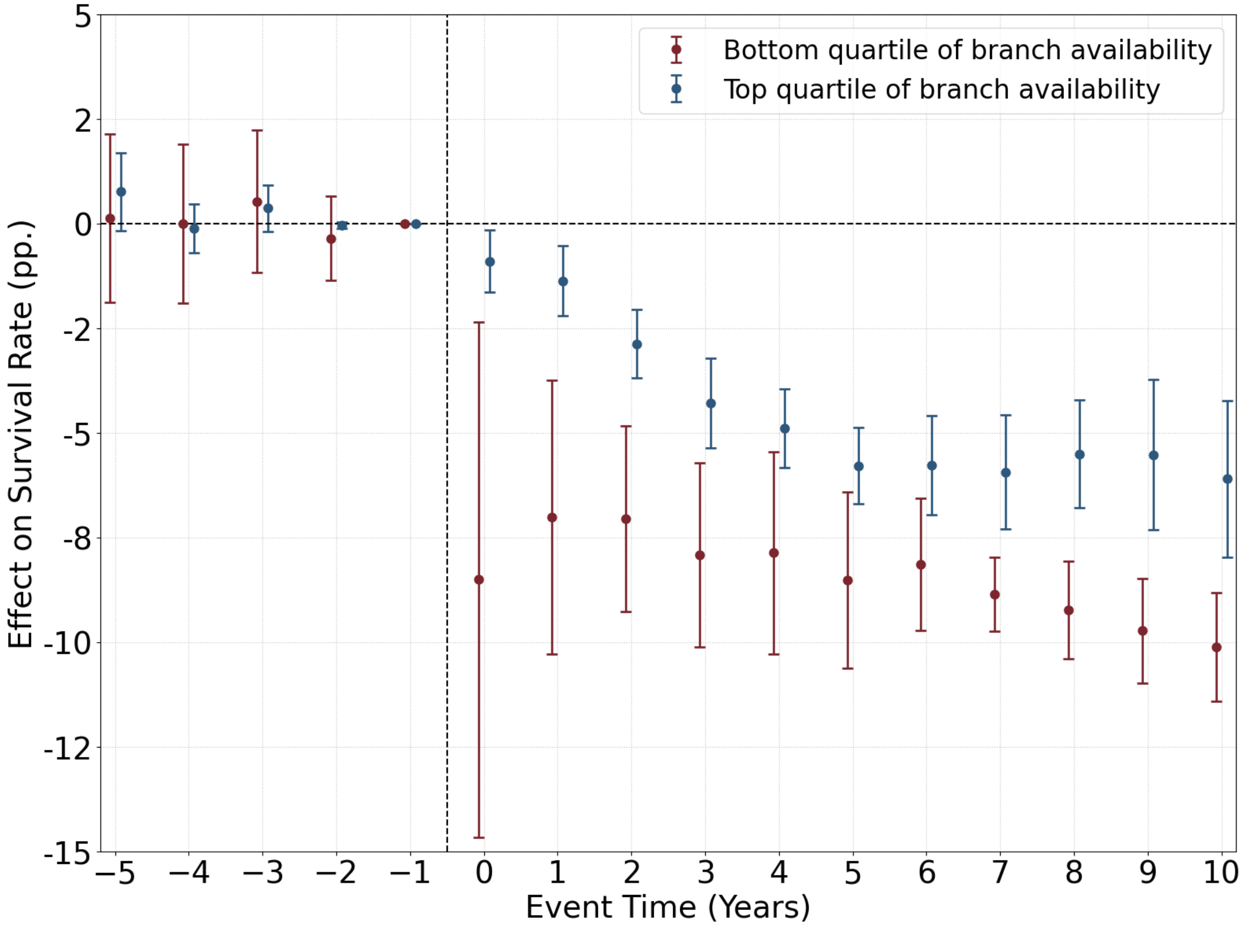

Note: This figure shows estimated effects of bank failure on firm-level survival rates using a matched local projection difference-in-difference (LP-DiD) estimator. Units are matched to balance across firm-size, industry, and county. The x-axis depicts event time (years) relative to bank failure. The y-axis displays the estimated effects. Estimates come from our baseline specification (see Equation 2). Burgundy (navy) estimates are obtained from an estimation of Equation 2 on a subset of firms that are in the bottom (top) quartile when ranking firms by the number of branches located within five miles of the firm . Controls include year-by-industry-by-firm-size fixed effects as well as firm age, firm age squared, and a lag of first difference in firm survival. Industry is defined at the two-digit SIC level. Standard errors are clustered by both bank-by-county and year levels, and bars depict the 95% confidence interval.

Table 1: Summary Statistics for UCC Filing to Firm Level Data Linkage

Statistic	**Value**
Number of unique firms	2,148,477
Number of unique filings	6,284,126
Cosine Similarity	
Average	0.97
Median	0.99
10th percentile	0.93
90th percentile	1.00
Geographic Distance (km)	
Average	0.03
Median	0.00
10th percentile	0.00
90th percentile	0.00
Match Probability	
Average	0.93
Median	0.94
10th percentile	0.89
90th percentile	0.97

Note: This table presents summary statistics of our final UCC filing dataset matched to firms. Geographic distance is measured in kilometers. Cosine similarity and match probability are based on the matching algorithm used.

Table 2: Descriptive Statistics of Firms

	Mean	SD	p10	Median	p90	N
Full sample						
Employees	11.8	71.0	1	3	20	1,545,143
No. of UCC Filings per Firm	3.3	9.6	1	1	6	1,545,143
Firm age at first UCC filing	8.3	11.8	0	4	21	1,545,143
Firm age at Exit	15.0	13.5	2	12	32	887,300
Treated sample (firms banking with failed banks)						
Years since start of banking relationship	7.6	5.8	1	6	16	25,470
Years since last contact with failed bank	6.8	5.6	1	5	15	25,470
Firm age at bank failure	17.7	15.0	3	15	34	25,470

Table 3: Descriptive Statistics of Banks

	All	Matched	Non-Failed	Failed
Assets ($ millions)				
Mean	1,395.4	5,120.8	5,238.5	2,329.9
P10	5.2	77.6	77.1	87.4
Median	105.3	367.5	367.5	368.2
P90	897.6	3,401.0	3,479.3	2,289.3
Employees				
Mean	182	614	631	216
P10	3	18	18	19
Median	25	74	74	69
P90	180	536	547	365
Number of banks	17,496	3,929	3,770	159

Note: This table presents summary statistics for different subsets of banks in our data. The sample consists of institutions classified under NAICS code 5221 (Depository Credit Intermediation) in the Orbis dataset. Assets are reported in millions of US dollars. Each statistic is calculated using bank-level averages over the sample period 1990 to 2023. All banks refers to all banks in the Orbis dataset. Unmatched banks are banks in the Orbis dataset that did not match to our firm-level panel. Matched banks include all banks that matched to at least one firm in our firm-level panel. Non-failed and failed banks are taken as subsets of all the banks matched to our firm-level panel.

Table 4: Effect of Bank Failure on Firm Survival

	(1)	(2)	(3)	(4)	(5)
Year 10	-0.092***	-0.111***	-0.069***	-0.092***	-0.074***
	(0.012)	(0.016)	(0.013)	(0.011)	(0.010)
Year 5	-0.081***	-0.094***	-0.069***	-0.081***	-0.057***
	(0.011)	(0.013)	(0.013)	(0.011)	(0.009)
Year 1	-0.029***	-0.038***	-0.042***	-0.029***	-0.025***
	(0.006)	(0.006)	(0.010)	(0.006)	(0.005)
Year 0	-0.014***	-0.020***	-0.033**	-0.014***	-0.015***
	(0.003)	(0.005)	(0.014)	(0.003)	(0.003)
Year -2	0.008	0.009	0.013	0.008	0.006
	(0.006)	(0.007)	(0.010)	(0.006)	(0.004)
Year -3	0.016**	0.018*	0.023	0.016**	0.011*
	(0.006)	(0.009)	(0.014)	(0.006)	(0.006)
Controls:					
Firm age	✓	✓	✓	✓	✓
Firm age squared	✓	✓	✓	✓	✓
Bank size	✓		✓	✓	
Fixed effects:	county × year × industry	county × year × industry	county × year × industry	county × year × industry	county × year × industry
Observations	348,628	10,734,870	509,372	348,628	31,414,061

Note: This table presents results from regressions analyzing the effect of bank failure on firm survival. The dependent variable is a binary indicator of whether a firm is still operating. Independent variables are event time dummies representing years relative to bank failure. The sample includes firms that had a relationship with a bank that eventually failed, with specific sample restrictions varying across columns. All regressions include county-by-year-by-industry fixed effects and control for firm age and age squared. Column 1 shows results for our headline regression specification. Column 2 shows results for a regression in which we do not control for bank size which significantly increases our sample size due only a fraction of the banks in our dataset reporting loan book values. Column 3 shows results for our headline regression where firms maintain treatment status even if the firm switches to another bank in the two years prior to failure and firms maintain control status even if they initiate a credit relationship with the failed bank within the two-year window preceding its failure. Column 4 shows results for our headline regression specification with SE clustering at the firm and year level, as opposed to the bank and year level. Column 5 show the results for our headline regression specification where the control group is expanded to include firms that borrow from non-bank lender. Standard errors in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Two-way standard error clustering by firm and year.

Table 5: Pooled Estimates of Effect on Firm Survival for Park National, Colonial, and All Banks

	Park National	**Colonial**	**All Banks**
Treatment	-0.018***	-0.068***	-0.045***
	(0.00367)	(0.00100)	(0.00542)
Year + Industry FE	Yes	Yes	Yes
Firm Age Control	Yes	Yes	Yes
SE Clustering	Firm and Year	Firm and Year	Firm and Year
N	96,554	143,944	215,043

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. This table presents pooled estimates of the effect of bank failure on firm survival for Park National Bank, Colonial Bank, and the aggregate of all banks. The dependent variable is a binary indicator of firm survival. The "Treatment" row shows the estimated average treatment effects, where negative coefficients indicate a decrease in the probability of firm survival following a bank failure. Standard errors are reported in parentheses. All models include year and industry fixed effects, control for firm age, and employ standard error clustering at the firm and year levels. "N" denotes the number of firm-year observations included in each specification.

Table 6: Top 25 Bank Failures by Number of Attached Businesses

Bank Name	Failure Date	Businesses Attached	Percentage (%)	Cumulative Percentage (%)	Total Assets at Failure ($ mns)	Resolution Type
First Republic Bank	1 May 2023	2,229	13.94	13.94	212,639	P&A
Sterling Bank	23 Jul 2010	2,000	12.50	26.44	355	P&A
Colonial Bank	14 Aug 2009	931	5.82	32.26	25,455	P&A
Silicon Valley Bank	7 Mar 2023	837	5.23	37.49	209,026	P&A
Hillcrest Bank	22 Oct 2010	580	3.63	41.12	1,584	P&A
United Commercial Bank	6 Nov 2009	576	3.60	44.72	10,895	P&A
Advanta Bank Corp.	19 Mar 2010	519	3.24	47.96	1,526	PO
Washington Mutual Bank	25 Sep 2008	514	3.21	51.18	307,022	P&A
First Community Bank	28 Jan 2011	470	2.94	54.12	2,188	P&A
Community National Bank	17 Dec 2010	404	2.53	56.64	32	P&A
Mirae Bank	26 Jun 2009	297	1.86	58.50	481	P&A
Bankfirst	17 Jul 2009	268	1.68	60.18	211	P&A
Main Street Bank	10 Oct 2008	263	1.64	61.82	112	P&A
California National Bank	30 Oct 2009	249	1.56	63.38	7,781	P&A
Guaranty Bank	5 May 2017	243	1.52	64.90	1,032	P&A
Western National Bank	16 Dec 2011	225	1.41	66.30	163	P&A
Texas Community Bank	13 Dec 2013	225	1.41	67.71	159	P&A
Irwin Union Bank And Trust	18 Sep 2009	204	1.28	68.98	2,840	P&A
San Diego National Bank	30 Oct 2009	201	1.26	70.24	3,595	P&A
First Southern Bank	17 Dec 2010	192	1.20	71.44	192	P&A
Fidelity Bank	30 Mar 2012	162	1.01	72.45	818	P&A
First Commercial Bank FL	7 Jan 2011	152	0.95	73.40	579	P&A
1St Centennial Bank	23 Jan 2009	147	0.92	74.32	798	P&A
First National Bank CF	29 Apr 2011	128	0.80	75.12	342	P&A
Independent National Bank	20 Aug 2010	128	0.80	75.92	156	P&A

Note: Resolution Type: P&A: Purchase and Assumption, PO: Payout

A Data appendix

A.1 Linking UCC filings to Orbis bank records

Since the UCC filing data does not contain an creditor identifier with which to match the lender on the form to our bank-level data, we implement a matching procedure based on name similarity and location. As a first step, we parse the 495,164 unique lender names in our dataset using OpenAI's `GPT4o` LLM to extract for each lender name string a "commonly referrered to" name. For example the raw string "'(1) FIRST NATIONAL BANK OF GILMER, MAIN BRANCH LOCATION" is parsed to "First National Bank of Gilmer". This step helps standardize the creditor names in the UCC text corpus, which improves matching quality when linking these data to our Orbis bank-level data.

As a next step, we create vector embeddings for each of the unique creditor names (post-parse) in our UCC dataset using the OpenAI `large3` embedding model with dimension 3,072. We then use these vector embeddings to train a Facebook AI Similarly Search (FAISS) index. FAISS is a set of tools developed by the Fundamental AI Research group at Meta (formerly Facebook) for efficient similarity search and clustering of dense vectors. It is particularly useful for large-scale nearest-neighbor searches in high-dimensional spaces, enabling very fast lookups of similar embeddings.[34]

We then also embed the 27,603 unique bank names in our bank-level data using the same `large3` embedding model to ensure consistency in the vector space. Once both datasets—UCC creditor names and bank names—are embedded, we use the FAISS index to perform approximate nearest-neighbor searches, identifying the closest matches between the creditor names in the UCC data and the bank names in the Orbis data. We only keep nearest neighbors with a cosine similarity of 0.8 or more.

Once we have established nearest-neighbor clusters between UCC creditors and Orbis bank names based on their vector embeddings, we enhance these candidate matches by merging in additional information. Specifically, we obtain branch-level addresses for 171,370 branches of banking organizations operating in the US. These data are obtained from the National Information Center (NIC) which is a repository of financial data collected by the Federal Reserve. The branch-level address data goes back to the 1970s. We enrich our candidate matches with branch level information so that for every candidate match we know whether the Orbis bank in the candidate match has ever had a bank branch in the same state, county, or ZIP code as the

[34] For more information on FAISS, see https://github.com/facebookresearch/faiss.

lender listed on the UCC form. We also include a count of the number of banks with the same bank name in the Orbis dataset. For example, only one bank (as defined by having a unique Bureau van Dijk identifier) is called “Rockwood Bank”, whereas 17 banks in the Orbis data share the name “First State Bank". This information captures the distinctiveness of each bank name, which is an important piece of information when training our matching model as described below. The resulting dataset contains 310 million candidate matches.

To develop our matching algorithm, we manually label 1,100 candidate matches to create a training set for a machine learning model. Using this labeled data, we estimate a LASSO logistic regression model that predicts the probability of a true match between a UCC filing and an Orbis bank. The model’s covariates are the cosine similarity of the names, whether the bank has ever had a branch in the state, county, or ZIP code as the lender listed on the UCC filing, and our proxy for name uniqueness as well as transformations—such as natural logs, standardization, squares, and interaction terms—for a total of 88 possible covariates. The LASSO method simultaneously selects relevant predictors and applies regularization, effectively reducing the coefficients of less influential predictors to zero. This approach aims to strike an optimal balance between model parsimony and predictive power. We use 5-fold cross-validation to determine the optimal C parameter for the LASSO regression, where C controls the strength of the regularization.

To further refine our model, we implement an active learning strategy using uncertainty-based sampling. This iterative process involves manually labeling the model’s most uncertain predictions, allowing us to efficiently enhance its performance. The outcome is a matching model that estimates the likelihood of a match between a creditor listed in a UCC filing and a bank in our bank-level data. We then apply this model to predict match probabilities for our 310 million candidate matches. Figure 13 visualizes the estimated matching model’s performance across a variety of metrics.

In a final step, we select for each UCC filing in our candidate match dataset the Orbis bank match with the highest predicted match probability and drop any matches with a predicted match probability less than 80%.

This process allows us to match 4,962,762 unique filings to 7,467 unique banks. Our entire corpus of filings in which a lender is listed consists of approximately 24M filings of which approximately 20M filings are made with lenders that are classed as “organizations”’—as opposed to natural persons. Gopal and Schnabl (2022) estimate that about 50% of lenders to small businesses are banks—as opposed to alternative lenders such as captive finance companies, FinTech companies, etc.—which suggests that about 10M of our filings are made by banks and that our match rate of UCC filings to the Orbis banks dataset is close to 50%.

Figure 13: Performance Metrics for Record Linking Model (UCC to Orbis)

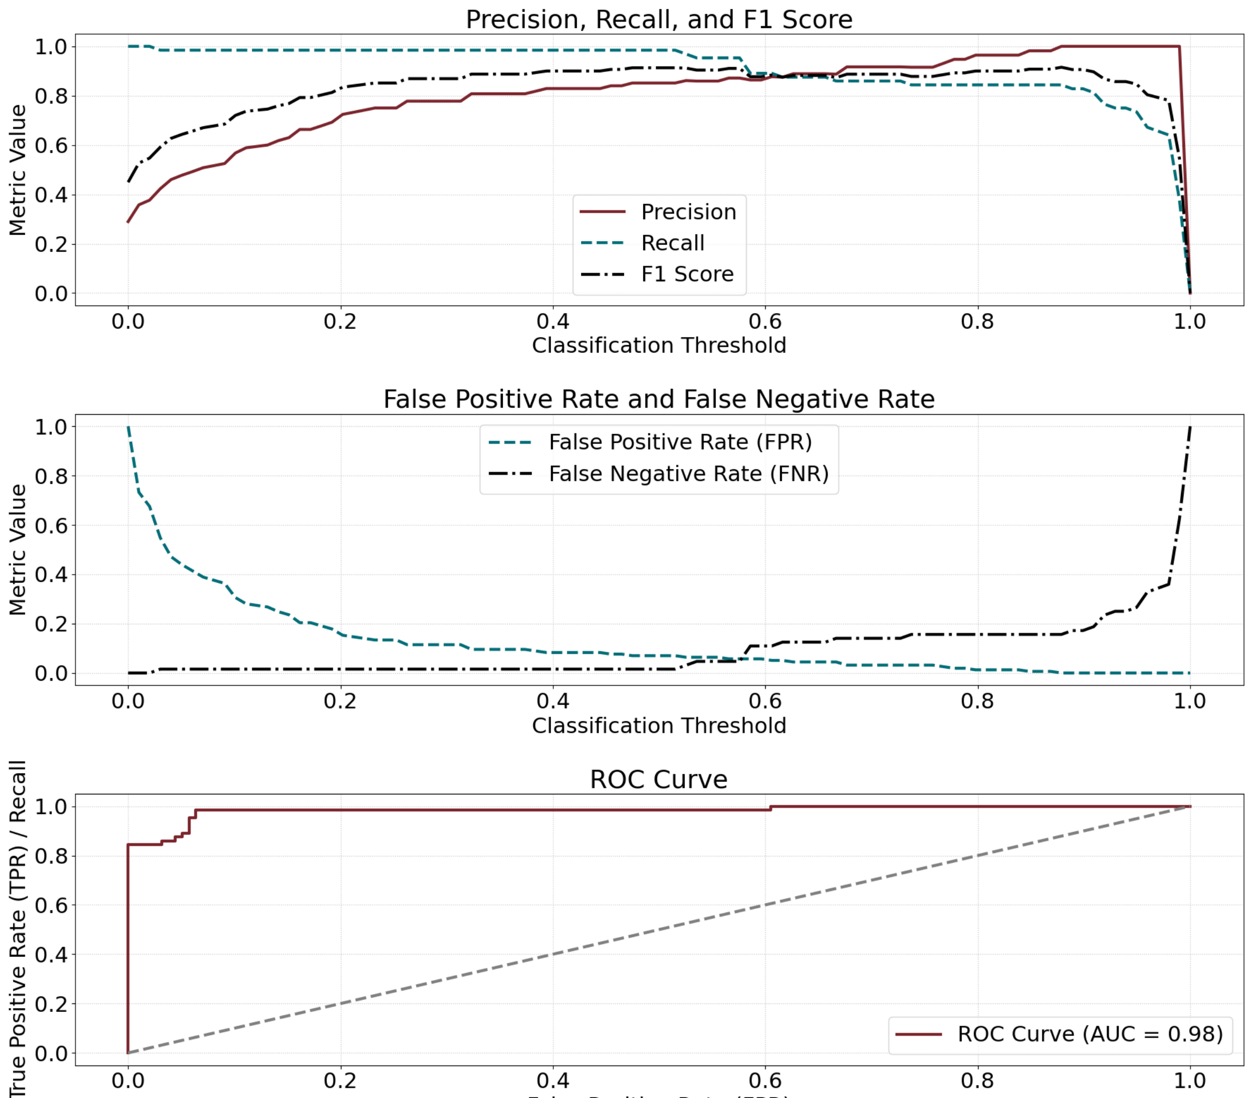

Note: This figure presents three key visualizations for the Record Linking Model (UCC to Orbis). **(A)** The top panel shows how Precision (burgundy solid line), Recall (blue-green dashed line), and F1 Score (black dashed-dotted line) change with varying classification thresholds. Precision is the probability that a match classified by the model is a true match, and recall is the probability that a true match is classified as a match by the model. The F1 score is the harmonic mean of precision and recall. **(B)** The middle panel displays the False Positive Rate (FPR, blue-green dashed line) and False Negative Rate (FNR, black dashed-dotted line) across different classification thresholds. The FPR is the fraction of records that are truly not a match but are classified as a match by the model. The FNR is the fraction of records that are truly a match but are classified as not a match by the model. **(C)** The bottom panel shows the Receiver Operating Characteristic (ROC) curve (burgundy line) for the model evaluated on the test set, with the gray dashed line representing a random classifier (AUC = 0.5). The Area Under the Curve (AUC) value is indicated to reflect the model's discriminative power.

Table 7: Summary Statistics for UCC Filings Matched to Banks

Statistic	Value
Number of unique filings	4,962,762
Number of unique banks	7,467
Cosine Similarity	
Average	0.98
Median	1.00
10th percentile	0.94
90th percentile	1.00
Match Probability	
Average	0.97
Median	0.99
10th percentile	0.94
90th percentile	0.99
Fraction with state match	1.00
Fraction with county match	0.89
Fraction with zip match	0.61

Note: This table presents summary statistics for UCC filings matched to banks. Cosine similarity and match probability statistics are based on the matching algorithm used.

A.2 Validation of UCC Filings Data Using CRA Small Business Loans Data

This section discusses the data validation exercise we conduct by comparing our loan-level UCC data to census-tract level lending data contained in the Community Reinvestment Act (CRA) dataset. The purpose of this exercise is to verify the extend to which UCC filings reflect small business lending.

The CRA, enacted by Congress in 1977, aims to encourage depository institutions to meet the credit needs of all segments of their communities, including low- and moderate-income neighborhoods. Under the CRA, banks with assets exceeding a specified threshold are required to report detailed information on their small business lending activities. The CRA data is considered a reliable and comprehensive source of information on small business loans, as it includes the number and dollar amount of loans originated, categorized by loan size and geographic location at the census tract level.

We obtain the CRA data from the Federal Financial Institutions Examination Council (FFIEC) for the years 2000 through 2022. The dataset includes information on:

- Number of loan originations.
- Total dollar amount of loans.
- Loan size categories (loans up to $100,000; $100,001 to $250,000; and $250,001 to $1,000,000).
- Geographic identifiers, including census tract, county, and state.

We filter the data to include only loan originations within our five UCC states. The data were then aggregated at the census tract-year level to create a panel dataset suitable for comparison with the UCC filings data.

From our collected UCC data, we extract the following relevant information:

- Filing date.
- Debtor's name and address.
- Secured party's (creditor's) name and address.
- Type of filing (e.g., initial financing statement, amendment, continuation).

We then geocode the addresses using the UC Census Bureau geocoding API to obtain geographic identifiers for each debtor, in particular for their census tract. Similar to the CRA data, we aggregate the UCC filings at the census tract-year level.

To compare the UCC filings count data against loan origination in the CRA data, we examin the correlation between the number of UCC filings and the number of CRA-reported small business loan origination at the census tract-year level. Figure 14 shows a binscatter of the log of UCC filing counts vs the log of CRA loan originations below $100,000 at the census tract-year level.

Figure 14: Binscatter Plot of UCC Filing Count vs CRA Originations

Note: This figure uses the binscatter technique to visualize the relationship between UCC filing counts and CRA loan originations at the census tract-year level. The binscatter method partitions the data into equal-sized bins based on the independent variable and plots the mean of the dependent variable within each bin. This non-parametric approach reveals the underlying relationship without imposing a specific functional form.

We performed regression analysis to quantify the relationship between UCC filings and CRA loan originations. The primary regression model is specified as:

$$\ln(\text{UCC Filings}_{ict}) = \beta_0 + \beta_1 \ln(\text{CRA Loans}_{ict}) + \gamma_c + \delta_t + \varepsilon_{ict} \tag{8}$$

where $\ln(\text{UCC Filings}_{ict})$ represents the natural logarithm of the count of UCC filings in census tract i, county c, year t. The term $\ln(\text{CRA Loans}_{ict})$ denotes the natural logarithm of the count of CRA-reported small business loan originations in the same census tract and year. We include γ_c to represent county fixed effects, which control for unobserved, time-invariant heterogeneity at the county level. Similarly, δ_t represents year fixed effects, accounting for temporal shocks common to all counties. Finally, ε_{ict} is the error term.

The results of the regression analysis are presented in Table 8. The model includes county and year fixed effects, and standard errors are clustered at the county level to account for potential serial correlation within counties over time.

Table 8: Regression Results: Log UCC Filings on Log CRA Loan Originations

	Coefficient (Standard Error)
$\ln(\text{CRA Loans}_{ict})$	0.795^{***}
	(0.005)
County Fixed Effects	Yes
Year Fixed Effects	Yes
Observations	325,040
R-squared	0.546

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table presents results from a regression analyzing the relationship between UCC filings and CRA loan originations. The dependent variable is the natural logarithm of UCC filings. The independent variable is the natural logarithm of CRA-reported small business loan originations. The regression includes county and year fixed effects. Robust standard errors, clustered at the county level, are reported in parentheses.

The coefficient on is 0.795 and is statistically significant at the 1% level. This indicates that a one percent increase in CRA-reported small business loan originations is associated with approximately a 0.8 percent increase in UCC filings, holding county and year effects constant. The R-squared value of 0.546 suggests that over half of the variation in UCC filings is explained by the model.

The observed positive correlation between UCC filings and CRA-reported small business loan originations—in addition to the institutional background of UCC filings discussed in the body of the paper—suggests that UCC loan filings are very likely to reflect small business lending.

A.3 Comparison with other US loan-level datasets

Studying the firm-bank relationships of very small firms in the US has traditionally been difficult because of lacking loan-level data for smaller businesses. While other countries maintain national credit registers which can be used to study firm-bank relationships of smaller firms, no such credit register exists in the US.[35].

The often-used Dealscan dataset predominantly covers syndicated loans, which are loans involving multiple participating banks and typically extended to larger, more maturel, and more sophisticated borrowers. The average firm in Dealscan borrows nearly $300 million per loan and employs nearly 3,000 employees (see e.g. Chodorow-Reich (2014)).

The Shared National Credit (SNC) progam operated jointly by the FDIC, the Federal Reserve, and the Office of the Comptroller of the Currency (OCC) is also not suitable for studying the impact of financial shocks on smaller firms. The SNC dataset exclusively covers large syndicated loans and the average loan commitment in the SNC is $189 million, with the average borrower having $3.5 billion in assets (Mian and Santos 2018). In a similar vein, regulatory data on loans collected by the Federal Reserve's Capital Assessments and Stress Testing information collection exercise (FR Y-14) only contains data on loans over $1 million made by bank holding companies with over $100 billion in assets. The median firm in this dataset has $21.5 million in assets, again meaning that small business lending is not covered (Greenwald, et al. 2023). Table 9 provides an overview of the three US loan-level datasets discussed above.

Table 9: Comparison of US loan-level datasets

Dataset	**Focus**	**Average Loan/Firm Size**
Dealscan	Syndicated loans	$300M / 3,000 employees
Shared National Credit (SNC)	Large syndicated loans	$189M loan / $3.5B assets
FR Y-14	Loans >$1M	Median $21.5M assets

[35] See Berton, et al. (2018) or Bentolila, Jansen, and Jiménez (2018) for examples of studies that leverage such credit registers in Italy and Spain respectively.

A.4 UCC Filing Matching Rate to Active D&B Firms

Figure 15: Fraction of UCC Filings Matching to Active D&B Firms by Time Margin

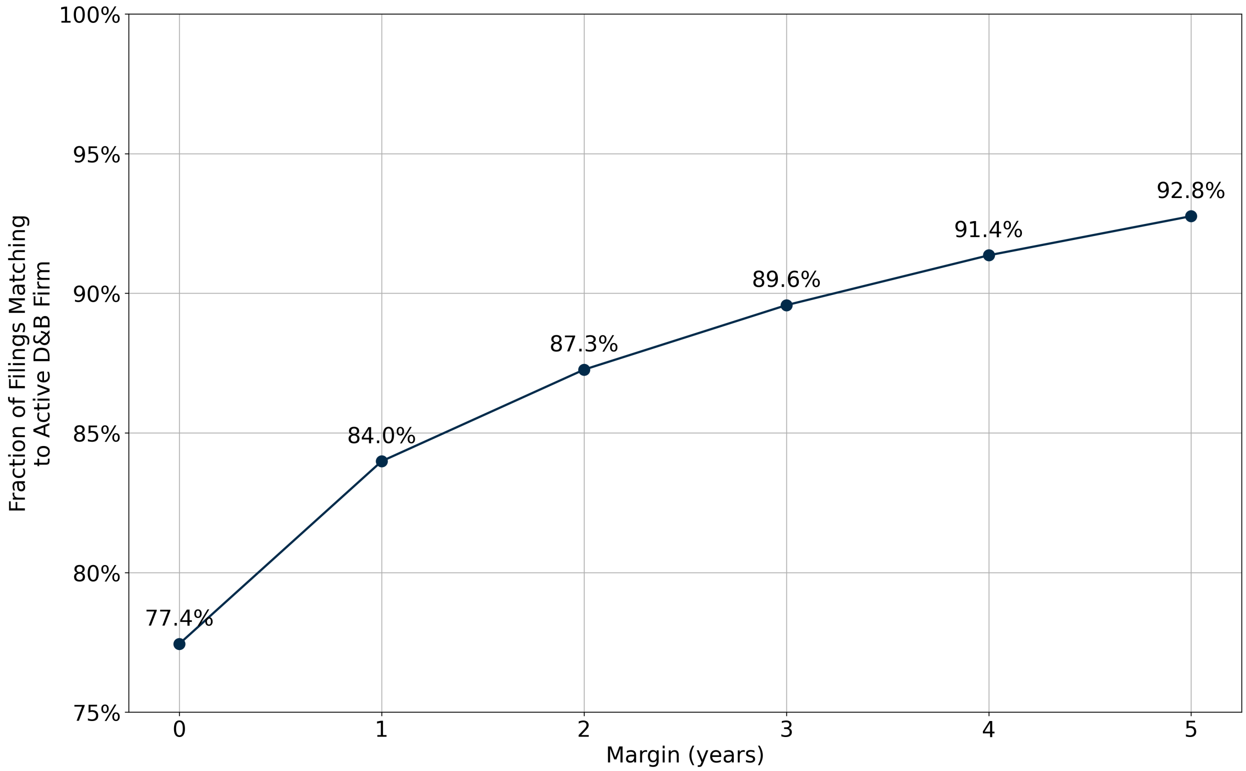

Note: This figure shows the fraction of UCC filings that match to firms active in the D&B database as the time margin for matching is expanded. The analysis verifies that exit from the D&B database indicates operational cessation and validates the UCC filing-to-firm matching procedure. The x-axis represents the time margin in years, allowing matches before a firm's entry or after its exit from the database. The y-axis shows the percentage of UCC filings that match to active firms. At margin 0, 77.4% of filings match firms active in the exact year of filing, increasing to 84.0% when allowing matches within a 1-year margin, suggesting high consistency between UCC filings and firm activity records in the D&B database.

B Event study analysis

To complement the DiD estimates presented in Figure 3, we conduct an event study of firm survival around the failure year of the bank using the following equation:

$$y_{ith} = \sum_{j \neq -1} \zeta^h \mathbb{1}\{h = j\} + \boldsymbol{\gamma} \mathbf{X}_{it} + \nu_{ith} \tag{9}$$

where y_{ith} is a binary variable indicating whether the firm is operational, $\mathbb{1}\{h = j\}$ represents an event time dummy for time horizon h, $\mathbf{X}_{it}$ is a set of controls, and ν_{ith} is an error term.

We control for the same set of variables as in our headline DiD regressions, namely county-by-year-by-industry fixed effects, firm age, and firm age squared. To include never-treated firms in the analysis, we assign them placebo "failure years" randomly drawn from the distribution of actual failure years observed among treated firms.

Figure 16: Event Study of Firm Survival Rates

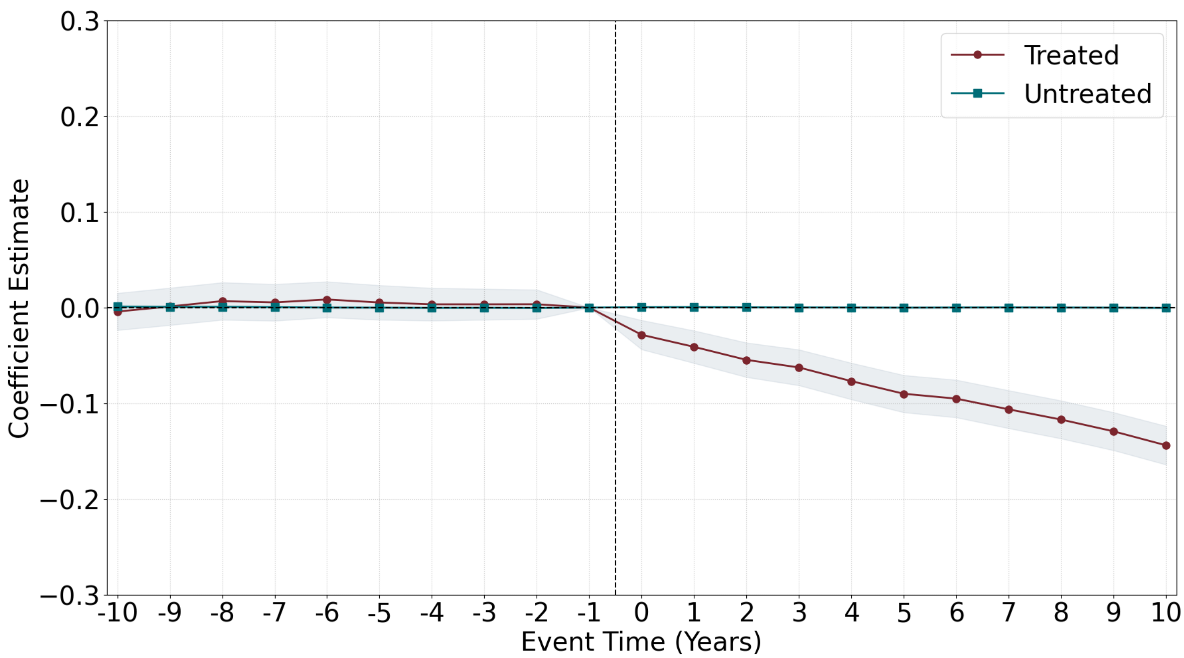

Note: This figure presents the event study results for firm survival rates. The x-axis shows event time relative to bank failure (or placebo failure for never-treated firms), with 0 representing the failure year. The y-axis displays the estimated coefficients, representing the difference in survival probability between treated and control firms. Shaded areas indicate 95% confidence intervals. The vertical dashed line at t=-1 marks the pre-treatment period.

Figure 16 presents the results of this event study. Prior to bank failure (event time < 0), both treated and control firms exhibit similar, flat trends in survival rates. The lack of pre-failure decline in survival rates among treated firms suggests that at least at the annual level and for this particular observable, firms did not begin to cease operating at increased rates prior to the bank failure.

After the bank failure (event time 0), the survival rates of treated firms show a negative trend compared to control firms, which maintain a flat trend. This divergence continues and widens over several years post-failure. The flat trend for control firms persists across the entire period, including after their assigned placebo failure dates.

C Switching analysis

Figure 17: Likelihood of switching to a healthy bank prior to bank failure

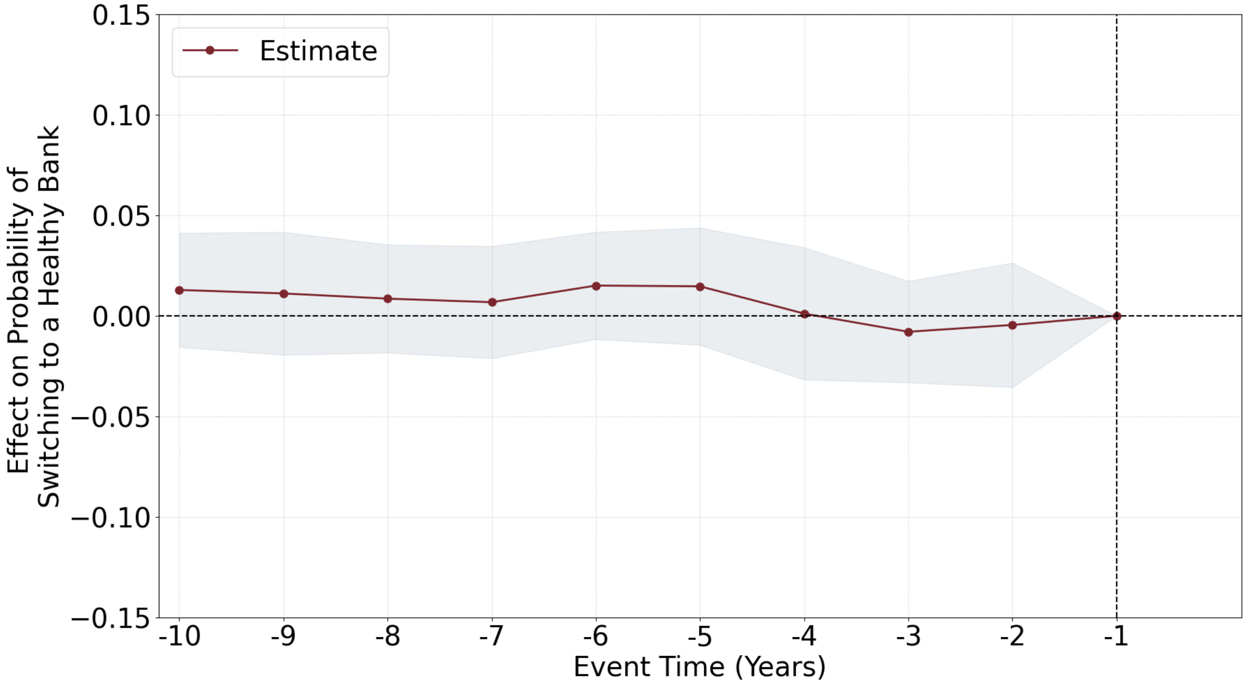

Note: This figure displays the estimated probability of a firm switching its credit relationship to a healthy bank in the years leading up to the failure of its original bank. The x-axis represents the number of years before bank failure, with 0 being the year of failure. The y-axis shows the change in probability of switching compared to the baseline year (Year 0). A positive value indicates an increased likelihood of switching to a healthy bank compared to the year of failure. For example, a value of 0.1 at Year -5 would mean that firms were 10 percentage points more likely to switch to a healthy bank five years before their original bank's failure compared to the year of failure. The shaded area represents the 95% confidence interval around the point estimates. The regression includes county-by-year-by-industry fixed effects and clusters standard errors at the firm and year level. The sample includes all firms that had a relationship with a bank that eventually failed. See Table 10 for detailed regression results.

Table 10: Likelihood of switching to a healthy bank prior to bank failure

	Estimate	**Confidence Interval**
Year -10	0.013*	(0.002, 0.024)
Year -9	0.011*	(0.001, 0.020)
Year -8	0.008	(-0.001, 0.017)
Year -7	0.007	(-0.002, 0.015)
Year -6	0.008*	(0.001, 0.015)
Year -5	0.004	(-0.004, 0.012)
Year -4	0.001	(-0.009, 0.011)
Year -3	0.001	(-0.010, 0.012)
Year -2	0.005	(-0.004, 0.013)
Year -1	0.002	(-0.002, 0.007)
Year 0 (Omitted)	0.000	(0.000)
Fixed effect controls: county × year × industry		
SE clustering: firm + year		
N	210,814	
Firms	17,083	

Standard errors in parentheses: $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Note: This table presents results from a regression analyzing the likelihood of firms switching their credit relationships to healthy banks prior to their original bank's failure. The dependent variable is a binary indicator of whether a firm has switched to a healthy bank. Independent variables are event time dummies representing years until bank failure. The sample includes all firms that had a relationship with a bank that eventually failed. We include county-by-year-by-industry fixed effects, consistent with our main regressions on firm survival and employment outcomes. See Figure 17 for a plot of these estimates.

D Appendix Figures

Figure D.1: FDIC Bank Failures and Assistance Transactions

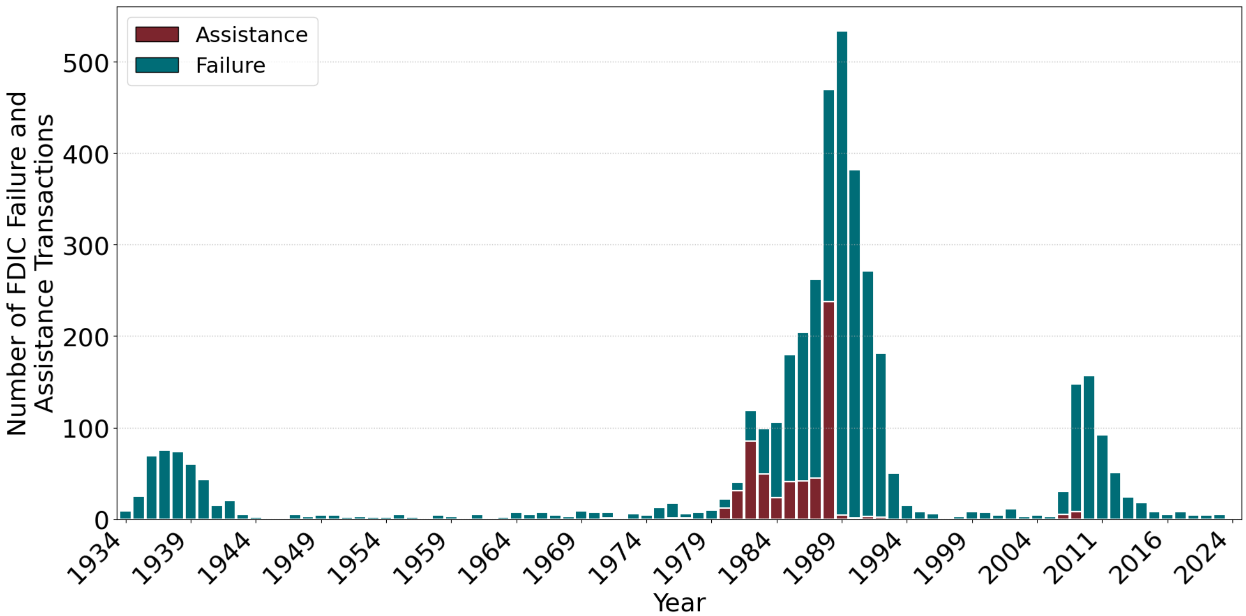

Note: This figure illustrates the number of FDIC-insured bank failures and assistance transactions from 1934 to 2024. The teal bars represent bank failures, while the maroon bars indicate assistance transactions. Notable peaks in bank failures are observed in the late 1930s (likely due to the Great Depression aftermath), the late 1980s to early 1990s (corresponding to the Savings and Loan Crisis), and around 2008-2010 (coinciding with the Global Financial Crisis). Assistance transactions, which were more common in the 1980s, have become less frequent in recent decades.

Figure D.2: Uninsured Deposits as a Share of Total Domestic Deposits (1990-2025)

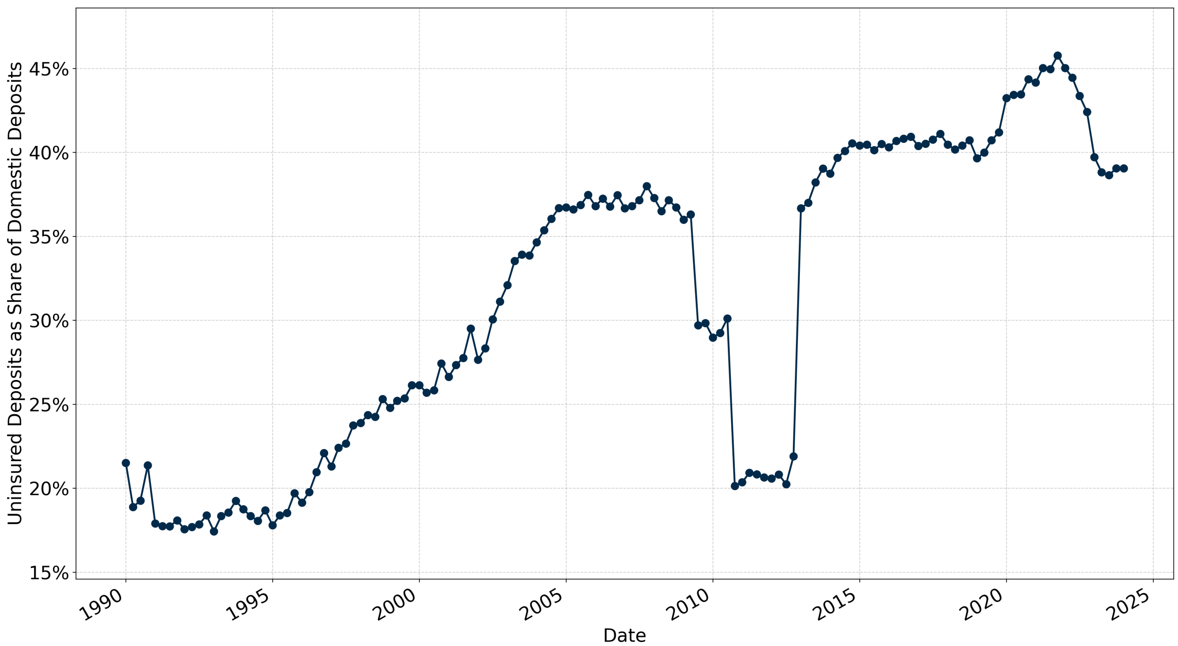

Note: This figure shows the percentage of uninsured domestic deposits in the US banking system from 1990 to 2025.

Figure D.3: Effect of Bank Failure on Firm-Level Survival Rates (by industry)

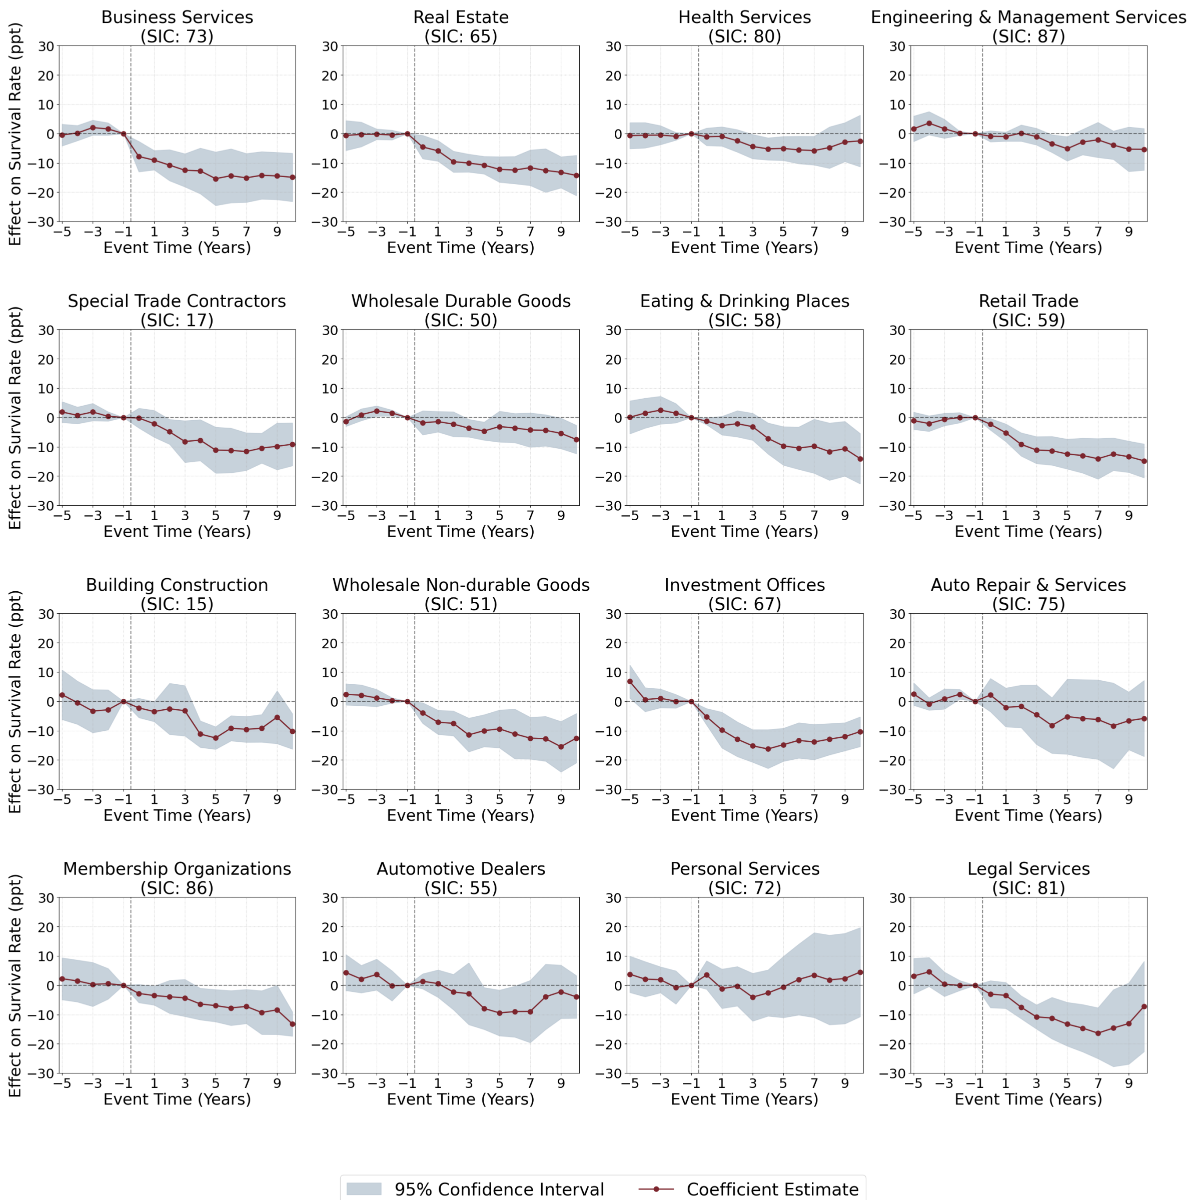

Note: This figure shows estimates of the effect of bank failures on firm survival rates across the 16 largest two-digit SIC industries, sorted by number of firms in our dataset. Each panel presents regression coefficients and 95% confidence intervals from a local projection staggered treatment DiD specification comparing firms whose bank fails (treated) to control firms (consisting of not-yet-treated firms and never-treated firms). The x-axis shows event time in years, where 0 marks the year of bank failure. The y-axis shows the effect on firm survival rates in percentage points. All specifications include firm age, firm age squared, and firm size (as proxied by employees) as controls, along with county-by-year fixed effects. The coefficients are estimated relative to event time -1. The shaded areas represent 95% confidence intervals based on standard errors clustered at the firm and year levels. Industry classifications follow the Standard Industrial Classification (SIC) system, with industry names and codes shown in each panel's title.

Figure D.4: Pooled Average Treatment Effects on the Treated (ATT) of Bank Failure on Firm-Level Survival Rates (by industry)

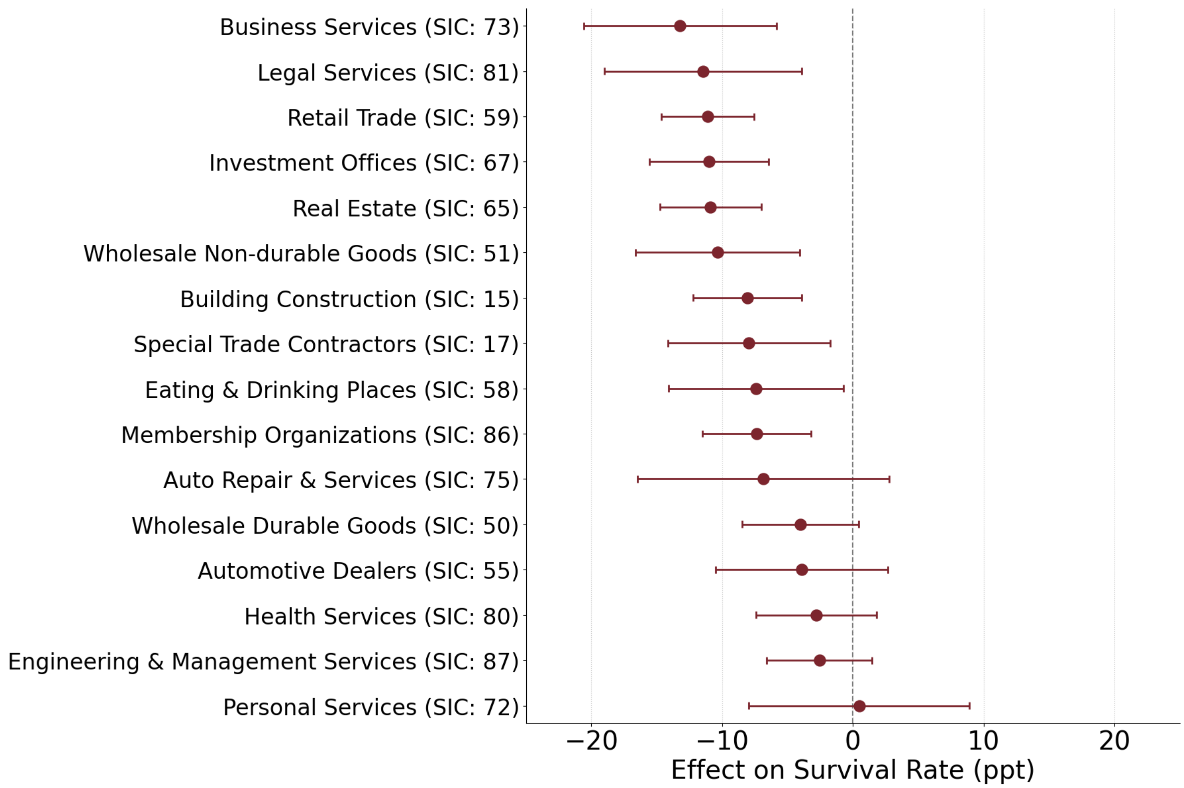

Note: This figure shows the average treatment effect on the treated (ATT) of bank failures on firm survival rates across the 16 largest two-digit SIC industries, sorted by number of firms in our dataset. The estimates reflect the pooled average effect across all post-treatment periods from a staggered treatment DiD specification comparing firms whose bank fails (treated) to control firms (consisting of not-yet-treated firms and never-treated firms). The x-axis shows the estimated effect on survival rates in percentage points. Each dot represents the point estimate for a specific industry, with horizontal bars indicating 95% confidence intervals based on standard errors clustered at the firm and year levels. All specifications include firm age, firm age squared, and firm size (as proxied by employees) as controls, along with county-by-year fixed effects. Industry classifications follow the Standard Industrial Classification (SIC) system, with industry codes shown in parentheses. The vertical dashed line at zero provides a reference for interpreting effect sizes.