NBER WORKING PAPER SERIES

WHY DO FIRMS PAY DIFFERENT INTEREST RATES ON THEIR BANK LOANS?

Mary Amiti
Anil K Kashyap
Anna Kovner
David Weinstein

Working Paper 34870
http://www.nber.org/papers/w34870

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026

We thank Sean Fulmer, Augustin Belin, Leeya Howley, Lillian Zhu, and Hriday Karnani for excellent research assistance, and seminar participants at the Federal Reserve Banks of New York and Richmond, the Harvard Business School, Chicago Booth, the London School of Economics, Stanford GSB, UCLA, and Wharton for helpful comments. The opinions expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Banks of New York and Richmond, the Federal Reserve System, or the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w34870

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Mary Amiti, Anil K Kashyap, Anna Kovner, and David Weinstein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Why Do Firms Pay Different Interest Rates on Their Bank Loans?
Mary Amiti, Anil K Kashyap, Anna Kovner, and David Weinstein
NBER Working Paper No. 34870
February 2026
JEL No. E51, G21, G32

ABSTRACT

We document significant variation in interest rates among similar commercial and industrial loans using confidential supervisory data on the largest US banks. This dispersion does not appear to be due to risk. We rationalize the data using a search cost model and find that search costs are highest for smaller and riskier borrowers and lower for public firms, consistent with predictable differences in the costs of screening and monitoring. We find that search costs are substantial. Over a third of firms behave as if they do not comparison shop; half of all firms appear to only obtain two quotes before picking a lender, while the remaining firms behave as if they search widely.

Mary Amiti
Federal Reserve Bank of New York
and CEPR
mary.amiti@gmail.com

Anil K Kashyap
University of Chicago
Booth School of Business
and NBER
anil.kashyap@chicagobooth.edu

Anna Kovner
Federal Reserve Bank of Richmond
anna.kovner@rich.frb.org

David Weinstein
Columbia University
Department of Economics
and NBER
dew35@columbia.edu

1 Introduction

There is a large literature in which economists regress loan interest rates on "sensible" explanatory variables such as credit ratings, relationship length, loan amount, and maturity (e.g. Petersen and Rajan (1994).) This literature finds that while these variables matter, they do not account for much of the variance in interest rates. While economists typically appeal to information or contracting frictions to explain the residuals in these regressions, precisely what these frictions are remains a mystery. Our lack of understanding is troubling, given the importance of these frictions in explaining why a bank's health matters for its customers' performance and in identifying the causes of capital misallocation.

We aim to fill this gap by analyzing confidential supervisory data on more than 160,000 business loans made in the U.S. between 2012 and 2024, which enables us to explain the cross-firm dispersion in interest rates by estimating the magnitudes of a particular friction: search costs. We show that our estimates are reasonable, in the sense that they closely match estimates derived from survey data and the patterns predicted by the screening-and-monitoring theory of intermediation.

Our data covers all loans of $1 million and more made by the largest U.S. banks. These data are collected for supervisory purposes and include a comprehensive set of risk proxies, in addition to the "loan spreads" (i.e., the difference between the interest rate on a newly originated loan and a reference rate, such as the London Interbank Offered Rate, LIBOR), paid by different borrowers. While existing studies have largely focused on small loans with values of less than $1 million (both in the US and in foreign countries) or the very largest syndicated loans, we study a part of the US loan market that is usually impossible to analyze. We exploit the richness of the data to establish four new facts and obtain estimates of search costs that can explain the dispersion in interest rates in the U.S. economy.

First, we show that the primary source of loan spread dispersion is not differences between risk categories, but differences within them. We demonstrate this by grouping loans into 722 bins with *identical* risk ratings, rate structures (i.e., floating or fixed), and nearly identical loan purposes, sizes, and maturities. By focusing on loans within the same bin, we allow for an *arbitrary* relationship between the loan spread and the characteristics used to form the bins. While we find a robust relationship in which riskier loans have higher spreads, when we regress loan spreads on bin fixed

effects, these fixed effects account for only one-quarter of the variation in spreads. These results are not driven by outliers. The mean difference between the ninetieth and tenth percentiles of the loan spread distribution within a bin is 241 basis points: much more than the difference in spreads between B-rated loans (that make up the bottom 16 percent of the loan-risk distribution in the sample) and A-rated ones (that constitute the top 1 percent). Adding other explanatory variables to the regression reveals that many plausible drivers of interest rate determination matter, but only a little. For example, when we add several other variables that the banking literature suggests should determine loan spreads, such as relationship length, borrower size, and origination date, listing status, collateral type, industry, Metropolitan Statistical Area, and bank-bin fixed effects, for a total of 7,000 controls, we still find that half of the variation in loan spreads within a bin remains unexplained.

Our second stylized fact is that differences in spreads within bins are not due to differences in banks' risk estimates. We demonstrate this by studying a subset of loans for which banks report their expected loss in the event of default. The estimates of expected losses within bins are much tighter than the distribution of loan spreads, making it implausible that differences in risk ratings could drive the results. Moreover, when we control for expected losses, the remaining variation in loan spreads is hardly affected. Controlling for banks' estimates of expected losses, however, helps rule out several other plausible factors that could account for loan spread dispersion. Since expected losses are the product of the probability of default, the loss given default, and the exposure at default, they arguably account for any differences in risk perceived by the lender.[1] In particular, banks' estimates of expected losses take into account any covenants that could alter a loan's riskiness. Likewise, expected losses are loan-specific and should capture the component of relationship strength that operates through perceived credit risk (or through any unobserved aspect of collateral).

Third, there is a pattern in the variance of loan spreads within bins: it declines with loan quality and the ease of obtaining information about borrowers. Bins containing only highly rated A loans exhibit about half the variation observed in bins containing lower-rated B loans. Loans to publicly traded firms, whose financial statements are audited and readily available, exhibit less

[1]Capital-related risk components could, in principle, vary within bins even when expected loss does not, but these are unlikely to explain the dispersion of the magnitude we document.

spread variation than loans to private borrowers.

Fourth, we demonstrate that debt instruments whose risk is well understood exhibit much less unexplained variation in spreads. We show this result by analyzing corporate bonds, whose risk is publicly assessed by rating agencies, making them an example of a debt instrument whose risk is easier for creditors to learn and monitor. We find that the spreads on corporate bonds, grouped into risk bins comparable to those used in our loan data, are more tightly distributed. Bin fixed effects explain over 75 percent of the variation in bond spreads: three times the amount for loans. In other words, the unexplained interest rate dispersion is much lower when creditors can easily access information about firms' risk and borrowers need not search for creditors.

We interpret these four facts through the lens of the Burdett and Judd (1983) model, which we use to endogenize the distribution of observed loan spreads within each bin. We argue that search costs are a plausible source of variation in interest rates for nearly identical loans, as survey evidence indicates that these costs are substantial for small and medium enterprises (SMEs). For example, (Wiersch et al., 2022, p. 11) find that 39 percent of SMEs complain of "difficult loan application processes" when borrowing from large banks and 37 percent complain of "long waits for credit decisions or funding." Similarly, (Federal Reserve Banks, 2015, p. 6) states that businesses spend on average 24 hours "researching and completing financing applications" when they apply for loans. Although 96 percent of SMEs borrow less than one million dollars and are therefore not covered in our data (Federal Reserve Banks, 2015, p. 10), the survey evidence indicates that search costs are substantial for certain tranches of the loan market.

We use the Moraga-González and Wildenbeest (2008) empirical implementation of the Burdett and Judd (1983) model to estimate these search costs. We assume that a borrower who obtains interest rate quotes for identical loans from different lenders always chooses the quote with the lowest interest rate. Borrowers also pick the optimal number of interest rate quotes to request as a function of the distribution of interest rate offers from bank loan officers, the expected reduction in borrowing costs from sampling more lenders, and the cost of search (i.e., filing an additional loan application and responding to lender questions). We assume that borrowers differ in their search costs; therefore, firms with high search costs will accept the first offer they receive, whereas those with low costs may obtain offers from all banks. Bank loan officers maximize profits by optimizing the various rates that they offer. In equilibrium, banks randomize their interest rate quotes such

that they are indifferent between raising the frequency of high interest rate quotes (which would cause them to make fewer loans but earn more interest income per loan) and lowering the frequency of high interest rate quotes (which would enable them to make more loans at a lower profit per loan).

To the best of our knowledge, we are the first to estimate borrower search costs. We find that, on average, approximately 37% of borrowers receive a single loan quote and accept it, about 48% search for a second quote, and the remaining borrowers seem to search intensively enough to see all available quotes. This search behavior is consistent with the idea that firms that accept their first loan offer have a search cost of at least $168,000 when originating a loan of the average size in our sample ($22.36 million). We estimate firms that search twice have a cost of $72,000 to $168,000 per quote, and most of the remaining firms, which search intensively, have a cost of no more than $6,700 per quote. A back-of-the-envelope calculation reveals that these search costs can account for about a fifth of the stock-market response to the bank failure studied in Slovin et al. (1993).

We validate our estimates in two ways. First, we establish that the *absolute* levels of our estimates of search frequencies and search costs in bins containing small loans match those reported in small business surveys. The fact that our estimates align closely gives us confidence that the magnitudes we estimate for small firms are plausible. Second, having pinned down the validity of our estimates for one part of the distribution of loans, we validate the *relative* patterns of search costs by comparing the variation in our search cost estimates with how a prominent theory of intermediation, the "screening and monitoring hypothesis," predicts they should vary. In particular, search costs should be higher for borrowers who are harder to screen and monitor. Because the screening and monitoring hypothesis predicts how search costs should vary *across* bins and our search cost estimates are identified *only* using the distribution of loan spreads *within* bins, there is no mechanical reason to believe that the cross-bin variation in our search cost estimates should align with the predictions of the screening and monitoring hypothesis. We also have a meaningful null hypothesis. Many papers have modeled the unexplained portion of loan rates as arising from idiosyncratic matching frictions orthogonal to standard explanatory variables, such as risk. If this view is correct, there would be no reason to expect our search cost estimates to rise and fall with the screening and monitoring hypothesis's predictions. In contrast, we find that estimated search costs fall with borrower credit quality and transparency as the screening and monitoring

hypothesis predicts.

The remainder of the paper proceeds as follows. Section 2 reviews the literature on interest rate determination, search frictions in credit markets, its implications for borrowers, and the theoretical papers that underlie the screening and monitoring hypothesis. Section 3 introduces the data used in the analysis and highlights new stylized facts. Section 5 presents our estimates of the search costs and discusses their external validity. We then present estimates of these costs for different types of borrowers and loans, demonstrate the external validity of our estimates and close with an analysis of their economic significance. Section 6 concludes.

2 Related Literature

There are many studies that analyze borrowing costs for small businesses around the world (e.g., for the U.S. (Petersen and Rajan, 1994; Berger and Udell, 1995), Belgium (Degryse and Ongena, 2005), Chile (Bordeu et al., 2025), and France (Mazet-Sonilhac, 2025)). These studies have focused on identifying the correlates of loan interest rates but have given little attention to the poor fit of the regressions.[2] Our paper is closer to the work of Martín-Oliver et al. (2008), Cerqueiro et al. (2011), and Mazet-Sonilhac (2025), who argue that the residuals of regressions of interest rates on explanatory variables covary with variables that could be correlated with search costs (e.g., the number of banks, loan amounts, and broadband internet access). In contrast to this earlier work, our paper examines much larger loans, estimates search costs, and demonstrates the external validity of the estimates.

Our estimates can help inform the large literature that studies the effects of loan supply shocks. We now have convincing evidence that different shocks can lead banks to reduce lending to certain customers in a variety of settings. To give just a few examples of the range of shocks that have been studied, Berger et al. (1998) and Peek and Rosengren (1998) document loan availability changes after bank mergers. Other studies such as Jiménez et al. (2017) and Cortés et al. (2020) study the impact of regulatory changes (e.g. the implementation of new macroprudential rules or the onset of stress testing) on bank lending. Many papers, such as Ashcraft (2005), Alfaro et al. (2021), Beck

[2](Bordeu et al., 2025) found a relatively high R^2, but this may reflect the fact that the median loan amount in their data is less than $2,200. Loans of this size are typically for factoring (i.e., self-liquidating loans collateralized by the accounts receivable of the firms' customers) and have very short maturities. As a result, there may not be much firm-level variation in their interest rates.

et al. (2021), and Chen et al. (2024), to name just a few, examine the effects of bank failures on lending. Finally, there is a vast body of literature, well surveyed by Sufi and Taylor (2022), that analyzes the effects of banking crises. In all papers of the sort that find some bank customers totally lose access to credit or see its price rise, our search costs could help translate these changes into estimates of how much funding might be made up elsewhere.

Other papers show how declines in bank credit affect real variables, such as bank customers' employment, export growth, and investment (Chodorow-Reich, 2014; Amiti and Weinstein, 2011, 2018; Alfaro et al., 2021). Our search cost estimates help generalize these findings by giving a sense of how shocks of this type could affect a wider range of businesses. The presence of search costs can also provide microfoundations for the premise behind many papers that examine the misallocation of capital (such as Hsieh and Klenow (2009) and Gopinath et al. (2017)). The existence of search costs also plays a central role in the literature on relationship lending in banking. Studies of relationship lending show that repeated interactions between borrowers and lenders are associated with lower loan spreads, particularly for more opaque firms (Bharath et al., 2011), and that these ties can be more valuable in periods of stress, e.g., Banerjee et al. (2021). These findings are consistent with the view that screening and monitoring costs decline as lenders accumulate private information about borrowers.[3]

The theoretical foundation for our analysis is drawn from classic models of asymmetric information and financial intermediation. Pyle and Leland (1977) and Diamond (1984) emphasize that lenders invest in screening and monitoring to mitigate information problems, creating informational advantages that are not easily transferable across lenders. As a result, incumbent lenders are better positioned than potential entrants to extend credit, thereby raising the cost of changing banks. Our findings support these predictions.

3 Data and Stylized Facts

In this section, we describe the data and present stylized facts about loan spreads. Our analysis uses a matched U.S. bank-firm data set collected by the Federal Reserve for Dodd-Frank Act-mandated stress testing (FR Y-14), which includes all the standard variables typically considered important

[3]Though, as Rajan (1992) explains, the accumulation of private information can also create lock-in that is costly for borrowers.

for interest rate determination. While the United States does not maintain a credit registry, these data are as close as possible to a registry for the set of large banks subject to stress testing because they are required to report detailed information on all loans with committed amounts greater than $1 million. One benefit of these data, relative to the extensive literature on syndicated lending, is that they include information on smaller businesses, which are more likely to be bank-dependent.

3.1 Construction and Characteristics of the Sample

The Y-14 data include information on 42 bank holding companies, which represent approximately 70% of all commercial bank lending in the United States (Chodorow-Reich et al. (2022)). We focus on the H.1 schedule, which includes a comprehensive dataset of all loans to commercial and industrial (C&I) borrowers, and limit our analysis to 36 banks with large numbers of C&I loans. We also limit our sample to loans to for-profit companies and to loans for corporate purposes or working capital, to minimize unobservable borrower heterogeneity. This supervisory data is collected quarterly for US bank holding companies, intermediate holding companies of foreign banks, and savings and loan holding companies with total assets of more than $100 billion, making it the most comprehensive source of data on US bank loans.[4] The Y-14 data include quarterly observations, beginning in the third quarter of 2011. Because reporting was inconsistent in the first year of the sample, we limit the analysis to loans made between 2012:Q4 and 2024:Q4.

The Y-14 data contain information on each loan's interest rate, reference rate, maturity, borrower location, industry, and collateral. For comparability for floating rate loans with different reference rates, we compute loan spreads by taking the difference between the reported rate and the 3-month LIBOR or, subsequently, its successor, the Secured Overnight Funding Rate, SOFR. For fixed-rate loans, we subtract the matched-maturity U.S. Treasury yield (or the one-month LIBOR or SOFR rate at origination for fixed-rate loans with maturities of less than six months).[5] Because there may be data entry errors, we drop reported interest rates exceeding 30 percent and observations with

[4]Most studies on U.S. bank lending rely on data from syndicated loans, using either the LPC DealScan data or the confidential Shared National Credit data, which covers only loans held by multiple banks, and mostly loans greater than $100 million. Some studies have also used data from the Community Reinvestment Act (CRA), which captures loans under $1 million for banks with more than $1 billion in assets, but do not include interest rate data. See https://www.federalreserve.gov/apps/reportingforms/Report/Index/FR_Y-14Q for detailed instructions on the Y-14 data collection, H.1 schedule. The FR Y-14Q data collection has been used in several papers. For example, Chodorow-Reich et al. (2022) compare the Y-14 data with Compustat, DealScan, and SNC, highlighting the differences between the large borrowers in these databases and the smaller borrowers in this sample. They also note that loans collected in the Y-14 account for more than 90% of total C&I and real estate-backed lending by stress-tested banks.

[5]Our results are not qualitatively affected by changing reference rates.

interest rates equal to zero. We track borrowers in the sample through their taxpayer identification numbers. In addition to the committed and outstanding loan amounts, we calculate the utilization rate, defined as the amount drawn (utilized) divided by the committed amount over the life of the loan in the dataset.

An advantage of this comprehensive supervisory data is that it includes a full set of fields designed to allow supervisors to estimate potential loan losses in a stressed economic environment. Banks directly report measures of loan risk, including an internal loan risk rating, which captures the bank's estimate of the loan's default risk. We use a confidential Federal Reserve concordance table that maps banks' internal loan risk ratings into a common rating scale that mirrors the one used for syndicated loans, where loans are rated from AAA to D, with AAA being the least risky. The AAA and AA categories account for less than one percent of all loans. We drop these observations from our sample because there are too few to implement our estimation procedure. We also dropped approximately 3 percent of loans with ratings below B for similar reasons, as well as because high-risk loans may involve unobserved heterogeneity that is difficult to quantify.

In addition to the common loan risk rating scale, a subsample of loans includes additional data because they are evaluated using a bank- and loan-specific "advanced approach" to rate the loan. The advanced approach is available to banks whose risk models have been approved by supervisors. After approval, the bank's own model is used to determine the required capital charge for a loan. For these loans, our data includes the bank's estimates of the probability of default, loan exposure at default, and loss given default, each expressed as a percentage of the committed loan amount. The expected loss ratio is the product of these numbers and, when multiplied by the committed amount, determines the expected dollar loss on a given loan. This second measure of loan risk allows us to observe a continuous measure of banks' evaluations of borrower risk, explicitly capturing many elements of loan risk that would otherwise be hard to observe.

We observe other loan characteristics that are likely to matter for interest rates, such as the loan type and purpose. The Y-14 reporting schedule lists 19 loan types and 30 loan purposes. We restrict our sample to corporate-purpose and working-capital loans to obtain a homogeneous set of loans. For example, this means we exclude loans used to fund acquisitions, which may have different risk characteristics depending on the acquisition's scale and scope. Similarly, loans

for project financing and real estate-related purposes are excluded because they may depend on collateral that is difficult to compare. We also drop loans made for other miscellaneous purposes.

Approximately a quarter of loan observations remain after we impose these restrictions on loan purpose. Within the corporate purpose and working capital loan purposes, we divide loan types into four categories, separating revolvers from term loans and loans for short-term or seasonal working capital from those for permanent working capital or general corporate purposes. Anticipating that we will do most of our analysis using bins of similar loans, we also drop loan types that are too rare to form a bin, i.e., when there are fewer than twenty originations of a particular loan type in our sample period or fewer than six banks originating that loan type. On average, there are approximately 3,750 originations per quarter, with a peak in 2012 and a minimum in Q3 2020.

Table 1 reveals several basic properties about the sample. Of the 160,627 loans in our sample, most are rated BBB or BB. The median loan spread monotonically increases with loan risk. Spreads fall with the amount committed and maturity, which likely reflects the fact that, all else equal, long-term borrowers and larger firms tend to have lower default risk. Revolvers, which account for approximately 72 percent of all loans and 82 percent of the value of all lending in our sample, have systematically higher spreads than term loans, which motivates our decision to divide these two types of loans into four separate categories when forming bins: revolvers for short-term capital, revolvers for general working capital, term loans for short-term capital, and term loans for general working capital. We also categorize loans according to whether they have fixed or variable interest rates.

Table 1: Loan Characteristics

	Percent of obs	Percent of committed	Mean spread	Median spread	Mean interest rate	Median interest rate	Mean utilization	SD spread	Mean Expected Loss
Rating:									
A	1.10	2.59	1.62	1.39	2.57	2.24	60.48	0.93	0.03
BBB	26.22	38.30	1.93	1.72	3.36	2.97	65.95	0.95	0.10
BB	56.65	47.23	2.57	2.43	4.05	3.66	69.36	1.07	0.28
B	16.04	11.89	3.15	3.00	4.66	4.25	69.09	1.26	0.95
Committed (millions):									
1-5	52.75	7.48	2.71	2.63	4.12	3.75	71.73	1.11	0.40
5-10	12.11	5.60	2.48	2.31	3.90	3.41	68.88	1.13	0.35
10-25	15.47	16.30	2.36	2.15	3.77	3.25	68.13	1.17	0.34
25-100	17.86	54.58	2.02	1.77	3.65	3.08	60.17	1.06	0.25
100-250	1.79	15.72	1.71	1.44	3.83	3.26	47.81	0.91	0.14
Maturity (years):									
0-1	32.29	11.10	2.77	2.75	4.18	3.75	61.67	1.07	0.40
1-4	32.31	35.80	2.44	2.31	3.92	3.44	67.63	1.13	0.35
4-5	28.25	50.35	2.26	2.00	3.75	3.38	71.45	1.18	0.30
5-8	5.77	2.58	2.32	2.12	3.78	3.60	87.87	1.17	0.34
8-15	1.36	0.16	2.11	2.00	4.28	4.20	94.81	1.02	0.48
>15	0.02	0.00	1.61	1.63	4.60	4.66	94.50	0.86	0.30
Loan type:									
Revolver, short-term capital	30.50	17.30	2.69	2.63	4.43	4.00	60.82	1.04	0.35
Revolver, general working capital	41.76	64.35	2.43	2.27	3.70	3.25	56.75	1.16	0.32
Term, short-term capital	6.38	2.20	2.23	1.95	3.34	3.18	96.75	1.19	0.36
Term, general working capital	21.36	16.14	2.37	2.17	3.95	3.64	93.18	1.20	0.39
Loan Rate type:									
Fixed	11.64	4.54	2.14	2.04	3.87	3.70	90.80	1.26	0.33
Floating	88.36	95.46	2.53	2.41	3.96	3.50	65.37	1.12	0.35

Note: Except for the last column, this table reports summary statistics for the sample of 160,627 loan originations. The units are in percentage points. We drop originations of revolver loans that were never drawn over the life of a loan, loans in bins with fewer than six unique banks, or bins with fewer than twenty loans. All values are calculated as of the origination date. Spread equals the difference between the interest rate of a newly originated loan and the reference rate. Mean utilization is the mean of the loan value in the category, defined as utilized divided by committed, calculated for each loan over the loan's lifespan in the dataset. The mean expected loss is calculated for the subset of loans that are rated under the advanced ratings approach.

Although 53 percent of loans are for amounts between $1 million and $5 million, small borrowers account for less than 8 percent of the total amount committed. Conversely, loans of $25 million or more account for about 20 percent of the observations but represent about 70 percent of the dollars lent. One way to say this succinctly is that the median loan is small, but the median dollar lent is part of a large loan. The large share of aggregate lending that arises from loans exceeding $5 in value highlights the importance of understanding this tranche of the loan market. Finally, just under 93 percent of all loans have a maturity of five years or less. So most of our loans are relatively short-term.

3.2 Documenting the Spread Dispersion Among Similar Loans

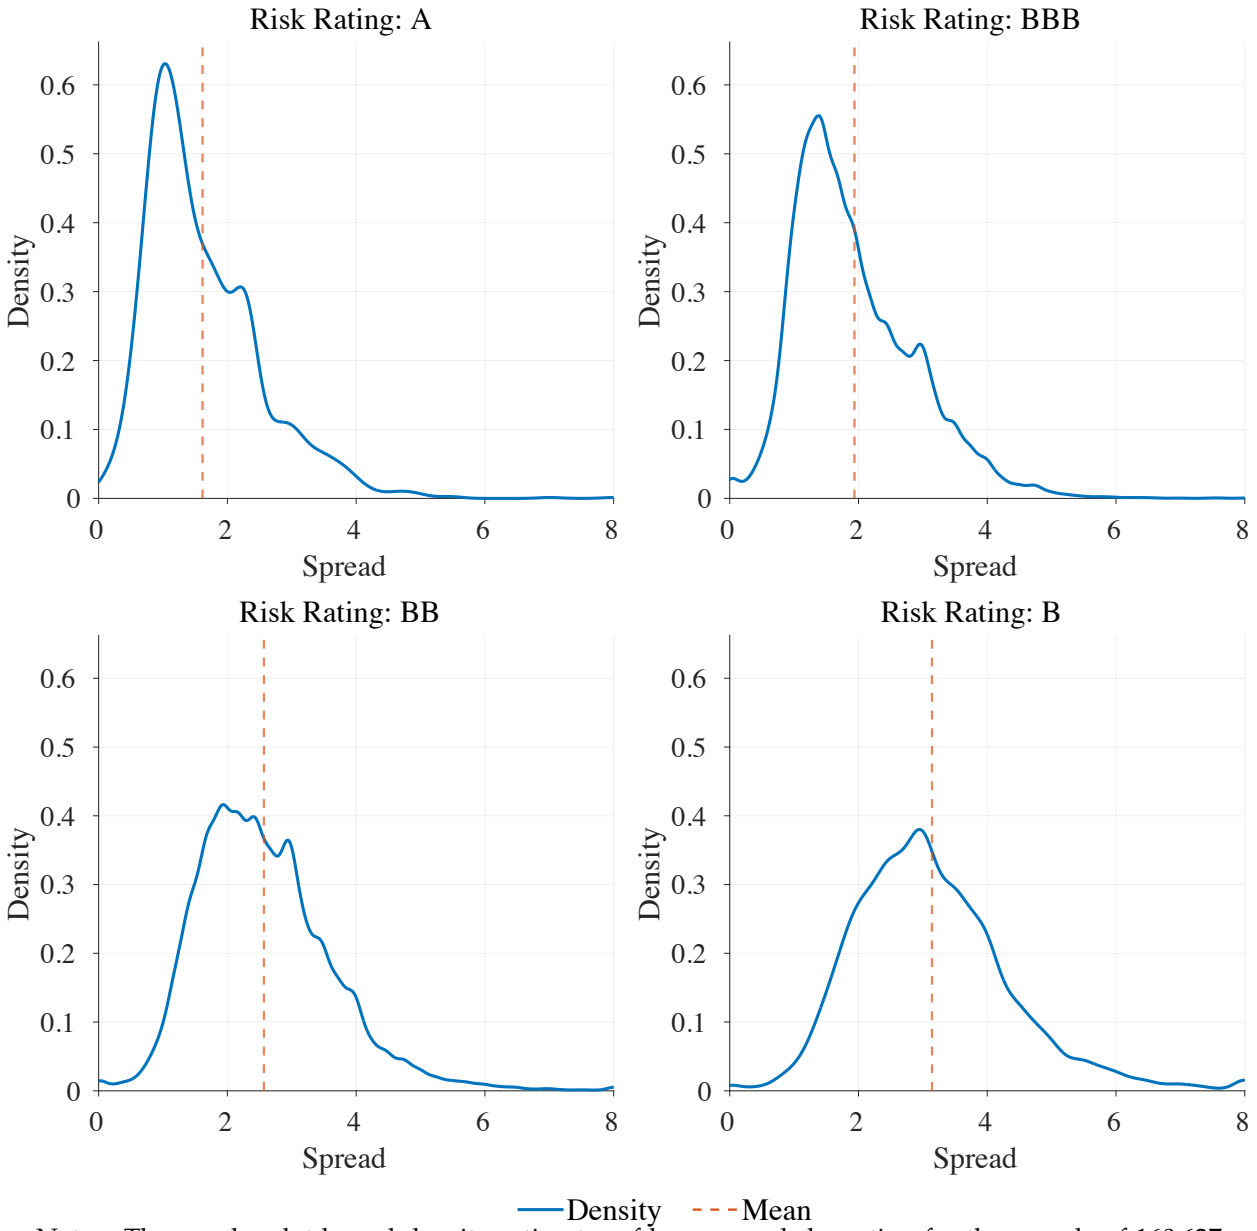
A salient pattern in the data is the substantial variation in loan spreads, even for loans with identical risk ratings. Although the median spread rises from 1.39 percentage points to 3.0 percentage points as we move from A-rated loans to B-rated ones, the standard deviation of spreads for loans of the same rating ranges from 0.93 to 1.26 percentage points. This means that, holding the risk rating constant, there is a high degree of variation in the spreads paid by borrowers with identical ratings and a substantial overlap in borrowing costs for firms with different risk ratings.

This overlap is more clearly seen in Figure 1, which presents kernel density plots of loan spreads at the time of origination, by rating. We observe considerable dispersion in these spreads, and the variance clearly rises with loan risk. The data make plain that although higher risk is associated with higher loan spreads, there is substantial variation in spreads that cannot be explained by a loan's risk rating.

To formalize the role of observable loan characteristics in explaining loan spread variation, Figure 2 plots the estimated coefficients and 95 percent confidence intervals from a multivariate regression of the loan spread on dummy variables for risk rating, committed amount, maturity, loan type, loan purpose, date, and industry fixed effects. There are five main takeaways from this regression. First, while there is some variation in spreads across time (panel a), loan type (panel f; term vs. revolver), and by NAICS industry code (panel g), there is no obvious pattern to the variation. Second, Panel b demonstrates that our earlier finding, that riskier loans have higher spreads, remains robust when controlling for other loan characteristics that might be correlated with risk rating. Otherwise identical A-rated loans have spreads that are about 150 basis points

lower than similar B-rated loans. Third, Panel c shows that maturity does not matter for spreads on loans with maturities under eight years, which account for 98.6 percent of all originations. Fourth, Panel d shows that the loan spreads fall with loan size.[6] Fifth, Panel h shows that floating-rate loans tend to have higher loan spreads.

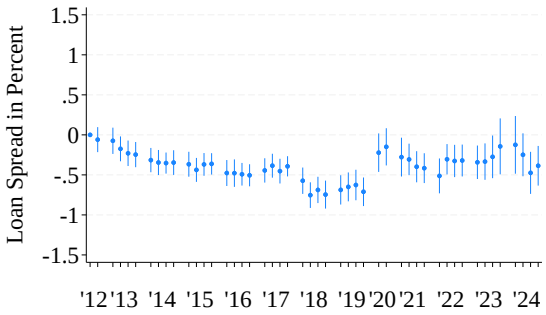
Figure 1: Kernel Density Estimates of Loan Spreads by Rating



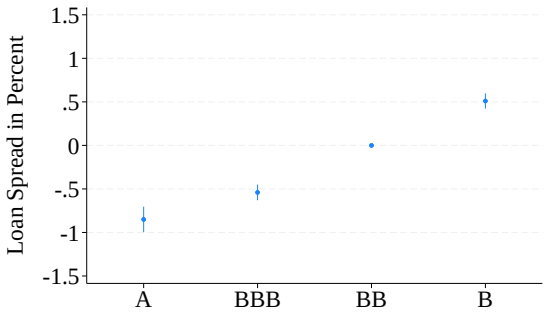
Notes: The graphs plot kernel density estimates of loan spreads by rating for the sample of 160,627 observations.

[6]There are very few loans of over $250 million, which accounts for the lack of precision about the magnitude of the corresponding fixed effect.

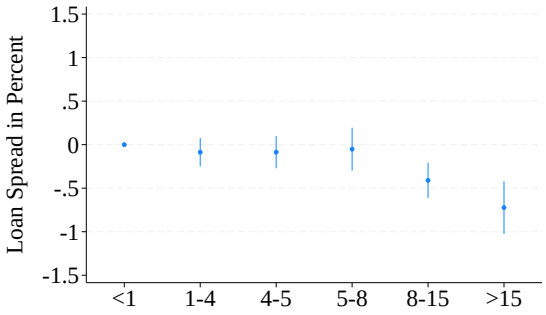
Figure 2: Impact of Loans Characteristics on Loan Spreads



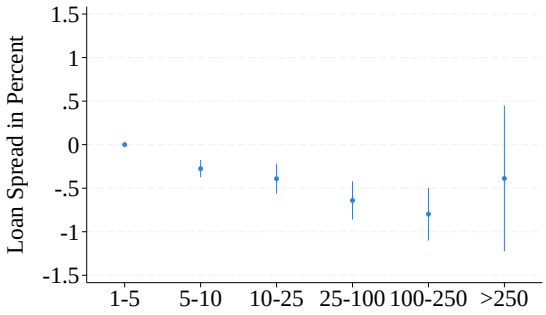
(A) Loan Origination Date Fixed Effect



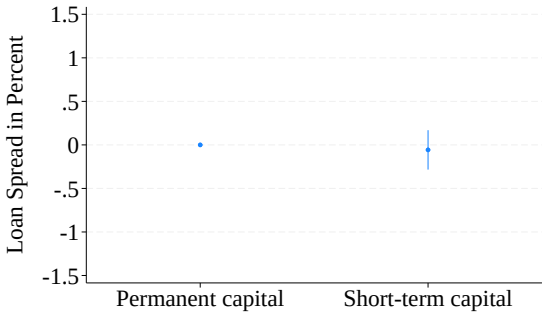
(B) Loan Rating Fixed Effects



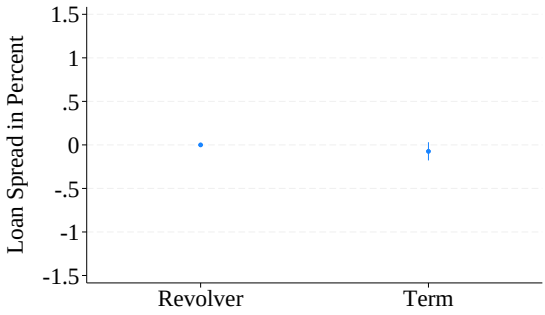
(C) Maturity Fixed Effect (years)



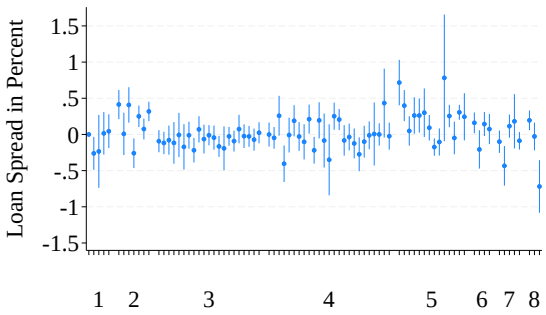
(D) Committed Loan Size Fixed Effect (millions)



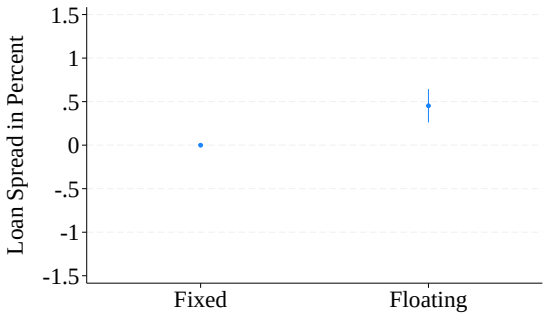
(E) Loan Purpose Fixed Effects



(F) Term Loan vs. Revolver Fixed Effects



(G) 3-Digit NAICS Fixed Effect



(H) Loan Rate Type Fixed Effect

Note: Notes: These graphs show the coefficients and 95 percent confidence intervals of a regression of the loan spread at the origination date of the loan on various fixed effects. The sample covers 160,027 originations. The horizontal axis in figure (g) displays 1-digit NAICS categories, but the fixed effects are estimated at the 3-digit level.

The most striking result from this regression is that the R^2 is only 0.25, indicating that these variables account for only a quarter of the variation in loan spreads. Put differently, three-quarters of the variance in spreads cannot be explained by our regression. This raises the question of whether there is substantial variation among loans with nearly identical characteristics or whether the linear model is mis-specified. For example, our explanatory variables may interact in complex ways, which could account for the observed variation in spreads.

We address these questions by examining variation in loan spreads across loans with essentially the same characteristics. To move forward, we first need to develop some notation. We denote the loan spread that firm j pays when originating loans in time t with a set of characteristics given by k from bank b by r_{jbkt}. In this setup, all loans denoted by k are in the same "bin" because they are identical in terms of their risk rating, interest rate regime, loan type, and loan purpose, and have nearly similar committed amounts and maturities. These variables are the primary ones banks collect to measure loan risk and were included in the Y-14 data to estimate loan performance under stress. When forming bins, we group time periods (quarters) with similar interest rate regimes to account for potential changes in loan pricing when monetary policy changes. By grouping loans by these dimensions, we reduce heterogeneity within each bin.

Table 2 presents bin-level statistics grouping the bins on various dimensions—note all bins have more than five banks originating loans and at least twenty loan originations. We have 722 risk bins, and the summary stats are shown in the first column of the table. The next four columns separate the bins by credit rating. We have 33 A-rated loan bins, 236 BBB-rated loan bins, 286 BB-rated loan bins, and 167 B-rated loan bins. We also have a single bin with originations that exceed $250 million, which is not tabulated due to limitations on data disclosure. There are typically many loan originations in each bin. The number of banks per bin ranges from around 9 for A-rated originations to around 14 for BB-rated originations.

Table 2: Average Loan Characteristics Across Bins

	All Bins	A	BBB	BB	B	1-5	5-10	10-25	25-100	100-250	Public	Private
Number of Banks	13.00	8.76	12.46	14.04	12.81	13.15	12.95	13.44	12.92	10.64	11.53	15.66
	(4.88)	(2.65)	(4.39)	(5.28)	(4.59)	(5.18)	(4.55)	(4.90)	(4.76)	(3.41)	(3.80)	(4.57)
Number of Loans	222.48	53.36	178.46	318.15	154.25	322.17	139.99	172.57	202.05	87.06	84.58	249.20
	(379.65)	(51.21)	(246.61)	(517.16)	(226.60)	(546.01)	(150.16)	(194.97)	(296.52)	(87.66)	(109.63)	(262.14)
Maturity (years)	3.72	3.34	3.69	3.91	3.50	4.02	3.42	3.57	3.52	3.99	4.08	4.03
	(2.33)	(1.69)	(2.20)	(2.58)	(2.14)	(3.03)	(1.98)	(1.83)	(1.70)	(1.08)	(1.37)	(1.44)
Loan Size (millions)	22.36	45.57	26.71	19.24	16.96	2.35	7.55	17.14	48.57	144.09	40.15	37.89
	(34.24)	(47.53)	(41.46)	(28.95)	(24.60)	(0.35)	(0.33)	(0.92)	(5.52)	(8.55)	(34.34)	(33.91)
Mean (spread)	2.33	1.52	1.76	2.42	3.15	2.58	2.40	2.30	1.96	1.76	1.88	2.17
	(0.71)	(0.51)	(0.44)	(0.43)	(0.55)	(0.62)	(0.64)	(0.74)	(0.72)	(0.53)	(0.57)	(0.61)
Median (spread)	2.19	1.43	1.64	2.26	2.98	2.47	2.24	2.15	1.80	1.59	1.73	2.00
	(0.70)	(0.56)	(0.46)	(0.43)	(0.51)	(0.57)	(0.63)	(0.74)	(0.70)	(0.55)	(0.56)	(0.61)
p90–p10 Range (spread)	2.41	1.52	1.88	2.52	3.14	2.56	2.47	2.44	2.16	1.83	1.88	2.28
	(0.93)	(0.69)	(0.59)	(0.71)	(1.10)	(1.06)	(0.79)	(0.80)	(0.88)	(0.68)	(0.74)	(0.65)
Interquartile Range (spread)	1.21	0.75	0.93	1.27	1.57	1.26	1.25	1.24	1.10	0.91	0.93	1.16
	(0.49)	(0.39)	(0.33)	(0.39)	(0.54)	(0.50)	(0.42)	(0.45)	(0.54)	(0.45)	(0.46)	(0.41)
Number of Bins	722	33	236	286	167	263	139	144	142	33	158	158

Notes: This table presents summary statistics for the sample of 160,627 loan originations divided into 722 bins. Each number is calculated by first computing the average of a variable within a bin, and then averaging the means across bins within a characteristic (e.g., rating or maturity). Columns 2-5 present averages for different ratings; columns 6-10 present averages for different committed amounts (in millions); columns 11 and 12 present averages for the 158 bins that have enough loans so that the bins they can sub-divided into bins with at least 6 banks of 20 publicly listed and at least 6 banks with 20 private borrowers. Standard deviations of the averages within bins are reported in parentheses.

One of the striking features of the table is the variance in loan spreads within each bin. To describe this variation without overemphasizing outliers, we compute measures of difference in the interest-rate spread within a bin by subtracting the spread at the 10th percentile from that at the 90th percentile. On average across all the bins, the "p90-p10 range," or spread gap, is 2.41 percentage points, indicating that loans with identical risk ratings and similar or identical loan characteristics have very different spreads. To get a sense of the magnitude of the within-bin heterogeneity, we compare it to the difference in mean spreads between A and B-rated bins reported in Table 1: 1.53 percentage points. Thus, there is, on average, more dispersion in spreads within a bin containing loans with identical risk ratings than there is across the mean of bins with very different ratings. Moreover, we see that this within-bin dispersion is inversely related to loan risk. The p90-p10 range for A-rated bins is slightly less than half that of B-rated loans. Columns 6 to 10 group the binds by committed amount, and the spread gap for loans of over $100-$250 million is about 30 percent smaller than that for loans of $1-$5 million.

3.3 Do Other Observables Explain the Spread Dispersion Among Similar Loans?

The first step in our analysis is to regress loan spreads on fixed effects for each bin k. We denote the loan spread for firm j originating a loan from bank b in bin k at time t by r_{jbkt}. So our specification is

$$r_{jbkt} = \alpha_k + \epsilon_{jbkt}, \tag{1}$$

where α_k is a bin k fixed effect. The error term, ϵ_{jbkt}, captures, among other things, unobserved loan heterogeneity, idiosyncratic variation in the determination of interest rates (e.g., bargaining power), or the role played by search costs in preventing firms from finding the cheapest loan rate. This specification allows for an arbitrary relationship among the variables used to construct the bins.

Table 3 shows the results of our estimation of equation 1. The first column contains the results obtained by only including bin fixed effects. The R^2 for this specification is 0.26, only 1 percentage point higher than when we included the loan characteristics as linear explanatory variables. Clearly, allowing for more complex interactions among loan characteristics provides little additional explanatory power. Another way of understanding the implication of this is by

Table 3: Loan Spread Regressions

	Loan Spread					
	(1)	(2)	(3)	(4)	(5)	(6)
Relationship Length					-0.018***	-0.007**
					(0.002)	(0.002)
Log Sales					-0.086***	-0.060***
					(0.002)	(0.002)
Observations	160,627	160,627	160,627	158,496	129,106	127,573
Number of Bins	722	722	722	722	722	722
R^2	0.26	0.27	0.31	0.42	0.44	0.53
RMSE	0.99	0.98	0.95	0.89	0.87	0.80
Mean of the p90 – p10 Range	2.40	2.38	2.31	2.01	1.95	1.83
p10 of the p90 – p10 Range	1.45	1.45	1.37	1.15	1.10	1.08
p25 of the p90 – p10 Range	1.87	1.86	1.78	1.56	1.43	1.36
p50 of the p90 – p10 Range	2.28	2.26	2.18	1.95	1.87	1.70
p75 of the p90 – p10 Range	2.75	2.72	2.65	2.35	2.26	2.14
p90 of the p90 – p10 Range	3.49	3.39	3.37	2.85	2.83	2.62
Number of Variables	722	768	803	7,300	6,436	6,903
Bin FE	Y	Y	Y			
Quarter FE		Y	Y	Y	Y	Y
Bank FE			Y			
Risk Proxies FEs						Y
Interest Rate Index FE						Y
Industry and Region FE						Y
Bank-Bin FE				Y	Y	Y

Notes: The dependent variable is the loan spread. RMSE stands for "Root Mean Squared Error". We construct the p90-p10 range by sorting the residuals within each bin and subtracting the 10th-percentile value from the 90th-percentile value. We then report the 10th, 25th, 50th, 75th, and 90th percentile of these p90-p10 ranges. Number of Variables denotes the number of regressors in each regression (including FEs). The risk proxy fixed effects are dummy variables equal to 1 if a stock ticker or a CUSIP number is present, a fixed effect for lien position, and a fixed effect for collateral type. The industry fixed effects are based on North American Industry Classification System (NAICS) 3-digit codes. The regional fixed effects denote MSAs. Robust standard errors are in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

comparing the standard deviation of the mean spread across bins given in column 1 of Table 2 with the root mean squared error from this regression, given in column 1 of Table 3. This comparison tells us that the standard deviation in spreads across bins given in Table 2 (0.710) is smaller than the within-bin dispersion in spreads in Table 3 (0.99). In other words, the primary source of dispersion in loan spreads for these loans is not differences across risk categories but differences within them. As a further diagnostic, we compute measures of dispersion within bins by defining the spread gap as the difference between the ninetieth and tenth percentiles of the residuals in the same bin. The lower portion of the table shows the distribution of spread gaps. Notice that the median adjusted gap is 2.28, so that within the typical bin, there is a 228 basis point difference in the 90th and 10th percentiles of the residuals after controlling for bin fixed-effects. This finding reinforces the conclusion that controlling for possible interactions among the characteristics does not help explain much of the dispersion in loan spreads.

We can obtain a sense for what is driving this dispersion by plotting the data. Figure 3 presents histograms of the loan spreads for the bins with the most loan originations by rating. These figures are strikingly similar to those in Figure 1, which portrayed the unconditional distributions of loan spreads for all loans with a given risk rating. For example, both the overall loan spread distribution of BBB loans and the spread distribution in the largest BBB bin range from approximately 1 to 5 percentage points. The fact that the within-bin variation is comparable to the overall level of variation, combined with the results from Table 3, suggests that observable characteristics cannot easily explain the loan spread dispersion in the largest bins.

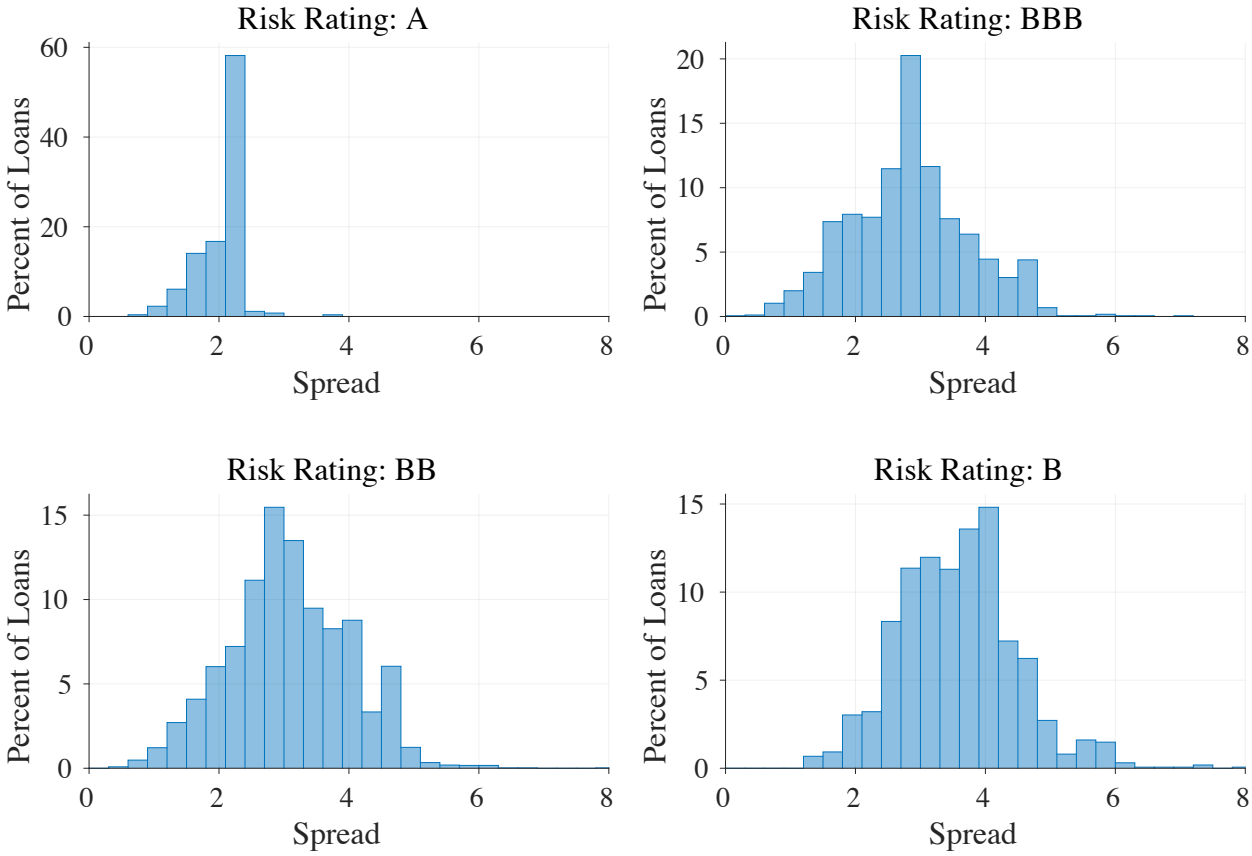
Since most of the loan spread dispersion cannot be explained by loan ratings and other standard risk variables, a natural question is whether a borrower-, bank-, or time-varying variable (X_{jbt}) can substantially improve the fit. We approach the problem by rewriting equation 1 as

$$r_{jbkt} = \alpha_k + \boldsymbol{X_{jbt}}\beta + \nu_{jbkt},, \tag{2}$$

where $\boldsymbol{X_{jbt}}$ is a vector of explanatory variables that vary by borrower, bank, and time; β is the corresponding vector of coefficients; and ν_{jbkt} is the new error term.

The remaining columns in Table 3 present the results of estimating equation (2) to determine whether various ways of specifying $\boldsymbol{X_{jbt}}$ can explain a significant portion of the residual variance

Figure 3: Loan Spreads for Largest Bins by Rating



Note: The figure plots histograms of loan spreads for the bins with the most loan originations by rating in the sample. Risk rating A has the following characteristics: Term, Short-term capital, 5-8 year maturity, 1-5 million committed. BBB's characteristics are: revolver, general working capital, 0-1 year maturity, 1-5 million committed bin. BB's characteristics are: revolver, general working capital, 0-1 year maturity, 1-5 million committed bin. B's characteristics are: revolver, short-term capital, 0-1 year maturity, 1-5 million committed bin. All four bins are from Regime 1.

in loan spreads. We look at two metrics: the R^2 and the spread gap between the ninetieth and tenth percentiles within bins. If additional explanatory variables explain the extent of the spread variation within a bin, we would expect the spread gap to close meaningfully. We begin by examining the explanatory power of finer time periods. Adding origination-quarter fixed effects to the first regression has almost no effect—the R^2 rises to 0.27. This result suggests that very little of the variation in loan spreads is due to temporal differences.

Martín-Oliver et al. (2008) argue that the variation in loan spreads within bins could arise if we assume banks differ in their marginal costs of providing loans within bins and search costs prevent borrowers from finding the cheapest provider. This explanation, which builds on the Carlson and McAfee (1983) search model, differs from the Burdett and Judd (1983) setup by assuming that some banks should offer loans at *consistently* lower rates than others. We can test this assumption directly by including bank fixed effects in our loan spread regressions to see if they explain a significant amount of the variance. The results from including bank fixed effects in column (3) of Table 3 increase the R^2 by only 0.04 and hardly affect the p90-p10 spread range. This implies that the spread dispersion is not present because some lenders offer loans that are consistently cheaper or more expensive.

One reason bank fixed effects might not reduce the variance in spreads much is bank specialization. A given bank may consistently offer high rates for certain types of loans and low rates for others. Alternatively, within a particular bin, different banks could have different amounts of bargaining power. For example, even for loans with a certain type of purpose with a floating rate, some banks could specialize in making short-term loans and avoid making longer-term loans, and *vice versa*. We address these concerns by replacing the bank fixed effects with bank-bin fixed effects. This substantially increases the number of regressors from 768 to 7,300 and raises the R^2 to 0.42, but does not overturn our finding that most spread dispersion occurs within, rather than across, bins.[7] The loan spread gap is still 195 basis points. Our finding that very little of the variation within individual bins can be explained by bank fixed effects is inconsistent with the Carlson and McAfee (1983) model, which predicts that each bank should charge the same rate to all customers.

The next column in Table 3 adds controls for factors that the banking literature finds are asso-

[7] The number of loans drops because we lose cases where a bank made one loan in a bin. In the subsequent columns in the Table, we have different numbers of loans in the bins because the large number of controls forces us to drop some observations. None of the results we emphasize is an artifact of the varying number of loans.

ciated with borrowing costs: relationship length and cross-selling.[8] If banks have soft information about borrowers that comes from repeated interactions, they may drop high-risk borrowers and retain low-risk ones, thereby offering lower interest rates on average to borrowers with whom they have a longer lending history. The fact that our bins already vary by loan size means we likely capture many cross-selling opportunities from borrowers, as banks may offer large borrowers low-interest loans to establish business ties and generate profits by selling other products. However, there could be important variation within bins that we are missing. We therefore also add the logarithm of borrower net sales as another control. This captures a broader range of cross-selling and other services, as firms with higher sales are likely to have more deposits and transactions. We find that banks with longer relationships tend to pay lower average rates, but the coefficient is small. The coefficient on log sales is negative, indicating that firms with higher sales have lower spreads, but it is small. More importantly, the R^2 rises only to 44 percent, and the mean spread gap remains at 187 basis points.

In the last column in Table 3, we also control for other potential risk proxies, such as whether the borrower is public (by adding a dummy variable that is one if the firm has a stock ticker), fixed effects for the presence of different types of loan collateral, and the base rate of the loan (typically 1-month or 3-month LIBOR, but occasionally the prime rate). We also control for the extent of bank competition (or other place-specific factors) and industry-related cross-selling or risk characteristics by including industry and MSA fixed effects. Even in this "kitchen-sink" specification with these additional 457 controls, the R^2 is only 53 percent, and the median loan spread gap is 170 basis points. In other words, adding more than 6,000 additional explanatory variables to our specification does not alter the finding that there is about as much variance that we cannot explain with observables as what we can.

Thus far, we have been operating under the assumption that the risk ratings submitted by banks capture risk that is priced in bank loans, but if this assumption is incorrect, our findings could be interpreted as an indictment of broader categories of bank risk captured in the Y-14 data itself, rather than telling us about the lack of a strong relationship between spreads and

[8]We set the length of a bank-firm relationship at the time of a loan's origination equal to the number of quarters between the time we first observe a bank lending to a firm and the loan origination date. We divide this measure by the average relationship length of all bank-firm pairs in that quarter so that the first years of the sample do not mechanically have shorter relationships.

measured risk. For example, the risk ratings may be too broad, and banks may perceive loans within a bin as having heterogeneous risk characteristics. In this case, banks' own risk models may yield pricing that is substantially different from estimates of loan risk within the same bins.[9] To assess the validity of our assumption about the underlying quality of the banks' risk rating data submitted to supervisors, we focus on the subset of firms that use advanced approaches and calculate continuous measures of loan risk. For these banks, our data includes the loan's expected loss ratio, defined as the expected amount of loan losses divided by the amount lent. In addition to adding a continuous measure of loan risk, this measure allows us to examine the impact of other loan terms on risk. For instance, if loan covenants influence repayment prospects, loss given default, or exposure at default, this will be captured by the expected loan loss estimate.[10] Likewise, if a bank has information indicating the relationship length or other information they can model based on historical payment patterns that matter for repayment, this should be taken into account as well.

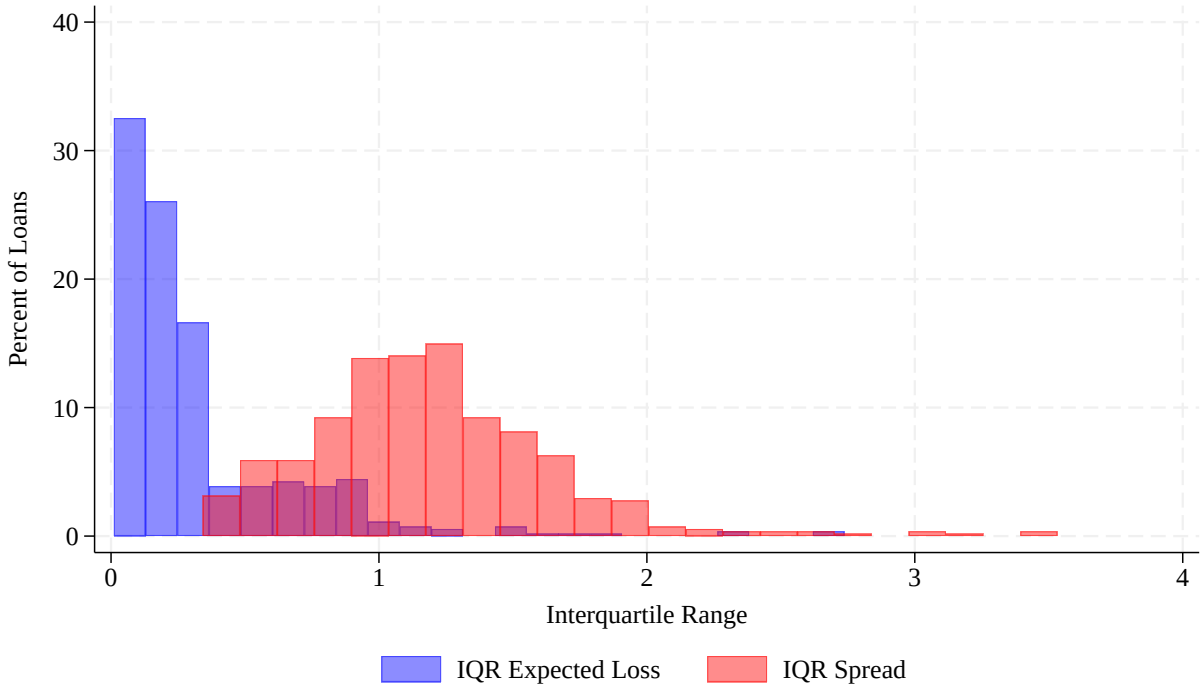
Figure 4 presents two histograms that cast doubt on the hypothesis that differences between banks' own risk assessments and those of the Y-14 data can explain our results. For each bin in our sample in which banks provide a loan-loss estimate, we compute the interquartile range for the expected loss ratio and the spread. The median interquartile range for the expected loss ratio is 21 basis points, while that for spreads within a bin is 115 basis points. As shown in the figure, there are very few bins in which the interquartile range of expected losses is as large as that of loan spreads. This plot makes a *prima facie* case that banks believe there is little dispersion of expected loss ratios within bins, and whatever dispersion does exist is too small to explain the substantial spreads we observe.

Table 4 investigates this hypothesis formally. The first column replicates our bin fixed effects regression presented in column 1 of Table 3). Column 2 restricts the sample to the 90,214 observations for which the expected loss ratio data is available to demonstrate that our choice of sample does not affect our ability to predict spreads. The third column adds the bank's expected loss

[9] For a bank to be able to use its own model for determining capital requirements, supervisors closely scrutinize the bank's risk models to ensure that they are reasonable. Banks' models incorporate a variety of factors, including historical defaults, industry and economic factors, quantitative data such as financial ratios, and qualitative information such as market position and industry conditions.

[10] Y-14 data do not contain information on loan covenants. Generally, these covenants are associated with bin characteristics, particularly rating and commitment size.

Figure 4: Distribution of Interquartile Ranges of Expected Loss and Loan Spreads



Notes: The histograms are constructed using 541 bins. To be included, a bin must contain at least six banks that report expected loss data, and each bank must originate at least 20 loans. The IQR Spread is defined as the interquartile range of the loan spread across all loan originations within a given bin; the IQR Expected Loss is calculated similarly.

ratio as an explanatory variable to determine if it provides additional explanatory power to the regression. While the expected loss ratio is a statistically significant predictor of the spread, it adds little to R^2. The small effect stems from the magnitude of the coefficient on the expected loss ratio and the fact that the expected loss ratio varies little within a bin. A move from the bottom to the top quartile of the range of the expected loss ratio only explains a two-basis-point movement in spreads: far less than the 115-basis-point interquartile range for spreads we observe in the typical bin.

Table 4: Loan Spreads and Expected Loan Losses

	Loan Spread		
	(1)	(2)	(3)
Expected Loss Ratio			0.115***
			(0.012)
Observations	160,627	90,214	90,214
Bin FE	Y	Y	Y
R^2	0.26	0.25	0.26
RMSE	0.99	0.97	0.96

Notes: The dependent variable is the loan spread. Column 1 regresses the loan spread over a set of bin fixed effects, using the full sample of 160,627 observations. Column 2 restricts the sample to observations for which expected loss data are available. Column 3 introduces the Expected Loss Ratio as a control variable; this ratio is calculated as the product of Probability of Default, Loss Given Default, and Exposure at Default, divided by the committed loan amount. RMSE stands for "Root Mean Squared Error". Robust standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

These results also help alleviate other potential concerns, as the expected loss ratio captures loan characteristics that banks use to predict loan performance but that may be hard to measure or not collected by supervisors. The limited impact of expected loan losses on loan spreads after controlling for risk bins suggests that bank-specific differences in risk models are not responsible for the dispersion.

3.4 Evidence from Bond Spreads

One potential concern with the results thus far is that our bins are mis-specified in some way that mechanically creates loan spread dispersion. We can assess this by estimating equation (1) with bond spreads (i.e., the bond yield minus the reference rate) as the dependent variable, rather than loan spreads. If the problem is predicting spreads using the characteristics we have identified,

then we should expect our specification to also fail to explain much of the dispersion in bond interest rate spreads. We do this by creating bond bins that group bonds by size, rating, and maturity. We report the result of estimating equation 1 in the first column of Table 5. The R^2 when estimating equation 1 using bond spreads is 0.77, triple what we found using loan spreads. Thus, the unexplained dispersion in loan spreads cannot be attributed to the difficulty of explaining spreads *per se*; it arises from a factor that is unique to loan originations.

Table 5: Bond Spread Regression

	Bond Spread				
	(1)	(2)	(3)	(4)	(5)
Log Offering				0.013	0.012
				(0.043)	(0.043)
Maturity (Years)					0.007**
					(0.002)
Observations	7,231	7,231	6,797	6,797	6,797
Number of Bins	272	272	264	264	264
R^2	0.77	0.80	0.81	0.81	0.81
RMSE	0.74	0.70	0.67	0.67	0.67
Mean(p90 - p10 Range)	1.40	1.33	1.28	1.28	1.27
p10(p90 - p10 Range)	0.30	0.43	0.40	0.40	0.39
p25(p90 - p10 Range)	0.59	0.60	0.62	0.62	0.61
p50(p90 - p10 Range)	0.93	0.88	0.86	0.86	0.86
p75(p90 - p10 Range)	1.65	1.47	1.40	1.40	1.40
p90(p90 - p10 Range)	3.23	2.98	2.66	2.66	2.66
Bin FE	Y	Y	Y	Y	Y
Quarter-Year FE	N	Y	Y	Y	Y
Industry FE	N	N	Y	Y	Y

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The dependent variable, the "Bond Spread," is the difference between the offering yield of a corporate bond and a comparable U.S. Treasury security. We construct the p90-p10 range by sorting the loan spread residuals within each bin and subtracting the 10th percentile from the 90th percentile. We then report the mean, 10th, 25th, 50th, 75th, and 90th percentiles of these p90-p10 ranges. RMSE stands for "root mean squared error." Robust standard errors in parentheses. * p < 0.05; ** p < 0.01 *** p < 0.001

The contrast between bond and loan spreads leads us to ask what distinguishes these two types of debt. Bonds are marketed to arm's-length investors and come with considerable information

about the borrower, including audited, publicly available financial statements. Rating agencies have reputational incentives to assess borrower risk accurately, making it easy to judge a bond's risk based on its rating. Thus, once a bond is rated, a bond buyer need not conduct further screening, and monitoring is easier. This makes bond markets relatively liquid, in that if one bondholder decides not to hold a firm's debt, it can easily sell its bonds.

By contrast, there is no standardized public mechanism for assessing the risk of small bank loans, so a borrower who has obtained a loan from one bank cannot switch lenders seamlessly. Instead, the borrower will need to search for a new bank to replace the lost credit, which incurs various costs.

This intuition leads us to investigate whether the importance of search costs for bank loans, which are largely absent for bonds, might be related to the unexplained high dispersion in loan spreads that we observe. Investigating this possibility requires us not only to model and estimate search costs, but also to examine the external validity of our estimates.

4 The Bank Search Model

Having failed to identify observable variables that explain loan spreads, we proceed by estimating a latent variable—search costs—to rationalize the data. We explain the methodology in this section. We discuss our approach to validating our estimates and assessing their economic significance in Section 5. Our approach to estimating borrower search costs closely follows the approach of Hong and Shum (2006) and Moraga-González and Wildenbeest (2008). While these papers consider the problem of estimating consumer search costs for low-cost sellers when the econometrician can only observe prices, we adapt this approach to a case in which we want to estimate borrower search costs when we can observe loan spreads. We assume it is possible to create loan bins in which loans offered by different banks are treated as equivalent.

We start by establishing some notation. We denote bins by k, but we suppress the k subscript on all variables in this section because the procedure we use is identical for all bins. We will bring this subscript back when we compare results across bins. Let ρ be the common per-dollar cost for a bank to offer a loan of amount L. Each borrower has a reservation loan spread of $\bar{r}$ above which it will not borrow, and we assume that every borrower costlessly observes the loan spread quote offered by one random bank, and those that search two or more banks choose loans from the bank

offering the lowest spread. Borrowers can also obtain $i-1 \geq 0$ additional loan spread quotes from other banks *before* choosing a lender at a cost of $(i-1)\,cL$, where c is the search cost per dollar borrowed. We assume that c differs across borrowers and is drawn from the distribution of search costs F_c.[11]

Following Burdett and Judd (1983), the loan spreads offered by the N banks arise from a mixed strategy equilibrium in which the equilibrium distribution of loan spreads that banks choose, F_r, makes each bank indifferent between charging any loan spread between the lower bound of the support for F_r ($\underline{r} > \rho$) and the upper bound ($\bar{r}$). This distribution of loan spreads ensures that the per-loan gains (losses) associated with offering a higher (lower) loan spread are exactly offset by the loss (gain) of customers associated with the higher (lower) loan spread. The optimal number of loan spread quotes a borrower will obtain is given by

$$i(c) = \arg\min_{i \geq 1} \left\{ (i-1)\,cL + \int_{\underline{r}}^{\bar{r}} irL\,(1 - F_r(r))^{i-1} f_r(r)\,dr \right\}, \tag{3}$$

where the second term is the expected annual interest payment conditional on doing i searches. As Honka et al. (2019) point out, the intuition for this expression can be obtained by noting that the probability that all loan spread quotes are greater than r after obtaining i loan spread quotes is $(1 - F_r(r))^i$, so the cumulative distribution function (CDF) for the minimum loan spread draw is $\left[1 - (1 - F_r(r))^i\right]$. The probability density function is just the derivative of the CDF $\left[i\,(1 - F_r(r))^{i-1} f_r(r)\right]$. Multiplying this by rL and integrating gives us the expected annual interest payment associated with obtaining i interest-rate quotes.

We next characterize the optimal number of interest-rate quotes a borrower will obtain, where we assume the maximum number of quotes they can obtain equals the number of banks lending in the bin, i.e., N. Since $i(c)$ is an integer, we divide borrowers into N subsets of size q_i, $i = 1, 2, ..., N$, with $\sum_{i=1}^{N} q_i = 1$, where q_i is the fraction of borrowers obtaining i interest-rate quotes. We denote the expected minimum loan spread in a sample of i loan spread quotes drawn from F_r by $E[r_{1,i}]$. We can write the search cost that would render a borrower indifferent (in expectation) between

[11] Morgan and Manning (1985) shows that simultaneous search may be optimal if agents need to make decisions rapidly and sequential search is time-consuming.

searching i and $i+1$ times as

$$\Delta_i \equiv E\left[r_{1,i}\right] - E\left[r_{1,i+1}\right], \quad i = 1, 2, \ldots, N-1 \tag{4}$$

Since Δ_i is a non-increasing function of i, the number of banks a borrower with a search cost of c obtains quotes from will be inversely related to its search cost. This property guarantees that we can write the fraction of borrowers searching i times is given by:

$$\begin{aligned} q_1 &= 1 - F_c\left(\Delta_1\right) \\ q_i &= F_c\left(\Delta_{i-1}\right) - F_c\left(\Delta_i\right), \quad i = 2, 3, \ldots, N-1 \\ q_N &= F_c\left(\Delta_{N-1}\right), \end{aligned} \tag{5}$$

where F_c is the CDF of the distribution of costs.

We can also use this structure to estimate the search-cost distribution, F_c, using only loan spread data. It is optimal for banks to set spreads using a mixed strategy for loan spreads in this setup. In equilibrium, one bank will charge the maximum loan spread, $\bar{r}$, which will equal the borrowers' reservation loan spread and gives us the upper value of the support for the equilibrium loan spread distribution, F_r. This bank will only lend to borrowers who obtain a single interest-rate quote (i.e., q_1), because borrowers who obtain two or more quotes could always do better. In equilibrium, other banks must be indifferent between charging the maximum loan spread and any other loan spread in the equilibrium support $\left[\underline{r}, \bar{r}\right]$, i.e.,

$$(\bar{r} - \rho) q_1 L = (r - \rho) L \left[\sum_{i=1}^{N} i q_i \left(1 - F_r(r)\right)^{i-1}\right], \tag{6}$$

where the left-hand side of this equation denotes the fraction of profits going to the bank charging the maximum loan spread, and the right-hand side denotes the fraction of profits going to a bank charging r. Hong and Shum (2006) solve the system by noting that the empirical counterpart of this equation is given by

$$(r_1 - \rho) q_1 L = (r_i - \rho) L \left[\sum_{\ell=1}^{N} \ell q_\ell \left(1 - \hat{F}_r(r_\ell)\right)^{\ell-1}\right] \quad i = 2, 3, \ldots, N-1, \tag{7}$$

where r_i is the observed loan spread charged by bank i, r_1 is the maximum observed loan spread, and $\hat{F}_r(r_i)$ is the fraction of banks offering loans with a spread below r_i. Since we also have $q_N = 1 - \sum_{i=1}^{N-1} q_i$, we have N equations and N unknowns (ρ and the q_i), so we can solve the system. Once we have estimates for the q_i, we can use equation 5, to solve for the CDF of the search-cost distribution $[F_c(\Delta_1), F_c(\Delta_2), ...F_c(\Delta_{N-1})]$.

Moraga-González and Wildenbeest (2008) rewrite the system of equations described above as a maximum-likelihood problem that is efficient and yields standard errors for the estimated parameters. They show that this method performs well in Monte Carlo trials, even when one does not observe prices offered by some suppliers, a common problem in estimating search costs.

5 Search Cost Estimates and Implications

In this section, we use our bank-cost search model to estimate borrowers' search costs. We then provide supporting evidence suggesting that our estimates do indeed reflect search costs and not just random noise. By leveraging predictions from the screening and monitoring hypothesis, we show that our search cost estimates are higher for loans that require additional information to assess risk and mitigate future monitoring costs. We also show that our search and search cost estimates closely match those from survey data. Finally, we estimate the potential impact of a bank failure in a counterfactual exercise.

5.1 Estimating Search Costs

We change our notation to label search cost estimates arising from different bins, k. Going forward, we add a k subscript to all variables so q_{ik} is the fraction of borrowers obtaining i interest rate quotes in bin k, and Δ_{ik} is the search cost that would render borrowers in bin k indifferent in expectation between obtaining i and $i+1$ interest rate quotes. It will be informative to discuss averages of q_{ik} and Δ_{ik} across bins with similar characteristics, e.g., all loans with a BBB rating. To do so, we define $q_{iK} \equiv \frac{1}{|K|} \sum_{k \in K} q_{ik}$ and $\Delta_{iK} \equiv \frac{1}{|K|} \sum_{k \in K} \Delta_{ik}$, so q_{iK} and Δ_{iK} are the bin averages of q_{ik} and Δ_{ik} in bins of characteristic K.

The first and sixth columns in Table 6 report the results from estimating the model using data for all of our bins, meaning that each element in the table represents the average value of the point estimate or standard error across all 722 bins. The standard deviations give us a sense of

the variation in estimates across these bins. We report our estimate for q_{iK}, the average share of borrowers searching i times in bins of type K, and the average lower bound of their search costs (Δ_{iK}) that would render them indifferent between searching i and $i+1$ times. The first row of the table indicates that we estimate 37 percent of borrowers accept the first quote they receive, implying a search cost per dollar borrowed of at least 75 basis points. Since the mean loan size across bins is $22.36 million, this estimate implies that, on average, 37 percent of borrowers in a bin incur search costs exceeding $167,700. We estimate that another 48 percent of borrowers obtain one additional quote. These borrowers have a search cost between 32 and 75 basis points per dollar borrowed, which implies that firms are behaving as if their search cost for a $22.36 million loan is at least $71,552. The last row indicates that, on average, 14 percent of borrowers have sufficiently low search costs to obtain loan quotes from every bank in the bin. For these borrowers, we estimate that their *maximum* search cost is three basis points or $6,700. This enormous heterogeneity in estimated search costs explains why some firms are bargain hunters and can obtain very low interest rates, while high-search-cost firms end up paying very high rates.

Table 6: Number of Searches and Average Search Cost Across All Bins and by Rating

Number of Searches	Share (q_{iK})					Search Cost (Δ_{iK})				
	All	A	BBB	BB	B	All	A	BBB	BB	B
1	0.37	0.36	0.38	0.36	0.36	0.75	0.44	0.59	0.82	0.91
	(0.16)	(0.17)	(0.17)	(0.16)	(0.13)	(0.29)	(0.23)	(0.24)	(0.25)	(0.27)
2	0.48	0.45	0.48	0.49	0.47	0.32	0.19	0.26	0.35	0.39
	(0.13)	(0.15)	(0.14)	(0.13)	(0.10)	(0.13)	(0.09)	(0.13)	(0.12)	(0.11)
3	0.00	0.02	0.01	0.00	0.00	0.18	0.11	0.15	0.20	0.22
	(0.04)	(0.07)	(0.05)	(0.02)	(0.04)	(0.08)	(0.06)	(0.09)	(0.08)	(0.07)
4	0.00	0.00	0.00	0.00	0.00	0.13	0.07	0.10	0.14	0.15
	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	(0.06)	(0.04)	(0.06)	(0.06)	(0.05)
N–1	0.00	0.00	0.00	0.00	0.00	0.03	0.03	0.03	0.03	0.04
	(0.01)	(0.00)	(0.00)	(0.02)	(0.00)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
N	0.14	0.17	0.14	0.14	0.14	0.00	0.00	0.00	0.00	0.00
	(0.07)	(0.10)	(0.07)	(0.06)	(0.08)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Note: This table presents the average estimates of q_{ik} and Δ_{ik} across bins containing loans of identical ratings for the sample of 160,627 loan originations grouped into 722 bins of loans of otherwise nearly identical characteristics. The units for Δ_{iK} are search cost per dollar borrowed, measured in percentage points, so numbers to the right of the decimal point are basis points. Standard deviations of the point estimates across bins within a category are in parentheses.

We observe an important pattern that provides insight into how the model rationalizes the data. We saw in Figure 1 that the spread distributions have a positive skew. Seen through the lens of the Burdett and Judd (1983) model, distributions of this form arise from distributions of search costs in which most borrowers search one, two, or N times. The intuition behind the model's tendency to yield this result is that the distribution's right-skewed tail is justified only if borrowers accept the first quote they receive (i.e., they face sufficiently high search costs to prevent them from obtaining more than one quote). Similarly, the left tail is justified by borrowers having very low search costs, so they comparison shop across all lenders, and the intermediate density by borrowers searching twice and thereby not usually receiving very high or very low quotes. In our case, we estimate that just over a third of borrowers accept the first quote they receive, just under a half obtain a quote from a second bank, and the remaining approximately fifteen percent of borrowers have sufficiently low search costs to comparison-shop across all banks lending in a given bin.

Recall that our bins pool loans within common monetary policy regimes. The estimates reported in Table 6 combine all four of these regimes. As a robustness check, we re-estimate the model separately for each of the regimes.[12] The relevant summary statistics for different regimes and the estimated search costs are shown in Tables A.4 and A.5 and are quite similar to the ones that pool over all four regimes. So in the validation exercises that follow, we pool all four regimes.

5.2 Validating our Estimates

One downside of the Moraga-González and Wildenbeest (2008) estimation procedure is that the cutoffs and the search cost estimates (i.e., the q_{ik} and Δ_{ik}) are latent variables that capture how the Burdett and Judd (1983) model rationalizes the distribution of loan spreads. Stated differently, the procedure generates estimates of the search cost cutoffs that are consistent with *any* distribution of observed loan spreads. One might be tempted to focus on the t-statistics of the point estimates, but these tell us nothing about whether our estimates are plausible because the null hypothesis that all search costs are zero is only consistent with a degenerate distribution of loan spreads—something that is clearly violated in our data.

[12] Recall that our first regime runs from 2012 Q4 to December 16, 2015, when rates were near the zero-interest-rate lower bound. The second interval spans December 17, 2015 to December 31, 2019, and includes a period of positive interest rates. We exclude the first two quarters of 2020 to avoid the initial period of the pandemic, when borrower risk may have been more difficult to assess. The third period begins on July 1, 2020, and continues until March 17, 2022, encompassing the second interval during which rates were again at the effective lower bound. The last period spans from March 18, 2022, through 2024 Q4, during which rates were non-negative.

However, there is a path forward. The screening and monitoring hypothesis provides clear predictions for how search costs should vary across bins.[13] Screening and monitoring costs are likely to appear as search costs for two reasons. First, a bank will need to screen (i.e., assess) a borrower's creditworthiness before issuing a loan quote. Second, future monitoring requires the bank to understand movements in financial and related information relative to a baseline, so it will also require the firm to provide similar information at origination to establish that baseline. Conditional on the bin, the provision of financial and other information constitutes a fixed search cost for a borrower seeking a loan quote. These costs will depend on the borrower's competence and record-keeping, which will vary *within* bins if firms differ in their ability to generate the required information. This is the variation we use obtain identification. Importantly, however, the costs will also vary *across* bins if the cost of providing the needed information depends on credit quality, loan size, etc. Our external validation exercise, therefore, consists of testing whether our estimates of search costs, obtained only using within-bin information, are consistent with the screening and monitoring theory's predictions of how they should vary across bins.

We also have a rigorous interpretation of the null hypothesis in these specifications. Much of the prior literature examining interest rate determination postulates (either explicitly or implicitly) that the error term in their version of equation (1) is independently and identically distributed (Petersen and Rajan, 1994; Berger and Udell, 1995; Degryse and Ongena, 2005; Mazet-Sonilhac, 2025). In this case, there should be no relationship between our search cost estimates and bin characteristics. Thus, it is meaningful to ask whether our search cost estimates vary according to the screening and monitoring hypothesis, or whether they are parameters that merely rationalize idiosyncratic variation in loan rates. If they are inconsistent with the former, we can formally reject the hypothesis that our search cost estimates capture only random variation in loan spreads.

[13]The screening and monitoring hypothesis is based on several assumptions. First, lenders do not expect to lose money by lending. They must recover the three costs of making a loan: their funding cost, the expected cost of borrower default, and the administrative costs of making the loan. These administrative costs include assessing credit risk ("screening") and ensuring that the borrower stays up to date with payments rather than absconding with the money ("monitoring"). Second, borrowers understand the lenders' problem and realize that bank administrative costs will be passed on to them through higher interest rates and greater information disclosure requirements (e.g., application procedures and ongoing information submission). Third, we assume that the costs of information disclosure vary across potential borrowers due to unobserved borrower characteristics such as management sophistication and capacity, the quality of the firm's financial records, prior releases of observable credit information, and time or monetary costs incurred in the loan application process. Finally, we assume that our estimates of search costs capture all the costs associated with providing the information that a bank needs to screen (and subsequently monitor) a borrower.

5.2.1 Search Costs and Credit Quality

Our first test is based on the screening and monitoring hypothesis's prediction that banks are likely to require more information disclosure and monitoring effort for a loan to a riskier borrower. Empirically, this amounts to comparing the estimated search costs for loans with different ratings. Table 6 also shows how the search cost estimates vary by rating. Average search costs clearly fall with credit quality. We estimate that, on average, A-rated borrowers incur search costs that are approximately half as large as those of B-rated borrowers. While we will conduct formal hypothesis testing in Section 5.3.1, this result suggests that our estimates vary according to the screening and monitoring hypothesis and are not being generated by random noise.

One way to measure the economic significance of the search cost estimates is to convert them into dollar amounts. A-rated firms that receive a second quote appear to act as if their search cost per dollar borrowed is between 19 and 44 basis points. Thus, the search cost for an A-rated firm seeking a second quote for an average-sized loan ($22.36 million) ranges from $42,000 to $98,000. By contrast, we estimate that B-rated borrowers of the same amount would incur search costs that are more than twice as high on average (between $87,000 and $203,000).

5.2.2 Search Costs and Loan Size

Our estimates are also consistent with the common-sense notion that there is a fixed cost associated with screening a firm, because the costs of many information-gathering activities (e.g., credit and background checks, document preparation) do not rise proportionally with loan size. If the cost per dollar lent falls with loan size, the screening and monitoring hypothesis implies that these cost savings should be passed on to borrowers in the form of lower paperwork costs. Since Δ_{ik} is our estimate of the cost of obtaining new loan quotes per dollar borrowed, we should expect it to fall with loan size. We see this pattern in Table 7, which shows that Δ_{iK} declines with loan size. For example, the search cost per dollar borrowed for borrowers who obtain a second quote ranges from 37 to 84 basis points for loans between $1 million and $5 million. However, search costs per dollar borrowed for a $100 million loan range are more than forty percent lower: between 21 and 50 basis points. Since there is nothing in our estimation procedure that should, *a priori*, produce this result, it provides additional suggestive evidence that our estimates are consistent with well-accepted theory and capture more than just noise.

Table 7: Number of Searches and Average Search Costs Across Bins by Committed

Number of Searches	Share (q_{iK})						Search Cost (Δ_{iK})					
	All	1-5	5-10	10-25	25-100	100-250	All	1-5	5-10	10-25	25-100	100-250
1	0.37	0.40	0.37	0.34	0.34	0.36	0.75	0.84	0.73	0.74	0.66	0.50
	(0.16)	(0.16)	(0.17)	(0.15)	(0.14)	(0.18)	(0.29)	(0.32)	(0.24)	(0.25)	(0.25)	(0.21)
2	0.48	0.45	0.48	0.51	0.50	0.47	0.32	0.37	0.31	0.30	0.27	0.21
	(0.13)	(0.13)	(0.13)	(0.12)	(0.12)	(0.17)	(0.13)	(0.15)	(0.11)	(0.10)	(0.12)	(0.09)
3	0.00	0.00	0.00	0.00	0.00	0.02	0.18	0.22	0.18	0.17	0.16	0.12
	(0.04)	(0.03)	(0.02)	(0.04)	(0.03)	(0.11)	(0.08)	(0.09)	(0.07)	(0.06)	(0.08)	(0.06)
4	0.00	0.00	0.00	0.00	0.00	0.00	0.13	0.15	0.12	0.11	0.11	0.08
	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.06)	(0.07)	(0.05)	(0.05)	(0.06)	(0.05)
N–1	0.00	0.00	0.00	0.00	0.00	0.00	0.03	0.04	0.03	0.03	0.03	0.03
	(0.01)	(0.00)	(0.00)	(0.00)	(0.03)	(0.00)	(0.03)	(0.04)	(0.02)	(0.02)	(0.02)	(0.02)
N	0.14	0.15	0.13	0.14	0.14	0.15	0.00	0.00	0.00	0.00	0.00	0.00
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.12)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Note: This table presents the average estimates of q_{ik} and Δ_{ik} across bins containing loans of various commitment sizes for the sample of 160,627 loan originations grouped into 722 bins of loans of otherwise nearly identical characteristics. The units for Δ_{iK} are search cost per dollar borrowed, measured in percentage points, so numbers to the right of the decimal point are basis points. Standard deviations of the point estimates across bins within a category are in parentheses.

5.2.3 External Validity of Our Search Cost Estimates

We can also assess whether our estimates align with those obtained from surveying firms on the frequency and costs of search. Wiersch et al. (2019, p. 8) report that among U.S. small and medium enterprises (SMEs) applying for loans from traditional lenders, 48 percent get only one quote, 40 percent get two or three quotes, and 13 percent get four or more quotes. According to Federal Reserve Banks (2015, p. 6), 72 percent of these firms borrow less than $250,000, making them substantially smaller borrowers than those in our sample. However, we can still assess the external validity of our estimates by comparing self-reported values from small-firm surveys with our estimates of search intensity among borrowers in the smallest loan bins. We estimate that in bins containing firms borrowing between $1 million and $5 million, 40 percent search once, 45 percent search two or three times, and 15 percent search four or more times. Remarkably, our estimates of search closely align with self-reported survey evidence of search intensity among borrowers.

We can also use survey evidence to assess whether our estimates of search costs align with those reported by firm managers. Federal Reserve Banks (2015, p. 4) find that SMEs, on average, spend 24 hours applying for credit when trying to originate a loan. This is likely a lower bound because it accounts only for the application process and does not include time spent identifying

potential lenders or calculating the financial data needed to complete a loan application. We can write the expected lower-bound search cost for firms in bin k as $\overline{\Delta}_k \equiv \sum_i q_{ik}\Delta_{ik}$ where q_{ik} is the value of q_i and Δ_{ik} in bin k. Our estimate of $\overline{\Delta}_k$ is 50 basis points for loans between \$1-5 million, which means that the expected lower-bound search cost for a loan of one million dollars is \$5,000. If a chief financial officer at an SME spends 24 hours working on the loan applications, this would imply that the CFO would have hourly compensation of \$208, which is well within the range of CFO compensation in small firms (Brewer, 2025). In other words, the opportunity cost of the self-reported time it takes a small firm to file loan applications yields a search cost estimate that is close to the values we estimate using the Burdett and Judd (1983) model. The fact that our estimates match both self-reported search frequencies and search costs gives us confidence that our estimates are plausible.

5.2.4 Search Costs and Information Disclosure

As we argued earlier, the dramatic difference in search and monitoring between bank lenders and bond investors provides a simple explanation for why the unexplained variation in loan spreads is so much higher than in bond spreads. We can explore whether information quality and availability also affect our loan search cost estimates by comparing them for public and private firms. Since listed firms are required to disclose a large amount of information in a standardized format, lenders can assess their risk more easily and monitor them with less effort. The greater information disclosure associated with being a public firm should mean that, all else equal, we should estimate lower search costs for public firms than for private ones. In contrast, if the unexplainable variance in loan spreads were due to random errors, we would not expect to see this pattern.

We test whether our search cost estimates capture this prediction by identifying 158 bins for which we have enough loans to split them into sub-bins containing either only listed or only private companies (limiting the analysis to sub-bins containing more than 5 banks originating loans and more than 20 originations). We report sample statistics for these two sets of sub-bins in the last two columns of Table 2. There is clearly greater dispersion in loan spreads for bins containing only private firms than for bins containing only listed borrowers. The mean p90-p10 range for private borrowers is forty basis points higher than for public ones. This greater dispersion in loan

spreads for private borrowers is consistent with higher search costs for private borrowers. Table 8 shows that our estimates confirm this intuition. The share of firms that search once, twice, and very intensively does not vary much across the two types of firms (about one-third, one-half, and one-sixth, respectively). However, the difference in the lower-bound search costs between private and public firms with identical characteristics is 14 basis points. This difference means that for a $39 million loan, roughly the average loan size for this set of borrowers, the lower-bound of the expected search cost for public firms is $55,000 less than for private firms. Thus, once again, our estimates align with the screening and monitoring hypothesis: firms that lower bank screening costs by releasing information have measurably lower search costs.

Table 8: Number of Searches and Lower-Bound Search Costs for Public and Private Firms

Number of Searches	Share (q_{iK})			Search Cost (Δ_{iK})		
	All	Public	Private	All	Public	Private
1	0.37	0.34	0.31	0.75	0.51	0.70
	(0.16)	(0.17)	(0.11)	(0.29)	(0.18)	(0.23)
2	0.48	0.47	0.54	0.32	0.21	0.28
	(0.13)	(0.16)	(0.08)	(0.13)	(0.09)	(0.10)
3	0.00	0.02	0.00	0.18	0.12	0.16
	(0.04)	(0.08)	(0.00)	(0.08)	(0.06)	(0.06)
4	0.00	0.01	0.00	0.13	0.08	0.11
	(0.01)	(0.06)	(0.00)	(0.06)	(0.05)	(0.05)
N–1	0.00	0.00	0.00	0.03	0.02	0.02
	(0.01)	(0.02)	(0.02)	(0.03)	(0.02)	(0.01)
N	0.14	0.16	0.14	0.00	0.00	0.00
	(0.07)	(0.11)	(0.05)	(0.00)	(0.00)	(0.00)

Note: This table presents the average estimates of q_{ik} and Δ_{ik} for a sample of sub-bins formed by dividing 158 bins into two sub-bins: one containing originations of public firms and the other containing originations of private firms. The units for Δ_{iK} are search cost per dollar borrowed, measured in percentage points, so numbers to the right of the decimal point are basis points. Standard deviations of the point estimates across bins within a category are in parentheses.

5.3 The Economic Significance of Search Costs

We can explore the economic significance of our estimates in two ways. First, we examine how much they vary across bins with different characteristics. Second, we conduct a counterfactual exercise to calculate the impact of a bank failure on the market value of client borrowers and compare this figure with stock market event studies examining the impact of rescuing a failing

bank.

5.3.1 Marginal Effects

One potential problem with the previous results is that when we examine differences across bins that just share one characteristic, it is difficult to isolate what is driving the result. For example, since A-rated firms tend to originate larger loans (see Table 1) and larger loans probably require more screening and monitoring, our loan rating results conflate the two effects. In Table 9, we address this problem by presenting results from regressing the expected lower-bound search cost in each bin ($\overline{\Delta}_k$) on rating, committed, and maturity fixed effects. Since these fixed effects always sum to one for each variable, all estimates are relative to A-rated loans, with committed amounts of \$1 million to \$5 million, and maturities of less than one year. These results provide a much clearer picture of how, say, increasing credit risk affects loan search costs, holding everything else constant.

The results suggest that the predictions of the screening and monitoring hypothesis are associated with large, statistically and economically significant movements in estimated search costs. In the first column, we show results for all bins in our sample. Relative to the omitted A-rated bins, search costs rise monotonically with credit risk. The magnitudes are particularly striking. Holding other loan characteristics fixed, a firm seeking a B-rated loan can expect to spend at least 25.3 basis points more on searching for every dollar borrowed. As we saw earlier, our lower-bound estimate of expected search costs per dollar borrowed decreases monotonically with loan size, consistent with a fixed cost associated with filing loan applications. All else equal, the lower-bound expected search cost per million dollars borrowed of a firm obtaining a \$1 million loan is \$1,220 more than for a firm obtaining a \$25 million loan.[14] These results enable us to formally reject the hypothesis that our search costs are generated by noise in favor of the hypothesis that our search cost estimates are consistent with the screening and monitoring hypothesis. The pattern of higher search costs for riskier loans persists across loan types. Columns (2) and (3) display the results for term and revolver loans. These display similar patterns.

[14] The screening and monitoring hypothesis does not have clear predictions about how search costs should vary with maturity. Nevertheless, for completeness, in Table A.2 and Table A.3, we present summary statistics and the search cost estimate for loans with different maturities. Search costs do not appear to vary much with maturity, at least for the 99 percent of originations of under eight years.

Table 9: Impact of Rating on Expected Lower-Bound Search Cost in a Bin ($\overline{\Delta}_k$)

	(1) All	(2) Term	(3) Revolver
Rating: BBB	0.104***	0.098*	0.116***
	(0.023)	(0.048)	(0.024)
Rating: BB	0.206***	0.214***	0.203***
	(0.022)	(0.046)	(0.023)
Rating: B	0.253***	0.275***	0.242***
	(0.023)	(0.051)	(0.023)
Committed: 5-10	-0.088***	-0.075**	-0.093***
	(0.018)	(0.028)	(0.024)
Committed: 10-25	-0.104***	-0.097***	-0.105***
	(0.017)	(0.026)	(0.021)
Committed: 25-100	-0.122***	-0.121***	-0.120***
	(0.018)	(0.030)	(0.022)
Committed: 100+	-0.177***	-0.275***	-0.156***
	(0.028)	(0.031)	(0.033)
Maturity: 1-4	0.009	-0.034	0.028
	(0.018)	(0.045)	(0.019)
Maturity: 4-5	-0.012	-0.051	0.004
	(0.017)	(0.045)	(0.018)
Maturity: 5-8	-0.032	-0.068	-0.014
	(0.023)	(0.047)	(0.041)
Maturity: 8+	-0.104***	-0.127*	-0.144*
	(0.028)	(0.049)	(0.068)
Observations	722	312	410
R^2	0.255	0.270	0.264

Note: The dependent variable is the expected lower-bound of the borrower search cost per dollar borrowed in bin k ($\overline{\Delta}_k$). The units for $\overline{\Delta}_k$ are search costs per dollar borrowed in percentage points, so numbers to the right of the decimal point are basis points. The sample in the first column includes all loans; the second column includes only term loans; and the third column includes only revolver loans. Standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

5.3.2 Stock Market Evidence

We can put these numbers into perspective by comparing them to stock market event studies of how the share prices of a troubled bank respond to news of a government capital injection. Bank bailouts benefit their borrowers in various ways, most obviously through continued lending, but another benefit is that they save the bank's customers from having to incur the search costs of finding a new lender.

Our estimates enable us to make a rough estimate of what fraction of the value of a bank bailout for a bank's borrowers could reasonably be attributable to saving them the cost of searching for new lenders. In a classic paper, Slovin et al. (1993, p. 257) found that the Federal Deposit Insurance Corporation's decision to intervene and prevent the liquidation of Continental Illinois Bank caused the share prices of its borrowers to rise by 1.2 percent.[15] Slovin et al. (1993) report that in their sample of borrowers, the average ratio of the value of bank debt to the value of equity was 0.89 (see their Table 1, Panel A). Our estimates in Table 8 for public companies imply that the lower bound of the search costs per dollar borrowed is 0.27 percent.[16] Therefore, if the ratio of search costs per dollar borrowed in their sample of public firms is the same as in ours, the implied savings for these customers would be .244 ($= 0.89 \times 0.002745$) percent of the value of its equity, or about one-fifth of the size of the estimated rescue value. Search costs for public firms that take their first loan quote have a lower bound of 51 basis points, so the impact on the share prices of these firms would be about twice as large. While this is obviously a crude calculation, it provides some credence for the idea that these search costs constitute a nontrivial fraction of the gains from bank bailouts.

6 Conclusion

We document that a large share of the dispersion in loan spreads arises among loans that are observationally nearly identical with respect to borrower risk, contract structure, and other characteristics commonly used in both academic and supervisory analyses. Using confidential supervisory data on a broad segment of the U.S. corporate loan market, we show that standard explanations based on observable risk and loan purpose, type, size, and maturity account for only a quarter of the

[15] More recently, Norden et al. (2013, p. 1650) examined the impact of Troubled Asset Relief Program announcements on the share prices of bank customers and obtained a remarkably similar estimate of 1.1 percent.

[16] Using the values for search costs for public firms given in Table 8, we have $0.2745 = 0.34 \times 0.51 + 0.47 \times 0.21 + 0.02 \times 0.12$.

variation in loan spreads.

We interpret this dispersion through the lens of a borrower search model and, to our knowledge, provide the first direct estimates of search costs in corporate credit markets. The estimated magnitudes are economically large and highly heterogeneous: a substantial fraction of firms behave as if they accept the first quote they receive, while others engage in limited or intensive comparison shopping. Search costs vary systematically with borrower risk, loan size, and information transparency, as predicted by the screening and monitoring hypothesis, and closely align with independent survey evidence on firms' reported borrowing experiences. Taken together, these results suggest that search frictions constitute a concrete, measurable financial friction that helps rationalize both price dispersion and borrowers' sensitivity to disruptions in bank lending relationships.

Our findings have broader implications for understanding the causes of financial frictions and for how shocks to the banking system propagate to the real economy. By tying dispersion in observed financing costs to underlying search and information frictions, we provide a complementary perspective to approaches that infer wedges indirectly from quantities or allocations. The results also suggest that policies or technological changes that reduce screening, monitoring, or information-acquisition costs—such as improved disclosure, standardization, or platform-based lending—could meaningfully compress borrowing costs and reduce firms' exposure to lender-specific shocks.

References

Alfaro, L., M. García-Santana, and E. Moral-Benito (2021). On the direct and indirect real effects of credit supply shocks. *Journal of Financial Economics 139*(3), 895–921.

Amiti, M. and D. E. Weinstein (2011). Exports and financial shocks. *The Quarterly Journal of Economics 126*(4), 1841–1877.

Amiti, M. and D. E. Weinstein (2018). How much do idiosyncratic bank shocks affect investment? evidence from matched bank-firm loan data. *Journal of Political Economy 126*(2), 525–587.

Ashcraft, A. B. (2005). Are banks really special? new evidence from the fdic-induced failure of healthy banks. *American Economic Review 95*(5), 1712–1730.

Banerjee, R. N., L. Gambacorta, and E. Sette (2021). The real effects of relationship lending. *Journal of Financial Intermediation 48*, 100923.

Beck, T., S. Da-Rocha-Lopes, and A. F. Silva (2021, April). Sharing the pain? credit supply and real effects of bank bail-ins. *The Review of Financial Studies 34*(4), 1747–1788.

Berger, A., A. Saunders, J. M. Scalise, and G. Udell (1998). The effects of bank mergers and acquisitions on small business lending. *Journal of Financial Economics 50*(2), 187–229.

Berger, A. N. and G. F. Udell (1995). Relationship lending and lines of credit in small firm finance. *The Journal of Business 68*(3), 351–382.

Bharath, S. T., S. Dahiya, A. Saunders, and A. Srinivasan (2011). Lending relationships and loan contract terms. *The Review of Financial Studies 24*(4), 1141–1203.

Bordeu, O., G. González, and M. Sorá (2025, May). Bank competition and investment costs across space. Working Paper 1042, Central Bank of Chile.

Brewer, K. (2025, June). The expected salary range for middle-market cfos.

Burdett, K. and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica 51*(4), 955–969.

Carlson, J. A. and R. P. McAfee (1983). Discrete equilibrium price dispersion. *Journal of Political Economy 91*(3), 480–493.

Cerqueiro, G., H. Degryse, and S. Ongena (2011). Rules versus discretion in loan rate setting. *Journal of Financial Intermediation 20*(4), 503–529.

Chen, B. S., S. G. Hanson, and J. C. Stein (2024). The decline of big-bank lending to small business: Dynamic impacts on local credit and labor markets. *Journal of Monetary Economics 144*, 103540.

Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008-9 financial crisis. *The Quarterly Journal of Economics 129*(1), 1–59.

Chodorow-Reich, G., O. Darmouni, S. Luck, and M. Ploser (2022). Bank liquidity provision across the firm size distribution. *Journal of Financial Economics 144*(3), 908–932.

Cortés, K. R., Y. Demyanyk, L. Lei, E. Loutskina, and P. E. Strahan (2020). Stress tests and small business lending. *Journal of Financial Economics 136*(1), 260–279.

Degryse, H. and S. Ongena (2005). Distance, lending relationships, and competition. *The Journal of Finance 60*(1), 231–266.

Diamond, D. W. (1984). Financial intermediation and delegated monitoring. *The Review of Economic Studies 51*(3), 393–414.

Federal Reserve Banks (2015, February). Joint small business credit survey report, 2014. Updated February 2015.

Gopinath, G., c. Kalemli-Özcan, L. Karabarbounis, and C. Villegas-Sanchez (2017). Capital allocation and productivity in south europe. *Quarterly Journal of Economics 132*(4), 1915–1967.

Hong, H. and M. Shum (2006). Using price distributions to estimate search costs. *The RAND Journal of Economics 37*(2), 257–275.

Honka, E., A. Hortaçsu, and M. Wildenbeest (2019). Chapter 4 - empirical search and consideration sets. In J.-P. Dubé and P. E. Rossi (Eds.), *Handbook of the Economics of Marketing, Volume 1*, Volume 1 of *Handbook of the Economics of Marketing*, pp. 193–257. North-Holland.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing tfp in china and india. *American Economic Review 99*(2), 140–165.

Jiménez, G., S. Ongena, J.-L. Peydró, and J. Saurina (2017). Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the spanish dynamic provisioning experiments. *Journal of Political Economy 125*(6), 2126–2177.

Martín-Oliver, A., V. Salas-Fumás, and J. Saurina (2008). Search cost and price dispersion in vertically related markets: The case of bank loans and deposits. *Review of Industrial Organization 33*(4), 297–323.

Mazet-Sonilhac, C. (2025). Search frictions in credit markets. *Sciences Po Economics Discussion Paper* (2024-14). Preprint; HAL ID: hal-04888145.

Moraga-González, J. L. and M. R. Wildenbeest (2008, July). Maximum likelihood estimation of search costs. *European Economic Review 52*(5), 820–848.

Morgan, P. and R. Manning (1985). Optimal search. *Econometrica 53*(4), 923–944.

Norden, L., P. Roosenboom, and T. Wang (2013). The impact of government intervention in banks on corporate borrowersâ stock returns. *Journal of Financial and Quantitative Analysis 48*(6), 1635–1662.

Peek, J. and E. 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.

Petersen, M. A. and R. G. Rajan (1994). The benefits of lending relationships: Evidence from small business data. *The Journal of Finance 49*(1), 3–37.

Pyle, D. and H. E. Leland (1977). Informational asymmetries, financial structure, and financial intermediation. *The Journal of Finance 32*(2), 371–387.

Rajan, R. G. (1992). Insiders and Outsiders: The Choice between Informed and Arm's-Length Debt. *The Journal of Finance 47*(4), 1367–1400.

Slovin, M. B., M. E. Sushka, and J. A. Polonchek (1993). The value of bank durability: Borrowers as bank stakeholders. *Journal of Finance 48*(1), 247–266.

Sufi, A. and A. M. Taylor (2022). Chapter 7 - financial crises: a survey. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics: International Macroeconomics*, Volume 6 of *Handbook of International Economics*, pp. 291–340. Elsevier.

Wiersch, A. M., S. Lieberman, and B. J. Lipman (2019, December). An update on online lender applicants from the small business credit survey. A publication of the Federal Reserve Bank of Cleveland.

Wiersch, A. M., B. J. Lipman, K. Wilson, and L. Misera (2022, August). Clicking for credit: Experiences of online lender applicants from the small business credit survey. Technical report, Federal Reserve Bank of Cleveland. A publication of the Federal Reserve Bank of Cleveland.

Online Appendix

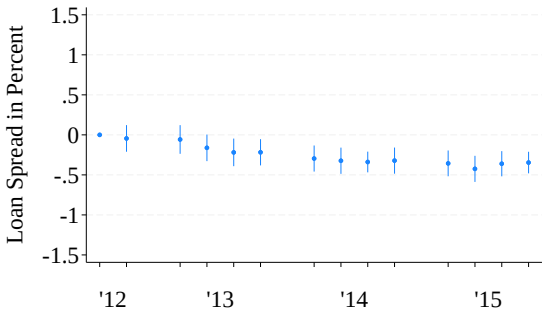
A Additional Results

Table A.1: Summary Statistics

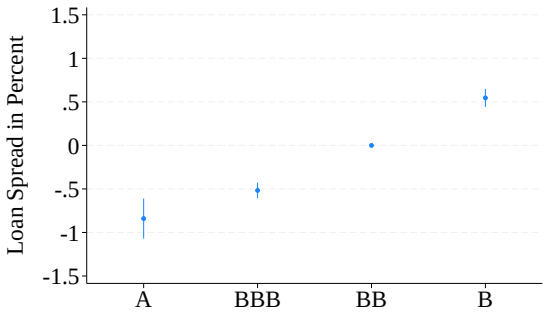
	Mean	Std. Dev.	p5	p25	Median	p75	p95	N. Obs.
Interest Rate	3.95	1.92	1.55	2.56	3.51	4.83	7.91	160,627
Spread	2.48	1.14	0.97	1.67	2.36	3.12	4.48	160,627
Relationship Length (quarters)	5.42	9.28	0.00	0.00	0.00	7.00	27.00	160,627
Normalized Relationship Length	0.82	1.17	0.00	0.00	0.00	1.68	3.06	160,627
Leverage	0.61	69.47	0.00	0.18	0.34	0.51	0.81	130,495
Short-Term Debt (millions)	205.53	24,119.96	0.00	0.00	0.04	3.12	72.93	136,117
Current Maturities of LT Debt (millions)	132.47	11,601.18	0.00	0.00	0.21	3.00	81.09	137,055
Long-Term Debt (billions)	1.29	62.34	0.00	0.00	0.00	0.11	2.76	138,073
Total Assets, Current (billions)	5.95	489.46	0.00	0.01	0.04	0.48	7.26	138,682
Log Sales	18.44	2.34	15.13	16.82	18.12	20.04	22.54	131,212
Net Sales, Current (millions)	2,786.91	93,817.84	0.05	16.14	62.65	438.19	5,765.00	138,855
Expected Loss	0.35	0.74	0.02	0.08	0.18	0.39	1.16	90,221
Interest Coverage	-311.61	278701.17	-3.17	2.20	6.10	17.48	124.41	111,887
Operating Income (millions)	383.44	22,539.35	-1.93	0.56	3.35	35.27	603.00	118,457
Interest Expense (millions)	61.81	2,670.03	0.00	0.06	0.38	6.07	130.31	137,642
Ticker (=1 if present)	0.13	0.34	0.00	0.00	0.00	0.00	1.00	160,627
Cusip (=1 if present)	0.14	0.34	0.00	0.00	0.00	0.00	1.00	160,627

Notes: The table presents summary statistics for variables used throughout the analysis. Spreads equal the difference between the interest rate of a newly originated loan and the reference rate. Relationship Length equals the difference in the number of quarters between the current loan origination and the earliest observed loan in that bank-firm pair. The variable is then normalized by dividing by the mean relationship length in each quarter. Leverage is equal to (Short Term Debt + Current Maturities of Long-Term (LT) Debt + Long-Term Debt)/Total Assets. Log Sales equals the log of Current Net Sales. Expected Loss equals the product of Probability of Default, Loss Given Default, and Exposure at Default, divided by the committed loan amount. Interest Coverage equals Operating Income divided by Interest Expense.

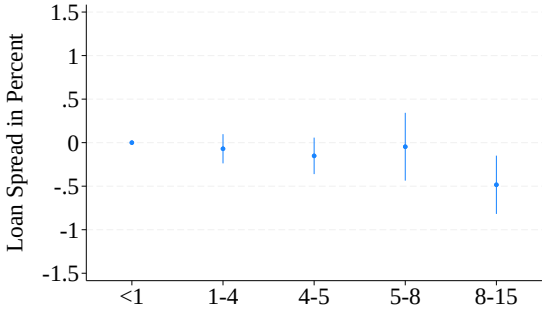
Figure A.1: Impact of Loans Characteristics on Loan Spreads (Regime 1)



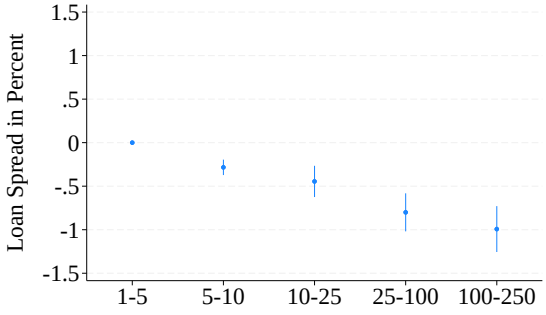
(A) Loan Origination Date Fixed Effect



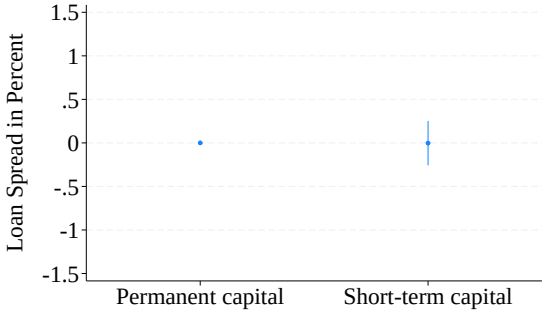
(B) Loan Rating Fixed Effects



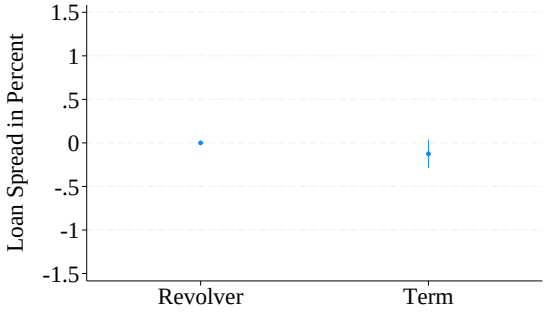
(C) Maturity Fixed Effect (years)



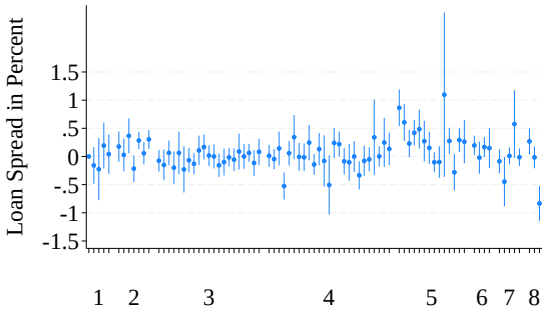
(D) Committed Loan Size Fixed Effect (millions)



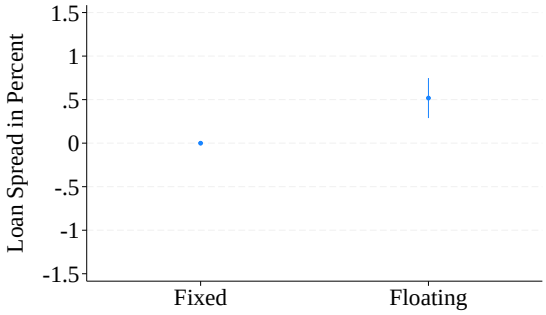
(E) Loan Purpose Fixed Effects



(F) Term Loan vs. Revolver Fixed Effects



(G) 3-Digit NAICS Fixed Effect



(H) Loan Rate Type Fixed Effect

Note: Notes: These graphs show the coefficients and 95 percent confidence intervals of a regression of the loan spread at the origination date of the loan on various fixed effects. The sample covers the 66,829 originations (the same sample as Table 1) from 2012:Q4 to 2015:Q4. Figure (G)'s x-axis displays 1-digit NAICS categories, but fixed effects are estimated at the 3-digit level.

Table A.2: Average Loan Characteristics By Maturity

	All Bins	0-1	1-4	4-5	5-8	8-15
Number of Banks	13.00	12.35	13.43	13.43	12.37	11.30
	(4.88)	(4.32)	(4.88)	(5.32)	(4.53)	(4.35)
Number of Loans	222.48	339.01	231.66	196.86	99.69	109.30
	(379.65)	(665.19)	(284.63)	(240.83)	(113.90)	(90.15)
Maturity (years)	3.72	0.75	2.70	4.91	6.39	10.54
	(2.33)	(0.12)	(0.31)	(0.07)	(0.42)	(0.50)
Loan Size (millions)	22.36	13.61	26.69	29.36	13.71	2.03
	(34.24)	(16.34)	(36.59)	(43.18)	(19.83)	(0.23)
Mean (spread)	2.33	2.41	2.35	2.26	2.37	2.21
	(0.71)	(0.67)	(0.68)	(0.76)	(0.77)	(0.57)
Median (spread)	2.19	2.31	2.21	2.08	2.23	2.06
	(0.70)	(0.66)	(0.68)	(0.73)	(0.72)	(0.53)
p90–p10 Range (spread)	2.41	2.38	2.38	2.42	2.48	2.41
	(0.93)	(1.02)	(0.74)	(1.00)	(1.03)	(0.94)
Interquartile Range (spread)	1.21	1.18	1.24	1.20	1.20	1.14
	(0.49)	(0.40)	(0.42)	(0.57)	(0.54)	(0.53)
Number of Bins	722	153	224	230	94	20

Note: This table presents summary statistics for the sample of 160,627 loan originations divided into 722 bins. Each number is calculated by first computing the average of a variable within a bin, then averaging the means across bins for a given characteristic. Columns 2-6 present averages for different maturity buckets. Standard deviations of the averages within bins are reported in parentheses.

Table A.3: Number of Search Costs and Average Search Costs Across Bins By Maturity

Number of Searches	Share (q_{iK})						Search Cost (Δ_{iK})					
	All	0-1	1-4	4-5	5-8	8-15	All	0-1	1-4	4-5	5-8	8-15
1	0.37	0.40	0.36	0.34	0.38	0.42	0.75	0.77	0.76	0.74	0.74	0.64
	(0.16)	(0.18)	(0.14)	(0.15)	(0.17)	(0.12)	(0.29)	(0.32)	(0.30)	(0.26)	(0.29)	(0.14)
2	0.48	0.45	0.49	0.50	0.46	0.43	0.32	0.33	0.32	0.31	0.32	0.29
	(0.13)	(0.14)	(0.11)	(0.13)	(0.15)	(0.12)	(0.13)	(0.15)	(0.14)	(0.12)	(0.14)	(0.07)
3	0.00	0.00	0.00	0.00	0.01	0.02	0.18	0.20	0.19	0.18	0.19	0.17
	(0.04)	(0.03)	(0.04)	(0.03)	(0.05)	(0.07)	(0.08)	(0.09)	(0.09)	(0.08)	(0.09)	(0.05)
4	0.00	0.00	0.00	0.00	0.00	0.00	0.13	0.13	0.13	0.12	0.13	0.12
	(0.01)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.04)
N–1	0.00	0.00	0.00	0.00	0.00	0.00	0.03	0.04	0.03	0.03	0.04	0.04
	(0.01)	(0.00)	(0.02)	(0.00)	(0.00)	(0.00)	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.02)
N	0.14	0.15	0.14	0.14	0.15	0.13	0.00	0.00	0.00	0.00	0.00	0.00
	(0.07)	(0.08)	(0.06)	(0.08)	(0.08)	(0.05)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Note: This table presents the average estimates of q_{ik} and Δ_{ik} across bins containing loans of identical maturity for the sample of 160,627 loan originations grouped into 722 bins of loans of otherwise nearly identical characteristics. Maturity buckets are expressed in years (e.g., 1-4 refers to 1-4 years). The units for Δ_{iK} are percentage points, so numbers to the right of the decimal point are basis points. Standard deviations of the point estimates across bins within a category are in parentheses.

Table A.4: Average Loan Characteristics By Regime

	All Bins	R1	R2	R3	R4
Number of Banks	13.00	12.65	14.47	12.03	12.29
	(4.88)	(4.70)	(5.73)	(3.71)	(4.24)
Number of Loans	222.48	277.30	236.53	124.69	193.60
	(379.65)	(494.91)	(373.33)	(152.39)	(273.85)
Maturity (years)	3.72	3.73	3.95	3.56	3.50
	(2.33)	(2.26)	(2.59)	(2.16)	(2.17)
Loan Size (millions)	22.36	16.65	20.48	24.15	32.72
	(34.24)	(22.91)	(30.03)	(34.80)	(49.19)
Mean (spread)	2.33	2.39	2.15	2.47	2.38
	(0.71)	(0.80)	(0.62)	(0.69)	(0.65)
Median (spread)	2.19	2.24	2.01	2.31	2.26
	(0.70)	(0.76)	(0.62)	(0.70)	(0.65)
p90–p10 Range (spread)	2.41	2.62	2.24	2.44	2.28
	(0.93)	(1.07)	(0.83)	(0.81)	(0.86)
Interquartile Range (spread)	1.21	1.32	1.11	1.23	1.14
	(0.49)	(0.54)	(0.44)	(0.46)	(0.43)
Number of Bins	722	241	210	121	150

Notes: This table presents summary statistics for the sample of 160,627 loan originations divided into 722 bins. Each number is calculated by first computing the average of a variable within a bin, then averaging the means across bins for a given characteristic. Columns 2-5 present averages for different monetary policy regimes. R1 stands for Regime 1, R2 for Regime 2, R3 for Regime 3, and R4 for Regime 4. Standard deviations of the averages within bins are reported in parentheses.

Table A.5: Number of Search Costs and Average Search Costs Across Bins By Regime

Number of Searches	Share (q_{iK})					Search Cost (Δ_{iK})				
	All	R1	R2	R3	R4	All	R1	R2	R3	R4
1	0.37	0.37	0.35	0.36	0.39	0.75	0.86	0.69	0.65	0.73
	(0.16)	(0.17)	(0.15)	(0.15)	(0.16)	(0.29)	(0.33)	(0.24)	(0.22)	(0.26)
2	0.48	0.48	0.50	0.48	0.46	0.32	0.36	0.29	0.27	0.32
	(0.13)	(0.14)	(0.12)	(0.13)	(0.12)	(0.13)	(0.15)	(0.11)	(0.10)	(0.13)
3	0.00	0.00	0.00	0.01	0.00	0.18	0.21	0.17	0.16	0.19
	(0.04)	(0.05)	(0.03)	(0.06)	(0.01)	(0.08)	(0.10)	(0.07)	(0.06)	(0.08)
4	0.00	0.00	0.00	0.00	0.00	0.13	0.14	0.11	0.10	0.13
	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.06)	(0.07)	(0.05)	(0.04)	(0.06)
N–1	0.00	0.00	0.00	0.00	0.00	0.03	0.04	0.03	0.03	0.04
	(0.01)	(0.00)	(0.00)	(0.03)	(0.00)	(0.03)	(0.04)	(0.03)	(0.02)	(0.03)
N	0.14	0.15	0.14	0.13	0.14	0.00	0.00	0.00	0.00	0.00
	(0.07)	(0.07)	(0.07)	(0.09)	(0.07)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Note: This table presents the average estimates of q_{ik} and Δ_{ik} across bins containing loans of identical monetary policy regimes for the sample of 160,627 loan originations grouped into 722 bins of loans of otherwise nearly identical characteristics. The units for Δ_{iK} are percentage points, so numbers to the right of the decimal point are basis points. R1 stands for Regime 1, R2 for Regime 2, R3 for Regime 3, and R4 for Regime 4. Standard deviations of the point estimates across bins within a category are in parentheses.

Table A.6: Loan Characteristics (Regime 1)

	Percent of obs	Percent of committed	Mean spread	Median spread	Mean interest rate	Median interest rate	Mean utilization	SD spread	Mean Expected Loss
Rating:									
A	1.7	3.0	1.78	1.67	2.07	1.95	63.91	0.90	0.04
BBB	28.1	42.4	2.05	1.90	2.42	2.20	65.98	0.99	0.10
BB	55.9	45.3	2.68	2.54	3.09	2.99	70.36	1.14	0.28
B	14.4	9.4	3.32	3.13	3.69	3.50	73.48	1.39	0.96
Committed (millions):									
1-5	55.8	9.4	2.82	2.74	3.28	3.25	73.10	1.17	0.38
5-10	12.7	7.1	2.59	2.42	2.92	2.70	70.89	1.24	0.31
10-25	15.4	19.5	2.40	2.17	2.72	2.45	68.15	1.23	0.30
25-100	15.2	54.8	1.92	1.70	2.24	2.00	57.69	1.03	0.20
100-250	0.9	9.2	1.54	1.42	1.80	1.68	46.04	0.64	0.11
Maturity (years):									
0-1	34.7	13.6	2.90	2.93	3.18	3.20	63.96	1.09	0.39
1-4	32.6	40.3	2.52	2.39	2.85	2.69	67.53	1.21	0.34
4-5	25.1	42.6	2.27	1.97	2.73	2.44	73.59	1.25	0.27
5-8	6.4	3.3	2.46	2.19	3.26	3.15	87.81	1.33	0.29
8-15	1.3	0.2	2.14	1.95	3.84	3.99	96.26	1.06	0.39
>15									
Loan type:									
Revolver, short-term capital	23.5	13.5	2.89	2.90	3.18	3.17	60.55	1.07	0.34
Revolver, general working capital	49.2	71.6	2.51	2.40	2.86	2.71	60.04	1.19	0.28
Term, short-term capital	8.0	2.6	2.34	2.08	2.82	2.65	95.78	1.32	0.26
Term, general working capital	19.4	12.2	2.47	2.27	3.09	2.92	93.42	1.31	0.41
Loan Rate type:									
Fixed	14.3	7.7	2.19	2.10	3.34	3.25	88.66	1.41	0.35
Floating	85.7	92.3	2.64	2.47	2.91	2.75	66.28	1.16	0.32

Note: This table reports summary statistics for the sample that only encompasses the first regime (from 2012Q4 until December 16, 2015). This sample includes 66,829 loan originations. We drop originations of revolver loans that were never drawn over the life of the loan, loans in bins with fewer than six unique banks, or bins with fewer than 20 loans. All values are calculated as of the origination date. Floating-rate loan spreads are calculated based on spreads to 3-month LIBOR or SOFR, depending on the period, and fixed-rate loan spreads are based on spreads to maturity-matched U.S. Treasury yield curves at origination. The mean utilization is the mean loan value in the category, defined as utilized divided by committed, calculated for each loan over its lifespan in the dataset.

Table A.7: Loan Spreads and Expected Loss (Regime 1)

	Loan Spread		
	(1)	(2)	(3)
Expected Loss Ratio			0.245**
			(0.078)
Observations	66,829	13,739	13,739
Bin FE	Y	Y	Y
R^2	0.25	0.26	0.27
RMSE	1.05	1.02	1.01

Note: Column 1 regresses the loan spread over a set of bin fixed effects, using the sample that only encompasses the first regime (from 2012Q4 until December 16, 2015). This sample includes 66,829 loan originations. Column 2 restricts the sample to observations for which expected loss data are available. Column 3 introduces the Expected Loss Ratio as a control variable; this ratio is calculated as the product of Probability of Default, Loss Given Default, and Exposure at Default, divided by the committed loan amount. RMSE stands for "Root Mean Squared Error". Robust standard errors in parentheses. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$