How Much Does It Cost for Firms to Switch Banks?*

Mary Amiti
Federal Reserve Bank of New York and CEPR

Anil K Kashyap
University of Chicago Booth School of Business, CEPR, and NBER

Anna Kovner
Federal Reserve Bank of Richmond

David E. Weinstein
Columbia University, CEPR, and NBER

March 21, 2025

Abstract

The cost for a firm to switch banks is a central factor in determining whether a contraction in lending translates into effects on the real economy. We use confidential supervisory data on the largest US banks and their commercial and industrial borrowers to determine borrowers' switching costs. Typically a borrower rolls over an expiring loan at their original lender, but in some cases doing so indicates that they are forgoing substantial savings. We estimate that a third of borrowers do not switch banks because the search costs are too high and that the lower bound of these search costs is 72 basis points. About half of borrowers have lower search costs, in the range of 30 to 72 basis points, leading them to a search for at least another loan quote. Only 14 percent of borrowers have search costs that are so low, around 3 basis points, that they search broadly. The search costs are highest for smaller and riskier borrowers. In a counterfactual exercise, we estimate that a bank failure that forced its borrowers to search for new lenders would result in a deadweight loss of at least 11 percent of the bank's loans.

Key Words: banking, credit supply, macrofinance
JEL: E51,G21,G32

*Amiti: Federal Reserve Bank of New York mary.amiti@ny.frb.org; Kashyap: University of Chicago Booth School of Business, anil.kashyap@ChicagoBooth.edu, Kovner: Federal Reserve Bank of Richmond, anna.kovner@rich.frb.org; Weinstein: Columbia University, dew35@columbia.edu. We thank Sean Fulmer and Augustin Belin for excellent research assistance. The opinions expressed in this paper are the authors and do not necessarily reflect those of the Federal Reserve Banks of New York, Richmond or the Federal Reserve System.

1 Introduction

Reductions in bank lending and bank failures generate negative externalities for their borrowers. This empirical regularity motivates bank supervision and regulation as well as policy interventions in the face of bank fragility. This paper uses unique supervisory data to quantify the costs to firms of changing banks. These costs can then inform estimates of the impact of changes in bank interest rates and the potential impact of a bank failure. We find that these costs vary with borrower characteristics, with smaller, riskier firms having higher costs to switch banks. The implied cost of bank failure affects borrowers differently, and the cost to society varies with bank characteristics.

We calculate our estimates from supervisory data on C&I lending compiled by the Federal Reserve for stress testing between 2012 and 2023. These comprehensive data on lending by the largest U.S. banks allow us to infer how much it costs different types of borrowers to switch banks. The key to this calculation is accurately accounting for the variation in interest rates that comes from credit risk. The data set we use was built by supervisors to measure credit risk, more specifically to determine how different loans would perform under stress, and thus naturally lends itself to this purpose. Hence, we have rich detail on the essential aspects of each loan in the banks' portfolios over $1 million, including banks' internal risk scores and, in many cases, numeric information on the expected loss in the event of a default, including the probability of default, loss given default, and exposure at default. We group loans into four sub-periods of the sample based on whether interest rates were near the effective lower bound and within each time period to compare loans that match on a host of loan features, including the lender's risk rating, loan size, maturity, loan purpose and loan type (revolving vs term). Our first major finding is that these risk "bins" can explain about 50% of the variance of observed interest rates.

We then develop a model of how lenders would set interest rates for a set of borrowers with homogenous levels of risk, but differences in search costs. As in Burdett and Judd (1983), the lenders target different types of borrowers, with some setting high rates that only attract borrowers with high search costs, and others posting lower rates that will attract more customers. In equilibrium, all strategies yield the same level of expected profits because loan offers with high rates are chosen less frequently than loans with lower rates. By observing the distribution of loan rates that are offered, we can infer the differences in the search costs. We follow the banking literature in interpreting these search costs as arising because a new lender must screen and monitor the borrower to assess creditworthiness. We model one component of these screening and monitory costs as being systematic, i.e., varying with observables like risk rating, and some of the screening costs as being idiosyncratic. These idiosyncratic costs vary across borrowers in a risk bin and reflect the difficulty a firm faces in finding a low-cost loan provider, preparing the appropriate paperwork to satisfy the bank, etc. Thus, we expect that low risk firms will pay lower interest rates because they are easier to screen and monitor than high-risk firms and also that they will have lower screening and search costs because banks can evaluate their risk relatively easily. We dub the idea that screening and monitoring costs underlie switching costs the "screening and monitoring hypothesis".

Our sample includes more than 148,000 loan originations, with the average loan size is around $21 million which we sort into 688 bins of borrowers with similar risk characteristics. We estimate that approximately

35% of borrowers get one quote and accept it, about 49% get a second a quote and the remaining borrowers seem to search intensively enough to see all quotes that are available. This search behavior is consistent with the single search borrowers (seeking a $21 million loan) having a search cost of at least roughly $153,000, the second group of firms having a cost of at least $63,000 and the other firms having a cost of at least $6,300. These averages, however, cover up a great deal of heterogeneity in imputed costs for borrowers in different bins.

The screening and monitoring hypothesis implies that borrowers whose loans are harder to evaluate or have loans that allow more discretion to the borrower require lenders to expend more resources on screening and monitoring. These loans should, therefore, have higher switching costs and this prediction is confirmed in the data. For example, holding fixed the loan size, a typical A-rated borrower who gets a second loan quote has a switching cost that is about half as large as a B-rated borrower who gets the same number of quotes. This finding is driven by the empirical regularity that as the borrowers' risk in a given bin decreases (as measured by higher credit ratings), the amount of dispersion in interest rates falls.

We use our estimates to compute the cost to the financial system of having a bank fail and forcing all of its customers to obtain loans from other banks. We make this calculation by cumulating the borrower switching costs incurred at origination and dividing this number by the amount of loans outstanding to express the deadweight loss of failure as a share of bank lending. Because the 35 banks in the sample differ in terms of their customer mix, the implied costs of a failure differ too. For the median bank, the estimated deadweight loss would be 11% of the loans and for the average bank the loss would be 15%.

Section 2 selectively reviews the main literatures where loan supply shocks are thought to be important. It also describes the theoretical papers that underlie the screening and monitoring hypothesis. Section 3 introduces the theory that we use to estimate the model. Section 5 introduces the data used in the analysis. We describe how the bins are constructed and assess the degree of interest rate variation in the different bins. We present evidence suggesting that the within-bin variation does not appear to be associated with observable variables related to credit risk and showing that substantial variation in interest rates, unexplained by loan risk characteristics collected by supervisors remains. Section 6 presents our estimates of the cost of switching. We first present some univariate comparisons of interest rates when we characterize loans by a single characteristic and then supplement that with a multivariate analysis of all the main characteristics. We then present the estimates of the switching costs for different types of borrowers and loans and close the section with a summary of the switching costs that a failing bank would incur. Section 8 concludes.

2 Literature Review

At least four well-developed literatures study shocks to lending and interpret the results as supporting the existence of switching costs. Given the depth of these literatures, we do not attempt to fully summarize all the related research and instead point to some prominent papers that give a sense of the relevant findings. The direct estimates of switching costs we provide could be used to revisit these topics to reconsider the implications of lending shocks for the economy.

One topic area concerns the impact of bank failures and bailouts on the economy. An early, influential

paper by Slovin, Sushka, and Poloncheck (1993) examined the fate of borrowers affiliated with Continental Illinois Bank around the time of its distress in 1984. They find that share prices for the customers fell after the bank ran into trouble (as might be expected if the borrowers cannot seamlessly replace any lost loans). They could not measure which borrowers made up any reduction in lending by switching to other banks.

Following the 2008/9 Global Financial Crisis, several studies examined the consequences of the failure of Lehman Brothers and the stress at other large banks. Chodorow-Reich (2014) studies the syndicated loan market and identifies borrowers from banks that incurred large real estate-related losses and estimates how much lending they lost. He then estimates employment outcomes for the firms that lost funding and shows that the credit losses translated into employment declines. The employment drops were largest for small and medium-sized firms and insignificant for large firms. This pattern is consistent with the idea that small borrowers faced disproportionately higher costs to switch lenders.[1]

Our counterfactual bank failure results produce larger results than in papers examining the direct administrative costs that are associated with bankruptcy filings. For instance, one well-known study by Warner (1977) estimates that bankrupt railroads between 1933 and 1955 incurred administrative expenses during bankruptcy of 4% of their pre-bankruptcy market value. In another widely-cited study, Weiss (1990) finds administrative costs of about 3% of market value for a set of firms that filed for Chapter 11 bankruptcy protection between 1980 and 1986. However, these studies mostly examine nonfinancial firms.

The observation that borrowers and lenders match in ways that are beneficial to both parties is also indirectly supported by the large literature on relationship lending. For instance, Bharath et al. (2011) find repeated loans between a borrower-lender pair are associated with lower loan spreads, particularly when the borrower is relatively opaque. These findings are also supportive of the view that switching costs are important. Our results also cast new light on the finding by Billett, Flannery, and Garfinkel (1995) that lenders with a higher credit rating are associated with larger abnormal returns upon receiving a loan, as it may also signal a lower probability of a firm being forced to incur fixed costs of switching.

A second set of analyses investigates lending disruptions that are created as a result of bank mergers. Berger et al. (1998) study over 6,000 mergers and find that small business lending drops at the time of a merger. In subsequent quarters, however, they find that lending recovers.[2] Karceski, Ongena, and Smith (2005) look at stock price reactions for publicly traded customers of Norwegian banks that are involved in mergers. Customers of acquired banks have negative share price reactions, though when borrowers are dropped by the bank, the share price reaction is less negative. They interpret this evidence as reflecting different switching costs for different customers, whereby the borrowers with low switching costs move and are less likely to be held up while the firms that stay behind are charged higher loan rates because their switching costs are prohibitive.

A third related area of study looks at the monetary policy transmission process, emphasizing the impact of policy changes on bank lending. Kashyap and Stein (2000) show that U.S. banks differ systematically in the extent to which they reduce lending after a tightening of monetary policy. Smaller banks with less liquid

1. Amiti and Weinstein (2011, 2018) find that a contraction in bank supply reduces their customers' export growth and investment, respectively.

2. This is reinforced by Peek and Rosengren (1998), who find that loan portfolios converge towards the structure of the acquiring bank. Because most mergers involve small banks, this will mean that small borrowers are not necessarily cast aside after a merger.

assets tend to cut lending more than similar-sized banks with high levels of liquid assets. They interpret this result as arising because smaller banks cannot raise market financing easily, so banks have to shrink when deposits are reduced (via the monetary policy tightening). Those with abundant liquidity can insulate their lending business, whereas the banks with few liquid assets have to reduce lending. Large banks that presumably can better raise alternative financing do not show this pattern. Because small banks tend to lend to small borrowers, the effect of a monetary tightening will differ for different firms.[3] Kashyap and Stein (2000) find that a given monetary policy tightening has different effects on different types of borrowers. Implicit in these cross-sectinoal analyses is that borrowers do not switch away from banks that are differently affected by changes, and unanswered in that literature is how many firms can switch and replace funding. One contribution of this paper is to offer numbers that are relevant to interpreting the literature on the impact of bank supply shocks

Using even better data from Spain, Jiménez et al. (2012) find similar results. They rely on the Spanish credit register that tracks all lending and loan applications. They find that when monetary policy tightens, loans granted by less well-capitalized banks with low levels of liquid assets decline significantly compared to strong-capitalized, highly liquid banks. They can also see if firms that were denied credit make successful applications elsewhere, and they find that about one-third of the borrowers that were turned down have another application successfully approved within the next three months. So, while there appears to be some substitutability, it seems to be imperfect. Therefore, this research also suggests that when bank-borrower relations are disrupted, the borrower incurs some costs.

Finally, there are many analyses of how changes in bank regulation, primarily capital regulation, affect bank customers. If only some banks are affected by regulatory changes and switching costs are low, the impact on borrowers should be limited. However, most studies do find significant costs to borrowers when bank capital standards are increased. Using Spanish data, Jiménez et al. (2017) document that when capital standards increase for some banks, their borrowers struggle to replace the funding. In the US, Cortés et al. (2020) examine lending to small businesses by the largest US banks after they become subject to stress testing and find that banks more affected by stress tests pass the implied increase in capital requirements into loan rates in markets where they have local knowledge, a price increase suggesting that borrowers face high costs of switching.

In addition to these empirical literatures, a large theoretical literature rationalizes the existence of switching costs and predicts how they might differ across borrowers. We use this literature to assess the plausibility of our estimates for different types of borrowers and loans. This theoretical literature highlights the problems created by asymmetric information between borrowers and lenders, whereby borrowers know more about their creditworthiness than lenders do. Financial intermediaries are valuable in such settings because they invest in information production (in particular screening borrowers before making loans and monitoring them afterward to ensure repayment).[4] Switching costs, however, are still likely to arise because once a lender becomes informed about a borrower, that lender has an edge in extending credit relative to uninformed

3. Drechsler, Savov, and Schnabl (2017) present complementary evidence that emphasizes differences in market power in the deposit market as a source of heterogeneity of monetary policy changes on lending.

4. Classic references include Pyle and Leland (1977) and Diamond (1984) who investigate how lenders can screen and monitor borrowers to address the information problems.

lenders.

This reasoning suggests that switching costs are likely to be highest in cases where information asymmetries are largest. Standard proxies for the extent of information problems include borrower size (where smaller firms are assumed to be more opaque), loan purpose (where loans that allow more discretion in their use are assumed to be less transparent) and the nature of the borrower's business (with more complex organizations or complicated competitive environments assumed to be more opaque).[5] We refer to these types of predictions as "the screening and monitoring hypothesis". We provide some comparisons of our estimates along these dimensions to assess this hypothesis.

Past research finds that loan interest rates tend to vary in ways consistent with this hypothesis. For example, smaller firms or firms without credit ratings pay more to borrow Hubbard, Kuttner, and Palia (2002). Similarly, loans for general corporate purposes are more expensive than those restricted to working capital. These kinds of findings, however, do not directly tie the loan pricing to the presence of switching costs.

3 The Bank Search Model

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 considered the problem of estimating consumer search costs for low-cost seller when the econometrician can only observe prices, we adapt it to a case in which we want to estimate borrower search costs when we can observe interest rates. We assume that it is possible to create loan categories such that within a loan category such that borrowers are homogeneous within the category and loans offered by different banks are viewed as homogeneous. The empirical counterpart to a category in the bank loan context might be all loans to firms with a particular risk rating, for the same purpose, with identical restrictions on interest-rate adjustment (i.e., fixed vs. floating), and in a narrow band of committed amounts, maturities, and origination dates. In what follows, we describe how we could identify search costs within one of these loan categories. We spend significant time in Section 5.2 documenting that loans we place in the same category have the same risk.

We start by establishing some notation. Let ρ be the common cost for a bank to offer a loan of amount L. Each borrower has a reservation interest rate of $\bar{r}$ above which it will not borrow, and we assume that every borrower costlessly observes an interest rate offered by one random bank. Borrowers can also obtain $i-1 \geq 0$ additional interest rate 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 for each borrower and is drawn from a distribution of search costs F_c.[6]

Following Burdett and Judd (1983), the interest rates offered by the N banks will correspond to a mixed strategy equilibrium in which the equilibrium distribution of interest rates that banks choose, F_r, makes each bank indifferent between charging any interest rate between the lower bound of the support for F_r ($\underline{r} > \rho$) and the upper bound ($\bar{r}$). This distribution of interest rates ensures that the per loan gains (losses) associated

5. Similar kinds of predictions can be made for consumer lending. In that case proxies might be based on the length of an individual's borrowing history, the variability in their income, and the exposure to unexpected spending requirements.

6. Morgan and Manning (1985) show that simultaneous search may be optimal if agents need to make decisions rapidly and sequential search is time consuming.

with offering a higher (lower) interest rate are exactly offset by the loss (gain) of customers associated with the higher (lower) interest rate. The optimal number of interest rate price quotes a borrower will obtain is given by

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

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

We also can also use this structure to characterize the optimal number of interest-rate quotes a borrower will obtain. Since $i(c)$ is an integer, we divide borrowers into B 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 interest rate in a sample of i prices drawn from F_r by $E[r_{1,i}]$. We can write the search cost that would render a borrower indifferent (in expectation) between sampling i interest rates and $i+1$ interest rates as

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

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

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

where F_c is the CDF of the distribution of costs.

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

$$(\bar{r} - \rho) q_1 L = (r - \rho) L \left[\sum_{i=1}^{N} i q_i \left(1 - F_r(r)\right)^{i-1} \right], \tag{4}$$

where the left-hand side of this equation denotes the fraction of profits going to bank charging the maximum interest rate, 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

$$(r_1 - \rho)q_1 L = (r_i - \rho)L \left[\sum_{i=1}^{N} i q_i \left(1 - \hat{F}_r(r_i) \right)^{i-1} \right] \quad i = 2, 3, \ldots, N-1, \tag{5}$$

where r_i is the observed interest rate charged by bank i, r_1 is the maximum observed interest rate, and $\hat{F}_r(r_i)$ is the fraction of banks offering an interest rate 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 3, 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) note that one can improve on the Hong and Shum (2006) estimation by using a maximum likelihood estimator because estimating $\hat{F}_r(r_i)$ using the empirical distribution is not efficient. They rewrite the system of equations described above as a maximum likelihood problem that is efficient and provides us with standard errors for the estimated variables. They show that this method performs well in Monte Carlo trials even when one does not observe loans offered by some banks (prices offered by firms in their implementation). We will therefore follow their approach.

4 Counterfactual Exercise

We consider a counterfactual exercise that enables us to estimate the cost of a bank's failure. In particular, we assume that the failure of a bank forces all firms to incur search costs by searching for a new lender in the same bin. This approach potentially overstates the cost of a bank failure because some its borrowers might choose not to borrow or because in some bank failures, the business is acquired through resolution, meaning borrowers may not need to seek new lenders.

We model a bank failure as an exogenous event that forces all firms associated with a bank to search again for new loans, i.e., we are not allowing borrowers to find alternative financing outside of the bin in response to a bank failure. In the base case, we consider the failed bank to be "small" in the sense that when a bank fails, it does not affect the marginal cost of lending of other institutions and the bank's failure induces another bank to enter the market so the equilibrium number of competitors is unaffected. The equilibrium interest-rate distribution will continue to be optimal given the distribution of search costs, so our estimates of q_i and Δ_i describe the properties of the new equilibrium. The only impact of the failure, then, is to force the bank's borrowers to incur a search cost in order to continue borrowing.

For $i < N$, a fraction q_i of the borrowers have a search cost in excess of Δ_i, and q_N borrowers have a lower-bound search cost of 0. Let L_b be the average value of loans issued by bank b to all its borrowers in the bin. We can write the average lower-bound search costs incurred by borrowers from that bank who in equilibrium search i times by $i\Delta_i L_b$. In general, we can write the average lower-bound for the search costs for borrowers searching i times as

$$C = \sum_{i=1}^{N-1} i q_i L_b \Delta_i. \tag{6}$$

Note that C and L_b are the average cost and loan amount in a bin, so to compute the total cost we need to sum C across all bins. We denote this sum by TC. Since L_b will tend to be larger for large banks, total costs will also tend to increase with bank size. It therefore makes sense to express these costs as a share of

a bank's total lending in all bins. If we denote a bank's total lending in all bins in in the estimation sample by TL_b, we can write the bankruptcy cost, expressed as a share of total lending by $SC = TC/TL_b$. Thus, the variables we want to know are TC and SC. Since these vary by bank, while we calculate the result at the individual bank level, we cannot report individual estimates due to data confidentiality concerns, but we can report the mean and standard deviation of the costs across banks in the sample.

5 Data

Our analysis uses a unique matched U.S. bank-firm data set collected by the Federal Reserve for Dodd-Frank Act mandated stress testing (FR Y-14). We focus on the H.1 schedule, which includes a comprehensive dataset of all loans to commercial and industrial borrowers (C&I loans) made by stress-tested banks. This supervisory data is collected quarterly for US bank holding companies, intermediate holding companies of foreign banks, and savings and loan holding companies with more than $100 billion in total assets, making it the most comprehensive data source for US bank loans. [7]

In aggregate, Y-14 data include information on 42 bank holding companies, accounting for approximately 70% of all US commercial bank lending. 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 we can gain information about smaller businesses, which are more likely to be bank dependent.

5.1 Sample Description

The Y-14 data include quarterly observations, with the dataset beginning in 2011:Q3. We limit the analysis to loans made between 2012:Q4 to 2023:Q3 because of potentially inconsistent reporting in the first year of the sample. 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. The data include the committed and outstanding amount of the loan as well as key characteristics such as the loan type and purpose, interest rate, base rate and interest rate spread, maturity, location, industry, and collateral. We track borrowers in the sample through their taxpayer identification number. Because there may be data-entry errors, we drop reported interest rates over 30 percent and dropping observations where interest rates are equal to zero. On average, there are approximately 3,750 originations per quarter, peaking in 2012 and with a minimum in 2020Q3.

Banks directly report measures of loan risk, including an internal loan risk rating, which captures the bank's estimated probability of default. We make use of a confidential Federal Reserve concordance table

7. Most studies on U.S. bank lending rely on data from syndicated loans, using either LPC DealScan data or the confidential Shared National Credit data, which cover only loans held by multiple banks, and mostly loans greater than $75 million. Some studies have also used Community Reinvestment Act (CRA) data, which capture loans smaller than $1 million for banks with more than $1 billion in assets and do not have data on interest rates. 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 to Compustat firms, DealScan, and SNC, highlighting the differences between the large borrowers in those 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.

that translates banks' internal loan risk ratings into a common rating scale where loans are rated from AAA to D that mimics the ratings scales used for syndicated loans. Some banks opt in to the "advanced approach" to assessing capital requirements whereby they make use of their own risk model approved by supervisors rather than being subject to the default capital requirements. For banks and loans subject to advanced approaches, our data also include information on a numeric probability of default, exposure at default and loss given default, each expressed as a percentage of the committed amount of the loan. 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. We exploit this information in some of our analyses.

5.2 Forming loan bins

Our theory requires us to be able to group firms so that all loan originations within a "bin" have identical risk characteristics so that the variation reflects switching costs and not differences in default risk. We take several steps to assure that this condition is met. We first reduce the sample to a more homogeneous set of loans by limiting firms to nonfinancial, for-profit borrowers. The Y-14 reporting schedule allows for 19 types of loans and 30 loan purposes. We focus on loans for corporate purposes and working capital, as loans made to fund acquisitions may have a different set of fixed costs depending on the scale and scope of an acquisition. We also drop project-financing and real-estate-related purposes, as well as other miscellaneous loan purposes.[8] Within the corporate purpose and working capital loans, we sort loan types into four loan categories, separating both revolving and term loans as well as loans for short-term or seasonal working capital from those for permanent working capital or general corporate purposes.

We also pool loans based on calendar periods that have similar overall interest rate environments, in case the price of aggregate risk varies according to whether interest rates are at the effective lower bound. Our first regime runs from 2012Q4 to December 16th 2015 when rates were near the lower bound. The second interval spans December 16 2015 to December 31, 2019 and includes a period of positive interest rates. We drop the first two quarters of 2020 to exclude the initial period of the pandemic. The third period starts on July 1, 2020 and continues until March 17, 2022 capturing the second interval where rates were again at the effective lower bound. The last period cover March 17, 2022 through 2023 Q3 when rates were non-negative. [The next draft of the paper will add data through 2024 that was not available when we began this project.]

8. . We drop real estate loans and other categories as these loans may be collateralized or have other characteristics that make it difficult to determine borrower risk from the standardized data we have in our sample.

Table 1: Loan Characteristics

	Share of obs (%)	Share of committed (%)	Mean interest rate (pp.)	Median interest rate (pp.)	Mean spread (pp.)	Median spread (pp.)	Mean utilization (%)	SD interest rate (pp.)
Rating:								
A	1.1	2.5	2.40	2.16	1.57	1.35	59.96	1.30
BBB	26.2	38.8	3.06	2.78	1.89	1.68	66.20	1.52
BB	56.8	47.2	3.75	3.50	2.54	2.40	69.54	1.57
B	15.9	11.5	4.35	4.11	3.15	3.00	69.62	1.70
Committed (millions):								
1-5	52.9	7.6	3.85	3.65	2.70	2.65	72.00	1.51
5-10	12.1	5.7	3.60	3.25	2.46	2.25	69.35	1.66
10-25	15.5	16.7	3.47	3.14	2.32	2.10	68.29	1.71
25-100	17.8	55.1	3.28	2.85	1.96	1.70	60.06	1.80
100-250	1.7	14.9	3.26	2.77	1.58	1.35	47.61	1.91
Maturity (years):								
0-1	31.8	10.9	3.84	3.52	2.76	2.80	61.72	1.57
1-4	32.3	35.2	3.59	3.25	2.40	2.25	67.80	1.73
4-5	28.6	51.1	3.49	3.24	2.22	1.95	71.75	1.66
5-8	5.9	2.7	3.62	3.51	2.32	2.10	88.06	1.33
8-15	1.3	0.2	4.08	4.05	2.23	2.10	95.24	1.13
>15	0.0	0.0	4.60	4.66	2.19	2.20	94.50	0.88
Loan type:								
Revolver, short-term capital	28.9	16.1	3.99	3.75	2.67	2.65	60.57	1.69
Revolver, general working capital	42.9	65.6	3.46	3.24	2.40	2.20	57.11	1.62
Term, short-term capital	6.8	2.4	3.27	3.15	2.20	1.93	96.83	1.37
Term, general working capital	21.5	15.9	3.70	3.50	2.36	2.15	93.28	1.59

Notes: This table reports summary statistics for the initial sample of 148,467 loan originations for 35 stress-tested banks from 2012:Q4 to 2023:Q3 (dropping 2020Q1 and 2020Q2). Bins are constructued by pooling loans within 4 separate interest rate regimes (described in the text) based on banks' internal risk rating, commitment amount, maturity, loan type, and loan purpose. We drop the revolver loans that were never drawn on over the life of the loan, bins with fewer than 6 unique banks or less than 20 overall loans. The first two columns show the share of originations by number and committed value. The next columns show the mean and median interest rate and interest rate spread, the mean amount of the commitment utilized and the standard deviation of the interest rate for all loan originations in each category calculated as of origination date. Interest rate spreads are calculated based on spreads to 3-month LIBOR and use forward curves at origination for fixed-rate loans. Mean utilization is the mean of the loans in the category, defined as utilized divided by committed, calculated for each loan over the loan's lifespan in the dataset.

Table 1 presents loan level sample statistics. The table reveals several stylized facts about lending. First, loans rated above A are rare which means that we typically do not observe enough originations in a bin to obtain switching-cost estimates. Therefore, we drop these loans from the analysis. Most of the usuable sample of loans are rated BB or BBB. The median loan spread, which we define as the difference between the interest rate on the loan and LIBOR, monotonically rises in loan risk. This variation motivates our bin structure, grouping firms with different risk ratings into different bins. Third, spreads fall with the amount committed and maturity, which probably reflects that all else equal, long-term borrowers and larger firms tend to be those with lower default risks. The fact that our theory suggests that loan costs should vary with loan size and maturity, along with this empirical variation, suggests that we should also allow bins to vary by size and maturity. Fourth, we see that revolvers, which account for about 72 percent of all loans and 82 percent of the value of all lending, have systematically higher spreads than term loans, which motivates our decision to divide these two types of loans into four separate categories: revolvers for

short-term capital, revolvers for general working capital, term loans for short-term capital, and term loans for general working capital. Since there are many other heterogeneous types of revolvers and term loans contained in the "other" category, we drop these types of loans in order to be conservative about what loan types we deem to be homogeneous. Finally, we also drop loans with ratings below B because we are concerned that banks are likely to assign different default probabilities to high-risk borrowers, which means that interest rate differentials within this bin are likely to reflect different risk assessments.

We also see the distribution of loans is skewed towards small borrowers. Although 52.9 percent of loans are for amounts between $1 million and $5 million, small borrowers account for less than 7.6% of the total amount committed by the banks. As a result, the median committed amount is just over $5 million, with a mean of just below $20 million. Conversely, loans of $25 million or more account for about 18 percent of the observations but represent 55 percent of the dollars lent. Just over 92% of all loans have a maturity of five years or less, and 7 percent of the sample are loans of more than five years.

Table 2: Average Loan Characteristics Across Bins

	All Bins	A	BBB	BB	B	1-5	5-10	10-25	25-100	100-250
min(interest rate)	0.72	0.58	0.56	0.71	0.99	0.85	0.75	0.69	0.52	0.58
	(0.43)	(0.33)	(0.33)	(0.39)	(0.49)	(0.44)	(0.38)	(0.43)	(0.35)	(0.34)
max(interest rate)	6.27	3.75	4.88	6.85	7.73	6.79	6.18	6.42	5.65	4.36
	(2.10)	(1.21)	(1.31)	(1.90)	(1.89)	(2.10)	(1.88)	(2.21)	(2.00)	(1.11)
number of banks	13.05	8.87	12.56	14.15	12.64	13.27	12.87	13.46	12.80	10.93
	(4.82)	(2.57)	(4.34)	(5.19)	(4.55)	(5.12)	(4.54)	(4.81)	(4.69)	(3.44)
number of loans	215.80	55.27	172.21	308.67	148.58	314.31	133.95	165.96	191.24	91.85
	(375.65)	(53.02)	(239.89)	(513.66)	(223.23)	(540.86)	(149.01)	(193.79)	(291.24)	(83.29)
interest rate	3.71	2.48	3.11	3.86	4.54	3.92	3.77	3.62	3.44	3.32
	(1.36)	(1.11)	(1.19)	(1.23)	(1.27)	(1.17)	(1.37)	(1.40)	(1.52)	(1.57)
maturity (years)	3.69	3.32	3.63	3.90	3.48	4.00	3.34	3.50	3.56	4.15
	(2.32)	(1.75)	(2.15)	(2.59)	(2.14)	(3.01)	(1.96)	(1.81)	(1.69)	(1.06)
loan size (millions)	21.24	40.91	24.84	18.82	16.56	2.36	7.56	17.18	48.58	145.13
	(30.51)	(44.02)	(34.38)	(27.77)	(23.57)	(0.36)	(0.35)	(0.92)	(5.45)	(6.37)

Notes: This table presents summary statistics for the estimation sample of 148,467 loan originations (the same sample as Table 1) for which we can compose bins of loans with similar characteristics. Numbers represent the mean across 688 bins for each variable. The minimum (maximum) interest rate is therefore an average of the minimum (maximum) interest rate across 688 bins. For observered characteristics we first calculclate the average of a variable within a bin and then average the means accross bins within the same matrurity (in years) range. Observable variables reported in the table are averages across bins by maturity (in years) for the following variables the lowest interest rate, highest interest rate, mean interest rate, number of banks, number of loan originations, mean maturity (years), and mean committed amount (millions). Column (1) shows the cross bin averages across all bins, columns (2-5) show the cross bin averages based on bin rating, and columns (6-10) show the cross bin averages for committed amount in millions.

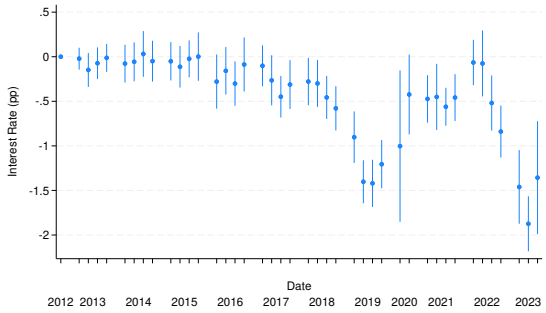
Table 2 shows some bin level statistics. In forming bins, we exclude any bins that would have fewer than 20 loans or fewer than 6 banks. The bins group loans based credit rating, amount committed, the four

purposes mentioned above, and maturity. The first column shows the typical bin has 215 loans that are originated by 13 banks. The average loan size at the bin level is $21 million. Data that will be presented later, shows that even within these narrowly defined bins there is considerable heterogeneity in interest rates. The next four columns show the bin characteristics when the bins are compared based on their rating. A rated loans are rare, there are only 30 bins that have enough loans to group the borrowers and the average number of loans in those bins is 55. There are at least 150 bins for the other three ratings. As would be expected, the average interest rate rises as ratings fall. The average loan size also decreases as ratings fall.

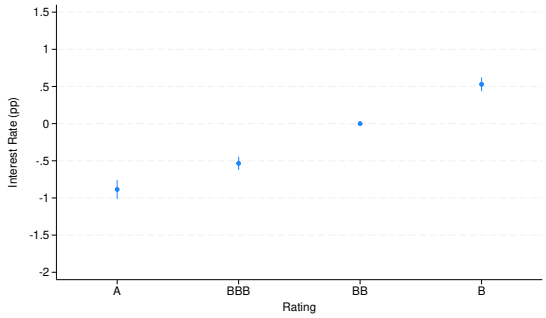
The final five columns in the table show the characterstics for the bins when they are grouped by the size of the committed amount of the loan. The smallest category, loans from commited amounts between $1 and $5 million, has the most bins (250) and those bins tend to have over 300 loans in them. There are only 27 bins for the loans that involve commitments of more than $100 million. The average interest rate falls as the loan size increases.

Figure 1: Impact of Loan Characteristics on Interest Rates (Adjusted by the Federal Funds Rate (FFR))

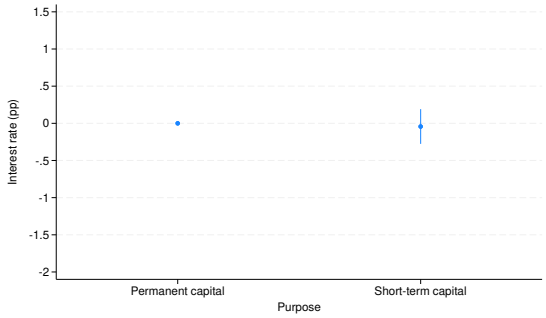
(a) Loan Origination Date Fixed Effect



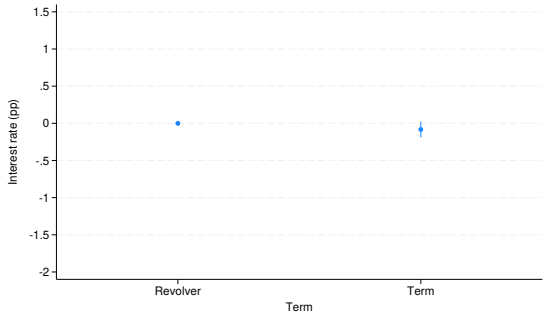
(b) Loan Rating Fixed Effects



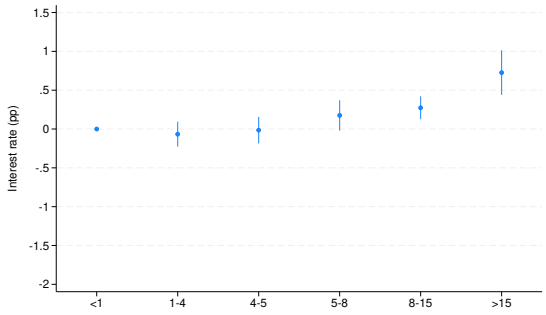
(c) Loan Purpose Fixed Effects



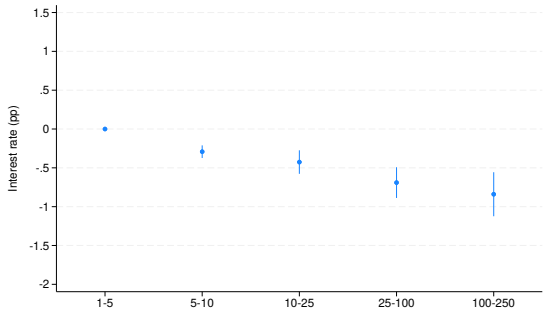
(d) Term Loan vs Revolver Fixed Effects



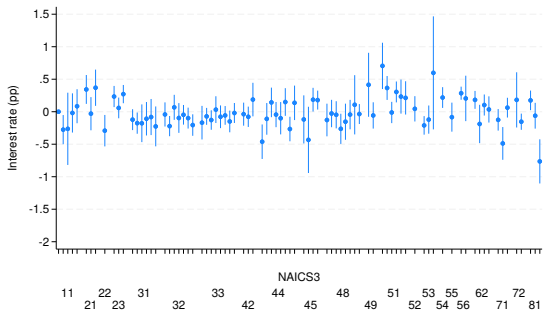
(e) Maturity Fixed Effect (years)



(f) Committed Loan Size Fixed Effect (millions)



(g) NAICS Fixed Effect



Notes: These charts show the coecients of a regression of interest rates minus the federal funds rate (at the origination date of the loan) on various controls. Controls include binary variables equal to one for categories of internal risk rating, the dollar amount of loan commitment, maturity in years, loan type, loan purpose, quarter of origination and North American Industrial Classification System (NAICS) 3-digit industry. The sample is the estimation sample of 148,127 originations (the same sample as Table 1) from 2012:Q4 to 2023:Q3. Blue bars represent 95% confidence intervals.

Turning to loan pricing, Figure 1 plots the estimated coefficients and 5/95 confidence levels from a multivariate regression of interest rates minus the federal funds rate on rating, committed amount, maturity, loan type, loan purpose, date, and industry fixed effects. The R^2 in this regression is , indicating that these variables explain most of variation in interest rates. The time fixed effects pick up any residual calendar variation that is not picked up by the federal funds rate. It is apparent that in the first three years of the sample, covering the first interest rate regime, the calendar effects are estimated relatively precisely to be zero. The key patterns that we report below all hold if we limit the sample to only that period. Based on that finding we believe changing levels of interest rates are not responsible for our main results. In these multivariate specifications, which include controls for all the characteristics shown in the various panels, the industry of the borrower does not significantly affect interest rates, which motivates our decision not to use industry as a characteristic when grouping borrowers into bins.

The pricing data is consistent with the screening and monitoring hypothesis. In particular, if we believe assessing creditworthiness is harder for smaller loans (which are most likely granted to smaller borrowers), revolvers and loans to low-rated borrowers, then we would expect to find higher spreads and loan rates for these customers. The data are consistent with Y-14 data used in other studies, such as Chodorow-Reich et al. (2022) that focus on small borrowers, as well as pricing with patterns identified in the empirical literature on syndicated loans and relationship banking. Overall, this regression establishes controlling for loan rating, amount committed, maturity, and loan type captures most of the important cross-sectional variation in interest rates.

Our estimation procedure depends crucially on bins having identical risk characteristics so that within-bin variation in interest rates is not dependent on risk. We test for this possibility by regressing the interest rate charged for each loan origination on bin fixed effects and the banks' expected loss ratio to see if interest rate differentials within bins are associated with higher risk.

Table 3: Interest Rates and Expected Loan Losses

	(1) Interest Rate	(2) Interest Rate	(3) Interest Rate	(4) Interest Rate Standardized
Expected loss ratio			0.106***	
			(0.006)	
Standardized Expected loss ratio				0.0449***
				(0.003)
Bin FE	Yes	Yes	Yes	Yes
Observations	148467	80610	80610	80610
R-Squared	0.509	0.516	0.518	0.518
RMSE	1.149	1.205	1.203	0.697

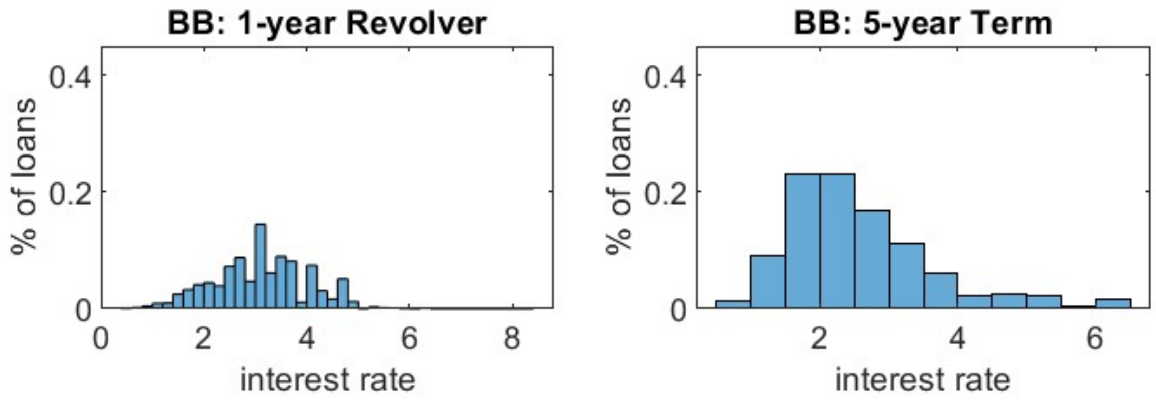
Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Results from regressing interest rates on an indicator variable for each loan bin where bins are based on monetary policy cycle, rating, maturity, size, type, and purpose. The first column estimates the regression for the estimation sample of 148,467 observations (the same sample as Table 1), while Column 2 limits the specication to the sample of 80,610 observations for which the expected loss ratio data is available. The third column adds the expected loss ratio (denoted as the probability of default times loss given default times the exposure given default, divided by the committed size of the loan) as an explanatory variable and the fourth column standardizes both the expected loss ratio and the interest rate to mean 0 and standard deviation 1. The median interest rate and expected loss ratio for column 3 is 3.75 percentage points and 0.19 percentage points, respectively. The standard deviation interest rate and expected loss ratio for column 3 is 1.73 percentage points and 0.73 percentage points, respectively.

We report the results in Table 3 We find that while there is a significant association between expected loan losses and interest rates, the economic significance of the impact of expected loan losses on interest rates is trivial. Adding expected loan losses to the regression only increases the adjusted R^2 by 0.002, which indicates that expected loan losses do not explain much of the variation in interest rates within bins. Moreover, the coefficient on expected loan losses is small—a one percentage point increase in expected loan losses as a share of the amount committed, which would correspond to a loan in the highest percentile of within-bin risk, is only associated with only a 4-basis point increase in interest rates. (Recall the median interest rate is 3.75%). Thus, the observed dispersion in interest rates within bins is not associated with differences in loan risk in an economically meaningful way.

Besides this direct evidence, it is helpful to recall that the dataset arises from an information collection mandated by supervisors to estimate default risk. The supervisors who developed the loan collection data requirements mandated the inclusion of the crucial variables associated with loan risk. To the extent that other variables related to loan risk are not included, it is unlikely that banks collect these variables systematically or use them to set loan pricing. This provides us with further comfort that the bins are suitably homogeneous.

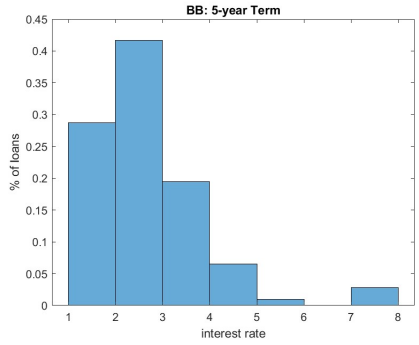
Figure 2: Interest Rate for Loans with a BB Rating Under Regime 1



Notes: The above figure represents two histograms for the interest rates of two bins that have a BB rating in regime 1 (2012Q4 - Dec, 16th, 2015). The BB: 1-year revolver with 4766 loan originations, has mean interest rate of 3.3479 a spread of 7.7125 (r_max - r_min) with a mean maturity of 0.84528 years, mean committed of 2.2068 million, and is a revolver, permanent capital. There are 20 banks represented in the sample. The BB: 5-year term with 375 loan originations, has mean interest rate of 2.8403 a spread of 5.769 (r_max - r_min) with a mean maturity of 4.902 years, mean committed of 7.5274 million, and is a term, permanent capital. There are 20 banks represented in the sample.

Figure 2 shows a pair of indcative bins. The picture shows three notable features that drive the econometric results. First, and most importantly, there is simply a great deal of variation in the interest rates that borrowers accept. The right tail of the distribution is important because in our framework the reason why these firms would settle for the realively high rates is because they do not expect to get a better deal if they searched. The left tail is also important. The mass of borrowers with very low rates determines how many firms have very low search costs. Through the lens of our model, the firms that are able to find the best rates are the ones with the lowest search costs. Whether this is a high or low percentage of firms with these relatively low rates identifies the set of firms that search intensively. Finally, the rest of distribution indicates whether the remaining firms search just once or more than that. For both of the distributions shown in the figure, the returns from getting more than two rate quotes is not very big unless a firm's search costs are very low.

Figure 3: Interest Rate for Loans with a BB Rating (5-year Term) Under Regime 3



Notes: The above figure represents a histograms for the interest rates of a bins that has a BB rating in regime 3 (Jul, 1st 2020 - Mar, 17th, 2022). The BB: 5-year term with 108 loan originations, has mean interest rate of 2.8686 a spread of 6.34 (r_max - r_min) with a mean maturity of 4.927 years, mean committed of 7.6231 million, and is a term, permanent capital. There are 19 banks represented in the sample.

Figure 3 shows a bin that has the kind of distribution where searching more than a second time is appealing even for firms with modest search costs. This distribution has a large mass at interest rates that are not extraordinarily low. This kind of distribution will only arise in our framework if firms assume that there are many firms that search more than twice.

6 Switching Costs

Table 4 reports the estimates when all bins are treatly equally. The search cost estimates are the minimum number that is needed to rationalize the data. So the 72 basis point estimate for the firms that search once is a lower bound, meaning that 35% of the borrowers are estimated to have switching costs that are at least that large. Since the averge bin has a mean loan size of $21.24 million, this estimate implies that the typical firm would have swtiching costs of at least $152,928. We estimate that another 49% of firms get one additional quote. The associated search cost estimate of 30 basis points implies that firms are behaving as if the switching cost for a $21.24 million loan is at least $63,720.

To estimate the model we must make an assumption about how many banks that a borrower could possibly consider. We set that parameter equal to the number of banks that are observed to make a loan in each bin. The last row in the table shows that for the average bin, 14% of the borrowers are found to have very low switching costs. being under $6,400. As indicated above, the size of this set of borrowers is determined by the number of borrowers that find the lowest rates in the bin.

Table 4: Number of Searches and Average Search Costs Across Bins

	share	search cost
number of searches=1	0.35	0.72
	(0.14)	(0.24)
number of searches=2	0.49	0.30
	(0.11)	(0.09)
number of searches=3	0.00	0.17
	(0.03)	(0.06)
number of searches=4	0.00	0.11
	(0.00)	(0.04)
number of searches=n-1	0.00	0.04
	(0.02)	(0.03)
number of searches=n	0.14	0.03
	(0.07)	(0.02)

Notes: This table presents summary statistics for the estimation sample of 148,467 loan originations (the same sample as Table 1) for which we can compose bins of loans with similar characteristics. Numbers represent the mean across 688 bins for each variable. In the search model interest rates are adjusted by subtracting the upper Federal Funds Rate at the origination date of the loan.

The next two tables are intended to explore whether the model estimates are consistent with the screening and monitoring hypothesis. That hypothesis implies that borrowers that require more ex-ante effort to judge their creditworthiness or loans that require more ex-post monitoring should come with higher rates and higher switching costs. As empirical proxies for borrowers that would require more screening and monitoring we use the credit rating and loan size.

Table 5: Number of Searches and Average Search Costs Across Bins by Rating

Rating	A	BBB	BB	B	A	BBB	BB	B
	share	share	share	share	search cost	search cost	search cost	search cost
number of searches=1	0.37	0.34	0.34	0.37	0.43	0.55	0.78	0.89
	(0.20)	(0.13)	(0.14)	(0.13)	(0.15)	(0.15)	(0.19)	(0.23)
number of searches=2	0.44	0.50	0.50	0.47	0.18	0.23	0.32	0.38
	(0.15)	(0.11)	(0.11)	(0.10)	(0.07)	(0.06)	(0.07)	(0.09)
number of searches=3	0.01	0.00	0.00	0.00	0.10	0.13	0.18	0.22
	(0.07)	(0.04)	(0.02)	(0.03)	(0.04)	(0.04)	(0.04)	(0.05)
number of searches=4	0.00	0.00	0.00	0.00	0.07	0.09	0.12	0.15
	(0.00)	(0.00)	(0.00)	(0.00)	(0.03)	(0.03)	(0.03)	(0.04)
number of searches=n-1	0.00	0.00	0.00	0.00	0.04	0.03	0.03	0.05
	(0.01)	(0.01)	(0.03)	(0.02)	(0.03)	(0.02)	(0.03)	(0.04)
number of searches=n	0.17	0.14	0.14	0.14	0.03	0.02	0.03	0.04
	(0.11)	(0.06)	(0.07)	(0.07)	(0.02)	(0.01)	(0.02)	(0.03)

Notes: This table presents summary statistics for the estimation sample of 148,467 loan originations (the same sample as Table 1) for which we can compose bins of loans with similar characteristics. Numbers represent the mean across N bins for each variable (N for A ratings = 30, BBB ratings = 226, BB ratings = 273, and B ratings = 159). In the search model interest rates are adjusted by subtracting the upper Federal Funds Rate at the origination date of the loan.

Table 5shows the results when we group bins according to their credit rating. The estimated search costs follow the same pattern as was found when the estimation was done with all bins. Specifically, for each rating, there are about of one third of the borrowers are estimated to accept the first rate they are presented with, while between 44 and 50 percent of the borrowers seek a second loan quote. The remaining 14 to 17 percent of the boorwers shop heavily and find the lowest rates available.

The estimated size of the search costs align with the predictions of the screening and monitoring hypothesis. The A rate borrowers have switching costs that are about half the size of the costs of the B rated borrowers. For instance, the A rated firms that get a second quote act as if their switching costs were at least 18 basis points, so that for a $10 mllion loan they would have costs of at least $18,000. The B rated borrowers are estimated to have costs of at least 38 basis points, implying a switching cost of at least $38,000 for the same $10 million loan. the BBB and BB borrowers have estimated costs that lie in between these estimates.

Table 6: Number of Searches and Average Search Costs Across Bins by Committed

Committed (millions)	1-5	5-10	10-25	25-100	100-250	1-5	5-10	10-25	25-100	100-250
	share	share	share	share	share	search cost	search cost	search cost	search cost	search cost
number of searches=1	0.37	0.36	0.33	0.33	0.33	0.77	0.71	0.72	0.65	0.49
	(0.14)	(0.14)	(0.13)	(0.13)	(0.18)	(0.24)	(0.22)	(0.24)	(0.22)	(0.15)
number of searches=2	0.47	0.49	0.51	0.51	0.51	0.33	0.30	0.29	0.27	0.20
	(0.11)	(0.11)	(0.11)	(0.10)	(0.13)	(0.09)	(0.09)	(0.09)	(0.09)	(0.08)
number of searches=3	0.00	0.00	0.00	0.00	0.01	0.19	0.17	0.17	0.15	0.11
	(0.03)	(0.03)	(0.04)	(0.03)	(0.03)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
number of searches=4	0.00	0.00	0.00	0.00	0.00	0.13	0.11	0.11	0.10	0.07
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
number of searches=n-1	0.00	0.00	0.00	0.00	0.00	0.04	0.04	0.03	0.03	0.03
	(0.00)	(0.03)	(0.02)	(0.03)	(0.00)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)
number of searches=n	0.15	0.13	0.13	0.15	0.15	0.03	0.03	0.03 0.03	0.02	
	(0.08)	(0.07)	(0.06)	(0.07)	(0.08)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)

Notes: This table presents summary statistics for the estimation sample of 148,467 loan originations (the same sample as Table 1) for which we can compose bins of loans with similar characteristics. Numbers represent the mean across N bins for each variable (N for 1-5 committed (millions) = 250, 5-10 committed (millions) = 134, 10-25 committed (millions) = 139, 25-100 committed (millions) = 138, and 100-250 committed (millions) = 27). In the search model interest rates are adjusted by subtracting the upper Federal Funds Rate at the origination date of the loan.

Table 6shows the estimates when the model is fit separately for the five size buckets. The estimated costs for the smallest loans, for commitments beetween $1 and $5 million are the highest, though the difference relative to the loans under $100 milllion are modest. However, the estimates are substantially higher than for the largest category of loans that covers the $100 million to $250 million commitments. For instance, among the borrowers that are assumed to get two loan quotes, the borrowers with the largest loans are estimated to have search costs of at least 20 basis points, while the ones with the smallest loans have a cost of 33 basis points. Of course, given the differences in the size of the loans, the absolute size of the costs would be much higher for the borrowers that take out the larger loans.As another way to view the model, we also run regressions that use the estimated switching cost, the deltas, on bin characteristics.

Table 7: Impact of Rating on Search Costs (Δ_1)

	(1) All	(2) Term	(3) Revolver	(4) Fixed	(5) Float
BBB	0.123***	0.186*	0.116**	0.447*	0.0991**
	(0.035)	(0.072)	(0.038)	(0.226)	(0.034)
BB	0.342***	0.443***	0.303***	0.612**	0.334***
	(0.035)	(0.072)	(0.038)	(0.225)	(0.033)
B	0.448***	0.623***	0.365***	0.724**	0.435***
	(0.036)	(0.074)	(0.039)	(0.227)	(0.034)
N	688	297	391	151	537
R^2	0.433	0.491	0.433	0.331	0.495

Standard errors in parentheses

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

Notes: Regression based on estimation sample in Table 1. Dependent variable is the search cost delta_1. The search cost delta_1 is calculated at the bin-level. The sample in the first column is all loans; in the second column is term loans; and the third is revolver loans; the fourth column are fixed loans; and the fith column are float loans. Standard errors are in parenthesis.

Table 7 shows the coefficients on the ratings (controlling for the loan maturity and sizes). In the first column, when all 688 bins are used in the estimation, the cost for the borrowers that accept the only quote that they get, show rising search costs as credit ratings fall. The B rated borrowers have about 45 basis points higher costs than the A rated ones (who are the reference group for the regression). The gradient is even steeper for the term loans that afford the borrowers more discretion in how the funds could be used. In this case, the B rated firms have switching costs that are 62 basis points higher than A rated borrowers.

Table 8: Impact of Rating on Search Costs (Δ_{avg})

	(1) All	(2) Term	(3) Revolver	(4) Fixed	(5) Float
BBB	0.328***	0.306**	0.262***	0.831**	0.281***
	(0.063)	(0.099)	(0.061)	(0.272)	(0.067)
BB	0.588***	0.652***	0.461***	0.940***	0.580***
	(0.063)	(0.098)	(0.060)	(0.271)	(0.066)
B	0.675***	0.854***	0.546***	1.005***	0.674***
	(0.065)	(0.101)	(0.062)	(0.273)	(0.068)
N	688	297	391	151	537
R^2	0.919	0.952	0.943	0.927	0.918

Standard errors in parentheses

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

Notes: Regression based on estimation sample in Table 1. Dependent variable is the search cost delta_all, which is calculated using equation in paper. The search cost delta_all is calculated at the bin-level. The sample in the first column is all loans; in the second column is term loans; and the third is revolver loans; the fourth column are fixed loans; and the fith column are float loans. Standard errors are in parenthesis.

Table 8 shows analagous estimates when the average search cost for each bin is used instead of the search cost for the borrowers that do no search. The patterns are again consistent with the screening and monitoring hypothesis in that the lower rated firms have higher estimated switching costs and the term

loans (for a given rated borrower) switching costs are higher than floating rate loans.

Because there are not previous estimates of switching costs, there is not an established benchmark against which our estimates can be compared. One way to judge them is to use them to deduce the deadweight loss that would be created if a bank were to fail. In such an event, the bank's customers would need to find a new lender. In this thought experiment, we suppose each borrower incurs the estimated mean cost of searching based on the estimates for that bank. For example, for each customer of a given bank, we know the bins in which their borrowers are located and we have an estimated costs for the borrowers in each bin. We assign each of the borrowers the mean cost for each bin and multiply that cost by the amount of the commitment (outstanding in 2023Q3) for that the borrower. We then add up all the costs and divide it by the total size of the commitments to arrive at a percentage of the loans that would be spent by the customers. Given that our estimates represent the lower bound of switching costs for each of the types of the borrowers, this calculation yields a lower bound on the switching costs. Also, because the customer mix of banks differ, each bank would have a different estimated cost. The loans that are assessed are those where the borrower renewed the loan at its existing lender.

Table 9: Bank Failure Counterfactual and Distribution of Switching Costs

	median	mean	p5	p95	standard deviation
TC (millions)	1536.69	2009.59	164.16	4865.27	1584.55
TL (millions)	19079.18	43220.69	389.11	140996.93	73687.17
SC	0.11	0.15	0.02	0.43	0.13

Notes: Results from an equation in the paper. The sample is made up of 85,474 loans. The sample is only made up of firms who renewed their loans in the sample of Table 1. TC (calculated by an equation in the paper) is short for total cost in millions, this is calculated for each bank in the sample. TL is short for total lending in millions, this is calculated for each bank by summing its loans over the sample. There are 35 unique banks represented in the sample.

Table 9 shows the distribution of the switching costs across the 35 banks. For the median bank, the costs are estimated to be 11 percent of the loan amounts. There is substantial heterogeneity across the banks, with the one at the 5th percentile of distribution having costs of only 2 percent of loans and the bank at the 95th percentile having costs of 43%.

7 Robustness and generalizability

The results depend critically on the loans within a single bin being of similar risk. We resticted the sample in various ways to see whether eliminating variation in different ways leads to different inferences about the switching costs. For example, one alternative that we considered was to study only fixed rate loans so that any variaton that might be due to reference rate would be removed. In another alternative, we limited the sample for include only loans originated between 2012 Q3 and December of 2015. During that whole period the Fed's policy rate was at its effective lower bound. In both these more limited samples, the estimated switching costs were quite similar to the ones derived from the full sample.

8 Conclusion

The cost for a borrower to switch banks is a central determinant of whether a contraction in lending translates into effects on the real economy. Without estimates of these costs, one cannot judge how various events like bank failures, bank mergers, regulatory changes, and monetary policy-induced changes in bank funding costs will play out. We believe this is the first paper to directly estimate these costs by combining a new methodology with comprehensive data on loan pricing and borrower risk.

While our estimates of the switching wedges are calculated for banks and their business customers, the patterns suggest they are tied to screening and monitoring costs. If this is correct, then non-bank lenders would likely face the same considerations in extending credit as banks. If this conjecture could be corroborated, it suggests that there could be externalities from non-bank failures in terms of credit provision.

Given that these are the first estimates of switching wedges in the literature, we have identified several areas for future research. First is to consider the debates over market power in banking by extracting markups from the estimates and examining if the estimated wedges differ depending on the degree of competition. Second is to measure the impact on loan pricing and credit availability after the onset of COVID. Understanding fixed costs offers a method to calculate the likely impact on borrowers and their lenders for loan pricing and credit availability following these shocks. Finally, a key subject of interest is understanding the extent of business cycle variation in the fswitching costs and the ability of firms to replace lost loans. One suspects that when macroeconomic uncertainty is higher, monitoring is more costly and screening is more challenging, in which case wedges would be expected to rise.

References

Amiti, Mary, and David E. Weinstein. 2011. "Exports and Financial Shocks." *The Quarterly Journal of Economics* 126 (4): 1841–1877.

———. 2018. "How Much Do Idiosyncratic Bank Shocks Affect Investment? Evidence from Matched Bank-Firm Loan Data." *Journal of Political Economy* 126 (2): 525–587.

Berger, Allen, Anthony Saunders, Joseph M. Scalise, and Gregory Udell. 1998. "The Effects of Bank Mergers and Acquisitions on Small Business Lending." *Journal of Financial Economics* 50 (2): 187–229.

Bharath, Sreedhar T., Sandeep Dahiya, Anthony Saunders, and Anand Srinivasan. 2011. "Lending Relationships and Loan Contract Terms." *The Review of Financial Studies* 24 (4): 1141–1203.

Billett, Matthew T., Mark J. Flannery, and Jon A. Garfinkel. 1995. "The Effect of Lender Identity on a Borrowing Firm's Equity Return." *The Journal of Finance* 50 (2): 699–718.

Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–969. ISSN: 00129682, 14680262, accessed March 19, 2025.

Chodorow-Reich, Gabriel. 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, Gabriel, Olivier Darmouni, Stephan Luck, and Matthew Ploser. 2022. "Bank liquidity provision across the firm size distribution." *Journal of Financial Economics* 144 (3): 908–932.

Cortés, Kristle R., Yuliya Demyanyk, Li Lei, Elena Loutskina, and Philip E. Strahan. 2020. "Stress tests and small business lending." *Journal of Financial Economics* 136 (1): 260–279.

Diamond, Douglas W. 1984. "Financial Intermediation and Delegated Monitoring." *The Review of Economic Studies* 51 (3): 393–414.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl. 2017. "The Deposits Channel of Monetary Policy." *Quarterly Journal of Economics* 132 (4): 1819–1876.

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

Honka, Elisabeth, Ali Hortaçsu, and Matthijs Wildenbeest. 2019. "Chapter 4 - Empirical search and consideration sets." In *Handbook of the Economics of Marketing, Volume 1,* edited by Jean-Pierre Dubé and Peter E. Rossi, 1:193–257. Handbook of the Economics of Marketing. North-Holland.

Hubbard, Robert, Kenneth Kuttner, and Darius N. Palia. 2002. "Are There Bank Effects in Borrowers' Costs of Funds? Evidence from a Matched Sample of Borrowers and Banks." *The Journal of Business* 75 (4): 559–581.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina. 2012. "Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications." *American Economic Review* 102 (5): 2301–2326.

———. 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.

Karceski, Jason, Steven Ongena, and David C. Smith. 2005. "The Impact of Bank Consolidation on Commercial Borrower Welfare." *The Journal of Finance* 60 (4): 2043–2082.

Kashyap, Anil K., and Jeremy C. Stein. 2000. "What Do a Million Observations on Banks Say about the Transmission of Monetary Policy." *American Economic Review* 90 (3): 407–428.

Moraga-González, José Luis, and Matthijs R. Wildenbeest. 2008. "Maximum likelihood estimation of search costs." *European Economic Review* 52, no. 5 (July): 820–848.

Morgan, Peter, and Richard Manning. 1985. "Optimal Search." *Econometrica* 53 (4): 923–944. ISSN: 00129682, 14680262, accessed March 19, 2025.

Peek, Joe, and Eric 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.

Pyle, David, and Hayne E. Leland. 1977. "Informational Asymmetries, Financial Structure, and Financial Intermediation." *The Journal of Finance* 32 (2): 371–387.

Slovin, Myron B., Marie E. Sushka, and John A. Poloncheck. 1993. "The Value of Bank Durability: Borrowers as Bank Stakeholders." *The Journal of Finance* 48 (1): 247–266.

Warner, Jerold B. 1977. "Bankruptcy Costs: Some Evidence." *The Journal of Finance* 32 (2): 337–347.

Weiss, Lawrence A. 1990. "Bankruptcy resolution: Direct costs and violation of priority of claims." *Journal of Financial Economics* 27 (2): 285–314.