

Informational Frictions and the Credit Crunch

OLIVIER DARMOUNI*

ABSTRACT

In this paper, I estimate the magnitude of an informational friction limiting credit reallocation to firms during the 2007-2009 financial crisis. Because lenders rely on private information when deciding which relationship to end, borrowers looking for a new lender are adversely selected. I show how to separately identify private information from information common to all lenders but unobservable to the econometrician by using bank shocks within a discrete choice model of relationships. Quantitatively, these informational frictions appears to be too small to explain the credit crunch in the U.S. syndicated corporate loan market.

Keywords: Informational frictions, Aggregate effects of credit supply shocks, Banking relationships.

JEL codes: G21, D82, L14

*Columbia Business School. Email: omd2109@columbia.edu. I am indebted to David Sraer, Jakub Kastl, Markus Brunnermeier, Motohiro Yogo, and Valentin Haddad for their continuous support. I would also like to thank Tobias Berg, Charles Calomiris, Adrien Matray, Atif Mian, Justin Murfin, Hoai-Luu Nguyen, Tomek Piskorski, Matthew Plosser, multiple reviewers, and seminar participants at Princeton University, Columbia, Harvard Business School, UT Austin, NYU Stern, Yale SOM, Wharton, Berkeley Haas, FDIC, HEC Paris, the Toulouse School of Economics, Wharton, Kellogg, the New York Fed, WFA, Copenhagen Business School, BYU, and UCLA Anderson for many discussions and comments that improved this paper. Special thanks to Gabriel Chodorow-Reich for his help and support at the early stages of this project. Previous versions of this paper were titled "The Effects of Informational Frictions on Credit Reallocation" and "Estimating Informational Frictions in Sticky Relationships." Lira Mota provided outstanding research assistance. I declare that I have no relevant or material financial interests that relate to the research described in this paper.

The defining feature of a lending relationship between a bank and a borrower is its stickiness: switching lenders is rare and costly.¹ In turn, credit markets are more vulnerable: a shock forcing a particular bank to cut lending can have *aggregate* effects if affected borrowers cannot easily find a new lender.² Understanding why relationships are sticky is important, as it can guide the design of institutions or policies to prevent breakdowns in lending markets.

In this paper, I estimate the effects of a key friction behind relationship stickiness: the *information gap* between a borrower's existing lender and its potential new lenders. Over the course of a relationship, lenders acquire abstract and hard-to-verify private ("soft") information about their borrowers that is unobservable to other lenders.³ The information gap represents the informational advantage that stems from relationship lending. The main contribution of this paper is to provide the first direct estimate of the magnitude of the information gap and its role in explaining the credit crunch that followed the 2007 to 2009 crisis (Chodorow-Reich, 2014).

The key identification challenge is that, empirically, *private* information is difficult to disentangle from *common* information that all lenders can observe but that the econometrician cannot. This paper shows that shocks to banks can be used to separately identify lenders' private information from information common to all lenders. Using loan-level data from the U.S. syndicated loan market, I find that lenders' private information appears to be too small to explain why relationships are sticky in this market, and therefore, cannot quantitatively account for much of the associated drop in lending documented in prior studies.⁴

The information gap reduces aggregate lending by creating adverse selection in the market for borrowers looking for a new relationship. Lenders' private information gives them the ability to selectively choose which relationships to end when scaling down lending after a shock, leaving their worst borrowers looking for funds elsewhere. This is the predominant view to rationalize relationship stickiness and the credit crunch observed the U.S. syndicated loan market. However, testing this private information channel directly has proven elusive. This paper offers a solution to this econometric challenge.

The above channel makes clear why shocks to banks' ability to lend can be useful in identifying the information friction. It implies an "inference hypothesis": borrowers leaving the most-affected

¹See Srinivasan et al. (2014) for a survey of the extensive literature on banking relationships.

²A large body of works studies the effect of bank shocks on firm borrowing and real outcomes after the financial crisis, including Ivashina and Scharfstein (2010), Chodorow-Reich (2014), Jiménez et al. (2019), Greenstone et al. (2020), and Schwert (2018).

³See for example Sharpe (1990), Rajan (1992), and Detragiache et al. (2000). Examples of soft information acquired during a relationship include the quality of management, potential future investment projects, as well as information whose public disclosure would hurt the firm.

⁴This is not to say that there is no asymmetric information between lenders and borrowers, but rather that there is no asymmetric information *across* lenders in this particular market.

lenders are less adversely selected, as described in Dell’Ariccia and Marquez (2004). Intuitively, these lenders cannot continue lending even to relatively good borrowers. Therefore, this inference hypothesis implies that a firm’s ability to borrow from a new lender after a breakup depends on the size of the shock faced by its previous lender.

There is evidence consistent with this effect in the U.S. syndicated loan market over the period 2004 to 2010. Exploiting the financial crisis that originated in the real estate sector, I use a lender’s exposure to this shock to measure its ability to lend in the corporate loan market. Conditional on leaving a relationship, a one standard deviation increase in the crisis exposure of a firm’s existing lender implies a 20% *increase* in the probability of borrowing from a new lender.⁵

However, this evidence does not solve the main identification challenge of isolating private information. In fact, the same reduced-form correlation would emerge if there were only common information that all lenders could observe but that the econometrician could not. In that case, new lenders would not learn any additional information from a relationship being ended. Rather, they simply would prefer lending to better borrowers, which are mechanically more likely to come from more-affected lenders, i.e. there is selection on common information.

The key idea to address this challenge is to exploit a comparison with the sample of borrowers who renewed their relationships. This comparison is useful because relationship renewal reflects how informed lenders lend to borrowers and introduces a benchmark against which new lenders can be compared. Models with and without private information make different predictions on the *joint* pattern of renewal and creation of relationships.

To this end, I introduce a two-stage discrete-choice model of firm borrowing. In the first stage, firms try to renew their relationship with their existing lender. Each lender faces a shock impacting its ability to lend. If a borrower fails to receive a new loan from its existing lender, it can turn to new lenders in the second stage. The main ingredient of the model is the existence of three layers of information: (i) all lenders have some information about borrowers, but (ii) each lender has private information about its existing borrowers, and (iii) the econometrician observes neither.

Empirically, the approach relies on the assumption that shocks to banks’ ability to lend are unrelated to the unobservable characteristics of its borrowers.⁶ The first stage can be estimated by regressing the probability that a firm renews its relationship with its existing lender on firm and lender characteristics. The information gap is estimated in the second stage, using the subsample of firms whose relationship ended in the first stage. In line with the "inference hypothesis" above, this stage estimates how the probability that a firm finds a new lender depends on the shock

⁵An equivalent finding in labor markets can be found in Gibbons and Katz (1991).

⁶Any such credit supply shock would do. I follow Chodorow-Reich (2014) to construct a variety of shocks to banks active in the syndicated loan market that appear to be orthogonal to borrowers’ characteristics.

faced by its previous lender. Unlike a purely reduced-form approach, it is possible to control for the mechanical selection on common information of firms that did not renew their relationship. Indeed, the first stage precisely characterizes how renewal depends on shocks to a firm's previous lender. The maintained assumptions are that the distribution of borrower unobservables and the lending rule (as a function of borrower and lender characteristics) are common across lenders, up to some matching error orthogonal to the previous lender's shock.⁷

In the context of U.S. syndicated loans, I find that the information gap is small and thus this friction is unlikely to explain relationship stickiness in this market. Quantitatively, the information gap cannot account for any substantial fraction of the reduced-form effect of bank shocks documented in Chodorow-Reich (2014). These estimates reveal that the reduced-form patterns consistent with the "inference hypothesis" are indeed almost completely driven by information common to all lenders but unobservable to the econometrician. This finding stands in contrast to the predominant view in empirical studies on this market that invokes the information gap as an important mechanism behind stickiness.⁸

Practically speaking, the result above is plausible as the syndicated loan market is dominated by large banks and among the most transparent firms. Nevertheless, there must be a different friction behind stickiness as there is clear reduced-form evidence that bank shocks matter in this market. A likely friction is the "covenant channel" recently documented by Chodorow-Reich and Falato (2019). They find that banks subject to worse shocks are more likely to act on covenant violations and push their borrowers into technical default. In turn, new lenders are reluctant to a borrower with an unresolved covenant violation because of the uncertain resolution of that violation. Moreover, firms that have breached an interest coverage or debt covenant often face a contractual prohibition on obtaining new lending. Their design cannot rule out informational frictions explicitly, but the size of their estimate is large and adds external validity to my result.

A number of additional results provide support for the main finding above. First, I provide model-free evidence consistent with the hypothesis that borrowers looking to switch lenders are worse on average. In particular, firms that are better along observable characteristics are more likely to renew their existing relationship. Moreover, borrowers that switch banks receive worse loan terms than those who stayed with their current bank. Second, I study a number of extensions. The main findings are quantitatively unaffected by alternative measures of bank shocks. Results are

⁷As opposed to the literature that identifies informational frictions by comparing firms with different degrees of opacity, this approach relies on comparing how lenders with different information would treat the same firm. The two-step approach bears some resemblance to econometric models in the line of Heckman (2013) but is used to account for differences in information among agents, a feature that is absent from these models.

⁸Examples include Saunders and Steffen (2011), Schenone (2010), Ferreira and Matos (2012), Bharath et al. (2011) or Dass and Massa (2011).

similarly unchanged when I introduce a comparative advantage for previous lenders that potentially varies with firm characteristics, alleviating the concern that the results are driven by other frictions whose magnitudes vary systematically across firms. Importantly, I show that a "naive" version of the model that ignores common information drastically overestimates the information gap and its effect on lending. I also provide Monte Carlo evidence suggesting that noise in the data is unlikely to explain away the main finding. Finally, I show that the information gap is smaller for larger and more transparent firms relative to others, whether classified based on public listing, inclusion in the Compustat database or other measures.⁹

This paper is organized as follows. Section I describes the U.S. syndicated corporate loan market and the data. Section II develops an empirical model of relationships in which lenders have different information about borrowers and discusses the identification strategy. Section III describes the estimation results. Section IV concludes.

Related work: This paper relates to prior studies that estimate the effect of credit supply shocks. It is closest to Chodorow-Reich (2014), given that the empirical application is based on his framework for measuring relationships and bank shocks in the U.S. syndicated loan market after the financial crisis.¹⁰ Relative to this literature, I focus on the explicit mechanism through which bank-specific shocks impact firm borrowing and isolate the effects of informational frictions. This paper offers a direct empirical test of the model of Dell’Ariccia and Marquez (2004). I also contribute to the growing literature on estimating informational frictions in credit markets by estimating a model. Crawford et al. (2018) quantify asymmetric information in lending markets by estimating a structural model of credit demand on small business loans data from Italy. Their main focus is on the interaction of market power and asymmetric information between borrowers and lenders in the market for first-time loans, abstracting from dynamic issues of lending relationships.

This paper also adds to the growing body of empirical studies examining the impact of asymmetric information in credit markets by developing new empirical strategies to measure private information. For instance, Adelino et al. (2019) study trade delays as a signal of quality in the mortgage market. Botsch and Vanasco (2019) study learning by lending in banking relationships in the syndicated loan market. Stroebel (2016) exploits the difference between vertically integrated and nonintegrated mortgage lenders to show that asymmetric information about collateral values is a significant source of adverse selection in this market. Kurlat and Stroebel (2015) find that asymmetric information with sellers and within buyers is substantial and has key implications

⁹Because the information gap is identified from a small sample of switchers, the statistical power to estimate cross-sectional heterogeneity is limited. In principle, the analysis could be extended to other sources of heterogeneity, such as borrower’s access to public debt markets (Schwert, 2018) or personal relationships with lenders (Karolyi, 2018).

¹⁰See also Peek and Rosengren (2000), Khwaja and Mian (2008) and Jiménez et al. (2019), among others.

for housing markets. Hertzberg et al. (2018) find evidence of screening on loan terms in online credit markets. Finally, a number of works study soft information in lending, including Hoshi et al. (1991), Mian (2006), Liberti and Mian (2009), Hertzberg et al. (2010), Rajan et al. (2015), Sutherland (2018), Iyer et al. (2016), Agarwal and Hauswald (2010), and Keys et al. (2010). For a sharp analysis of lending relationships and matching in the syndicated loan market, see Schwert (2018).

My focus on market breakdown and aggregate lending complements the banking literature that studies interest rates, motivated by the increase in lenders' bargaining power that comes with an information monopoly and a "hold-up" problem, including Schenone (2010), Petersen and Rajan (1994), Degryse and Van Cayseele (2000), Berger and Udell (1995) or D'Auria et al. (1999). I also relate to works on the impact of competition on interest rates, as studied for instance in Petersen and Rajan (1995) or Ruckes (2004).

I. Relationships in the U.S. Corporate Loan Market

A. *Data*

In order to trace relationships between borrowers and lenders, it is necessary to have loan-level data describing who borrowed from which bank at each point in time. As a result, balance sheet data on borrowers or banks are not sufficient on their own. The data come from the DealScan database, which covers the syndicated corporate loan market in the United States. Studying the corporate loan market is interesting in itself given the role it plays in driving economic growth.¹¹ Syndicated loans play a central role in the American corporate loan market, and the Federal Reserve's Terms of Business Lending survey estimates that in recent years they accounted for about 50% of commercial and industrial lending with a maturity of more than one day and 60% of loans with a maturity of more than one year. These loans are typically large, and the median loan is about \$300 million in my sample. These loans are typically not made by a single lender but by a consortium referred to as a syndicate. Lenders in a syndicate are typically large banks that are divided between lead lenders and participants. Lead lenders provide a larger share of the funds and have more responsibilities in terms of reporting and monitoring.

A significant body of work shows that relationships matter in this market. Idiosyncratic shocks to banks have negative effects on borrowers through a variety of channels.¹² This is reassuring

¹¹Corporate borrowing has been shown to impact investment (Peek and Rosengren, 2000; Chaney et al., 2012) and firm employment (Chodorow-Reich, 2014; Greenstone et al., 2020), as well as innovation activity (Hombert and Matray, 2017).

¹²For instance, issuance of new loans (Ivashina and Scharfstein, 2010), employment (Chodorow-Reich, 2014), loan

given the relatively large size of these firms. However, there are reasons to believe that estimates of information frictions from DealScan would constitute a lower bound compared to the average smaller, more opaque firm. Nevertheless, the sample is not restricted to public firms and includes about 60% of private firms that are more dependent on lending relationships. Only about half of the firms in my sample are in the Compustat database. About 90% of the lending agreements signed before the crisis include credit lines, which are flexible liquidity management tools that resemble credit cards offered to households. These credit lines allow borrowers not to commit ex ante to any loan size and are thus difficult to replace with other forms of financing, such as bond issuance. Following Chodorow-Reich (2014), the timeframe centers on the bankruptcy of Lehman Brothers in September 2008. I divide the sample into a pre-crisis period spanning January 2004 to August 2008 and the crisis period spanning October 2008 to December 2010.¹³ I include only loans made to non-financial, American firms that were used to finance the operations of the firm.¹⁴ Table AI displays summary statistics for the borrowers in the sample.

B. The Transmission of Bank Shocks through Relationships

How do bank shocks impact credit markets? A key idea is that sticky relationships make credit markets more vulnerable. If a friction prevents some borrowers of a distressed bank from switching to a new, relatively healthier lender, even idiosyncratic bank shocks can have aggregate effects. The goal of this paper is to estimate the aggregate effects of an informational switching friction. This friction is the predominant explanation for relationship stickiness in the empirical literature studying the syndicated loan market.¹⁵

The key premise is that not all lenders share the same information about borrowers. Lenders that have lent to a firm in the past have acquired private information over the course of this relationship that is unknown to other lenders. I dub this difference in information across lenders the *information gap*. The information gap affects credit reallocation through the following key channel: because lenders have private information about their existing borrowers, they are able to selectively choose which relationships to end when faced with a shock that forces them to reduce lending. This "cherry-picking" implies that borrowers whose relationship ended face stigma: potential new

pricing (Santos, 2011), loan contract strictness (Murfin, 2012).

¹³Results are robust to changes in the exact time window.

¹⁴That is when the purpose of the loan is declared as "working capital" or "corporate purposes" as opposed to M&A activity or debt restructuring. In general, it is difficult to gauge how accurate these labels are. In principle, for public firms corporate filings are possible source of additional information about funds use, as in Korteweg et al. (2019).

¹⁵See for instance, Peek and Rosengren (2000), Schenone (2010), Petersen and Rajan (1994), Murfin (2012), Chodorow-Reich (2014), Degryse and Van Cayseele (2000) and Berger and Udell (1995). Classical theoretical references include Sharpe (1990), Rajan (1992), Detragiache et al. (2000) and Dell'Ariccia and Marquez (2004).

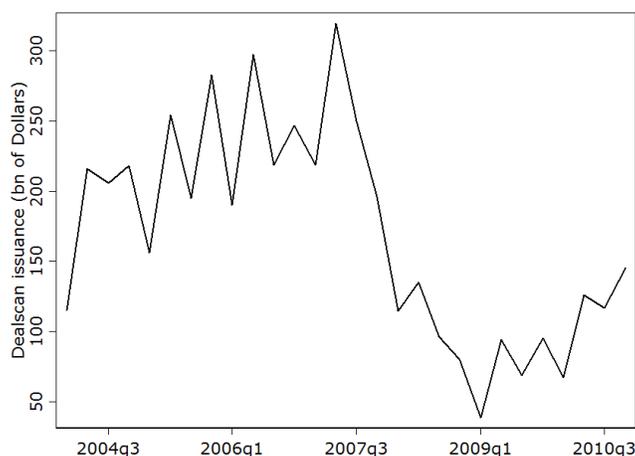


Figure 1. Aggregate corporate issuance in DealScan (in 2008 U.S.\$) 2004–2010.

lenders are wary that these borrowers are of lower quality. This negative signal makes it difficult for borrowers to switch lenders and the information gap leads to imperfect credit reallocation.

Can we find evidence of this channel in the data? Like many other lending markets, the syndicated loan market suffered an unprecedented collapse after the financial crisis. Figure 1 shows that the issuance of new loans was cut in half in the sample, from an average of about \$200 billion of new loans per quarter before the crisis to only \$100 billion afterwards. An important feature of this market is that this collapse occurred at the *extensive margin of credit*. In fact, loan size remained stable; it was a sharp decrease in the number of firms receiving new loans that depressed lending volume after the crisis.

To study this drop in the extensive margin of credit, I focus on firms with an existing loan in the pre-crisis period spanning January 2004 to August 2008 then ask: how many of these firms obtained a new loan in the crisis period spanning October 2008 to December 2010? If they did, from whom did they borrow? Did they renew their existing relationship, or did they find a new lender and form a new relationship? This distinction is key for the empirical strategy developed in this paper.¹⁶ Table I displays the decrease in lending over this period. The share of firms obtaining a new loan fell drastically after September 2008, to almost half its level in normal times.¹⁷ Changes in loan terms are shown in Table AII in the Appendix. The last two columns of Table I reveal how

¹⁶The pre-crisis time length is chosen to match the typical maturity of loans before the crisis, which was just less than four years. The crisis time length is chosen to be two years because firms typically sign a new loan two years before their existing loan expires.

¹⁷Lending accounts for new loans that are occasionally misclassified as loan modifications. For instance, the renewal of a two-year credit line can sometimes be reported as a two-year extension of an existing credit line. In all that follows, I classify a firm as borrowing after the crisis if it received a new loan or a modification of an existing loan granting extra funds. See also Roberts and Sufi (2009) for issues of misclassification in DealScan.

Table I

The extensive margin of credit: 2004–2010.

Year t	% firms with new loan by $t + 2$	Renew	New lender
2004	42.58%	37.55%	5.03%
2008	25.05%	21.17%	3.88%
2010	36.23%	31.01%	5.22%

A loan in the post-period is classified as made by a new lender if no lead lender of its lending syndicate was a lead lender of its last pre-period syndicate. The sample is restricted to U.S. nonfinancial firms that list the reason for borrowing as "working capital" or "corporate purposes."

sticky lending relationships are in this market. The share of firms forming a new relationship is strikingly small at 3.88%; that is only about one-fifth of the share of firms renewing their existing relationship.¹⁸

Table I clearly illustrates the precise question this paper seek to answer: How many firms were not able to form a new relationship because potential new lenders knew less than existing lenders? In other words, how much larger than 3.88% would the share of firms finding a new lender be if all banks had the same information? This counterfactual represents the aggregate effect of this information friction, via imperfect credit reallocation.

Note that this question focuses on the *extensive* margin of lending — firms' access to credit — and most of the paper relegates loan terms to the background. Although in general, data on loan terms are informative about information frictions, there are three specific reasons for this choice in this particular application. First, in terms of welfare loss, market breakdown is typically thought to have more devastating effects than higher interest rates or stricter covenants. Second, few firms switch lenders in my sample, which causes the analysis of loan terms offered to new borrowers to lack statistical power.¹⁹ Finally, loan terms are not only affected by informational frictions, but also by the distribution of bargaining power between borrower and lender. Conceptually, this is a complex issue. For example, one can expect new borrowers to pay a higher interest rate because new lenders are less informed. On the other hand, because relationships are sticky going forward,

¹⁸Because loans are made by a syndicate of lenders, one needs to take a stance of how to define a "new lender". The classification used in Table I compares the the syndicate of the last pre-crisis loan received by a firm to that of its first new crisis loan (if any). Because of their special role as information gatherers, I restrict attention to lead lenders when classifying new relationships: a firm is classified as "borrowing from a new lender" if no lead lender in its first post-crisis loan syndicate was a lead lender of its last pre-crisis loan syndicate. Reclassifying loans arranged by former participants do not materially affect the results. I treat mergers in the same way as Chodorow-Reich (2014), see that article's footnote 23.

¹⁹Unfortunately, this sample size problem makes estimating the rates offered by other potential lenders as in Crawford et al. (2018) impossible in this application. More generally, in practice, the large degree of stickiness in many settings makes it difficult to estimate full demand systems for switchers.

new lenders have incentives to offer low rates in an attempt to lock-in borrowers, as suggested by the literature on switching costs (Farrell and Klemperer, 2007). Although there are a number of applications for which the first and second points can be addressed, this last point requires a clear identification strategy to control for relative bargaining power. The empirical model of lending I develop in this paper can be estimated for any distribution of bargaining power.

C. *Measuring Bank Shocks*

Heterogeneous bank shocks are at the heart of the empirical strategy. I follow the construction of Chodorow-Reich (2014), who meticulously argues that these are valid supply shocks in that particular setting. The main measure of lenders' exposure to the crisis δ is defined by the relative change in lending at each bank after the Lehman bankruptcy that occurred in September 2008. For each lender, I then count the number of loans made in the crisis period to firms that received a loan pre-crisis (from this particular lender or any other lender in the sample). I divide this number by the total number of loans made in the pre-crisis period by this lender, adjusting for the asymmetrical time window between the two periods. Moreover, because these loans are syndicated across multiple lenders, I weight each element in the numerator and the denominator by the loan share of that particular lender:²⁰

$$\delta^l = 1 - 2 \times \frac{\#\text{post-crisis loans made by lender } l \text{ to firms that borrowed pre-crisis}}{\#\text{pre-crisis loans made by lender } l} \quad (1)$$

A larger δ^l implies that fewer loans were made during the crisis and indicates a more affected lender; a constant loan supply at the bank level would result in a δ^l of zero. Because syndicated loans are made by multiple lenders, I transform this lender-level measure into a firm-level measure by exploiting the structure of the firm's pre-crisis syndicate in the same way as Chodorow-Reich (2014). For each firm f , I compute a weighted average of these lender δ^l , using as weights the loan shares $\omega^{l,f}$ of each lender in the syndicate s^f of the last pre-crisis loan of this particular firm. This yields a clear measure of the credit supply shock faced by this firm $\delta^{s^f} = \sum_{l \in s^f} \omega^{l,f} \delta^l$.

As shown in Figure A1 in the Appendix, the mean of this measure is about 50%, which is in line with the aggregate dollar figure presented in Figure 1. Moreover, with a standard deviation of 13%, firms face a variety of supply shocks consistent with the idea that the need for reallocation arises after a crisis. Section III considers other measures of bank shocks, including exposure to Lehman Brothers, real estate charge-offs, or stock price correlation with the ABX mortgage-backed

²⁰Because the data on loan shares are occasionally missing, I follow the method introduced in Chodorow-Reich (2014) to recover them via imputation. This measure excludes loan modifications, as there is too little data to consistently recover loan shares in that case.

securities index. Finally, for ease of exposition, in the remainder of the paper I will often refer to a syndicate as a "bank" or a "lender" and write δ^b instead δ^s , even though it is understood that firms borrow from multiple lenders at once.

II. Estimating Informational Frictions

In principle, a number of frictions could explain why relationships are sticky in this market. The predominant view points to the informational advantage of existing lenders over potential new lenders. However, estimating this information gap has proven elusive. In particular, asymmetric information is by nature difficult to measure directly. This paper introduces an empirical strategy to estimate the information gap from the observed patterns of renewals and new relationships in the data. In particular, it addresses the key identification challenge of information common to all banks but not present in the econometrician's dataset. Below, I present a discrete choice of relationship formation that can be taken to the data while explicitly allowing for differences in information between lenders and also between lenders and the econometrician. It leverages two sources of cross-sectional variation in the data: (i) some lenders have lent to a firm in the past, while other have not; and (ii) lenders face different shocks to their propensity to make loans.

A. Setup

Consider a firm f with an existing relationship with lender b . After a shock to banks, the firm has a new project that requires financing (or an older project that requires new funds) and can ask for a new loan. The firm's type, possibly different from its past type, is characterized by $\{x^f, \nu^f\}$.

- *Firm observables x^f* : This term includes all the controls available to the econometrician. In the DealScan sample, this includes firm characteristics such as public ownership, sales, and industry, as well as rich information on loans received before the crisis. Loan terms include pre-crisis loan size and spread, whether it was collateralized, and whether it covered the crisis period, as well as whether the firm had multiple pre-crisis loans.
- *Firm unobservable type ν^f* : This term includes all firm characteristics that are unobservable to the econometrician. In other words, it corresponds to the residual in a regression framework. However, lenders have varying degrees of information about this ν^f . This information gap is the source of the friction that this paper wants to estimate.

To make this *information hierarchy* transparent, I decompose the firm type as follows:

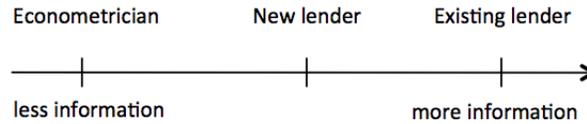


Figure 2. Information hierarchy.

$$\nu^f = \nu_1^f + \mathcal{W}\nu_2^f$$

Both ν_1^f and ν_2^f vary across borrowers such that:

1. ν_1^f is *common information*, observed by all lenders
2. ν_2^f is *private information* of the firm's previous lender
3. The econometrician can observe neither ν_1^f nor ν_2^f

The main parameter of interest is the *information gap* \mathcal{W} that represents the weight on the previous lender's private information.²¹ Intuitively, there is an informational hierarchy and three levels of information, as depicted in Figure 2. At one extreme is the firm's previous lender, who knows both ν_1^f and ν_2^f , as well as observables x^f . At the other extreme is the econometrician, who knows only x^f . \mathcal{W} measures how informed new lenders are relative to these two extremes. If $\mathcal{W} = 0$, all lenders share the same information about the firm and the information gap is zero. As \mathcal{W} increases, so does the information gap between the firm's previous lender and its potential new lenders.²²

Note the broad interpretation of the firm's type, defined as its propensity to receive a loan absent informational frictions. Firm characteristics can impact the propensity to receive a loan in equilibrium through: (i) creditworthiness, and (ii) demand for bank loans. Banks are naturally less willing to lend to firms with a poor track record or in a fledging industry. However, some "good" firms may be unwilling to borrow at the rate offered by banks because they have enough financial slack, or other funding opportunities outside of the banking sector. This broad interpretation is

²¹The factor \mathcal{W} is not separately identified from the variance of ν_2^f , therefore for the rest of the paper I adopt the normalization that the $Var[\nu_2^f] = 1$. The standard deviation of the privately observable component is therefore \mathcal{W} .

²²The main specification estimates a single information gap \mathcal{W} across all firms, but in principle it can vary with firm characteristics. For instance, the last section studies how \mathcal{W} varies in the cross-section by comparing public and private firms, as well as firms included in the Compustat database versus others. In principle, the analysis could be extended to other sources of heterogeneity. However, the statistical power to estimate cross-sectional heterogeneity is limited because the information gap is identified from a small sample of switchers in the DealScan data.

the correct one in the context of estimating market breakdown. The question is whether firms with high enough type ν^f fail to receive a loan because of informational frictions. In terms of inefficiency, it makes no difference whether a firm has a low type that reflects low creditworthiness or low demand for bank loans. In both cases, the firm would not borrow in the counterfactual of no information friction.

Besides firm type, two other forces drive lending:

- *Aggregate shock* μ_0 : This term captures other factors that reduced lending after the crisis, independently of informational frictions. It accounts for the financial turmoil that affected all lenders equally, as well as demand shocks for end products that affected all firms. It also captures other types of aggregate shocks to lending, such as "uncertainty" shocks or events in other lending markets. This term ensures that counterfactual lending is calibrated properly to the new period: total lending can be much lower than before the shock to the banking sector. Moreover, the fact that it captures both aggregate demand and supply shock is not a problem as this paper focuses on a *reallocation* friction, i.e. whether bank-specific shocks have aggregate effects because of the information gap.
- *Bank-specific shock* δ^b : This term captures the credit reallocation problem: Beyond the aggregate shock μ_0 , some lenders were more affected than others. The key question is whether borrowers that saw their relationship end were able to reallocate towards new, relatively healthier lenders. I follow Chodorow-Reich (2014) to construct δ^b as the bank's exposure to the real estate crisis, as described in Section 2.

The timing of the model assumes two stages. The firm first tries to renew its existing relationship and negotiates for a new loan with its pre-crisis lender. If the lender instead chooses to end the relationship, the firm has the possibility of trying to form a new relationship and obtaining a loan from a new lender.²³ Figure 3 illustrates the setup. The next section solves for the equilibrium of each stage in turn.²⁴

²³Institutional details justify this timing assumption. Shopping around for lenders is difficult as corporate loans must be tailored to the specific borrower, and loan terms offered by potential lenders are not publicly available. Firms tend to first bargain with their existing lender to save on transaction costs associated with identifying a potential new lender, as commencing a negotiation takes time and involves substantial communication costs.

²⁴A two-stage model is appropriate given my focus on the inference drawn by new lenders after the termination of an existing relationship. Another possible modeling approach relies on pairwise stability of matches, such as in Chen and Song (2013) and Schwert (2018). This approach can successfully fit relationship data in the syndicated loan market over the past decades, but does not explicitly describe the underlying information structure.

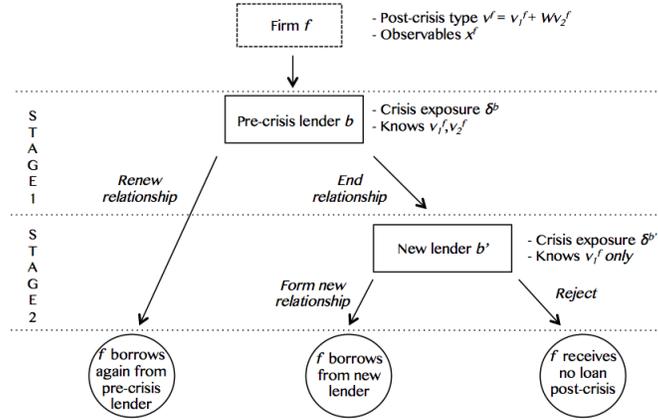


Figure 3. The model of firm post-crisis borrowing.

B. Equilibrium Lending

I model lending as the outcome of a bargaining between firm and lender. Lending generates a surplus $s(\nu^f, x^f, \mu_0, \delta^b)$ that depends on both firm and lender characteristics and that can be divided between the pair. For ease of exposition, I adopt the following sign convention: surplus increases in firm type ν^f and decreases in bank shock δ^b . A loan is made if there is positive expected surplus, i.e. lending has a positive NPV, given the information of the lender:

$$\mathbb{E}_b[s(\nu^f, x^f, \mu_0, \delta^b)] > 0 \tag{2}$$

As emphasized above, the key friction is that not all banks b have the same information, reflected in the operator \mathbb{E}_b . The firm’s previous bank observes the complete firm type ν^f , while other lenders only observe its common component ν_1^f .

Note that the model makes no prediction on how the surplus is shared in the pair, i.e. interest rate or strictness of loan terms. In fact, it is consistent with any distribution of relative bargaining power between firm and lender. As explained in Section B, this has the advantage of not confounding informational frictions with issues related to bargaining power.²⁵ Moreover, while simple, this lending rule can be microfounded in the spirit of Levin’s (2003) model of relational contracts. While his model is fully dynamic and includes moral hazard, hidden information, and unverifiable performance, his Theorem 1 shows how to characterize optimal contracts through a very similar condition.

In much of the paper, I analyze the case of linear surplus and normally distributed unobservable firm type. This choice helps with power issues when taking the model to the data and simplifies

²⁵Schwert (2018) also takes the division of surplus as unobserved for similar reasons.

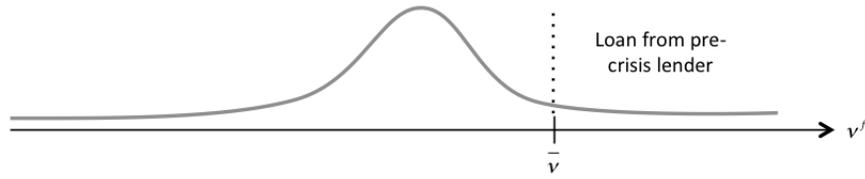


Figure 4. Stage 1: renewing existing relationships.

exposition somewhat.²⁶ However, in the Internet Appendix I show that all comparative statics derived in this section are valid under much weaker technical conditions. More specifically, I make the following parametric assumptions:

- Linear surplus: $s = \nu^f + x^f \mu + \mu_0 + \delta^b \beta$
- Normality: $\nu_1^f \sim \mathcal{N}(0, \sigma_1^2)$, $\nu_2^f \sim \mathcal{N}(0, 1)$, so that $\nu^f \sim \mathcal{N}(0, \sigma_1^2 + \mathcal{W}^2)$

B.1. First Stage: Relationship Renewal

First, the firm can try to renew its relationship with its existing lender and obtain a new loan. This lender knows ν^f and the loan is granted if there is positive surplus between the pair given this information. The relationship is renewed for firms with a sufficiently high type, relative to the bank shock δ^b . Firm f renews its relationship with lender b if:

$$s(\nu^f, x^f, \mu_0, \delta^b) \geq 0 \iff \nu^f \geq \bar{\nu}(\mu_0, x^f, \delta^b) \quad (3)$$

This "cherry-picking" corresponds to a simple cutoff rule for renewing relationships and Figure 4 illustrates the equilibrium of the first stage. Firms above the cutoff renew their relationship, while firms below are left looking for a new lender. Because the lender has access to private information, it selectively chooses to renew its relationship with its best borrowers.²⁷

In the linear-normal case, this cutoff rule $\bar{\nu}$ is linear in firm observables and bank shocks: $\bar{\nu} = -\delta^b \beta - x^f \mu - \mu_0$. The probability that firm f renews its relationship with lender b is thus

²⁶The linearity and normality assumptions are widely used in empirical work estimation asymmetric information such as Einav et al. (2012) and Crawford et al. (2018).

²⁷In practice, banks sometimes "evergreen" loans and renew some of their worst borrowers to prevent recognizing losses. A cutoff rule cannot imply both cherry-picking and evergreening, so it is important to establish which forces dominates empirically. Below, I provide multiple model-free facts that support cherry picking. Nevertheless, widespread evergreening would lead to a bias when estimating the informed lender decision rule.

given by:

$$\mathbb{P}(\text{borrow from pre-crisis lender}) = \mathbb{P}(\nu^f \geq -\delta^b \beta - x^f \mu - \mu_0) \quad (4)$$

$$= \Phi(\mu_0 + \delta^b \beta + x^f \mu) \quad (5)$$

where $\Phi(\cdot)$ is the normal cdf, which is equivalent to a standard probit model. The coefficients β and μ of the surplus function can therefore be recovered via a probit regression.

The cutoff $\bar{\nu}$ naturally depends on firm and lender characteristics as well as the aggregate shock μ_0 . In particular, lenders that are more affected by the crisis renew fewer relationships: the cutoff moves to the right. This comparative statics is the origin of the selection effect that plays an important role in the second stage equilibrium. This cutoff $\bar{\nu}(\mu_0, x^f, \delta^b)$ represents the *informed lender decision rule* and plays a crucial role in the estimation of the information gap in the second stage.

Model-free evidence of cherry-picking: There is reduced-form evidence consistent with the hypothesis that borrowers looking to switch lenders are worse on average. In particular, firms that are better along observable characteristics are more likely to renew their existing relationship as shown in Table II in Section III. Moreover, borrowers that switch banks receive worse loan terms relative to those who stayed with their current bank. Controlling for lender and borrower characteristics, Table IAIII in the Internet Appendix shows that loans made to new borrowers are significantly smaller, carry a larger spread, and have shorter maturities.²⁸ To the extent that these loan terms are informative about borrower quality, these results suggest that new borrowers are less creditworthy than repeat borrowers.

B.2. Second Stage: New Relationship Formation

Firms that saw their relationship end in stage 1 can try to form a new relationship and borrow from a new lender b' . The new lender cannot observe the full type ν^f but has two sources of information. First, it can directly observe the common information component ν_1^f . Second, leaving a relationship is a signal in itself and lets the lender make an inference about the private information component ν_2^f . In particular, the set of firms looking for a new lender is selected:

$$\nu^f \leq \bar{\nu}(\mu_0, x^f, \delta^b) \iff \nu_2 \leq \frac{\bar{\nu}(\mu_0, x^f, \delta^b) - \nu_1}{\mathcal{W}} \quad (6)$$

Firms that saw their relationship end can try to form a new relationship with a new lender

²⁸The Internet Appendix is available in the online version of this article on the *Journal of Finance* website.

b' . However, this lender knows only ν_1^f and lends only if there is a positive expected surplus, conditional on its information:

$$\mathbb{E}[s|\nu_1, \nu_2 \leq \frac{\bar{\nu} - \nu_1}{\mathcal{W}}] = \nu_1^f + \mathcal{W}\mathbb{E}_{b'}[\nu_2^f] + \mu_0 + x^f\mu + \delta^{b'}\beta + \epsilon \quad (7)$$

This expression reveals that there are two reasons new lenders might impose different lending standards than existing lenders. The first is the information gap: new lenders do not observe ν_2^f and must instead draw an inference about it. Second, the term ϵ captures all matching frictions that are not informational in nature. For instance, borrowers might have preferences for specific banks because of specialization across regions or industries. It also captures other switching costs, such as upfront due diligence costs or search costs.²⁹

New lenders also use a cutoff rule to form a new relationship. Firm f coming from lender b forms a new relationship with lender b' if:

$$\mathbb{E}[s|\nu_1, \nu_2 \leq \frac{\bar{\nu} - \nu_1}{\mathcal{W}}] \geq 0 \iff \nu_1^f \geq \nu^*(\mu_0, x^f, \delta^b, \delta^{b'}, \epsilon) \quad (8)$$

The key difference from the first stage is that the cutoff rule is different: $\nu^*(\mu_0, x^f, \delta^b, \delta^{b'}, \epsilon)$ represents the *uninformed lender decision rule*. In general, this rule is stricter than the informed decision $\bar{\nu}$ because borrowers looking for a new lender are adversely selected.³⁰ This is consistent with the reduced-form evidence discussed above that switchers face harsher loan terms relative to borrowers renewing their relationship.

Inefficiencies and Market Breakdown: The information gap only matters to the extent that it impacts the allocation of credit. Figure 5 illustrates the effect of the information gap on lending. The x-axis represents the commonly observed component of firm's type ν_1^f , and all firms to the right of the cutoff ν^* receive a loan from a new lender. The y-axis represents firms' true type $\nu^f = \nu_1^f + \mathcal{W}\nu_2^f$. If the information gap \mathcal{W} were zero, all firms, denoted by black dots, would lie on the 45-degree line. Instead, when new lenders have less information, firms are scattered around the diagonal. If the information gap were zero, lenders would lend to firms whose *true* type is above the cutoff $\nu^{\text{full info}}$.

The two cutoff rules delineate three areas. The dotted area corresponds to firms that are good enough along all dimensions to receive a loan even if new lenders have less information. On the other hand, the grey area corresponds to underfunding: these firms are unable to receive a new

²⁹As modeled by Boualam (2018).

³⁰The proof is in the Appendix. Moreover, as the information gap goes to zero, this rule converges to the informed lender cutoff $\bar{\nu}$.

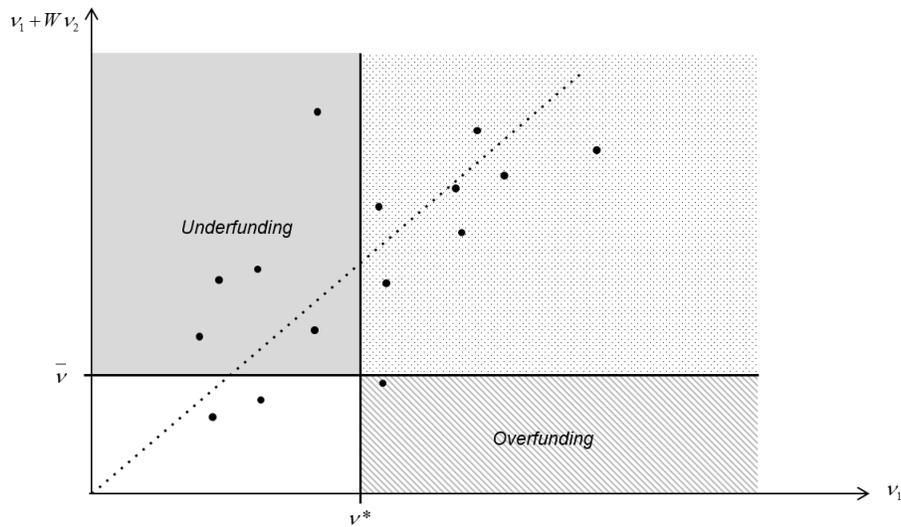


Figure 5. The effect of the information gap on lending.

loan because of the information gap. The firms are good overall, but happen to be worse along the commonly observed dimension ν_1 . Interestingly, there is a third region of overfunding: some firms happen to be particularly good only along the dimension ν_1 . In fact, overfunding can dominate if the information gap \mathcal{W} is close to zero.

C. The Inference Hypothesis

Importantly, the uninformed lending rule ν^* depends on the size of the firm's previous lender shock:

Inference hypothesis: If $\mathcal{W} > 0$, ν^* decreases with the firm's previous lender crisis shock δ^b .

In other words, new lenders apply looser lending standards to borrowers that ended a relationship with the most-affected lenders. Because new lenders learn something from a prior relationship being ended, borrowers coming from the most-affected lenders face *less* stigma. Intuitively, the most-affected lenders have to end a large number of relationships, including some with relatively good borrowers. The inference hypothesis is the key empirical prediction of the information gap. The key idea to keep in mind is that private information affects the *lending rule* used by new lenders.

This hypothesis is the starting point of the empirical strategy used this paper: in principle, one can exploit bank shocks, i.e. cross-sectional variation in banks' propensity to lend, to estimate the information gap. The inference hypothesis corresponds to the following empirical prediction:

conditional on its previous relationship having ended, the probability that a firm borrows from a new lender *increases* with the size of the shock faced by its previous lender.

At first glance, there is evidence for this inference hypothesis in the U.S. corporate loan market during the crisis. Consistent with the inference hypothesis, Figure 6 shows that among firms that saw their relationship end, firms coming from more-affected lenders are more likely to obtain a loan from a new lender relative to firms coming from less-exposed lenders. Table AIII in the Appendix replicates these findings in a regression framework controlling for borrower characteristics and past loan terms. Conditional on leaving a relationship, a one standard deviation increase in the crisis exposure of a firm's existing lender implies a 20% increase in the probability of borrowing from a new lender.³¹ Figure A2 in the Appendix shows that this result is robust to alternative measures of bank shocks. Note also that the pattern in Figure 6 is not driven by the matching of bad borrowers with less healthy lenders. Indeed, if that were the case, the correlation between pre-crisis lender health and new relationships would go the other way: borrowers coming from less-affected lenders would be more likely to borrow from a new lender.³²

However, the reduced-form correlation in Figure 6 alone cannot identify the information gap. In fact, the same correlation would arise even if all banks had the same information as long as some of this common information were unobservable to the econometrician. To see this, note that in this case lenders do not *learn* anything from a relationship having ended. They simply want to lend to borrowers that are good enough. But, mechanically, there are more good borrowers leaving more-affected lenders, because these lenders have to end a larger number of relationships. This channel alone, independent of private information and the information gap, could explain the correlation. Empirically, it is crucial to isolate private information because common information by itself would not lead to any inefficiency: relationships appear sticky, but with no aggregate effects.

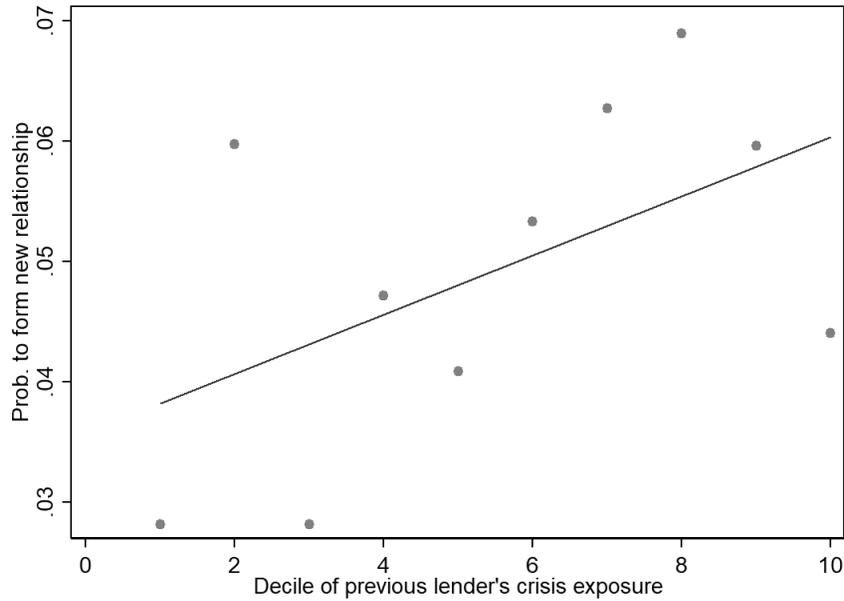
To see this more precisely, note that the model equivalent of the reduced-form correlation is given by:

$$\mathbb{P}(f \text{ borrows from a new lender}) = \mathbb{P}(\nu_1^f \geq \nu^*(\mu_0, x^f, \delta^b, \delta^{b'}, \epsilon) | \nu^f \leq \bar{\nu}(\mu_0, x^f, \delta^b)) \quad (9)$$

where ν^* is the *uninformed lender decision rule* given information gap \mathcal{W} . This expression

³¹In absolute terms, this corresponds to a 0.8pp increase, compared to a mean of 3.9%. An equivalent finding in labor markets can be found in Gibbons and Katz (1991).

³²While there is evidence of matching of bank-dependent firms with highly-capitalized banks in the past decades in the syndicated loan market (Schwert, 2018), this matching does not seem to invalidate the measure of bank shocks used in this application. The correct interpretation is that random matching would lead to larger observed effects of bank shocks on credit. See Schwert (2018) Section IV.C for an insightful discussion.



Note: This sample includes only borrowers who did not renew their previous relationship after the crisis. A borrower is classified as renewing a relationship if at least one lead lender of its post-crisis lending syndicate was a lead lender of its last pre-period syndicate. Other borrowers receiving a new loan after the crisis are classified as forming a new relationship. The crisis exposure of a firm's pre-crisis lender is computed as the weighted average of the relative drop in lending between 2004 to 2008 and 2008 to 2010 of each lender in the firm's last pre-crisis lending syndicate, weighted by the loan shares of each lender. The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes."

Figure 6. Pre-crisis lender crisis exposure and the formation of new relationships.

makes clear that previous lender shock δ^b is correlated with the probability of forming a new relationship through two channels:

1. Inference about private information ν_2^f : how new lenders adjust their lending rule ν^* with δ^b
2. Selection on common information ν_1^f : only the subsample with $\nu^f \leq \bar{\nu}(\mu_0, x^f, \delta^b)$ reaches the second stage

To illustrate, consider an extreme case with no private information, that is $\mathcal{W} = 0$. In this case, new lenders do not learn anything from a relationship having ended. Formally, they use the same lending rule as an informed lender: $\nu^*(\delta^b, \delta^{b'}) = \bar{\nu}(\delta^{b'})$. There is no longer any inference: this lending is independent of the shock to the firm's previous lender δ^b . However, the sample reaching stage 2 is selected: $\nu_1^f \leq \bar{\nu}(\delta^b)$. The model equivalent of the reduced-form correlation is given by:

$$\mathbb{P}(f \text{ borrows from a new lender}) = \mathbb{P}(\nu_1^f \geq \bar{\nu}^f(\delta^{b'}) | \nu_1^f \leq \bar{\nu}(\delta^b)) \quad (10)$$

which implies a positive correlation even without private information.

D. Identification Strategy

It is important to observe that the confounding force of common information stems from the fact that first-stage renewal is naturally correlated with the previous lender's shock δ^b through the informed lender decision rule $\bar{\nu}$. However, $\bar{\nu}$ can be directly estimated by looking at first-stage renewal probabilities. This is a key idea that suggests how to isolate the effect of the information gap. Indeed, data on relationship renewal is the missing piece because lenders are informed about borrowers when deciding whether to renew. This introduces a benchmark against which new lenders can be compared. In fact, both models, with and without private information, can explain Figure 6, but they make different predictions on the *joint* pattern of renewal and creation of relationships.

Specifically, I estimate the full two-stage discrete-choice model and use the first stage estimates to control for selection on common information in the second stage. The inference hypothesis above links the information gap \mathcal{W} to a cross-sectional moment: does the uninformed lending rule ν^* depend directly on δ^b , after controlling for selection in the first stage? If $\mathcal{W} = 0$, there should be no dependence. A stronger dependence indicates a larger information gap \mathcal{W} . In the context

of the model, the equivalent of the reduced-form correlation in Figure 6 is

$$\mathbb{P}(\text{borrow from a new lender}) = \mathbb{P}(\nu_1^f \geq \nu^*(\delta^b) | \nu^f \leq \bar{\nu}(\delta^b)) \quad (11)$$

$$= \int_{\nu^f \leq \bar{\nu}(\delta^b)} \Phi\left(\frac{1}{\sqrt{\mathcal{W}}}(\nu^f - \nu^*(\delta^b))\right) \frac{\phi(\nu^f)}{1 - P_1^f} d\nu^f \quad (12)$$

The left-hand-side variable is observable, and the regressor of interest is the previous lender shock δ^b . Because the coefficients governing $\bar{\nu}$ have been estimated in the first stage, one can isolate the effect of the shock δ^b that works through the uninformed lending rule ν^* . The information gap can therefore be estimated via non-linear least squares on the sample of firms that did not renew their relationship. The rest of the section presents a more precise discussion of this idea and stresses two important identification issues: the role of parametric assumptions and how to account for other non-informational frictions in matching.

Role of parametric assumptions: I implement this idea in three steps, which require a parametric assumption with varying degrees.

Step 1: Estimate informed lender decision rule $\bar{\nu}$

This is a classic discrete-choice model, and in principle, some semi-parametric estimation methods can be used if the data are rich enough. However, due to power concerns, I follow common practice and estimate the normal-linear case, which is simply a standard probit model: the distribution function F of firm type is assumed to be normal, and the latent index is linear in observable characteristics. These parametric assumptions are often thought to be relatively harmless and are widely accepted.

Step 2: Characterize the distribution of types looking for new lenders $F(\nu^f | \nu^f \leq \bar{\nu}(\mu_0, x^f, \delta^b))$

This step requires $\bar{\nu}$ and F , which have been estimated or chosen in Step 1. It is therefore as parametric as Step 1.

Step 3: Estimate \mathcal{W} from the sensitivity of ν^* to δ^b

The last step relies on the inference hypothesis: $\frac{\partial \nu^*}{\partial \delta^b} < 0$ if $\mathcal{W} > 0$ but equal to zero if $\mathcal{W} = 0$. Remember that this proposition is valid under weak technical conditions (described in the Appendix). They apply not only to the normal-linear case but to other parametric choices as well.

In terms of data, the last step leverages the cross-sectional variation in Figure 6: firms that end their relationship with more-affected lenders are more likely to form a new relationship than others. Controlling for selection through Step 2, the data are therefore directly informative about

the slope $\frac{\partial \nu^*}{\partial \delta^b}$. The parametric assumptions on F and surplus are used to go from this slope to \mathcal{W} . In the model, $\frac{\partial \nu^*}{\partial \delta^b} = f(\mathcal{W})$, where f is implicitly determined by equation (8). In the actual estimation, ν^* (and therefore f) is calculated numerically.

It is clear that the shape of f depends the exact parametric assumptions and thus that \mathcal{W} does as well. I do not believe that \mathcal{W} can be recovered fully non-parametrically. Nevertheless, the estimation of \mathcal{W} is also driven by two cross-sections of the data: the cross-section of renewal probabilities (captured by $\bar{\nu}$) and the cross-section of new relationship probabilities described in Figure 6.

Non-informational matching errors: Recall that the model allows for other factors that drive new relationship formation beyond the information gap. The expected surplus in the second stage is given by $\mathbb{E}[s] = \nu_1^f + \mathcal{W}\mathbb{E}_\nu[\nu_2^f] + \mu_0 + x^f\mu + \delta^b\beta + \epsilon$. These "matching errors" are captured by ϵ . For example, borrowers might have preferences for specific banks because of specialization across regions or industries. These matching errors induce a level difference between the lending rules $\bar{\nu}$ and ν^* , of informed and uninformed lenders respectively. However, this difference is not a threat to identification per se. Indeed, as explained above, the information gap is not identified from the level difference $\bar{\nu} - \nu^*$ but rather from whether ν^* depends directly on δ^b .

As a result, the condition necessary for identification is much weaker than assuming no matching errors: it requires that ϵ be orthogonal to the previous lender's shock δ^b . In that case, the uninformed lending rule depends directly on δ^b only if the information gap is positive. In the current setting, this assumption seems plausible because exposure to the real estate market appears to be unrelated to specialization in commercial lending activities. Chodorow-Reich (2014) provides direct evidence that specialization is not a key driver of relationships in the sample period I consider. In the estimation below, I model matching errors in three different ways for robustness: a constant across borrowers, heterogeneity depending on borrowers' observable characteristics, and a constant plus iid noise.

III. Estimation and Results

In this section, I estimate the parameters of the surplus function and the information gap \mathcal{W} that underlies the distribution of firms types:

- Surplus: $s = \nu^f + x^f\mu + \mu_0 + \delta^b\beta$
- Firm types: $\nu_1^f \sim \mathcal{N}(0, \sigma_1^2)$, $\nu_2^f \sim \mathcal{N}(0, 1)$, so that $\nu^f \sim \mathcal{N}(0, \sigma_1^2 + \mathcal{W}^2)$

In this setting, the following decomposition makes interpretation of the information gap clear:

$$\underbrace{\text{Var}[\nu^f]}_{\text{unobservable to the econometrician}} = \underbrace{\sigma_1^2}_{\text{observed by all lenders}} + \underbrace{\mathcal{W}^2}_{\text{unobservable to new lenders}} \quad (13)$$

More specifically, the information gap corresponds to a decomposition of residual variance into how much of what is unobservable to the econometrician is also unknown to new lenders. Because the units of the model are arbitrary, only a relative measure of information is identified:

$$\frac{\text{Std. deviation of unobservables to new lenders}}{\text{Std. deviation of unobservables to econometrician}} = \frac{\mathcal{W}}{\sqrt{\sigma_1^2 + \mathcal{W}^2}} \in [0, 1] \quad (14)$$

As the next section shows, the natural normalization is to impose $\sigma_1^2 + \mathcal{W}^2 = 1$. \mathcal{W} thus captures the relevant measure of relative information.

The parameters $(\mathcal{W}, \beta, \mu)$ are estimated using the two-stage discrete-choice model of firm borrowing described in the previous section. In the first stage, I estimate the probability that a firm renews its relationship with its pre-crisis lender. In the second stage, I estimate the probability that a firm borrows from a new lender, conditional on seeing its relationship ended.

A. First Stage: Relationship Renewal

The firm's previous lender knows the type ν^f and renews its relationship based on a linear cutoff rule³³

$$\nu^f \geq -\delta^b \beta - x^f \mu \quad (15)$$

The probability that firm f renews his relationship with lender b is thus given by

$$\mathbb{P}(\text{borrow from pre-crisis lender}) = \mathbb{P}(\nu^f \geq -\delta^b \beta - x^f \mu) \quad (16)$$

$$= \Phi(\delta^b \beta + x^f \mu) \quad (17)$$

where $\Phi(\cdot)$ is the normal cdf. With the normalization $\sigma_1^2 + \mathcal{W}^2 = 1$, this first-stage estimation equation corresponds to a standard probit model. I therefore estimate the coefficients β and μ on lender and firm characteristics via probit regression. The measure of bank-specific exposure δ^b is the relative change in lending after the crisis by the firm's pre-crisis syndicate, as described in Section C. The vector of firm observables x^f contains indicators for whether the borrower is public, reports sales over the median, is in the manufacturing sector, has an existing loan that

³³For more compact notation, hereafter I include the aggregate shock μ_0 in the vector of coefficients μ by extending the firm observables x^f to include a constant term.

covers the crisis period, and has multiple pre-crisis loans. In addition, I include rich information on terms of the last pre-crisis loan: spread, size, whether it was secured by collateral, and whether there were multiple lead lenders or two or fewer participants in its syndicate.

Table II presents the results. The normalized marginal effect of pre-crisis lender exposure $\hat{\beta}$ is significant and equal to -3.02. The interpretation is that a one-standard-deviation increase in pre-crisis lender exposure decreases the probability of renewing its relationship by 3 percentage points, or about 15% of the average renewal rate. The economic and statistical significance of this measure of bank crisis exposure justifies its use in the second stage to estimate the information gap. Moreover, public or large firms find it easier to renew their relationship.

Identifying β . The identification assumption for the first stage is that δ^b is orthogonal to unobservable crisis characteristics of borrower quality and demand ν^f . This identifying assumption is common to the literature on the firm-level effects of credit supply shocks. The identification threat is that banks that were most-affected by the crisis were matched with corporate borrowers that received a negative shock to their creditworthiness at the same time. Because the crisis originated in the real estate market, it is plausible that these shocks are orthogonal to banks' corporate loan portfolio. Chodorow-Reich (2014) gives a detailed discussion supporting the validity of this assumption in this particular sample. Chodorow-Reich and Falato (2019) confirm this result in a different sample of syndicated loans. Nevertheless, I support this assumption in four different ways, replicating a number of these tests.

First, the evidence presented in Figure 6 goes directly against the idea that bad lenders were matched with bad borrowers. If that were the case, the graph would be *downward* sloping: borrowers coming from the most-affected lenders would be of lower quality and find it *more difficult* to form a new relationship. Second, I run a regression at the level of the borrower-lender pair and compare a specification including firm fixed effects, which absorb any unobservable characteristics related to borrower quality and demand, to OLS estimates based on a full set of borrower controls. Table IAIII in the Internet Appendix shows that the two coefficients on pre-crisis lender's health are virtually identical. This comparison suggests that the bank health measure is indeed orthogonal to unobserved borrower characteristics driving post-crisis loan demand. Third, Figure IA1 in the Internet Appendix shows that the sample is relatively well-balanced on observable firm characteristics. Finally, estimation results are robust to using three other measures of bank exposure to the crisis that have been used as instruments in the literature, as I show in Section B. Overall, these results are consistent with the findings of Chodorow-Reich (2014) and Chodorow-Reich and Falato (2019) that bank-level shocks following the financial crisis are orthogonal to unobservable borrower characteristics in this market.

Table II

First-stage Estimates: Relationship Renewal.

	Outcome: Borrow from pre-crisis lender		
Pre-crisis lender's exposure	-3.86***	-3.52***	-3.02***
	(0.61)	(0.63)	(0.67)
Public		4.54**	2.96**
		(1.45)	(1.46)
High sales		5.82**	3.87**
		(1.48)	(1.59)
Existing loan covers the crisis		-1.78	-1.09
		(1.31)	(1.33)
Multiple pre-crisis loans		9.95***	8.83***
		(1.30)	(1.32)
Manufacturing		0.98	0.77
		(1.33)	(1.32)
Pre-crisis loan terms	-	-	Yes
Mean of dependent variable	21.17%	21.17%	21.17%
R squared	0.83%	4.17%	5.59%
Number of observations	4,044	4,044	4,044

Probit regression. Reported coefficients are marginal effects at the mean of the other covariates, multiplied by 100. A borrower is classified as borrowing from its pre-crisis lender if at least one lead lender of its post-crisis lending syndicate was a lead lender of its last pre-period syndicate. The crisis exposure of a firm's pre-crisis lender is computed as the weighted average of the relative drop in lending between 2004 to 2008 and 2008 to 2010 of each lender in the firm's last pre-crisis lending syndicate, weighted by the loan share of each lender. A firm is classified as having high sales if it reports sales over the median. A firm has an existing loan covering the crisis if the maturity of its last pre-crisis loan is after December 2010. Pre-crisis loan terms include: spread, size, whether it was secured by collateral, and whether there were multiple lead lenders or two or fewer participants in its syndicate. The sample is restricted to U.S. nonfinancial firms that list the reason for borrowing as "working capital" or "corporate purposes."

B. Second Stage: Formation of New Relationships

Firms that saw their relationship end can try to form a relationship with a new lender b' . The uninformed lending rule ν^* is implicitly defined by a positive expected surplus condition

$$\mathbb{E}_{b'}[s] = \nu_1^f + \mathcal{W}\mathbb{E}_{b'}[\nu_2^f] + x^f\mu + \delta^{b'}\beta + \epsilon > 0 \quad (18)$$

The last two terms are not directly observable, as the identity of the new lender is unrecorded if no new loan is made. I thus estimate $\delta^{b'}\beta + \epsilon = \delta^{MAX}\beta + \gamma$ flexibly. The first term $\delta^{MAX}\beta$ represents a simple benchmark, where firms approach the least-affected lender for a new loan. The second term γ represents possible deviation from this benchmark, with γ to be estimated. This term reflects a composite error of different forces that can affect new relationship formation independent of the information gap: match-specific preferences for certain types of banks, search frictions that can lead borrowers to apply to lenders other than the least affected one, and any other measurement error in the second stage surplus. A negative γ implies a comparative disadvantage of new lenders relative to previous lenders overall. As discussed in Section 3, it is necessary that γ be uncorrelated with the shock to previous lenders δ^b . Below, I discuss three cases. The main specification estimates a constant γ for all firms. I show robustness to allowing it to vary with firm observable characteristics ($\gamma^f = z^f\Gamma$), and finally discuss the impact of i.i.d. noise.

The probability that a firm borrows from a new lender, conditional on its existing relationship being ended, is given by

$$\begin{aligned} \mathbb{P}(\text{borrow from a new lender}) &= \mathbb{P}(\nu_1^f \geq \nu^*(\delta^b, x^f, \mathcal{W}, \gamma) | \nu^f \leq \bar{\nu}(\delta^b, x^f)) \\ &= \int_{\nu^f \leq -\delta^b\beta - x^f\mu} \Phi\left(\frac{1}{\sqrt{\mathcal{W}}}(\nu^f - \nu^*(\delta^b, x^f, \mathcal{W}, \gamma))\right) \frac{\phi(\nu^f)}{1 - P_1^f} d\nu^f \end{aligned} \quad (19)$$

The information gap \mathcal{W} , as well as γ , is estimated via non-linear least squares regression:

$$(\hat{\mathcal{W}}, \hat{\gamma}) = \operatorname{argmin} \sum_f \left[\mathbb{1}(f \text{ borrows from a new lender}) - \mathbb{P}(\nu_1^f \geq \nu^*(\delta^b, x^f, \mathcal{W}, \gamma) | \nu^f \leq \bar{\nu}(\delta^b, x^f)) \right]^2 \quad (21)$$

The estimated information gap between lenders is 1.11%; detailed estimation results are presented in Table III. This gap is very small and implies that there is little private information among lenders in this market. New lenders thus appear to know as much as existing lenders about their borrowers. Recall that the information gap measures information *relative* to the econometrician, and \mathcal{W} is scaled by the variance of the unobservable ν^f . Thus, a magnitude close to zero reflects the fact that a large amount of information is unobservable to the econometrician but is known to

all lenders.

Table III also presents results for several alternative specifications. In particular, I reestimate the model using three different measures of bank exposure to the crisis that are used in prior literature. The first measure is the fraction of loans co-syndicated with Lehman Brothers before the crisis as introduced by Ivashina and Scharfstein (2010). The idea behind this measure is that lenders that had joint obligations with Lehman Brothers had to step in after its collapse, reducing their ability to finance new loans. The second measure is the loading of a bank's stock price on the mortgage-backed security ABX index as introduced by Chodorow-Reich (2014). The last measure is the bank ratio of real estate charge-offs to assets following the crisis computed with balance sheet data, in the spirit of Murfin (2012) and Chodorow-Reich (2014).

These different measures based on a variety of data sources help alleviate the concern that the main measure of bank health is contaminated by borrower characteristics. When I reestimate the model using each of these three measures instead, I find the estimated information gap is relatively stable. Compared to a baseline of $\mathcal{W} = 1.11\%$, the model using Lehman exposure estimates it to be 1.10%, the model using the ABX loading 2.34% and the model using real estate charge-offs 2.00%. These results confirm that there is virtually no information gap in this market.

As a further check, I also run an extended specification that estimates the information gap separately for private and public firms. Consistent with the idea that private firms are less transparent to outside financiers, I find that their information gap is 1.96% relative to 0.65% for public firms. Similarly, I estimate a gap of 2.87% for firms not included in the Compustat database versus 0.85% for others. Although these differences are just below statistical significance, the direction of the cross-sectional heterogeneity provides additional support that the estimate of \mathcal{W} accurately reflects informational frictions. Table AIV in the Appendix summarizes all cross-sectional results, including the sample splits based on firm size, pre-crisis syndicate structure and sector.

While the main specification estimates a constant intercept γ for all firms, Table III shows that the results are robust to allowing this term to vary with firm characteristics. In this specification, I assume that $\gamma = z^f \Gamma$, with the vector z^f including indicators for whether the firm is public, received multiple loans in the pre-crisis period, or is in the manufacturing sector, as well as an intercept. The estimated information gap is virtually unchanged at 1.09%, with firm characteristics having no significant effect on the composite error term γ .

Random Matching Errors: The specifications above assume that matching errors are a function of observable characteristics. It is possible however that unobservable noise in the second stage could drive the formation of new relationships. This would imply that the estimated model is misspecified: it assumes that unobservable matching errors are identically zero, instead of mean zero with some variance. Although misspecification bias is difficult to sign in general, it is a

legitimate concern in this setting as a number of potential factors other than the information gap drive the formation of new relationships. To investigate the performance of the estimator in the presence of unmodeled noise in ϵ , I conduct a series of Monte Carlo simulations.

Towards this end, I create a large number of artificial data sets, adding progressively more matching error noise. Each data set has two important characteristics: (i) the sample size is the same as that of the true data (to keep fixed the numerical properties of the estimation code) and (ii) it leads to the same first stage probit that the true data (the same cross-sectional distribution of renewal probabilities). These two properties are important because we want to isolate the performance of the estimation of \mathcal{W} in the second stage. More specifically, for each observation (x^f, δ^b) I draw three random variables: the two components of firm types ν_1 and ν_2 , normally distributed mean zero and standard deviation $\sqrt{1 - \mathcal{W}^2}$ and \mathcal{W} respectively; and a matching error ϵ , distributed normally with mean zero and standard deviation σ . I then simulate the model to obtain the joint probability of renewal and new relationships on which I can run the estimation routine. The simulation procedure is described in more detail in the Internet Appendix.

Figure A3 in the Appendix illustrates the results graphically for three low to moderately high levels of the true information gap (5%, 10% or 20%). I create 50 artificial data sets to smooth out some simulation noise, estimate the mode and report the median \mathcal{W} estimate, as well as the 25th and 75th percentiles. Overall, it appears that misspecification alone is unlikely to be able to explain a small value of \mathcal{W} . While there seem to be some misspecification bias when noise is added, this bias does not always have a negative sign and does not seem to grow fast with additional noise. The estimated information gap remains far from the 1% that I estimate in the data. Nevertheless, the potential for misspecification is an important caveat that must be kept in mind when interpreting the estimates.

C. Aggregate Effects of the Information Gap

The model estimates can be used to address the following counterfactual question: how many loans were not made after the crisis because of the information gap and imperfect reallocation? Table IV reports the effects on aggregate lending of assuming that all lenders have the same information, that is assuming that $\mathcal{W} = 0$. More precisely, in this counterfactual, new lenders apply the same lending rule used by informed lenders, estimated in the first stage. Formally, the counterfactual probability that firm f borrows in the second stage is given by:

$$\mathbb{P}_{\mathcal{W}=0}(f \text{ borrows from a new lender}) = \mathbb{P}(\nu_1^f \geq \bar{\nu}(\delta^b, x^f) | \nu_1^f \leq \bar{\nu}(\delta^b, x^f)) \quad (22)$$

Table III
Second-Stage Estimation Results: Information Gap.

	Information Gap (\mathcal{W})	Intercept γ
Main Model	1.11% [0.75%, 2.94%]	-0.29 [-0.47, -0.13]
<i>Other Bank Shocks</i>		
Lehman Exposure	2.28% [1.32%, 2.93%]	-0.11 [-0.22, -0.03]
ABX Loading	2.34% [1.15%, 5.15%]	-0.06 [-0.18, -0.00]
RE Chargeoffs	1.99% [1.03%, 4.47%]	0.00 [-0.50, -0.00]
<i>Other Extensions</i>		
Heterogeneous comp. advantage $\gamma = z^f \Gamma$	1.09% [1.05%, 2.26%]	-0.28 (constant) [-0.49, -0.13]
		-0.006 (public) [-0.04, 0.10]
		-0.016 (multiple pre-crisis loans) [-0.09, 0.05]
		-0.004 (manufacturing) [-0.08, 0.06]
N. observations	3,188	

Coefficients are estimated via nonlinear least squares on the subsample of firms that did not renew their relationship after the crisis. 95% confidence intervals are bootstrapped to account for the fact that the first stage is estimated. A borrower is classified as renewing its relationship if at least one lead lender of its post-crisis lending syndicate was a lead lender of its last pre-crisis syndicate. The sample is restricted to U.S. nonfinancial firms that list the reason for borrowing as "working capital" or "corporate purposes."

where the informed lender decision rule $\bar{\nu}$ is estimated in the first stage. For ease of exposition, I measure the extensive margin response of lending by the share of firms with a loan in 2004 to 2008 that received a new loan in 2008 to 2010.

Consistent with the very small estimate of the information gap above, Table IV reveals that the effect on aggregate lending is marginal. Counterfactual lending in the absence of the information gap is virtually identical to the data: the difference is 0.31 percentage points, or about 1% in relative terms. This magnitude is similar when using other specifications. All three alternative measures of bank shocks imply a counterfactual effect below one percentage point. Introducing a comparative advantage that varies with firm characteristics yields nearly identical results. If anything, these results point to slight underfunding, as Section 2 shows is possible when \mathcal{W} is very small, although the magnitudes are extremely close to zero.

D. *Bias from Ignoring Common Information*

The information ν_1^f that is common to all lenders but unobservable to the econometrician leads to bias when estimating the information gap; this is the key reason behind using the identification strategy introduced in this paper. To make this point clear, I estimate a "naive" model that ignores this common information. This naive model assumes that the distribution of firms that see their relationship end is independent of pre-crisis lender exposure, that is $\nu_1^f \perp \delta^b$. In other words, it does not condition on selection from the first stage when estimating second stage new relationship probabilities. More precisely,

$$\mathbb{P}_{\text{naive}}(\text{borrow from a new lender}) = \mathbb{P}(\nu_1^f \geq \nu^*(\delta^b, x^f, \mathcal{W}, \gamma)) \quad (23)$$

$$= 1 - \Phi\left(\frac{\nu^*(\delta^b, x^f, \mathcal{W}, \gamma)}{\sqrt{1 - W^2}}\right) \quad (24)$$

The information gap \mathcal{W} , as well as γ , is estimated via non-linear least squares regression:

$$(\hat{\mathcal{W}}, \hat{\gamma}) = \operatorname{argmin} \sum_f [\mathbb{1}(f \text{ borrows from a new lender}) - \mathbb{P}_{\text{naive}}(\delta^b, x^f, \mathcal{W}, \gamma)]^2 \quad (25)$$

Table V shows that the naive model dramatically overestimates the information gap and its effects. The naive estimate of \mathcal{W} is over 50 times larger, at a very large 63%. Accordingly, estimates of the aggregate effects are accordingly very large, over fifteen times that of the main specification. The naive model suggests dramatic underfunding during the crisis, while the main model points to, if anything, marginal overfunding. This large bias demonstrates the value of the empirical strategy introduced in this paper, and the necessity of controlling for common information ν_1 observed by

Table IV
Aggregate Effects of the Information Gap.

	Crisis Period Outcome		
	% firms receiving a loan (model fit)	Counterfactual with no information gap	Difference (pp)
Data	25.05%	-	-
Main Model	25.06% [23.60%, 26.06%]	24.75% [23.10%, 25.71%]	-0.31pp [-0.72,-0.21]
<i>Other bank shocks</i>			
Lehman Exposure	24.87% [23.24%, 26.25%]	24.14% [22.26%, 25.56%]	-0.73pp [-0.93,-0.31]
ABX Loading	24.98% [23.59%, 26.11%]	24.17% [22.37%, 25.43%]	-0.81pp [-1.61, -0.36]
RE Chargeoffs	24.94% [21.40%, 25.86%]	24.29% [20.94%, 25.28%]	-0.65pp [-1.59, 0.00]
Heterogenous comp. advantage γ	25.12% [23.26, 25.94%]	24.81% [22.93%, 25.58%]	-0.31pp [-0.56, -0.24]
N. observations	4,044	4,044	

The outcome variable is the share of firms with a loan in 2004 to 2008 that received a new loan in 2008 to 2010. In the main model, the crisis exposure is the firm's lenders pre-crisis average of the relative fall in lending between 2004 to 2008 and 2008 to 2010 of crisis lending syndicate, weighted by the loan share of each lender. 95% confidence intervals are bootstrapped to account for the fact that the first stage is estimated. Lehman exposure is defined as the fraction of loans co-syndicated with Lehman Brothers before the crisis, ABX loading is the loading of bank stock price on the mortgage-backed security ABX index, and real estate charge-offs is the ratio of real estate charge-offs to assets following the crisis. These last three bank variables can be found on Gabriel Chodorow-Reich's website. For each model, counterfactual lending is computed by assuming that the information gap \mathcal{W} is equal to zero. The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes."

Table V

Estimates from Naive Model Ignoring the Bias from Common Information ν_1 .

	Main model	Naive model
Information gap	1.11%	63.36%
% firms receiving a loan (model fit)	25.06%	24.34%
Counterfactual with no information gap	24.75%	30.27%
Difference	-0.31pp	5.52pp
N. observations	4,044	4,044

Note: The outcome is the share of firms with a loan in 2004 to 2008 that received a new loan in 2008-2010. The naive model ignores the fact that only a subset of firms reaches the second stage. For each model, counterfactual lending is computed by assuming that the information gap \mathcal{W} is equal to zero. The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes."

all lenders but unobservable to the econometrician. The naive model attributes the entire cross-sectional variation from Figure 6 to the information gap, while in reality, most of the variation is due to selection on common information.

The main takeaway is that the information gap is likely too small to explain relationship stickiness in the U.S. syndicated loans market. Quantitatively, the information gap cannot account for a significant fraction of the reduced-form effect of bank shocks documented in Chodorow-Reich (2014). These estimates reveal that virtually all the reduced-form patterns consistent with the "inference hypothesis" are driven by information common to all lenders but unobservable to the econometrician. This is plausible as the syndicated loan market is dominated by large banks and among the most transparent firms. It is not obvious what could constitute soft information that is known to existing lenders but unobservable to other lenders.³⁴

Therefore results suggest that there must exist a different friction behind stickiness as there is clear reduced-form evidence that bank shocks matter in this market. A particularly plausible friction is the covenant channel recently documented by Chodorow-Reich and Falato (2019). They show that banks subject to worse shocks are more likely to act on covenant violations and push

³⁴Nevertheless, the predominant view in empirical works on this market is to invoke the information gap as an important mechanism behind stickiness. Examples include Santos and Winton (2008), Santos (2011), Saunders and Steffen (2011), Schenone (2010), Ferreira and Matos (2012), Murfin (2012), Bharath et al. (2011) or Dass and Massa (2011).

their borrowers into technical default. In turn, this debt overhang makes it difficult to find a new lender. They estimate this covenant channel to be large. Their results thus add external validity to my finding that the bulk of the aggregate effects of bank shocks in the syndicated loan market do not come from the information gap. Another source of friction could stem from imperfect matching between borrowers and lenders. Boualam (2018) highlights how the destruction of credit relationships and their slow build-up after an adverse aggregate shock can generate slow subsequent recoveries. The evidence in Schwert (2018) shows that borrower-lender matching is quantitatively important to explain the amount of lending during the recent crisis. This result suggests that reducing matching frictions could have further prevented the reduction in lending.

E. Implications for Policy

Understanding the nature of the reallocation frictions, whether informational or not, is important from a policy perspective. In particular, during a crisis, policy makers often implement *targeted* interventions that aim to provide public support for the most-affected lenders. For instance, the Capital Purchase Program, which was part of the Troubled Asset Relief Program authorized by Congress in the fall of 2008, provided over \$200 billion in equity injections to many large U.S. financial institutions. Institutions receiving CPP funds were among the most affected by the crisis according to my measure or purchased some of the most-affected banks. A key question is how these targeted interventions impact aggregate lending.³⁵

If the information gap is large, a potential concern is that these interventions can have unintended consequences on credit reallocation that reduce their effectiveness. More specifically, direct support for the weakest lenders has two distinct effects. The first is a positive *bank-level* effect: the share of firms able to renew their relationship with these lenders increases. However, this intervention hurts those borrowers that are not able to renew their relationship: by amplifying stigma, it dilutes the positive signal that comes from new lenders' inference. This logic is related to the models of Uhlig (2010) and Malherbe (2014). However, if the underlying reallocation friction is not informational in nature, such negative equilibrium effects are absent, and supporting weak banks is more effective. This is another reason why a new approach to estimating the magnitudes of informational frictions is valuable. The results of this paper suggest that intervention should be mostly positive for large firms but in segments of the population of firms with greater informational frictions, the unintended negative effects could be more important.³⁶

³⁵These interventions also often have multiple objectives, such as preventing bank runs or domino effects.

³⁶I thank an anonymous referee for suggesting this cross-sectional implication.

F. Generalizability and Other Applications

The approach of this paper is well-suited to study the effect of the large credit shock that occurred during the 2007-2009 financial crisis. Some the key requirements for this empirical methodology to be used in other contexts are as follows. First, this approach requires loan-level data to trace relationships and specifically distinguish between stayers and switchers. Second, it relies on credit supply shocks that affect banks' propensity to lend independently of borrower characteristics. Finally, the power of the estimation approach is driven by the number of switchers in the data: the numbers of firms that do not renew their relationship after a shock, but are able to find a new lender. This is natural given that the focus is to estimate the magnitude of a switching friction.

However, relationships tend to be sticky and therefore the share of switchers is often small relative to total borrowers. There are two cases in which there is a sufficient sample size of switchers. First, large shocks, such as the Great Recession, can induce a large amount of credit reallocation. Second, smaller shocks in settings with comprehensive loan-level data with the near universe of borrowers, akin to the credit registries maintained in some countries. In the U.S., until the recent Dodd-Frank act loan-level data for corporate loans were limited to the syndicated loan market dominated by large borrowers. In other countries, smaller credit supply shocks could in principle be used as sample sizes tend to be an order of magnitude or two larger. Future research could potentially apply a similar empirical approach in order to determine, for example, how these frictions vary across market segments or the business cycle.

IV. Conclusion

This paper estimates the effect of an informational switching friction in explaining the credit crunch during the 2007 to 2009 crisis. Lenders have private information about their borrowers and this information gap makes switching lenders difficult. At the same time, there exists common information that all lenders can observe but the econometrician cannot. I argue that this information hierarchy implies that reduced-form estimates of the information gap are biased and I show that naive information models dramatically overestimate the effects of this friction on lending. In this paper, I address this empirical challenge by estimating a discrete-choice model of relationships with three explicit layers of information. I show how to use bank shocks to identify this private information separately from information common to all lenders. The model estimates can also be used for quantification: beyond documenting its existence, how much does the friction matter? There are many settings in which economists have a reasonable idea of which frictions are at play,

but have little sense of which is the most relevant empirically.

I find that lenders' private information is too small to explain why relationships are sticky in the U.S. syndicated corporate loan market and therefore it cannot quantitatively account for the associated drop in lending documented in prior studies. This result suggests that different frictions behind stickiness can quantitatively drive the effects of bank shocks. A plausible alternative in this market is the covenant channel recently documented by Chodorow-Reich and Falato (2019). Of course, the result above does not imply that the information gap is not a relevant friction in other markets. In fact, given the types of large borrowers and lenders active in the syndicated loan market, one might expect that it represents a lower bound for the magnitude of informational frictions. Future research could potentially study these frictions in other settings.

Initial submission: December 20th 2017 Accepted: March 27th 2019

Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

References

- Adelino, Manuel, Kristopher Gerardi, and Barney Hartman-Glaser, 2019, Are lemons sold first? Dynamic signaling in the mortgage market, *Journal of Financial Economics* 132, 1–25.
- Agarwal, Sumit, and Robert Hauswald, 2010, Distance and private information in lending, *The Review of Financial Studies* 23, 2757–2788.
- Berger, Allen N, and Gregory F Udell, 1995, Relationship lending and lines of credit in small firm finance, *Journal of Business* 351–381.
- Bharath, Sreedhar T, Sandeep Dahiya, Anthony Saunders, and Anand Srinivasan, 2011, Lending relationships and loan contract terms, *The Review of Financial Studies* 24, 1141–1203.
- Botsch, Matthew, and Victoria Vanasco, 2019, Learning by lending, *Journal of Financial Intermediation* 37, 1–14.
- Boualam, Yasser, 2018, Credit markets and relationship capital, Technical report, Working paper, University of North Carolina.
- Chaney, Thomas, David Sraer, and David Thesmar, 2012, The collateral channel: How real estate shocks affect corporate investment, *American Economic Review* 102, 2381–2409.
- Chen, Jiawei, and Kejun Song, 2013, Two-sided matching in the loan market, *International Journal of Industrial Organization* 31, 145–152.
- 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, 1–59.
- Chodorow-Reich, Gabriel, and Antonio Falato, 2019, The loan covenant channel: How bank health transmits to the real economy .
- Crawford, Gregory S, Nicola Pavanini, and Fabiano Schivardi, 2018, Asymmetric information and imperfect competition in lending markets, *American Economic Review* 108, 1659–1701.
- Dass, Nishant, and Massimo Massa, 2011, The impact of a strong bank-firm relationship on the borrowing firm, *The Review of Financial Studies* 24, 1204–1260.
- D’Auria, Claudio, Antonella Foglia, and Paolo Marullo Reedtz, 1999, Bank interest rates and credit relationships in italy, *Journal of Banking & Finance* 23, 1067–1093.

- Degryse, Hans, and Patrick Van Cayseele, 2000, Relationship lending within a bank-based system: Evidence from european small business data, *Journal of Financial Intermediation* 9, 90–109.
- Dell’Ariccia, Giovanni, and Robert Marquez, 2004, Information and bank credit allocation, *Journal of Financial Economics* 72, 185–214.
- Detragiache, Enrica, Paolo Garella, and Luigi Guiso, 2000, Multiple versus single banking relationships: Theory and evidence, *Journal of Finance* 55, 1133–1161.
- Einav, Liran, Mark Jenkins, and Jonathan Levin, 2012, Contract pricing in consumer credit markets, *Econometrica* 80, 1387–1432.
- Farrell, Joseph, and Paul Klemperer, 2007, Coordination and lock-in: Competition with switching costs and network effects, *Handbook of industrial organization* 3, 1967–2072.
- Ferreira, Miguel A, and Pedro Matos, 2012, Universal banks and corporate control: Evidence from the global syndicated loan market, *The Review of Financial Studies* 25, 2703–2744.
- Gibbons, Robert, and Lawrence F Katz, 1991, Layoffs and lemons, *Journal of Labor Economics* 9, 351–380.
- Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen, 2020, Do credit market shocks affect the real economy? quasi-experimental evidence from the great recession and "normal" economic times, *American Economic Journal: Economic Policy* 12, 200–225.
- Heckman, James, 2013, Sample selection bias as a specification error, *Econometrica* 47, 153–161.
- Hertzberg, Andrew, Andres Liberman, and Daniel Paravisini, 2018, Screening on loan terms: Evidence from maturity choice in consumer credit, *The Review of Financial Studies* 31, 3532–3567.
- Hertzberg, Andrew, Jose Maria Liberti, and Daniel Paravisini, 2010, Information and incentives inside the firm: Evidence from loan officer rotation, *Journal of Finance* 65, 795–828.
- Hombert, Johan, and Adrien Matray, 2017, The real effects of lending relationships on innovative firms and inventor mobility, *The Review of Financial Studies* 30, 2413–2445.
- Hoshi, Takeo, Anil Kashyap, and David Scharfstein, 1991, Corporate structure, liquidity, and investment: Evidence from Japanese industrial groups, *The Quarterly Journal of Economics* 106, 33–60.

- Ivashina, Victoria, and David Scharfstein, 2010, Bank lending during the financial crisis of 2008, *Journal of Financial Economics* 97, 319–338.
- Iyer, Rajkamal, Asim Ijaz Khwaja, Erzo FP Luttmer, and Kelly Shue, 2016, Screening peers softly: Inferring the quality of small borrowers, *Management Science* 62, 1554–1577.
- Jiménez, Gabriel, Atif Mian, José-Luis Peydró, and Jesús Saurina, 2019, The real effects of the bank lending channel, *Journal of Monetary Economics* .
- Karolyi, Stephen Adam, 2018, Personal lending relationships, *Journal of Finance* 73, 5–49.
- Keys, Benjamin J, Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, 2010, Did securitization lead to lax screening? Evidence from subprime loans, *The Quarterly Journal of Economics* 125, 307–362.
- Khwaja, Asim Ijaz, and Atif Mian, 2008, Tracing the impact of bank liquidity shocks: Evidence from an emerging market, *American Economic Review* 98, 1413–42.
- Korteweg, Arthur G, Michael Schwert, and Ilya A Strebulaev, 2019, Proactive capital structure adjustments: Evidence from corporate filings .
- Kurlat, Pablo, and Johannes Stroebel, 2015, Testing for information asymmetries in real estate markets, *The Review of Financial Studies* 28, 2429–2461.
- Levin, Jonathan, 2003, Relational incentive contracts, *American Economic Review* 93, 835–857.
- Liberti, Jose M, and Atif R Mian, 2009, Estimating the effect of hierarchies on information use, *The Review of Financial Studies* 22.
- Malherbe, Frédéric, 2014, Self-fulfilling liquidity dry-ups, *Journal of Finance* 69, 947–970.
- Mian, Atif, 2006, Distance constraints: The limits of foreign lending in poor economies, *Journal of Finance* 61, 1465–1505.
- Murfin, Justin, 2012, The supply-side determinants of loan contract strictness, *Journal of Finance* 67, 1565–1601.
- Peek, Joe, and Eric S Rosengren, 2000, Collateral damage: Effects of the Japanese bank crisis on real activity in the united states, *American Economic Review* 90, 30–45.
- Petersen, Mitchell A, and Raghuram G Rajan, 1994, The benefits of lending relationships: Evidence from small business data, *Journal of Finance* 49, 3–37.

- Petersen, Mitchell A, and Raghuram G Rajan, 1995, The effect of credit market competition on lending relationships, *The Quarterly Journal of Economics* 110, 407–443.
- Rajan, Raghuram G, 1992, Insiders and outsiders: The choice between informed and arm’s-length debt, *Journal of Finance* 47, 1367–1400.
- Rajan, Uday, Amit Seru, and Vikrant Vig, 2015, The failure of models that predict failure: Distance, incentives, and defaults, *Journal of Financial Economics* 115, 237–260.
- Roberts, Michael R, and Amir Sufi, 2009, Renegotiation of financial contracts: Evidence from private credit agreements, *Journal of Financial Economics* 93, 159–184.
- Ruckes, Martin, 2004, Bank competition and credit standards, *The Review of Financial Studies* 17, 1073–1102.
- Santos, João AC, 2011, Bank corporate loan pricing following the subprime crisis, *The Review of Financial Studies* 24, 1916–1943.
- Santos, Joao AC, and Andrew Winton, 2008, Bank loans, bonds, and information monopolies across the business cycle, *Journal of Finance* 63, 1315–1359.
- Saunders, Anthony, and Sascha Steffen, 2011, The costs of being private: Evidence from the loan market, *The Review of Financial Studies* 24, 4091–4122.
- Schenone, Carola, 2010, Lending relationships and information rents: Do banks exploit their information advantages?, *The Review of Financial Studies* 23, 1149–1199.
- Schwert, Michael, 2018, Bank capital and lending relationships, *Journal of Finance* 73, 787–830.
- Sharpe, Steven A, 1990, Asymmetric information, bank lending, and implicit contracts: A stylized model of customer relationships, *Journal of Finance* 45, 1069–1087.
- Srinivasan, Anand, et al., 2014, Long run relationships in banking, *Foundations and Trends® in Finance* 8, 55–143.
- Stroebel, Johannes, 2016, Asymmetric information about collateral values, *Journal of Finance* 71, 1071–1112.
- Sutherland, Andrew, 2018, Does credit reporting lead to a decline in relationship lending? Evidence from information sharing technology, *Journal of Accounting and Economics* 66, 123–141.
- Uhlig, Harald, 2010, A model of a systemic bank run, *Journal of Monetary Economics* 57, 78–96.

Appendix: Additional Results

Table AI

Borrower Summary Statistics.

	N	Mean	Std. dev.	p25	p50	p75
Public (indicator)	4044	.3652324	.4815548	0	0	1
Loan due during crisis (indicator)	4044	.3872404	.4871796	0	0	1
Manufacturing sector (indicator)	4044	.3325915	.4711998	0	0	1
Pre-crisis spread (bps)	4044	181.2427	124.4589	87.5	166.6667	250
Pre-crisis loan size (\$M)	4044	481.7738	976.1442	125	228.5714	467.0833
Pre-crisis maturity (days)	4044	1545.399	555.9556	1096	1826	1827
Pre-crisis loan secured (indicator)	4044	.3632542	.4809967	0	0	1

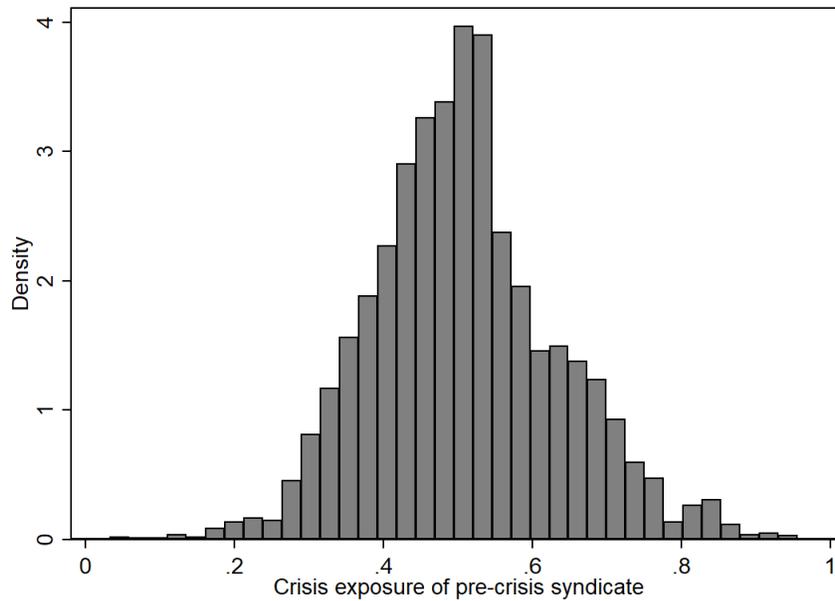
Summary statistics for the sample used in estimation.

Table AII

Change in Loan Terms After the Crisis.

	Pre-crisis (Jan 04-Sept 08)		Post-crisis (Oct 08-Dec 10)		% change in mean
	Mean	Std. Dev.	Mean	Std. Dev.	
Loan size (\$M)	482	976	470	818	-2.5%
Spread (bp)	181	124	294	128	62.4%
Maturity (years)	4.2	1.5	3.2	1.3	-23%
#Lenders in syndicate	5.1	3.8	5.1	3.7	0%

This table includes only firms that received a loan in the pre-crisis period. The crisis loan is the first new loan received in the crisis period (if any). The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes."



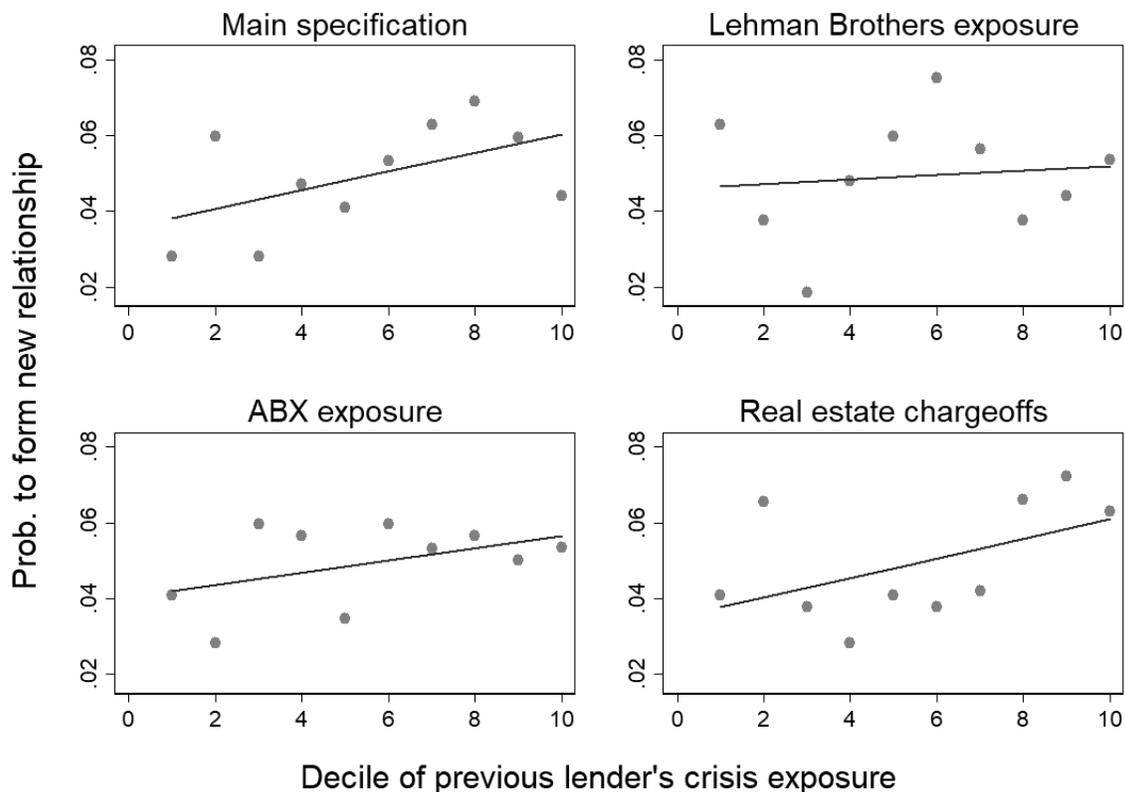
The crisis exposure of a firm's pre-crisis lender is computed as the weighted average of the relative decrease in lending between 2004 to 2008 and 2008 to 2010 of each lender in the firm's last pre-crisis lending syndicate, weighted by the loan shares of each lender. The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes."

Figure A1. Distribution of Firms' Pre-Crisis Lender Crisis Exposure

Table AIII**The Inference Hypothesis: Reduced-Form Regressions.**

	Outcome: Borrow from a new lender, conditional on not borrowing from pre-crisis lender		
	(1)	(2)	(3)
Pre-crisis lender exposure	0.51 (0.34)	0.64* (0.35)	0.77** (0.40)
Firm characteristics	No	Yes	Yes
Pre-crisis loan characteristics	No	No	Yes
N	3,188	3,188	3,188
Pseudo R-squared	0.15%	1.8%	4.2%

Probit regression results. Reported coefficients are marginal effects multiplied by 100. A borrower is classified as borrowing from its pre-crisis lender if at least one lead lender of its post-crisis lending syndicate was a lead lender of its last pre-period syndicate. The crisis exposure of a firm's pre-crisis lender is computed as the weighted average of the relative fall in lending between 2004 to 2008 and 2008 to 2010 of each lender in the firm's last pre-crisis lending syndicate, weighted by the loan share of each lender. Firm characteristics include: public ownership, high sales, manufacturing sector, received multiple pre-crisis loans, has a pre-crisis loan that extends throughout the crisis period. Pre-crisis loan characteristics include: spread, size, whether it was secured by collateral, and whether there were multiple lead lenders or two or fewer participants in its syndicate. The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes." *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.



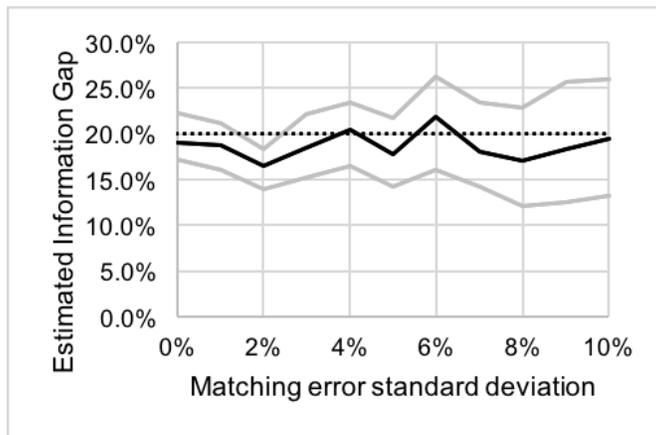
This sample includes only borrowers who did not renew their previous relationship after the crisis. A borrower is classified as renewing a relationship if at least one lead lender of its post-crisis lending syndicate was a lead lender of its last pre-period syndicate. Other borrowers receiving a new loan after the crisis are classified as forming a new relationship. In the main specification, the crisis exposure of a firm's pre-crisis lender is computed as the weighted average of the relative drop in lending between 2004 to 2008 and 2008 to 2010 of each lender in the firm's last pre-crisis lending syndicate, weighted by the loan shares of each lender. Alternative for bank crisis exposure: fraction of loans co-syndicated with Lehman, stock price loading on ABX index, real estate charge-offs over assets, as reported on Gabriel Chodorow-Reich's website. The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes."

Figure A2. Pre-Crisis Lender Crisis Exposure and the Formation of New Relationships: Other Bank Shocks

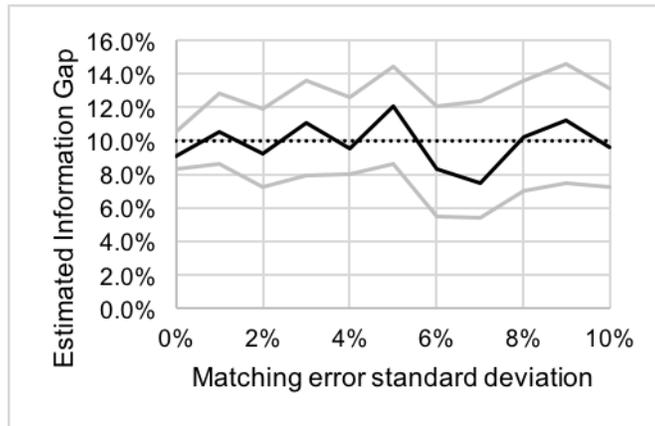
Table AIV
Information Gap by Sub-samples

	Information Gap	N. observations
Baseline (all firms)	1.16%	4,044
<i>By size quartile (pre-crisis loan size)</i>		
Q1	3.10%	1,041
Q2	3.74%	985
Q3	4.24%	1,007
Q4	0.91%	1,011
<i>By pre-crisis syndicate size</i>		
Larger than 4	2.22%	2,184
Smaller than 4	1.10%	1,860
<i>By Compustat status</i>		
In Compustat	0.85%	2,134
Not in Compustat	2.87%	1,910
<i>By public status</i>		
Public	0.65%	1,477
Private	1.90%	2,567
<i>By sector</i>		
Manufacturing	1.32%	1,345
Non-manufacturing	3.36%	2,699

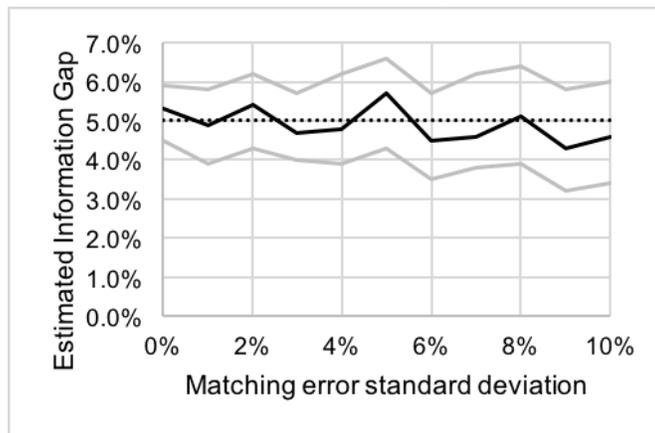
For each sample, coefficients are estimated via non-linear least squares on the subsample of firms that did not renew their relationship after the crisis. A borrower is classified as renewing its relationship if at least one lead lender of its post-crisis lending syndicate was a lead lender of its last pre-crisis syndicate. The sample is restricted to U.S. non-financial firms that list the reason for borrowing as "working capital" or "corporate purposes."



Panel A. Information Gap of 20%



Panel B. Information Gap of 10%



Panel C. Information Gap of 5%

Figure A3. Monte Carlo Simulations with Random Matching Errors.