

Recent Developments in Consumer Credit and Payments*

BY MITCHELL BERLIN

On September 24-25, 2009, the Research Department and the Payment Cards Center of the Federal Reserve Bank of Philadelphia held their fifth joint conference to present and discuss the latest research on consumer credit and payments. Sixty participants attended the conference, which included seven research papers on topics such as securitization and distressed loan renegotiation, consumer disclosure, data breaches and identity theft, and the effects of the U.S. financial crisis on global retail lending. In this article, Mitchell Berlin summarizes the papers presented at the conference.

In his opening remarks, Mitchell Berlin noted that the flow of high-quality papers on consumer finance and payment issues has increased steadily since the first joint conference of the Research Department and Payment Cards Center in 2001.

SECURITIZATION AND RENEGOTIATION

In the first paper, Tomasz Piskorski of Columbia University reported the results of a study (with Amit

Seru and Vikrant Vig) that provided evidence that frictions impeded the renegotiation of certain types of distressed mortgages during the recent financial crisis.¹ In particular, Piskorski and his coauthors showed that loans that banks packaged into mortgage-backed securities and placed in trusts — securitized loans — were foreclosed on more often than otherwise similar mortgages that remained in bank portfolios.

¹The conference agenda along with links to most of the papers presented can be found on the Philadelphia Fed's website at: <http://www.philadelphiafed.org/research-and-data/events/2009/consumer-credit-and-payments/program.cfm>.

*The views expressed here are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

The authors examined a large sample (about 327,000) of first-lien, nonagency mortgages originated between 2005 and 2006. This sample is drawn from the LPS database, which includes both loans held in the originating bank's portfolio and loans that were securitized. The sample was restricted to distressed mortgages, that is, loans that were at least 60 days' delinquent. They also considered a subsample of high-quality loans: loans with full documentation (full doc) with credit scores above 680.

First, the authors presented descriptive statistics showing that distressed portfolio loans were foreclosed less often than securitized loans. Then they used a logit model to estimate the probability of foreclosure, depending on whether the loan was securitized or held in portfolio, but also taking account of loan and borrower characteristics that might affect this probability. These included the borrower's credit score, the loan-to-value (LTV) ratio, the size of the loan at origination, the loan maturity, and the original interest rate, among other variables.

The logit results provided evidence that delinquent portfolio loans were less likely to be foreclosed and that the effect was both statistically significant and economically large. Evaluated at the mean values for all of the other variables, the probability of being foreclosed was between 3.8 and 7 percentage points lower for portfolio loans than for securitized loans, depending on the precise model specification. These corresponded to an 18 percent and 32 percent relative decline in the mean foreclosure rate.

Piskorski and his coauthors



Mitchell Berlin is a vice president and economist and head of the Banking section in the Philadelphia Fed's Research Department. This article is available free of charge at

www.philadelphiafed.org/research-and-data/publications/.

sought to address concerns that lenders may have learned information about borrower quality subsequent to origination but prior to delinquency. If this were true, the sample of securitized loans might disproportionately include loans that the originating bank had sold after receiving information, suggesting a higher probability of default. First, they replicated their basic results using a subsample of high-quality loans, arguably a set of loans for which it is less likely that underwriters might have subsequently learned more information about borrower quality. They also addressed this concern more directly by examining a subsample of loans for which borrowers' credit scores or LTVs at the time of delinquency were available. In these tests, the authors continued to find that portfolio loans were less likely to be foreclosed than securitized loans.

The authors also addressed concerns that their results were driven by differences in variations in state laws that affect the ease of foreclosure. Specifically, they divided states into tough and weak states, depending on whether average foreclosure times were low or high. They found that securitized loans were significantly more likely to be foreclosed in both tough and weak states, a finding inconsistent with the view that differences in state laws were the source of their results.

The authors also considered the possibility that the differences in foreclosure rates were not driven by a bias against renegotiation for securitized loans but by a bias against foreclosures by banks postponing recognizing losses on portfolio loans. Piskorski and coauthors found that delinquent borrowers resumed making payments on loans held in portfolio at a higher rate than for securitized loans, a finding inconsistent with this view.

LIAR'S LOANS

Ashlyn Aiko Nelson of Indiana University discussed the results of a study (with Wei Jiang and Edward Vytlačil) that explored the effects on loan performance of origination channel and level of documentation for mortgage loans. Nelson and her coauthors found that loans originated by brokers and loans that required little or no documentation by the borrower (low doc loans) were particularly

Piskorski and coauthors found that delinquent borrowers resumed making payments on loans held in portfolio at a higher rate than for securitized loans.

prone to agency problems. Nelson argued that broker-originated loans performed badly because brokers had incentives to make loans to low-quality borrowers, while low doc loans were more likely to perform badly because borrowers overstated income.

Nelson and her coauthors examined the mortgage loans made between January 2004 and February 2009 by a large national mortgage bank. The bank's files contained a wealth of data about borrower characteristics, permitting the authors to take account of many borrower-specific factors that might affect loan performance. The authors divided their sample into four subsamples: loans originated by brokers requiring full documentation, loans originated by brokers requiring low or no documentation, loans originated by banks requiring full documentation,

and loans originated by brokers requiring low or no documentation.

Nelson explained that the loan sample was not representative of the mortgage loan market as a whole, which raised some questions about the extent to which the results could be generalized. The bank made a disproportionately large share of loans originated by brokers and a disproportionately large share of low doc loans. Nonetheless, the authors' view was that this bank represented an extreme example of tendencies that were common to many banks.

In their main econometric treatment, the authors used two models of loan delinquency: (1) a probit model, which estimates the probability of delinquency, taking into account loan and borrower characteristics; (2) a duration model, which estimates the average time to delinquency, taking into account the same characteristics. They estimated separate regressions for each of the four subsamples and found that both the origination channel and the level of documentation affected delinquency. Specifically, they found that bank-originated full doc loans had a delinquency rate of 13 percent and a five-year survival rate of 86 percent, while the comparable numbers for bank-originated low doc loans were 18 percent and 68 percent; broker-originated full doc loans were 24 percent and 65 percent; and broker-originated low doc loans were 32 percent and 36 percent.

Nelson explained that particular types of borrowers might have selected into particular types of loans. To address this issue, the authors estimated a model to explain which types of borrowers chose particular types of loans. Broadly, they found that borrowers using brokers had lower credit quality and less experience with mortgages. In contrast, in addition to self-employed borrowers, low doc borrowers were

typically more experienced and of higher credit quality.

The authors then sought to examine the separate effects of observable and unobservable differences in borrower risk on delinquency. According to the authors, many of the low doc loans looked good on paper but did poorly, suggesting that the poor performance of low doc loans was due to factors that did not appear on the loan applications. In contrast, more than half of the poor performance of broker loans was due to observable borrower characteristics.

To explain the poor performance of low doc loans, Nelson and coauthors explored the evidence for falsified loan application information. In one of their approaches to this question, they hypothesized that falsified application information would reduce the predictive power of the empirical model. To create an out-of-sample test, they estimated the model over six-month periods to predict the probability of delinquency for the subsequent six-month period. The model's predictive power was substantially lower for the low doc subsamples, a result consistent with falsified application information.

The authors also performed a more direct test for falsified application information in low doc loans. They identified the particular borrower attributes most likely to be falsified: whether the home was a primary residence, employment information, and information about income, wealth, and existing debt. The results for income provided the most striking evidence of overstatement.

In the full doc samples, income was negatively related to delinquency. However, in the low doc samples, stated income was positively related to delinquency, and the effect was strongest when the loan had been originated by brokers without an ongoing relationship with the bank. The authors also

attempted to quantify the extent of the overstatement, comparing stated income with alternative measures of the borrower's true income, for example, average income in the borrower's ZIP code. They found that the ratio of stated income to average income in the ZIP code was significantly higher for the low doc sample and estimated that income was overstated by between 15 and 20 percent.

[Borrowers] may focus on the cost of borrowing once, even though they are likely to borrow repeatedly.

DISCLOSURE AND PAYDAY LENDING

Adair Morse of the University of Chicago reported the results of a carefully designed field study (with Marianne Bertrand) that attempted to determine whether payday borrowers' decisions were affected by particular types of cognitive bias and whether their borrowing decisions might be affected by particular types of disclosures at the point of the transaction.

In a payday loan, borrowers sign over their next paycheck and pay a fee (\$15-17 for each \$100 borrowed) in exchange for a loan. Payday loans are quite expensive compared with other types of loans, with annual percentage rates (APRs) typically exceeding 400 percent, and customers typically borrowing repeatedly. According to the authors, one explanation for why borrowers use such an expensive source of borrowed funds is cognitive bias. For example, borrowers may not realize how high the APR on the payday loan is, if they don't compare it to relevant alternatives. Alternatively, they may focus on the cost of borrowing once, even though they are likely to borrow repeatedly.

The authors considered three different disclosures. In one, the APR on a payday loan was compared with the APRs on other types of loans likely to be familiar to payday borrowers. The second disclosure displayed the dollar cost of repeated borrowings up to three months and compared this with the dollar cost of repeated borrowings on a credit card. The third disclosure illustrated the likelihood that a

typical payday borrower will engage in repeated borrowings, e.g., how many borrow once, how many renew once or twice, etc.

The goal of the experiment was to determine whether a disclosure reduced either the likelihood or the amount of subsequent borrowings. Furthermore, Morse and her coauthor examined whether the disclosures worked for particular types of borrowers, for example, whether a borrower's level of education or some measure of the borrower's degree of self-control might affect the outcome of a particular disclosure.

The authors performed this experiment with the cooperation of a large payday lender operating in 10 states. They had access to all customers entering 77 stores over a two-week period. Crucially, the experiment was designed so that the disclosures (including no disclosure at all) were randomly assigned over borrowers. The goal was to minimize the possibility that factors other than the actual disclosures might affect borrower behavior; for example, the disclosures were equally distributed across different days of the week because a Monday borrower might differ from a Thursday borrower. Morse re-

ported that subsequent empirical tests by her and her coauthor verified that the experimental design had successfully randomized across customers.

The lender's records also included demographic data about each borrower, for example, level of education; financial data; and information about past and subsequent transactions with the customer, provided by the lender. In addition to these data, participating borrowers also answered survey questions about the intended use of the loan and the borrower's own view of his or her planning and spending habits. Using the survey answers, the authors designed an index of borrower self-control and a gratification index measuring whether the loan was for discretionary expenditures.

The authors' main result was that there was a statistically significant and economically large effect on subsequent borrowing behavior for the disclosure that added up the costs of subsequent borrowings *in dollars*. The effect was to reduce both the likelihood of further borrowing and the subsequent amounts borrowed. In particular, borrowers receiving this disclosure were 5.5 percentage points less likely to borrow in subsequent pay cycles (10 percent less likely to borrow compared with the control group) and they borrowed nearly \$40 less (17 percent less than the control group). Morse noted that this effect was large for this type of experimental study, especially since the disclosure was made only once.

The effects of the other disclosures on the likelihood of subsequent borrowings were relatively weaker, both statistically and economically. In particular, disclosing relative APRs seemed to have little effect. In both cases, however, there was evidence of some reduction in the amounts borrowed.

The authors also found that the effects differed across different types

of customers. The decline in the probability of borrowing occurred mainly among individuals without a college degree. The decline was also stronger for those borrowers who reported higher self-control, those who were not borrowing for discretionary purposes, and those with lower debt-to-income ratios. This last result is broadly consistent with the authors' other finding that reduced borrowing occurred only with a lag. According to Morse, borrowers may have needed time to adjust their financial situation, while others in more financially strained circumstances may simply have been unable to adjust, at least in the time frame considered in the study.

In conclusion, Morse and coauthor suggested that the success of the disclosure in the payday setting justifies further explorations of policies that might reduce consumer biases in other contexts.

IDENTITY THEFT

William Roberds of the Federal Reserve Bank of Atlanta presented the results of a theoretical model (with Stacey Schreft) that examined the incentives for competing networks — for example, credit card networks — to adopt policies to reduce identity theft. The underlying questions were whether networks collect too much information and whether they adopt appropriate levels of security to protect that data. In general, Roberds argued that competing networks have incentives to collect too much personally identifiable information (PII)— for example, name, address, Social Security number, and so forth— while spending too few resources to protect it from theft.

In their model, many individuals are honest and joined one of two competing networks to facilitate making transactions to purchase goods. But some individuals are

fraudulent types; that is, they seek to join a network and then default on their payments.

The authors identified two types of identify theft, both involving opening new accounts, rather than stealing an existing customer's account, to purchase goods at the customer's network. First, skilled identity thieves can use sophisticated techniques — for example, hacking the network's database — to steal PII from one network to join another network (high-tech fraud). A second type of identity theft (low-tech fraud) simply requires someone to assemble enough information to create a viable identity to join a network, for example, by stealing a wallet and impersonating that person. This type of theft requires no skill, but it does require time and effort, and it is more costly for an impersonator to join a network if he or she must provide more information.

Networks have two potential security strategies. The first is to collect more PII about a customer. By keeping this information on record, the network can increase the likelihood that fraudulent customers will be detected if they attempt to impersonate a new customer applying for credit. The second security strategy is for the network to spend resources to protect its database. In particular, it can make it more difficult for skilled frauds to steal PII.

Roberds explained that data security involves an externality: By demanding a lot of PII to join, a network makes it more difficult for fraudulent customers to join. But keeping very detailed information about the network's own customers in its database makes it easier for skilled identity thieves to use stolen data to join the other network, because this tends to increase the likely overlap in the types of PII required to join each network. And a network's costly measures to secure

its own database from skilled identity thieves reduce fraud at the *competing* network.

To provide a benchmark for evaluating the choices of competing networks, the authors performed a thought experiment. They asked: How much information and data security would a benevolent social planner instruct the networks to choose? This planner would take into account all of the costs and benefits to individuals, and these hypothetical choices are termed the *optimal* outcome.² After calculating the optimal outcome, the authors examined market outcomes in successively more general examples and compared these with the optimal outcomes.

In the first example, they assume that firms do not secure their data at all. In this case, firms collect too much information. Collecting more information makes it harder for a thief to construct a viable identity to join but also makes it easier for a thief to steal data that can be used at the other network. In this example, networks collect too much information and data breaches occur more often than in the optimum, but interestingly, the prevalence of identity theft is lower than the optimal level. The more overlap between the PII collected by the competing networks, the greater the inefficiency.

In their second example, they assume that the proportions of skilled and unskilled frauds are identical. Again the basic externality arises: Networks collect too much information and invest too little in data security. The main insight from this example is that although identity theft is lower than in the optimum, it is unskilled

²It is important to note that while there are real costs to identity theft, the optimal level of identity theft is not zero. This is because it is costly to deter theft, and these costs are ultimately borne by individuals.

theft that is mainly deterred. Networks make it very difficult for unskilled thieves to join, but their excessive data collection and inadequate data security makes skilled identity theft relatively attractive.

For the most part, these insights carry over to the most general version of the model, in which security levels are freely chosen by the networks. In this setting, the authors showed that when networks require substantially similar types of information, security levels were too low and networks collected too much data. The authors argued that, in effect, competing

Broadly, [White, Li, and Zhu] argued that the passage of the Bankruptcy Reform Act was associated with a statistically significant and economically large increase in mortgage defaults.

networks substitute information collection for data security. As in the simpler examples, there was too much skilled identity theft and too little unskilled identity theft compared with optimal levels. Furthermore, there is less identity theft in equilibrium than in the optimum, even though networks collect too much information.

Finally, the authors examined the effects of public policies that might improve market outcomes. Since the model is quite complicated, they used simulations to evaluate the effects of these policies. One possibility is to increase civil liability for data breaches. They found that this improved incentives to increase security, but networks still collected too much information. A second approach is for some government agency to set minimum data security standards. This nearly attained the optimal outcome in their simulations. A third approach is to limit data collection. This policy did as well as

imposing civil liability but significantly increased identity theft.

BANKRUPTCY REFORM AND MORTGAGE DEFAULT

Michelle White of the University of California-San Diego reported on the results of an empirical study (with Wenli Li and Ning Zhu) of the effects of the Bankruptcy Reform Act of 2005 on mortgage default. Broadly, they argued that the passage of the Bankruptcy Reform Act was associated with a statistically significant and economically large increase in mortgage defaults. According to the authors, the

act may have contributed to the severity of the subsequent crisis in housing markets.

White argued that bankruptcy law helped people save their homes, at least temporarily, but also reduced the costs of ultimately defaulting on the mortgage. Under both Chapter 7 (liquidation) and Chapter 13 (restructuring) proceedings, homeowners can protect exempted assets, in particular, homes in which a household's equity does not exceed the state-mandated homestead exemption. Both bankruptcy procedures also give a delinquent homeowner some time to pay back missed mortgage payments (*arrears*), but Chapter 13 provides a substantially longer period (three to five years) and also permits delinquent homeowners who cannot pay back arrears a significant amount of time before foreclosure.

The Bankruptcy Act of 2005 had three main effects that might affect delinquent homeowners: First,

it raised filing costs, thus making it less attractive for a household to use bankruptcy to save a home either permanently or temporarily. The act also placed a cap on the homestead exemption at \$125,000, a provision that affected 10 states with high homestead exemptions. Third, the act imposed a means test for homeowners to use Chapter 7. Specifically, it required consumers with incomes above the state median to file using Chapter 13.

The authors' empirical strategy was to use a difference-in-difference approach. In particular, the authors examined the differential effects of the change in bankruptcy law on certain households living in different states. The empirical tests exploited variation in consumers' circumstances, state median incomes, and state homestead exemptions to determine whether the act affected homeowner delinquency.

The authors used data from LPS on the performance of first-lien 30-year mortgages. These data also included information about customer credit quality, notably credit scores at the time the loan was originated. By merging the LPS data with data from the Home Mortgage Disclosure Act, the authors could also take account of information and homeowner demographic characteristics, most notably household income. The final sample included about 381,000 prime mortgages and 268,000 subprime mortgages.

Using these data, the authors examined whether mortgage defaults increased following passage of the act for those homeowners for whom the provisions of the act were actually binding, in other words, in cases in which consumers had equity in excess of the now lower exemptions or where the consumer was now required to use his or her excess cash flow to pay nonmortgage debts. Specifically, they

examined a window of three months before and after passage of the act.

Descriptive statistics showed that following passage of the act, default rates were 15 percent higher for prime loans and 9 percent higher for subprime loans. In addition, default rates increased even more for homeowners who were subject to the new cap

The Bankruptcy Act of 2005 had three main effects that might affect delinquent homeowners: it raised filing costs, placed a cap on the homestead exemption at \$125,000, and imposed a means test.

on the homestead exemption and for prime homeowners who failed the means test. However, defaults decreased for homeowners with subprime loans who failed the means test. The authors interpreted these findings as largely consistent with the view that the act increased mortgage defaults.

The authors then estimated a logit hazard model, which estimates the probability of defaulting with the passage of time, taking into account the changes in the bankruptcy law that might affect a particular household, and controlling for a large number of demographic variables. They found that homeowners subject to the cap on the homestead exemption were more likely to default following passage of the act; specifically, White and her coauthors found that prime mortgage holders subject to the cap were 53 percent more likely to default and subprime mortgage holders were 44 percent more likely to default. Prime mortgage borrowers subject to the means test were 14 percent more likely to default, but there was no effect for subprime mortgage borrowers. White suggested that the result for subprime mortgage borrowers may have been

due to income having been overstated by these borrowers.

Results were largely the same when the authors reestimated their model using a six-month window before and after passage of the act. Thus, the authors concluded that the effects of act were not temporary.

CONSUMER PROTECTION LAWS

Simon Gervais of Duke University presented the results of a theoretical model (with Bruce Carlin) that examined the role of the legal system when customers are poorly informed about the appropriate type of financial product to buy. In their model, households depend on brokers to match them to financial products for which they are best suited.

The main assumption of the model is that particular products are better suited for particular types of consumers. For example, a household with moderate savings and a student in high school might be better advised to invest its savings for college in a fixed income product, rather than a stock index fund, but the household might not have the sophistication to know which product is most appropriate. In Gervais and Carlin's model, both brokers and the producers of financial products must exert costly effort: At some cost, brokers can direct consumers to those products that suit them best, although there is an unavoidable probability of a mistake. Similarly, at some cost,

the producers of financial products can develop higher quality products that are suited to a wider range of consumers. Crucially, no court can observe their effort directly, nor can the court disentangle the reason why a particular product turned out poorly for a particular customer: Was it a bad match or a bad product?

In this setup, Gervais explained that product quality and effort by brokers are partial substitutes; that is, more effort by a broker reduces losses to consumers and this, in turn, reduces the producer's incentive to develop the highest quality product. In a similar fashion, when higher quality financial products are developed, there is less chance of losses to consumers, and this reduces brokers' incentives to identify the most suitable products for their customers.

The law's design must take this interaction into account. For example, while legal penalties for a broker will tend to increase broker effort, thereby increasing the probability of a good match for the customer, this will tend to decrease the provision of quality products because the firm realizes that the broker's effort will make up for the lack of quality.

First, the authors demonstrated that without legal penalties, the market leads to a serious underprovision of effort both by brokers and firms. Indeed, in their stylized setup, consumers are not willing to buy the product at all and the market breaks down altogether. Intuitively, without penalties, consumers who pay a price up-front expecting an appropriate financial product will always be disappointed because producers and brokers always have an incentive to chisel once they have been paid.

Gervais then presented their main results in a version of the model in which a customer can seek redress through the legal system only when he

or she has followed a broker's advice. In this context, the authors showed that to achieve an efficient outcome, total legal penalties can't merely be compensatory; they must be punitive. This conclusion is jointly the result of the substitutability of effort by producers and brokers and of the court's inability to assign blame to one or the other for a bad outcome. When the law seeks to push, say, the broker to increase effort by penalizing him or her when a match turns out poorly,

[Gervais and Carlin] showed that to achieve an efficient outcome, total legal penalties can't merely be compensatory; they must be punitive.

the producer of the financial product responds by reducing effort. Intuitively, this means that total penalties must exceed the losses imposed on the borrower for having been mismatched to induce full effort by both brokers and producers.

The authors then enriched the basic model to include the realistic possibility that the firm pays the broker for each sale. In this setting, they show that the optimal penalty structure places higher penalties on the broker than when brokers do not receive direct payments from producers.

Gervais then explained how legal penalties changed if customers were permitted to seek legal redress even when they have ignored a broker's advice. In this setting, Gervais and Carlin demonstrated that the optimal penalties were no longer punitive; customers only received compensation for having made a poor decision. Intuitive-

ly, punitive penalties reward customers for ignoring their broker's advice and then seeking redress through the courts whenever they make a bad decision. Accordingly, the optimal legal scheme can't reward customers for making bad decisions.

GLOBAL RETAIL LENDING

In the final paper, Jörg Rocholl of the ESMT European School of Management and Technology presented the results of an empirical study (with Manju Puri and Sascha Steffen) of the effects of the crisis in U.S. mortgage markets on German banks. Rocholl and his coauthors used the unique structure of the German banking system as a natural experiment for distinguishing supply-side effects from demand-side effects.

Rocholl explained that there are 11 Landsbanken in Germany, jointly owned by state governments and the savings banks in those states. The savings banks provide financial services only for the customers in their municipality, primarily small and medium-sized firms, as well as retail customers. A key feature of the system is that the Landsbanken can rely on both formal and informal support from the savings banks with an ownership share. Thus, losses at the Landsbanken will impose losses on the savings banks, which are significant owners.

In Rocholl's account, Germany experienced growth well into 2008 and avoided the housing bubble occurring in the U.S. and other European countries. Nor did it experience the housing bust. But a number of Landsbanken were heavily exposed to risky U.S. housing assets and had experienced large losses by the third quarter of 2007.

A key feature of the German banking market provided the setting for a natural experiment. Only some of the savings banks were owners of

affected Landsbanken, while others had no exposure to U.S. real estate losses. And since the national housing market was largely homogeneous, Rocholl and his coauthors argued that troubles at a savings bank's Landsbank may be viewed as a pure shock to the supply of loans.

Using a difference-in-difference analysis, the authors compared the change in lending behavior at affected and unaffected savings banks before and after August 2007, when the U.S. housing crisis began to affect assets owned by certain Landsbanken. The authors also had information about loan applications at these banks, which made possible a clear distinction between changes in supply and demand. For example, if the authors observed that loan applications were similar across savings banks, but fewer loans were booked at affected banks post-crisis, this is strong evidence that the underlying source of the change in lending was supply driven.

The authors had data on all consumer and mortgage loans by savings banks in Germany between July 2006 and June 2008. They also had data on loan applications and the bank's risk rating of the consumer, as well as information about any preexisting financial relationships with the consumer, for example, credit lines and assets held at the bank.

In the central findings of the paper, Rocholl and his coauthors estimated a linear probability model of loan acceptance rates. They found that loan acceptance rates at affected banks declined significantly after August 2007, while acceptance rates increased insignificantly at unaffected banks. The decline at affected banks was economically large; across all types of consumer lending, acceptance rates declined 8.2 percent. The decline was strongest for customers that were assigned low credit ratings by the banks, suggesting a flight-to-quality effect. The results were consistent across loan categories, although the effects were larger for mortgage loans. The authors argued that this is because mortgage loans represent a larger commitment by the bank than other types of consumer loans.

Rocholl and his coauthors also estimated a cross-sectional regression to examine how bank characteristics affected lending behavior. They found that the declines were most dramatic for smaller banks and that for such banks the declines were particularly large for mortgage loans. They also found that the effects were greatest for banks that were relatively illiquid entering the crisis.

The authors then examined the demand for loans. In a regression framework, the authors found that loan

applications declined at both affected and unaffected banks. There was no statistically significant difference in the trend for these two groups of banks. The authors suggest that the decline in applications reflected a decline in demand, as consumers became less certain about future economic conditions, and the decline was not bank-specific. Nor did the authors find any significant difference in the amount of loan requested. This reinforced the authors' view that the declines in lending by affected banks were driven by the supply shock rather than effects on demand.

The authors then examined the effects of relationships in a linear probability model using a triple difference approach, in which loan applicants were further differentiated according to whether they had an existing relationship with the bank. Customer relationships with a bank increased acceptance rates, and the effect was strongest at affected banks after August 2007. Thus, pre-existing customer relationships mitigated the negative supply shocks at affected banks, perhaps because lenders have more information about such borrowers. 