The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXI, NO. 6 • DECEMBER 2016

A Tale of Two Runs: Depositor Responses to Bank Solvency Risk

RAJKAMAL IYER, MANJU PURI, and NICHOLAS RYAN*

ABSTRACT

We examine heterogeneity in depositor responses to solvency risk using depositor-level data for a bank that faced two different runs. We find that depositors with loans and bank staff are less likely to run than others during a low-solvency-risk shock, but are more likely to run during a high-solvency-risk shock. Uninsured depositors are also sensitive to bank solvency. In contrast, depositors with older accounts run less, and those with frequent past transactions run more, irrespective of the underlying risk. Our results show that the fragility of a bank depends on the composition of its deposit base.

WHO RUNS ON A BANK, AND why? We know that runs are related to bank solvency in aggregate (Saunders and Wilson (1996), Calomiris and Mason (1997)). Yet deposits are not a homogeneous mass—they are held by people with different histories and different relationships to their banks. A person with only a modest checking account, for example, may not bother to learn about her bank's financial health, whereas those with higher balances or a broader relationship (e.g., they also hold a loan) may know more about their bank and as a result have more reason to act on that knowledge, since their financial well-being is tied up with their bank's.[1] Following this line of thought, if some kinds of depositors are more or less sensitive to the solvency risk of their bank, then the make-up of a bank's deposit base may be an important determinant of stability. Treating deposits as held by heterogeneous depositors, each with their own

*Rajkamal Iyer is with MIT Sloan. Manju Puri is with Fuqua School of Business, Duke University, and NBER. Nicholas Ryan is with Yale University. We are grateful to Mr. Gokul Parikh and the staff of the bank for all their help and to Anup Roy and Pramod Tiwari of IFMR for supervision of the depositor survey. We thank the Editor, Michael Roberts, and four anonymous referees for comments that greatly improved the paper. We thank Nittai Bergman, Doug Diamond, Mark Flannery, Xavier Giroud, Ali Hortaçsu, Daniel Paravisini, Antoinette Schoar, Andrei Shleifer, and Tavneet Suri for comments. We thank seminar and conference participants and discussants at the ABFER (Singapore), ASSA meetings (San Diego), Corporate Finance Conference (Bristol), CAFRAL, Reserve Bank of India, Columbia University, GSE Summer Forum, Barcelona, UC Berkeley, CEPR-EBRD-EBC-ROF Conference, Duke University, European Central Bank, FDIC/JFSR, FIRS (Croatia), Indiana University, Lingnan University, Minneapolis Fed, MIT, NBER Summer Institute, New York Fed, Riksbank, Tel Aviv, and the World Bank. The authors declare that they have no relevant or material financial interests related to the research in this paper.

[1] We use female pronouns throughout though depositors may be of any gender.

DOI: 10.1111/jofi.12424

notion of solvency risk, may help us understand the nature of runs and aid in the design of banking regulation.

Despite the importance of understanding the micro-level response to solvency risk, evidence on this subject is scarce, for several reasons. First, and most plainly, it is hard to obtain detailed microdata on depositors, their relationships with a bank, and their withdrawal behavior during a run. Second, interpretation of most shocks is not clean as to date we lack a clean ex ante measure of banks' solvency risk that would allow us to examine whether depositors respond to that risk independently from the actions of other depositors or the outcome of a run. Third, it is difficult in practice to compare the response of depositors to shocks with different degrees of underlying solvency risk.

In this paper, we study the behavior of depositors across two shocks with different degrees of solvency risk that were experienced by a single bank. To do so, we use a new data set from a bank in India with microlevel depositor data. This data set allows us to identify depositor characteristics along with the timing of every depositor transaction. We use this data set to study the behavior of depositors with different characteristics across the two shocks, which were eight years apart and each triggered runs on the bank. We define a high-solvency-risk shock as a shock that renders the bank insolvent and a low-solvency-risk shock as one that does not affect the bank's solvency, *absent any further response by depositors.* Of course, depositors may not be aware of the nature of a shock at the time they decide whether to run—but whether the actions of different types of depositors reflect a bank's underlying solvency risk is precisely the question of interest.[2] We study depositor withdrawals for the bank's entire depositor base under both shocks and, among the subset of depositors who hold accounts at the time of both shocks, for the same individual depositors in the two events.

The bank we study experienced a high-solvency-risk shock and was subject to runs in early 2009, during and after a regulatory intervention that ultimately placed the bank in receivership. We first examine depositor behavior during this high-solvency-risk shock and then compare it with a prior low-solvency-risk shock. The timeline of high-risk shock that we exploit is as follows. The bank had a build-up of bad loans. This build-up was uncovered by the central bank during an audit. While the bank's negative net worth was documented by the central bank, it remained private information. This audit was followed, after several months, by public news that the central bank was severely restricting the bank's activity.

We find that there is a large run by depositors immediately following the public news of the high-solvency-risk shock. Uninsured depositors are far more likely to run than insured depositors. Depositors that have loan linkages with

[2] If a large fraction of depositors run, then even when the initial shock does not affect the solvency of the bank, the run can become self-fulfilling, put the solvency of the bank into question, and bring about failure. We therefore define underlying solvency risk as a threat to the solvency of the bank as a result of the initial shock, without the response of depositors (while acknowledging that panics can also bring down banks).

the bank or who are bank staff are also more likely to run. Further, depositors are more likely to run if a member of their network has already done so, or if they have a higher volume of transactions with the bank. In contrast, depositors with a longer relationship with the bank are less likely to run than others. Thus, while loan linkages appear to increase the likelihood of running, account age reduces the likelihood of running, even in the presence of high solvency risk. These results suggest that, beyond the mere fact of a relationship, how relationships are established matters for depositor behavior.

We next broaden the event window to study whether some types of depositors run even before negative news becomes public. We find that there was indeed a silent run, beginning at the time of the regulatory audit but prior to the public release of information that was driven by uninsured depositors, depositors with loan linkages, and staff members. Staff of the bank withdrew first in response to the audit, followed closely by uninsured depositors and depositors with loan linkages. Thus, while in principle a regulatory audit is private information available only to the bank, in practice uninsured depositors, depositors with loan linkages, and bank staff withdraw more immediately following an audit.

The results above suggest that there are sharp differences in the responses of different depositor types to a high-solvency-risk shock. Observing how withdrawals respond to this one shock, however, leaves two important questions unanswered. First, is it depositor relationships themselves that matter, or do those relationships reflect omitted characteristics of depositors, such as education or financial literacy, that drive withdrawals? Second, are depositors responding to the fundamental nature of the shock, or would they withdraw just the same in response to a low-solvency-risk shock?

To address the first of these questions, we focus on a sample of depositors that hold accounts during the high-solvency-risk shock and collect household survey data on demographics, financial literacy, and assets. We find that each of these sets of depositor characteristics matter for explaining which depositors run after the shock. Depositors are significantly more likely to run if they are more educated, are engaged in a business or professional occupation, are more financially literate, or hold more assets. However, when we add these additional characteristics to the set of factors that explain why depositors run, we find that the strong effects of depositor banking relationships on liquidation are unchanged.

To address the second question on whether depositors respond to the nature of the shock, we contrast the depositor response to the high-solvency-risk shock with the depositor response to a low-solvency-risk shock that hit the same bank eight years earlier. At this time our bank experienced a run in response to the idiosyncratic failure of another bank in the same city due to fraud. Our bank had no fundamental linkages to the failed bank and the run lasted only a few days. Our bank was solidly solvent at the time, but depositors' beliefs about its solvency risk could have been very different from the true state.

We find that, during a low-solvency-risk shock, depositors with loan linkages are *less* likely to run. The behavior of depositors with loan linkages is thus sensitive to the nature of the shock, in a direction that suggests they are informed about the bank's true solvency—they are more likely to run when the bank's solvency is at risk. Bank staff are less likely to run under the low-solvency-risk shock than the high-solvency-risk shock, and uninsured depositors are again more likely to run as compared to insured depositors, but to a much lesser extent than in the case of the high-risk shock. Some depositors, however, are not sensitive to solvency risk. Depositors with a longer duration relationship with the bank are less likely to run and those with a higher volume of transactions with the bank are more likely to run, regardless of the type of shock.

Though education, financial literacy, and the other observables collected do not alter the effect of banking relationships on withdrawal, other unobservable characteristics of depositors may. We test for such unobservables by estimating the determinants of running among the pool of depositors that held accounts during both shocks, which allows us to add depositor fixed effects to control for time-invariant unobservable characteristics of depositors. It is fairly remarkable to observe the behavior of the same depositors in response to different shocks outside of a laboratory setting. We find that the results reported above are all robust to adding depositor fixed effects. Note that this constant sample across shocks is subject to survivorship bias, in that any depositor present in the constant sample saw that the bank survived the earlier low-solvency-risk shock and still kept some deposits at the bank. When we address this selection using a reweighting procedure, we find that the results are again unchanged.

Our interpretation of the differential response of depositors to shocks of differing solvency risk is that, due to their banking relationships, some types of depositors are informed about solvency risk and have an incentive to act by withdrawing in a crisis. Depositor heterogeneity in the response to a single shock may be due to information or depositor incentives. For example, we find a negative coefficient on loan linkages in the case of the low-solvency-risk shock. These depositors might know there is little risk of failure and therefore stay with the bank. Alternatively, loan-linked depositors may not run because they face higher costs of switching banks or have greater trust in the bank. Under the high-solvency-risk shock, however, we find that loan-linked depositors are more likely to run. This suggests that, even if they do have greater trust or face higher switching costs, they must also be informed in order to change their behavior in a way that is responsive to the nature of the shock. Similar behavior differences across shocks apply to the bank staff and uninsured depositors. We argue that these depositors may be informed about solvency risk through personal networks of bank staff, loan officers, and other depositors.

This paper adds to a large theoretical and empirical literature on bank runs. Our results are consistent with theoretical models of coordination problems where fundamentals play an important role in coordinating beliefs (Goldstein and Pauzner (2005)).[3] Our findings also provide an empirical basis for the

[3] Bank fundamentals, either directly by acting on depositors' incentives or indirectly by acting on their beliefs about others' actions, shape depositor actions and whether the bank survives or fails. See Bryant (1980), Diamond and Dybvig (1983), Postlewaite and Vives (1987), Goldstein and Pauzner (2005), and Rochet and Vives (2004) for models based on coordination problems. See

heterogeneity in signals received by different depositors, which is an important building block in these theoretical models.

The empirical literature on bank runs focuses on whether bank runs are justified by fundamentals or are best characterized as panics. The literature finds that banks with worse fundamentals experience greater deposit withdrawals in a crisis (Gorton (1988), Saunders and Wilson (1996), Calomiris and Mason (1997)).[4] Looking at bank-level data, these withdrawals act as a form of depositor discipline on risky banks (Park and Peristiani (1998), Billett, Garfinkel, and O'Neal (1998), Martinez-Peria and Schmukler (2001), Goldberg and Hudgins (2002), Bennett, Hwa, and Kwast (2014)).[5] However, the empirical literature also finds that some runs are driven in part by panic, not just fundamentals (Calomiris and Mason (1997), Iyer and Puri (2012)). Our study takes this question to the microlevel to identify what types of depositors respond to the true solvency risk of a bank. By using microdata on responses across two well-understood shocks, this paper offers sharp evidence that depositors do indeed respond to bank fundamentals, and do not withdraw only due to coordination problems or shocks common to depositors and their banks.[6]

A smaller set of papers considers the responses of individual depositors to bank runs (Davenport and McDill (2006), Iyer and Puri (2012), Brown, Guin, and Morkoetter (2014)). We combine rich administrative and survey data to identify the effects of a wide range of banking relationships and depositor characteristics on actual withdrawals for the universe of depositors at a failed bank. In comparison to other studies of depositor behavior during panics (Iyer and Puri (2012), Brown, Guin, and Morkoetter (2014)), our paper is unique in being able to contrast depositor behavior across shocks with differing degrees of solvency risk. This contrast matters greatly for the interpretation of depositor behavior after a shock. Suppose that depositors with longer lived accounts or loan linkages run less during shocks that look like panics. Are these deposits stable, or are they informed about actual solvency risk? Our findings clarify that it depends on the type of banking relationship: long-lived deposits are stable and not sensitive to the true solvency risk, whereas the opposite is true for deposits held by depositors with loans.

Chari and Jagannathan (1988), Jacklin and Bhattacharya (1988), Chen (1999), Calomiris and Kahn (1991), and Diamond and Rajan (2001) for information-based models of runs.

[4] See also Gorton and Metrick (2012) and Chen, Goldstein, and Jiang (2010).

[5] Flannery and Sorescu (1996) find that spreads on bank-subordinated debentures reflect bank risk relatively more following policy changes that increased the default risk on subordinated bank debentures.

[6] The bank-level literature on market discipline often cannot distinguish these alternatives, for at least two reasons. First, the solvency risk posed by a shock is often determined ex post, by which banks ultimately fail, with market discipline measured by whether these doomed banks saw early withdrawals (Saunders and Wilson (1996), Goldberg and Hudgins (2002)). This test does not sharply distinguish market discipline in response to solvency risk from a self-fulfilling panic, in which we would also expect banks that saw early withdrawals to fail. Second, in the study of a crisis, banks and their depositors may be subject to common shocks, so that depositors at distressed banks are withdrawing not in response to perceived bank solvency but, for example, because they have lost their own jobs.

The different depositor behavior across different shocks that we document has implications for the design of policy to mitigate bank fragility without sacrificing depositor discipline. For example, our results suggest that loan linkages strike this balance, since loan-linked depositors will withdraw more only in the case of a high-solvency-risk shock. The liquidity coverage ratio in Basel III requires that banks have enough high-quality liquid assets to cover total expected cash outflows during a 30-day shock (Basel Committee on Banking Supervision (2013)). Under this rule, cash flows are based on anticipated run-off rates for "stable" and "less stable" deposits.[7] Our results support this characterization in broad terms, but suggest several modifications or caveats. First, older accounts are an example of a type of relationship that leads to stability. Second, some depositor relationships, like having a loan, may be stable during a panic but not following a fundamental shock to asset values. This instability may be a good thing, however, in the sense that it incentivizes banks to accumulate stable deposits, and *conditionally* stable deposits may in turn preserve market discipline. Third, this rule allows deposits covered by an "effective" deposit insurance scheme to be considered stable. This proviso is important when, as in India and many other developing countries, insurance payouts may be delayed: following the fundamental shock, we find that the run-off rate of insured deposits is 20%, well above the Basel III assumption. The distinction between stable and less stable deposits is nonetheless still justified, since runs from uninsured depositors are greater still. Fourth, transactional accounts with a high frequency of transactions may not be presumed stable. In general, we find that liquidity coverage ratios based on depositor characteristics are sound in principle but may be fine-tuned, taking into account how depositor heterogeneity interacts with solvency risk.

Our results also speak to other policies for financial stability. We find that depositors with more frequent past transactions with the bank are more likely to run, regardless of a bank's solvency risk. This suggests that, during a crisis, regulators can selectively target certain classes of depositors that are most prone to run. Indeed, during the recent crisis in the United States, the transaction account guarantee program was targeted in this way.[8] There are different rationales in the literature for why deposit-taking and lending should come under the same institutions (Diamond and Rajan (2001), Kashyap, Rajan, and Stein (2002), Hanson et al. (2014)). Our finding on the response of loan-linked deposits to solvency risk provides a new reason, based on financial stability: depositors who are also borrowers are more likely to discipline banks, and withdraw mainly in response to high-solvency-risk shocks, providing stable deposits during a panic. A final policy implication of our results pertains to regulatory disclosures. Though the change in depositor behavior across shocks

[7] See Basel Committee on Banking Supervision (2013, p. 27). Stable deposits are categorized as retail deposits that are fully insured or where depositors have another established relationship with the bank that makes withdrawal highly unlikely. Deposits in transactional accounts where salaries are automatically deposited are also considered stable.

[8] For instance, in the United States, non-interest-bearing accounts that have high transaction activity had unlimited deposit insurance coverage during the recent crisis.

is consistent with the market discipline of banks, a strong regulatory signal and subsequent action play an important role in sparking withdrawals during a high-solvency-risk shock. Improving regulatory supervision and information disclosure is therefore complementary to market discipline by depositors.[9]

The rest of the paper is organized as follows. Section I describes the institutional environment, the shocks we study, and the data. Section II presents the empirical results on depositor behavior in response to the high-solvency-risk shock. Section III compares the two shocks and interprets the differences in depositor behavior across the two types of shocks. Section IV concludes.

I. Institutional Environment and Event Description

A. Institutional Details

The Indian banking system consists mainly of public sector banks, private banks, and cooperative banks. The Reserve Bank of India (RBI) is the main regulatory authority of the banking system and monitors bank portfolios and capital requirements for all three types. Cooperative banks are additionally supervised by the state government on matters of governance but not finance.

Deposit insurance exists but coverage is incomplete. The Deposit Insurance and Credit Guarantee Corporation, part of the RBI, provides deposit insurance up to INR (Indian rupees) 100,000 (roughly USD 2,000) for each depositor at a bank. The deposit insurance is funded by a flat premium charged on insured deposits and required to be borne by the banks themselves. Though deposit insurance is present, there are several delays in processing depositors' claims. The central bank first suspends convertibility when a bank approaches failure and then decides whether to liquidate a bank or arrange a merger with another bank. During this period, depositors are allowed a one-time nominal withdrawal, up to a maximum amount that is stipulated by the central bank.[10] If a bank fails, the deposits held by a depositor cannot be adjusted against loans outstanding. The stipulated cash reserve ratio and statutory liquidity ratio (SLR) that banks are required to maintain are 5% and 25%, respectively.[11]

Cooperative banks are not different in kind from banks with other ownership structures. Depositors at cooperative banks are not required to hold an equity claim in the bank, and shareholders of cooperative banks have limited liability and generally do not receive dividends.[12] Thus, the nature of cooperative banks does not select depositors with different characteristics from those

[9] See also Flannery and Houston (1999), Berger, Davies, and Flannery (2000), and DeYoung et al. (2001) for evidence supporting the importance of regulatory information for banks.

[10] In most cases, depositors are allowed a withdrawal of up to INR 1,000 (USD 20) per account.

[11] The SLR is the minimum allowable ratio of liquid assets, given by cash, gold, and unencumbered approved securities, to the total of demand and time liabilities.

[12] The bank issues shares at face value. To borrow from the bank, the bank asks a borrower to buy shares worth 2% of the loan principal amount, which can be redeemed at face value at the end of the loan. The implied interest payment foregone by borrowers is equivalent to processing fees charged by other banks for loan originations. In general, the bank does not pay dividends, as reserves are used to meet capital adequacy requirements.

at banks with other ownership structures. One of the main reasons depositors prefer cooperative banks is that they offer more customized services than larger private banks. In the United States, the closest analogues to Indian cooperative banks are community banks, which play an important role in the U.S. economy (Kroszner (2007)).[13]

B. Event Description

We now describe the events that we study in this paper. First, we describe the high-solvency-risk shock. The bank we study is a cooperative bank that functioned well until 2005, when the management changed and the bank took heedless and possibly corrupt risks. In May 2007, an RBI inspection privately noted that the bank had introduced proscribed insurance products and made two unsecured loans far in excess of the exposure ceiling. These two loans totaled INR 230 million (USD 5 million), or 60% of the bank's total nonperforming assets as of March 31, 2008. The fundamental reason for the bank's collapse was the nonperformance of these large loans. After a routine inspection for the financial year showed the poor state of the bank's finances, the RBI brought the bank under greater scrutiny and conducted a further audit of the bank's books beginning on November 4 and lasting through November 15, 2008. This audit found that, due to a large volume of nonperforming assets, the bank was insolvent with a negative net worth of INR –313 million (USD –6.25 million). Further, the public balance sheets of the bank in 2007 and 2008 did not reflect the true extent of nonperforming assets. This audit by the central bank was private information and not announced to depositors. In response to the findings of the audit, the central bank ordered restrictions on bank activity including the partial suspension of convertibility. Information about the restrictions imposed on the bank by the regulator was widely covered in the press on January 28, 2009. Depositors were prevented from prematurely liquidating their term deposits. Critically for this study, there was no restriction on withdrawals from transaction accounts. The bank was also forbidden to take new deposits, make new loans, or pay dividends. On May 13, 2009, the central bank finally decided that the bank should be placed under receivership and mandated a withdrawal limit of INR 1,000 for all depositors from all accounts, including transaction accounts. There were long delays in processing deposit insurance claims.

We characterize this event as a high-solvency-risk, or fundamental, shock since the bank was insolvent even absent depositor runs. Importantly, however,

[13] In a speech on March 5, 2007, Federal Reserve Governor Randall Kroszner stated that, "Community banks play an important role in the United States economy, as they have throughout our history . . . many community banks continue to thrive by providing traditional relationship banking services to members of their communities. Their local presence and personal interactions give community bankers an advantage in providing financial services to those customers for whom, despite technological advances, information remains difficult and costly to obtain . . . I believe that the most significant characteristics of community banks are: (1) their importance in small-business lending; (2) their tendency to lend to individuals and businesses in their local areas; (3) their tendency to rely on retail deposits for funding; and (4) their emphasis on personal service." Cooperative banks display the same four significant characteristics as community banks.

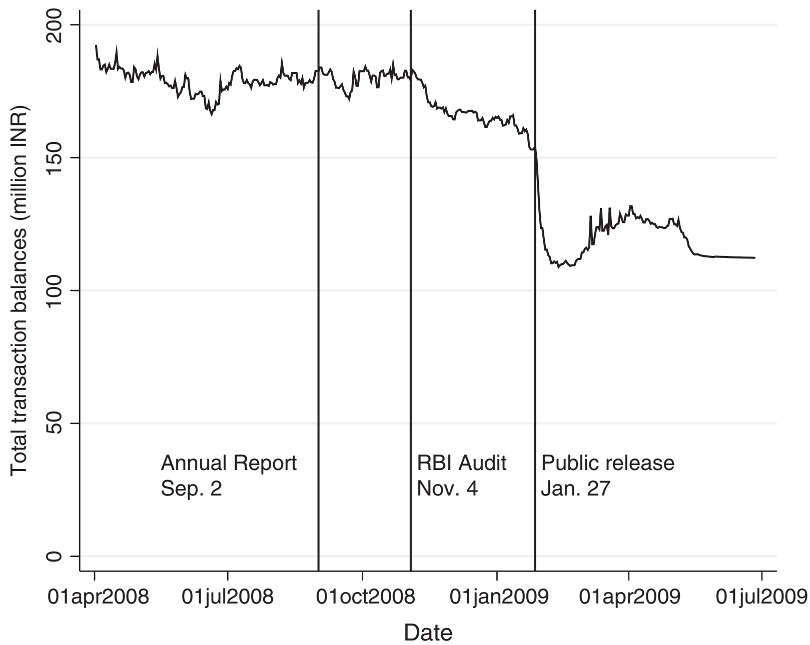

Figure 1. Transaction balances, high-solvency-risk shock. The figure shows aggregate transaction account balances for depositors in the bank from 300 days before the public release of information on regulatory action against the bank, which occurred on January 27, 2009, through 150 days after. The vertical lines indicate the dates of (i) the bank's annual report, (ii) the Reserve Bank of India's (RBI, the primary regulator) audit of the bank's finances, and (iii) the public release of information on RBI's actions following this audit. The lines are labeled with the date of the event itself but are drawn to intersect the closing balance of the day before the event.

this failure was idiosyncratic in nature and not due to weak macroeconomic conditions. For example, the state economy grew by just over 9% during the year the bank was under scrutiny, no other banks failed during the event window, and most banks in the region were gaining deposits. Depositors at the bank under study were aware of other bank failures in the state in the recent past, and that uninsured depositors had not been made whole. Further, because the bank was located in a major city with numerous other cooperative, private, and public bank branches nearby, the physical transaction costs of relocating deposits were small.

Figure 1 presents the aggregate pattern of withdrawals by depositors during the high-solvency-risk shock. Significant dates during the crisis are marked by vertical lines in the figure. Prior to the RBI inspection, which began on November 4, 2008 and lasted until November 15, transaction balances had been largely stable over the fiscal year to date. After the regulatory audit by the central bank we see a gradual but significant run, whereby deposits declined 16% from November 4, when the audit began, to January 27. On January 28, newspapers reported on the regulatory action against the bank including partial suspension of convertibility. In the week following the public release of information, we see a large run on the bank, with transaction balances declining a further 25%, for a total decline of 37%, since the day prior to the audit.[14]

We now turn to the description of the low-solvency-risk shock. We classify shocks as having low solvency risk when they do not materially affect a bank's solvency *absent any further response by depositors*.[15] The bank under study experienced a run in 2001 that was triggered by fraud at another large bank in the same city that had branches nearby (henceforth, Bank Two).[16] On March 8, 2001, some major brokers defaulted on their pay-in obligations to the stock exchange. Rumors were afoot that Bank Two had lent heavily to a broker who then suffered huge losses from stock holdings in badly performing sectors (information technology, communication, and entertainment). This led to a run on Bank Two on March 9, and then again on March 12, 2001. When Bank Two failed to repay depositors on March 13, the central bank temporarily suspended convertibility and restrained the bank from making payments above INR 1,000 per depositor. The failure of Bank Two triggered runs at several other cooperative banks in the state (Iyer and Peydro (2011)), including the bank that we study here. We characterize this shock as involving low solvency risk, since our bank had no fundamental linkages with Bank Two through interbank loans outstanding or a correspondent relationship. Further, our bank did not have any investments in the stock market, and its lending portfolio, which consisted of individual and small business loans, was performing fine. Our bank faced runs for only a few days after the failure of Bank Two, with activity returning to prerun levels afterwards. Note that the RBI made no statements regarding the solvency of other banks after the failure of Bank Two—the runs on our bank stopped on their own. Again, at the time of the shock, the economy of the state was growing (at a 9.8% annual rate) and nearby public sector banks saw an increase in deposits over this period.

C. Data

We use data from two sources: administrative data on balances, transactions, and loans from the bank that experienced the two shocks described above, and

[14] Note that these dates were *not* inferred ex post by looking at the time series of balances, but rather by examining documentary evidence on information about the crisis—both private information, from bank records, and public information, from newspaper accounts. Still, we use a Chow test to verify whether the dates documented for the event mark statistically significant structural breaks in the time series of balances. The break on the day of and the day after the public release of information is the largest in the time series by far and highly statistically significant (Figure IA.1 of the Internet Appendix). The second largest break, also statistically significant, occurs in the week after the RBI audit began.

[15] As econometricians, we know whether the bank is solvent based on information from the central bank. Depositors, however, do not observe whether a shock poses high or low solvency risk. Rather, they have to form expectations about the threat to the bank's solvency posed by the shock.

[16] Iyer and Puri (2012) study another bank (Bank Three) that was also affected by this shock and describe the shock in greater detail.

household survey data on depositor characteristics from a survey we conducted of a subset of depositors. We describe each of these data sources in turn.

The administrative data cover both the low-solvency-risk (2001) and high-solvency-risk (2009) shocks. This bank had eight branches around the city. The data record all deposit balances, transactions, and loans at all branches from January 2000 through December 2005 and from April 2007 through June 2009.[17] We describe the variables we use below; the Appendix summarizes the definitions of these variables.

Transaction accounts are defined as current (i.e., checking) or savings accounts, both of which hold demandable deposits. We calculate daily transaction account balances and withdrawals or deposits between days.[18] Liquidation in the cross-section is defined as the withdrawal of 50% of transaction balances over the seven days beginning the day before the shock. (We often refer to this group as "runners," as opposed to "stayers," and vary this definition in a robustness check.) We also estimate hazard models at a daily frequency, where liquidation is more stringently defined as the withdrawal of 50% of transaction balances in any single day. Transaction balances 90 days prior to the shock (120 days prior in hazard specifications) are used to measure depositor balance levels ex ante and to classify depositors by their deposit insurance coverage. We classify depositors with total deposits greater than INR 100,000, the deposit insurance threshold, as "above insurance cover" or uninsured and compare this group of depositors to those with lesser balances. To measure past account activity, we use the share of days over the year prior to the information release, excluding the 90 days immediately prior, during which the depositor liquidated 50% of her balances (i.e., the mean of the lagged dependent variable from the hazard specifications). Account age is defined as the duration an account has been opened in years as of the day before the shock (March 13, 2001 for the low-solvency-risk shock or January 27, 2009 for the high-solvency-risk shock). We top-code account age at seven years, as the age of accounts older than seven years were apparently not recorded or missing when the bank computerized its records.

Family identifiers and depositor loan linkages are defined based on depositor surnames and addresses. We compare each depositor to all others based on surname and address to classify them as belonging to families.[19] We also have

[17] The bank changed its database format and computer system in the interval between these periods. We define variables such as loan linkages to agree across the two events and note the few instances when the change in database may affect the analysis in Section IV.

[18] Daily transaction account balances are directly available from the bank's database for the later period. For the earlier period, daily balances are calculated from monthly balance and daily transactions files at the account level. We confirmed the reliability of this calculation by matching balances at month-end to the opening balance for the same account the next month.

[19] We calculate the ratio $R = 1 - L/MaxOps$, where L is the Levenshtein edit distance between strings, the minimal number of character operations required to change one string into another, and *MaxOps* is the maximum number of character operations that could be required to change one string into another given the lengths of each. Accounts are declared as linked if $R_{\text{Surname}} > 0.75$ and $R_{\text{Address}} > 0.80$ for the surname and address, respectively. We consider these criteria fairly conservative.

data on borrowers from the bank. We define loan linkages for depositors by matching on customer surname and address across depositor and borrower files. Accounts are compared on surname and address using the same criteria as the family match and taken as belonging to the same customer if there is a match. Depositors matched in this manner are defined as having a loan linkage in each crisis if they, or any member of their family, have a current or past loan from the bank as of the date of each run. The definition of loan linkage excludes overdraft accounts against fixed deposits as such accounts may impose restrictions on the withdrawal of deposits. Note that depositors with loans are generally not allowed to offset loans outstanding against deposits in the case of failure.[20] Accounts held by staff members are marked with distinct account codes in the data, although they are identical in substance to the accounts held by nonstaff. We define depositors as having a staff linkage if either they themselves or a member of their family holds an account with a staff code.

We define the introducer network of depositors based on depositor references when opening an account. It is commonplace in India for banks to ask a person opening an account to be introduced by an acquaintance who already holds an account with the same bank, in order to verify their identity. We define a depositor's introducer network as consisting of anyone who introduced that depositor, anyone introduced by the same person as that depositor, and anyone that the depositor him- or herself introduced. This definition is reciprocal in that each depositor is a member of the network of those who belong to her network. To capture network linkages, we define a dummy variable equal to one for a depositor on each date if any member of a depositor's introducer network has liquidated her balance by that date, during the event window of 90 days before to 30 days after each run. We also identify depositor neighborhoods by drawing up a list of 292 precise neighborhoods in the bank's city and fuzzy-matching these neighborhoods to depositor addresses.

Some specifications use data on depositors present during both runs. Since account numbers changed between the runs, this constant sample is determined using a match following the same procedure as above on depositor name, surname, and address.

The second source of data is a household survey of depositors' education, occupation, financial literacy, and assets. This survey was specifically designed to collect information on omitted factors that may be correlated with the primary variables of interest on banking relationships. The sampling therefore overweights depositors with loan linkages (sampled with probability one), those with any balance above the insurance cover (probability one), staff members (probability 0.5), and those with accounts less than one year old (probability 0.5), relative to a randomly sampled group of other depositors (probability 0.18). A total of 6,008 depositors were assigned to be sampled and 4,634 surveys (or 77%) were completed. The primary reason for not completing the survey was

[20] In some cases the central bank makes an exception.

that people were not found at their last known address; only 17 depositors declined to complete the survey.[21]

The survey questionnaire covered three broad areas: demographics, in which we include occupation and education, financial literacy, and asset holdings. The occupation and education categories used in the instrument follow those of India's National Sample Survey. To capture financial literacy, we ask mainly about knowledge of various prices and interest rates, such as the current rate on 12-month fixed deposit accounts, the current rate of inflation, the level of the stock market, or the price of gold, a common household asset in India. We code a depositor as knowing each price if she is within 30% of the true value in the month in which she was surveyed. We also ask questions on newspaper subscription and the time spent reading the paper, since this is a primary source of local news and the events surrounding the runs were widely covered in the local newspapers. Last, we ask about common assets such as vehicle and land ownership, in order to gauge household socioeconomic status.

The survey was conducted in February and March of 2015, well after the high-solvency-risk shock in 2009. Since the survey data postdate the shock event, one may be concerned that these are poor controls, in the sense that asset holdings or other variables may have changed depending on whether a depositor ran. We believe that the survey timing is not a concern for the demographic variables because education and occupation decisions would largely predate the runs. It may be a concern for measures of financial literacy or assets, however, to the extent that these characteristics are endogenous to having run. We address this concern by considering separate specifications for liquidation with each of the three groups of factors as explanatory variables.

D. Depositor Banking Relationships and Other Characteristics

Table I presents summary statistics on depositor balances and transaction activity for all depositors (columns (1) and (2)) and for the survey sample of depositors (columns (3) and (4)). Across all 29,852 depositors, 4% liquidate their accounts during the run week (column (1), first row). The extent of the run among the insured is modest, with 4% of depositors liquidating and the average withdrawal equal to 19% of the balance ex ante.[22] On average, depositors hold a transaction balance of INR 5,460, and approximately 1% have a balance above the deposit insurance limit of INR 100,000. With respect to additional relationships with the bank, 1.6% of depositors have a loan linkage and 3.2% have a staff linkage. Account activity is modest, with depositors making a transaction only 1.5% of days, on average, and an unconditional mean transaction size of about INR 140 (USD 3).

[21] Compare to Brown, Morkoetter, and Guin (2014), who get a 16% survey response rate for European depositors.

[22] These numbers are comparable to those from other bank runs. For example, Kelly and Ó Gráda (2000) document that in the bank run on Emigrants Industrial Savings Bank that occurred between December 11, 1854 and December 30, 1854, 234 account holders (7%) closed their accounts. Similarly, the number of depositors that ran in the recent IndyMac case was less than 5%.

Table I
Summary Statistics on Balances and Transactions in Administrative Data

The table shows summary statistics from a survey of a subsample of depositors holding accounts at the time of the fundamental run. This survey sample of 4,635 depositors was selected to overweight depositors with relationships of interest with the bank; see text for details of the sampling procedure. The three panels represent categories of variables related to depositor demographics (education and occupation), financial knowledge, and assets, as recorded in survey interviews in January through March of 2015. Age, newspaper subscription, and asset ownership questions have sample sizes of 4,578, 4,615, and 4,615, respectively, due to refusals or lack of knowledge.

	Full Sample		Survey Sample	
	Mean (1)	*SD* (2)	Mean (3)	*SD* (4)
Liquidation dummy (withdraw 50% = 1)	0.039	0.193	0.058	0.233
Transaction balance, ’000s, 90 days prior	5.462	32.597	12.980	61.310
Balance above 100k, 90 days prior	0.009	0.096	0.039	0.194
Age of account in years at run	6.302	1.699	5.943	2.136
Depositor or family has loan	0.016	0.124	0.069	0.254
Depositor or family is staff	0.032	0.175	0.067	0.250
Mean daily liquidation dummy, year prior	0.003	0.012	0.004	0.012
Mean daily transaction dummy, year prior	0.015	0.054	0.021	0.060
Daily withdrawal, year prior to run	142.254	1,332.555	243.596	1,702.593
Daily deposit, year prior to run	140.861	1,318.174	243.524	1,709.209
Observations	29,852		4,634	

By design, the survey sample of 4,634 depositors includes a greater fraction of depositors with balances above the insurance cover, depositors who are staff, or depositors who hold a loan (column (3)). Since these types of depositors sampled with higher probability are more likely to run, the liquidation rate in the survey sample is also higher, at around 6% instead of 4%. In the empirical results section below we compare the determinants of withdrawal across the two samples in much greater detail.

Table II provides summary statistics on characteristics as captured in the survey. The statistics in the first two columns are weighted by the inverse of the probability of sampling to reflect the characteristics of depositors in the full sample, whereas the statistics in columns (3) and (4) are unweighted and therefore reflect the characteristics of the survey sample. We report both sets of results for completeness; however, in practice, the sampling weights barely change the estimated depositor characteristics (column (1) versus column (3)), which suggests that these characteristics are not highly correlated with the banking relationship variables used to determine sampling probabilities. For brevity, therefore, we discuss the characteristics of the full sample using the weighted estimates.

In the full sample (column (1)), the average age of depositors is 47 at the time of the survey. Depositors are quite educated, with 37% completing secondary school (up through the United States equivalent of 10th grade), 17%

Table II

Summary Statistics on Demographics and Financial Knowledge from Survey Data

The table shows summary statistics from a survey of a subsample of depositors holding accounts at the time of the fundamental run. This survey sample of 4,635 depositors was selected to over-weight depositors with relationships of interest with the bank; see text for details of the sampling procedure. The three panels represent categories of variables related to depositor demographics (education and occupation), financial knowledge, and assets, as recorded in survey interviews in January through March of 2015. Age, newspaper subscription, and asset ownership questions have sample sizes of 4,578, 4,615, and 4,615, respectively, due to refusals or lack of knowledge.

	Weighted to Reflect			
	Full Sample		Survey Sample	
	Mean (1)	*SD* (2)	Mean (3)	*SD* (4)
Panel A: Demographics				
Depositor age	46.883	11.664	46.87	11.87
Education, completed primary (=1)	0.073	0.261	0.069	0.254
Education, completed middle (=1)	0.113	0.317	0.106	0.308
Education, completed secondary (=1)	0.371	0.483	0.363	0.481
Education, completed higher secondary (=1)	0.165	0.371	0.172	0.377
Education, beyond higher secondary (=1)	0.255	0.436	0.269	0.444
Occupation other/missing (=1)	0.047	0.211	0.047	0.212
Occupation wage labor (=1)	0.075	0.264	0.071	0.256
Occupation retail (=1)	0.064	0.246	0.060	0.237
Occupation work at home (=1)	0.225	0.418	0.228	0.420
Occupation salaried (=1)	0.264	0.441	0.271	0.445
Occupation business (=1)	0.324	0.468	0.323	0.468
Panel B: Financial Knowledge				
Newspaper, whether subscription (=1)	0.730	0.444	0.740	0.438
Newspaper, hours reading	0.374	0.461	0.375	0.457
Knows RBI governor (=1)	0.091	0.287	0.092	0.290
Interest rate, savings account, known (=1)	0.176	0.381	0.179	0.383
Interest rate, fixed deposit account, known (=1)	0.291	0.454	0.300	0.458
Inflation rate, last 12 months, known (=1)	0.049	0.215	0.050	0.217
Sensex index value, known (=1)	0.064	0.245	0.065	0.246
Gold price, known (=1)	0.631	0.482	0.642	0.479
Panel C: Assets				
Scooter, whether owned (=1)	0.409	0.492	0.421	0.494
Motorcycle, whether owned (=1)	0.731	0.444	0.729	0.445
Car, whether owned (=1)	0.125	0.330	0.137	0.344
House/flat, whether owned (=1)	0.963	0.190	0.965	0.184
Ancestral land, whether owned (=1)	0.244	0.430	0.252	0.434
Holiday mode, bus (=1)	0.857	0.350	0.847	0.360
Holiday mode, train (=1)	0.420	0.494	0.413	0.492
Holiday mode, car (=1)	0.125	0.331	0.136	0.343
Holiday mode, plane (=1)	0.004	0.060	0.005	0.070
Observations	4,634		4,634	

completing higher secondary (high school diploma), and 26% having some education beyond higher secondary school. The most common occupations are business (32%), salaried professional employment (26%), and work at home (23%). Nearly three-quarters of depositors subscribe to the newspaper (Panel B), and they spend on average 0.37 hours (22 minutes) reading it each day. Most depositors know the current price of gold (63%), some know the current rate of interest on term deposits (29%), but very few know the current inflation rate (5%) or the value of the most common stock index (6%). The asset holdings of depositors, shown in Panel C, reflect a broadly middle-class and urban depositor population. Most households own a scooter or motorbike, but few own a car; most own their own house or flat but few own ancestral land, a marker of wealth and family lineage. People take holidays, but travel by bus more than train or car.

II. Empirical Results from the High-Solvency-Risk Shock

We present the empirical results going backwards in time, first for the high-solvency-risk shock at the time the shock became public, then before the public release of information and after the private RBI audit, and then before the private audit. We next present results for the earlier low-solvency-risk shock and contrast these findings with those for the high-solvency-risk shock.

A. Liquidation under the High-Solvency-Risk Shock after the Public Information Release

We start by documenting heterogeneity in depositor response to the high-solvency-risk shock. The tendency of depositors to withdraw after the public information release depends strongly on depositor characteristics. Table III compares the balances and banking relationships of depositors, in the administrative data, by whether a depositor ran in the week after the public release of information on the shock. Columns (1) through (3) present the means for depositors who ran, depositors who stayed, and the difference between the two groups. (Again, depositors that withdrew more than 50% of their transaction balance over the week beginning at the information release are classed as runners.) Runners and stayers differ significantly on all observable dimensions. Runners have seven times larger transaction balances, are 10 times more likely to have balances above the deposit insurance limit, and are much more active in terms of the number and size of transactions over the past year. In addition, runners have held their accounts for approximately one year less and are much more likely to have a loan or a staff linkage.

During the run week, we use both linear probability and probit models for the likelihood of liquidation to estimate the determinants of liquidation in a multivariate framework. We apply the linear probability model, though liquidation is a binary outcome, in part because it allows the inclusion of a large number of fixed effects in later specifications that use data on depositors present under both shocks. Table IV presents these estimates with liquidation (withdrawing

Table III

Summary Statistics by Run Status, High-Solvency-Risk Shock

Summary statistics for depositor characteristics by whether the depositor liquidated during the run (column (1)) or not (column (2)). Column (3) shows the difference between columns (1) and (2) and the *SE* of the difference. Liquidation is a dummy for withdrawing 50% of transaction balances in the week of the run. Transaction history is a dummy for whether the transaction balance changed on a given day, whereas daily withdrawal and daily deposit are the withdrawal and deposit amounts. All variables are defined in the Appendix. Standard deviations are in square brackets and *SE*s in parentheses with $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

	Sample Mean [*SD*]		
	Run (1)	Stay (2)	Run-Stay (3)
Run (withdraw 50% = 1)	1	0	1
	[0]	[0]	(0)
Transaction balance	31.1	4.43	26.6***
	[77.7]	[28.9]	(0.97)
Above insurance cover	0.068	0.0069	0.061***
	[0.25]	[0.083]	(0.0029)
Account age	5.29	6.34	−1.05***
	[2.31]	[1.66]	(0.051)
Loan linkage (=1)	0.048	0.014	0.034***
	[0.21]	[0.12]	(0.0037)
Staff (=1)	0.059	0.031	0.028***
	[0.24]	[0.17]	(0.0052)
Liquidation history	0.016	0.0027	0.014***
	[0.029]	[0.010]	(0.00035)
Transaction history	0.093	0.012	0.081***
	[0.13]	[0.046]	(0.0016)
Daily withdrawal, year prior to run	996.7	107.8	888.9***
	[3,883.5]	[1,099.6]	(39.6)
Daily deposit, year prior to run	1,011.7	105.7	906.0***
	[3,762.1]	[1,098.0]	(39.2)
Observations	1,157	28,695	

50% of balances) as the outcome variable. Columns (1) through (3) report results for the full sample of depositors using different specifications: the first two columns provide results from linear probability models with alternate controls for ex ante transaction account balances, and column (3) reports estimates from a probit model. Finally, column (4) reports results using the same specification as (2) but for the much smaller survey sample. In each specification, the explanatory variables are depositor characteristics, variables capturing their transaction history, and variables capturing their relationship to the bank.

The estimates in Table IV show that banking relationships are strongly associated with liquidation. Looking at column (1), depositors with loan linkages are 4.7 percentage points more likely to run, which is statistically significant at the 5% level. Recall that about 4% of depositors run, so this amounts to a doubling of the tendency to liquidate. Each additional year a depositor has an

Table IV

Who Runs after the Public Release? High-Solvency-Risk Shock

The table shows estimates for linear probability and probit models for the probability of liquidation during the week following the public release of information on the fundamental shock. Liquidation is defined as withdrawing at least 50% of one's prior balance. Balance is the transaction balance in '00,000s of INR. For definitions of the remaining variables, please see the Appendix. Linear probability model estimates are coefficients and estimates from the probit model are marginal effects. The sample in the first three columns is the population of depositors, and the sample in column (4) is the survey sample of depositors for which the household survey was completed. *SE*s are in parentheses with $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

	LPM (1)	LPM (2)	Probit (3)	LPM (4)
Loan linkage (=1)	0.047**	0.046**	0.035**	0.061**
	(0.021)	(0.020)	(0.015)	(0.026)
Account age	−0.0072***	−0.0074***	−0.0058***	−0.0045**
	(0.0011)	(0.0010)	(0.00051)	(0.0021)
Staff (=1)	0.019**	0.018**	0.019**	0.021
	(0.0092)	(0.0092)	(0.0077)	(0.018)
Liquidation history	3.12***	3.25***	1.14***	3.35***
	(0.23)	(0.22)	(0.067)	(0.54)
Transaction balance	0.077***			
	(0.017)			
Above insurance cover		0.21***	0.18***	0.21***
		(0.030)	(0.028)	(0.038)
Observations	29,852	29,852	29,852	4,634
Sample	Full	Full	Full	Survey

account with the bank decreases the tendency to run by about 0.72 percentage points. Being a staff member increases the tendency to run by about two percentage points. The mean daily liquidation dummy gives the average share of days over the prior year, excluding the 90 days immediately prior, during which a depositor withdrew 50% of her balances, as a control for past account activity. The mean of this variable is 0.003, since most depositors do not liquidate 50% of their balances on most days. We can get a better sense of the size effect by scaling the coefficient of 3.12 downwards by a factor of 30: having liquidated on average one more day per *month* increases the likelihood of running by a significant and large 10 percentage points.[23] A one-standard-deviation (about INR 32,000) increase in transaction balances prior to the run increases the tendency to liquidate by $0.077 \times 32 = 2.5$ percentage points, comparable to the effect of being a member of the bank staff.

These conclusions hold in the models with categorical controls for the ex ante balance, columns (2) and (3). The effect of higher balances comes largely through depositors with balances above the insurance limit that are 21 percentage points more likely to run than fully insured depositors. Depositors

[23] Using alternative transaction controls, such as the mean of a dummy for past transactions, does not change the results.

with high balances may be better informed and also stand to lose more in the event of a failure due to the temporary loss of funds below the insurance limit and a permanent loss above the limit. The incentive to withdraw is in principle continuous around INR 100,000, as depositors with balances just above the limit remain mostly insured, with only the marginal balance above the threshold at risk. Table IA.II of the Internet Appendix tests for a discontinuity at the insurance limit, and does not find evidence that liquidation changes discretely at that point. The coefficient on being above the insurance cover remains large and significant with separate linear balance controls on either side of the insurance threshold, but the coefficient grows smaller and is not statistically different from zero with cubic or more flexible controls. These results support the idea that the effect of having a balance above the insurance cover is the effect of having a high balance, and not due to any discrete change such as a change in attention associated with having an uninsured balance.

The magnitudes of the effects of the other depositor characteristics are generally steady across the specifications shown as well as in alternative specifications where liquidation is defined as withdrawal of 25% or 75% of balances instead of 50% (Table IA.I of the Internet Appendix). The results are also not affected by adding fixed effects for eight branches or for 292 detailed geographic neighborhoods to control for unobserved depositor characteristics that are correlated with the tendency to run. Finally, the results are qualitatively and quantitatively very similar in the much smaller survey sample of depositors (Table IA.III, column (2) of the Internet Appendix). None of the coefficient estimates in that regression are outside the confidence intervals for the coefficients in the analogous specification in the full sample, and most estimates are nearly identical.[24] This finding is important to establish that sample selection does not drive the results of the next section, which compares the relative importance of depositor characteristics and banking relationships as determinants of running.

Depositor balances and relationships with the bank are important correlates of the tendency to run. Consistent with their relationships providing information about the bank, depositors with loan linkages and staff linkages are more likely to withdraw during the run. Depositors who hold balances above the deposit insurance threshold are far more likely to run, and depositors with high transaction volume with the bank are also more likely to run. In contrast, having an account with the bank for a longer duration reduces the likelihood of running.

B. Running and Depositor Characteristics

A concern with the above analysis is that depositor balances or relationships may predict running because they proxy for omitted variables, such as

[24] The Internet Appendix is available in the online version of the article on *The Journal of Finance* website. Internet Appendix Table IA.II of the confirms that the estimates in the survey sample are also similar to the main sample when the regression is weighted by the inverse of sampling probabilities.

education, occupation, or financial literacy, that themselves are responsible for liquidation behavior. It is plausible that more educated depositors both hold loans and follow the news, for example. This section addresses this concern by relating liquidation to depositor characteristics from the household survey that are grouped into the three broad themes of demographics (age, education, and occupation), financial literacy, and assets (see Table II for the full set of survey variables).

Table V shows that these factors are, in fact, strong predictors of the tendency to run, and in an economically sensible manner. In column (1), the specification includes demographic determinants of liquidation. Older depositors run significantly more than others. Relative to a depositor with a primary school education, the omitted category, a depositor with an education beyond higher secondary (U.S. high school equivalent) is 2.4 percentage points ($SE = 1.4$ percentage points) more likely to run, which is statistically different from zero at the 10% level. Occupation is the strongest determinant of running among these factors. Relative to a depositor working in wage labor, the omitted category, a depositor who reports business as her occupation is 4 percentage points ($SE = 1.1$ percentage points) more likely to run. This coefficient is statistically different from zero at the 1% level and comparable in magnitude to the effect of loan linkages (as shown in Table IV). Salaried and work-at-home occupations, also indicators of relatively higher class depositors, are positively and significantly associated with running. Not surprisingly, the p-value of an F-test for the joint significance of these demographic factors is less than 0.001.

In column (2), the specification of Table V tests whether running is related to financial knowledge, where the knowledge measures are newspaper subscription and readership and actual knowledge of various asset prices and interest rates at the time of the survey. Depositors that have a newspaper subscription are 2.4 percentage points ($SE = 0.79$ percentage points) more likely to run, and reading the newspaper for one additional hour each day is associated with a 1.9 percentage point ($SE = 0.96$ percentage points) increase in the probability of running. These results are sensible given that those who read the newspaper would have seen stories reporting on the high-solvency-risk shock. Knowledge of asset prices is generally weak (Table II). However, if depositors know the interest rate on fixed deposit accounts, they are more likely to run by 1.6 percentage points ($SE = 0.86$ percentage points, $p < 0.10$). This measure of knowledge may be more powerful because fixed deposit accounts are directly related to banking, unlike stock indices or inflation, which are related to more general economic activity. The financial knowledge indicators are jointly significant.

In column (3), asset holdings are also significant predictors of liquidation. In column (4), we report a specification using all controls together. We report these last two specifications for completeness but do not emphasize the column (3) and (4) results, as we believe that asset variables are far more likely than demographics or financial literacy to have been affected by the run itself.

Table VI combines these depositor characteristics with administrative data on banking relationships to address the main question of interest: are

Table V

The Effect of Depositor Characteristics and Financial Knowledge on Runs

The table shows coefficient estimates for linear probability models of the probability of running during the week following the public release of information on the high-solvency-risk shock. Running is defined as withdrawing at least 50% of one's prior balance. Explanatory variables are from the household survey and the regression sample is the survey sample of 4,634 depositors; sample sizes are smaller because of refusals to answer some questions. Explanatory variables are grouped into categories of demographics (column (1)), financial knowledge (column (2)), and assets (column (3)), and *F*-tests for the joint significance of the explanatory variables are shown by column. *SE*s are in parentheses with $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

	Survey (1)	Survey (2)	Survey (3)	Survey (4)
Depositor age	0.0011***			0.00087**
	(0.00039)			(0.00041)
Education, completed middle (=1)	−0.0092			−0.017
	(0.014)			(0.014)
Education, completed secondary (=1)	0.0017			−0.0079
	(0.013)			(0.013)
Education, completed higher secondary (=1)	0.021			0.0068
	(0.015)			(0.015)
Education, beyond higher secondary (=1)	0.024*			0.0034
	(0.014)			(0.015)
Occupation, other/missing (=1)	0.037			0.031
	(0.023)			(0.022)
Occupation, retail (=1)	0.011			0.0026
	(0.015)			(0.015)
Occupation, work at home (=1)	0.032***			0.025**
	(0.011)			(0.011)
Occupation, salaried (=1)	0.022**			0.018
	(0.011)			(0.012)
Occupation, business (=1)	0.040***			0.032**
	(0.011)			(0.013)
Newspaper, whether subscription (=1)		0.024***		0.0081
		(0.0079)		(0.0097)
Newspaper, hours reading		0.019*		0.020**
		(0.0096)		(0.0100)
Knows RBI governor (=1)		−0.0012		−0.013
		(0.013)		(0.014)
Interest rate, savings account, known (=1)		−0.015		−0.011
		(0.010)		(0.011)
Interest rate, fixed deposit account, known (=1)		0.016*		0.016*
		(0.0086)		(0.0088)
Inflation rate, last 12 months, known (=1)		−0.0095		−0.0095
		(0.017)		(0.017)
Sensex index value, known (=1)		0.0025		0.0064
		(0.015)		(0.015)
Gold price, known (=1)		−0.00059		−0.0017
		(0.0084)		(0.0095)

(*Continued*)

Table V—*Continued*

	Survey (1)	Survey (2)	Survey (3)	Survey (4)
Scooter, whether owned (=1)			0.028***	0.020**
			(0.0088)	(0.0097)
Motorcycle, whether owned (=1)			0.0076	0.0036
			(0.0091)	(0.0096)
Car, whether owned (=1)			0.016	0.011
			(0.015)	(0.015)
House/flat, whether owned (=1)			0.028*	0.020
			(0.015)	(0.014)
Ancestral land, whether owned (=1)			0.0058	0.0089
			(0.0090)	(0.0096)
Holiday mode, bus (=1)			0.0020	0.0027
			(0.011)	(0.012)
Holiday mode, train (=1)			−0.012*	−0.013
			(0.0075)	(0.0086)
Holiday mode, car (=1)			−0.0059	−0.011
			(0.014)	(0.015)
Holiday mode, plane (=1)			0.10	0.094
			(0.077)	(0.074)
F-test *p*-value	0.000	0.001	0.008	0.000
Observations	4,578	4,615	4,615	4,578

relationships only a proxy, or are they meaningful on their own? Column (1) replicates the main specification of Table IV in the survey sample, and columns (2) through (4) progressively add the explanatory depositor characteristics from Table V to this specification. Remarkably, though depositor characteristics are themselves significant predictors of running, including these variables in the main specification does not alter the strong and statistically significant effects of banking relationships. Loan linkages, account age, liquidation history, and having balances above the insurance cover all remain critical determinants of running, with nearly the exact same coefficients as in the specification without these additional controls. For example, the effect of loan linkages is 0.061 (SE = 0.026) in the main specification, and 0.060 (SE = 0.026) in the preferred specification of column (3), which includes demographic and knowledge controls but not asset controls. The estimated coefficient on being a member of the bank staff is 0.016 (SE = 0.018). This is slightly smaller than the earlier estimates of 0.021/0.018 (survey sample/full sample), and, because the estimate is imprecise, we cannot reject the possibility that the coefficient is equal to these estimates or that it is equal to zero. The effect of education becomes smaller and statistically insignificant when banking relationships are introduced. Occupation and newspaper readership remain strong predictors of liquidation (columns (2) and (3)).

The above evidence strongly supports the view that banking relationships are not a proxy for omitted characteristics such as depositor education or financial literacy, but rather matter on their own accord. The survey measures of depos-

Table VI

Effects of Banking Relationships and Depositor Characteristics on Runs

The table shows coefficient estimates for linear probability models of the probability of running during the week following the public release of information on the high-solvency-risk shock. Running is defined as withdrawing at least 50% of one's prior balance. Explanatory variables are from both administrative data on banking relationships and the household survey. The regression sample is the survey sample of 4,634 depositors; sample sizes are smaller because of refusals to answer some questions. See Table II for a complete listing of explanatory variables from the survey. *SE*s are in parentheses with $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)
Loan linkage (=1)	0.061**	0.060**	0.060**	0.059**
	(0.026)	(0.026)	(0.026)	(0.026)
Account age	−0.0045**	−0.0049**	−0.0053**	−0.0052**
	(0.0021)	(0.0021)	(0.0022)	(0.0022)
Staff (=1)	0.021	0.017	0.016	0.016
	(0.018)	(0.018)	(0.018)	(0.018)
Liquidation history	3.35***	3.32***	3.33***	3.32***
	(0.54)	(0.54)	(0.54)	(0.55)
Above insurance cover	0.21***	0.21***	0.21***	0.21***
	(0.038)	(0.038)	(0.038)	(0.038)
Depositor age		0.0011***	0.0010***	0.00099**
		(0.00037)	(0.00038)	(0.00039)
Education, completed middle (=1)		−0.0065	−0.012	−0.012
		(0.014)	(0.014)	(0.014)
Education, completed secondary (=1)		0.000069	−0.0065	−0.0077
		(0.012)	(0.012)	(0.013)
Education, completed higher secondary (=1)		0.0057	−0.0024	−0.0037
		(0.014)	(0.014)	(0.014)
Education, beyond higher secondary (=1)		0.010	−0.0021	−0.0042
		(0.013)	(0.014)	(0.014)
Occupation, other/missing (=1)		0.032	0.026	0.023
		(0.022)	(0.021)	(0.021)
Occupation, retail (=1)		0.0077	0.0042	0.0019
		(0.015)	(0.015)	(0.015)
Occupation, work at home (=1)		0.026**	0.022**	0.020*
		(0.011)	(0.011)	(0.011)
Occupation, salaried (=1)		0.014	0.010	0.0096
		(0.011)	(0.011)	(0.011)
Occupation, business (=1)		0.026**	0.019	0.016
		(0.010)	(0.012)	(0.013)
Newspaper, whether subscription (=1)			0.013	0.0070
			(0.0088)	(0.0095)
Newspaper, hours reading			0.022**	0.024**
			(0.0093)	(0.0095)
Knowledge controls	No	No	Yes	Yes
Asset controls	No	No	No	Yes
Observations	4,634	4,578	4,578	4,578

itor characteristics predict liquidation but do not displace the effect of banking relationships. Of course, other omitted variables not collected in the survey may still can found the estimates of banking relationships. In Section II.E, we conduct further tests to control for other unobservable but time-invariant characteristics of depositors.

C. Liquidation before and after the Public Information Release

The models above consider liquidation in the cross-section after the public release of information. We now examine the timing of depositor withdrawals before the public release of the negative information to see which depositors start running and when, paying particular attention to the possible private release of information about RBI's audit of the bank.

As shown in Figure 1, balances declined significantly prior to the public release of information. To examine what types of depositors run in the period before the public release of information, we estimate Cox hazard models, both strictly proportional models and with time-varying coefficients. Failure is defined as the withdrawal of 50% of balances during any given day.[25,26] The model with time-varying coefficients holds the ex ante depositor characteristics fixed over the event window, from 120 days before to 30 days after the shock, and estimates how the effects of these characteristics change over time. This model specifies the hazard as

$$\begin{aligned}\Lambda_i(t) = {} & \Lambda_0(t)\exp\{\beta_1(t)_i\text{AccountAge}_i + \beta_2(t)\,\text{StaffLinkage}_{i+}\beta_3(t)\,\text{LoanLinkage}_i \\ & + \beta_4(t)\,\text{NetworkMemberHasRun}_{it} + \beta_5(t)\,\text{AboveInsuranceCover}_i \\ & + \beta_6(t)\,\text{DailyTransactions}_i\} \qquad (1)\end{aligned}$$

The only difference from the baseline Cox proportional hazard model is that each coefficient is allowed to vary over time. Each time-varying coefficient is modeled with a basis of cubic B-splines with knots every 30 days from 120 days before to 30 days after the day of the public information release, for a total of nine parameters for each variable. This specification allows the coefficient on each characteristic to change smoothly as a cubic function within each 30-day window and constrains the first and second derivatives of each $\beta(t)$ to be constant at the knots that mark the boundaries between 30-day windows.

Hazard ratios from the base hazard model, reported in Table VII, column (1), agree with the cross-sectional models that focus on the run in the week after the public disclosure of the high-solvency-risk shock. (Note that, because of the

[25] As the unconditional likelihood of transactions on any given day is very low, in practice this definition is similar to the definition employed in the cross-section of withdrawal of 50% over the run week.

[26] We exclude depositors with balances less than INR 100 as of 120 days before the run to make the model simpler to estimate by maximum likelihood. As these accounts generally have very low activity, the omission has little effect, but the omitted category for balances in the hazard models should be taken as INR (100 to 100,000).

Table VII

Who Runs Prior to the Public Release? High-Solvency-Risk Shock

The table shows exponentiated coefficient estimates (i.e., hazard ratios) for Cox proportional hazard models of the probability of liquidation from 90 days before to 30 days after the public release of information on January 27, 2009. The model in the first column assumes that the coefficients on each characteristic have a constant effect on liquidation over time. The model in the second column allows the coefficient on each characteristic to vary according to a cubic spline function with knots at 30-day intervals over the event window. The hazard ratios reported for the model in the second column are the effect of each variable evaluated as on the date of the public release of information. The path of the full time-varying hazard ratios over time are shown in Figure 4 for select variables. *SE*s are in parentheses with $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$ indicating significant differences from a hazard ratio of one.

	Cox (1)	Time Varying (2)
Loan linkage (=1)	1.56***	1.46**
	(0.12)	(0.28)
Account age	0.80***	0.83***
	(0.01)	(0.01)
Staff (=1)	2.56***	3.06***
	(0.16)	(0.45)
Transaction history	985.96***	409.31***
	(90.29)	(147.38)
Above insurance cover	1.07	3.79***
	(0.09)	(0.75)
Network member has run	2.81***	3.38***
	(0.25)	(0.56)
Time-varying splines	No	Yes
Observations	2,867,291	2,867,291

differences in event window, the time horizon of the dependent variable, and the reporting of hazard ratios, the magnitude of these estimates is not directly comparable to the coefficients reported in Table IV.) Having an older account decreases the likelihood of liquidation. Staff linkages increase the propensity to liquidate by a factor of 2.56 ($p < 0.01$ against the null of a unit hazard ratio) and loan linkages increase it by a factor of 1.56 ($p < 0.01$). The relative strength of these effects is reversed, as compared to the cross-sectional analysis, where loan linkages are more powerful than staff linkages. The staff effect is larger in the hazard model because this model covers a broader window than just the run week and staff were more likely to run earlier in this period than other depositors. Given the extended hazard window, we also introduce a time-varying explanatory variable for whether a member of the depositor's network has run by a given date. We find that a network member having run increases the hazard that a depositor will run by nearly threefold, the same increase in hazard as being a member of the bank staff.[27] Having a balance prior to the event window that is above the insurance limit is not associated with a

[27] Kelly and Ó Gráda (2000) also document the importance of network effects in bank runs. See also He and Manela (2012) for a theory of information acquisition in rumor-based runs.

higher hazard—this result, seemingly contravening the importance of being uninsured in the cross-section, is due to model misspecification; we reconcile the two findings below. Daily volume of transactions is highly predictive of liquidation.

Table VII, column (2) reports hazard ratios from the time-varying hazard model on the day of the public information release. Because the coefficient on each variable is a function, it can be evaluated at different times in the event window. Formally, these are the exponentiated coefficients on the constant value for each characteristic, which can be interpreted as the effect of that characteristic on the run date, because the B-spline corresponding to the knot at that date has been omitted from each coefficient basis. Staff are more likely to liquidate around the run, relative to other depositors and to the hazard ratio estimated over the event window. Depositors with uninsured balances are far more likely to liquidate relative to the proportional specifications. The hazard ratio for depositors above the deposit insurance limit is about four, relative to the fully insured. This ratio is far larger than the ratio of around one reported in the proportional hazard model, and indicates that high balance depositors, such as staff, are more likely to liquidate at times when information about the bank's solvency is revealed. Thus, the strictly proportional hazard model is not well specified because it does not account for the fact that the effect of depositor characteristics on liquidation changes with the information available over time. As this coefficient difference suggests, a likelihood-ratio test of the alternative time-varying model against the null proportional hazards model rejects the null model with a p-value < 0.001 ($\chi^2_{(42)} = 261.74$).

Looking at the full path of coefficients over the event window shows that staff, and possibly uninsured depositors, are more responsive even before the public release of information. For the same time-varying hazard specification reported in Table VII, column (2), Figure 2 plots the coefficients on staff linkages, loan linkages, and uninsured depositors continuously on each date over the event window. The hazard ratio corresponding to staff linkages, shown in Panel A, is around four and significantly different from one both at the time of the private audit by the central bank and just before the public release of information, but staff are no more likely to run than other depositors in the middle of the event window. This Bactrian camel-backed pattern suggests that staff respond to private information about the fundamentals of the bank and are not merely more likely to withdraw for whatever reason. Panel B shows that, while depositors with loan linkages are generally more likely to withdraw over the event window, this effect is not any stronger at a particular time around the shock. Panel C plots the time-varying hazard of liquidation for depositors above the insurance limit. Given their high balances, these depositors are typically unlikely to withdraw 50% of their balances in one day, as shown by the low hazard ratios in October and December 2008. However, like staff, they are more likely to withdraw than usual during the period after the central bank audit, with the hazard ratio rising to about one, on a par with depositors with much smaller balances. After a lull in the middle of the event window, where

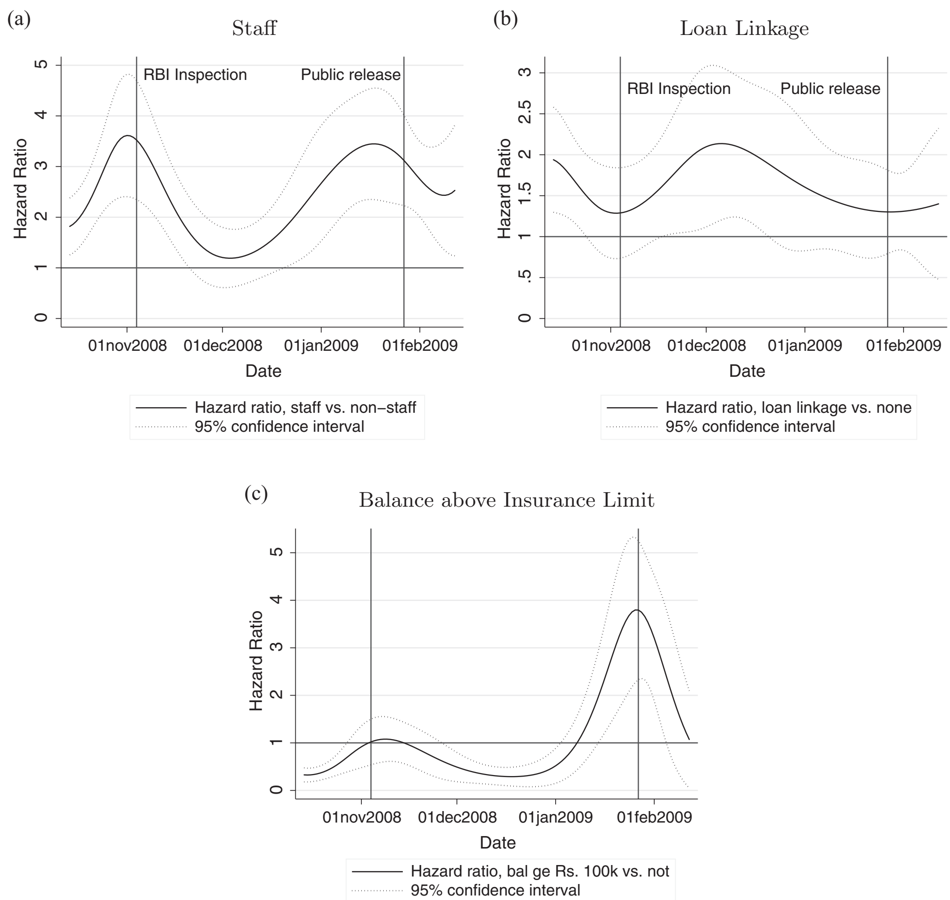

Figure 2. Who runs before the public release? Time-varying hazard ratios. The figure shows estimated time-varying hazard ratios for depositor characteristics from a Cox proportional hazard model of liquidation (withdrawal of 50% of transaction balance in one day) on depositor characteristics. The event window is 90 days before the public release of information on January 27, 2009 through 30 days after. The coefficient on each depositor characteristic is allowed to vary smoothly over time according to a cubic spline with knots at 30-day intervals. The resulting hazard ratio and confidence intervals for the coefficient are plotted here for three coefficients of interest.

the uninsured are significantly *less* likely to withdraw than others, the hazard associated with high balances increases steeply just before the public release of information, reaching the factor of 3.79 as reported in Table VII, column (2).

The hazard specifications show significant effects of both depositors holding balances above the insurance threshold and depositor ties to the bank, via staff and loan linkages. We find a pecking order of withdrawals in response to the private information of the regulatory audit: the staff of the bank withdraw first, followed closely by uninsured depositors.

D. Reaction of Depositors Prior to the Regulatory Audit

Did depositor runs begin before the regulatory audit? The regulatory audit showed that the financial position of the bank was deteriorating over the prior fiscal year, in spite of the fact that the annual reports of the bank did not reveal the true extent of the solvency risk. To understand whether some depositors were running before the regulatory audit, we examine depositor withdrawals around the release of the bank's annual report for the prior fiscal year, ending March 31, 2008, which was released on September 2, 2008, which was about two months before the audit. We do not find any significant depositor withdrawals in this period, except for some by the bank staff. These results suggest that the regulatory audit was an important shock that revealed information about bank fundamentals and acted as a coordinating signal.

As shown in Figure 1, aggregate balances were roughly flat in the period after the annual report was released on September 2. To measure the response of different depositors, we replicated our earlier cross-sectional regression for liquidation in the week following the release of the annual report. Staff are a significant 1.6 percentage points more likely to liquidate than other depositors over this week, a response that is considerably weaker than their relative tendency to liquidate during the run. Depositors with loan linkages and uninsured balances show no response to the annual report. The coefficient on loan linkages is not significantly different from zero in any specification and point estimates are always less than 1.1 percentage points. Uninsured depositors have point estimates of –0.02 (2 percentage points) and 0.009 (1 percentage point) in the linear probability and probit models, respectively. These coefficients are both small and not statistically different from zero. Thus, depositor runs primarily begin after the regulatory audit. Recall that there is also a statistically significant structural break in the time series of depositor balances in the week after the regulatory audit began (Figure IA.1 of the Internet Appendix), but it is not nearly as large as the break at the public run. We also do not find any significant increase in interest rates paid by the bank in this period that could have compensated depositors for higher risk.[28]

III. Depositor Behavior across Shocks

A. Liquidation under the Low-Solvency-Risk Shock

While the results above suggest that, in the case of a high-solvency-risk shock to the bank, there are significant differences in the likelihood of depositors running based on depositor characteristics, the evidence is not sufficient to conclude that these depositors are responding to the true solvency risk of the

[28] Interest rates were steady or declining over the year and a half prior to the run. The interest rates paid on fresh term deposits were around 10% over this period and were declining slightly leading up to the run. Interest rates on demandable savings deposits are not recorded at a high frequency in the data. Bank management told us that these rates were constant at 8.5% over the same period.

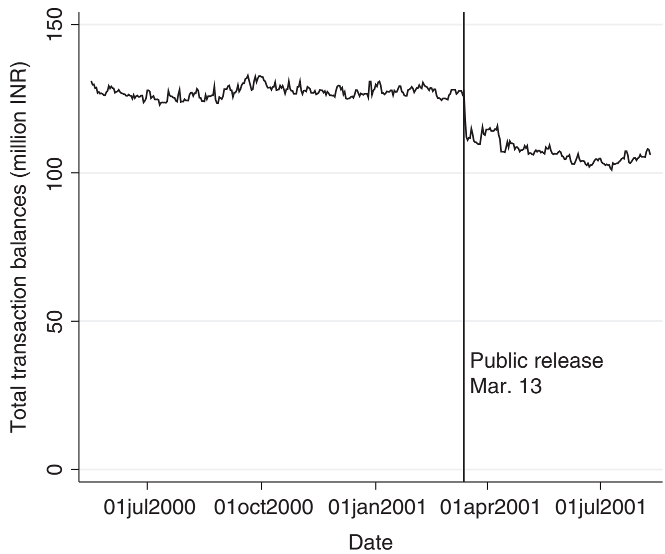

Figure 3. Transaction balances, low-solvency-risk shock. The figure shows aggregate transaction account balances for depositors in the bank from 300 days before the public release of information about a fraud at another bank, which occurred on March 13, 2001, through 150 days after. The vertical line indicates the date of the failure of another cooperative bank to which the bank under study had no exposure. The line is labeled with the date of the event itself but is drawn to intersect the closing balance of the day before the event.

bank. For instance, these depositors may withdraw in the same way and to the same degree in response to a low-solvency-risk shock, due to coordination failures or their relationships to the bank rather than to solvency risk. The question that remains, therefore, is whether these depositors behave differently when there is shock that does *not* put the solvency of the bank at risk.

To address this question, we contrast the behavior of depositors in response to the high-solvency-risk shock with the response to an earlier low-solvency-risk shock to the same bank, as described in Section III above. Recall that our bank was solvent and had no fundamental linkages with the bank that failed in this earlier shock, although depositors may have believed that the bank was at risk at the time. The magnitude of the response to the low-solvency-risk shock was smaller. Figure 3 plots the time series of aggregate transaction balances around the low-solvency-risk shock. Balances are roughly steady until the public shock, then decline by 11% in the week after the shock and are flat again. During the high-solvency-risk shock, by contrast, they declined by 25% in the same week, on top of the 16% decline that had already occurred after the regulatory audit (Figure 1). Four percent of insured depositors run in the week after the low-solvency-risk shock, as in the later event, and they withdraw similar amounts on similar ex ante balances.

Table VIII presents regression results for liquidation using the entire sample of depositors present at the time of the low-solvency-risk shock, analogous to

Table VIII

Who Runs in a Low-Solvency-Risk Shock?

The table shows coefficient estimates for linear probability and probit models of the probability of liquidation during the week following the public release of information on the low-solvency-risk shock on March 13, 2001. Liquidation is defined as withdrawing at least 50% of one's prior balance. Balance is the transaction balance in '00,000s of INR. For definitions of the remaining variables, please see the Appendix. Estimates from probit models are marginal effects. *SE*s are in parentheses with $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

	LPM (1)	LPM (2)	Probit (3)
Loan linkage (=1)	-0.012^{**}	-0.011^{**}	-0.0083^{*}
	(0.0052)	(0.0053)	(0.0048)
Account age	-0.0017^{***}	-0.0018^{***}	-0.0020^{***}
	(0.00041)	(0.00041)	(0.00044)
Staff (=1)	-0.025^{**}	-0.026^{***}	-0.021^{**}
	(0.0098)	(0.0098)	(0.0083)
Liquidation history	3.09^{***}	3.20^{***}	1.57^{***}
	(0.26)	(0.26)	(0.098)
Transaction balance	0.11^{***}		
	(0.016)		
Above insurance cover		0.090^{***}	0.078^{***}
		(0.030)	(0.027)
Observations	23,729	23,729	23,729

the Table IV (columns (1) to (3)) specifications for the high-solvency-risk shock. Having a younger account or a higher volume of transactions with the bank makes a depositor more likely to run, as is true for the high-solvency-risk shock. However, unlike for the high-solvency-risk shock, depositors with loan linkages and staff are *less* likely to run than other depositors. For example, having a loan decreases the likelihood of running by 1.2 percentage points, or about 30%, of the baseline four percentage points. Staff are 2.5 percentage points less likely to run than other depositors, which is the same magnitude but opposite sign of their response around the high-solvency-risk shock. Thus, the relative tendency of depositors with loan linkages and of staff depositors to withdraw is different across the two shocks, with both types withdrawing more around the high-solvency-risk event.

To bear down on the difference between shocks, we estimate liquidation specifications similar to those in Tables IV and VIII for a sample of depositors that were present both during the high-solvency-risk shock of 2009 and the earlier, low-solvency-risk shock of 2001. Table IX presents coefficients from linear probability models analogous to those in Table IV but estimated in pooled samples of depositors using observations from both runs. Column (1) includes all depositors present in either event, column (2) restricts attention to the constant sample of slightly over 10,000 depositors present in both events, and column (3) uses the column (2) sample and adds fixed effects. In each specification, the coefficients in the upper half of the table show the main effects of each

Table IX

Comparison of Depositor Runs across High- and Low-Solvency-Risk Shocks

The table shows coefficient estimates for linear probability models of the probability of liquidation during a bank run pooling depositor-level data across both shocks. Column (1) is a pooled regression of all depositors observed in either shock, column (2) is restricted to a constant sample of depositors observed in both shocks, and column (3) is the constant sample and includes fixed effects in the specification. Liquidation is defined as withdrawing at least 50% of one's prior balance. For definitions of the remaining variables, please see the Appendix. *SE*s are in parentheses with $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

	Pooled (1)	Constant (2)	Constant, Fixed Effects (3)
Loan linkage (=1)	-0.017^{***}	-0.022^{***}	-0.017
	(0.0050)	(0.0080)	(0.011)
Account age	-0.0023^{***}	-0.0022^{***}	-0.0018^{**}
	(0.00039)	(0.00054)	(0.00073)
Staff (=1)	-0.028^{***}	-0.042^{***}	0.0023
	(0.0097)	(0.016)	(0.034)
Transaction history	0.93^{***}	1.47^{***}	1.29^{***}
	(0.039)	(0.088)	(0.14)
Above insurance cover	0.00073	-0.061	-0.056
	(0.030)	(0.040)	(0.057)
High-risk shock ×			
Loan linkage (=1)	0.073^{***}	0.069^{**}	0.12^{***}
	(0.016)	(0.029)	(0.031)
Account age	-0.00011	0.00051	-0.000031
	(0.00033)	(0.00046)	(0.00056)
Staff (=1)	0.050^{***}	0.047^{**}	0.054^{*}
	(0.013)	(0.021)	(0.029)
Above insurance cover	0.14^{***}	0.22^{***}	0.21^{***}
	(0.040)	(0.064)	(0.077)
Constant	0.035^{***}	0.028^{***}	0.029^{***}
	(0.0018)	(0.0026)	(0.0031)
Depositor fixed effects	No	No	Yes
Observations	53,581	21,726	21,726

variable around the low-solvency-risk shock, and those in the lower half of the table show interaction terms between a depositor characteristic and the high-solvency-risk shock. The determinants for liquidation in the upper half of the table are similar to those from our prior analysis of the low-solvency-risk shock, so we focus on the interaction terms here.

Across all specifications, loan linkages, belonging to the staff, and having uninsured balances predict a higher tendency for depositors to liquidate during the high-solvency-risk shock relative to the low-solvency-risk shock. In the pooled regression of column (1), depositors with loans are a highly significant 7.3 percentage points more likely to liquidate during the high-solvency-risk shock than the low-solvency-risk shock, as compared to other depositors. This result is basically unchanged (6.9 percentage points) when we restrict the estimation to the constant sample of depositors, in column (2). Adding fixed effects to control for unobserved depositor characteristics in column (3), the difference in the effect of loan linkages across shocks is somewhat larger, rising to 12 percentage points, and remains statistically significant. This change in the probability of withdrawal is very large, compared to the 4% of depositors that run overall, and is common to both shocks. Moreover, the change in sign, perhaps even more starkly than the change in the magnitude of withdrawals by the uninsured, shows that the nature of the shock matters.

The effect of staff status under the high-solvency-risk relative to low-solvency-risk shock is fairly steady across specifications at 5.0, 4.7, and 5.4 percentage points in the pooled, constant sample, and fixed effects specifications, respectively. Uninsured depositors in the constant sample (column (2)) are 22 percentage points more likely to run during the high-solvency-risk shock than the low-solvency-risk shock, which is greater than the estimated effect of 14 percentage points in the full pooled sample (although the two estimates are within two *SE*s of each another). Adding fixed effects in column (2) does not further change the effect of having uninsured balances, as it stays at 21 percentage points ($SE = 7.7$ percentage points, $p < 0.01$).

B. Additional Robustness Checks

Above we show that the cross-sectional specifications for liquidation in Table IV are not affected by the introduction of observable depositor characteristics on demographics, financial knowledge, and assets. We also find that these results are invariant to the addition of bank branch or neighborhood fixed effects, which control for unobserved time-invariant characteristics of depositors common at the branch or neighborhood level.

The constant sample controls for time-invariant unobservable characteristics at the depositor level; even in specifications with depositor fixed effects, we find no significant change in the results reported earlier. The constant sample is subject to survivorship bias, however, since a depositor must have stayed with the bank after the earlier shock to be present in this constant sample. This bias could go either way. Depositors who saw the bank survive the low-solvency-risk shock might be less likely to run, during the high-solvency-risk shock, as they have seen the bank experience and survive a shock. On the other hand, depositors who stayed during the low-solvency-risk shock might be more informed and therefore more likely to run during a high-solvency-risk shock.

We study selection into the constant sample in two ways. First, we compare the estimated coefficients from the pooled sample with the constant sample. As shown in Table IX, these coefficients are very similar. If an unobservably different group of depositors stayed with the bank, then this selection should change the coefficients on observable characteristics if the unobservable selection factor is correlated with our explanatory variables. For example, if only inattentive depositors stayed, then under the restriction to the constant sample, we would expect that the main-effect coefficients on depositor characteristics during the panic would be smaller, which is not the case. Second, we apply a reweighting procedure in order to make the constant sample resemble the full sample under the low-solvency-risk shock on observable characteristics (DiNardo, Fortin, and Lemieux (1996)).[29] The results are again qualitatively unchanged.

The main results pertain to the withdrawal behavior of depositors during crises. One may be concerned that the differential results across depositors reflect different volatility in the deposits of these depositor types, which would be equally visible in noncrisis times but may not be perfectly captured by the control for depositors' liquidation history. To test this hypothesis, we run the same models for liquidating 50% of one's prior balances in two placebo periods, namely, one year and eight years before the high-solvency-risk shock (see Table IA.IV of the Internet Appendix). In these periods, we find that loan linkages, staff status, and being above the insurance cover have small and statistically insignificant effects, unlike in the crisis events studied. Account age and liquidation history, the effects of which do not vary across shocks of differing solvency risk, do predict withdrawals in normal times. The effect of liquidation history is expected, since it is intended to control for differences in transaction volatility across depositors. Both account age and liquidation history have much weaker effects on withdrawals in normal times than during crises with the magnitude of the coefficients on these variables being reduced approximately two-thirds and one-third, respectively, relative to the high-solvency-risk shock. Thus, we conclude that our main results reflect the behavior of depositors in response to solvency shocks, and not persistent differences in the withdrawals or volatility across depositor classes.

We note that the data here are observational, and our empirical strategy has been to control for observable and time-invariant unobservable factors that may be correlated with banking relationships and also affect the tendency to run. It remains possible that a time-varying factor, such as unmeasured depositor attention or intelligence, may be correlated with both banking relationships (like loan linkages) and the tendency to run during high-, but not low-, solvency-risk shocks. The relevant question for the validity of our finding on the importance of banking relationships would then be whether depositors with such relationships are always more attentive, or whether this correlation is unique to the events here.

C. Interpretation of Empirical Results

The main questions of the paper are whether depositors can distinguish shocks of differing solvency risk, and what kinds of depositors do so. The differ-

[29] The procedure corrects for the probability of selection into the constant sample as follows. First, we estimate a probit model for selection into the constant sample using all depositors present under the low-solvency-risk shock, with depositor banking relationships as explanatory variables. Second, we use this model to form the odds ratio of the likelihood of a depositor not surviving into the constant sample. Third, we use this ratio as a depositor-level weight in the constant sample regression. This procedure overweights depositors that were less likely to survive into the constant sample, in order to estimate the effect of the fundamental shock in the constant sample if there had been no selection on observables.

ential response of depositors to the two types of shocks strongly suggests that depositors are informed about the degree of their bank's solvency risk.

Administrative and other evidence from the bank suggests that our results capture a real change in depositor behavior across shocks, and not mechanical responses to changing actions of the bank. A question that may remain is whether the bank withdrew credit or laid off staff under the high-solvency-risk shock, forcing loan-linked depositors or staff to withdraw to meet liquidity needs. We find that bank staff were not laid off until February 10, 2010, more than a year after the high-solvency-risk shock. Regarding loan-linked depositors, only 3.2% of loans were modified or closed out in the period between the RBI audit and the public release of information under the high-solvency-risk shock, and the effects of loan linkages on withdrawal behavior are robust to estimating a main effect of loan linkage that excludes this group.

Additional results support the idea of a group of depositors that are informed about solvency risk (Chari and Jagannathan (1988)). During the high-solvency-risk shock, informed classes of depositors (staff, loan-linked depositors) withdraw prior to the public release of information in response to a regulatory intervention that was only privately observable. The withdrawal of other members of one's network is predictive of early withdrawal over this window. Word-of-mouth or in-network communication is a plausible channel to explain the actions of early movers. Direct measures of informedness, such as newspaper readership and (more weakly) education, predict a greater probability of withdrawal in the public run. Some of the strongest responses to the high-solvency-risk shock are observed by staff members and loan-linked depositors, who have direct contact with the bank through their own employment and through loan officers, respectively. These relationships may be a source of information on the bank's finances. We also estimate a direct effect of one's network members having run in hazard models.[30] In addition to being informed about solvency risk, loan-linked depositors may have greater incentive to withdraw under a high-solvency-risk shock. This incentive arises because a high-risk crisis may create greater income or liquidity shocks for loan-linked depositors, since they will lose the value of their lending relationship with the bank if it fails, which is more likely in a shock of higher risk.

Depositors with a higher frequency of transactions with the bank are also more likely to run, irrespective of the nature of the shock. While one may expect that depositors transacting frequently would also be informed, the evidence suggests that the liquidity needs or lower transaction costs of these depositors urge them to run in response to all shocks to the bank. Indeed, depositors with a history of liquidation tend to withdraw more even during noncrisis periods.

Finally, depositors with older accounts run less under both shocks. Under a low-solvency-risk shock, this could be because they are informed that the bank

[30] The informational advantage of loan-linked depositors is not likely due to their status as member-borrowers, who, in a cooperative bank, hold some voting rights. Borrower voting rights pertain to the election of bank directors, and borrowers have no direct role in lending policy or the supervision of bank finances (Rao (1999)).

is solvent or because they are trusting but not informed about solvency. Looking at the high-solvency-risk shock, however, suggests that, unlike loan-linked depositors, the second channel is more plausible for longer lived accounts—old deposits are not informed, but rather are stable regardless of the shock.

IV. Conclusion

In this paper, we use a unique setting—a single bank that experienced shocks of different solvency risk—to examine, using microlevel administrative and survey data, whether depositors' actions depend on the underlying solvency risk posed by a shock.

We find that there is substantial heterogeneity in depositor responses to the true solvency risk facing a bank. Depositors with loan linkages or who are staff of the bank display different behavior across types of shocks. In particular, they are more likely to run when the true solvency risk of the bank is high, and less likely to run when the true solvency risk is low. Uninsured depositors are more likely to run under both shocks, but again are relatively more likely to do so when the true solvency risk is high. We also find that depositors with more transaction activity and younger accounts are more likely to run regardless of the solvency risk of the bank. The results support the idea that some types of depositors are, at least partly, informed about solvency risk. Our results speak to the fragility of banks, suggesting that banks with otherwise identical balance sheets can be differently fragile depending on their relationships with depositors.

The overarching goal of banking regulation is to provide stability without sacrificing market discipline of risky banks. Our results suggest that depositor-bank relationships, at least some kinds, can provide just this kind of *conditional* stability. For example, depositors with loan linkages run based on the true solvency risk of the bank, disciplining a bank when needed but not sparking an unjustified panic. Much debate around the Basel III standards involves what should count as stable deposits for the purposes of liquidity coverage ratios. Stable deposits, from our results, are (i) older, (ii) insured, (iii) infrequently transacted upon, regardless of the nature of the underlying shock. Loan-linked accounts are stable in a low-risk shock, but are not stable when solvency risk is high. This suggests that coverage ratios may be fine-tuned based on depositor characteristics, taking into account variation in solvency risk.

To what extent is our study informative about banking and bank regulation in general? We make several observations about the external validity of our findings, both across banks and across shocks. First, small credit unions and community banks in the United States, Germany, and other countries, as in India, are both vital to lending and vulnerable to shocks (Kroszner (2007), Gilbert, Meyer, and Fuchs (2013)). We expect the dynamics of information acquisition and depositor withdrawals in these banks to be similar to what we observe in our study.

Second, there is some evidence that the behavior of retail depositors, even at large banks, is consistent with our findings. Brown, Guin, and Morkoetter

(2014) survey depositors on withdrawals from large European banks in the recent crisis and find runs of a similar extent (5% of depositors) to those studied here. They also find that depositor relationships matter, in that borrowers are less likely than other depositors to leave a distressed bank (UBS). This finding is consistent with our findings from the panic or low-solvency-risk shock, which is the appropriate comparison since, in these too-big-to-fail banks, depositors are insulated from any losses (UBS was ultimately bailed out). By contrast, the weak implementation of deposit insurance in India enables us to study shocks of a more fundamental nature, for depositors, than regulation in the United States or Europe would allow to play out.

The depositor response to any shock will ultimately depend on the bank involved and the regulatory regime in force. Basel III itself applies only to large banks, and we would not apply our quantitative results to calculate run-off risk for a European bank, for example. Rather, our results establish depositor heterogeneity in the response to solvency risk, and thereby support the idea of using banking relationships in the categorization of stable and less stable deposits. It is not enough to look at the balance sheet of a bank to assess fragility—one also needs to account for the composition of its deposits.

Finally, what can we hope to learn about crises, in the worst of times, from studying two shocks that occurred in, economically speaking, the best of times? Being able to isolate the nature of the shocks studied here is beneficial for understanding what drives depositors to run. In a crisis, especially if there is significant interbank lending, the difference between high- and low-solvency-risk shocks would not be clearcut. Calomiris and Mason (1997) study bank failures during the Chicago bank panic of 1932 and find that, while failed banks saw greater deposit withdrawals before the crisis, solvent banks did not fail, making the case for informed market discipline across banks. The microlevel evidence here supports the idea that some depositors are informed about bank fundamentals, and further shows that depositors' response to a shock depends on their banking relationships and the nature of the shock itself. An important topic for future research is how the depositor heterogeneity established here may attenuate or amplify initial shocks to solvency in a systemic crisis.

Initial submission: July 15, 2014; Accepted: September 16, 2015
Editors: Bruno Biais, Michael R. Roberts, and Kenneth J. Singleton

Appendix

Table A.I
Variable Definitions

Variable	Definition
Run (=1)	Dummy variable equal to one if a depositor withdrew 50% of transaction account balances in the week beginning from the close the day before the run. In hazard models, running is defined as withdrawal of 50% of balances in a single day.
Transaction balance	Total transaction balances in thousand INR, 90 days prior to the run. (Regression specifications use balance in ’00,000s of INR as indicated in the table notes.)
Above insurance cover (=1)	Dummy variable equal to one for balances above the deposit insurance limit, 90 days prior to the run.
Account age	Time an account has been open in years on the day before the shock.
Loan linkage (=1)	A dummy indicating that a depositor or a member of the depositor’s family has a current or past loan from the bank on the date of the run, excluding overdraft accounts against fixed deposits.
Staff (=1)	A dummy indicating that a depositor or a member of the depositor’s family is a staff member.
Liquidation history	The mean of a dummy equal to one if a depositor withdrew 50% of balances on a given day over the year prior to the run, but excluding the 90 days immediately prior.
Transaction history	The mean of a dummy equal to one if a depositor had a transaction on a given day over the year prior to the run, but excluding the 90 days immediately prior.
Daily transactions, year prior to run	Mean number of transactions per day over the year prior to the run, but excluding the 90 days immediately prior.
Daily withdrawal, year prior to run	Mean withdrawal amount per day over the year prior to the run, but excluding the 90 days immediately prior.
Daily deposit, year prior to run	Mean deposit amount per day over the year prior to the run, but excluding the 90 days immediately prior.
Network member has run (=1)	A depositor’s introducer network consists of anyone who introduced that depositor, anyone introduced by the same person as that depositor, and anyone that the depositor himself or herself introduced. The variable network member has run is equal to one during the long-event (hazard model) window if a member of the depositor’s network has run by each given date.

The table gives definitions for variables shown in Table I.

REFERENCES

Basel Committee on Banking Supervision, 2013, Basel III: The liquidity coverage ratio and liquidity risk monitoring tools, Bank for International Settlements. Available at: http://www.bis.org/publ/bcbs238.pdf. Accessed December 19, 2013.

Bennett, Rosalind L., Vivian Hwa, and Myron L. Kwast, 2014, Market discipline by bank creditors during the 2008–2010 crisis, Working paper, FDIC Center for Financial Research.

Berger, Allen N., Sally M. Davies, and Mark J. Flannery, 2000, Comparing market and supervisory assessments of bank performance: Who knows what when? *Journal of Money, Credit and Banking* 32, 641–667.

Billett, Matthew T., Jon A. Garfinkel, and Edward S. O'Neal, 1998, The cost of market versus regulatory discipline in banking, *Journal of Financial Economics* 48, 333–358.

Brown, Martin, Benjamin Guin, and Stefan Morkoetter, 2014, Switching costs, deposit insurance and deposit withdrawals from distressed banks, Working paper, University of St. Gallen.

Bryant, John, 1980, A model of reserves, bank runs, and deposit insurance, *Journal of Banking & Finance* 4, 335–344.

Calomiris, Charles W., and Charles M. Kahn, 1991, The role of demandable debt in structuring optimal banking arrangements, *American Economic Review* 81, 497–513.

Calomiris, Charles W., and Joseph R. Mason, 1997, Contagion and bank failures during the Great Depression: The June 1932 Chicago banking panic, *American Economic Review* 87, 863–883.

Chari, Varadarajin, and Ravi Jagannathan, 1988, Banking panics, information, and rational expectations equilibrium, *The Journal of Finance* 43, 749–761.

Chen, Qi, Itay Goldstein, and Wei Jiang, 2010, Payoff complementarities and financial fragility: Evidence from mutual fund outflows, *Journal of Financial Economics* 97, 239–262.

Chen, Yeh-Ning, 1999, Banking panics: The role of the first-come, first-served rule and information externalities, *Journal of Political Economy* 107, 946–968.

Davenport, Andrew M., and Kathleen M. McDill, 2006, The depositor behind the discipline: A micro-level case study of Hamilton Bank, *Journal of Financial Services Research* 30, 93–109.

DeYoung, Robert, Mark J. Flannery, William W. Lang, and Sorin M. Sorescu, 2001, The information content of bank exam ratings and subordinated debt prices, *Journal of Money, Credit and Banking* 33, 900–925.

Diamond, Douglas W. and Philip H. Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401–419.

Diamond, Douglas W., and R. G. Rajan, 2001, Banks and liquidity, *American Economic Review* 91, 422–425.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux, 1996, Labor market institutions and the distribution of wages, 1973–1992: A semiparametric approach, *Econometrica* 64, 1001–1044.

Flannery, Mark J., and Joel F. Houston, 1999, The value of a government monitor for U.S. banking firms, *Journal of Money, Credit and Banking* 31, 14–34.

Flannery, Mark J., and Sorin M. Sorescu, 1996, Evidence of bank market discipline in subordinated debenture yields: 1983–1991, *Journal of Finance* 51, 1347–1377.

Gilbert, R. Alton, Andrew P. Meyer, and James W. Fuchs, 2013, The future of community banks: Lessons from banks that thrived during the recent financial crisis, *Federal Reserve Bank of St. Louis Review* 95, 115–143.

Goldberg, Lawrence G., and Sylvia C. Hudgins, 2002, Depositor discipline and changing strategies for regulating thrift institutions, *Journal of Financial Economics* 63, 263–274.

Goldstein, Itay, and Ady Pauzner, 2005, Demand-deposit contracts and the probability of bank runs, *Journal of Finance* 60, 1293–1327.

Gorton, Gary, 1988, Banking panics and business cycles, *Oxford Economic Papers* 40, 751–781.

Gorton, Gary, and Andrew Metrick, 2012, Securitized banking and the run on repo, *Journal of Financial Economics* 104, 425–451.

Hanson, Samuel G., Andrei Shleifer, Jeremy C. Stein, and Robert W. Vishny, 2014, Banks as patient fixed income investors, NBER Working Paper No. 20288.

He, Zhinguo, and Asaf Manela, 2012, Information acquisition in rumor-based bank runs, NBER Working Paper No. 18513.

Iyer, Rajkamal, and Jose-Luis Peydro, 2011, Interbank contagion at work: Evidence from a natural experiment, *Review of Financial Studies* 24, 1337–1377.

Iyer, Rajkamal, and Manju Puri, 2012, Understanding bank runs: The importance of depositor-bank relationships and networks, *American Economic Review* 102, 1414–1445.

Jacklin, Charles J., and Sudipto Bhattacharya, 1988, Distinguishing panics and information-based bank runs: Welfare and policy implications, *The Journal of Political Economy* 96, 568–592.

Kashyap, Anil K., Raghuram Rajan, and Jeremy C. Stein, 2002, Banks as liquidity providers: An explanation for Kelly, Morgan, and Cormac Ó Gráda, 2000, Market contagion: Evidence from the panics of 1854 and 1857, *American Economic Review* 90, 1110–1124.

Kelly, Morgan, and Cormac Ó Gráda, 2000, Market contagion: Evidence from the panics of 1854 and 1857, *American Economic Review* 90, 1110–1124.

Kroszner, Randall S., 2007, Community banks: The continuing importance of relationship finance, Speech at America's Community Bankers Government Affairs Conference.

Martinez-Peria, Maria Soledad, and Sergio L. Schmukler, 2001, Do depositors punish banks for bad behavior? Market discipline, deposit insurance, and banking crises, *Journal of Finance* 56, 1029–1051.

Park, Sangkyun, and Stavros Peristiani, 1998, Market discipline by thrift depositors, *Journal of Money, Credit and Banking* 30, 347–364.

Postlewaite, Andrew, and Xavier Vives, 1987, Bank runs as an equilibrium phenomenon, *The Journal of Political Economy* 95, 485–491.

Rao, Madhav, 1999, Report of the High Power Committee on Urban Cooperative Banks. Reserve Bank of India. Available at: https://www.rbi.org.in/Scripts/PublicationReportDetails.aspx?FromDate=12/07/99&SECID=7&SUBSECID=0. Accessed September 8, 2015.

Rochet, J. C., and Xavier Vives, 2004, Coordination failures and the lender of last resort: Was Bagehot right after all? *Journal of the European Economic Association* 2, 1116–1147.

Saunders, Anthony, and Berry Wilson, 1996, Contagious bank runs: Evidence from the 1929–1933 period, *Journal of Financial Intermediation* 5, 409–423.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.