EVALUATING THEORIES OF BANK RUNS WITH HETEROGENEITY RESTRICTIONS

Ferre De Graeve
Sveriges Riksbank

Alexei Karas
University College Roosevelt
and Utrecht University School
of Economics

Abstract
This paper empirically tests theories of bank runs. We use a structural panel VAR to extract runs from deposit market data. Identification exploits cross-sectional heterogeneity in deposit insurance: we identify bank runs as adverse deposit market supply shocks hitting uninsured banks harder compared to insured. Conditional on a run, we study the behavior of uninsured banks with bad and good fundamentals. We find that both experience runs, but deposit outflows at the former are more severe. Panic effects, which affect all uninsured deposits alike, irrespective of fundamentals, dominate in the aggregate. Insured banks partially absorb the outflow of uninsured deposits. (JEL: C3, E5, G01, G21)

1. Introduction

The recent crisis is a forceful reminder that bank runs are a constant threat to financial systems. While runs can take place in different markets, the prevalence of bank runs in costly banking crises makes understanding their determinants of critical importance. This is all the more true since the two main theories on the cause of runs imply substantially different policy responses. While *panic* runs (as in e.g. Diamond and Dybvig 1983) can be forestalled by deposit insurance, suspension of convertibility or

The editor in charge of this paper was Fabio Canova.

Acknowledgments: We would like to thank three anonymous referees, Fabio Canova (the editor), Charles Calomiris, Elena Carletti, Russell Cooper, Dean Corbae, Olivier De Jonghe, Gianni De Nicoló, Giovanni Dell'Ariccia, Huberto Ennis, Jonas Fisher, Simon Gilchrist, Alejandro Justiniano, Robert Kollmann, Roland Meeks, Mattias Villani, Raf Wouters, and seminar participants at the Federal Reserve Board, Sveriges Riksbank, De Nederlandsche Bank, Swiss National Bank, Cardiff Business School, ECARES, Maastricht University, Queen Mary, Tilburg University, the Federal Reserve System Committee Meeting on "Financial Structure and Regulation" and the Wharton conference on "Liquidity and Financial Crises" for useful comments and valuable suggestions. The views expressed herein are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank.

E-mail: ferre.de.graeve@riksbank.se (De Graeve); a.karas@ucr.nl (Karas)

Journal of the European Economic Association August 2014 12(4):969–996
© 2014 by the European Economic Association
DOI: 10.1111/jeea.12080

liquidity provision, *fundamental* runs (as in e.g. Allen and Gale 1998) involve policy measures such as balance sheet constraints, or even laissez-faire.

Traditional tests of bank runs (and financial crises more generally) are, by and large, variations on the following approach. First, define periods which constitute a crisis. Subsequently, verify if aggregate fundamental factors correlate with these dates, or whether fundamental variables explain cross-sectional differences across banks during crisis times. The reduced-form nature of these tests along with some of the assumptions involved in crisis-period identification pose significant challenges to definitive structural inference on the importance of different theories of bank runs.

In this paper, we propose a twist to standard macroeconometric methods (structural vector autoregressions, SVAR) by incorporating cross-sectional information into identification. This enables the use of novel, heterogeneity restrictions to single out bank runs from other concurrent events. Our approach has important advantages relative to earlier evidence on banks runs. Firstly, it eliminates the need to subjectively label particular periods as a crisis, as well as the need to assume that everything occurring in those periods is exogenous and due to the bank run. Secondly, our approach goes beyond establishing reduced-form correlations by providing structural impulse responses. The results we provide are conditional on a bank run, while alternative, concurrent sources of variation in the data are explicitly controlled for. These features address inconvenient properties of traditional tests of bank runs which have rendered structural interpretations difficult.

We identify bank runs as deposit market supply shocks in which insurance matters. They are characterized by an outflow of uninsured deposits relative to insured deposits, and are not driven by a (relative) decrease in uninsured deposit interest rates. Conditional on bank runs, we quantify the contribution of both the fundamental and the panic view by investigating cross-sectional differences across uninsured banks with good and bad fundamentals. We also evaluate the effect of being insured in the face of a run. We apply our approach to Russian deposit market data for the period 2002–2007. Russian micro bank data are not only of high quality and detail, they are also very informative about the topic at hand: our sample includes at least one severe market disruption, dozens of bank failures, cross-sectional heterogeneity in deposit insurance, and more.

Our results bear on the policy debate. In particular, we show that there is merit in both the fundamental and the panic view. On the one hand, fundamentally flawed banks face substantially larger deposit outflows during a bank run, relative to banks with strong fundamentals. This corroborates the fundamental view. On the other hand, even banks with solid fundamentals face significant outflows. This finding provides support for the panic view, especially since such outflows are not observed at banks that have deposit insurance. Importantly, particularly from a policy perspective, we quantify the relevance of both theories in the aggregate. In our sample, panic effects substantially outweigh fundamental effects.

With very few exceptions, empirical studies have attributed bank runs to the fundamental view and downweigh the role of panics (see e.g. Gorton 1988; Saunders and Wilson 1996; Schumacher 2000; Calomiris and Mason 2003b). However, due to its reduced-form nature, finding fundamentals to be important is subject to different possible interpretations. Our results, which are structural, attribute a much larger role to the panic view of bank runs. This has important policy implications. In particular, fundamentals-based regulation may prove insufficient to curb transmission of banking crises through deposit markets. Rather, policies geared towards effectively shielding depository institutions from panic effects may be required to do so effectively.

The paper is organized as follows. We start with a brief overview of theories of bank runs, how they have been validated, and how our approach relates to them. In Section 3, we lay out our empirical approach and motivate our identification restrictions. After describing events in our sample, Section 4 discusses the occurrence and impact of bank runs, validates the empirical approach and provides tests that discriminate between the different views. Following an overview of robustness checks in Section 5, we compare our approach to alternative ones and conclude in Section 6.

2. Tests and Theories of Bank Runs

Theories of bank runs have traditionally been cast in one of two strands: the panic and the fundamental view. The panic view (e.g. Diamond and Dybvig 1983; Peck and Shell 2003; Ennis and Keister 2009) sees bank runs as a result of coordination problems among agents. Depositors run on their bank solely because they expect other depositors to withdraw. As a result, bank runs can arise as sunspot equilibria. Since all that matters in this view is what other depositors do, runs can occur irrespective of whether a bank has good or bad fundamentals. Effective policies are therefore those that reduce the incentive for depositors to withdraw, even if others withdraw. Deposit insurance, for instance, readily accomplishes that. By contrast, the fundamental view (e.g. Allen and Gale 1998; Chari and Jagannathan 1988; Jacklin and Bhattacharya 1988) posits that depositors run on banks because of information on fundamentals that makes them question particular banks' solvency. Therefore, in this view, bank runs correlate with fundamentals. Effective policies ensure adequate balance sheet health and thus reduce depositors' rationale for withdrawals.

The early strand of empirical tests of bank runs aimed to discriminate between these two theories using one of two tests on aggregate data. A first type of tests is exemplified by the narrative of Friedman and Schwartz (1963) which verifies if any fundamental variable correlates with the incidence of crises. Many variations of this test exist—typically less descriptive, more rigorously econometric—and include Calomiris and Gorton (1991), Canova (1994), Demirgüç-Kunt and Detragiache (1998).

A second type of tests aims to verify whether the data generating process of crisis episodes is the same as that during normal times. If they are the same, then one typically concludes that the usual *fundamental* relationships are at work during crises. Gorton (1988) is perhaps the earliest example of this test, while a similar approach is embedded in, for example, Donaldson (1992) and Nason and Tallman (2012).

Much of the more recent empirical literature has looked beyond aggregate data and shown that during crises there are cross-sectional differences in deposit flows or failure probabilities between banks that correlate with fundamentals. Key examples are Saunders and Wilson (1996), Schumacher (2000), Martinez Peria and Schmukler (2001), Calomiris and Mason (2003b) and Iyer and Peydró (2011).

The present paper adds to this evidence by proposing a procedure that naturally enables crisis identification based on a priori restrictions and allows structural inference conditional on such crises. We thereby avoid both the need to (subjectively and exhaustively) identify certain dates as crises and the reduced form nature of much of the earlier tests.

In our approach to identification and conditional testing, we essentially apply the insights of the above theories to the cross-section of banks. First, to distinguish bank runs from other events, we use the fact that—all else equal—having deposit insurance reduces depositors' incentive to run. For instance, in Diamond and Dybvig (1983), the existence of deposit insurance eliminates coordination failures. In our analysis we use this insight by requiring that, during a bank run, banks with deposit insurance cannot be more severely affected than banks without. In addition, we study that part of deposit flows which is supply-driven. We do so by conditioning on interest rate movements. Our identifying restrictions make sure correlations are not driven by demand concerns (which have different interest rate implications) or macro considerations (which affect all banks in a similar way). Combined, our identifying restrictions focus attention on shocks which induce a drain of funds at uninsured banks, relative to insured ones, which is not driven by a (absolute or relative) reduction in their interest rate. In times of crisis, deposit volumes change for a variety of factors. Our approach allows one to refrain from studying unconditional deposit flows, and instead centers attention on the bank run itself.

Figure 1 provides a simplified graphical description of our approach to identifying bank runs, as well as the tests we perform conditional on a bank run. Our conditional empirical analysis can be seen as a series of questions which jointly evaluate the theories: How severely are banks with the poorest fundamentals (henceforth "bad banks") affected? Are banks with good fundamentals ("good banks") ran on too? If good banks face runs as much as bad banks, this is evidence in favor of the panic view:

			$\xrightarrow{\hspace{6em}}$	Deposits		
Identification:	Uninsured banks	$<$	Insured banks			
Tests:	bad = good	$<$	0		$\Longrightarrow$	panic, no fund
	bad $<$ good	$=$	min(0,insured)		$\Longrightarrow$	fund, no panic
	bad $<$ good	$<$	min(0,insured)		$\Longrightarrow$	fund & panic
			$\lesseqgtr 0$		$\Longrightarrow$	DI effective?

FIGURE 1. Identification and conditional tests.

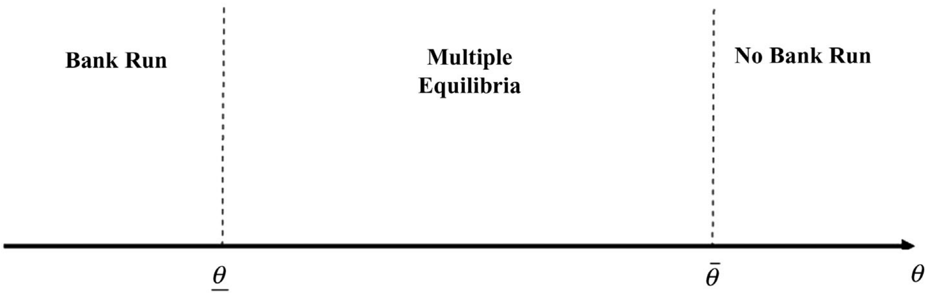

FIGURE 2. Bank runs and fundamentals. Source: Goldstein and Razin (2012).

depositors do not discriminate based on fundamentals. If, contrary to bad banks, good banks do not face runs, then the fundamental theory holds with no scope for panic. If good banks face runs, but less severely so than bad banks, then both theories have merit. Moreover, our tests also evaluate policy effectiveness. For instance, do insured banks face runs, too? Do they merely withstand runs, or do they actually absorb the outflow of uninsured deposits during the run?

Our tests also relate to more recent global-games models of financial crises, as initiated by Morris and Shin (1998), and applied to bank runs in Goldstein and Pauzner (2005) and many others since. By and large, contemporary theory nests both the fundamental and the panic view as summarized in Figure 2. The axis represents a bank's fundamental soundness. On the one extreme, if a bank's fundamentals are excellent ($\bar{\theta} < \theta$), this indicates zero probability of failure, and thus eliminates any incentive for depositors to run. On the other extreme, if a bank's fundamentals are extremely poor ($\theta < \underline{\theta}$), depositors are sure the bank will fail and therefore withdraw. Between these extremes ($\underline{\theta} < \theta < \bar{\theta}$), failure is uncertain, which induces strategic complementarities between depositors: whether withdrawing is optimal depends on whether others withdraw. This region essentially reflects the Diamond and Dybvig (1983) environment, in which the failure to coordinate among depositors not to run gives rise to multiple equilibria. There is a clear potential for correspondence between our tests and this recent theoretical perspective on crises. One way of reading our tests is as evaluating whether, in the cross-section of banks, the different regions of Figure 2 (co-)exist. However, and this is important to keep in mind, the mapping between contemporary theory and our testing framework is not exact.[1]

1. In part, that is because the theory does not address aspects of the data that certainly are relevant in our setting. For instance, in Figure 2, there is no intensive margin during a run. Given that a run occurs, it is equally intense irrespective of whether it is fundamental or due to a coordination failure. Also, the theory really describes agents' decision problem involving fundamentals of a single entity, while our analysis studies an entire cross-section of entities.

3. Empirical Strategy

3.1. Reduced Form

Our empirical approach starts from a reduced-form model of (log) deposit quantities (D) and interest rates (R). Let y_{it}^g be the $(J \times 1)$ vector of endogenous variables in period t for bank i belonging to group g; here $y_{it}^g = (D_{it}^g, R_{it}^g)'$. The model we estimate is a panel VAR of the form

$$y_{it}^g = b_g X_{it}^g + e_{it}^g, \tag{1}$$

where $i = 1, \ldots, N_g$; $g = 1, \ldots, G$; $t = 1, \ldots, T$; $X_{it}^g = (y_{it-1}^{g\prime}, \ldots, y_{it-p}^{g\prime}, 1)'$, p is the number of lags and X_{it}^g is a vector of length $k = Jp + 1$. Units are indexed by i; groups are indexed by g. A group g consists of N_g individual units. The total cross-sectional dimension of the data is $N = \sum_{g=1}^{G} N_g$. Lag length is the same across groups and determined using the Akaike Information Criterion. The $(J \times k)$ coefficient matrix b_g is homogeneous across units within a group g, but is allowed to differ between groups: $b_g \neq b_h$, $g \neq h$. Each bank-unit i is subject to reduced-form shocks e_{it}^g. These shocks can be unit-specific, u_{it}^g, or group-specific, v_t^g:

$$y_{it}^g = b_g X_{it}^g + v_t^g + u_{it}^g. \tag{2}$$

The reduced-form nature of the shocks implies they may be correlated (a) across equations within a unit, captured by a nondiagonal covariance matrix of the unit-specific shocks $u_{it}^g \sim N(0, \Sigma_u)$, where Σ_u is of dimension $J \times J$, and (b) across equations across groups, with $v_t = (v_t^{1\prime}, \ldots, v_t^{G\prime})' \sim N(0, \Sigma_v)$ and Σ_v a nondiagonal matrix of dimension $GJ \times GJ$. Finally, unit-specific shocks are orthogonal to group-specific ones, $u_{it}^g \perp v_t^g, \forall i, g$.

Using the terminology of Canova and Ciccarelli (2013), equation (2) describes a panel VAR with static but no dynamic interdependencies and with group-level cross-sectional heterogeneity. Specifically, while individuals are correlated through group-specific shocks v_t^g (i.e. static interdependence), lagged endogenous variables of unit i do not appear in equations of unit j (i.e. no dynamic interdependence). Coefficients are the same for all units within the same group but are heterogeneous across groups.

The motivation for using this particular reduced form is as follows. Panel VARs are becoming a widely used tool in empirical macro. The reason is largely twofold and applies equally to our analysis of micro data. Not only is the empirical fit of reduced-form panel VARs substantial, they are also flexible enough to maintain consistency with various underlying structural models.

Given the choice of a panel VAR, the minimum flexibility we wish to allow in our reduced form is that different groups of banks, such as insured and uninsured banks, can behave differently and be related. Clearly these are potentially essential features during bank runs. Dynamic interdependencies are less important for our study. On the one hand, our data are at a quarterly frequency, whereas one may expect between group shifts to occur mostly at higher frequencies. On the other hand, in our data block exogeneity tests typically do not reject the null of no dynamic group interdependencies. We allow for heterogeneity across different groups of banks, as group stratifications are based on characteristics theory deems essential in determining depositor behavior. While allowing for further within-group heterogeneity is possible (e.g. Canova 2004), it is beyond the scope of the present analysis. In sum, while several aspects of equation (2) can be generalized (see Canova and Ciccarelli 2009), heterogeneity and static interdependence are key here.

3.2. Estimation and Inference

In keeping with the majority of recent applications of structural panel VARs, our approach is Bayesian. The setting is hierarchical since it has individual as well as group-specific features. As in most hierarchical settings, there are multiple equivalent ways to set up the model (see e.g. Canova 2007; Gelman et al. 2004). Concerning reduced-form estimation, perhaps the simplest approach is to add group-specific time-dummies to X^g_{it}, one per period per group. Then the group-specific residual component v^g_t is absorbed by the matrix of explanatory variables. As a result, when stacked, equation (2) fits into the standard linear regression framework (with block-diagonal residual covariance matrix $I_N \otimes \Sigma_u$). Then, given a noninformative conjugate prior, OLS retrieves the posterior mean. The coefficient estimates in b_g corresponding to the group-specific time-dummies are then estimates of v^g_t, from which one can construct an estimate for Σ_v.

Let us now turn to structural inference. Stacking the model (1)–(2) across all units within a group and consequently across all groups we get

$$y_t = X'_t\beta + e_t, \qquad e_t \sim N(0, \Sigma) \tag{3}$$

$$= X'_t\beta + Mv_t + u_t, \quad Mv_t + u_t \sim N(0, M\Sigma_v M' + I_N \otimes \Sigma_u), \tag{4}$$

where

$$\underset{NJ\times 1}{y_t} = \left(y^{1\prime}_t, \ldots, y^{G\prime}_t\right)', \quad \underset{N_gJ\times 1}{y^g_t} = \left(y^{g\prime}_{1t}, \ldots, y^{g\prime}_{N_gt}\right)', \quad \underset{GJk\times 1}{\beta} = (\beta'_1, \ldots, \beta'_G)',$$

$$\underset{Jk\times 1}{\beta_g} = vec(b'_g), \quad \underset{GJk\times NJ}{X_t} = \operatorname{diag}\left(X^1_t, \ldots, X^G_t\right),$$

$$\underset{Jk\times N_gJ}{X^g_t} = \left(I_J \otimes X^g_{1t}, \ldots, I_J \otimes X^g_{N_gt}\right),$$

and M, which has dimension $NJ \times GJ$, is a matrix that places group-specific shocks in the relevant equation of all the units in that group.[2]

2. To be precise, $M = m \otimes I_J$ where m is a $(N \times G)$ matrix with entries

$$m(i, g) = \begin{cases} 1 & \text{if unit } i \in \text{ group } g, \\ 0 & \text{if unit } i \notin \text{ group } g. \end{cases}$$

Observe that (4) is nested within the generalized linear model (3). Inference is therefore standard. Particularly, given a noninformative conjugate prior, the posterior distribution $p(\beta, \Sigma | y)$ is of the Normal-Inverse Wishart form. For each draw (B, S) from that reduced form posterior, we take a random Givens rotation (with angles drawn from the uniform distribution) to decompose S and compute impulse responses. If the responses satisfy the identifying assumptions (discussed in what follows, see equation (5)) the draw is accepted as coming from the joint posterior of reduced-form parameters and rotation matrices. All results take into account both sampling and identification uncertainty and are based on 1000 accepted draws. Our priors are invariably conjugate and noninformative, as in Uhlig (2005, p. 410). Finally, whenever alternative group stratifications are used, the model is re-estimated in its entirety. We provide additional detail on estimation and simulation in Appendix A.

3.3. Structural Identification

3.3.1. Approach. Structural panel VAR studies typically identify shocks by imposing restrictions common to all cross-sectional units. Examples are Love and Zicchino (2006), Canova and Pappa (2006), Jarociński (2010) and Calza, Monacelli, and Stracca (2013). Such identifying assumptions restrict the impulse response functions of certain endogenous variables and take the form

$$\frac{\partial y_{i,t+s}}{\partial \varepsilon_t} \lesseqgtr 0, \quad \forall i = 1, \ldots, N,$$

where ε denotes a particular structural shock, s is the horizon at which the restrictions hold, $s \subseteq \{0, 1, \ldots, \infty\}$, and the notation implies traditional zero restrictions if equality holds and sign restrictions otherwise. Different from this traditional approach, we exploit the existence of heterogeneity in the cross-section for the purpose of identification. Particularly, we relax the assumption that identifying restrictions have to hold for all units of the cross-section. This allows imposing restrictions on

$$\frac{\partial y_{\Lambda,t+s}}{\partial \varepsilon_t},$$

where Λ denotes a subset of the cross-sectional dimension, $\Lambda \subset \{1, \ldots, N\}$. Dispensing with the requirement that identifying restrictions are common to all cross-sectional units naturally gives rise to *heterogeneity restrictions*, which contrast the behavior of different subsets of the cross-section:

$$\frac{\partial y_{\Lambda_1,t+s}}{\partial \varepsilon_t} \lesseqgtr \frac{\partial y_{\Lambda_2,t+s}}{\partial \varepsilon_t}.$$

These can be implemented with sign or exclusion restrictions, depending on the preference of the researcher and the question at hand. Note that the subsets can, but need not, be exhaustive. In analogy to traditional SVARs, it may be natural to constrain the behavior of some subsets of the cross-section, while leaving others free.

Heterogeneity restrictions are particularly rich when applied to panel VARs on micro data. The reason is obvious: heterogeneity prevails in micro data. The availability of numerous cross-sections allows for many possible stratifications. Importantly, economic variables can underlie the stratifications. This implies that theory can be linked to empirical identification not just in the time series dimension (as in the typical macro SVAR), but also through the cross-section. It is in this cross-sectional use of theory that our restrictions are distinct from other applications of sign restrictions in models with both time and cross-sectional dimensions, such as Amir and Uhlig (2009), Eickmeier and Ng (2011) or Helbling et al. (2011).

3.3.2. Identifying Bank Runs. We now propose identifying restrictions that filter out a bank run from concurrent developments in the deposit market. To identify bank runs we contrast the behavior of two types of banks: those with deposit insurance and those without. Thus, in terms of our empirical model (2) we now discriminate between two groups: U = Uninsured, I = Insured.

We define a bank run as *a supply shock in which insurance matters*:

$$\begin{cases} \Delta D_t^U & < 0 \\ \Delta R_t^U & \geqslant 0 \\ \Delta D_t^U & < \Delta D_t^I \\ \Delta R_t^U & \geqslant \Delta R_t^I, \end{cases} \tag{5}$$

where Δ is shorthand for an impulse response, $\partial(.)/\partial \text{run}_t$.[3] The absence of a unit-specific cross-sectional index conveys these correspond to the group-specific responses of deposits and interest rates—that is, the average response over the cross-section of uninsured (resp. insured) banks. The time index t indicates the restrictions are imposed only contemporaneously, at the time the shock occurs.

The first restriction says, quite uncontroversially, that a bank run lowers the quantity of deposits at uninsured banks. In addition, the second restriction confines attention to supply-driven deposit outflows. After all, our interest is in analyzing bank runs (which are a particular type of supply shock), rather than, for instance, a demand-driven deposit outflow. The latter could follow from uninsured banks lowering the interest rate they pay on their deposits, for example in response to lower loan demand. To exclude such cases, we add a restriction on the interest rate. These first two restrictions combined effectively rule out demand shocks in the uninsured deposit market.

We additionally impose heterogeneity restrictions, contrasting the behavior of different types of banks. In particular, the third restriction requires that a bank run is

3. Thus, Δ measures the change relative to baseline, where the latter is measured by the dynamics of the system (2) in the absence of the structural shock. This implies, for instance, that if uninsured banks pay substantially higher interest rates relative to insured banks on average (and they do), this is picked up by the baseline. The focus on impulse responses to (structural) shocks is important in that it allows ruling out endogenous responses to other, concurrent, events. Examples include responses to earlier as well as alternative structural shocks. Incorporating these would confound the estimate of the pure bank run effect.

not characterized by a worse deposit outflow at insured banks compared to uninsured banks. In other words, we focus on those supply shocks where insurance matters. This restriction ensures that the reason for the outflow is depositor-fear of losing their funds. Observe that the restriction does not require the deposit insurance scheme to be fully credible. Deposits at insured banks can decrease, remain stable, or increase; the restrictions are agnostic in this respect. It only requires that having deposit insurance does not aggravate the deposit outflow relative to banks that are not covered by the deposit insurance scheme.

The fourth restriction rules out relative demand shocks between insured and uninsured banks. These are filtered out by additionally requiring that the relative outflow at uninsured banks is not driven by a relative increase in the interest rate of insured banks. Finally, the runs we consider should be thought of as systemic, since we impose the restrictions on the group-wise behavior.

While these restrictions have a lot of intuitive appeal, one can also think of them as having a direct analogue in theoretical models of bank runs. Consider, for instance, the model of Diamond and Dybvig (1983). This stylized model contains two supply shocks: bank runs and depositor preference shocks for early liquidation. Our first two identifying restrictions jointly isolate supply shocks in the data. In Diamond and Dybvig (1983) preference shocks are not of concern; the bank is able to cope with normal deposit withdrawals. The cross-sectional restriction we impose adds a concern for solvency. This is exactly what represents a bank run in Diamond and Dybvig (1983): depositors withdraw not because they wish to consume, but out of fear of losing their deposits. The joint set of restrictions in equation (5) establish this by requiring that the supply-driven outflow does not occur (as strongly) at banks where depositors' funds are (more) safe. While the Diamond and Dybvig model as such does not deal with different types of banks simultaneously, it does deal with equilibria in the presence and absence of deposit insurance, which is true for most of the literature on bank runs. We view our heterogeneity restrictions as the logical extension of different equilibria in these kinds of models to a cross-sectional setup.

As a final remark, note that equation (5) does not explicitly rule out heterogeneous elasticities to common shocks or potential demand changes at insured banks. In Appendix B we discuss these possibilities, apply alternative identification restrictions and show these are not of particular concern for our conclusions.

3.4. Data

We apply our approach to the Russian deposit market over the period 2002–2007. The Russian deposit market provides a very useful case. The reason is twofold. First, there is cross-sectional variation across banks in the degree to which their household deposits are guaranteed by the government. The insured nature of deposits at state-owned banks in Russia has varied from implicit to explicit but was always there. Before 2004, state-owned banks exclusively enjoyed the explicit state guarantee backing their retail deposits (Civil Code art. 840.1). This guarantee was removed at the end of 2003 (Federal Law No. 182-FZ). Despite that, implicit insurance for state banks continued to exist. Throughout our sample, state-owned banks have enjoyed privileged access to state funds, de facto exemption from some regulatory norms and, on occasion, financial support from the state. Their cost of capital is reduced by the perception that the state will stand behind them (Tompson 2004). Private banks, by contrast, did not have the state backing them (or their deposits). Our method exploits such heterogeneity in insurance between state and private banks to identify bank runs. Particularly, irrespective of whether insurance is implicit or explicit, or fully credible or not, the relative restrictions in equation (5) ought to hold. Moreover, this heterogeneity will allow us to assess the value of having insurance in the face of a bank run.

A second reason why the Russian deposit market is of particular interest is that it has witnessed substantial turbulence in our sample period. In May 2004 the Central Bank of Russia (CBR) closed a bank accused of money laundering while the Federal Service for Financial Monitoring (Federalnaya Sluzhba po Finansovomu Monitoringu) announced that it suspected about a dozen other banks of being involved in money laundering and sponsoring terrorism, without naming the "dirty dozen" (Tompson 2004; Zykova 2004). Several inconsistent black lists began circulating as people tried to guess which banks were suspected by the FSFM. Mutual suspicion led to a drying up of liquidity on the interbank market, putting pressure on the hundreds of smaller banks that are highly dependent on it. The crisis provoked runs on lots of banks, among which major players such as Guta Bank and Alfa Bank. Thus, there is narrative evidence suggestive of (at least one) bank runs occurring in our sample period. We will confront the timing of runs identified by the method to evidence extraneous to the model.

Quarterly financial indicators for all Russian banks are obtained from two private financial information agencies, Interfax (www.interfax.ru) and Mobile (www.mobile.ru). Both data sets cover the same major components of bank balance sheets and income statements but differ in details: some indicators are provided by one agency but not the other. Karas and Schoors (2005) provide a detailed description of the two data sets and establish consistency between the different data sources. The two main variables used in our analysis are household deposits and the implicit interest rate paid on those deposits; the latter is constructed by dividing the interest paid on household deposits by their level. In addition, we use various measures of risk, performance and balance sheet structure to stratify banks into good and bad. The analysis comprises a total of 1,225 banks. Appendix C contains additional detail on the data, as well as summary statistics on variables used to estimate the reduced form and to stratify banks into different groups.

4. Results

4.1. A First Look at the Effects of a Bank Run

Figure 3 plots the effect of a bank run, identified with the restrictions in equation (5), on deposits and interest rates across insured and uninsured banks. The responses are

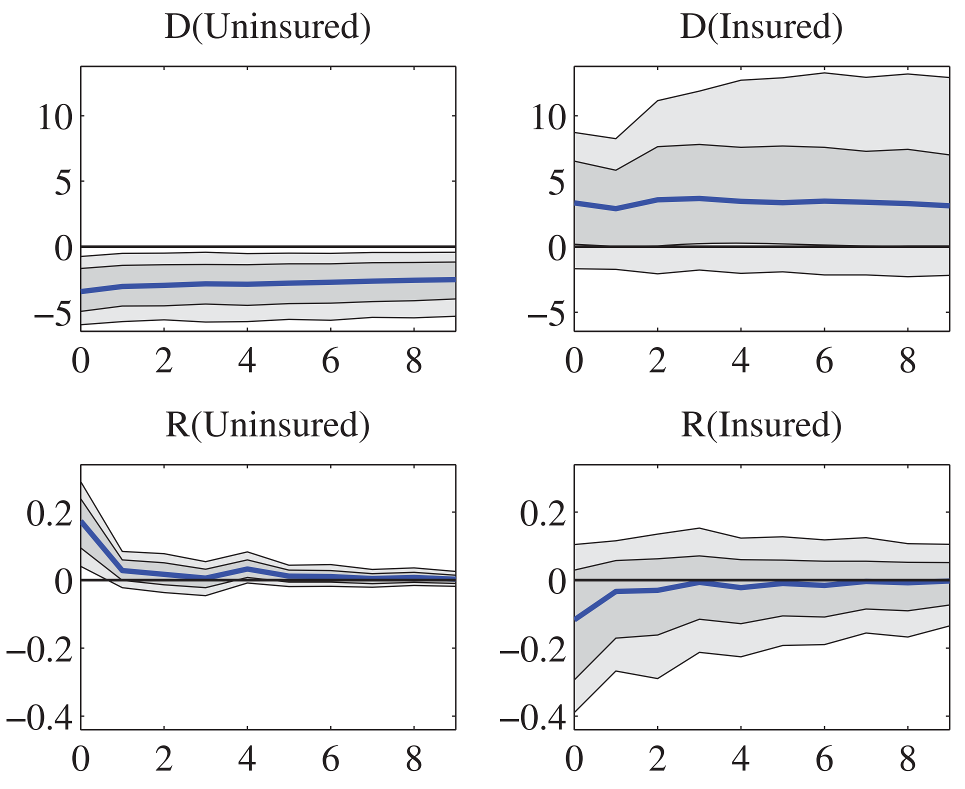

FIGURE 3. Impulse responses to a bank run (1 std. impulse). Thick line: median; dark gray area: 68% credible interval; light gray: 90%.

to a one standard deviation impulse and measure the effect on the average insured and uninsured bank. These reveal a first set of *qualitative* results. Let us start by restating the identifying restrictions imposed on these graphs: that a bank run is a structural shock that implies an outflow of deposits at uninsured banks, that deposit flows are less severe (or even positive) for insured banks, and that these flows are not associated with a (absolute or relative) decrease in the interest rate offered by uninsured banks. These are the only restrictions imposed. All other features in the figure are left unrestricted, and therefore the object of study. We now turn to these.

First, the restrictions are imposed only on impact, at $t = 0$. The apparent persistence of the reduction in uninsured deposits is substantial. The effects of a bank run on the volume of uninsured deposits hardly dissipate.

Second, the figure reveals that insured banks do not face an outflow of deposits, rather to the contrary. While the average uninsured bank experiences a reduction in deposits, the average insured bank sees its deposit base increased. This happens without the insured banks increasing interest rates, or the uninsured banks lowering theirs. Note that while the inflow at insured banks is close to permanent, it is not the most significant. Crucially, however, insured banks are not subject to the run.

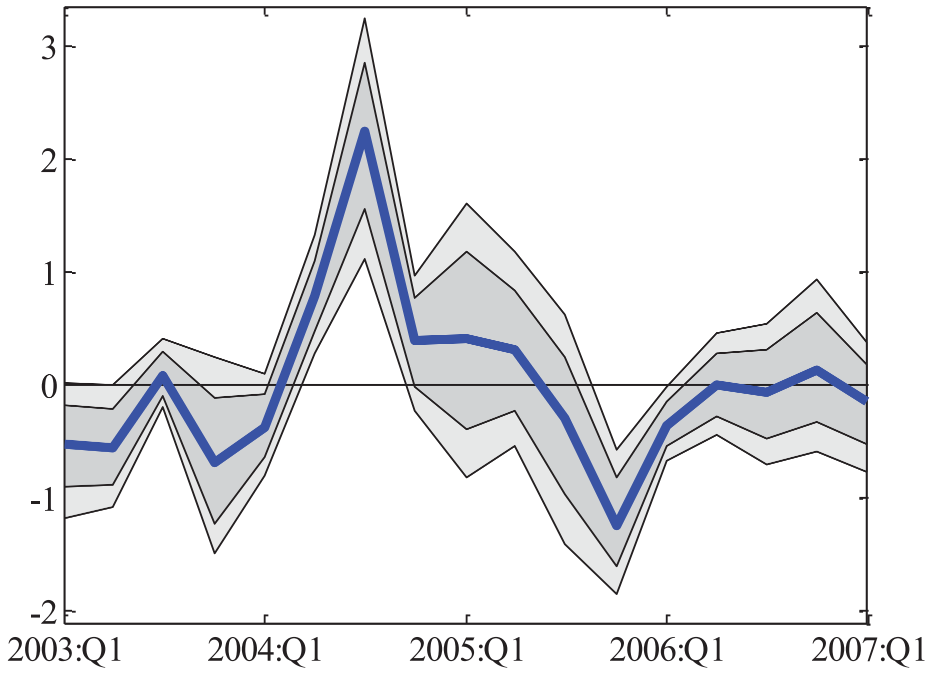

FIGURE 4. A time series of the identified shock.

Third, Figure 4 plots the identified shock over our sample period. The single largest shock is observed during the summer of 2004 (both Q2 and Q3 are significant). The positive sign of the shock implies it pertains to an outflow of deposits at the uninsured banks (and the corresponding signs for the other restrictions). Thus, our approach identifies the 2004 summer (and essentially no other period) as a bank run. One can cross-validate that finding with information outside the model. As a measure of external validation we use press coverage. We perform a computerized search in the article databases of *The Economist* and *The New York Times* for our sample period using the terms “Russia”, “deposit”, and “run”. Out of all hits, three pertain directly to the present paper’s subject. All three are dated summer of 2004 and each of them suggests the possibility of a bank run.[4] We interpret this to be evidence for the fact that, first, a run was very likely in the 2004 summer, and second, there were no other episodes in our sample period suggestive of bank runs. Both are consistent with the results in Figure 4.

4. The articles referred to are “Don’t run for it” (*The Economist*, 6/26/2004), “There’s always Sberbank” (*The Economist*, 7/10/2004), and “Depositors’ jitters increasing as some Russian banks close” (*The New York Times*, 7/9/2004). While this validation is meant as indicative rather than literal, it is interesting that the articles appeared both in 2004:Q2 and 2004:Q3, the same two periods the shock is significantly positive. We perform a similar search in the news database of the Russian news agency Lenta.ru using the terms “deposit”, “bank”, and “crisis”. Out of 16 hits, nine directly pertain to the present paper’s subject. Those nine are dated July through September of 2004 and each suggested the possibility of a bank run.

4.2. Evaluating the Theories

We here provide substantially more detail on the previous results. In particular, we (i) quantify the effects of the 2004 run, (ii) perform hypothesis tests in the cross-section that assess the significance of the panic and the fundamental view on bank runs, and (iii) quantify the contribution of the two theories to the total effect of the run.

Let us first dwell briefly on the additional cross-sectional heterogeneity that is dealt with here. We decompose the uninsured group of banks further into banks with sound fundamentals and banks with flawed fundamentals. This additional cross-sectional stratification allows us to disentangle the panic and fundamental views on bank runs. From the perspective of theory, what matters for depositors is their ex-ante evaluation of banks' solvency.

As any assumption on depositors' information sets is likely an incomplete characterization of their actual information, we approximate depositor information sets in different ways, analogous to stratifications employed in earlier tests of bank runs. We start by using real time bank balance sheet information to assess solvency. We verify whether depositors distinguish banks on the basis of their degree of capitalization. Another frequently analyzed characteristic of bank balance sheets is their liquidity position, which we take as a second measure to stratify banks. Of course, solvency is not determined solely by a bank's degree of capitalization, or liquidity, but rather by an amalgam of factors. Accordingly, we also split banks using a more comprehensive measure: their ex-ante probability of failure. These are determined by estimating a default prediction model similar to, for example, Park and Peristiani (1998) and Calomiris and Mason (2003b). While this logit model may be of independent interest, we refer the reader to Appendix C for details. We here focus on assessing differences in deposit flows during a run across banks with a high and a low probability of default. For each of these stratifications, we use the median as the cutoff value. As a final way to distinguish solvent from insolvent banks, we assume that ex-post actual solvencies are known in real time. This approximates the case of perfect information, as if depositors were able to perfectly predict which banks would fail. This stratification is analogous to the one used in Saunders and Wilson (1996). For each classification, one can stratify on the basis of the entire sample or based on a particular time period. We stratify at every point in time, but our conclusions are insensitive to this choice. Finally, the previous balance sheet information is only used to stratify banks into good and bad banks. In later robustness checks we also incorporate such information directly in the reduced form. Summary statistics on stratifications are contained in Appendix C.

Table 1 contains the main results. For each of the stratifications used panel A measures the impact of the 2004 run on the quantity of deposits for good, bad, and insured banks. The coefficient in the upper panel can be interpreted directly as the percentage change in the deposit base of the different groups of banks. We focus on the contemporaneous impact.

Panel B evaluates a number of hypotheses of interest. Particularly, we evaluate the significance of the difference in deposit response between different types of banks.

TABLE 1. Response to the 2004 run: deposits.

	Capital	**Liquidity**	**Ex ante**	**Ex post**
Panel A: Impact				
Bad banks	–0.06	–0.06	–0.07	–0.10
Good banks	–0.07	–0.06	–0.05	–0.05
Insured banks	0.05	0.05	0.05	0.05
Panel B: Cross-sectional tests				
Fundamental: *Bad* − *Good*	0.01	–0.00	–0.02	–0.05
% *Bad* − *Good* < 0	0.40	0.51	0.81	0.82
Panic: *Good* − min(0, *Insured*)	–0.06	–0.05	–0.04	–0.05
% *Good* − min(0, *Insured*) < 0	0.93	0.94	0.93	0.95
Insurance: *Insured*	0.05	0.05	0.05	0.05
% *Insured* > 0	0.81	0.81	0.83	0.79
Panel C: Aggregate effects				
Outflow uninsured:[a]	–0.06	–0.06	–0.06	–0.06
*due to fundamentals[b]	0%	2%	23%	2%
*due to panic[c]	100%	98%	77%	98%
*absorbed by insured[d]	27%	23%	24%	25%

Notes: Panel A: posterior median. Panel B: odd rows contain posterior median, even rows indicate percentage of posterior that confirms the hypothesis. Panel C: let capital letters B, G, I denote the impact coefficients from panel A for bad, good and insured, respectively, and V_B, V_G, and V_I the pre-crisis volume of deposits in the respective groups.
a. Percentage change in aggregate uninsured deposits is calculated as $(BV_B + GV_G)/(V_B + V_G)$, which is then decomposed into two parts: b + c.
b. The fundamental part $(B - G)V_B/(BV_B + GV_G)$.
c. A panic-driven part $G(V_B + V_G)/(BV_B + GV_G)$.
d. The inflow at insured banks is $IV_I / \mid BV_B + GV_G \mid$.

The first test verifies whether the outflow at bad banks is more severe than the response at good banks. This provides a test of the fundamental view. The second test evaluates the response at good banks. The panic view on bank runs predicts that depositors will run on fundamentally sound banks, too. The test verifies whether good banks lose deposits, outright or relative to insured banks. The third test verifies whether insured banks absorb deposits during bank runs. For all these hypotheses, the table reports the median estimate and the associated acceptance rate: the fraction of the posterior consistent with the hypothesis.

We are now ready to quantify the contribution of the two competing views to bank runs. The first two rows of panel A show the effect of the 2004 run on uninsured good and bad banks. First, irrespective of the measure used to stratify, good banks invariably are subject to the run. The effect is substantial and amounts to at least 5% of good banks' deposit base. Such an outflow is not observed at insured banks (panel A, row 3), as corroborated by the according acceptance rate on the difference between good and insured (panel B, row 4). This establishes the relevance of the panic view. Banks that do not have deposit insurance but have sound fundamentals face significant deposit outflows. Hence, solid fundamentals are not a substitute for being insured.

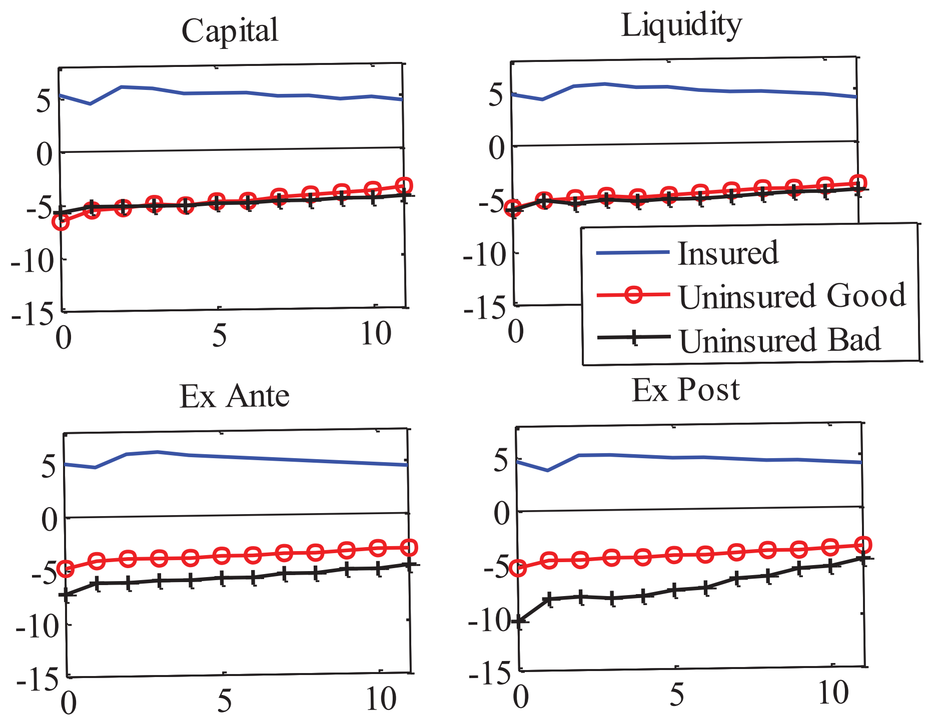

FIGURE 5. Deposit responses to the 2004 run (medians).

Second, bad banks also lose deposits, at least 6%. For the ex-ante and ex-post stratifications, we find significantly stronger outflows at bad banks relative to good banks (with acceptance rates >80%, panel B, row 2). Thus, fundamentally flawed banks face even more significant outflows during a bank run. This difference can be quantitatively large: the table indicates that bad banks can face runs twice as severe as those observed at good banks (panel A, last column). This finding establishes the relevance of the fundamental view on bank runs. Thus, importantly, we find evidence in support of both views on bank runs. Figure 5 plots the median response of deposits across the different types of banks for the various stratifications.

For these results to have policy relevance, however, the relative importance of the two views needs to be assessed. Therefore, in addition to the impulse responses to the 2004 run, the bottom panel of the table computes the implied aggregate effects. These are calculated by taking into account the volume of deposits in the respective groups of banks. They make it possible to quantify the aggregate importance of deposit flows between the different types of banks, as well as effects on the deposit market as a whole. In panel C, the first row calculates the total outflow of uninsured deposits. In aggregate terms, the uninsured deposit market shrinks by 6%. Recall from Figure 3 that this outflow is close to permanent. The next two rows of panel C decompose the aggregate outflow into the part driven by fundamentals and the part

TABLE 2. Response to the 2004 run: interest rates.

	Capital	**Liquidity**	**Ex ante**	**Ex post**
Panel A: Impact				
Bad banks	0.27	0.29	0.35	0.38
Good banks	0.31	0.26	0.20	0.25
Insured banks	–0.19	–0.21	–0.22	–0.12
Panel B: Cross-sectional tests				
$Bad - Good$	–0.04	0.02	0.14	0.14
$\%\ Bad - Good > 0$	0.41	0.58	0.76	0.59
$Good - \max(0, Insured)$	0.28	0.23	0.18	0.20
$\%\ Good - \max(0, Insured) > 0$	0.93	0.96	0.87	0.93
$Insured$	–0.19	–0.21	–0.22	–0.12
$\%\ Insured < 0$	0.75	0.76	0.78	0.70

caused by panic. It turns out that the panic view is the primary contributor to the run in our sample. Fundamental effects—that is, the more severe outflows at bad banks—explain no more than 25% of the total deposit outflow.[5] Whichever way one classifies good and bad banks, good banks always lose a significant fraction of their deposits. From an aggregate perspective this outflow is the main contributor. The final row of panel C measures the inflow of deposits at insured banks as a fraction of the outflow of uninsured deposits. Insured banks absorb about one quarter of the outflow from the uninsured deposit market. Hence, while insured banks are not subject to the run, they are not necessarily viewed as a safe haven. The fact that such a large part of the outflow disappears from the deposit market suggests the potential severity of bank runs for the real economy. While the money flows out of the deposit market, we do not know whether it ends up in "socks or stocks". We refrain from quantifying the impact beyond the deposit market. For evaluations of the real effects of bank runs, see for example Bernanke (1983) and Calomiris and Mason (2003a).

From a policy perspective, these results suggest that the primary concern is shielding fundamentally solvent banks from bank runs. In our sample, this is readily achieved by deposit insurance: insured banks were not subject to depositor runs. Since poor fundamentals can severely aggravate runs, there is scope for fundamentals-based regulation, too. In our sample, however, this seems to be of second-order importance.

A final result of interest can be observed in Table 2, which contains the interest rate response for the different types of banks. We know from Figure 3 that the inflow of deposits at insured banks is not demand-driven: there is no increase in the deposit interest rate of insured banks, while uninsured banks increase theirs. First, note that the

5. The aggregate fundamental effect is small in the ex-ante case because the outflow at bad banks is not much worse than that at good banks, while good and bad banks alike lose a lot. For the ex-post case, the outflow for the average bad bank is much more severe than that of good banks, but it now applies to a relatively small fraction of banks/deposits.

responses in the table are in percentage points per year. The increase in the uninsured interest rate is not very large—though it may mask some intra-group heterogeneity. In cases where we observe significant fundamental effects in Table 1, Table 2 shows that there is a tendency for bad banks (that face larger deposit outflows) to increase their deposit rate by more than good banks. Again, this increase does not appear too big quantitatively. Moreover, the results do not establish a causal link from the increased interest rate to the drop in quantities, or vice versa—they occur simultaneously. Nonetheless, two related explanations for this phenomenon are particularly plausible. A first interpretation sees the interest hike as a response: banks in trouble increase their deposit rates as a "gamble for resurrection", an attempt to keep deposits from flowing away. A second interpretation reverses that logic and sees the interest rate hike as a cause: it signals to depositors that the bank is in trouble, and depositors therefore run (more). These types of effects are suggested by, among others, Hellmann, Murdock, and Stiglitz (2000).

5. Robustness

The previous results are very robust. We here summarize the different types of tests we ran and indicate the extent of variation we find. Detailed results are available in an Online Appendix. A first series of tests concerns alternative controls (incorporating balance sheet information in the reduced form, regional subsample). A second set of tests considers variations on the way bank runs are identified. Rather than imposing the restrictions on the entire group of uninsured banks, we impose them on only one of the subgroups (bad, resp. good). We also split good and bad banks based on quartiles rather than the cross-sectional median. In addition, Appendix B discusses two alternative ways of identification, one which ensures that the response to bank runs is not merely driven by heterogeneous responses to common shocks, another which rules out contemporaneous demand shifts at insured banks. A third batch of checks uses alternative variable specifications (log-differences instead of log-levels, alternative interest rate measure). A final collection of tests verifies alternative model specifications (allowing for more heterogeneity within and interaction across groups).

Across this variety of tests, we find that the summer of 2004 is always identified as a bank run; the fundamental effect is present in about half the tests; and the panic effect is always there. Regarding the aggregate effects, the maximum contribution of fundamentals to the uninsured outflow is 36% (and thus panic minimally 64%), while the insured banks typically absorb 20%–30% of the entire uninsured outflow. In the Appendix, we also address some alternative interpretations and evaluate a number of extensions. For instance, we elaborate on how the approach only requires relative credibility of insurance and how that relates to the variation between implicit and explicit deposit insurance during our sample period. Among the extensions, it turns out that there are no substantially different responses between state-owned and foreign banks, or between insured banks with good and bad fundamentals.

6. Discussion and Conclusions

There are important empirical challenges that substantially impair inference based on traditional tests of bank runs. Before concluding, we here describe these hurdles and discuss how our approach addresses them. A first issue is the definition of a crisis. Calomiris and Gorton (1991) point out the difficulty of defining what constitutes a bank run in the data. Canova (1994), Gorton and Winton (2003), and Ennis (2003) show how this has led authors to disagree on whether a particular period constitutes a crisis, thus introducing *subjectivity* in empirical tests and their results.

Subjectivity aside, a second difficulty in any empirical analysis of runs lies in the fact that the *exogeneity* of the bank run is questionable. This is especially relevant in the context of assessing the fundamental nature of bank runs. For instance, it is not because deposit outflows correlate with recessions (i.e. a fundamental factor) that bank runs are due to recessions (as held by the fundamental view). Recessions themselves should lower deposit demand of banks, which are faced with a lower loan demand schedule during recessions. Thus, the observation that deposit flows exhibit a reduced-form correlation with fundamentals does not, in itself, necessarily constitute evidence for the fundamental view. In part, this endogeneity concern is the basis of the more recent work which studies the effect of events that are—arguably—exogenous. Examples are Iyer and Peydró (2011) and Iyer and Puri (2012), who investigate the effects of a bank fraud discovery in India.

A related complicating factor in assessing the relevance of the different theories underlying bank runs is an implicit *exhaustivity* assumption present in the aforementioned studies. Similarly to event-studies, they necessarily assume that the run is the only thing that occurs during the period of interest. Even if the event under consideration is truly exogenous, deposit responses can be convoluted by concurrent events, such as endogenous demand responses in anticipation of a recession caused by the event. The effect of the exhaustivity assumption can also be seen in a model context, such as Diamond and Dybvig (1983). Extant empirical strategies in the field of bank runs typically cannot distinguish liquidity preference type shocks from bank runs.

Our empirical approach addresses the previously mentioned issues of subjectivity, exogeneity and exhaustivity. Particularly, in our approach crisis identification relies on explicit a priori restrictions. These force one to be very specific about definitions, which reduces the scope for subjectivity. Applying those restrictions to a reduced-form empirical model, the approach extracts exogenous structural shocks from raw data. This makes the method more generally applicable and obfuscates the need for restricting attention to data which are by themselves assumed exogenous. Contrary to traditional tests of bank runs, the present method also does not require making an exhaustivity assumption, namely that the run is the only thing that happens during the particular period of interest. Rather, the method restricts attention to the run, while controlling for earlier and contemporaneous alternative shocks, such as liquidity preference or demand shocks. In sum, the method we propose in this paper addresses the issues of subjectivity, exogeneity, as well as exhaustivity. Related to our work, Boyd, De Nicolò, and Loukoianova (2010) use SVAR techniques to address similar issues in the analysis of aggregate crisis data. Finally, and in accordance with contemporary financial crisis theory, our tests are such that the different views need not be mutually exclusive: panic and fundamental effects can coexist.

In sum, we propose a cross-sectional twist to SVAR methods. Imposing theory-implied heterogeneity restrictions on Russian deposit market data suggests that there was one bank run during the sample period 2002–2007, in line with narrative evidence. Conditional on a run, we quantify the effects of the two main theories of bank runs. For our sample, we find that the panic view is more important than the fundamental view. While we do find evidence that the fundamental view matters, its aggregate effects are relatively small. Though effects observed in the Russian deposit market in our sample period may not generalize to always and everywhere, they do have important policy implications. Foremost, our results suggest that panic effects are a real concern. Contrary to uninsured banks with good fundamentals, banks with deposit insurance were safeguarded against runs. This suggests that purely fundamentals-based regulation may not be sufficient to shield solvent banks from bank runs.

On a methodological note, our results suggest that the cross-section provides valuable information both in the process of identification (with few identifying restrictions being required, and external validation successful) and testing (with cross-sectional differences discriminating between otherwise observationally equivalent theories).

Appendix A: Estimation and Inference

This appendix provides additional detail on estimation and simulation of the empirical model. We first discuss the equivalence between the Maximum Likelihood, Ordinary, and Generalized Least Squares estimators. Rewriting the model (4) such that group-specific shocks are absorbed by the matrix of explanatory variables—as discussed in the text—implies the model has block-diagonal (kronecker) residual covariance. It is well known that GLS is the ML estimator in this case. For the model written in this form the level at which the covariance is not diagonal is the unit level (i.e. Σ_u measures the covariance across equations within a unit). At the unit level, the right-hand-side variables are the same in each equation. OLS is therefore equivalent to GLS. As a result, given a noninformative conjugate prior, OLS retrieves the posterior mean.

We now turn to posterior simulation. The primary interest of the present study lies in uncovering group-level behavior. This implies that while the reduced-form equation (4) is estimated on individual-level data, our identification and results are based on group-wise impulse responses. It is convenient to write the model in its smaller within-group-averaged form (where we no longer need to track individual units). This system is sufficient for the computation of group-specific impulse responses. Within-group averaging is equivalent to premultiplying the system (4) with the known $GJ \times NJ$ matrix $A = (M'M)^{-1}M'$:

$$Ay_t = AX_t'\beta + v_t + Au_t, \quad v_t + Au_t \sim N(0, \Sigma_v + A(I_N \otimes \Sigma_u)A'),$$

where Ay_t is the $GJ \times 1$ vector of group-wise averaged variables, $Ay_t = (\bar{y}_t^{1\prime}, \ldots, \bar{y}_t^{G\prime})'$, and analogous definitions apply to AX_t' and Au_t. Such linear transformations preserve the posterior family but change its scale. Therefore, one can sample from the posterior by drawing from the Normal-Inverse Wishart distribution: $B \sim N(\beta | A\Sigma A')$, $AS \sim IW(A\Sigma A', GJ, T)$, where the $GJ \times GJ$ covariance matrix $A\Sigma A' = \Sigma_v + A(I_N \otimes \Sigma_u)A' = \Sigma_v + \mathrm{diag}(N_1^{-1}\Sigma_u, \ldots, N_G^{-1}\Sigma_u)$, and IW is indexed by scale, dimension, and degrees of freedom, respectively.

Considering the group-wise system is a convenient way of ensuring draws (B, S) are consistent with the particular structure of Σ in equations (3)–(4). That is, an unconstrained draw $S \sim IW(\Sigma, NJ, T)$ need not have the block structure $\Sigma = M\Sigma_v M' + I_N \otimes \Sigma_u$. Drawing AS from $IW(A\Sigma A', GJ, T)$ and subsequently transforming the system back to unit-level form (i.e. multiplying with M) will satisfy this structure. Note that, since inference is entirely based on group-wise statistics, the last step is not required for the present analysis.

Appendix B: Alternative Identification Restrictions

This appendix discusses two possible constellations of responses that the baseline restrictions do not rule out explicitly. First, are the cross-sectional differences found

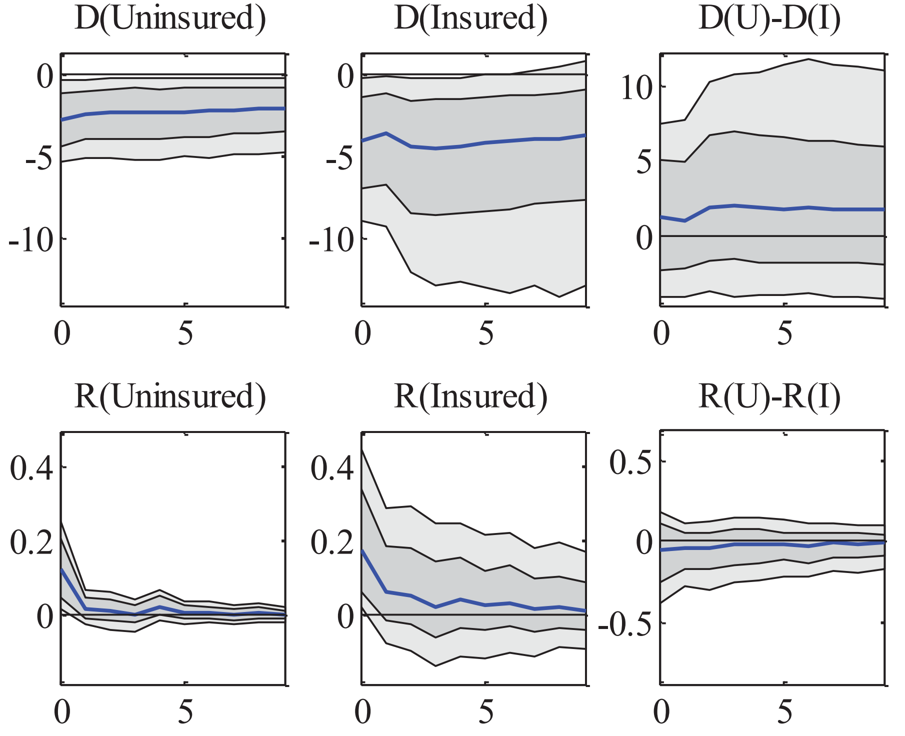

FIGURE B.1. Regular supply shock: impulse responses.

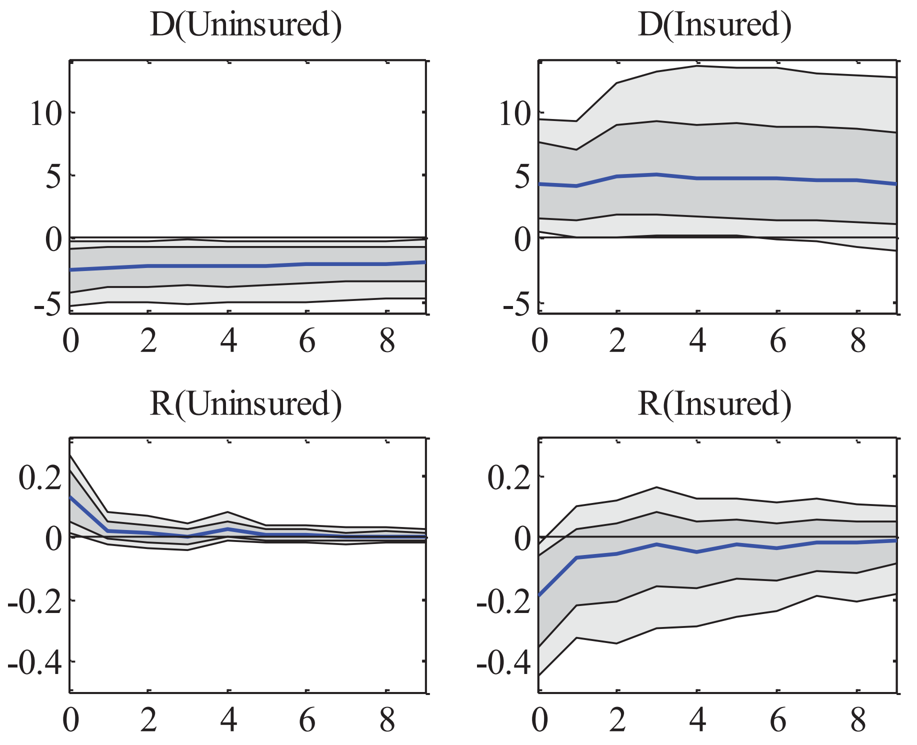

FIGURE B.2. Negative supply in uninsured, positive supply in insured (impulse responses).

really particular to bank runs, or do they reflect general differences in responsiveness across banks? As a simple way of establishing this, we here apply an alternative identification scheme: we identify a negative supply shock in both the insured and uninsured group, by imposing the restrictions

$$\begin{cases} \Delta D^U & < 0 \\ \Delta R^U & \geqslant 0 \\ \Delta D^I & < 0 \\ \Delta R^I & \geqslant 0. \end{cases}$$

The results are contained in Figure B.1. It is apparent that there is no significant difference between insured and uninsured banks conditional on such "regular" supply shocks. Further, for specifications based on three groups, the cross-sectional differences conditional on runs are different from those conditional on regular supply shocks. This suggests that any cross-sectional difference found during runs is not just a consequence of regular differences in elasticities between insured and uninsured banks.

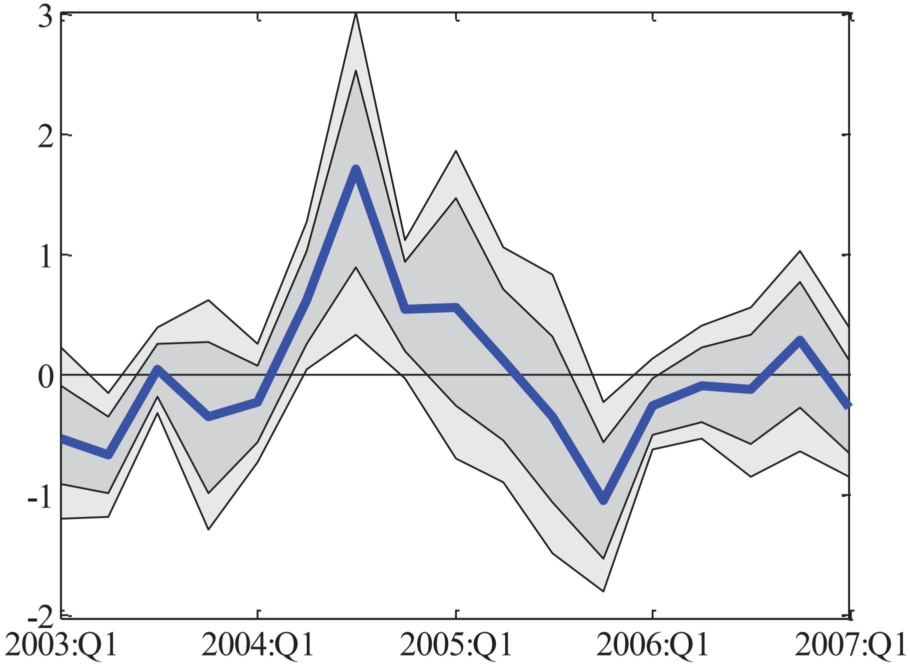

FIGURE B.3. Negative supply in uninsured, positive supply in insured (shock).

Second, if one is willing to assume a higher degree of credibility of deposit insurance, one can impose the more stringent restrictions

$$\begin{cases} \Delta D^U & < 0 \\ \Delta R^U & \geqslant 0 \\ \Delta D^I & > 0 \\ \Delta R^I & \leqslant 0. \end{cases}$$

These restrictions identify bank runs as an outflow of uninsured deposits that (in part) results in a supply-driven inflow of deposits at insured banks. Note that these restrictions explicitly rule out insured bank demand shifts, something the baseline restrictions do not do. Applying these more stringent restrictions allows us to evaluate whether insured demand shifts are a concern quantitatively. The results thus obtained are presented in Figures B.2 and B.3 and turn out to be quantitatively similar to those in the paper, both in terms of impulse responses as well as the time series of the shock.

Appendix C: Data

The original data cover the period from 1999 to 2007. Because the constructed interest rate series exhibit a break in 2001 (due to changes in variable definitions) we limit our sample to observations after the break. We further lose four quarters of observations

TABLE C.1. Summary statistics—different stratifications.

	(I) Insured		(II) Uninsured		(III) Ex-ante bad		(IV) Ex-ante good		(V) Ex-post bad		(VI) Ex-post good	
	Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.
Log-difference (household deposits)	0.128	0.401	0.118	0.434	0.115	0.417	0.122	0.450	0.021	0.628	0.122	0.425
Implicit interest rate	0.046	0.037	0.073	0.040	0.074	0.039	0.072	0.041	0.064	0.047	0.073	0.040
Log (Assets)	8.758	2.431	6.586	1.763	6.839	1.849	6.332	1.634	6.060	1.556	6.605	1.767
Capital/Assets	0.176	0.114	0.230	0.156	0.239	0.156	0.222	0.154	0.306	0.190	0.228	0.154
ROA	0.007	0.016	0.006	0.014	0.004	0.010	0.008	0.017	0.006	0.041	0.006	0.012
Liquid assets/Assets	0.187	0.157	0.219	0.160	0.127	0.067	0.312	0.172	0.209	0.193	0.220	0.159
Bad loans/Assets	0.008	0.017	0.008	0.017	0.011	0.021	0.006	0.013	0.015	0.040	0.008	0.016
Nongovernment securities/Assets	0.019	0.037	0.075	0.118	0.106	0.143	0.043	0.074	0.151	0.179	0.072	0.115
Term deposits: firms/Assets	0.113	0.141	0.065	0.087	0.058	0.075	0.071	0.096	0.042	0.071	0.066	0.087
Term deposits: households/Assets	0.099	0.119	0.154	0.139	0.138	0.119	0.170	0.155	0.092	0.124	0.156	0.139
Bank-quarter observations	705		15005		7508		7497		522		14483	

TABLE C.2. Default prediction model (logit).

Variables	
Log (Assets)	–0.13*
	(0.08)
Capital/Assets	–1.37
	(0.85)
ROA	–22.12***
	(5.53)
Liquid Assets/Assets	–7.33***
	(2.19)
Bad Loans/Assets	7.39***
	(2.38)
Non-Government Securities/Assets	3.13***
	(0.65)
Term Deposits of Firms/Assets	–5.71**
	(2.22)
Term Deposits of Households/Assets	–5.72***
	(2.12)
Observations	21193
Pseudo-R^2	0.28
AUR	0.867

Notes: Standard errors in parentheses. AUR measures the percentage of correctly classified events relative to one minus the percentage of correctly classified non-events. Values above 0.8 are typically considered very successful (see e.g. Hosmer and Lemeshow 2000).
***Significant at 1%; **significant at 5%; *significant at 10%.

because of the four lags in our reduced form. As a result, the final sample covers 17 quarters (2003:Q1–2007:Q1).

Bank panels are unbalanced because some banks fail, some merge, and some are founded during the sample period. If a bank merged or was acquired we treat the resulting larger bank as a new entity. There are a total of 1,225 banks (15,710 bank-quarters) in our sample, of which 60 (705) correspond to insured and 1,165 (15,005) to uninsured.

The list of state-owned banks is compiled from Matovnikov (2002) and Mamontov (2005); the state ownership definition does not vary over time. The foreign ownership definition relies on Bank of Russia's quarterly lists of 100% foreign-owned banks (available at http://cbr.ru/analytics/bank_system/) and varies over time.

Table C.1 presents descriptive statistics for all the variables, including those that are used to stratify banks into good and bad. Specifically, columns (I)–(II) contain sample averages and standard deviations for the insured and the uninsured group of banks separately.

Columns (III)–(IV) in Table C.1 provide similar statistics within the uninsured group of banks, when it is stratified along the ex-ante dimension. Table C.2 contains the estimated default prediction model for the full sample (2002:Q1–2007:Q1). The ex-ante stratification in the paper is based on a recursive estimate of the same specification, where the estimate is updated every period.

Finally, columns (V)–(VI) of Table C.1 split the uninsured group according to the ex-post stratification. Here, a bank is stratified as bad at date t if it fails between t and $t+6$. Regarding the two remaining stratifications, we measure capitalization as the ratio of capital over assets, and liquidity as the ratio of reserves and government securities over assets.

References

Allen, Franklin and Douglas Gale (1998). "Optimal Financial Crises." *Journal of Finance*, 53, 1245–1284.

Amir Ahmadi, Pooyan and Harald Uhlig (2009). "Measuring the Dynamic Effects of Monetary Policy Shocks: A Bayesian FAVAR Approach with Sign Restrictions." Working paper, University of Chicago.

Bernanke, Ben S. (1983). "Non-monetary Effects of the Financial Crisis in the Propagation of the Great Depression." *American Economic Review*, 73(3), 257–276.

Boyd, John, Gianni De Nicolò, and Elena Loukoianova (2010). "Banking Crises and Crisis Dating: Theory and Evidence." CESifo Working Paper No. 3134.

Calomiris, Charles W. and Gary Gorton (1991). "The Origins of Banking Panics, Models, Facts, and Bank Regulation." In *Financial Markets and Financial Crises*, edited by Glenn R. Hubbard. University of Chicago Press.

Calomiris, Charles W. and Joseph R. Mason (2003a). "Consequences of Bank Distress During the Great Depression." *American Economic Review*, 93(3), 937–947.

Calomiris, Charles W. and Joseph R. Mason (2003b). "Fundamentals, Panics and Bank Distress during the Depression." *American Economic Review*, 93(5), 1615–1647.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca (2013). "Housing Finance and Monetary Policy." *Journal of the European Economic Association*, 11, 101–122.

Canova, Fabio (1994). "Were Financial Crises Predictable?" *Journal of Money, Credit and Banking*, 26, 102–124.

Canova, Fabio (2004). "Testing for Convergence Clubs in Income per Capita: A Predictive Density Approach." *International Economic Review*, 45, 49–77.

Canova, Fabio (2007). *Methods for Applied Macroeconomic Research*. Princeton University Press.

Canova, Fabio, and Matteo Ciccarelli (2009). "Estimating Multicountry VAR Models." *International Economic Review*, 50, 929–961.

Canova, Fabio and Matteo Ciccarelli (2013). "Panel Vector Autoregressive Models: A Survey." *Advances in Econometrics*, 31, forthcoming.

Canova, Fabio and Evi Pappa (2006). "Does it Cost to be Virtuous? The Macroeconomic Effects of Fiscal Constraints." NBER Working Paper No. 11065.

Chari, V. V. and Ravi Jagannathan (1988). "Banking Panics, Information, and Rational Expectations Equilibrium." *Journal of Finance*, 43, 749–761.

Demirgüç-Kunt, Asli and Enrica Detragiache (1998). "The Determinants of Banking Crises: Evidence from Developed and Developing Countries." *IMF Staff Papers*, 45, 81–109.

Diamond, Douglas W. and Philip H. Dybvig (1983). "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy*, 91, 401–419.

Donaldson, Glen R. (1992). "Sources of Panics: Evidence from the Weekly Data." *Journal of Monetary Economics*, 30, 277–305.

Eickmeier, Sandra and Tim Ng (2011). "How Do Credit Supply Shocks Propagate Internationally: A GVAR Approach." Bundesbank Discussion Paper No. 27/2011.

Ennis, Huberto M. (2003). "Economic Fundamentals and Bank Runs." *Federal Reserve Bank of Richmond Economic Quarterly*, 89, 55–71.

Ennis, Huberto M. and Todd Keister (2009). "Bank Runs and Institutions: The Perils of Intervention." *American Economic Review*, 99(4), 1588–1607.

Friedman, Milton and Anna J. Schwartz (1963). *A Monetary History of the United States, 1867–1960*. Princeton University Press.

Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin (2004). *Bayesian Data Analysis*. Chapman and Hall.

Goldstein, Itay and Ady Pauzner (2005). "Demand-Deposit Contracts and the Probability of Bank Runs." *Journal of Finance*, 60, 1293–1327.

Goldstein, Itay and Assaf Razin (2012). "Review of Theories of Financial Crises." NBER Working Paper No. 18670.

Gorton, Gary (1988). "Banking Panics and Business Cycles." *Oxford Economic Papers*, 40, 751–781.

Gorton, Gary and Andrew Winton (2003). "Financial Intermediation." Chapter 8 in *Handbook of the Economics of Finance I*, edited by George M. Constantinides, Milton Harris, and Rene M. Stulz. Elsevier Science.

Helbling, Thomas, Raju Huidrom, M. Ayhan Kose, and Chris Otrok (2011). "Do Credit Shocks Matter? A Global Perspective." *European Economic Review*, 55, 340–353.

Hellmann, Thomas F., Kevin C. Murdock, and Joseph E. Stiglitz (2000). "Liberalization, Moral Hazard in Banking and Prudential Regulation: Are Capital Requirements Enough?" *American Economic Review*, 90(1), 147–165.

Hosmer, David W. and Stanley Lemeshow (2000). *Applied Logistic Regression*. Wiley.

Iyer, Rajkamal and José Luis Peydró (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(4), 1414–1445.

Jacklin, Charles J. and Sudipto Bhattacharya (1988). "Distinguishing Panics and Information-Based Bank Runs: Welfare and Policy Implications." *Journal of Political Economy*, 96, 568–592.

Jarociński, Marek (2010). "Responses to Monetary Policy Shocks in the East and the West of Europe: A Comparison." *Journal of Applied Econometrics*, 25, 833–868.

Karas, Alexei and Koen Schoors (2005). "Heracles or Sisyphus? Finding, Cleaning and Reconstructing a Database of Russian Banks." Working paper, Ghent University (http://www.feb.ugent.be/nl/Ondz/wp/Papers/wp_05_327.pdf).

Love, Inessa and Lea Zicchino (2006). "Financial Development and Dynamic Investment Behaviour: Evidence from Panel VAR." *Quarterly Review of Economics and Finance*, 46, 190–210.

Mamontov, Alexei (2005). "Gosudarstvo v Bankakh: Zlo ili Blago? (State Banks: Good or Bad?)" *Natsional'nyi Bankovskii Zhurnal*, 12, 24.

Martinez Peria, Maria S. and Sergio L. Schmukler (2001). "Do Depositors Punish Banks for Bad Behavior?: Market Discipline, Deposit Insurance, and Banking Crises." *Journal of Finance*, 56, 1029–1052.

Matovnikov, Mikhail (2002). "Nadezhnost' Banka Tesno Svyazana so Strukturoi ego Aktsionernogo Kapitala (A Bank's Reliability is Closely Connected with the Structure of its Equity Capital)." Tsentr Ekonomicheskogo Analiza-Interfax, Moscow.

Morris, Stephen and Hyun S. Shin (1998). "Unique Equilibrium in a Model of Self-fulfilling Currency Attacks." *American Economic Review*, 88(3), 587–597.

Nason, James M. and Ellis W. Tallman (2012). "Business Cycles and Financial Crises: the Roles of Credit Supply and Demand Shocks." FRB Cleveland Working Paper No. 1221.

Park, Sangkyun and Stavros Peristiani (1998). "Market Discipline by Thrift Depositors." *Journal of Money, Credit and Banking*, 30, 374–364.

Peck, James and Karl Shell (2003). "Equilibrium Bank Runs." *Journal of Political Economy*, 111, 103–123.

Saunders, Anthony and Berry Wilson (1996). "Contagious Bank Runs: Evidence from the 1929–1933 Period." *Journal of Financial Intermediation*, 5, 409–423.

Schumacher, Liliana (2000). "Bank Runs and Currency Run in a System without Safety Net: Argentina and the 'Tequila' Shock." *Journal of Monetary Economics*, 46, 257–277.

Tompson, William (2004). "Banking Reform in Russia: Problems and Prospects." OECD Working Paper No. 410.

Uhlig, Harald (2005). "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics*, 52, 381–419.
Zykova, Tatyana (2004). "Chernyi Spisok Zubkova" (Zubkov's black list). *Rossiyskaya Gazeta*, 3478, 18 May.

Supporting Information

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

Online Appendix