Information-Based Bank Runs or Panics?*

Ferre De Graeve†
Federal Reserve Bank of Dallas, Dallas, TX, USA

Alexei Karas‡
Roosevelt Academy, Middelburg, The Netherlands

November 20, 2008

Abstract

This paper proposes a novel approach to distinguish between different theories of bank runs. Our method identifies bank runs from raw deposit market data, while controlling for the effects of alternative shocks. We construct hypothesis tests, conditional on bank runs, which quantify the role of panic-induced deposit withdrawals as well as the extent to which withdrawals are based on information. We provide an application for the Russian deposit market. The method identifies one severe bank run. Information external to the model confirms the timing of the run. We find evidence in favor of both the panic and the information-based views of bank runs. In particular, panic effects induce significant runs at solvent banks with uninsured deposits, relative to banks whose deposits are insured. In addition, insolvent banks face runs that are four times as harsh, which establishes the significance of information-based theories. In the aggregate, panic accounts for almost 90% of the deposit run, while information is responsible for the remaining 10%. The fact that, in our application, both the panic and fundamentals view exert effects has strong policy implications. While market discipline adds to the stability of the deposit market, it does not eliminate the possibility of panic runs. Our approach builds on structural identification techniques frequently used in macroeconomics. We extend these time-series methods by extracting information from cross-sectional heterogeneity. We show that such information is very useful for identifying structural shocks, cross-validating the model and extracting testable implications. The method can easily be extended to apply to alternative types of shocks, as well as other markets.

Keywords: Bank run, Panic, Information-based, Identification, SVAR, panel-VAR
JEL: E5, G01, G21

---

*We would like to thank Olivier De Jonghe, Roland Meeks, Larry Wall, and participants at the Federal Reserve System Committee Meeting on "Financial Structure and Regulation" in Boston for useful comments and valuable suggestions. The views expressed herein do not necessarily reflect those of the Federal Reserve Bank of Dallas, or the Federal Reserve System.

†ferre.degraeve@dal.frb.org.

‡a.karas@roac.nl.

1 Introduction

Banking crises and bank runs are among the more severe events economies have had to face in the past century. The costs involved in resolving crises of this type amount to several percentages of GDP.[1] The incentive for depositors to run on their bank is clear: fear of losing their money. In order to judge the appropriate policy response, however, it is of utmost importance to know whether bank runs occur on the basis of fundamental information or are rather a manifestation of panic. If bank runs are panic-driven, economy wide action is called for. Policy responses include the provision of liquidity facilities or deposit insurance. Alternatively, when bank runs are based on fundamentals, market-discipline-type regulation (imposing constraints on bank behavior) or government recapitalization may be more valuable.

In this paper, we propose a method that identifies bank runs and simultaneously evaluates the extent to which they are driven by fundamentals and how much is attributable to panics. The identification of bank runs is obtained by contrasting the flows of deposits among banks whose deposits are insured, relative to bank whose deposits have no guarantees. Importantly, our approach controls for endogenous dynamics, simultaneity and the presence of other contemporaneous influences, such as demand shocks. The hypothesis tests we perform to discriminate among the information-based (Jacklin and Bhattacharya 1988, Allen and Gale 1998) and panic theories (Diamond and Dybvig 1983, Postlewaite and Vives 1987) are therefore purely conditional on the bank run. To arrive at this framework for validating the different theories, we apply structural identification methods from macroeconomics to microeconomic data.

A major difficulty in assessing the driving forces of bank runs is singling out the run from other factors. The literature can be divided into two different strands. First, Friedman and Schwartz (1963) *narratively* classify particular episodes in US history as bank runs. Their subsequent analysis suggests that these runs are characterized by panic effects, without fundamentals driving them. This panic view has been contested by many, including Gorton (1988), Saunders and Wilson (1996) and Calomiris and Mason (2003). By and large, the approach taken in this strand of the literature is to take the episodes identified by Friedman and Schwartz as given and show that there are fundamental factors which can explain substantive parts of the observed deposit outflows. The underlying fundamental factors can be international, national, regional, sector or bank-specific in nature (see, in particular, the overview in Calomiris and Mason 2003).

Second, the response to the narrative approach is able to assess whether or not fundamentals are at work during the narratively identified episodes. However, it does not necessarily single out other driving forces. For instance, if deposit runs are more likely to occur in the face of bad fundamentals, then banks' demand for deposits will be accordingly low. Finding a correlation between fundamentals and deposit flows can then -spuriously- lead to the conclusion that depositors run based on fundamental information. As a result, there is scope for concern as to whether the narratively identified periods are indeed solely driven by a run. Indeed, in any given period, observed deposit flows are an equilibrium result of responses to multiple sources of shocks, and a bank run may only be one of them.[2] One way the litera-

[1] For banking crises in general, World Bank (2001) estimates the fiscal costs between 3 and 50% of GDP.

[2] Note that this concern becomes more prominent the lower the frequency of the data is, and the frequency

ture has tried to control for other types of shocks is to analyze arguably *exogenous* events. Iyer and Peydro-Alcalde (2007) and Iyer and Puri (2008), for instance, study the outflow of deposits following a case of (unanticipated) fraud at a bank in India. Khwaja and Mian (2008) investigate lending responses conditional on unanticipated nuclear tests in Pakistan. The requirement of the presence of an exogenous event poses a very stringent constraint on the data. As a result, it is difficult to assess how the results of these studies translate to other countries or periods.

The method we propose in this paper addresses both these issues. First, our approach filters out bank runs from raw deposit market data. This makes the method more generally applicable, and obfuscates the need of searching for data which contain an exogenous event. Second, our approach does not require to make the assumption that the run is the only thing that happens during any given observation period, contrary to both the narrative and the exogenous event approaches. This makes our tests fully conditional on bank runs (not confounded by alternative shocks), which is ultimately what is of interest to policymakers.

Similar to most of the literature testing the effects of bank runs, we start by estimating a reduced form model for the data of interest. However, we then make an additional step and transform this reduced form into a structural form. This enables us to study the effect of structural shocks (i.c. bank runs) and identify endogenous interactions between deposits and interest rates across various types of banks. Our approach builds on the structural VAR (SVAR) literature in macroeconomics. We borrow some of the tools developed there and apply them to micro data. We go beyond the macroeconomic literature by extracting information from cross-sectional heterogeneity among distinct groups of banks. In turn, we also contribute to the SVAR literature. Heterogeneity, it turns out, is highly informative both for the identification of structural shocks as well as for hypothesis testing. Additionally, and among other things, we show that our model's implications are consistent with information external to the model, which macroeconomic SVARs typically do not accomplish, as shown by Rudebusch (1998).

We implement our method using Russian deposit market data. This market is particularly interesting for at least two reasons. First, there exists cross-sectional variation in the degree of deposit insurance among different groups of banks. Second, during the course of our sample period, this market faced severe distress. Taken together, the Russian deposit market is an ideal environment to test the power of the method as well as validate the various theories on the incentives to run on a bank.

Our results show that over the period 2002-2007, there was one (and only one) bank run in the Russian household deposit market. This is consistent with narrative evidence on our sample period. In addition to timing bank runs, the method simultaneously measures the deposit and interest rate movements across different types of banks. Conditional on a bank run, depositors run on solvent banks. They maintain, and even increase deposit holdings at insured banks. This is a clear manifestation of panic. In addition, insolvent banks face even more severe deposit outflows. The fundamentally weak banks try to counter the run by offering higher interest rates on deposits, but fail to convince depositors of maintaining their accounts. Quantitatively, the deposit outflow at insolvent banks is approximately four

in much of the literature is not very high. Even with high-frequency data (as in Donaldson 1992), the presence of other factors cannot be completely ruled out.

times as large as the panic run on solvent banks. Therefore, the bank run identified in our application also offers strong support for the information-based view. However, in terms of the aggregate deposit outflow, panic withdrawals are much more important. Our results attribute almost 90% of the withdrawals to panic, while information-based withdrawals account for about 10% of the total drain in deposits.

The paper is organized as follows. In Section 2, we provide a simple model of the deposit market which articulates our characterization of a bank run. Section 3 documents the details of the method, and explains how to translate the model's intuition to a data environment. In Section 4 we discuss why the Russian deposit market may be of particular use to investigate the properties of the method as well as validate the effects of the different theories of bank runs. Section 5 describes the properties of the data we use in our application. The results, both in terms of bank run identification and its effects are investigated in Section 6. In Section 7, we assess the scope for possible alternative interpretations of the results, discuss some caveats and perform a variety of robustness checks. Section 8 concludes, discusses policy implications and suggests possible avenues for future applications of the method.

2 Bank runs in a simple model economy

We start by presenting a graphical model for the deposit market. The model is purposely very stylized, in order to highlight the few basic mechanisms our method relies on. Consider an economy in which there are two types of banks. We distinguish between insured and uninsured banks. Insured banks reside under a deposit insurance scheme, while no such guarantees exist for the uninsured group of banks. In such a setting, one can think of a bank run as an instant in time where deposits flow away from the uninsured banks, without being driven by the interest rate they pay. Depositors withdraw their money from uninsured banks because neither solvency (ever getting their deposits back) nor liquidity (retrieving their deposits at the desired moment) is guaranteed. The outflow of uninsured deposits may become an inflow of deposits for the insured banks. These banks are, by definition, able to pay back their depositors. However, since liquidity is not necessarily guaranteed, the pool of uninsured and insured deposits may also seek alternative investments (implying an absolute outflow of the deposit market).

Figure 1 provides a simple representation of the market for deposits in equilibrium in this economy (equilibrium $\{A(I), A(U)\}$). Demand equals supply for both groups of banks.[3] Figure 2 shows the equilibrium situation following a bank run in our prototype economy (equilibrium $\{B(I), B(U)\}$). At a minimum, a bank run implies a reduction of the volume of deposits supplied to uninsured banks. This is depicted by the inward shift of the supply of deposits to uninsured banks, for a given demand curve. For the reasons outlined above, the insured banks' supply curve could shift either down or up. What is relevant here is that it will not exceed the reduction that uninsured banks face. For the uninsured banks, the uncertainty with respect to solvency aggravates possible outflows due to liquidity.

In short, bank runs are instances in which the uninsured banks' equilibrium shifts to the northwest of both its initial equilibrium and of the insured banks' equilibrium. Equilibria in

[3] The overlap of both groups' schedules is merely to keep the graphs easy to read. Our argument does not require the schedules for the two groups to overlap, nor that they are of the same shape.

Figure 1: Equilibrium in the deposit market

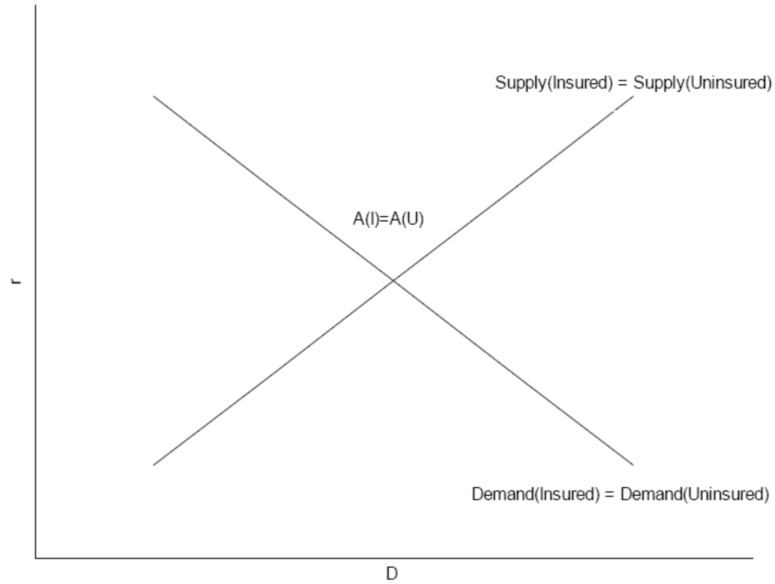

the southwest are ruled out because they are demand driven.

3 The method

Our empirical approach starts with a reduced form model of both insured and uninsured banks' deposits and the respective interest rates they pay. There are a number of reasons why we choose for a flexible reduced form model, rather than a more structural model.

First, the empirical fit of reduced form panel-VARs is substantial for micro data. Panel-VARs were introduced by Chamberlain (1983) and Holtz-Eakin et al. (1988). Second, the majority of structural models have a reduced form representation which is encompassed by this model. So, essentially, nothing is lost by analyzing reduced forms. Third, while maintaining consistency with the variety of structural models, our approach does not need to make strong and highly debatable assumptions regarding the functional form of demand and supply equations.

The particular reduced form model we consider takes the following form:

$$\begin{pmatrix} D(j)_{i,t} \\ r(j)_{i,t} \end{pmatrix} = c^j + A^j * \begin{pmatrix} D(j)_{i,t-1} \\ r(j)_{i,t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon^D(j)_{i,t} \\ \varepsilon^r(j)_{i,t} \end{pmatrix} \tag{1}$$

This is a panel-VAR on (log) deposit quantities D and deposit interest rates r.[4] The index and superscript j refers to different types of banks. For now, $j = I, U$ which denote insured and uninsured banks, respectively. Subscripts i and t denote (group-specific) cross-sectional

[4] The panel-VAR differs from reduced forms typically considered in empirical studies of market discipline, such as Park and Peristiani (1998) or Martinez Peria and Schmukler (2001). These studies typically ignore dynamics and include lagged bank-specific variables instead. The presence of lagged dependent variables in (1) takes up the role of these variables (see, e.g., Wooldridge 2002).

Figure 2: Post-run equilibrium in the deposit market

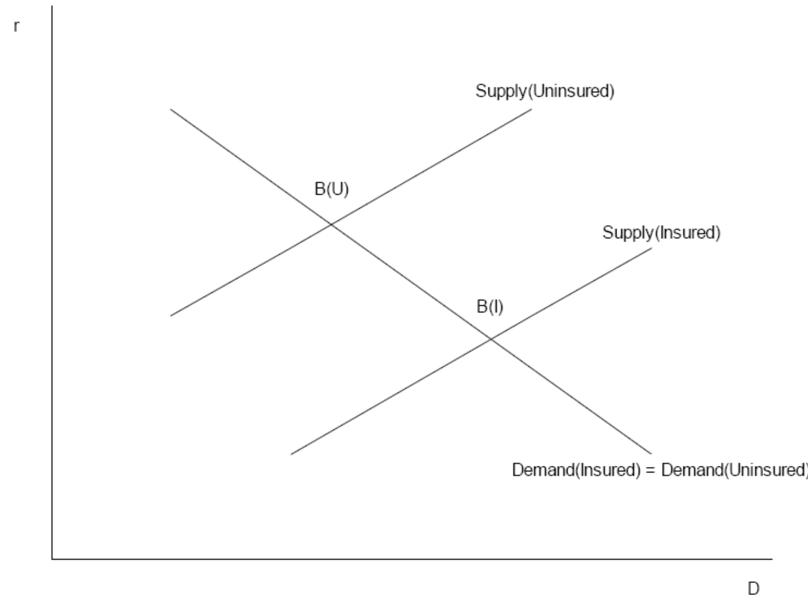

units and time, respectively. A^j is a coefficient matrix and c^j is a vector of constants. Below, we allow different types of banks to exhibit different reduced form coefficients.[5] $\varepsilon^D(j)_{.,t}$ and $\varepsilon^r(j)_{.,t}$ contain reduced form shocks for all banks.

As such, reduced form systems such as (1) tell us little or nothing about the economics in the data. The covariance matrix of the reduced form residuals is non-diagonal and has no structural interpretation. Moreover, there are no contemporaneous interactions between the different endogenous variables. This is what distinguishes such a reduced form from structural models. The movements in D and r observed in the data are an amalgam of all types of shocks affecting demand and supply in the deposit market. The aim of our approach is to filter out one particular shock of interest, *viz.* a bank run.

Put differently, what we are interested in are the effects of a structural shock, as described in the model economy. We therefore put additional structure on the reduced form. To do that, we make use of identification techniques which are frequently adopted in macroeconomic analyses using VARs. Identification in macro VARs typically imposes restrictions on different variables over time. In our approach, we also make use of an additional, cross-sectional dimension to obtain structural identification. That is, we distinguish between different types of banks, and contrast differences in the groups' behavior in the identification process. We will show later that this additional dimension is informative both in terms of shock identification, as well as in terms of hypothesis testing.

The most practical way to understand our approach is to apply it to our model economy. We identify a bank run as a structural shock which contemporaneously lowers the deposits

[5]To keep the presentation simple, we write the system with only one lag for all variables and without additional control variables. These do not change anything conceptually. The actual estimation typically uses more lags. The empirical implementation also allows for incorporating additional control variables.

of uninsured banks:[6]

$$\Delta D(U)_t < 0.$$

In order to make sure that we are analyzing a bank run, rather than, for instance, a general drain on deposits, we impose that the drain of insured deposits is smaller than the one for uninsured deposits (recall the absence of the solvency concern for insured deposits). We also leave open the possibility that the uninsured outflow gets deposited at the insured banks. This amounts to the additional restriction:

$$\Delta D(I)_t \geqslant \Delta D(U)_t.$$

So far, the restrictions only pertain to the volume of deposits. However, the above set of restrictions is not only consistent with a bank run, the relative deposit outflow could also result from demand side pressures. For instance, the observation that $\Delta D(I)_t \geqslant \Delta D(U)_t$ and $\Delta D(U)_t < 0$ may result from the fact that insured banks raise their deposit rate, relative to the uninsured banks. To exclude such cases, we additionally impose that

$$\Delta r(U)_t \geqslant \Delta r(I)_t.$$

These restrictions suffice to disentangle a bank run from other structural shocks in the deposit market. In terms of Figure 2, the restrictions guarantee that after a bank run, the uninsured banks' equilibrium outcome will lie to the northwest of the insured banks' equilibrium.[7]

Computationally, the approach consists of a search for matrices S, which are orthogonal decompositions of the variance-covariance matrix of the reduced form group-wise average residuals $[\ \bar{\varepsilon}^D(U)_t;\ \bar{\varepsilon}^r(U)_t;\ \bar{\varepsilon}^D(I)_t;\ \bar{\varepsilon}^r(I)_t\ ]$, and which satisfy the restrictions above conditional on the run. When premultiplying (1) with such a matrix, it partially identifies the structural form. That is, the reduced form shocks are then transformed into a set of structural shocks, i.c. bank runs. They also identify the simultaneous structural relations between the endogenous variables in the system. In addition to searching among the many possible rotations of the shock variance-covariance matrix, coefficient uncertainty of the estimated reduced form is also taken into account. The procedure results in impulse response functions for all variables in the system after a bank run, as well as a time series of the structural shock, and associated confidence bands.[8]

4 A real world analogue for the model economy

We consider the Russian economy to be a useful case study to apply our approach to. The reason is twofold. First, there is cross-sectional variation among banks in the degree to which

[6]Here, Δ denotes changes relative to baseline, where the latter is measured by the dynamics of the system in the absence of the shock.

[7]The additional restriction that $\Delta r(U)_t > 0$ is necessary to completely rule out demand shocks: it ensures that uninsured banks' new equilibrium lies to the northwest of their own initial equilibrium. In our application, we find that this restriction never binds.

[8]For details on implementation within a macroeconomic framework, see e.g. Uhlig (2005).

their household deposits are guaranteed by the government. Our method will exploit such heterogeneity in identifying bank runs. Second, there is some prima facie evidence that there has been (at least one) bank run over the period we have data for. On the one hand, this allows us to evaluate the effects of bank runs. On the other hand, we will confront the timing of runs identified by the method to evidence extraneous to the model.

4.1 The existence of insured and uninsured banks

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.[9] In addition, 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, an implicit guarantee that is little affected by recent legal changes (Tompson, 2004). Favorable state policy with respect to state-owned banks was particularly evident during the summer of 2004, when the Russian central bank (CBR), having refused to extend a stabilization credit to Guta-Bank, which had come under pressure, instead lent money to state-owned Vneshtorgbank, enabling it to buy Guta and thus to nationalize one of the country's larger private banks (Tompson, 2004).

4.2 The occurrence of bank runs

In May 2004 the CBR closed a bank accused of money laundering while the Federal Service for Financial Monitoring (FSFM) announced it suspected about a dozen banks in money laundering and sponsorship of terrorism, without naming the "dirty dozen". Several inconsistent black lists began circulating the banking community as bankers 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 of confidence provoked runs on several large banks among which were Guta Bank and Alfa Bank. Thus, there is narrative evidence suggestive of (at least one) bank runs occurring in our sample period.

5 The data

5.1 Source

The bank-specific variables used in this paper include deposits and interest rates as well as measures of risk, performance and balance sheet structure. The data were made available to

[9] Federal Law No. 182-FZ "On the introduction of changes and amendments to the Civil Code of the Russian Federation" removed the state's subsidiary liability for the retail deposits of banks in which the Russian Federation or subjects of the federation held majority stakes (i.e. it ended the 100 per cent state guarantee hitherto enjoyed by such banks under Civil Code art. 840.1).

Figure 3: Average growth rates of consumer deposits at insured and uninsured banks

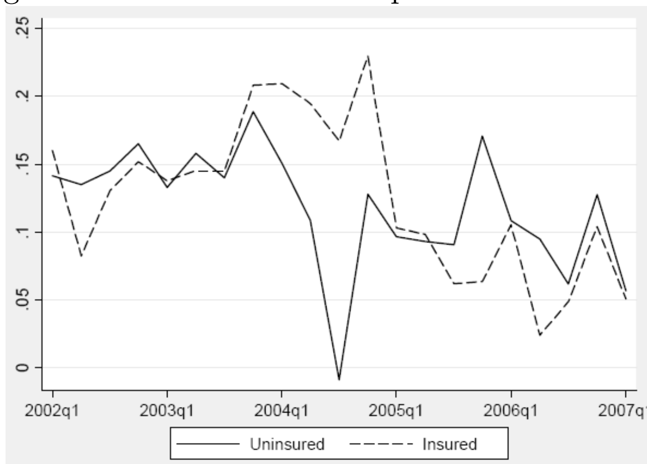

the authors by three established and highly respected private financial information agencies, Interfax, Banksrate.ru and Mobile.[10]

Interfax provides quarterly bank balances and income statements from 1999 through 2007. The average implicit interest rate that a bank offers on its deposits is calculated by dividing interest expenses by the corresponding level of deposits.[11] Since our dataset disaggregates both interest expenses and deposits by the legal status of the depositor, the variables measuring deposit flows and interest rates are computed separately for household deposits. The constructed interest rate series exhibit a break in 2001, presumably due to changes in variable definitions. We limit our sample to observations after the break. 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 "new" from the standpoint of our sample. Lists of banks with the state as a majority owner are available at two points in time, February 2002 (Matovnikov, 2002) and July 2005 (Mamontov, 2005). These lists reveal that the state ownership category remains rather stable over our sample period.

5.2 Evolution

Figure 3 shows that the average growth rates of consumer deposits in both insured and uninsured banks are comparable through the major part of our sample period. The only substantial divergence takes place during the turbulent summer of 2004 with the uninsured category experiencing a much lower (and even negative) average deposit growth.

[10] For more information on the data providers see their respective websites at www.interfax.ru, www.banksrate.ru and www.mobile.ru. Karas and Schoors (2005) provide a detailed description of the datasets and confirm the consistency of different data sources.

[11] By using an implicit rate, there may arise a concern that its fluctuations are mainly driven by the fluctuations in the quantity variable in the denominator, thereby generating spurious movements in our interest rate variable. All our results carry through, however, if we divide the interest rate expenses by the bank-specific average quantity of deposits.

Figure 4: Average interest rates at insured and uninsured banks

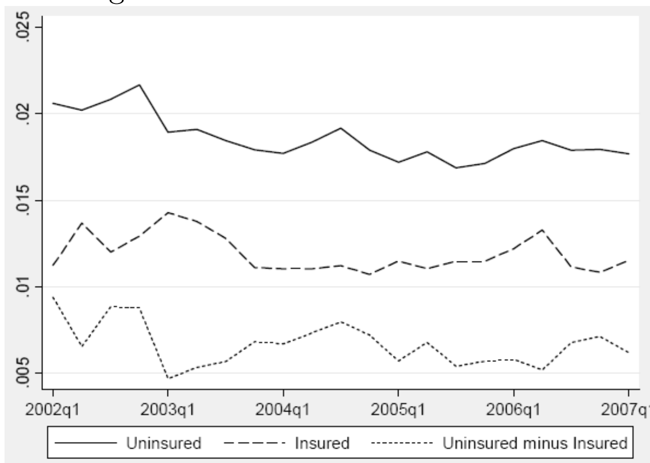

As expected, uninsured deposits on average pay higher interest. As depicted in Figure 4 the average spread on a quarterly basis decreased from around 100 basis points in the beginning of the sample period to around 60 basis points at the end. The spread was notably higher through 2004 compared to both 2003 and 2005.

6 Application

This section provides the results from applying our method to the Russian household deposit market data. We start by discussing the setup of the empirical model. In particular, we detail the types of heterogeneity the method uses to identify bank runs, and how it allows us to discriminate among the different theories. We then present the bank runs identified by the model. We contrast these episodes to narrative evidence as a form of external validation. Subsequently, we turn to testing the relative importance of the different views on bank runs.

6.1 Setup

First, in view of what follows, it is useful to explicitly distinguish between two concepts: runs and failures. While at some point we will also look at the causes of bank failures, our prior interest lies in uncovering what depositors' motives are to run on a bank. A run on a bank is not quite the same as, and may have different causes than, a bank's default. For instance, depositors can run on a bank without that bank having to declare failure. Therefore, below, we shall take "run" to refer to depositor withdrawals (under the additional conditions specified before). "Failure", by contrast, will be understood to reflect the fact that a bank ceases to exist.[12]

Second, the way in which one distinguishes between fundamentally sound and weak banks is important in assessing whether fundamental information is used by depositors. Calomiris

[12]Failures can be a long administrative process. Our dating of failure measures the starting date of that process. This can pertain to, e.g., the license revokal or the transfer of control to a judge.

and Mason (2003) perform a test of panic effects during the Great Depression. In particular, they estimate default prediction models based on a wide variety of fundamental factors. Subsequently, they ask how much unexplained variation there was during certain narratively identified crisis episodes. As Calomiris and Mason (2003) describe, that estimate provides an upper bound for the importance of panic effects. Such a test is one-sided because it can exclude, but not confirm the presence of panic in the data. In other words, it is not because the econometrician is not able to project all defaults based on observable fundamentals, that this proves that there are no fundamentals able to capture the remaining unexplained variation in failures. Perhaps the test misses one (or more) crucial fundamental that explains all observations.

In order to remedy this feature, we avoid the two inherent restrictions of the Calomiris and Mason (2003) approach: data availability and the ex ante point of view. Rather than specifying and estimating a default model to try to replicate agents' information set during the crisis, we -the econometrician- are perfectly informed over ex post outcomes. That is, we test whether outcomes differ for banks that fail relative to those that survive. Distinguishing between actually failed and non-failed banks provides a way to control for all fundamentals (both observed and unobserved), and investigate *whether* and *how much* information matters. The contribution of using actually observed data on fundamentals, as in Calomiris and Mason (2003), is to verify *which* information matters. We therefore split the group of uninsured banks according to whether they actually fail or not. This cross-sectional split is also used in Saunders and Wilson (1996). The rationale for this assumption is as follows. Isolating those banks that did fail following the crisis is tantamount to assuming the maximum amount of information. Even if there is a variable that we could not consider that is able to predict the bank failures, our sample split allows for the possibility that depositors had it in their information set.[13] As a result, at every point in time, the "low" group of uninsured banks consists of banks that fail in the year and a half subsequent to that observation.[14]

We thus consider three groups of banks: insured banks (indexed by I), uninsured banks who fail ex-post (U, low) and uninsured surviving banks $(U, high)$.[15] Given this cross-sectional split, we estimate a three-group reduced form system and impose the restrictions on the insured group relative to the failed, non-insured group of banks.[16] Following the logic of the model economy, a bank run is identified as a shock which implies a deposit outflow at uninsured failed banks, relative to insured banks. The restrictions on the relative interest rates control for the possibility of demand shocks. All these restrictions are imposed for only one period -contemporaneously-, which is sufficient to achieve identification. This is less than typically needed in macroeconomic VARs and exemplifies the informational content of the cross-sectional restrictions.

[13] While splitting the banks based on observed defaults controls for solvency as the reason for the run, there is still a potential role for the liquidity argument. However, this likely plays only a limited role relative to solvency issues (also see Calomiris and Mason (2003), who focus on solvability alone). Moreover, the lack of liquidity possibly applies to all banks, even the insured ones.

[14] Small variations in the time span during which the bank fails (between one and two years following the observation period) do not alter the conclusions. Too small a window makes estimation uncertainty dominant. Expanding the failure window to the fullest, such that, e.g., a bank that fails in 2007 is already included in the low group in 2003, does not change our qualitative conclusions.

[15] There were no state banks that failed during our sample period.

[16] Thus, the reduced form is -maintaining the one lag assumption for brevity-:

Once the model is identified, the hypothesis tests verify whether there exists a different response for the uninsured banks that did not fail, relative to those that did.[17] Throughout, the uninsured surviving group's behavior is completely unrestricted. This provides an ideal framework for hypothesis testing, as we document below. First, we describe (and validate) the historical episodes that the method identifies as bank runs.

6.2 Identifying bank runs

6.2.1 A time series of the identified shock

Figure 5 plots a confidence interval for the bank run shock over our sample period. One should be careful not to overinterpret such a series, due to the inherent presence of noise. To eradicate part of the noise, the figure plots a two-period moving average of the shock.[18] The single largest shock is observed during the summer of 2004. 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).[19]

6.2.2 External validation

The previous section shows how our approach identifies the 2004 summer (and essentially no other period) as a bank run. This section is devoted to corroborating that finding with information outside of the model. This type of external validation resolves a number of issues some have raised as a criticism to the use of structural VARs. First, one critique of identification in (S)VARs is the apparent arbitrariness of the structural shocks. Rudebusch

---

$$\begin{pmatrix} D(I)_{i,t} \\ r(I)_{i,t} \\ D(U,low)_{k,t} \\ r(U,low)_{k,t} \\ D(U,high)_{l,t} \\ r(U,high)_{l,t} \end{pmatrix} = \begin{pmatrix} c^{D(I)} \\ c^{r(I)} \\ c^{D(U,low)} \\ c^{r(U,low)} \\ c^{D(U,high)} \\ c^{r(U,high)} \end{pmatrix} + A * \begin{pmatrix} D(I)_{i,t-1} \\ r(I)_{i,t-1} \\ D(U,low)_{k,t-1} \\ r(U,low)_{k,t-1} \\ D(U,high)_{l,t-1} \\ r(U,high)_{l,t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon_{i,t}^{D(I)} \\ \varepsilon_{i,t}^{r(I)} \\ \varepsilon_{k,t}^{D(U,low)} \\ \varepsilon_{k,t}^{r(U,low)} \\ \varepsilon_{l,t}^{D(U,high)} \\ \varepsilon_{l,t}^{r(U,high)} \end{pmatrix}$$

where i, k and l index cross-sectional units per group of banks. In a traditional panel-VAR (e.g. Holtz-Eakin et al. 1988), cross-group elasticities in A (i.e. elements off the block diagonal) are zero. In the structural form, by contrast, cross-elasticities need not be zero.

[17] Also note that in the current setting, contrary to the narrative approach, the crisis date is endogenously and simultaneously estimated.

[18] Calculation of the shocks and impulse response functions requires a normalization to be performed. Therefore, at least one of the restrictions needs to incorporate an actual sign restriction, in addition to the relative restrictions. Here, a positive bank run shock corresponds to one in which uninsured deposits at a posteriori failed banks fall.

[19] Note that the level of the shock seems to be biased below zero. This results from the fact that the run itself is so large. Put differently, if there were no other fluctuations but the 2004 run (and assuming a balanced number of observations in each group of banks), the shock series would be constant below zero, except for the period of the run, which would make that average zero again. This also implies that the cases in which the shock appears significantly negative in Figure 5 (2003:Q3 and 2006:Q4) do not necessarily have economic significance. Rather, they do not appear too different from the other negative instances. The same clearly does not hold for the 2004 run.

Figure 5: Estimated bank runs

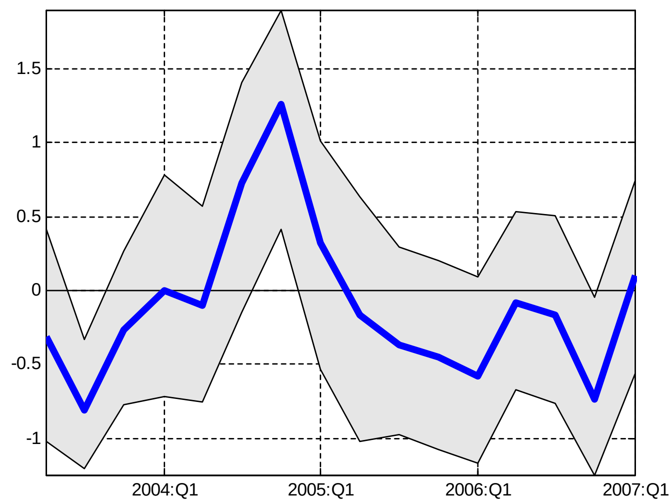

(1998), for instance, shows how the monetary policy shocks identified through different VARs are largely unrelated (whereas they are supposed to measure the same thing), both among themselves and when compared to alternative measures of monetary policy shocks. Second, an issue with interpreting shocks from partially identified models is that the significance of the shock may be due to the simultaneous occurrence of several other shocks -which the model does not explicitly identify. The validation we propose compares our results to information extraneous to the model.

We perform this external validation in various ways. First, we skim the literature trying to find evidence for bank runs in our sample period. One source we have on this issue is Tompson (2004). The latter describes the run-like features of the Russian deposit market during the summer of 2004. A second measure of external validation is based on press coverage. We performed an -a priori specified- computerized search in the article databases of The Economist and the NY Times for our sample period using the terms “Russia”, “deposit” and “run”. Out of all hits, three directly pertained to the present paper’s subject. All three were dated summer of 2004 and each of them suggested the possibility of a bank run.[20] We interpret this to be evidence for the fact that 1) a run was very likely in the 2004 summer, and 2) there were no other episodes in our sample period suggestive of bank runs.

Thus, the run identified by the model is in accordance with outside information. Hence, the results of our external validation exercise strongly suggest that neither of the problems particular to identified VARs apply in our test. Thus, this type of external validation provides additional confidence with respect to the validity of our approach.

[20]The articles referred to are “There’s always Sberbank” (Economist, 7/10/2004), “Don’t run for it” (Economist, 6/26/2004), and “Depositors’ jitters increasing as some Russian banks close” (NY Times, 7/9/2004).

Table 1: Hypothesis tests

	Identifying restrictions
Bank run	$\begin{cases} \Delta D(U,low)_t < 0 \\ \Delta D(U,low)_t \leqslant \Delta D(I)_t \\ \Delta r(U,low)_t \geqslant \Delta r(I)_t \end{cases}$
Hypothesis	Test
$H1:$ Information-based	$\Delta D(U,low) < \Delta D(U,high)$
$H2:$ Panic	$\Delta D(U,low) \geqslant \Delta D(U,high)$
$H3:$ Panic	$\begin{cases} \Delta D(U,high) < 0 \\ \Delta D(U,high) < \Delta D(I) \end{cases}$

6.3 Testing the theory: Information-based or panic runs?

6.3.1 Hypothesis tests

With the run identified, we can now address depositors' motives for running on particular banks. The behavior of depositors at surviving uninsured banks allows to quantify the importance of the information-based and panic views. Recall that the interest rate and deposit response of the surviving group is completely unrestricted. In terms of deposits, one expects the high (surviving) group's response to lie between one of two extremes. On the one hand, if bank runs are characterized by depositor panic, then the drain of deposits at the failed and surviving uninsured banks should be very similar. This would corroborate Diamond-Dybvig type panic runs. On the other hand, in case runs are information-based as suggested by Allen and Gale (1998), the surviving banks would exhibit deposit flows similar to those of the insured group of banks. Because we consider a multivariate system, however, the results are richer. For instance, one may observe that the uninsured banks increase their interest rates strongly in an attempt to prevent the deposits from draining. Alternatively, it may be the case that the high group is able to withhold the run, mainly by promising higher interest rates. In short, a wide variety of responses is possible.

In Table 1 we provide a summary of the hypotheses of interest and the tests that corroborate them. Note that these tests are conditional on a bank run, the conditions for which are restated in the table.[21]

A first hypothesis test determines the effect of fundamentals in a bank run. A necessary and sufficient condition to conclude that bank runs are information based is that deposit outflows at fundamentally weak banks are larger than those at solvent banks ($H1$).

There is more than one way in which panic effects can occur. A first possible manifestation of panic ($H2$) would be that the deposit outflow is the same or even larger at solvent banks relative to the outflow at failed banks. Second, if depositors withdraw their funds at solvent banks ($\Delta D(U,high) < 0$), this may (but need not) be driven by panic, too. Either this

[21] Observe that we drop the time subscripts in the hypthesis tests. The reason is that the various effects could occur at any point in time. They could be instantaneous, lagged, turn out to be persistent, or even change over time.

outflow is larger than the response of the insured deposits, which then implies panic effects are at work. This is summarized in hypothesis $H3$.

Alternative hypotheses can be constructed but are not necessarily conclusive. For instance, when the outflow is the same at solvent banks and insured banks ($\Delta D(U, high) = \Delta D(I)$) this could suggests there is no panic. This will be so under the additional condition that $\Delta D(U, high) \geqslant 0$. Alternatively, $\Delta D(U, high) = \Delta D(I)$ could be accompanied by a general outflow out of the deposit market ($\Delta D(U, high) = \Delta D(I) < 0$). In this case, however, more than one interpretation is possible. First, this can be due to a general outflow of deposits, without there being panic, as mentioned before. Second, it could mean that there is panic (causing $\Delta D(U, high) < 0$) and at the same time insurance is not fully credible (causing $\Delta D(I) < 0$). A third possible interpretation is that there occurs a fully random panic run in which depositors withdraw without distinguishing among insured and uninsured banks.

As a result, the list of hypotheses posited in Table 1 is not exhaustive. We restrict attention to those tests that are important to validate the different theories. Note that not all these hypotheses are necessarily mutually exclusive. More than one of the hypothesized effects can be at work at the same time. The impulse responses provided below will test which of these effects are at work, and quantify them.

6.3.2 Results

Figure 6 plots the impulse responses to a bank run for the three bank-types' deposits and interest rates. The graph shows the effect of a one standard deviation bank run. Dotted lines indicate confidence bands. To enable hypothesis testing, Figure 7 provides confidence bands for inter-group differences.

First note the restrictions that have been imposed in identifying these impulse responses (Figure 6, columns 1 and 2). Failed, non-insured banks experience a drain in deposits. Insured banks, by contrast, do not suffer as large a drain. This relative restriction is one-sided, requiring that insured banks do not face a worse deposit outflow. In fact, insured banks may even experience a deposit inflow. The response of the insured group's deposits shows that this is the case, although the significance is marginal. Contrary to the interest rate of failed uninsured banks, the insured banks' deposit rate does not show any significant movement.

Let us now focus on the group of uninsured banks that did not fail (Figure 6, column 3). Apparently, the deposits of these banks fall significantly, while their interest rates remain unaffected. Importantly, does the response of these uninsured banks resemble that of the failed uninsured banks? In Figure 7, the error bands on the difference between the failed and surviving uninsured banks show that failed banks experience a significantly larger outflow of deposits than those banks that did not fail. Note that this occurs despite the fact that the failing banks raise their interest rates, also in relative terms. The impulse responses in Figures 6 and 7 plot the effect of a one standard deviation innovation in the shock. Taking into account the actual value of the shock in the 2004 crisis, the identified bank run in the summer of 2004 accounts for a deposit drain at the uninsured banks that fail ex-post of about 15% and an interest rate rise of around 30 basis points.[22]

[22] The 2004 bank run is almost 2.5 standard deviations away from its mean. So to an approximation, one

Figure 6: Impulse responses to a bank run across bank types

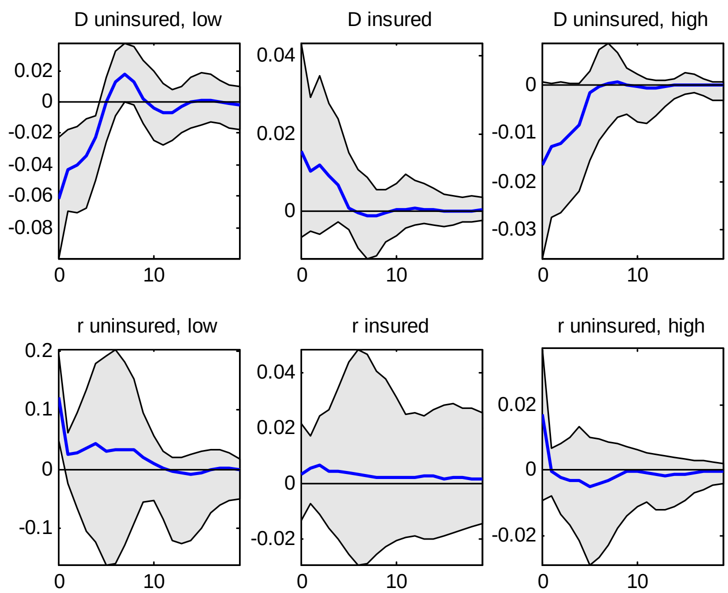

Figure 7: Impulse responses to a bank run across bank types: group-differences

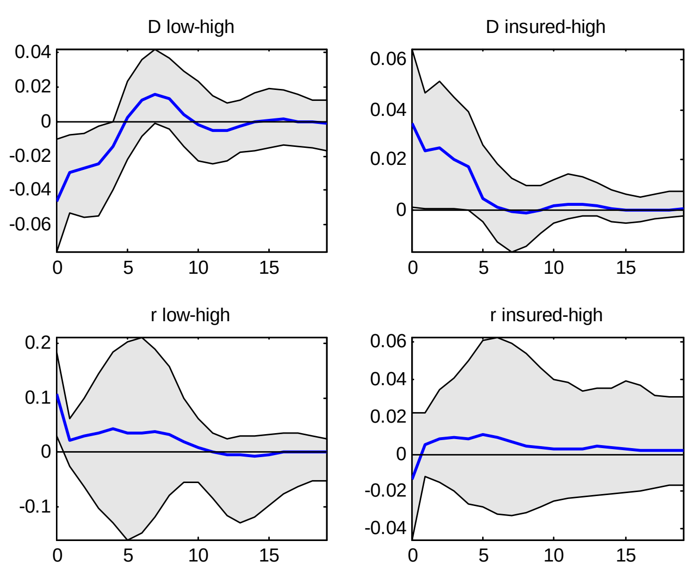

Based on these results, we conclude that $H1$ holds. Bank runs are, to a degree, information based. Solvent banks face less significant deposit outflows, along the lines of Allen and Gale (1998), among others. The quantitative importance of this effect is large: the deposit outflow at the failed (low) banks is four times (!) as large as that of the solvent banks.

Next, we consider the scope for panic effects. Since $H1$ and $H2$ are mutually exclusive, and we already established that $H1$ holds, $H2$ cannot hold. Deposit outflows at solvent banks are less severe compared to failed banks. However, there are other ways in which panic can manifest itself. Consider $H3$, for instance. It clearly holds that $\Delta D(U, high) < \Delta D(I)$, and -though borderline significant- it also holds that $\Delta D(U, high) < 0$. As a result, to the extent that the latter finding is statistically significant, there are panic effects, too. Depositors run on banks which did not fail. In economic terms, the median response *is* significant: uninsured solvent banks experience a deposit outflow of 1.6%, which translates to 4% conditional on the 2004 crisis.

There is a subtle difference in interpretation with the results of Calomiris and Mason (2003). In their work, the unexplained variation is an upper bound for the degree of panic. In particular, they attribute non-forecasted failures to either panic or fundamentals that they do not have data on. Our results are not prone to such a one-sided interpretation. Specifically, observing a run on solvent banks necessarily follows from depositors inability to forecast failures correctly. By contrast, in Calomiris and Mason (2003) the forecasting ability that is investigated is that of the econometrician.

Finally, observe that there is not much action in terms of interest rates. Only the uninsured failed banks raise their deposit rates, and they do so at the time the shock hits. The hike in their deposit rate is significant, yet still deposits flow away from these banks. Such effects strongly resemble the mechanisms at work in various theories of asymmetric information. For instance, the hike in uninsured failed banks' deposit interest rates in the wake of the deposit outflow could be interpreted as "gambling for resurrection". Hellman, Murdock and Stiglitz (2000), for instance, show how moral hazard at troubled banks may create the incentive to engage in this type of risky behavior. Alternatively, the increase in the deposit rate may well reveal the bank as being a "lemon", and thereby cause or worsen the drop in the volume of deposits. Such interpretations are supported by the finding that the interest rate at the insured solvent banks does not rise, while their deposit outflows are less severe.

6.3.3 The relative and aggregate importance of panic and information

The impulse responses in Figure 6 plot the effect of a bank run on the average insured, uninsured failing and uninsured surviving bank. They establish, among other things, that the run on insolvent banks is much more severe at insolvent banks and thereby support the relevance of the information-based view on bank runs. Quantitatively, this suggests that 75% of the run on insolvent banks is information driven, and 25% is due to panic. This does not necessarily imply, however, that the information-based view is more relevant on aggregate. In order to evaluate the aggregate importance of both views, one needs to take into account the number of banks in each group. At the time of the 2004 bank run, 4.5%

can multiply the impulse responses by that factor to obtain the imputed effects conditional on the 2004 bank run. Exact inference is also conditional on the size of the shock. When we re-run the model using the 2004 shock size, the above approximation is very good and inference is not affected.

of the uninsured banks are classified as failing ex post (47 banks). As a result, though the information-based effects are quantitatively strong, they apply to a smaller fraction of banks. Taking into account these fractions suggests that, on aggregate, panic caused about 88% of the total outflow of uninsured deposits, while information effects were responsible for about 12% of the total deposit withdrawals. This implies that while both information and panic are significant on aggregate, panic is the most important driving factor behind the run in our application. Finally, of the entire deposit outflow that the 2004 bank run caused, only about 8% was deposited at the insured banks. The remaining 92% was drawn out of the deposit market.

7 Alternative interpretations, caveats and robustness

7.1 Information structure

7.1.1 Uninformed depositors?

First, note that in the general approach to identify bank runs, it is not the case that we assume that depositors are completely uninformed, as could be the case in a full-blown panic. Our approach requires depositors to know whether or not their deposits are secured. We view this as a very minimalist informational assumption. It is straightforward to see that this type of information is from an entirely different nature than being able to judge the health of a bank or its balance sheet. Moreover, if this information were not known to depositors, it is puzzling why the level of interest rates at state banks is consistently below that of the other banks (or why the summer of 2004 is identified as a crisis).

Second, one could argue that full-blown panic runs (where depositors have absolutely no idea about which banks are fundamentally sound) are not fully allowed for when the restrictions are imposed on the uninsured failing group of banks relative to the insured group. In other words, complete randomness implies a positive probability for runs that are local to the fundamentally sound banks. We therefore also perform the analysis by imposing restrictions on the solvent uninsured group of banks relative to the insured. The results of this exercise indicate that there is no run identified, except for the one in 2004 (with similar relative impulse responses). As a result, though a theoretical possibility, a panic run in which depositors only withdraw their money at solvent banks and not at the insolvent ones did not occur in our sample.

7.1.2 Deposit outflows as the cause of bank failures

Another caveat is the relation between deposit withdrawals and solvency. In particular, the results of Section 6 could look very similar if all banks were equally solvent in the high and low groups, but where withdrawals are the only source of failure. This would imply there is no informational difference between the two groups that can discriminate good from fundamentally weak banks. In order to verify whether such is the case, we perform a set of default prediction (logit) regressions.

Table 2 summarizes the results. In addition to a standard set of bank fundamentals, deposit growth does help forecast future defaults. However, deposit growth contributes only

Table 2: Default prediction and household deposit growth

	2002-2007	2002-2007
Ln(Assets)	-0.17**	-0.16**
	(0.07)	(0.07)
Capital / Assets	-1.58*	-1.81**
	(0.82)	(0.80)
ROA	-14.85***	-13.16***
	(3.48)	(3.04)
Liquid Assets / Assets	-8.93***	-8.87***
	(2.33)	(2.21)
Bad Loans / Assets	6.07**	5.77***
	(2.47)	(2.22)
Non-government Securities / Assets	3.07***	2.99***
	(0.63)	(0.62)
Term Deposits of Firms / Assets	-6.39***	-6.23***
	(2.30)	(2.26)
Term Deposits of Households / Assets	-6.00***	-5.46***
	(2.13)	(2.02)
Household Deposit Growth		-0.47***
		(0.11)
Observations	22155	22155
Pseudo-R^2	0.27	0.29
AUR	0.864	0.871

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Time dummies not reported.

marginally to the overall explanatory power of the model in terms of pseudo-R^2. Importantly, there is hardly any improvement in the cross-sectional predictive power of the model, as can be judged from comparing the AUR of the two models.[23] Hence, while some essentially healthy banks may have failed due to deposit outflows, there are far too little of them to be able to affect our results meaningfully.

In addition to the results of the logit regressions, there is a timing difference which further reduces the concern of runs being the cause of default. In particular, if the method identifies a run at failed banks at date t, this pertains to failures that occur after date t (hence ex-post failures). Therefore, deposit outflows that cause *contemporaneous* failures are no cause for concern. Deposit outflows that give rise to *subsequent* defaults, too, are of little importance, as the logit results suggested.[24]

7.1.3 A real-time information set

We here use a depositor information set that is different from the full information approach used earlier, where we used actual failures to identify fundamentally flawed banks. In particular, we now use a perspective analogous to Calomiris and Mason (2003). More specifically, we ask how our results would change if we -the econometrician- were to forecast, in real time, whether or not a bank is to fail in the next period. The results in Figures 8 through 10 show the estimated time series of runs as well as the impulse response functions. Here, at each date, the low uninsured group consists of those banks for which the logit of Table 2 predicts a failure with probability above the time-specific median.

Interestingly, in Figure 8, the summer of 2004 is still identified as a crisis. In addition, following the bank run of 2004, Figure 8 also suggests there was a negative shock in late 2005 - early 2006. We provide a possible interpretation for that finding in the next paragraph. Let us here focus on the impulse responses of Figures 9 and 10. Qualitatively, the effects are very similar to the baseline results of Figures 6 and 7. Quantitatively, however, we observe substantial differences. In terms of the median response, the difference in responses between the low and high groups are roughly half of what they are in the baseline result with full information. A comparison of Figures 6 and 9 reveals that both components of the difference change. The deposit response of the actually failed banks is more negative than the response of the group of banks predicted to fail. The reverse pattern is observed for the solvent group of banks: a lower outflow at actually surviving banks, and a more pronounced outflow at banks predicted to survive. These results suggest the real time information set is indeed incomplete, missing a variable that succesfully forecasts default. Relative to the results in Figure 6, the responses of the high and the low group are closer to one another because of banks which did fail, but were predicted not to fail (creating a stronger outflow at the high group) and due to banks that survive, but were predicted to fail (inducing a less severe outflow at the low group of banks).

[23] The 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 succesful (see, e.g., Hosmer and Lemeshow 2000).

[24] In order to further reduce this concern, we re-run the model for a group-split in which the low group at date t contains banks that fail after *t+1*. All conclusions remain valid for these alternative results.

Figure 8: Estimated bank runs, forecast-based groups

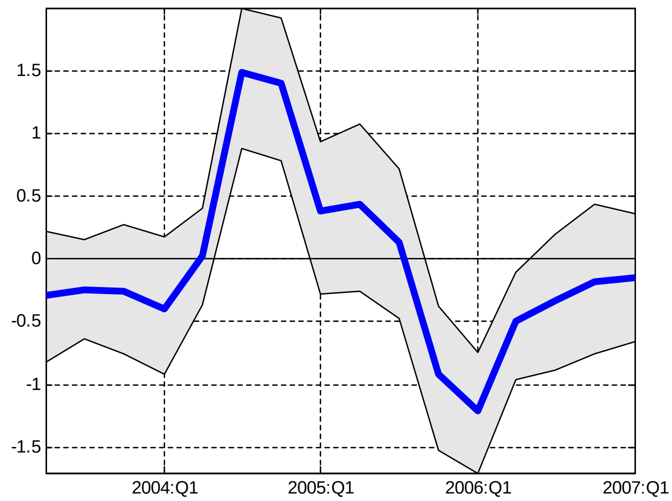

Figure 9: Impulse responses to a bank run across forecast-based groups of banks

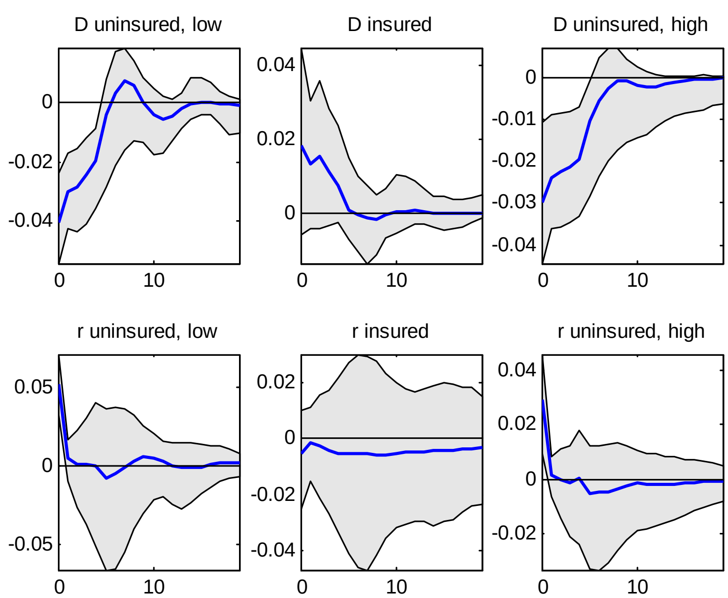

Figure 10: Impulse responses to a bank run across forecast-based groups of banks: group-differences

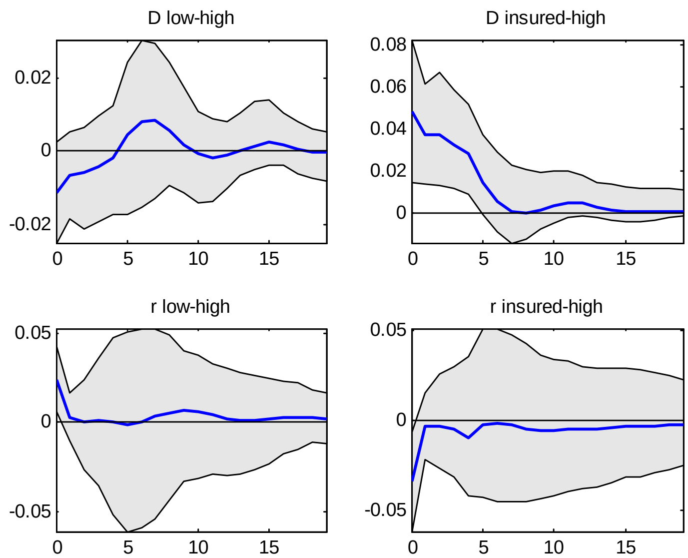

7.2 Credibility of deposit insurance

One could think of the deposit insurance mechanism that we describe as a non-credible one. Non-credible deposit insurance is not necessarily problematic for our method. On the one hand, in case deposit insurance is only partially credible, one would still expect the deposit outflow at the insured banks to be less harsh than that of the uninsured banks, or at least the failed ones. On the other hand, if deposit insurance were not credible at all, we should find that the summer of 2004 was not a bank run, or -in the extreme case where the run is on the insured banks- a negative shock. Neither seems to be the case. Additional circumstantial evidence supporting the credibility of the insurance mechanism is (the title of) an Economist article during the 2004 crisis: "There's always Sberbank", Sberbank being the largest state-owned bank, and thus having insured deposits.

Related to this issue, if there were fluctuations in credibility over time, then our method would classify such dynamics as shocks. In other words, if there were a boost in the credibility of deposit insurance, increasing the relative supply of funds to insured banks, this would end up being classified as a bank run. In some sense, one can interpret the series of structural shocks in Figure 5 as "bank preference shocks".[25]

[25] In the period Sep 2004 - Sep 2005 about three quarters of Russian banks entered the newly adopted deposit insurance scheme. This development could potentially diminish the difference in the perceived safety of deposits in state-owned versus private banks, increasing the relative supply of funds to private banks. Our method would classify such a case as a negative "bank preference shock". This interpretation is consistent with the results in Figure 7.

7.3 Identification is evident?

From Figures 3 and 4 the identification of 2004 as a crisis episode may seem evident. As a result, the entire approach may seem too involving to begin with. There are a number of reasons why this logic does not apply. First, even if the raw data may suggest the summer of 2004 is the only one in which a crisis occurred, this is not quite the same as assuming that this is the only thing that happened during that time. Especially in lower frequency data, event-study-type of assumptions which attribute all movements in that particular episode to the run alone, are particularly hard to defend. Our method does not need to make such an assumption, and allows other shocks to have hit banking markets during that time, as well as during any other time period. Second, and conversely, the approach also allows for bank runs to have occurred, yet for them not to be directly visible in data aggregates. Although the results indicate that such a run did not occur, one can not exclude this a priori.

7.4 Robustness checks

Data construction The baseline results measure the interest rate by an implicit measure, calculated as the interest rate expenses on households deposit accounts relative to the volume in those accounts in the corresponding period. As a result, there may arise a concern that interest rate variations are mainly driven by the fluctuations in the quantity variable in the denominator, thereby generating spurious movements in our interest rate variable. All our results carry through, however, if we divide the interest rate expenses by the bank-specific average quantity of deposits. Importantly, the increase in the interest rate of failing banks -where the effect of using implicit interest rates may affect our results the most- is still observed.

Estimation method Whether the reduced form is estimated using Ordinary Least Squares, a fixed effects or a General Method of Moments estimator has virtually no effects on our identified shocks or impulse response functions. The fact that the data have a substantial time dimension in addition to the cross-section is probably one of the factors contributing to such stability.

Macroeconomic developments The baseline results are based on estimation of equation (1), with four quarterly lags. One might envisage that the macroeconomic environment also plays a role in the evolution of interest rates and deposit flows. The results shown are robust to a variety of ways to control for such evolutions. For instance, incorporating time dummies leaves all conclusions unaffected.[26]

Foreign banks One issue we have not addressed yet is the presence of foreign banks. In the baseline results, these are contained in the insured group of banks. One can think of a couple of reasons to do so. The most important one is, in our view, that while foreign banks are not backed by the state, it is highly unlikely that the mother organisations in the home

[26] The press coverage cited before provides some anecdotal evidence for the macroeconomic environment playing a rather limited role, with one of the articles going as far to suggest that it must have taken the authorities substantial effort to cause a crisis in such a favorable economic environment (NY Times 2004).

country will allow their foreign subsidiaries to fail. The main results continue to hold both when we drop the foreign banks out of the analysis altogether and when we lump the foreign banks together with the high uninsured group of banks. The former is most likely due to the fact that foreign banks exhibit a similar response as that of the state-banks. The latter presumably follows from the fact that there are far too little foreign banks to significantly alter the estimated reduced form dynamics of the high group of uninsured banks. That said, because the response of foreign banks may be of independent interest, we also expand the reduced form with foreign banks as a separate category. The results of this exercise suggest that the response of foreign banks is never significantly different from that of the state banks (with the same time period identified as the only bank run). So the amount of deposits that is withdrawn at the uninsured banks and remains in the deposit market, flows both to the insured banks as well as the foreign banks. To that extent, both these types of banks are viewed as equally safe stores of value.

Real-time logit The cross-sectional distinction between high and low types using the real-time data set in Section 7.1.3 is based on an estimate over the entire sample period, shown in Table 2. We also consider a cross-sectional split based on a logit estimated over a gradually expanding sample. This variation ensures that for a forecast for date t the information set of the econometrician only contains information available up to date $t-1$. Therefore, such a specification can be interpreted as an actual real-time forecast. The results in Figures 8 through 10 are robust when failing banks are identified using a logit that is reestimated every time period.

In sum, the results are, typically also quantitatively, robust to a variety of changes in specification, estimation and controls. Therefore, alternative explanations do not seem to impair the validity of our results and interpretation thereof.

8 Conclusion

The method proposed in this paper identifies bank runs from raw data and simultaneously enables hypothesis tests which discriminate among various theories of depositor and bank behavior. On the one hand, the method disentangles bank runs from other possible events occurring simultaneously. On the other hand, the method can be applied to a wide range of data, which enables to assess these theories for other countries and different episodes. The method can therefore contribute to a wider understanding of the effects of bank runs, beyond the application we investigate.

With respect to our application to the Russian deposit market, we find that a bank run occurred once (and only once) during our sample period (the period 2002-2007). The results indicate that, in case of the Russian deposit market, bank runs are driven by both panic and information. Depositors run on solvent banks, a manifestation of panic. In particular, we observe a deposit outflow from solvent banks relative to insured banks. The conclusion that information-based theories also play their role follows from the finding that deposit outflows during a bank run are less severe at solvent banks compared to insolvent banks. The latter effect is quantitatively large. The run on fundamentally flawed banks is four times as harsh as the (panic) run on solvent banks. As a result, there is scope for market discipline to mitigate the effects of bank runs. However, these do not appear sufficient to prevent panic runs. In particular, at the aggregate level, random panic withdrawals account for almost 90% of the total deposit outflow at uninsured banks. Information-based withdrawals explain only slightly more than 10% of the 2004 run on Russian banks.

Since depositors run on solvent banks too, the run identified in this paper is likely to be suboptimal from a welfare perspective. In particular, unless the drain in deposits does not affect, for instance, the lending supply of solvent banks, real consequences are likely to follow. Bernanke (1983) provides some evidence on the harsh real effects (in terms of length and depth) of such crises, for the particular case of the US during the Great Depression.

In the same vein as the analysis performed in this paper, similar types of heterogeneity are present in deposit markets of more developed economies. For instance, while state banks do not (or no longer) exist in US and European banking markets, there are differences in the degree of deposit insurance for different types of deposits. Alternatively, one could think of deposits at banks that are too-big-to-fail as carrying insurance (albeit implicit). All these types of heterogeneity could be used to identify bank runs and assess the driving forces behind them.

Methodologically, the paper builds on the structural VAR literature in macroeconomics. Our focus on micro data adds to that literature in a couple of dimensions. First, we show that cross-sectional restrictions are very informative for structural identification. In addition to enabling the identification of relative shocks, they also allow relative hypothesis testing along fully unrestricted dimensions. Cross-sectional implications are, as we show in our application, a promising avenue to discriminate among theories of bank runs. The method can be adapted to also distinguish between various macroeconomic theories, which frequently have similar aggregate effects, yet different distributional implications. Second, we provide a way of external validation that is typically absent in macroeconomic VAR analyses. Rudebusch (1998) criticizes the use of VARs because of this lack of validation. He argues that identified shocks typically bear little or no relation to outside evidence on those shocks (e.g. unexpected monetary policy interventions). Our external validation exercise, by contrast, is successful and thereby underscores the validity of our approach. In addition, it suggests that this type of external validation need not be a prohibitive hurdle for SVAR studies.

Finally, the method is promising to disentangle other types of shocks to evaluate alternative theories. While this paper identifies one shock of interest within the deposit market, one can use this method to identify a more exhaustive set of shocks that determines all the fluctuations in deposit flows and the interest rates thereon, across different types of banks. An avenue for future research beyond the deposit market could address the ongoing controversy with respect to the strength of the bank lending channel relative to the financial accelerator in propagating real and nominal fluctuations. The method can go some way toward disentangling supply and demand shocks in both bank lending and firm borrowing, and investigate how different types of banks and firms are affected. This could shed light on the relative importance of the two channels as well as their strength and their distributional implications.

References

[1] Allen, F., Gale, D., 1998. "Optimal Financial Crises", *Journal of Finance* 53, 1245-84.

[2] Bernanke, B., 1983. "Non-Monetary Effects of the Financial Crisis in the Propagation of the Great Depression", *American Economic Review* 73, pp. 257-276.

[3] Calomiris, C., Mason, J., 2003. "Fundamentals, Panics and Bank Distress during the Depression", *American Economic Review* 93, 1615-47.

[4] Chamberlain, G., 1983. "Panel Data", Chapter 22 in *The Handbook of Econometrics* Volume II, ed. by Z. Griliches and M. Intrilligator, Amsterdam: North-Holland Publishing Company.

[5] Diamond, D.W., Dybvig, P.H.,1983. "Bank Runs, Deposit Insurance, and Liquidity", *Journal of Political Economy* 91, 401-19.

[6] Donaldson, R.G., 1992. "Sources of Panics: Evidence from the Weekly Data", *Journal of Monetary Economics* 30, pp. 277-305.

[7] Friedman, M., Schwartz, A.J., 1963. *A Monetary History of the United States, 1867-1960*, Princeton, New Jersey, Princeton University Press.

[8] Gorton, G., 1988. "Banking Panics and Business Cycles", *Oxford Economic Papers* 40, 751-81.

[9] Hellman, T.F., Murdock, K.C., Stiglitz, J.E., 2000. "Liberalization, Moral Hazard in Banking and Prudential Regulation: Are Capital Requirements Enough?", *American Economic Review* 90, 147-65.

[10] Holtz-Eakin, D., Newey, W., Rosen, H.S., 1988. "Estimating Vector Autoregressions with Panel Data", *Econometrica* 56, 1371-95.

[11] Hosmer, D.W., Lemeshow, S., 2000. *Applied Logistic Regression*, Wiley, New York.

[12] Iyer, R., Peydro-Alcalde, J.-L., 2007. "The Achilles' Heel of Interbank Markets", mimeo.

[13] Iyer, R., Puri, M., 2008. "Understanding Bank Runs: The Importance of Depositor-Bank Relationships and Networks", NBER Working Paper 14280.

[14] Jacklin, C.J., Bhattacharya, S., 1988. "Distinguishing Panics and Information-Based Bank Runs: Welfare and Policy Implications", *Journal of Political Economy* 96, pp. 568-592.

[15] Karas, A., Schoors, K., 2005. "Heracles or Sisyphus? Finding, Cleaning and Reconstructing a Database of Russian Banks", Ghent University Working Paper 05/327.

[16] Khwaja, A. I., Mian, A, 2008."Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market", *American Economic Review* 4, pp. 1443-1442.

[17] Matovnikov, M., 2002. "Nadezhnost' Banka Tesno Svyazana so Strukturoi ego Aktsionernogo Kapitala", Tsentr Ekonomicheskogo Analiza-Interfax, Moscow.

[18] Mamontov, A., 2005. "Gosudarstvo v Bankakh: Zlo ili Blago?", *Natsional'nyi Bankovskii Zhurnal* 12, 24.

[19] Martinez Peria, M.S., Schmukler, S., 2001. "Do Depositors Punish Banks for Bad Behavior?: Market Discipline, Deposit Insurance, and Banking Crises", *Journal of Finance* 56, pp. 1029-1052.

[20] Park, S., Peristiani, S., 1998. "Market Discipline by Thrift Depositors", *Journal of Money, Credit and Banking* 30, pp. 374-64.

[21] Postlewaite, A., Vives, X., 1987. "Bank Runs as an Equilibrium Phenomenon", *Journal of Political Economy* 95, pp. 485-491.

[22] Rudebusch, G.D., 1998. "Do Measures of Monetary Policy in a VAR Make Sense?", *International Economic Review* 39, pp. 907-31.

[23] Saunders, A., Wilson, B., 1996. "Contagious Bank Runs: Evidence from the 1929-1933 Period", *Journal of Financial Intermediation* 5, pp. 409-423.

[24] Tompson, W., 2004. "Banking Reform in Russia: Problems and Prospects", OECD Economics Working Paper 410.

[25] Uhlig, H., 2005. "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure", *Journal of Monetary Economics* 52, 381-419.

[26] Wooldridge, J., 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA, MIT Press.

[27] World Bank, 2001. "Finance for Growth: Policy Choices in a Volatile World", Policy Research Report, Washington D.C.