Identifying VARs through Heterogeneity: An Application to Bank Runs*

Ferre De Graeve Alexei Karas

Sveriges Riksbank Working Paper Series No. 244

June 7, 2010

Abstract

We propose to incorporate cross-sectional heterogeneity into structural VARs. Heterogeneity provides an additional dimension along which one can identify structural shocks and perform hypothesis tests. We provide an application to bank runs, based on microeconomic deposit market data. We impose identification restrictions both in the cross-section (across insured and non-insured banks) and across variables (as in macro SVARs). We thus (i) identify bank runs, (ii) quantify the contribution of competing theories, and, (iii) evaluate policies such as deposit insurance. The application suggests substantial promise for the approach and has strong policy implications.

Keywords: Identification, SVAR, panel-VAR, Heterogeneity, Bank run

JEL: C3, E5, G01, G21

*We thank Charles Calomiris, Fabio Canova, Russell Cooper, Dean Corbae, Olivier De Jonghe, Marco Del Negro, Simon Gilchrist, Alejandro Justiniano, Robert Kollmann, Roland Meeks, Mattias Villani, Larry Wall, 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, and the Federal Reserve System Committee Meeting on "Financial Structure and Regulation" (Boston Fed) for useful comments and valuable suggestions. An early version of this paper circulated under the title "Information-Based Bank Runs or Panics?". 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. Corresponding author: ferre.de.graeve@riksbank.se. Research Department, Sveriges Riksbank, Brunkebergstorg 11, SE-103 37 Stockholm, Sweden. A.karas@roac.nl. Roosevelt Academy, Middelburg, The Netherlands.

1 Introduction

We incorporate heterogeneity into the structural VAR methodology pervasive in empirical macroeconomics. Specifically, we propose to take micro heterogeneity on board in the process of estimation, structural identification and hypothesis testing. We introduce heterogeneity restrictions. These work in the cross-sectional dimension, whereas traditional restrictions work in the time domain. Our approach substantially broadens the scope of SVAR methods. It can thereby contribute to the empirical validation of structural models with heterogeneity, the identification of distributional shocks and testing implications in the cross-section. It also adds to microeconometric reduced form methods by enabling more structural interpretations. The method proves particularly useful when combined with the richness of microeconomic data, where heterogeneity prevails.

Applying the method to real data indicates substantial gains relative to traditional macro VARs. The cross-section adds to the informational content of the model, both in terms of identification and testing. Regarding identification, relative to macro applications, a small number of identifying restrictions suffice to obtain sharp predictions. With respect to tests, various stratifications of the cross-section powerfully discriminate between competing views, often observationally equivalent on an aggregate level. In our application, external validation of the model is successful and the model is robust to various changes in variables, specification and estimation procedure. Traditional VAR studies, by contrast, often fail external validation (Rudebusch, 1998), and entire fields exist in part due to the lack of robustness (e.g. the technology-hours debate following Galí, 1999).

Our application is on bank runs. 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. The *panic* view (e.g. Diamond and Dybvig, 1983; Peck and Shell, 2003; Postlewaite and Vives, 1987) sees bank runs as a result of coordination problems among agents, implying they can arise as sunspot equilibria. In this case, the policy spectrum consists of aggregate measures, such as deposit insurance, suspension of convertibility or liquidity provision. The alternative *fundamental* or *information-based* 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. In such an environment, policy options include ex ante imposition of balance sheet constraints, ex post recapitalization, or even laissez-faire.[1]

In our application, we identify bank runs with heterogeneity restrictions, and quantify the contribution of competing theories in the cross-section and on aggregate. Identification of bank runs exploits variation in deposit insurance across banks in a multivariate system of deposit interest rates and quantities. The results we provide are structural, i.e. conditional on a bank run. They add to earlier reduced form evidence on bank runs. The application is on Russian deposit market data for the period 2002-2007. Russian micro bank data are not only of exceptional quality, they are also very informative: our sample includes at least one severe market disruption, dozens of bank failures, cross-sectional heterogeneity in deposit insurance, and more. While the application helps to highlight the richness of the approach, our results also 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 information-based view. On the other hand, even banks with solid fundamentals face significant outflows. This finding, in turn, 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 from an aggregate perspective. 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

[1]For recent theoretical insights on the policy implications of bank runs, see, e.g. Goldstein and Pauzner (2005) and Ennis and Keister (2009). For general equilibrium perspectives on bank runs, see Cooper and Corbae (2002) for an example of the panic view and Uhlig (2009) for a fundamental view.

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 toward effectively shielding depository institutions from panic effects may be required to do so effectively.

Our broad conclusions align well with recent experimental evidence that finds support for the coordination failure view of bank runs (Madiès, 2006; Garratt and Keister, 2009; Schotter and Yorulmazer, 2009). Our findings also do not appear to be inconsistent with events observed during the recent crisis. Worldwide, one has witnessed plenty of arguably solvent banks facing problems and heard many calls for systematic measures. One type of policy adopted by many countries in response to the recent financial turbulence is an increase in the coverage rate of deposit insurance. Examples include the U.S., where the FDIC increased the coverage limit from \$100,000 to \$250,000, and many of the European member states, where some countries (e.g. Germany, Ireland) even went as far to fully cover deposits, without limit. In addition to disentangling fundamental and panic effects, our results provide an estimate for the effectiveness of deposit insurance.

The paper is organized as follows. Section 2 starts with a short review of the macro-econometric approach to structural identification. We then lay out how heterogeneity can be incorporated in such a setting. Section 3 describes our application. We first provide details on events in our sample period, and discuss our identification strategy. Next, we present results on the effect of bank runs, quantify the importance of the competing views and discuss the relation to other empirical approaches. After analyzing the scope for alternative interpretations, and verifying the robustness of our results, we conclude in Section 4.

2 Identification through heterogeneity

We start with a brief review of structural identification in vector autoregressions (VARs). Consider a reduced form VAR:

$$Y_t = A(L)Y_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim N(0, \Sigma), \tag{1}$$

where Y is a vector of endogenous variables $Y^{(m)}$, $m = \{1, ..., M\}$, t indexes time, and $A(.)$ is a matrix polynomial in the lag operator L. A reduced form such as (1) does not allow structural interpretations (all variables are endogenous, the reduced form residuals are an amalgam of structural shocks). In other words, the economist's interest is typically in structural models such as (2):

$$CY_t = B(L)Y_{t-1} + u_t \qquad u_t \sim N(0, D). \tag{2}$$

Crucially, such a model is characterized by simultaneous interactions between variables in Y, through C. The driving forces in models of this kind are structural, exogenous shocks, u_t. The latter is manifested by u_t having a diagonal covariance matrix, D. Dynamic stochastic general equilibrium (DSGE) models, among many others, fit this kind of structure. Note that any particular reduced form such as (1) is consistent with multiple structural models; the data do not allow us to distinguish between them. To pin these down, restrictions from economic theory are typically imposed. Imposing such restrictions serves to identify the VAR, making it structural (hence, SVAR). The power of structural VARs lies in the fact that they allow the recovery of interesting patterns in the data using a minimal amount of theory. This is especially useful in fields where there is little theoretical consensus, or where models are less than fully specified.

The entertained identifying restrictions take different forms. They constrain the impulse response functions of variables to shocks, and the most pervasive types are:

- Short-run restrictions (Sims, 1980; Bernanke, 1983; Christiano et al., 1999):

$$\frac{\partial Y_t^{(m)}}{\partial u_t^{(k)}} = 0 \tag{3}$$

- Long-run restrictions (Blanchard and Quah, 1989; Galí, 1999):

$$\frac{\partial Y_\infty^{(m)}}{\partial u_t^{(k)}} = 0 \tag{4}$$

- Sign restrictions (Faust, 1998; Canova and De Nicoló, 2002; Uhlig, 2005):

$$\frac{\partial Y_{s\in S}^{(m)}}{\partial u_t^{(k)}} \lesseqgtr 0 \tag{5}$$

where $u_t^{(k)}$ is a particular structural shock, $\frac{\partial Y^{(m)}}{\partial u_t^{(k)}}$ denotes the impulse response of variable $Y^{(m)}$ to that shock, and s is time, with S a set of time periods. Imposing these restrictions on reduced form VARs allows one to recover the structural shocks and how the endogenous variables respond to them. Such identified models then inform us how different variables behave across the set of models that satisfy the imposed restrictions.

We propose to fully incorporate the cross-sectional dimension into the structural VAR method. Consider first the following generalization of the reduced form (1), that takes account of both the time and cross-sectional dimension:

$$Y_{i,t} = A_i(L)Y_{i,t-1} + \varepsilon_{i,t} \qquad \varepsilon_{i,t} \sim N(0, \Sigma_i), \tag{6}$$

where the symbols are the same as in (1), but are now also indexed by i $(= 1, ..., N)$, denoting cross-sectional units. Equation (6) is a reduced form panel-VAR. It embeds traditional VARs as a special case, in which there is only one cross-sectional unit. Panel-VARs are introduced by Chamberlain (1983) and Holtz-Eakin et al. (1988). The reduced form (6) can capture additional complexity, such as time-varying coefficients, factor structures and more (e.g. Binder et al., 2005; Canova and Ciccarelli, 2009). The point we wish to make does not hinge on the inclusion or absence of those, and for ease of exposition we leave them out of what follows.

To our knowledge, there are very few examples of studies that identify panel-VARs. Those that do (e.g. Canova and De Nicoló, 2002), achieve identification by imposing traditional restrictions. Put differently, these models achieve identification by imposing restrictions common to all cross-sectional units. Instead, we suggest taking advantage of **heterogeneity restrictions**. These extend traditional identification to the cross-sectional dimension and impose restrictions on subsets of the cross-section:

$$\frac{\partial Y_{\Lambda,s}^{(m)}}{\partial u_t^{(k)}},$$

where $m \in \{1, ..., M\}$, s denotes time (depending on the type of restriction imposed) and $\Lambda \subseteq \{1, ..., N\}$ is an index set (indexing cross-sectional units).[2] Such restrictions can take various forms. First, note that this class of restrictions nests traditional restrictions of the

[2]While, in principle, Λ can consist of a single element of the cross-section, considering sets is useful as

type (3)-(5). These are characterized by $\Lambda = \{1, ..., N\}$ and impose a restriction on the cross-section as a whole. Second, and more importantly, incorporating subsets of the cross-section opens up a new array of possible restrictions, and thereby a new array of structural models that can be evaluated. Basically, heterogeneity restrictions require different implications for different cross-sections. They include restrictions:

- within variables, across subsets, e.g.:

$$\frac{\partial Y^{(m)}_{\Lambda_1,s}}{\partial u^{(k)}_t} \lesseqqgtr \frac{\partial Y^{(m)}_{\Lambda_2,s}}{\partial u^{(k)}_t}$$

- across variables, within subsets, e.g.:

$$\frac{\partial Y^{(m_1)}_{\Lambda,s}}{\partial u^{(k)}_t} \lesseqqgtr \frac{\partial Y^{(m_2)}_{\Lambda,s}}{\partial u^{(k)}_t}$$

- across variables, across subsets, e.g.:

$$\frac{\partial Y^{(m_1)}_{\Lambda_1,s}}{\partial u^{(k)}_t} \lesseqqgtr \frac{\partial Y^{(m_2)}_{\Lambda_2,s}}{\partial u^{(k)}_t}.$$

All these types 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 an analogy to traditional SVARs, it may be natural to constrain the behavior of some subsets of the cross-section for some variables, while leaving others free.[3]

2.1 Discussion

As we show in our application, these restrictions also appear to be very informative. In particular, the application has three characteristics that many macroeconomic VARs do not

---

the reduced form estimation may involve a substantial amount of dimension reduction. In addition, from an identification perspective, it may often be more natural to impose restrictions on a subset of the cross-section rather than on individual units.

[3]Identification restrictions are traditionally accompanied by an orthogonality assumption on the structural shock covariance matrix as well as an invertibility condition. For the latter, see e.g. Fernández-Villaverde et al. (2007). There is no reason for non-fundamentalness to be more of an issue in the current setup relative to macro VARs. By contrast, information originating in the cross-section may, in some cases, serve to achieve fundamentalness.

have. First, there is a substantially reduced need for alternative restrictions. For instance, contemporaneous sign restrictions suffice to achieve identification and deliver sharp predictions. By contrast, extant applications of sign restrictions invariably impose restrictions over longer time spans. Second, identified structural shocks appear consistent with information external to the model, thereby overcoming earlier critiques of SVARs, such as Rudebusch (1998). Third, the results are extremely robust. This holds in various dimensions, including variable and model specification, and contrasts with many macro VAR applications.

Some approaches are related to ours. For instance, Peersman (2009) considers a two-country macro VAR and identifies symmetric and asymmetric shocks. This gives rise to a similar structure. The setup we consider is, however, a lot more general. In part, this is because estimation can take advantage of the panel dimension. As a result of the large amount of data, there is less of a curse of dimensionality relative to standard VARs, as there is a lot of scope to consider factor structures (as in, e.g. Boivin et al., 2009; Canova and Ciccarelli, 2009). Crucially, however, numerous cross-sections allow for many possible stratifications. Economic variables can underlie the stratifications. This implies that theory can be linked to the empirics not just in the time series dimension (as in the typical macro SVAR), but also through the cross-section.

There are a couple of studies that incorporate the cross-section into the testing stage. Examples include Canova and Pappa (2006), who separately identify state-specific fiscal shocks and analyze whether their effects differ depending on budgetary characteristics of those states. Another example is Gertler and Gilchrist (1994), who study cross-sectional implications of (monetary) shocks identified at the aggregate level. Boivin et al. (2009) is similar in perspective. We argue that the cross-sectional advantage extends beyond the testing stage. There is scope for using multiple, and possibly different, stratifications in all three stages of the analysis: estimation, identification and testing. Moreover, these restrictions are particularly rich when applied to panel-VARs on micro data. The reason is obvious: heterogeneity prevails in micro data.[4]

[4] While there is other work applying VAR techniques to micro data, e.g. the early work of Chamberlain, or Franco and Philippon (2007), our approach effectively takes advantage of the cross-section in a structural manner.

From a broader perspective, heterogeneity prevails in much of modern macroeconomic analysis. This holds true for models, shocks and empirical tests. With respect to models, for instance, the past decade has seen a proliferation of models with heterogeneous agents following Krusell and Smith (1998). In terms of shocks, much attention has been devoted to distributional shocks. Examples include non-neutral technology shocks (Greenwood et al., 1997) and distribution risk (Danthine and Donaldson, 2002). Concerning tests, multiple macroeconomic theories are observationally equivalent on an aggregate level. One way to resolve such macroeconometric equivalence is to study the implications for the cross-section (Levin et al., 2008). While distributional consequences are interesting in their own right, they also allow to discriminate between macro theories. The empirical literature on the credit channel of monetary policy is one example (e.g. Gertler and Gilchrist, 1994) of that approach. While early empirical tests concentrate on aggregate fluctuations, researchers eventually turned to the cross-sectional dimension in search of answers. The literature on bank runs is another example. Our cross-sectional take on SVARs has the scope for empirically validating DSGE models with heterogeneity. From the microeconometric point of view, it allows one to draw more structural inference in fields where empirical evidence is typically reduced form in nature.

3 An application to bank runs

We now operationalize heterogeneity restrictions by applying our approach to the field of bank runs. Following recent financial turbulence across the globe, bank runs have taken center stage again. This holds true both from an academic and a policy perspective. The state of affairs in the academic literature makes it a prime candidate for our method. On the theoretical front, there are two competing views on the causes of bank runs: the panic and the fundamental view. From a model perspective, there is little consensus in the field. Put differently, there is no workhorse model in banking (and bank runs in particular), whose structural estimation one could put faith in.[5] Rather, the field consists of a large amount

[5]For instance, the recent overview of empirical research on bank runs in Degryse et al. (2009, Chapter 7) contains no structural models. In macro, by contrast, there appears to be a somewhat broader consensus,

of less than fully specified theoretical models, each geared to highlight particular important features. On the empirical front, virtually all the evidence is reduced form in nature. As a consequence, definitive structural distinctions between the different views are rather scarce. Our approach comes natural in such a field: (i) it allows structural inference with minimal use of theory, and (ii) cross-sectional differences prove quintessential, both in identifying the event of interest (runs) as well as in discriminating between the competing views.

3.1 Background

Our application focuses on 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). 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. Private banks, by contrast, do not have the state backing them (or their deposits). Our method will exploit such heterogeneity in insurance between state and private banks to identify bank runs. 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 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

which has contributed to the estimation of DSGE models (e.g., Smets and Wouters, 2007).

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 of confidence 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.

3.2 Identification

Our application starts with a reduced form model of deposits and the interest rates paid on them. The estimation uses data at the bank level, and detailed data characteristics are contained in Appendix A. The particular reduced form we entertain is a panel-VAR. There are a number of reasons that advocate 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, and our data is no exception in that respect. Second, the majority of structural models have a reduced form representation which is encompassed by this model. Third, while maintaining consistency with the variety of structural models, there is no need to make strong and debatable assumptions regarding the functional form of demand and supply equations. Fourth, it provides a flexible way of dealing with heterogeneity, where deemed necessary.

The reduced form model we consider takes the following form:

$$\begin{bmatrix} D(U)_{i,t} \\ R(U)_{i,t} \\ D(I)_{j,t} \\ R(I)_{j,t} \end{bmatrix} = c + A \begin{bmatrix} D(U)_{i,t-1} \\ R(U)_{i,t-1} \\ D(I)_{j,t-1} \\ R(I)_{j,t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon(U)^D_{i,t} \\ \varepsilon(U)^R_{i,t} \\ \varepsilon(I)^D_{j,t} \\ \varepsilon(I)^R_{j,t} \end{bmatrix}. \tag{7}$$

This is a panel-VAR on (log) deposit quantities, D, and deposit interest rates, R.[6] The

[6]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 (7) takes up the role of these variables. We write the system with one lag and without additional control

indices I and U refer to different types of banks, Insured and Uninsured banks, respectively. Subscripts i and j denote (group-specific) cross-sectional units and t indexes time. A is a coefficient matrix and c is a vector of constants. We allow different types of banks to exhibit different reduced form coefficients. The vector, ε, contains reduced form shocks for all banks.

Given this structure, reduced form systems such as (7) tell us little or nothing about the economics in the data. Significant coefficients in such a system do not admit a structural interpretation. The covariance matrix of the reduced form residuals is non-diagonal and also has no particular structural interpretation. Moreover, there are no contemporaneous interactions between the different endogenous variables. These features are 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 extract one particular shock of interest, *viz.* a bank run. To learn about its effects, we put additional structure on the reduced form.

Our identifying restrictions filter out a bank run from concurrent developments in the deposit market. We define a bank run as *a supply shock in which insurance matters*:

$$\begin{cases} \Delta D(U)_t < 0 \\ \Delta R(U)_t \geqslant 0 \\ \Delta D(U)_t < \Delta D(I)_t \\ \Delta R(U)_t \geqslant \Delta R(I)_t \end{cases}, \tag{8}$$

where Δ is shorthand for an impulse response, $\frac{\partial}{\partial run_t}$.[7] The absence of the cross-sectional index i (resp. j) conveys these correspond to the average responses of deposits and interest rates, over the cross-section of uninsured banks (resp. insured).

variables, for conciseness. Our baseline results are based on a specification that includes four lags, the choice preferred by standard lag length criteria.

[7] Thus, Δ measures the change relative to baseline, where the latter is measured by the dynamics of the system (7) 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 impulse responses are concerned with changes in response to a particular shock, relative to that 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.

The first restriction says, quite uncontroversially, that a bank run lowers the quantity of deposits at uninsured banks. In addition, the second restriction restricts 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 the fact that uninsured banks lower the interest rate they pay on their deposits, e.g. in response to lower loan demand. To exclude such cases, we add a restriction on the interest rate. These two restrictions combined effectively rule out demand shocks in the deposit market. We additionally impose a *heterogeneity restriction*, contrasting the behavior of different types of banks. In particular, the third restriction requires that a bank run is 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. 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 an even larger increase of the interest rate at insured banks.

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 to that requirement. 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 (8) 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 (1983) 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 restriction as the logical extension of different equilibria in these kinds of models to a cross-sectional setup.

Finally, note that the restriction on the response of insured banks is only *relative* to the uninsured banks. As a result, 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 does require that having deposit insurance does not aggravate the deposit outflow relative to banks that are not covered by the deposit insurance scheme.[8]

The combined set of identifying assumptions filter out bank runs from other forces that affect supply and demand in the deposit market. We identify bank runs as supply shocks that are worse for uninsured banks than for insured banks. Moreover, the runs we consider should be thought of as systemic, as we impose the restrictions on the group-wise behavior.[9,10]

[8] A number of papers, among which Demirgüç-Kunt and Huizinga (2004), find that deposit insurance increases the probability of a banking crisis. The rationale is that this occurs because insurance reduces market discipline on behalf of the depositors, thereby increasing moral hazard on behalf of the banks. Our assumption that the deposit outflow at insured banks is less harsh compared to the uninsured may seem at odds with that literature. First, however, the type of data here is substantially different: our restrictions pertain to within-country variation in deposit insurance (not cross-country) and to deposit outflows (not banking crises). Second, if insurance leads to less market discipline, then it must be that insurance is credible, which renders our identifying assumption uncontroversial.

[9] Note that the stratification level used here is not at the bank-specific level, but at an intermediate level. First, while identification is achieved at this level, this does not imply equivalence with identification based on group-wise aggregated data. From an estimation perspective, the obtained reduced form is substantially more precise by incorporating micro information. From an identification and testing perspective, additional (sub-)levels of stratification provide additional information. Second, this is not a restrictive feature of the method, but of particular interest in the present application.

[10] Computationally, the approach consists of a search for orthogonal decompositions of the variance-covariance matrix of the reduced form group-wise average residuals, $[\bar{\varepsilon}(U)_t^D;\ \bar{\varepsilon}(U)_t^R;\ \bar{\varepsilon}(I)_t^D;\ \bar{\varepsilon}(I)_t^R]$, which satisfy a particular set of restrictions common to a variety of structural models. For details on implementation within a macroeconomic framework, see e.g. Uhlig (2005). In addition to searching among the many possible roots of the shock variance-covariance matrix, coefficient uncertainty of the estimated reduced form is also taken into account. Drawing exact confidence bands in the present framework requires the development of additional econometric theory. Confidence bands will not only depend on the relative length of time and cross-sectional dimensions, but potentially also need to take into account attrition over the heterogeneous groups, unbalancedness across groups, and more. Developing that theory is beyond the scope of this paper. Instead, for drawing confidence bands, we stick to the macroeconomic approach, treating the panel-VAR as

3.3 A first look at the effects of a bank run

Figure 1 plots the effect of a bank run, identified with the restrictions in (8), on deposits and interest rates across insured and uninsured banks. The responses are to a one standard deviation impulse and measure the responses for 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, recall that the restrictions are imposed only on impact, at $t = 0$. The apparent persistence of the reduction in uninsured deposits is substantial. It takes more than a year for the effects of a bank run on the volume of deposits to dissipate. On a methodological note, the fact that contemporaneous sign restrictions suffice to achieve identification speaks to the informational content of heterogeneity restrictions. The fact that our confidence bands are conservative -i.e. they overestimate uncertainty- adds to that. Typical macro VARs, by contrast, require sign restrictions to hold over substantially longer horizons.

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 the inflow at insured banks is not particularly significant. Crucially, however, insured banks do not face a deposit outflow. Hence, insured banks are not subject to the run.

Third, Figure 2 plots a confidence interval for the identified shock over our sample period. The single largest shock is observed during the summer of 2004 (both Q2 and Q3 are

---

if it were an aggregate VAR. The fact that the dynamics are estimated using an additional cross-sectional dimension tends to lower estimation uncertainty, thereby reducing confidence band width. Neglecting this reduced estimation uncertainty when we compute confidence regions therefore works (strongly) against finding significant differences. More on this in Section 3.6.

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.[11] 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 NY Times for our sample period using the terms "Russia", "deposit" and "run". Out of all hits, three pertained directly to the present paper's subject. All three were dated summer of 2004 and each of them suggested the possibility of a bank run.[12] 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. This type of external validation resolves a number of issues some have raised as a criticism to the use of structural VARs. Foremost, our approach is not subject to the Rudebusch-critique. Rudebusch (1998) 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.[13] The time series of bank runs identified by the model appears to be in accordance with outside information. Therefore, this type of external validation provides additional support for the validity of our approach. Moreover, the time series of bank runs is in agreement across a variety of robustness checks, also contrary to the case of Rudebusch (1998).

[11] The negative shock in 2005:Q4 is somewhat less robust across specifications. For an interpretation of the negative shock, see Section 3.6.

[12] 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" (NY 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 sixteen hits, nine directly pertain to the present paper's subject. Those nine are dated July through September of 2004 and each suggested a possibility of a bank run.

[13] Sims (1998) argues that the impulse responses are of interest even if external validation of this kind fails. Our results here serve to show that the shock series itself makes sense, even though that is not a strict requirement.

3.4 Evaluating the theories

We here provide substantially more detail on the above 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 (henceforth "good banks") and banks with flawed fundamentals ("bad banks"). 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. The fundamental view predicts depositors will run precisely on those banks they deem at risk. In this view, depositors have no incentive to withdraw from banks for which there is no insolvency concern. According to the panic view, by contrast, depositors run on banks, irrespective of their fundamentals.

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 characterizations 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 e.g. Park and Peristiani (1996) and Calomiris and Mason (2003b). While this logit model may be of independent interest, we refer the interested reader to Appendix B 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).[14]

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.[15] The coefficient in the upper panel can be interpreted directly as the percentage change in the deposit base for the different groups of banks. We focus on the contemporaneous impact.

Panel B provides p-values on two particular hypothesis tests. These tests evaluate the significance of the difference in deposit response between different types of banks. A 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 whether the response at good banks is more severe than that of the insured banks. The panic view on bank runs predicts depositors will run on healthy banks, too. Hence, a p-value below the conventional significance levels, along with finding a significant outflow at good banks, provides support for the panic view.

We are now ready to quantify the effect of the two competing views on 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 quantitatively large and amounts to at least 10% of good banks' deposit base. Such an outflow is not observed at insured banks (Panel A, row 3), as corroborated by the according p-value on the difference between good and insured (Panel B, row 2). This

[14] One can compute these differences in different ways: 1) by expanding the panel-VAR with the additional (sub-) groups and perform the identification step again, or 2) by performing a panel regression of the variables of interest on the shock series resulting from the two-group panel-VAR. The results presented are those based on the latter approach, but our conclusions are insensitive to this choice (Appendix C for the former approach). Moreover, for each classification, one can stratify on the basis of the entire sample or based on a particular time period. This, too, leaves results unaffected.

[15] The 2004 response provides a quantitatively more appealing measure of the impact of the run. The impulse responses in Figure 1 (and Appendix C for the subgroups) in analogy to macroeconomic VARs, measure the effect to a one standard deviation structural shock. We view the quantitative response to the 2004 run, observed in Figure 2 and confirmed by external evidence, as a more relevant one in the current setting. To compute that impact, we rescale the shock to have unit value in 2004:Q3.

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.

Second, bad banks also lose at least 10% of their deposits. Importantly, for the ex ante and ex post stratifications, we find significantly stronger outflows at bad banks relative to good banks. Thus, fundamentally flawed banks face even more significant runs. 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 3 plots the first year response of deposits across the different types of banks for two of our stratifications.

For these results to have policy relevance, however, the relative importance of the two views needs to be assessed.[16] Therefore, in addition to the impulse responses to the 2004 run, the bottom panel of the table computes the implied aggregate effects. These enable the quantification of the aggregate importance of deposit flows between the different types of banks, as well as effects on the deposit market as a whole.[17] In panel C, the first row calculates the total outflow of uninsured deposits. In aggregate terms, the uninsured deposit market shrinks by 10 to 15% (panel C, row 1). The next two rows decompose the aggregate outflow into the part driven by fundamentals (panel C, row 2) and the part caused by panic (panel C, row 3). It turns out that the panic view is the primary contributor to the run in our sample. Fundamental effects, i.e. the more severe outflows at bad banks, explain no more than 15% of the total deposit outflow.[18] Whichever way one classifies good and bad banks, good banks always lose a significant fraction of their deposits. From an aggregate

[16] See e.g. Calomiris and Mason (2003b) for an alternative empirical assessment and Goldstein and Pauzner (2005) for a theoretical one.

[17] Aggregate effects are calculated based on the point estimates in the upper panel of the table by taking into account the average number of banks of the various types. Similar aggregate results are obtained when the 2004:Q3 number of banks is used.

[18] 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 the good banks, but it now applies to a relatively small fraction of banks.

perspective, this outflow is the main contributor. The final row of panel C measures the inflow of deposits at insured banks as a proportion of the outflow of uninsured deposits. Insured banks absorb only a small fraction of the outflow from the uninsured deposit market (1-3%, panel C, row 4). 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.[19]

From a policy perspective, this suggests that the primary concern is shielding fundamentally solvent banks from bank runs. In our sample, this is readily achieved by deposit insurance. Insured banks withstood runs by depositors. Since poor fundamentals can severely aggravate runs, there is scope for fundamentals-based regulation, too. In our sample period, 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 1 that the inflow of deposits at insured banks is not demand-driven: there is no change in the deposit interest rate of insured banks, while uninsured banks increase theirs. First, note that the responses in the table are in percentage points. The increase in the uninsured interest rate, while significant, is not very large - though it may mask some intra-group heterogeneity. The table shows that (and this is confirmed in most, but not all, of the robustness checks), in cases where we observe significant fundamentals, there is a tendency for the 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 (absolute or relative) 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

[19] 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 e.g. Friedman and Schwartz (1963), Bernanke (1983) and Calomiris and Mason (2003a).

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 et al. (2000).

3.5 Discussion

A major difficulty in assessing the driving forces of bank runs is singling out the run from other factors. We here provide an overview of the more recent contributions to the empirics of bank runs, and how our approach relates to those. For an overview of earlier empirical analyses, see Calomiris and Gorton (1991) and Gorton and Winton (2003). We focus on three particular issues that complicate empirical analysis of bank runs. These are subjectivity, exhaustivity and endogeneity. We here discuss how the approach taken in this paper deals with them.

Most of the empirical research on bank runs relies on a *subjective* form of identification. In particular, it studies the effect of particular periods that have been characterized as a bank run. For instance, 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 interpretation has been contested by many, including Gorton (1988), Saunders and Wilson (1996) and Calomiris and Mason (2003b). 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, 2003b). However, especially in an area where the definition of the object of study -bank runs- is so elusive (see e.g. Calomiris and Winton, 1991) this subjective nature is of major concern. As a consequence of subjectivity, Gorton and Winton (2003) and Ennis (2003) point out how different authors disagree on whether or not a particular period constitutes a bank run. In our approach, identification relies on a priori restrictions. These force one to be very specific about definitions, which reduces the scope for subjectivity.

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 (i.e. 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, in itself, does not 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 (2010) and Iyer and Puri (2008), who investigate the effects of a bank fraud discovery in India. In a related area of research, Khwaja and Mian (2008) analyze the effects of an unexpected nuclear test in Pakistan. Our 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 contain an exogenous event.

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. As in event-studies, they necessarily assume the run is the only thing that occurs during that particular period. 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.[20] 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 can not distinguish liquidity preference type shocks from bank runs. Contrary to both the narrative and the exogenous event approaches, the present method does not require making an exhaustivity assumption, *viz.* 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

[20] Note that the restrictive nature of this exhaustivity assumption increases the more aggregate in nature the event is and the lower the frequency of the data. However, especially for aggregate events, which are more likely to affect expectation formation, such convolution could well be instantaneous.

addresses the issues of subjectivity, exogeneity, as well as exhaustivity. There is some recent experimental evidence on bank runs (Madiès, 2006; Garratt and Keister, 2009; Schotter and Yorulmazer, 2009) to which these issues do not apply (by construction). Interestingly, our broad conclusion aligns well with those studies; there is an empirical role for the panic view of bank runs.

3.6 Alternative interpretations, extensions and robustness

Identification In the baseline results, the restrictions are imposed on the uninsured group as a whole. However, it is possible that runs occur in subgroups of the uninsured pool of banks, but do not result in deposit outflows across the entirety of uninsured banks. To verify whether such is the case, we re-run the analysis for all stratifications twice: once with the restrictions imposed only on bad banks versus insured and once with the restrictions imposed only on good banks versus insured. Table 3 shows the results of that exercise, both for the ex ante (columns I and II) and ex post stratification (columns V and VI). Invariably, the estimates confirm the baseline results: both panic and fundamentals are at work and the former dominates in the aggregate. In all cases, the timing of the run remains very similar.

From looking at raw deposit market data, 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 period 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, as explained in the discussion on the exhaustivity assumption. Moreover, the longer sample allows to control more efficiently for other types of shocks important for the deposit market. Second, and conversely, the approach also allows for bank runs to have occurred, yet for them not to be immediate from data aggregates. Although the results indicate that such runs did not occur, one can not exclude this a priori.

Deposit outflows and stratifications The fundamental effect for the ex post stratification could look very similar if all banks were equally solvent in the high and low groups, and withdrawals were the only source of failure. This would imply there is no prior informational difference between the two groups that can discriminate good from fundamentally weak banks. In this case, we would incorrectly attribute effects to the fundamental view. Note that the relative importance of the fundamental view is not large to begin with in our results, at least from an aggregate perspective. Hence, this concern does not apply to the evaluation of panic effects. Thus, if anything, this would suggest that the baseline estimate might underestimate the scope for panic. The ex ante stratification, for which the fundamental effect is also present, has two features that reduce the above concern.[21] On the one hand, the information used to forecast default does not contain deposit growth. Thus, more is happening on the bank's balance sheet. On the other hand, there is a timing difference which reduces the concern of runs being the cause of default. The impulse responses in Table 1 are for stratifications determined prior to the shock. In other words, the impulse responses measure the response to a run across good and bad banks, where the latter stratification is based on information that predates the run. Thus both the type of information used and its timing reduce the concern for reverse causation.[22]

Panic In what preceeds we label the outflow at good banks as panic-driven. This characterization is similar in spirit to that of Saunders and Wilson (1996). It is, however, more

[21] While the ex ante point estimate for the fundamental effect in Table 1 is smaller than in the ex post case, it applies to more banks. This occurs because using the logit along with the median as the cutoff between bad and good, it overpredicts the number of defaults. Irrespective of the stratification, however, the total effect and its decomposition are in agreement.

[22] From a methodological perspective, our analysis studies responses for a *given, discrete* stratification. Ultimately, however, one may want to think about incorporating dynamics in stratifications, as well as continuous stratifications. The method could deal with that, in principle. In particular, one could envisage a model with reduced form coefficients exhibiting systematic heterogeneity. Identified shocks could then affect both the variables of primary interest as well as those determining heterogeneity. Within standard macro VARs, exogenous switching is already hard to deal with (for some recent contributions, see Rubio-Ramirez et al., 2010). Endogenous switching of the type alluded to above creates additional challenges beyond the scope of the current analysis.

precise. In particular, the fact that such outflows do not occur at insured banks reduces the scope for alternative explanations. In the absence of insured banks as a control, macro effects arise as a particular concern. For instance, Covitz et al. (2009) classify runs (on asset-backed commercial paper) as either discriminate or indiscriminate. There, indiscriminate runs are those that are not related to fundamentals (of the commercial paper program). But they may well be driven by macro effects, rather than be manifestations of panic. Along the lines of Calomiris and Wilson (2004), for instance, one could attribute such outflows to an overall increased depositor risk aversion. In our results, the response at good banks is not part of a general outflow of the deposit market, but rather particular to uninsured banks. This is what reduces the concern for alternative (macro) explanations.

Regional fundamentals In view of the evidence provided in Calomiris and Mason (2003b), one may wonder whether regional fundamentals could drive the above results. To check whether that is the case, we redo the analysis for a subset of banks, viz. those located in Moscow. Columns III and VII in Table 3 show that there is some variation in point estimates relative to the baseline results. For instance, the evidence in favor of the fundamental view is no longer significant for the ex post stratification, but turns out to be somewhat stronger for the ex ante case. For the ex ante stratification, the aggregate fundamental contribution to the run now almost reaches 30%. Overall, however, panic is invariably significant and predominant at the aggregate level.

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 organizations in the (typically Western) home country will allow their foreign subsidiaries to fail. The main results continue to hold when we drop the foreign banks from the analysis altogether. 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 result of this exercise, contained in Table 4, suggests that the response of foreign banks is not significantly different from that of the state banks. 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.

Fundamentals of insured banks In principle, one could also test whether depositors distinguish between good and bad state banks during a run. However, the classifications we use in the analysis, in particular the ex ante and the ex post ones, are hard to apply to state banks. The reason is that there were no failures of state banks in our sample. So both the ex ante and ex post stratification would result in the bad insured bank group being empty. Table 4 checks whether depositors distinguish between insured banks on the basis of their capitalization or liquidity. We do not find such differences to be significant.

Depositor characteristics Kelley and Ó Gráda (2000) and Iyer and Puri (2008) show that, at a given bank, depositor characteristics matter for the decision to withdraw. While we study withdrawals across rather than within banks, these, too, may be affected by differences in the pool of depositors at different banks. There may be depositor characteristics that explain why depositors withdraw more at bad banks than at good banks, and a lot more at good banks relative to insured banks. To control for such differences, we combine our approach with difference-in-difference techniques. We ask whether the effect of the run is different from other cases in which depositor characteristics matter. Depositor characteristics are supply factors. Therefore, we ask whether there is a significant difference across banks in the response to a run and other (non-run) supply shocks. To answer that, we construct the following test statistics:[23]

$$T^{FUND} = \left[\frac{\partial D(U,Bad)}{\partial run} - \frac{\partial D(U,Good)}{\partial run}\right] - \left[\frac{\partial D(U,Bad)}{\partial non-run\ supply} - \frac{\partial D(U,Good)}{\partial non-run\ supply}\right]$$

$$T^{PANIC} = \left[\frac{\partial D(U,Good)}{\partial run} - \frac{\partial D(I)}{\partial run}\right] - \left[\frac{\partial D(U,Good)}{\partial non-run\ supply} - \frac{\partial D(I)}{\partial non-run\ supply}\right].$$

Note that the terms in the first brackets are the baseline results of Table 1. The second brackets contain the controls and measure the respective responses to alternative supply

[23] We compute the non-run supply shock as the part of unconstrained supply shocks that is orthogonal to the bank run.

shocks, as a way of keeping depositor characteristics constant. Thus, if our earlier results are not driven by depositor characteristics, one would expect T^{FUND} and T^{PANIC} to be significant, as before. Table 4 contains the results for the different stratifications used and broadly confirms the earlier conclusion: there is some evidence for the fundamental view, while strong indications of panic effects.[24]

Uninformed depositors In our approach to identifying bank runs, it is not the case that we assume that depositors are completely uninformed, as could be the case in a fully random panic. Our approach requires depositors to know whether or not their deposits are insured. We view this as a very minimalist informational assumption. It is obvious 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.

Credibility of deposit insurance One may argue that deposit insurance, which we use in identification, is not credible. Non-credible deposit insurance is not necessarily problematic for our method. On the one hand, if deposit insurance were 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 uninsured 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. Related to this issue, in some of our results we find a negative shock in 2005:Q4. While it is not as large nor as robust as the 2004 positive peak, it does deserve some discussion. A negative shock implies an inflow to uninsured banks relative to insured banks which is not driven by (a relative rise in) the interest rate. One possible interpretation for a negative shock consistent with our identifying assumptions is that it measures reductions in the credibility of insurance. Interestingly, the negative shock occurs a couple of quarters following the crisis. The factual regulatory response to the crisis was to adopt a general deposit insurance

[24] A concern with T^{PANIC} is that it intertwines with our (relative) identifying restrictions. To make sure that is not the case, we run the test on shocks identified using bad and insured banks only, leaving the good banks' response unconstrained.

scheme. To the extent that banks' enrollment in the new insurance scheme is deemed credible by depositors, their money transfers from insured to (previously) uninsured banks could be reflected by a negative shock.[25] In that sense, the negative shock a year after the crisis can indicate the time it took for the general deposit insurance program to gain credibility. Again, one can look at narratives to infer the plausibility of such a scenario. Indeed, the deposit insurance agency did publish reports suggesting a slow response.[26]

Credibility of non-insurance One may also wonder whether it is truly credible that the government will allow non-insured banks to go bust. Recent experience in many Western countries, for instance, shows the resilience of governments to let banks go bust. Russia, by contrast, has witnessed many bank failures: ten percent over the course of our sample period. Thus, at least from an ex post perspective, non-insurance is clearly credible.

Other robustness checks In addition to the extensions discussed above, we performed a wide variety of robustness checks. First, all results carry through when the deposit variable used in the estimation is specified in log-differences rather than in log-levels. Moreover, the baseline results are based on a panel-VAR with four quarterly lags and without additional controls. Different lag length, incorporating time dummies or including bank balance sheet variables directly in the reduced form leaves all conclusions unaffected. In addition, the fact that incorporating fundamentals in the reduced form does not alter our results reduces the concern for anticipated fundamental shocks contributing to our results.

Second, 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

[25] The classification between state and private banks in our empirical exercise is fixed. Thus, it is maintained after the introduction of insurance for non-state banks.

[26] In particular, by the end of 2004:Q4 a relatively small fraction of Russian banks had enrolled in the deposit insurance program (31% of all banks, 22% of system-wide retail deposits). Enrollment increased to 67% of banks by the end of 2005:Q2 that had retail deposits comprising more than 99%. Note that not all banks have retail deposits and that firm deposits are not covered by the insurance program. Calculations are based on data from the Deposit Insurance Agency ("Sostoyaniye Rynka Vkladov Grazhdan v 2005 Godu", Agentstvo po Strahovaniyu Vkladov, 2006), CBR and Interfax.

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 could affect our results the most- is still observed. Also, the substantially different time patterns in the responses of deposits and interest rates in Figure 1 also suggests that this effect, if at all present, does not have a quantitatively important impact.

Third, whether the reduced form is estimated using Ordinary Least Squares, a Fixed Effects, a General Method of Moments, Mean-group or Swamy estimator has little effect 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 one likely factor contributing to such stability. We have also considered different specifications, including heteroscedasticity across groups, different cross-group reduced form interactions, and more. None of these affected our baseline results.

Finally, concerning inference, the baseline confidence bands ignore the fact that the reduced form estimation is based on panel data, treating the reduced form as if it were estimated on (group-wise) aggregated data. This procedure substantially overstates the width of the confidence bands. Experiments which take into account the additional cross-sectional dimension confirm this. In particular, in addition to the baseline results, we re-run the procedure, adjusting the degrees of freedom in drawing confidence bands to take account of both the n and T dimensions, rather than just T. Moreover, we also perform a bootstrap exercise, performing identification on the set of bootstrapped reduced forms, and constructing confidence bands based on these draws. Each variant invariably narrows the confidence bands drawn in the baseline results, with very similar median estimates. All conclusions remain, and typically turn out to be much more significant than in the (conservative) baseline estimates.

4 Concluding remarks

We propose a cross-sectional approach to standard macroeconometric methods. Applying our method to Russian deposit market data suggests there was one bank run during the sample period 2002-2007, which is in line with narrative evidence. Our approach has the advantage that it allows controling for effects that go hand in hand with crises and that make it difficult to disentangle the effect of the bank run itself. This should prove especially useful in view of analysis of data on the recent crisis. While institutional details may differ, similar heterogeneous features exist in banking markets in other countries, such as banks that have both insured and uninsured deposits (e.g. relative to a coverage limit) or too-big-to-fail institutions.

We quantify the effects of the two main theories of bank runs. For our sample in particular, we find that the panic view is much more important than the information-based view. While we do find evidence that the fundamental view matters, its aggregate effects are always 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-induced bank runs are a real concern. This implies that purely fundamentals-based regulation is not a panacea. While our conclusion may seem to sit awkwardly with the literature establishing the importance of market discipline in deposit markets (e.g. Flannery, 1998), it does not. These studies show how bank fundamentals determine depositor behavior, while our results may seem to suggest otherwise. One crucial difference is that our results are *conditional* on a bank run, whereas the market discipline result is an unconditional one. While it is certainly useful from a regulatory point of view to know that depositors punish (reward) banks for bad (good) behavior in normal times, it is quintessential to acknowledge that they may not make that distinction during a financial crisis.

From a methodological perspective, our approach can serve to take macro models with heterogeneity to the data. It thereby adds to reduced form microeconometric approaches. Our application suggests that relative to traditional macro VARs, 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).

References

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

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

[3] Blanchard, O.J., Quah, D., 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances", *American Economic Review* 79, 655-673.

[4] Binder, M., Hsiao, C., Pesaran, H., 2005. "Estimation and Inference in Short Panel Vector Autoregressions with Unit Roots and Cointegration", *Econometric Theory* 21, 795-837.

[5] Boivin, J., Giannoni, M.P., Mihov, I., 2009. "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data", *American Economic Review* 99, 350-384.

[6] Calomiris, C., Gorton, G., 1991. "The Origins of Banking Panics, Models, Facts, and Bank Regulation", in *Financial Markets and Financial Crises,* ed. by G. Hubbard. Chicago: University of Chicago Press.

[7] Calomiris, C., Mason, J., 2003a. "Consequences of Bank Distress during the Great Depression", *American Economic Review* 93, 937-947.

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

[9] Calomiris, C., Wilson, B., 2004. "Bank Capital and Portfolio Management: The 1930s 'Capital Crunch' and the Scramble to Shed Risk", *Journal of Business* 77, 421-455.

[10] Canova, F., Ciccarelli, M., 2009. "Estimating Multicountry VAR Models", *International Economic Review* 50, 929-961.

[11] Canova, F., De Nicoló, G., 2002. "Monetary Disturbances Matter for Business Fluctuations in the G-7", *Journal of Monetary Economics* 49, 1131-1159.

[12] Canova, F., Pappa, E., 2006. "Does it Cost to be Virtuous? The Macroeconomic Effects of Fiscal Constraints", in NBER International Seminar in Macroeconomics 2004, 11065.

[13] Chamberlain, G., 1983. "Panel Data", Chapter 22 in *Handbook of Econometrics* II, ed. by Z. Griliches and M. Intrilligator. Amsterdam: Elsevier Science.

[14] Chari, V.V., Jagannathan, R., 1988. "Banking Panics, Information, and Rational Expectations Equilibrium", *Journal of Finance* 43, 749-761.

[15] Christiano, L.J., Eichenbaum, M., Evans, C.L., 1999. "Monetary Policy Shocks: What Have We Learned and to What End?", Chapter 2 in *Handbook of Macroeconomics* IA, ed. by J.B. Taylor and M. Woodford. Amsterdam: Elsevier Science.

[16] Cooper, R., Corbae, D., 2002. "Financial Collapse: A Lesson from the Great Depression," *Journal of Economic Theory*, 107, 159-190.

[17] Covitz, D.M., Liang, N., Suarez, G.A., 2009. "The Evolution of a Financial Crisis: Panic in the Asset-Backed Commercial Paper Market", Federal Reserve Board FEDS 36.

[18] Danthine, J.-P., Donaldson, J.B., 2002. "Labor Relations and Asset Returns." *Review of Economic Studies*, 69, 41-64.

[19] Degryse, H., Kim, M., Ongena, S., 2009. *Microeconometrics of Banking: Methods, Applications, and Results.* New York: Oxford University Press.

[20] Demirgüç-Kunt, A., Huizinga, H., 2004. "Market Discipline and Deposit Insurance", *Journal of Monetary Economics* 51, 375-399.

[21] Diamond, D.W., Dybvig, P.H., 1983. "Bank Runs, Deposit Insurance, and Liquidity", *Journal of Political Economy* 91, 401-419.

[22] Ennis, H.M., 2003. "Economic Fundamentals and Bank Runs", Federal Reserve Bank of Richmond *Economic Quarterly* 89, 55-71.

[23] Ennis, H.M., Keister, T., 2009. "Bank Runs and Institutions: The Perils of Intervention", *American Economic Review* 99, 1588-1607.

[24] Faust, J., 1998. "The Robustness of Identified VAR Conclusions about Money", *Carnegie-Rochester Conference Series on Public Policy* 48, 207-244.

[25] Fernández-Villaverde, J., Rubio-Ramírez, J.F., Sargent, T.J., Watson, M.W., 2007. "ABCs (and Ds) of Understanding VARs", *American Economic Review* 97, 1021-26.

[26] Flannery, M.J., 1998. "Using Market Information in Prudential Bank Supervision: A Review of the U.S. Empirical Evidence", *Journal of Money, Credit and Banking* 30, 273-305.

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

[28] Franco, F., Philippon, T., 2007. "Firms and Aggregate Dynamics", *Review of Economics and Statistics* 89, 587-600.

[29] Galí, J., 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?", *American Economic Review* 89, 249-271.

[30] Garratt, R., Keister, T., 2009. "Bank Runs as Coordination Failures: An Experimental Study", *Journal of Economic Behavior & Organization* 71, 300-317.

[31] Gertler, M., Gilchrist, S., 1994. "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms", *Quarterly Journal of Economics* 109, 309-340.

[32] Goldstein, I., Pauzner, A., 2005. "Demand-Deposit Contracts and the Probability of Bank Runs", *Journal of Finance* 60, 1293-1327.

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

[34] Gorton, G.B., Winton, A., 2003. "Financial Intermediation", Chapter 8 in *Handbook of the Economics of Finance* I, ed. by G.M. Constantinides, M. Harris and R.M. Stulz. Amsterdam: Elsevier Science.

[35] Greenwood, J., Hercowitz, Z., Krusell, P., 1997. "Long-Run Implications of Investment-Specific Technological Change", *American Economic Review* 87, 342-362.

[36] Hellmann, 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.

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

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

[39] Iyer, R., Peydró, J.-L., 2010. "Interbank Contagion at Work: Evidence from a Natural Experiment", *Review of Financial Studies*, forthcoming.

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

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

[42] Karas, A., Schoors, K., 2005. "Heracles or Sisyphus? Finding, Cleaning and Reconstructing a Database of Russian Banks", mimeo.

[43] Kelley, M., Ó Gráda, C., 2000. "Market Contagion: Evidence from the Panics of 1854 and 1857", *American Economic Review* 90, 1110-1124.

[44] Khwaja, A. I., Mian, A., 2008."Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market", *American Economic Review* 98, 1443-1442.

[45] Krusell, P., Smith, A.J., 1998. "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy* 106, 867-896.

[46] Levin, A.T., López-Salido, J.D., Nelson, E., Yun, T., 2008. "Macroeconometric Equivalence, Microeconomic Dissonance, and the Design of Monetary Policy", *Journal of Monetary Economics* 55, S48-S62.

[47] Madiès, P., 2006. "An Experimental Exploration of Self-Fulfilling Banking Panics: Their Occurrence, Persistence, and Prevention", *Journal of Business* 79, 1831–1866.

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

[49] 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, 1029-1052.

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

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

[52] Peck, J., Shell, K., 2003. "Equilibrium Bank Runs", *Journal of Political Economy* 111, 103-123.

[53] Peersman, G., 2009. "The Relative Importance of Symmetric and Asymmetric Shocks: The Case of United Kingdom and Euro Area", mimeo.

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

[55] Rubio-Ramirez, J., Waggoner, D., Zha, T., 2010. "Markov-switching Structural Vector Autoregressions: Theory and Application", *Review of Economic Studies*, forthcoming.

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

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

[58] Schumacher, L., 2000. "Bank Runs and Currency Run in a System without Safety Net: Argentina and the 'Tequila' Shock", *Journal of Monetary Economics* 46, 257-277.

[59] Schotter, A., Yorulmazer, T., 2009. "On the Dynamics and Severity of Bank Runs: An Experimental Study", *Journal of Financial Intermediation* 18, 217-241.

[60] Sims, C., 1980. "Macroeconomics and Reality", *Econometrica* 48, 1-48.

[61] Sims, C., 1998. Comment on "Do measures of Monetary Policy in a VAR Make Sense?", *International Economic Review* 39, 933-941.

[62] Smets, F., Wouters, R., 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach", *American Economic Review* 97, 586-606.

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

[64] 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.

[65] Uhlig, H., 2009. "A Model of a Systemic Bank Run", *Journal of Monetary Economics* 57, 78-96.

[66] Zykova, T., 2004. "Chernyi spisok Zubkova" (Zubkov's black list), *Rossiyskaya Gazeta* 3478, 18 May 2004.

Figure 1: Impulse Responses to a Bank Run (1 std. impulse)

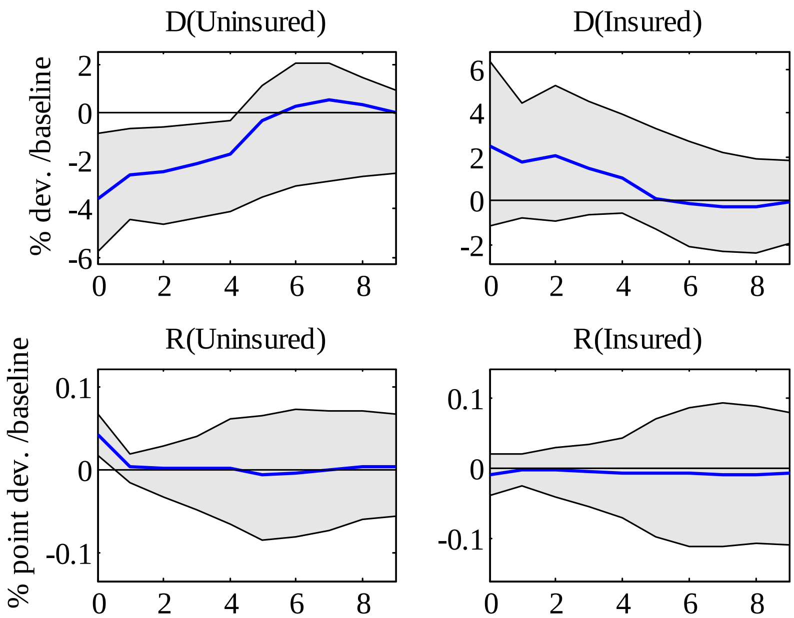

Figure 2: A Time Series of the Identified Shock

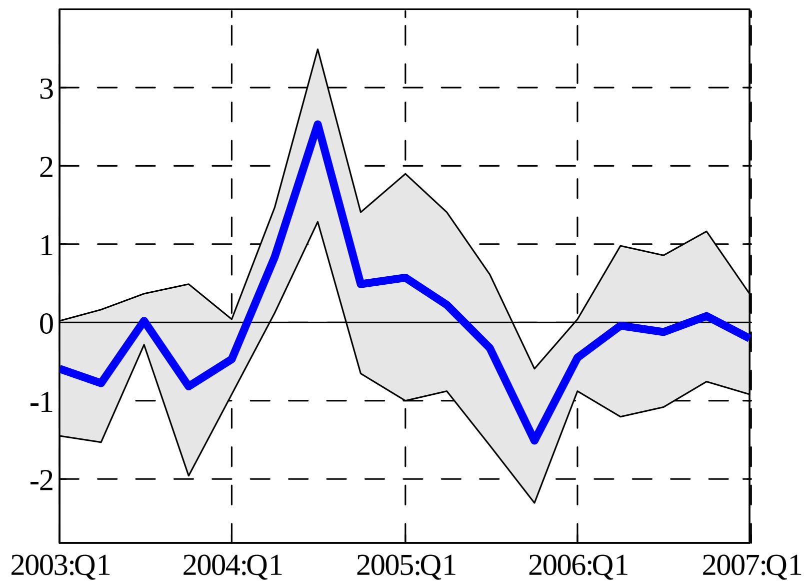

Figure 3: Response to the 2004 Run: Deposits

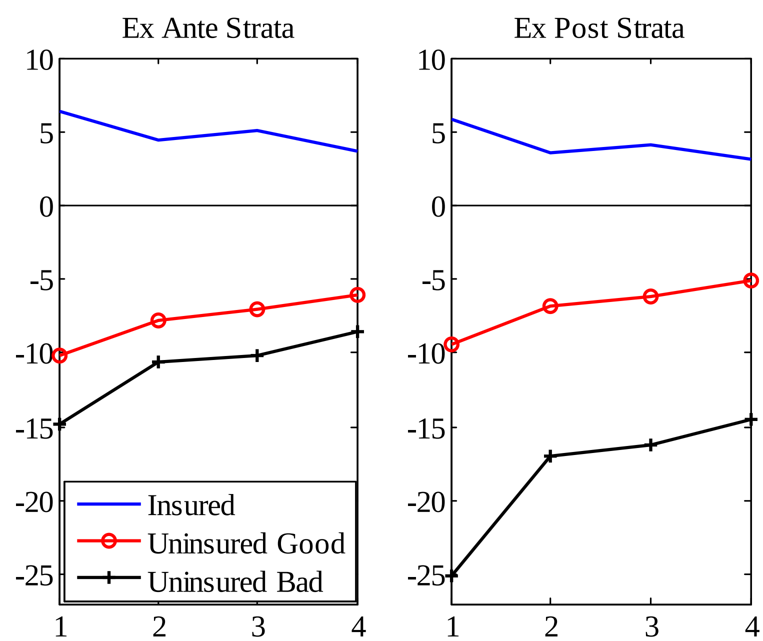

Table 1: Response to the 2004 Run: Deposits

	Capital	**Liquidity**	**Ex ante**	**Ex post**
Panel A: Impact				
Bad banks	-0.12***	-0.13***	-0.15***	-0.25***
	(0.01)	(0.02)	(0.01)	(0.08)
Good banks	-0.15***	-0.14***	-0.11***	-0.10***
	(0.02)	(0.01)	(0.01)	(0.01)
Insured banks	0.06	0.06	0.07	0.06
	(0.04)	(0.04)	(0.04)	(0.04)
Panel B: Cross-sectional tests				
Fundamental: $Bad - Good$	0.03	0.01	-0.04**	-0.15**
H_0: No Fundamental effect (p-value)	0.96	0.62	0.02	0.03
Panic: $Good - min(0, Insured)$	-0.15***	-0.14***	-0.11***	-0.10***
H_0: No Panic effect (p-value)	0.00	0.00	0.00	0.00
Panel C: Aggregate effects				
Outflow uninsured[a]:	-13.4	-13.3	-13.3	-11.3
* due to fundamentals[b]	0%	0%	15.3%	13.8%
* due to panic[c]	100%	100%	84.7%	86.2%
* absorbed by insured[d]	1.7%	1.3%	1.5%	2.4%

Note: *** (**,*) significant at the 1% (5%, 10%) level. Standard errors in parenthesis. Let capital letters B, G, I denote the impact coefficients from Panel A, for bad, good and insured, respectively, and N_B, N_G, and N_I the respective volumes of deposits prior to the run. Then, [a] percentage change in total volume of uninsured deposits is calculated as $(BN_B + GN_G)/(N_B + N_G)$, which is then decomposed into [b] a fundamental part: $(B - G)N_B/(BN_B + GN_G)$, and [c] a panic-driven part: $G(N_B + N_G)/(BN_B + GN_G)$. The inflow at insured banks [d] is $(IN_I)/(BN_B + GN_G)$.

Table 2: Response to the 2004 Run: Interest Rates

	Capital	**Liquidity**	**Ex ante**	**Ex post**
Panel A: Impact				
Bad banks	0.13***	0.16***	0.19***	0.38***
	(0.02)	(0.02)	(0.03)	(0.12)
Good banks	0.15***	0.14***	0.12***	0.09***
	(0.03)	(0.03)	(0.03)	(0.02)
Insured banks	-0.01	-0.01	-0.01	0.02
	(0.03)	(0.03)	(0.03)	(0.03)
Panel B: Cross-sectional tests				
$Bad - Good$	-0.02	0.02	0.07**	0.29**
H_0: $Bad \geqslant Good$ (p-value)	0.76	0.33	0.03	0.01
$Good - \max(0, Insured)$	0.15***	0.14***	0.12***	0.07**
$H_0 : Good \geqslant \max(0, Ins)$ (p-value)	0.00	0.00	0.00	0.01

Table 3: Alternative stratifications, identification and robustness

	Ex ante stratification				**Ex post stratification**		
	(I) Restrict Bad	**(II) Restrict Good**	**(III) Moscow**	**(IV) B⩾75 G⩽25**	**(V) Restrict Bad**	**(VI) Restrict Good**	**(VII) Moscow**
DEPOSITS							
Bad banks	-0.14***	-0.13***	-0.20***	-0.17***	-0.22***	-0.23***	-0.20**
Good banks	-0.09***	-0.11***	-0.11***	-0.14***	-0.07***	-0.13***	-0.10***
Insured banks	0.08*	0.08*	0.12**	0.08*	0.06	0.08*	0.13**
Fundamental	-0.05**	-0.02*	-0.09**	-0.03	-0.15**	-0.10*	-0.10
Panic	-0.09***	-0.11***	-0.11***	-0.14***	-0.07***	-0.13***	-0.10***
Outflow uninsured:	-11.8	-11.9	-15.1	-15.6	-8.3	-13.8	-11.6
* fundamentals	20.3%	10.0%	28.8%	8.5%	18.8%	7.6%	0%
* panic	79.7%	90.0%	71.2%	91.5%	81.2%	92.4%	100%
* to insured	1.9%	2.0%	3.2%	1.9%	3.3%	2.7%	6.8%
INTEREST RATES							
Bad banks	0.19***	0.15***	0.23***	0.19***	0.38***	0.24*	0.33***
Good banks	0.10***	0.12***	0.12***	0.09***	0.06***	0.14***	0.09***
Insured banks	-0.03	-0.01	-0.04	-0.01	0.02	-0.01	-0.04*
$Bad - Good$	0.09**	0.03	0.11**	0.10**	0.32***	0.10	0.24**
$Good - \max(0, Insured)$	0.10***	0.12***	0.12***	0.09**	0.04***	0.14***	0.09***

Note: Rows: see Table 1. Columns: Restrict bad (good) = identification is based on bad (good) subgroup of all uninsured banks, relative to insured banks; Moscow = only banks located in Moscow included; B⩾75 G⩽25 = banks with default probability above the 75th percentile classified as bad, below 25th as good, while middle two quartiles are dropped. Note that aggregate effects for Moscow and quartile stratifications pertain to a subset of the total uninsured market.

Table 4: Further cross-sectional differences

Foreign banks				
	D(insured)	D(foreign)	p-value on difference	
	0.03 (0.04)	0.01 (0.04)	0.40	
Fundamentals insured banks				
	D(insured bad)	D(insured good)	p-value on difference	
* Capital	0.05* (0.03)	0.14* (0.07)	0.87	
* Liquidity	0.09** (0.04)	0.11 (0.09)	0.56	
Depositor characteristics				
	T^{FUND}	p-value	T^{PANIC}	p-value
* Ex ante	-0.01	0.69	-0.19***	0.00
* Ex post	0.09**	0.01	-0.35**	0.01

Note: *** (**,*) significant at the 1% (5%, 10%) level. Standard errors in parenthesis.

Appendix A: Data

The bank-specific variables used in our analysis include deposits and interest rates as well as measures of risk, performance and balance sheet structure. Quarterly data on bank balance sheets and income statements is obtained from two established private financial information agencies, Interfax and Mobile, and covers the period from 1999 till 2007.[1] The average implicit interest rate that a bank offers on its deposits is calculated by dividing interest expenses by the corresponding level of deposits. 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, 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 a new entity. 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 stable over our sample period. Figure A shows that the growth rates of consumer deposits in both insured and uninsured banks are comparable through the major part of our sample period. As expected, uninsured deposits generally pay higher interest.

Figure A: Deposit growth and interest rates: 25th, 50th and 75th percentiles (in %)

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

Appendix B: Default prediction model

Table B contains the estimated logit for our full sample. The ex ante stratification in the paper is based on a recursive estimate of the same specification, where the estimate is updated every period.

Table B: 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[2]	0.867

[2] 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).

Appendix C: Impulse responses

For the different stratifications used, we here plot the (confidence bands on) impulse responses to a 1 std. bank run. Identifying restrictions pertaining to the uninsured group are imposed on the average across good and bad.