REAL-TIME IDENTIFICATION AND HIGH FREQUENCY ANALYSIS OF DEPOSITS OUTFLOWS

Edoardo Rainone*

January 2021

Abstract

We propose a method based on control charts to identify in real-time sudden deposits outflows through the payment system. The performance of the methodology is assessed with both Monte Carlo simulations and real transaction-level TARGET2 data for a large sample of Italian banks. We identify a set of idiosyncratic bank stress episodes. Using high frequency payment system data, we provide new evidences on the interaction between retail, wholesale and central bank funding in the post global financial crisis period.

Keywords: bank runs, interbank networks, payment systems, money, control charts, digital economy, financial stability.

JEL Classification Codes: E50, E40, G01, G10, G21.

*Directorate General of Economics, Research and Statistics, Bank of Italy, Via Nazionale 91, 00184 Rome, Italy. email: edoardo.rainone@bancaditalia.it. I thank Salvatore Alonzo, Elena Andreou, Paolo Bramini, Luisa Carpinelli, Luigi Chiusolo, Francesco Columba, Domenico Depalo, Alessio De Vincenzo, Giorgio Gobbi, Giovanni Guazzarotti, Kasperi Korpinen, Robin Lumsdaine, Silvia Magri, Gaetano Marseglia, Alexander Muller, Fabrizio Palmisani, Niki Papadopoulou, Giandomenico Scarpelli, Anatoli Segura, Massimiliano Stacchini, Massimo Valentini, Stefano Vespucci, for their comments as well as participants at the Bank of Italy seminar, the 15th Bank of Finland Simulator seminar, and the IAAE 2019 conference. All the errors are my own. The views expressed in the paper are solely those of the author and do not necessarily represent those of the Bank of Italy.

1 Introduction

When depositors withdraw their funds from a bank and move them to another one or convert them in cash, they leave digital footprints in the payment system. The recent development of financial and payment services, the fintech revolution and the spread of smartphones, mobile devices and financial applications can drastically transform the speed with which changes in depositors' trust materialize and thus the stability of various funding sources, affecting the liquidity indexes under Basel III. Despite the importance of this measurement, evidence on the subject is scarce.

Along these lines, this paper's contribution is twofold. The main contribution is methodological. We propose a method based on control charts that can be used by supervisory and resolution authorities (in a SupTech perspective), commercial and central banks to monitor depositors' behavior in real-time and save relevant social and private costs. A control chart is a statistical tool, invented by Walter A. Shewhart while working for Bell Labs in the 1920s (Shewhart, 1926), designed to monitor the quality of industrial products in the continuum and to real-time detect anomalies in the production process. Upper and lower bounds of the chart define when the process is under control. When the output diverges from its normal behavior and appears more frequently outside the bounds, the chart signals that the process is not under control. We adapt this statistical tool to be used effectively with payment and financial data. The methodology proposed can be extended to monitor other type of balance sheet items, like wholesale funding (Gertler et al., 2016; Martin et al., 2014). First, we show how payment system data can be used to measure deposits' flows, second we design a method based on control charts to track sudden changes in real-time, and third we test its performance using simulated and real data. Importantly for financial stability, our methodology can provide information on potential contagion and systemic effects in real-time, when applied to the entire banking system.

Our second contribution is in terms of economic insights. We apply the methodology to unique real data from TARGET2, the pan-European large value payment system. Analyzing the Italian banking system during the post global financial crisis period, from August 2012 to August 2019, we identify and characterize a set of distress episodes. The estimated average length of distress episodes is about four weeks. The average liquidity drain from the distressed bank is significant and equal to about 3 percent of the bank's deposits, close to the run-off rate of stable deposits in the liquidity cover ratio (LCR, Basel III, 2013). Combining the output of our methodology with high frequency payment data, we uncover novel features of depositors' behavior during distress episodes, complementing the empirical micro literature which uses aggregate deposit data (Oliveira et al., 2014) and granular deposit accounts data (Iyer and Puri, 2012; Iyer et al., 2016). Unexplored information about the type of digital transfers tells us that about half of them were done using real-time settlement (instead of the usual deferred settlement), putting an additional liquidity pressure on the stressed bank.

Differently from other papers, the relational nature of payment data allows us to study the destination of outflows. Deposits are almost entirely directed to other banks, and not massively converted in cash, leaving almost unchanged the deposit currency ratio. Contrary to the usual negative contagion effects found in the literature (see Goldsmith-Pinkham and Yorulmazer, 2010, for example), we provide evidence of positive spillovers (deposits inflows) to other banks. Deposits from distressed banks shift mainly to large domestic banks. We are also able to analyze deposits that moved across the border, and provide evidence that only a small portion of funds moved abroad, which is reassuring from a (national) financial stability perspective. In addition, analyzing the full banking system we are able to inform on contagion and systemic effects, for which we do not find any evidence in the episodes identified.

As payment data gives us full visibility on banks' liquidity, we can also study the interaction between retail and wholesale funding with high frequency data for the first time. Interestingly, we find that outflows in retail deposits are not exacerbated or anticipated by wholesale funding. Oppositely to unsecured money market gridlocks (e.g. Freixas et al., 2000) and the shortage of interbank liquidity observed during the global financial crisis (Afonso et al., 2011; Angelini et al., 2011; Heider et al., 2015), we show that the reliance on secured money markets and central bank funding, implied by the crisis and the subsequent regulation and policies, prevented a sudden dry up of wholesale funding and provided immediate funding during distress episodes in which even retail depositors shifted away. Under the fixed-rate full allotment regime, banks offset the liquidity outflow mainly through open market operations. These results highlight the importance of having high quality collateral at disposal to quickly obtain liquidity, and challenge the idea that "modern" mechanisms of bank run concern wholesale funding only. On the contrary, the increased stability of wholesale funding (due to the shift from unsecured to secured) and increased instability of retail funding (due to technological innovations in customers' transfers and deposits) can change their relative volatility and risk from the standpoint of the aggregate funding of a bank.

The rest of the paper is organized as follows. Section 2 describes how to track deposits' flows with payment system data and the methodology proposed to identify distress episodes in real-time. Section 3 shows the performance of the algorithm with simulated and real TARGET2 data. Section 4 presents economic insights from the distress episodes identified. Section 5 concludes. In the reminder of this section, we relate our paper to the relevant literature.

1.1 Related Literature

Our paper is related to a number of strands of the literature. Firstly, the literature on the use of payment system data/information to design early warning systems. Secondly, the micro empirical literature concerned with depositors' behavior; thirdly, the large literature focused on liquidity shocks, wholesale funding and policy tools to prevent distress episodes.

Our paper relates to the literature using payments data, especially large value payment system (LVPS) data, to study interbank liquidity. Among the others, Denbee et al. (2014) describe methods for measuring liquidity provision using data from CHAPS, the UK LVPS. Denbee et al. (2021) use a spatial error model to estimate the liquidity multiplier of an interbank network and banks' contributions to systemic risk. Heijmans et al. (2014) show how large payment data provide early warning information, with case studies from the Dutch part of the Eurosystem's LVPS, TARGET2. Heijmans and Heuver (2014) develop a battery of indicators for signs of liquidity shortages of banks using TARGET2 transaction data and collateral management data. Rosati and Secola (2006) analyze the distribution of LVPS cross-border interbank payment flows using gravity models. Soramäki et al. (2007) analyze the connectivity of Fedwire, the American LVPS, describing the network topology of the interbank payment flows and the response of such complex network to pertubations, such as September 11th. Craig et al. (2018) analyze the behavior of banks with respect to the settlement of interbank claims and focus on payment delay. They use a instrumental variable approach to analyze whether delays in incoming transactions could cause delays in outgoing transactions. All these papers focus mainly on wholesale or aggregate liquidity, our paper differs from them because it (i) is the first to use LVPS data to detect and study retail depositors behavior and (ii) develops a new early warning system using payment data on control charts.

By studying deposits outflows with payment system data, we complement the existing empirical micro literature in banking and finance which is concerned with depositors' behavior and perception (Chen and Hasan, 2006; Keister and Narasiman, 2016; Kiss et al., 2014). Our results on a *size premium* in idiosyncratic shocks during not-turbulent times complements the exiting evidences on turbulent ones (Oliveira et al., 2014). Iyer et al. (2016) and Iyer and Puri (2012) look at micro data from one bank in India and show how heterogeneity in depositors' responses to solvency risk and in bank-depositors relationships can generate different type of distress. Our study is somehow complementary as we look at payments instead of accounts and consider an environment with mainly digital instead of physical banking. While the advantage of accounts data is that it is possible to observe heterogeneous behaviors across types of depositors, with payment data we can understand whether deposits are converted in cash or moved to another bank, and better assess contagion by observing transfers of clients of other banks, in real-time. These aspects are particularly important to understand (i) if idiosyncratic shocks turn to systemic, (ii) what is the preferred safe option for depositors, (iii) if depositors accelerate deposits outpouring using new technologies, like instant payments.

By analyzing the interaction between retail, wholesale and central bank funding, we relate to a large literature that looks at financial and money markets during distress episodes (see Brown et al., 2016; Cañón and Margaretic, 2014; Cooper and Ross, 1998; Ennis and Keister, 2006; Ferrante, 2018; Gertler and Kiyotaki, 2015; Gertler et al., 2016; Waldo, 1985, among others). Banks that rely on funding from wholesale markets may be significantly affected by a crisis of another bank (Goldsmith-Pinkham and Yorulmazer, 2010). After the global financial crisis, the literature has focused on a so-called "modern" mechanisms of bank run, which is concerned with wholesale funding, especially in the interbank segment where other banks exert monitoring and market discipline, rather than customers' deposit withdrawals that are typically thought to be sticky. By providing new evidences on the role of wholesale funding and central bank facilities during idiosyncratic distress episode, we complement the large empirical literature developed after the global financial crisis, which is mostly focused on macro shocks and unsecured interbank markets (see Afonso et al., 2011; Angelini et al., 2011; Heider et al., 2015, among the others). From this perspective, the paper provides evidences to the existing theories predicting how the effects of an idiosyncratic liquidity shock are different from those of an aggregate shock, and how central bank interventions impact them (e.g. Allen et al., 2009; Brunetti et al., 2010; Garcia-de Andoain et al., 2016), in the post-crisis policy framework. Our results show that these tools are effective in smoothing out bank-specific liquidity shocks in the post-crisis environment, in which the unsecured money market almost disappeared in favor of the secured one. Our evidences can also be informative for the large theoretical literature on policy tools to prevent distress episodes (see Bryant, 1980; Dávila and Goldstein, 2020; Diamond and Dybvig, 1983; Engineer, 1989; Ennis and Keister, 2009; Gropp et al., 2010; Keister and Narasiman, 2016; Martin, 2006, among the others).

2 Identifying Deposits Outflows in Payment Systems

In this section we describe how to identify deposits outflows using payment system data. In what follows, the term deposits refers to non financial firms and households, and not interbank claims. Before getting into the details of our algorithm, we give a brief and simplified overview of payment systems and the related data.[1]

2.1 How to Track Deposits' Flows in the LVPS

If a customer of a bank wants to withdraw her deposits, she has two options: (i) withdraw them in cash, (ii) transfer them to another bank.[2] When the deposits are transferred to another bank (liabilities decrease), they have to move from the reserve account of their bank to another one (assets decrease) at a certain point in time. Section B in the online appendix describes in detail the relationship between deposits and reserves. In modern financial systems, interbank positions are settled in central bank money on the LVPS, they can be settled directly or through a retail payment system (RPS). The latter aggregates many single transactions, calculate interbank exposures and finally sends the position of each bank to the LVPS, netting the gross positions. It follows that a transaction reaches a final settlement in central bank money following different routes, depending on its nature and customers' preferences on the speed of settlement. If customers prefer a real-time settlement, they can buy this service from the debtor's bank and the payment is sent directly to the LVPS, otherwise the payment is settled through

[1] See Kokkola (2010) and Haldane et al. (2008) among others for a more detailed description of payment and settlement systems.

[2] Here we assume that the depositor does not want to convert the deposits into something different and potentially illiquid, like securities, gold or other financial activities different from cash or deposits, which can be used also for payments. Even in this case, if the account of the counterparty is in a different bank we will observe it anyway.

the RPS. Usually this process ends with a delay of one day w.r.t. a real-time settlement, which settles in the same day in which the client sends the payment instruction.

Nowadays, there is also the possibility of using 'instant payments', that are retail payments settled in few seconds. Banks have to dedicate part of their reserves or some collateral to pre-fund them. These systems can settle basically the same type of transactions of classic RPSs plus mobile and peer-to-peer instant transfers.With a classic RPS it takes up to one business day for a payment in euro to reach the beneficiary. With instant payments, the funds are available immediately (in the order of seconds) for use by the recipient, 24/7/365. At affordable prices, this feature is appealing for depositors in good and bad times. But it would be particularly attractive, even at relatively high prices, if they want to move their deposits immediately because they think that the bank is going to fail soon. We may call this situation a sort of 'instant distress'.

Also an higher demand for paper cash is visible in the LVPS. Holding constant its inventory, if customers want to transform their deposits in cash, the bank has to get the banknotes from the central bank. The latter gives the cash to the bank and debits its reserve account in the LVPS.

Figure A.1 in the online appendix gives a simplified view on how customers can move their deposits in a stylized payment system. As a result, the LVPS represents an optimal perspective from which monitoring and studying customers' behavior .If this information is promptly available, it can timely detect difficulties of a bank before the stability of the entire financial system is threatened.

In Section 3.2, we use data from TARGET2 to track deposits outflows. TARGET2 is the LVPS system owned and operated by the Eurosystem, operated under real time gross settlement (RTGS) basis. More than 1,700 banks use TARGET2 to initiate transactions in euro, either on their own behalf or on behalf of their customers. Taking into account branches and subsidiaries, more than 55,000 banks worldwide (and all their customers) can be reached via TARGET2. A key feature of TARGET2 is that customer payments can be isolated from interbank payments. This is important in order to distinguish between banks versus non-banks deposit holders, and thus between wholesale versus retail deposit outflows, as we are interested in the latter here. Unfortunately we have no information on the final receiver of the funds, the purpose of the transaction and we can not distinguish between non-financial firms versus households. To have such information we should obtain proprietary data from all the commercial banks and match them with payments data, without a unique ID, an exercise that is impossible in the current institutional environment.[3] In Section C of the online appendix we detail the information on customer payments available in TARGET2.

[3]Even if it was possible, the purpose of the transaction is something that is not always known even to the bank itself. It follows that part of the outflows may also be generated by the conversion of deposits in assets, bitcoins, metal or gemstones, that imply an interbank settlement. Assuming that such amount is limited implies that customers mostly do not want to change drastically the nature of their deposits and keep holding the most liquid assets with a stable value. Another limitation of our data is that we can not identify within banking group transfers. This seems to be a minor issue for our scope given that it is unlikely that customers move deposits within the same banking group in case of distress.

2.2 Methodology

From a methodological perspective, the goal of our analysis is twofold. First, we want to identify distress episodes from payment system data. Second, we want to do it in real-time. The first problem consists in finding a structural break in depositors behavior with high frequency data. The second problem is more peculiar and close to a early warning method. On this topic, the economics and finance literatures have developed several methods mostly for banking, currency and balance of payments crises. A not fully exhaustive list of methods used in this context includes the minimization of the noise-to-signal ratio (Kaminsky and Reinhart, 1999), fully-parametric (Logit/Probit) models (Berg and Pattillo, 1999), semiparametric models (Arduini et al., 2012), classification trees and random forest (Alessi and Detken, 2017).

In all these approaches the econometrician (i) observes ex-post $W_{ct} = 1$ if a crisis occurred at time t in country c, (ii) chooses a set of covariates X_{ct} that should predict a specific type of crisis, (iii) specifies $f(X_{ct})$ as a function of the observables, (iv) selects a crisis threshold τ for f and (v) predicts a crisis if $f > \tau$. Observe that in this case it is also possible to assess the quality of the method and possibly type I and II errors.

Unfortunately, our problem is different from a classical early warning exercise. Usually the crisis is known and researchers try to use some information to predict the event, based on previous cases. For idiosyncratic distress episodes there may be no evidence for two reasons: (i) the deposits can fly away and get back in between two observations of supervisory data if they are not very frequent,[4] (ii) the digital component does not allow to physically observe lines at the bank branches.[5] See Section D of the online appendix for a discussion on the tempestivity of supervisory reports. In practice, episodes are not necessarily public knowledge. It means that we have to uncover them. In this sense, the exercise is substantially different. The coexistence of these two goals, identifying episodes and doing it in real-time, makes our task more difficult.

To achieve these goals, we propose to adapt a statistical tool that is very popular in industrial production: the control charts. The control chart was invented by Walter A. Shewhart while working for Bell Labs in the 1920s (Shewhart, 1926). Shewhart framed the problem in terms of common- and special-causes of variation. The control charts is now the most used tool to control industrial production processes. The tool was designed to monitor the quality of products in the continuum and to real-time detect anomalies in the production process. The aim was to minimize the cost of the production of wrong pieces. Here the logic is close, as we are interested in minimizing the social costs of a distress episode by identifying it as soon as possible.[6]

[4] System-wide crises are usually longer and thus rarely invisible to supervisory data, if the latter is reasonably frequent.

[5] This feature characterizes also wholesale funding idiosyncratic crises. If the wholesale funding crisis is systemic, market indicators can send timely signals.

[6] An alternative to the Shewhart control chart would be the cumulative sum (CUSUM) control chart. The cumulative sum in this type of chart is the sum of deviations of individual sample results from the target. See Andreou and Ghysels (2009) for a review of this and other type of sequential tests used in financial time series. CUSUM charts are more sensitive to small and temporary shifts, which is less desirable in our context. Note that if there is a true shift in the

In practical terms, a basic control chart consists of points representing a statistic (y_t) that measures a characteristic of a sample taken from the production process over time. This statistic must be close to the 'desired value' (y^*), which is given by the industrial process. It constitutes the center line of the chart. The mean and the standard deviation of the statistic (σ_y) are calculated from the process under control -i.e. without anomalies- and used to determine upper and lower control limits (respectively UCL and LCL) that indicate the thresholds at which the process output is considered statistically unlikely and are drawn typically at 3 standard deviations from the center line, under normality assumption. By the way UCL and LCL are constructed, observations can exceed these limits in rare cases if the process is under control, for example $P(y^* - 3\sigma_y = LCL < y_t < ULC = y^* + 3\sigma_y) = 0.997$ of the observations occur within 3 standard deviations of the mean. This approach is also called the 'sigma approach'. In Figure 1 we show how simulated process under control (panel (a)) and not under control (panel (b)) look like. For some examples and a review see Lowry and Montgomery (1995).[7]

Figure 1: Processes Under and not Under Control

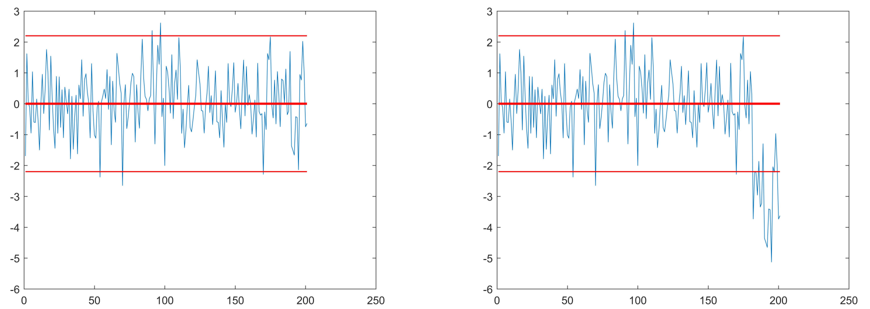

(a) Under Control (b) Not Under Control

Notes. Time series of simulated data from a normal distribution. x-axis: days. y-axis: deposits net outflows. The DGP is from our pivotal simulation setting, as described in Section 3.1. The red bold line represents the expected value of the process, the red light lines track the UCL and LCL, here set equal to +/-2.2 σ. The last 20 observations of the series 'Not Under Control' are shocked with expected outflow equal to 3σ, where σ is the standard deviation of the simulated normal process.

For our problem, control charts have the following appealing features: (i) they are designed to be applied on real-time data (as in RTGS systems), without any ex-post and known definition of previous crises (W_{ct}); (ii) they allow us to detect timely 'special-causes' with high-frequency and firm-specific data.Nevertheless, they also have some undesired features: (i) they are designed for controlled processes of a specific firm with very standardized outputs, so the 'desired value' (y^*) is given; (ii) they rely on quite strong distributional assumption (usually normality) and (iii) do not consider seasonality and other common and problematic features of financial time series.

These are substantial limitations in our context for the following reasons. Customer payments

process average, the Shewhart chart will pick it up eventually. Observe that this problem is similar to finding a structural break in the drift. See Andreou and Ghysels (2009) for a view on historical and sequential tests.

[7] See Montgomery (1980) and Lorenzen and Vance (1986) for a discussion on the economic design of control charts. See Scheffe (1947), Nelson (1989), Iacobini (1994), Nelson (1984), Lowry et al. (1992), Roberts (1959) among others for interesting discussions, interpretations and extensions of control charts.

are not normally distributed and it is in general difficult to find realistic parametric assumptions. Banks have often structural unbalances generated by their clients' heterogeneity, some banks have more merchants than buyers among their clients or the other way around, thus they can have persistently or cyclically positive or negative net positions on customer payments.[8] It follows that the 'desired value' (y^*) is not given and its UCL and LCL are more difficult to establish.

A way to see this problem could be in terms of an omitted unobserved treatment variable. Suppose that the deposits outflows are generated by the following model,

$$y_t = f(z_t) + \tilde{y}_t. \tag{1}$$

Deposit variation is fully explained by $f(z_t)$, that captures the observable component, and a unobservable component $\tilde{y}_t$. $f(z_t)$ can be seen as the term that captures seasonality, customers composition and any other systematic source of variation. This model differs from standard early warning models described at the beginning of this section because we do not know ex ante which t is assigned to a distress period. If a shock occurs, but we can not observe it, a term adds to the random component in the unobserved term,

$$\tilde{y}_t = -d_t + \epsilon_t, \tag{2}$$

where d_t captures the additional net withdrawals if depositors do not trust their bank at time t and ϵ_t is a random component with a certain cumulative distribution function m_ϵ. When there is no shock $\tilde{y}_t = \epsilon_t$, thus we expect $P(\tilde{y}_t < \tau) = m_\epsilon(\tau)$, where τ is a threshold. If a shock occurs, i.e $d_t > 0$, then $\tilde{y}_t = -d_t + \epsilon_t$, thus we expect $P(\tilde{y}_t < \tau) = \rho > m_\epsilon(\tau)$. It follows that when a shock occurs, the probability of observing significant deposit withdrawals above the threshold τ increases. Our goal is to identify the time when the distress kicks in and d_t turns greater than zero. If the empirical frequency of $\tilde{y}_t < \tau$ increases in sequence, significantly exceeding $m_\epsilon(\tau)$ in a persistent way, we can interpret it as a shift generated by distress (d_t) and not as a single anomalous outflow generated by some banks' clients.The greater the intensity, both in terms of size and persistence, of the shock the higher the probability of detecting it. From this perspective it is also easy to see the importance of distributional assumptions. If ϵ is assumed to be normal, then we have $m_\epsilon(\tau) = F(\tau) = \frac{1}{\sigma_\epsilon\sqrt{2\pi}} \int_{-\infty}^{\tau} e^{\frac{-\epsilon^2}{2\sigma_\epsilon^2}} d\epsilon$. If this assumption is not correct, we can severely misspecify $P(\tilde{y}_t < \tau)$. What is particularly dangerous is if we overestimate $P(\tilde{y}_t < \tau)$ and set a potential critical value τ too low. In such situation we may have abnormal situations not detected. Given the high frequency of the data, we can use nonparametric methods to estimate the actual density m_ϵ when d_t is zero. In Figure A.2 in the online appendix, the blue lines report the empirical density of ϵ, which is estimated from our sample of pre-shock periods for

[8] From interbank payments data we do not observe the portfolio of assets and the actual sender/recipient of any given transaction, so outflows can be related to transfers of assets that are not related to distrust. There are many technical features and intermediary chains that make these time series less predictable, banks may change their participation in RPS and RTGS or change the way they route their payments. There is high and specific seasonality and complex idiosyncratic time patterns in cash withdrawals (like holidays and weekends) and in payments related to taxes or fiscal dates.

the channels described above (more details are provided Section 3.2). The orange lines depict instead the estimated normal distribution computed with the same sample mean and variance. While for cash withdrawals the density is very close to the normal, for digital transfers this is not the case. If we set a critical threshold to correspond with the fifth percentile of the normal distribution, we will not label as a warning a value largely below the real fifth percentile, thus not identifying timely any episode. Such evidence highlights the importance of a nonparametric step. In the next section we use simulated data to show the benefits of a nonparametric approach.

Regularized Nonparametric Shewhart Chart

To offset the undesired features of classic control charts, we introduce two important steps, one at the beginning and one at the end. The first consists in regularizing the input of the control chart (CC), the second uses nonparametric methods to identify anomalous situations in the CC framework. We name the method 'Regularized Nonparametric Shewhart Chart' (ReNoSCh).

In the first step we regularize the time series adding knowledge about the monetary phenomena under analysis:

$$\min_f \sum_{T}^{t=1} V(f(z_t), y_t) + \lambda R(f), \tag{3}$$

where V is an underlying loss function that describes the cost of predicting $f(z)$ when the label is y, such as the square loss or hinge loss; and λ is a parameter which controls the importance of the regularization term. $R(f)$ is typically chosen to impose a penalty on the complexity of f. After this step the target variable becomes $\tilde{Y} = Y - f(Z)$. [9] A very simple specification of (3) would be a simple linear model with important dummies included in Z, thus having $f(z_t) = z_t\beta$, $\lambda = 0$ and $V(.)$ equal to the squared difference. The algorithm can accommodate more sophisticated specifications with minimal changes. After having chosen the model, the target variable becomes the residual, i.e. the difference between the observed and the predicted. The resulting time series is similar to the classic CC input.

Now that we have a well-behaving series we need to define a way to assess timely when customer payments are not 'under control'. The issue here is that normality is far from reality in financial time series and especially in customer payments, as showed above. For this reason in the second step we derive nonparametrically critical thresholds and warnings, avoiding inadequate distributional assumptions. The idea is that if the monitored variable is not normal and its distribution is not known a-priori, we can nonparametrically estimate it and then assess whether the new observations are concentrating in unlikely regions. Given our interest in depositors' trust, we are worried only by a divergence towards the LCL, i.e. when the bank starts to have significant outflows of deposits. More

[9] Z can be chosen in several ways. In our practical experience, when daily payment system time series are used, it is important to add (or let the method add) many day, week and month dummies in addition to the constant, the trend and cyclical effects. An additional difficulty associated with daily time series is that there are also non calendar-constant effects. For example, many transfers and cash withdrawals are made around Easter, whose date varies. It is important to include it to not have false alarms that are just holiday-implied abnormal reductions of deposits. In Section 3.2 we describe the set of controls included in our application when euro payments are modeled.

formally, we propose the following procedure.

Let t_c be the last observed day without distress. It coincides with the day in which the algorithm is run in normal times. It is a rolling date, as we move forward at daily frequency.

ReNoSCh algorithm

1. Regularize the target variable Y with model M on a big time support T_B, which ends at time t_c, and set the new target variable $\tilde{Y}$ equal to the obtained residuals;

2. Use a time interval $T_S \leq T_B$, which ends at time t_c, to estimate the distribution of $\tilde{Y}$ with a nonparametric method D;

3. Estimate a threshold ψ_p for a selected probability p such that $P(\tilde{Y} < \psi_p | M(T_B), D(T_S)) = p$;

4. Set s and k, with $s < k$, where s is the critical number of days in which the observations exceeded the critical threshold in the last k days. A warning at time $j \geq t_c$ is defined as a binary variable: $W_j^{k,p} = 1$ if $\sum_{t=j-k}^{j} I(\tilde{y}_t < \psi_p) > s$ and 0 otherwise;

5. A distress episode occurs if U consecutive warnings are observed;

6. The episode ends when the bank fails or when we have E consecutive non-warnings days.

$\tilde{y}_t$ is the t^{th} element of the vector $\tilde{y}$, which contains all the observations of $\tilde{Y}$ sorted by their time index. T_B and T_S are two time intervals, which can have constant length (being rolling windows) or increasing length (being expandind windows). Step 1 transforms the target in a stationary variable to be used in the CC environment. We suggest to regularize the variable on wide interval T_B to capture as much as possible of its systematic component, and use eventually a smaller interval T_S in step 2, to capture potential small changes in the distribution of the random component occurred in the very last period when it is nonparametrically estimated in step 3. This approach is nonparametric because it labels "abnormal" observations (below ψ_p) using the observed density, instead of assuming a parametric distribution, like the normal distribution (commonly used with control charts), and define abnormal observations as those below a threshold equal to a certain number of estimated standard deviations (sigma) far from the estimated mean. It follows that if data is not normal, or not generated by a specific pre-set distribution, our approach is expected to be more effective to label abnormal observations, we show in Section 3.1 that this is the case. In what follows we use a daily rolling window, keeping T_B and T_S constant, which has the main advantage that the dimension of the data does not increase over time.[10] Nevertheless, the algorithm can accommodate an expanding window in case it is preferred. In

[10] This allows for the comparison of the algorithm performance in different points in time with a constant number of observations. Furthermore, in the practical implementation of the algorithm, running this type of algorithm over the entire banking system could be onerous, thus having a rolling window limits the related costs. Note that if the reference period is too large, including observations very far in time, the ability of the regularization step to capture more recent and relevant patterns in the payment activity of the bank decreases and the possibility of technical changes in payments routing increases. On the other hand if the reference period is too short, it could be less effective to capture the systematic

step 4 the researcher has to decide p, the probability that determines the critical threshold. Here the trade-off is standard, the higher the threshold the more false positive, the lower the threshold the more false negative. In step 5 we propose a criterion for labeling a day as a warning day, which is observing s days below the threshold over the last k days, if there are U warning days in a row (step 5) then we have an alert in place. Step 6 sets a similar criterion to establish the end of the episode. Such criteria can be changed and even made continuous, like using $W_j^{k,p} = \sum_{t=j-k}^{j} I(\tilde{y}_t < \psi_p)/k$, in this case U must be a continuous threshold, which does not represent the number of warning days anymore. The researcher may also avoid the consecutiveness of warning days and set U as a relative frequency in a reference period.

Observe that the number of choices involved is inevitably pretty big $C = (M, T, t, D, p, s, k, U, E)$, and there is not strictly preferable ones a priori. For this reason the practitioner has to fine tune these choices depending on the environment and constraints she faces. In the next section we provide some examples using Monte Carlo simulations. In Section 3.2 we give additional insights using real data. If we have multiple target variables, the algorithm can be used for each target time series separately, or they can be jointly considered using multivariate control charts (see Lowry and Montgomery, 1995). In the empirical application, we stick with the simplest set of choices, using a linear model in step 1 and sorting and counting in step 4. It is shown that even with such a not very sophisticated methodology the algorithm works pretty well in detecting distress episodes.[11]

In the practical use of ReNoSCh, after a sequence of warnings for a bank, it is suggested to check with the payment system operator whether some technical change in payments routing took place.[12] In Section E of the online appendix we discuss other potentially useful sources of information to detect deposits outflows in real time and discuss why payment system data is better to identify true distress episodes.

part of the bank's payment activity. For example, if only one year is considered, the systematic payments during Christmas could not be properly captured if the only observation period in the sample is not representative of what usually happens during Christmas, if for example in that year the bank's clients had some special business. The inclusion of more years allow to represent the systematic part better. We suggest to set the reference period between 2 and 7 years, guided by our experience with real data.

[11] The MATLAB code for the construction of the control chart is available at the following link: ReNoSCh.m.

[12] It could happen that the bank did some technical change in the way it routes outgoing payment messages into the system or a new technical arrangement was introduced by the operator and changed the road of ingoing payments. In this case, it is possible that from the implementation date of these changes net outflows in the subset of payments considered increase, because relevant inflows are not included in this subset anymore, creating a spurious warning. The good thing in this rare case is that the payment system operator, which is the data provider and the entity in charge with the running of the system, knows exactly when such changes happen and can then cross check it, also because the banks' treasuries are always in contact with and communicate technical changes to her. Note that this is a data feature and does not only apply to ReNoSCh.

3 Evaluation of the Performance of the Methodology

3.1 Monte Carlo Study

In this section simulated data is used to study the properties of our algorithm. We set the numerical experiment parameters using some real world features, which are described in Section 4. We simulate deposits net outflows using the following model:

$$y_t = z_t\beta - d_t r_t + \epsilon_t, \ with \ \ t = 1, ..., T. \tag{4}$$

Where $T = 960 = 20 * 12 * 4$, a length close to the number of working days in 4 years, the average period in our empirical sample. ϵ_t is normally distributed with mean equal to zero and variance equal to σ. For simplicity let z_t be a $T \times k$ matrix capturing monthly seasonal effects and β be a $k \times 1$ vector, where $k = 12$, the number of months. Let $\beta_{1:11} = 0$ and $\beta_{12} \geq \beta_{1:11}$, resembling higher customer payments outflows in December, during Christmas. r_t is an indicator function that switches to one when a distress episode triggers, d_t is the relative outflow. We generate 500 samples characterized by distress and other 500 generated without any. To mimic the features of real episodes observed, we set $r_t = 1$ if $t \in [T - 20, T]$ when a it occurs, so the episode is four weeks long in the last period of the sample.[13] We set $d_t = 3\sigma$, which means that the expected daily net outflow during distress days is slightly below the first percentile of the distribution of customer payments.

Using this setting we play with our algorithm's parameters. To move along the different dimensions, we use a pivotal setting with $\beta_{12} = 0$, $s = 3$, $k = 5$, $U = E = 5$, $M = linear$, $T_S = T_B = T$ and $p = 0.075$. As the βs are all set equal to zero we skip the regularization step of ReNoSCh, to then quantify the bias that emerges if it is not performed when seasonal effects matter. We focus on four outcomes generated by the simulations. In order to understand the sensitivity of the efficacy and the fallacy of our algorithm to the main parameters we look at the frequency of true positive over the total number of real distress episodes generated and the frequency of false positive over the total number of no-distress episodes generated. To assess its reactivity under different specifications we compute the average distance from the beginning of an episode for the first day in which a warn is observed and for the day in which a sequence of warning days is labeled as a distress period.

First, we want to study the sensitivity of our method to the ratio between s and k, which is the number of days with an extreme value (s) over the last k observations. We set $k = 5$, the working days of a week and change $s = 1, 2, 3, 4, 5$. The first panel of Table 1 reports our results, where we can see that when s is very low we have a high number of false positives, equal to 74 percent when $s = 1$, while every true case is detected. The false positives dramatically decrease already when $s = 2$ to totally

[13] We place the distress period at the end of the time interval for illustrative purposes, starting from a longer observation period in normal times and then having distress kicking in. It mimics what a practitioner does and the logic of the algorithm, which leverages observations before the distress. The results are similar if we place the distress period in a random position, as long as there are enough observations before the distress to run the regularization steps properly.

disappear when $s = 3$. When s increases further the algorithm is less able to intercept true cases, about 15 percent of them are not recognized when $s = 5$. This results is implied by the fact that a too strict sequence of observations below ψ_p is less likely to be observed, because some outflows may not be that negative and alternate with values above the critical threshold even if the bank is in distress. From the last two columns we can also see that the reactivity of the algorithm is higher for smaller s.

Another important parameter to set is p, the critical probability. Choosing it extremely low would not let the algorithm detecting critical episodes if the average outflows are not that extreme, while many false positive may appear if it is set too high. In the second panel of Table 1 we report the results of simulations with $p = 0.01, 0.025, 0.05, 0.075, 0.15$. When p is set equal to 0.01, about 12 percent of the true positive is not detected. The percentage of correctly recognized distress episodes sharply increase already at $p = 0.025$, to reach zero at $p = 0.075$. When this probability is set too high, in the table equal to 15 percent, some false positive cases start to appear. With an higher p the algorithm is faster in detecting the true distress episodes.[14]

In the third panel of Table 1 we study how the presence of systematic outflows, here represented by seasonal effects, can bias the identification of distress episodes if the regularization step is not performed. We first introduce the seasonal outflows in December and set them equal to $3/4\sigma$, a fourth of $E(d_t)$, and then double them. From the table we can see that the distortion is only on false positives, which move from zero to 74 percent. In other words, the algorithm signals a distress episode almost every time there are seasonal outflows, if they are comparable. In our practical experience, this is often the case. Especially during holidays the value of outflows can be even higher than those observed in distress episodes for some banks. In the fourth panel of the table we change U, the number of consecutive warning days. Given that we defined a warning day as a signal which may be induced by more than one day below ψ_p, we can see that even with low values of U the number of false positives is pretty small, but not negligible. The minima of both type I and type II errors are reached when it is equal to five. Increasing it further implies a sharp raise of type II error, with basically no true distress episode detected. This is again because the probability of observing very long sequences of warning days is low even when the expected outflow is high.

As a last exercise, it is shown how the nonparametric approach proposed performs with respect to the "sigma approach" when the distribution of the random component is not normal. The sigma approach labels as "abnormal" observations below a certain threshold, which is a certain number of estimated standard deviations (sigma) far from the estimated mean, using two moments of an imposed distribution. The nonparametric approach in step 3 of ReNoSCh does not impose normality nor use estimated means and standard deviations. Instead, it uses the full distribution of observations to label abnormal observations. If an observation is below the ψ_p percentile of the empirical distribution then

[14] Such result depends on $E(d_t)$, the expected outflow during a distress episodes. Given that we set it equal to 3σ and we simulated ϵ with a normal distribution, the median of $\tilde{y}$ when the shock occurs is below its first percentile before the shock occurs. If $E(d_t)$ were higher we would have preferred to set p higher.

Table 1: Simulation Study - Playing with ReNoSCh Parameters

	True positive	False positive	First warning day	First alert
s				
1	1.000	0.744	0.060	4.060
2	1.000	0.076	0.854	4.854
3	1.000	0.000	1.968	5.982
4	0.996	0.000	3.152	7.290
5	0.868	0.000	7.136	11.680
p				
0.010	0.884	0.000	6.610	10.854
0.025	0.992	0.000	2.626	6.738
0.050	0.997	0.000	2.262	6.288
0.075	1.000	0.000	1.968	5.982
0.150	1.000	0.020	1.582	5.590
β				
0	1.000	0.000	1.968	5.982
$3/4\sigma$	1.000	0.050	1.824	5.824
$3/2\sigma$	1.000	0.742	1.794	5.794
U				
1	1.000	0.042	1.944	4.000
3	1.000	0.014	1.944	4.140
5	1.000	0.000	1.968	5.982
7	0.000	0.000	30.000	30.000
9	0.000	0.000	30.000	30.000

Notes. Results based on 1000 replications, with 500 processes with distress episodes and 500 without. The column 'True positive' reports the percentage of detected true distress episodes over total true distress episodes. The column 'False positive' reports the percentage of no distress episodes erroneously labeled as distress episodes. The column 'First warning day' reports the average distance between the beginning of a distress episode and the first day in which a warn is observed, according to the algorithm's parameters. The column 'First alert' reports the average distance between the beginning of a distress episode and the day in which a sequence of warning days is labeled as a distress episodes, according to the algorithm's parameters. When no distress episode is recognized by the algorithm even if present, we set such distance to 30 days.

it is abnormal. Not implicitly assuming normality, it is thus expected to be more effective if data is not normal. This time ϵ is generated as a mixture of two different distributions:

$$\epsilon = I(\omega < \kappa)N(0,\sigma) + I(\omega \geq \kappa)P, \text{ with } \kappa \in [0,1], \omega \sim U(0,1), \tag{5}$$

and P equal to a not normal distribution. In practice κ observations have the same distribution of before and $1-\kappa$ have a different one. We set P equal to χ^2_ζ or T_υ, where ζ and υ are the degrees of freedom of the χ^2 and T distributions respectively. The T distribution is the closest approximation to normality, among the standard distributions, but is useful to study the performance of the two approaches when the distribution has fatter tails (when υ is small). The χ^2 instead departs more dramatically from the normal distribution, especially in terms of skewness and kurtosis, which are simple functions of ζ, respectively equal to $\sqrt{8/\zeta}$ and $12/\zeta$, so that it is easy to play only with the degrees of freedom to change these features of the resulting distribution of ϵ. Results do not change qualitatively if consider other distributions. Our pivotal setting is used to generate replications, the critical threshold ψ_p for the sigma approach is set such that $P(z < \psi_p) = 1 - F(z) = p$, with $z \sim N(\hat{\bar{y}}_t, \hat{\sigma}_{y_t})$, where $\hat{\bar{y}}_t$ and $\hat{\sigma}_{y_t}$ are the sample estimates of y_t mean and standard deviation in the pre-distress period. We set $\kappa = 0.6$, to have slightly more than the majority of observations being normal, and $d_t = \nu$ such that $P(y_t < \nu) = 0.13$, which has the same probability on its left of $d_t = 3\sigma$ when the random component has a standard normal distribution, as in our former experiments. Given our interest in understanding whether a nonparametric approach detects distress episodes more effectively than the sigma approach, we focus on the power of these two procedures. In the second and third columns of Table 2 we report the frequency of true positives (on all true cases) detected respectively by the sigma and the nonparametric approach. The performance is explored along two dimensions: the share of normal observations, κ (the mixing parameter), and the degrees of freedom, ζ and υ. On average the nonparametric outperforms

Table 2: Simulation Study - Sigma vs Nonparametric Approach

Distribution	Normal share	Dof	Sigma	Nonparametric
T_υ				
	κ	υ		
	0.10			
		4	1.000	1.000
		2	0.996	1.000
		1	0.910	1.000
	0.05			
		4	1.000	1.000
		2	0.996	1.000
		1	0.898	1.000
	0.00			
		4	1.000	1.000
		2	0.998	1.000
		1	0.886	1.000
χ^2_ζ				
	κ	ζ		
	0.90			
		1	1.000	0.970
		3	0.998	0.942
		6	0.150	0.966
	0.80			
		1	1.000	0.970
		3	0.956	0.882
		6	0.340	0.872
	0.60			
		1	1.000	0.964
		3	0.784	0.678
		6	0.000	0.548

Notes. Results based on 1000 replications, with 500 processes with distress episodes and 500 without. The column 'Sigma approach' reports the percentage of detected true distress episodes over total true distress episodes for the sigma approach. The column 'Nonparametric approach' the percentage of detected true distress episodes over total true distress episodes for the nonparametric approach.

the sigma approach by far with values always greater than 50 percent. From the first panel, where we consider the T distribution, nonparametric always outperforms sigma with true positives always detected. With fatter tails, setting $\upsilon = 1$, and no normal observations the sigma approach is not able to detect more than 10 per cent of true positives. Given the similarity of the T with the normal, the frequency of true positives of the sigma is not so low. When we instead turn to the χ^2 the sigma approach suffers much more, reaching its higher values only when ζ is low and κ is high and not being able to detect positive cases at all already when ζ is equal to 6 and 40 per cent of observations are not normal. When skewness and kurtosis are significantly different form normality and outliers inflate the variance estimated by the sigma approach, the latter suffers the most in detecting true positive cases. It has to be noted that the power of both approaches decreases with ζ, this is because the fatter right tail of the random component makes more likely to have some observations above the critical threshold, as witnessed in Figure A.3 in the online appendix, where the histograms for the nonparametric, sigma and observations are plotted.

In the practical implementation, we suggest to observe the distribution of the random component and compute the outcomes of the algorithm under different combinations and monitor them in parallel.

3.2 Empirical Application

After having described the source of information to keep track of deposits outflows and the methodology to timely identify shocks, we move to the empirical part. We consider the Italian banking system operating in TARGET2 from August 2012 to August 2019. The period is particularly suited for our analysis because many idiosyncratic distress episodes characterized the Italian banking system, which emerged from the global financial crisis and the sovereign debt crisis with some vulnerabilities, and there were several changes of regulation related to banks' resolution. Some banks received aid under the applicable EU rules on resolution. The national insolvency proceedings facilitated the market exit of some banks. Other banks were precautionary recapitalized. Furthermore, many banks were often spotlighted for the high incidence of non-performing loans in their balance sheets.

With respect to the subjective choices mentioned before, let us first detail our settings in what follows. We constructed the daily variation of bank i's net position on each of the following channels by subtracting the total amount of outgoing payments to the total amount of ingoing payments related to the respective payment category.[15] *Deferred net settlement transfers* ($DNST$) are the bank's multilateral position settled by domestic and international RPSs. *Real-time settlement transfers* ($RTST$) are bilateral gross interbank transfers settled in real-time in central bank money for customer payments. *Cash withdrawals* ($CASH$) are cash operations by commercial banks exchanged with reserves in TARGET2. See Section C of the online appendix for a detailed description. We regularize payment data for the distressed banks on a time support $T_B = 4$ years. We included a constant, a trend, monthly and daily dummies, pre/post/during holiday periods fixed effects -namely Christmas, Easter and Summer-, start/middle/end of the month dummies, fiscal and tax payment dates dummies (around the 20th of each month some taxes are paid, generating significant payment burden for the banks). This setting gives up to about almost 1000 observations per bank and less than 100 controls, which are by construction not linearly dependent, thus we can use simple OLS to regularize the data.[16] We then take a smaller time interval T_S of six months; we estimate the distribution of $\tilde{Y}$ by simple sorting and counting. Kernel density estimate does not provide superior results. We estimate a threshold such that $P(\tilde{y}_t < \psi_p) = p = 0.075$, which gives us more sensitivity (than 0.05) to departures from a control state and does not capture too small deviations (like 0.15, see Table 1).

Based on our Monte Carlo study in Section 3.1, we set a warning as a sequence $s = 3$ of observations

[15] In principle inflows of deposits during distress episodes attracted by the combination of credible deposit insurance and above-market rates (Martin et al., 2017) and outflows implied by asset side runs (Ippolito et al., 2016; Ivashina and Scharfstein, 2010) from firms' increased draw-down on available credit lines offered by the distressed bank may be included in the net deposits outflows.

[16] In case the degrees of freedom are too small other techniques like LASSO can be use to treat the dimensionality problem.

below the threshold on $k = 5$ consecutive days. Three over five days gives a good compromise if we think that a distress episode implies continuity of observations below the critical threshold (as shown below). Table A.1 in the online appendix reports some simple numerical examples that show that $s = 3$ outperforms $s = 2$ in terms of type II error and $s = 4$ in terms of type I error, when distress episodes last at least five consecutive days of outflows below the critical threshold. Finally, we have to set U, the number of consecutive warning days to trigger an alert. Our Monte Carlo study suggests that $U = 5$ minimize type I and II errors in the simulation settings explored, but with real data this is not guaranteed. Given that this is probably the most critical choice, we run the algorithm in parallel for $U = 0, 1, .., 7$.

In the implementation, we do two exercises to study the performance of our algorithm. First, we run the algorithm systematically for the whole banking system for each day in the sample period, keeping track of the warnings for each bank in each day. Second, we identify bank-specific potential distress episodes that involved banks in our sample, which may have caused deposits outflows. To define this set, we take all the cases in which the supervisory authority start a principal procedure and all the cases in which newspapers start to diffuse 'bad news' about a bank. Below we detail the type of measures considered and a formal definition of bad news diffusion.

The principal measures concerning entities subject to supervision considered here are: (i) special administration procedures; (ii) liquidation and withdrawal of authorization procedures; (iii) prohibitions, closures and injunctions.[17] We define bad news here as the release of negative information about the soundness of the bank, i.e. any information that put the solvency of the bank into question and may impact the risk associated with its deposits. More formally, if all the major domestic newspapers, namely *Il Corriere della Sera, La Republica* and *il Sole 24 Ore*, and eventually international newspaper, namely the *Financial Times* and *The Wall Street Journal*, publish an article with the words 'crisis', 'failure', 'distress' associated with the name of the bank, the latter enters into our sample with the date of the publication of the article that appears first. We can not exclude that outflows could happen without the diffusion of bad news, but, as retail deposits outflows are triggered by generalized distrust, they are usually accompanied by bad news and if a rumor diffuse for example through social networks in a so widespread form to generate a distress episode, it is difficult to imagine that the press does not recognize it (especially given the participation of the press itself in the social networks). Following these critera, we construct a sample of potential distress episodes composed by bank-day pairs.

[17] In more detail, special administration procedures occur when the supervisory authority detects violations or irregularities or it foresees serious losses of assets. The commissioners are then appointed and responsible for ascertaining the company's situation, removing irregularities and promoting useful solutions in the interests of depositors and the sound and prudent management of the bank. The compulsory administrative liquidation of the bank can also be ordered when it is undergoing extraordinary administration. The compulsory liquidation can be ordered, upon motivated request of the administrative bodies, the extraordinary shareholders' meeting, the extraordinary commissioners or the liquidators. From the date of issue, the functions of the administrative, control and shareholders' meeting bodies, as well as any other body of the bank cease. The supervisory authority may also take extraordinary measures and impose a ban on new operations or order the closure of branches of authorized banks due to violation of legislative, administrative or statutory provisions governing their activities, management irregularities or even insufficient funds. For these cases we include in our sample the bank and the day in which the first principal measure took place.

For these cases, we first check whether the algorithm signals the existence of a distress episode from the first exercise and then check whether there was a real outpouring of reserves (see Section F of the online appendix). In this case we take the episodes as true positives. In case the algorithm signals a warning in the first exercise outside these cases in the second exercise, we check for the existence of rumors or news regarding the bank and its balance sheet data from supervisory reports. If nothing is anomalous, we label it as a false positive. From the comparison between the first and the second exercises, we do not observe any true deposits outflow starting without the diffusion of bad news by the press, but we do observe diffusion of bad news without anomalous deposits outflows, which highlights that we can not just rely on public information to identify real outflows. The exercise reveals the existence of three anomalous deposits outflows cases, all among the ten cases at the intersection of the two criteria. The outflows all started from the day of the triggering events (supervisory actions and bad news), see Figure 2.[18] The episodes did not involve always the same bank and were not clustered, meaning that they did not happen all in a narrow time interval.[19] The involved banks had on average 2.5 per cent of the total amount of deposits held at the banking system before the shock. The algorithm is able to real-time detect them, with a very small portion of false alarms if an appropriate number of consecutive days of outflows is set. Table 3, shows type I and II errors when different numbers of consecutive warning days (U) are labeled as distress episodes.

Table 3: Empirical Type I and II Errors

U = 0	Warning	No warning	U=1	Warning	No warning
No distress	1.000	0.000	No distress	0.527	0.473
Distress	1.000	0.000	Distress	1.000	0.000
U=2	Warning	No warning	U=3	Warning	No warning
No distress	0.163	0.837	No distress	0.033	0.967
Distress	1.000	0.000	Distress	1.000	0.000
U=4	Warning	No warning	U=5	Warning	No warning
No distress	0.006	0.994	No distress	0.001	0.999
Distress	1.000	0.000	Distress	1.000	0.000
U=6	Warning	No warning	U=7	Warning	No warning
No distress	0.000	1.000	No distress	0.000	1.000
Distress	0.667	0.333	Distress	0.667	0.333

Notes. Warnings are generated with the ReNoSCh algorithm, described in Section 2.2. U is the number of consecutive warning days set by the researcher, according to the specifications in the empirical application. The other parameters of ReNoSCh are described in Section 3.2. A warning triggers when the number of consecutive warning days is greater or equal to U. A distress episode is defined as the event in which a significant outpouring of liquidity was observed ex-post as a reaction of bad news in many newspapers. For U = 6 and 7 in the distress cases, some non warning day appeared in between consecutive warning days.

For few consecutive warning days ($U < 5$), we can see that type I error is very big, while we always get properly the true distress episodes. When U is set to 5 days, a working week, the algorithm reaches its best, with type II error equal to zero and type I error equal to 0.1 percent. For values greater than 5, the type I error is still on its minimum but the type II error gets bigger. This is because we may have

[18]Observe that if large outflows were the reason for the supervisory intervention and the bad news in the press, they should have happened after day 0, probably after day 20 in Figure 2, which is not the case.

[19]For high confidentiality of the data used we can not give the dates of the identified episodes neither the characteristics of the banks involved.

alternate warning/no warning days sequences. We think that setting $U = 5$ (a working week) is the best choice, as it shows a remarkably good performance with real and simulated data.[20] For E we set the same length, thus an episode ends after five consecutive non-warnings days. Under this specification, the raw and regularized time series of daily changes of net customer payments is stationary before and after the shock, passing the augmented Dickey-Fuller test (see Table A.2).

In Section F of the online appendix we backtest ReNoSCh with ex-post methods, using structural exogenous (Chow, 1960) and endogenous (Bai and Perron, 1998) break tests, and show that it is able to identify the break date with high precision. It is important to note that the high performance of the algorithm is also generated by the good quality of the information available. Given that we are able to identify exactly customer payments to other banks and cash withdrawals, if deposits are flying away, we see it. Clearly if the measurement is poor and, for instance we are not able to disentangle customer payments from interbank payments, or cash withdrawals are confounded with other operations of the bank with the central bank (like open market operations), these numbers may worsen dramatically and even invert. Indeed, in Section 4 we show how the liquidity drain is offset by open market operations and some interbank payments, if these flows were not separable from customer payments, we would just have had a flat line (the balance of the reserve account of the bank under distress).

4 Economic Insights from the Identified Distress Episodes

After having shown that our real-time identification scheme works with simulated and real data, let us move to the description of the features of the identified episodes. In this section, we exploit the high frequency and granularity of payment system data, firstly describing the features of outflows, and then studying the interactions between retail and wholesale funding. The goal of the first part is to characterize the anatomy of outflows, while the goal of the second part is to assess the consequences of idiosyncratic shocks on wholesale funding and the strategies adopted by the banks to eventually absorb the shock and remain liquid. Both contain new evidences that can be captured with payment system data and can inform policy makers by establishing links to monetary policy, market operations, regulation, and supervision, answering questions like: how much of deposits is converted in cash? To what extent are real-time payments used to move deposits in distress episodes? How much do they accelerate the speed and depth of banks' distress? Do wholesale funding worsen or offset the liquidity drain generated by outflows in retail funding? What is the share of funding compensated by central bank liquidity and money markets?

In what follows, we mainly focus on changes occurred around the identified shock. Given that the methodology developed in Section 2.2 is able to identify the beginning of anomalous deposits outflows, and this is taken as the starting date of the identified shock. It follows that the coefficients in this

[20] Alternatively, we can also allow for non consecutive days criteria, like observing average frequency of warning days non necessarily in a row, but this could increase the type II error.

section can be interpreted as causal effects of this type of events on different types of deposits outflows (for the first part), and on different types of reserves' inflows (for the second part).

The analysis includes also non-distressed banks as a control group, exploiting a panel of payment activities for all the banks in the sample at daily frequency. The inclusion of bank and day fixed effects allow to control for unobserved heterogeneity both in time and cross-section. Clearly, pure exogeneity and external validity can not be claimed because these type of events can not be randomized and are idiosyncratic and highly heterogeneous. Nevertheless, evidences reported in this section are informative and new in many dimensions and in the second part it is fairly reasonable to assume that deposits outflows caused reserves' inflows and not the other way around.

4.1 Features of Outflows

In what follows, we use the sample of distress episodes identified by ReNoSCh (averaging the single cases). When not specified differently, we report the amount of outflows as a percentage of the banks' deposits before the shock.[21] Figure 2 provides an aggregate overview of the sudden decline of deposits when depositors' trust is undermined. In this plot, we average changes in net positions for digital transfers ($RTST + DNST$) and cash withdrawals ($CASH$) and cumulate them over time, starting from 100 days before the episode begins.[22] The vertical line is the starting date estimated by ReNoSCh. We can see that it gets the break quite precisely, and more importantly in real-time. The drop in digital transfers is pretty impressive as well as that for cash withdrawals. Nevertheless, the magnitude is different. The distress episodes last on average four weeks, 20 working days, with an initial more intense phase of about two weeks. There is a following less intense period of other two weeks. The liquidity drain is significant and equal to about 3 percent of the deposits of the bank.

[21] Deposits are calculated from supervisory reports data with monthly frequency, more precisely we consider the same type of deposits reported for the calculation of the LCR in the last observation available before the first day of distress identified.

[22] In these episodes, instant payments were not used by the banks.

Figure 2: Visual Evidence from Identified Episodes

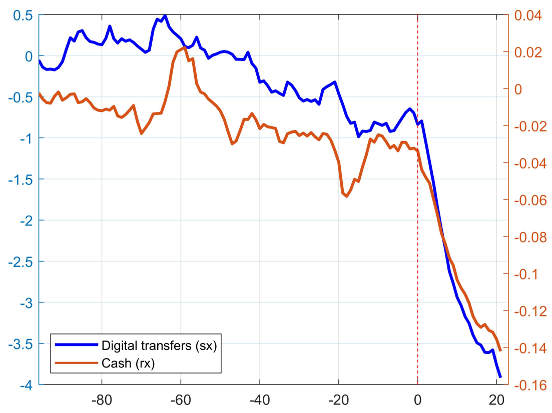

Notes. x-axis: days. 100 days before and 20 after the beginning of the episode. Day 0 is the first day of deposits outflows estimated by ReNoSCh (red vertical line). y-axis: cumulated outflow expressed as a percentage of deposits. The cumulated outflow is computed as the net unexpected position for each channel averaged across all the episodes identified. Left axis: value for digital transfers, right axis: value for cash withdrawals.

In what follows we analyze in more detail the features of identified episodes. First, cash withdrawals and digital transfers to other banks are quantified and compared. Second, a closer look at both is provided, exploiting variation in banknotes denominations and characteristics of digital transfers, in particular destination and speed of settlement. Finally, we study the interaction between retail and wholesale funding.

Routes to Safety

Let us now explore in more detail the features of the identified episodes. After having identified these cases, we may be interested in understanding how they occurred and what are their most salient features. Thanks to the detailed information available in payments data, we can uncover the behavior of depositors as never done before. Figure A.6 in the online appendix depicts the potential routes to safety available to a depositor. She can convert deposits into cash at the bank teller or at the ATM, alternatively she can move her deposits to another bank, domestic or foreign. She may prefer a big bank or she may even move the deposits abroad. The increasing development of financial services and financial integration in Europe makes nowadays very easy to open and transfer euro funds from one participating country to another. In addition the fast spread of new mobile and fintech products is eliminating many frictions. This is a particularly important aspect for the eurozone, as the entity ultimately backing money could change across countries.

In order to reduce the incentive to withdraw deposits from a bank the government can guarantee deposits (Bryant, 1980; Diamond and Dybvig, 1983). In Europe, every country has a fund that guarantees deposits up to a certain amount. Therefore, at least for eurozone countries, the institution guaranteeing banknotes (ECB) is different from the one guaranteeing deposits (depending on the country). This difference, together with the disutility represented by holding a credit with the fund instead of deposits (in case of default of the bank), produces a clear disparity between cash and deposits at the bank under distress, in same-country banks or in different-country banks.

Interbank Digital Transfers vs Cash Withdrawals

Disentangling these two different choices helps us understand the nature of depositors' fears. As banknotes are central bank money, depositors shall prefer it over deposits at other banks, especially if they trust less the entire banking system or they believe that there will be negative spillovers to other banks (Goldsmith-Pinkham and Yorulmazer, 2010).[23] On the other hand, if they do not trust only a specific bank anymore, they may prefer moving the deposits to other banks that they see as sounder, using digital transfers. In this case their credits remain in commercial bank money and deposit currency ratio holds constant. Timely having information on these preferences is very helpful from a financial stability and a monetary policy perspective (Waldo, 1985).

We pool together the payment activity data at daily frequency for all banks (distressed and non-distressed) in the system around the identified episodes and estimate the following simple regression models,

$$y_{i,t} = \delta r_{i,t} + \alpha_i + \mu_t + u_{i,t}, \tag{6}$$

where the dependent is the daily (t) change in the net position (credits minus debits) of bank i for interbank digital transfers ($y_{i,t} = RTST_{i,t} + DNST_{i,t}$) or cash operations ($y_{i,t} = CASH_{i,t}$). $r_{i,t}$ is a dummy that switches to one if bank i is distressed in day t, according to ReNoSCh, α_i is a bank fixed effect, μ_t is a day fixed effect and $u_{i,t}$ is the error term.[24] The pre-shock period is 20 settlement days before the shock. The post-shock period starts at the first day identified by ReNoSCh and is 20 days long, with a symmetrical interval. Table 4 reports the estimated coefficients with and without fixed effects (respectively in the second and in the first column). In the upper panel we use the raw payment data, in the lower panel we consider the regularized data used in ReNoSCh, $\tilde{y}_{i,t}$ instead of $y_{i,t}$, thus controlling also for bank-specific trend, monthly and day of the week dummies, pre/post/during holiday periods, start/middle/end of the month, fiscal and tax payment effects. It follows that the latter provides the cleanest effect of the shock on bank's cash flows. Let us focus on the bottom of the second column, where we compare distress days with non distress days controlling for the widest set of factors. As we can see, digital transfers are much bigger in magnitude. The average daily drain is about 0.12 percentage points w.r.t. less than 0.01 percent outflow generated by cash withdrawals.

Depositors seem to be worried about their bank's solidity, moving their funds mainly to other banks, and thus keeping the deposit currency ratio constant. This is good news for the financial stability of

[23] They may also withdraw cash and then deposit it to another bank, but this hypothesis is quite unlikely as the direct transfer is much easier and cheaper (the transfer costs few euro).

[24] We do not include lagged terms in this and in the following specifications as the daily changes analyzed in this section do not show significant autoregressive features, see Table A.3.

Table 4: Digital Transfers vs Cash Withdrawals

Dependent: daily bank-specific net position in			
	Raw data		
Interbank digital transfers			
	Distress	-0.0386	-0.1324 ***
		(0.0228)	(0.0255)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	1.990	2.043
	Observations	9,804	9,804
Cash withdrawals			
	Distress	0.0096 ***	-0.0060 ***
		(0.0018)	(0.0020)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	1.929	2.006
	Observations	9,804	9,804
	Regularized data		
Interbank digital transfers			
	Distress	-0.1408 ***	-0.1176 ***
		(0.0226)	(0.0275)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	1.996	2.043
	Observations	9,804	9,804
Cash withdrawals			
	Distress	-0.0052 ***	-0.0047 ***
		(0.0011)	(0.0014)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	1.972	2.016
	Observations	9,804	9,804

Notes. * : $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$. OLS estimates of δ. Heteroscedasticity and autocorrelation consistent estimates for all the standard errors. The dependent is the daily change of the net position in euro respectively for interbank digital transfers and cash withdrawals computed as the sum of credits minus the sum of debits for each bank daily. The coefficients on different lines refer to separate regressions and are reported in percentage points of deposits of the shocked banks. R^2 and other coefficients of each regression are not reported for brevity. The pre-shock period is the 20 settlement days before the shock. The post-shock period is the 20 days after the shock. DW is the Durbin-Watson statistic.

the system. Taking other financial intermediaries' perspective it can also be seen as a positive spillover. In 20 days almost 3 percent of deposits were transferred to other banks. Such positive spillovers may alleviate negative ones, like those reported in Goldsmith-Pinkham and Yorulmazer (2010). While negative spillovers were explored in the literature, our evidence is new and brings relevant insights on retail funding during idiosyncratic distress episodes. Even though the wholesale funding of the rest of the banking system can shrink during an idiosyncratic distress episode, the retail funding can raise because of inflows from the stressed bank. Below we also shed lights on which intermediaries benefit the most. In Section G of the online appendix we provide a detailed analysis of cash withdrawals using denomination-level data, which shows a prominent role of the demand for small denomination available at the ATM.

Positive Liquidity Spillovers

Another interesting aspect to investigate is which banks received the deposits. Being a cheap and relatively stable liability, they generate a positive spillover. Banks are not identical and depositors may have heterogeneous preferences. Furthermore, as a unique feature, the Eurosystem is composed by many nations having the same currency.

When trying to address these curiosities we face a small constraint imposed by retail payment systems' netting mechanisms. Given that these systems send net multilateral positions, we do not have bilateral interbank flows, so we can not identify exactly the receiver of the funds. Nevertheless, this is possible for real-time bilateral transfers settled directly in the RTGS. Even if this is a portion of all the digital transfers, it accounts for about a half of the total liquidity drain and can thus provide us with useful information.

We focus on $RTST_{ij,t}$, where i is the distressed bank and j is another bank not hit at time t. We run simple pairwise regressions where the dependent is the change of the net bilateral position of bank i versus bank j (credits minus debits). Such variation is computed as the difference of the daily average net position before and after the shock. The pre-period is 20 settlement days preceding the start of the episode. The post-period coincides with 20 days after the episode started. On the right hand side we put a dummy taking value equal to one if bank j has the same nationality of bank i. We also interacted this dummy with a proxy of the size of the bank, computed as the sum of payments sent in TARGET2 in the previous five years.[25] Our regression model takes the following form,

$$l_{kj} = \theta g_{kj} + \gamma g_{kj} * s_j + \alpha_i + w_{kj}, \tag{7}$$

where $l_{kj} = \sum_{t<t_r} RTST_{ij,t} / \sum_{t<t_r} 1 - \sum_{t \geq t_r} RTST_{ij,t} / \sum_{t \geq t_r} 1$ is the increase in average net liquidity received by bank j from bank k after the shock to bank k occurs (at time t_r), g_{kj} is a dummy equal to one if bank k and j belong to the same country, s_j is the size of bank j, α_i is a distress episode fixed effect and w_{kj} is the error term. Table 5 reports our results. In column (1) we take all the variations, in column (2) we restrict the sample to negative variations -i.e. when $l_{kj} < 0$-. From the first row of

Table 5: Bilateral Digital Transfers - Nationality and Size

Dependent: Δ pair-specific customer payments net position	All (1)	Negative (2)
Domestic banks	-0.0007 ***	-0.0013 ***
	(0.0003)	(0.0006)
Size of domestic banks	-1.61E-09 ***	-4.47E-09 ***
	(8.86E-10)	(8.26E-10)
Run FE	Yes	Yes
Observations	1,263	475

Notes. * : $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$. Estimated coefficients of a dummy switching from zero to one when the shock occurs. Heteroscedasticity and autocorrelation consistent estimates for all the standard errors. The dependent is the change in the daily net position in euro for each bank hit by a shock via-a-vis with other banks. The coefficients are reported in percentage points of deposits of the shocked bank. Only the RTST are considered. We restrict this analysis to such transactions because for DNST the counterparty is not identifiable as ACH send multilateral positions to the RTGS. The pre-shock period is the 20 settlement days before the shock. We pooled together all the identified episodes. In the 'All' column all the bilateral positions are taken, in the 'Negative' only positions with negative deltas are considered.

the table we can see that the change in bilateral interbank flows to domestic banks was much bigger.

[25] Even if we can identify the country of destination of deposits outflows, we do not exactly know the characteristics of the foreign bank that is receiving the funds. Thus the analysis of cross-bank spillovers can not include bank size in the level, not only interacted with the domestic bank dummy.

These outflows were more likely to be directed to bigger banks. This evidence can also reconcile to a premium generated by the perception of depositors of a possible 'too-big-to-fail' policy (e.g. Oliveira et al., 2014). An alternative explanation could also be that big banks have more reach and visibility than small ones, and thus people willing to open a new account is more likely to be exposed to their advertising and marketing. For this reason we label it simply a *size premium.* We can exclude that the migration of deposits is generated by a drop in interest rates offered by the distressed banks. Indeed, as witnessed by Figure A.7, distressed banks did not lower interest rates on deposits after the identified date.[26]

An additional interesting aspect to explore is the destination of funds once they cross the border. Table 6 reports the top four countries in terms of negative change of cross-border net inflows to the distressed banks in percentage points of their deposits. The first column contain the country name, the second (third) column reports the pre-shock (post-shock) net inflows from the country to the distressed banks, the fourth column reports the difference, and the last the relative percentage change. The top four countries are, in order of net negative change: Germany, Belgium, Great Britain and Luxembourg.

Table 6: Cross-border Digital Transfers - Top Four Countries

Country	Pre-shock	Post-shock	Δ	% change
DE	0.013	0.005	-0.008	-65%
BE	0.001	-0.004	-0.005	-424%
GB	0.009	0.007	-0.003	-27%
LU	0.001	-0.001	-0.002	-153%

Notes. Top four countries in terms of negative change of net inflows to the distress banks. All RTST payments are included. We restrict this analysis to such transactions because for DNST the counterparty is not identifiable as ACHs send multilateral positions to the RTGS. The second, third and fourth columns are expressed in percentage points of distressed banks deposits. The fifth is a change in percentage points. The pre-shock column reports the net inflows from the country to the distressed bank before the shock. The post-shock column reports the net inflows from the country to the distressed bank after the shock. The pre-shock period is the 20 settlement days before the shock. The post-shock period is 20 days long. We pooled together all the identified episodes.

The Speed of Digital Transfers

The time to settlement could be an important discriminant during distress periods. Advances in technologies used to transfer money create the possibility to move them from one account to another faster and without significant spatial and temporal frictions. Today, with the rise of instant and mobile payments is even possible to have full disposal of funds in seconds from wherever there is an internet connection.

It follows that there is a potential for faster and continuous outpouring of deposits and no moment of respite for bankers to manage the draining. Clearly, such possibility can be exploited by banks'

[26]Of course we do not rule out the possibility that interest rates matter in the shifting of deposits among banks in general terms. If adapted to identify less disruptive outflows and match them with inflows of other banks, our algorithm could be used also to assess the sensitivity of deposits to interest rates.

clients, especially when the fear of loosing money kicks in suddenly. A series of important questions are related to this argument and connect to the potential risks of faster payments (Weyman, 2016). Can these technological innovations accelerate the speed and depth of distress? Can it significantly increase the volatility of banks' reserves? What is the additional liquidity pressure created by the possibility of use real-time settlement?

Despite the prominent role of new technologies in payments, there are not evidences on the effects of the speed of settlement because of the scarce visibility on critical episodes. Having the possibility to identify real-time and deferred settlement payments one by one with their timestamp offers us the possibility to quantify the additional pressure generated by real-time settlement. The real-time settlement, anticipating the liquidity outpouring of about one day, puts additional pressure on the stressed bank.

As a first step, we are interested in quantifying the magnitude of real-time vs deferred settlement. We regress the daily change of the net position of the bank for real-time and deferred transfers, pooling together all the identified episodes and estimating the following panel regression models including stressed and non-stressed banks as above,

$$y_{i,t} = \delta r_{i,t} + \alpha_i + \mu_t + u_{i,t}, \tag{8}$$

where the dependent is the daily (t) change in the net position (credits minus debits) of the bank for real-time transfers ($y_{i,t} = RTST_{i,t}$) or deferred transfers ($y_{i,t} = DNST_{i,t}$). $r_{i,t}$ is a dummy equal to one after the shock hits the distressed banks and zero otherwise, α_i is a bank fixed effect, μ_t is a day fixed effects, and $u_{i,t}$ is the error term. The pre-shock period is 20 settlement days before the shock. The post-shock period coincides with 20 days after the episode started. Estimated coefficients are reported in Table 7.

In the first column we do not include fixed effects, in the second column we do. In the upper panel we use the raw daily net position of the bank. In the lower panel we consider the regularized time series used in ReNoSCh, $\tilde{y}_{i,t}$ instead of $y_{i,t}$. If we look at the raw data without including episode fixed effects (the upper left panel) it seems that most of the additional outflows are generated by the deferred transfers, real-time payments are not even significant. When we control for regular pattern in payment data and include fixed effects, the magnitudes are pretty comparable (the lower right panel). When the shock triggers banks experience an average outflow of 0.06 and 0.05 percent of deposits per day respectively for deferred and real-time payments.

Subsequent natural question would be: how much does real-time settlement accelerate the liquidity pressure? Given that we exactly know the time stamp of each transaction, we can change them to construct counterfactuals, averaging data across the episodes detected. To give a flavor from our sample, we depict in Figure 3 two counterfactuals.

Table 7: Speed of Digital Transfers - Time to Settlement

Dependent: daily bank-specific net position in			
		Raw data	
Real-time Settlement			
	Distress	0.0083	-0.0602 ***
		(0.0146)	(0.0180)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	2.016	2.0721
	Observations	9,804	9,804
Deferred Net Settlement			
	Distress	-0.0469 ***	-0.0722 ***
		(0.0173)	(0.0195)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	2.004	2.0503
	Observations	9,804	9,804
		Regularized data	
Real-time Settlement			
	Distress	-0.0683 ***	-0.0529 ***
		(0.0163)	(0.0215)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	1.9682	2.022
	Observations	9,804	9,804
Deferred Net Settlement			
	Distress	-0.0725 ***	-0.0647 ***
		(0.0153)	(0.0181)
	Bank FE	No	Yes
	Day FE	No	Yes
	DW	1.995	2.0415
	Observations	9,804	9,804

Notes. * : $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$. OLS estimates of coefficients of a δ. Heteroscedasticity and autocorrelation consistent estimates for all the standard errors. The dependent is the daily net position in euro respectively for real-time settlement transfers and deferred net settlement transfers computed as the sum of credits minus the sum of debits for each bank hit by a shock. The coefficients on different lines refer to separate regressions and are reported in percentage points of deposits of the shocked bank. R^2 and other coefficients of each regression are not reported for brevity. The pre-shock period is the 20 settlement days before the shock. The post-shock period is 20 days after. We pooled together all the identified episodes.

The light blue line depicts the cumulated outflows generated in the first 20 days of a episode if all the payments were settled on a deferred basis. The blue line depicts the cumulated outflows generated in the same time span if all the payments were settled real-time. To construct the first (second) counterfactual we postpone (anticipate) the settlement time of all real-time (deferred) payments by one day, leaving unchanged the other category, average across all the episodes and cumulate over time. The difference between the cumulated outflows is represented by the dark blue area. Real-time settlement increases the immediate liquidity pressure by about 0.33 percent of deposits in the first 10 days, about 10 percent of the total average drain generated. To then decrease by about a half in the following 10 days. This is a significant magnitude, and if the bank has not excessive reserves to cover such outflows, it has to resort on the money market or on the central bank or on the liquidation of assets. In the next section we provide evidences on the liquidity sources used by banks in our sample.

Figure 3: Counterfactuals - Pure Real-time vs Deferred Settlement

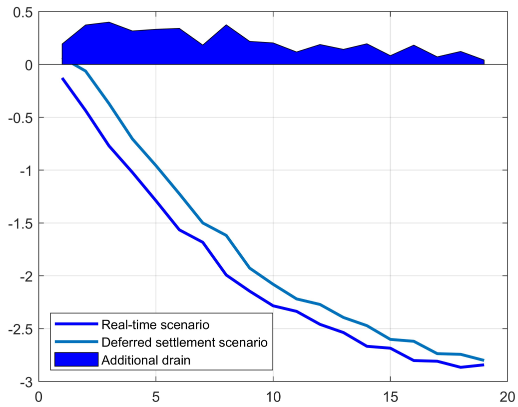

Notes. x-axis: days. y-axis: cumulated outflow expressed as a percentage of deposits. Day 0 is the first day of deposits outflows estimated by ReNoSCh. The lines represent the average (across episodes) cumulated net positions in two scenarios. The blue line depicts a scenario where everything is settled real-time. The light blue line depicts a scenario where everything is settled deferred. In the first scenario the observed DNST are shifted one day before, in the second the observed RTST are shifted one day after. The dark blue area is the difference between the two scenarios.

4.2 Wholesale Funding and Central Bank Operations

So far our analysis focused on the the so-called "traditional" concept of bank run (Diamond and Dybvig, 1983). After the global financial crisis, the literature has focused on a so-called "modern" mechanisms of bank run, which is concerned with wholesale funding, especially in the interbank segment where other banks exert monitoring and market discipline, rather than customers' deposit withdrawals that are typically thought to be sticky. After having shown with high frequency data that such withdrawals are not so sticky, we now study the interaction between retail and wholesale funding, to see whether outflows in retail deposits are compensated, exacerbated or anticipated by wholesale funding.

In practice, deposits' distress implies a drain of bank's reserves. If the drain is significant, the bank can not ultimately be able to honor its obligations. Thus when distress occurs the bank has to seek for liquidity. From a balance sheet perspective, the bank needs to substitute deposits with other liabilities or to shrink its assets, both implies injections of central bank money in its reserve account to compensate the drain. If also funds on the wholesale market fly away, the liquidity crisis of the bank can be more dramatic and let it finally fail.

In the identified episodes, bank's reserves did not dramatically drop even for one day -and this is also why the level of reserves is not a good indicator to timely identify them-, meaning that banks did offset the liquidity drain immediately. Here we are interested in understanding how.

In principle, the bank has several options. Here we discuss three main ways: the unsecured money market, the use of collateral in money and financial markets and operations with the central bank.

The unsecured money market used to be the most important channel to reallocate liquidity among banks. Even if the market was dramatically hit by the 2007-08 global financial crisis and the Sovereign debt crisis, it did not totally freeze (Afonso et al., 2011; Angelini et al., 2011; Rainone, 2017).[27] In TARGET2 we can identify precisely the unsecured money market operations of the e-MID.

Another way to get liquidity in money and financial markets is from securities. The bank can sell or pledge them to get a secured loan. The interest rate in the unsecured money market can be much higher or it would even be difficult to find a lender during distress episodes. Collateral can be a remedy. Indeed, we observed a shift from unsecured interbank market to central counterparty clearing (CCP) during the global financial crisis and the sovereign debt crisis.[28] Since it is quite complex to disentangle securities selling from repo activity of banks on a daily basis,[29] and we yearn to use payment system data to easily follow up on ReNoSCh outputs, here we consider them jointly and proxy the liquidity obtained from securities (selling and pledging) using the cash leg of securities exchanges on central securities depositories in TARGET2.

Alternatively, the commercial bank can go to the central bank and borrow funds. A core function of central banks is to act as a 'lender of last resort' to the banking system (Garcia-de Andoain et al., 2016). In the US, the Federal Reserve uses the Discount Window (DW) to fulfill this task, while the Eurosystem uses the Marginal Lending (ML) against collateral. Historically, both the DW and ML have been little used, even when banks faced acute liquidity shortages.[30] In October 2011, during the sovereign debt crisis, the Eurosystem adopted additional monetary policy measures in the form of a commitment to continue the fixed-rate full allotment policy initiated during the financial crisis of 2007-2008.[31] Under fixed-rate full allotment, counterparties have their bids fully satisfied, against adequate collateral, and on the condition of financial soundness. The Eurosystem's regular open market operations (OMO) consist of one-week liquidity-providing operations in euro (main refinancing operations, or MROs) as well as three-month liquidity-providing operations in euro (longer-term refinancing operations, or LTROs).[32] These operations are then much more attractive for a bank, especially because the rate is lower than the ML rate.The only drawback is that MROs are done weekly, and not daily like the ML. This limit can be particularly problematic in the case of a sudden and fast or 'instant' distress.

Usually all these operations are settled in the RTGS, as they involve large value payments that need to be settled in real-time on reserve accounts. From this standpoint, using payment system data offers a wide view on funding sources. Transactional data in TARGET2 allows us to identify payments related to these three liquidity channels.

To understand which source was most used, we regress the net daily bank positions on these three

[27]Since the 2008 financial crisis the overall interest in the linkages between banks has risen (Ashcraft and Duffie, 2007; Cocco et al., 2009; Furfine, 2003; Hartmann et al., 2001; Iori et al., 2008; Rainone, 2020; Soramäki et al., 2007).

[28]See Mancini et al. (2015) and Piquard and Salakhova (2019) for an analysis of the determinants.

[29]A repurchase agreement (repo) is a form of short-term borrowing for dealers in securities. The dealer sells the securities to investors, usually on an overnight basis, and buys them back the following day.

[30]Although other explanations may exist, this lack of borrowing is commonly attributed to stigma (Armantier et al., 2015; Bernanke, 2009).

[31]This tool is significantly different from the term auction facility (TAF, see McAndrews et al., 2017; Taylor and Williams, 2009; Wu, 2011, for more details) implemented by the FED.

[32]MROs serve to steer short-term interest rates, to manage the liquidity situation and to signal the monetary policy stance in the euro area, while LTROs provide additional, longer-term refinancing to the financial sector.

channels separately. For the e-MID, we use the daily variation in the outstanding position of the bank as a borrower minus the position as a lender (UM). To capture the amount of liquidity got from the use of collateral, we used the daily net position of the bank in TARTET2-Securities (CO), including the transactions with the CCP.[33] If positive, it means that the bank sold securities or borrowed money through a repo. For monetary policy operations, we computed the daily change in the outstanding amount of net liquidity borrowed from the central bank (CB). We pool together all the identified episodes again and regress the change in daily positions of the whole banking community (distressed and non-distressed banks) on a dummy equal to one after the distress bank is hit by the shock,

$$h_{i,t} = \delta r_{i,t} + \alpha_i + \mu_t + a_{i,t}, \tag{9}$$

where the dependent is the change in daily (t) net positions of the bank in the unsecured interbank market ($h_{i,t} = UM_{i,t}$) or collateral operations ($h_{i,t} = CO_{i,t}$), or central bank operations ($h_{i,t} = CB_{i,t}$) or their sum. $r_{i,t}$ is a dummy that switches to one after the shock hits bank i and zero otherwise, α_i is a bank fixed effect, μ_t is a day fixed effect, and $a_{i,t}$ is the error term. The pre-shock period is 20 settlement days before the distress. The post-shock period coincides with 20 days after the distress started. Table 8 reports our results. On average the bank got a daily liquidity inflow from all these channels equal to 0.33 percent of deposits during distress days. Almost all of the offsetting liquidity came from OMO.

To better get a sense of the timing of these operations, we leverage again the high frequency of our data. In Figure 4 we plot the average daily cumulated net position of banks for each channel. Day 0, tracked by a horizontal red line, is the first day in which ReNoSCh identified a shock. As before we considered 100 days before and 20 after. The yellow line reports the net position of the banks in the unsecured money market. The orange line is the collateral source. The light blue line represent the outstanding position in the OMO. Also visually, it is clear that the latter played the major role in offsetting the liquidity drain. It looks as a step function because the auctions are not available everyday, as mentioned before.

This figure is somehow specular to Figure 4 and represents a sort of reassuring picture in terms of correct identification of timing and offsetting liquidity sources as the numbers roughly square: an amount of liquidity inflows equal to about 3 percent of deposits in 20 days to counterbalance a 3 percent of deposits flowing out in the same 20 days.

[33] Both general collateral (GC) and special repo (SR) transactions are included. The CCP margins are excluded.

Table 8: Offsetting the Liquidity Drain - Sources of Funding

Dependent: daily net position in		
All sources	Distress	0.3280 ***
		(0.1523)
	Run FE	Yes
	Day FE	Yes
	DW	2.973118795
	Observations	9,804
OMO	Distress	0.3223 ***
		(0.1491)
	Run FE	Yes
	Day FE	Yes
	DW	2.9541
	Observations	9,804
Collateral	Distress	0.0034
		(0.0629)
	Run FE	Yes
	Day FE	Yes
	DW	2.0083
	Observations	9,804
Unsecured	Distress	0.0025
		(0.0073)
	Run FE	Yes
	Day FE	Yes
	DW	2.0423
	Observations	9,804

Notes. * : $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$. OLS estimates of δ. Heteroscedasticity and autocorrelation consistent estimates for all the standard errors. The coefficients on different lines refer to separate regressions and are reported in percentage points of deposits of the shocked bank. R^2 and other coefficients of each regression are not reported for brevity. We pooled together all the identified episodes, all the specifications include fixed effects. Collateral is computed by taking the net cash position of each bank on transactions settled through the CSDs, it includes the the transactions with the CCPs. The net cash position in the unsecured channel is computed by using all the e-MID transactions settled by each bank.

Figure 4: Offsetting the Liquidity Drain

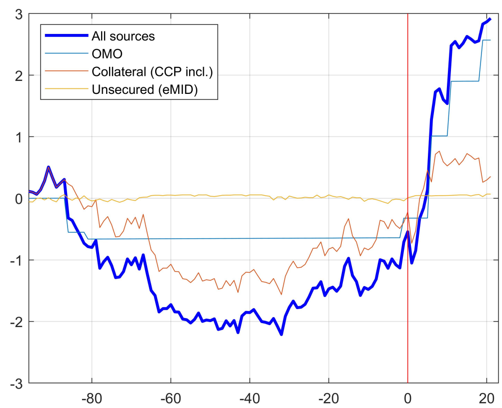

Notes. x-axis: days. y-axis: cumulated flow expressed as a percentage of deposits. Day 0 (the red vertical line) is the first day of deposits' outflow estimated by ReNoSCh. Cumulated net positions for each channel averaged across all the cases identified. OMO stands for open market operations, the outstanding position is represented. 'Collateral' represents the net position for securities related settlements. It includes trades and repo contracts, CCP included. 'Unsecured' represents the net outstanding position in the unsecured money market in e-MID. The blue bold line is the sum of the all the sources.

It highlights that these monetary policy tools are effective in smoothing out bank specific liquidity shocks timely in the post-crisis environment in which the unsecured money market almost disappeared in favor of the secured one. Such evidences complement evidences on the effects of central bank interventions on unsecured wholesale funding during macro shocks, such those in Brunetti et al. (2010).

One could have expected a more prominent role played by the secured money market. At the end of the day, a repo has features similar to a OMO. The bank needs collateral for both operations, and the rate of GC repos has often been below the MRO rate. In addition if there is a stigma effect, it is better to get money from the market than from the central bank, especially if the counterparty is anonymous and the contract is secured. Nevertheless, there are important differences between repo and OMO. The first is the potential length of the maturity. The second is the barrier to the market, a bank has to be a member of the CCP and have proper infrastructures to participate to the secured market. The third is the type of collateral. While the repo market is thick for GC, especially for government bonds, it is much more thin for less popular securities. In this sense, the rate for GC may not be fully comparable with the MRO rate, because the GC is for specific securities, while for the MRO the bank can use a wider set of eligible assets.[34] Another fundamental aspect is the computation of the liquidity cover ratio (LCR) within the Basel III framework. When the banks have to estimate the total net cash outflows over the next 30 calendar days, in the computation of the secured funding run-off, the amount to add to cash outflows for outstanding maturing secured funding transactions with the central bank is 0 percent (Basel III, 2013). The same amount is reachable only for secured money market funding backed by Level 1 assets, which may be scarce. Level 2A assets add 15 percent, while Level 2B add 50 percent. This feature reflects the different loan roll-over probability. While the secured market may suddenly dry-up, this does not hold for the central bank, which plays the role of lender-of-last-resort (see Garcia-de Andoain et al., 2016).

The evidences collected in this section help us to better characterize distress episodes in post global financial crisis environment. Oppositely to shortage of unsecured interbank liquidity observed before and during the global financial crisis (Afonso et al., 2011; Angelini et al., 2011; Heider et al., 2015), the reliance on secured money markets and central bank funding, implied by the crisis and the subsequent regulation and policies, provided immediate funding, offsetting (instead of exacerbating) the negative liquidity shock. This does not mean that the banks were not distressed, because the amount of deposits flying away was significant and close to Basel run-off rates. It means that distress could have been much higher, potentially impacting other institutions, if banks did not have enough high quality collateral and access to secured market platforms and central bank facilities.

5 Conclusion

This paper's contribution is twofold.

First, we propose a new methodology to identify distress episodes in real-time using payment system

[34]The comparability of the two baskets depends on the CCP policy at a specific time. Some CCPs offer secured transactions using the same standardized baskets as the ECB, but others might not.

data. More specifically, we illustrated (i) how to measure deposits' flows in RTGS systems; (ii) an algorithm that is able to identify bank runs and quantify their severity in real-time; (iii) its good performance in numerical simulations and (iv) with real data from TARGET2.

Second, we show (i) the existence of distress episodes and their significance in terms of liquidity risk for banks; (ii) the major role played by digital transfers to other banks w.r.t. cash withdrawals; (iii) the positive liquidity spillovers to institutions not distressed and in particular to large domestic banks; (iv) the importance of real-time settlement in accelerating the speed of distress and finally; (v) the interaction between retail and wholesale funding during idiosyncratic distress episodes. In particular, we show how banks offset the liquidity drain generated by a retail bank run with wholesale funding.

The results are relevant from several policy standpoints. Our estimates of deposits outflows are informative for the calculation of 30 days potential outflows due to retail deposit run-off in the LCR (Basel III, 2013) and their potential evolution in an increasingly digital financial world. The diffusion of technological innovations, like real-time settlement, instant payments and central bank digital currency (CBDC) have to be monitored and taken into account. We showed that the majority of depositors preferred digital balances at other intermediaries to cash conversion. If such preference was mainly driven by the inconvenience of storing banknotes, the introduction of CBDC would constitute an appealing alternative to deposits at other intermediaries (Panetta, 2018b; Weidmann, 2018), removing the direct positive spillovers to big domestic intermediaries uncovered in this study. For these reasons, some kind of disincentive to convert deposits to CBDC could be envisaged (Bindseil, 2020; Bindseil et al., 2021; Panetta, 2018a).

The paper shows how the post-crisis policy framework and interbank market structure interact with bank-specific idiosyncratic distress. Oppositely to pre-crisis unsecured wholesale funding, post-crisis collateralized wholesale funding helped absorbing, instead of aggravating, the shock. Central bank facilities guaranteed the timely provision of liquidity against high quality collateral. Even if not specifically designed for this purpose, fixed-rate full allotment open market operations were mostly used to replenish the liquidity drain generated by deposits outflows, avoiding penalizing fire sales. The weekly auctions proved to be a quite efficient parachute, even in the presence of fast deposit outpouring. Such evidence informs the debate on the role of 'automatic' lender-of-last resort (LOLR) that the central bank can play for example with the introduction of CBDC (Brunnermeier and Niepelt, 2019; **?**). Most pragmatically, the paper provides an operational tool to trigger a real-time LOLR, and highlights the idea that the introduction of CBDC is not the only method to observe deposits outflows in real-time. Indeed, we showed that RTGS data, owned by and readily available to central banks, can be used for this purpose. The proposed algorithm can be used to detect in real-time distress episodes and save significant social costs. This can allow supervisors to act promptly in case of idiosyncratic and systemic risk. If cloud computing and API (application programming interface) are available, such big payment data can be operational and become an effective SupTech tool (Broeders and Prenio, 2018).

We showed that real-time payments have side effects, they are not only an important tool to mitigate credit risk in payments between banks (Kahn and Roberds, 1998), but they can also increase liquidity risk of banks during distress episodes by decreasing more quickly their reserves.

In this regard, the recent diffusion of instant payments, which offers the possibility to settle continuously, can be a structural break. While in an increasingly digital world banks can just shut their digital doors to prevent massive outflows, depositors may anticipate this reaction and withdraw deposits in advance. Regulators have to monitor these phenomena closely to mitigate new risks coming from the digitalization of financial services.

As a final policy remark, our evidence that depositors withdraw small denomination banknotes is surprising and should raise concerns, because most of the accounts of these depositors are likely to be covered by the deposit insurance. While this fact can be just due to unawareness of the deposit guarantee scheme, it may also point to distrust of the national banking system and the ability of the government to refund them in a short time. This is particularly possible in the Eurosystem where banknotes are guaranteed by a supranational authority, the Eurosystem, while deposits are not.

References

Accornero, M. and Moscatelli, M. (2018). Listening to the buzz: social media sentiment and retail depositors' trust, *Bank of Italy Temi di Discussione (Working Paper) No* **1165**.

Afonso, G., Kovner, A. and Schoar, A. (2011). Stressed, not frozen: The federal funds market in the financial crisis, *The Journal of Finance* **66**(4): 1109–1139.

Alessi, L. and Detken, C. (2017). Identifying excessive credit growth and leverage, *Journal of Financial Stability* .

Allen, F., Carletti, E. and Gale, D. (2009). Interbank market liquidity and central bank intervention, *Journal of Monetary Economics* **56**(5): 639–652.

Andreou, E. and Ghysels, E. (2009). Structural breaks in financial time series, *Handbook of financial time series*, Springer, pp. 839–870.

Angelini, P., Nobili, A. and Picillo, C. (2011). The interbank market after august 2007: what has changed, and why?, *Journal of Money, Credit and Banking* **43**(5): 923–958.

Arduini, T., De Arcangelis, G. and Del Bello, C. L. (2012). Balance-of-payments crises during the great recession: Is this time different?, *Review of International Economics* **20**(3): 517–534.

Armantier, O., Ghysels, E., Sarkar, A. and Shrader, J. (2015). Discount window stigma during the 2007–2008 financial crisis, *Journal of Financial Economics* **118**(2): 317–335.

Ashcraft, A. B. and Duffie, D. (2007). Systemic illiquidity in the federal funds market, *American Economic Review* **97**(2): 221–225.

Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes, *Econometrica* pp. 47–78.

Bai, J. and Perron, P. (2003). Computation and analysis of multiple structural change models, *Journal of Applied Econometrics* **18**(1): 1–22.

Bartiloro, L. (2011). Is your money safe? what italians know about deposit insurance, *Bank of Italy Questioni di Economia e Finanza (Occasional Paper) No* **104**.

Basel III, B. (2013). The liquidity coverage ratio and liquidity risk monitoring tools, *Bank for International Settlements* .

Berg, A. and Pattillo, C. (1999). Are currency crises predictable? a test, *IMF Staff papers* **46**(2): 107–138.

Bernanke, B. (2009). The federal reserve's balance sheet: an update. in: Speech at the federal reserve board conference on key developments in monetary policy, washington.
URL: *http://www.federalreserve.gov/newsevents/speech/bernan ke20091008a.htm*

Bindseil, U. (2020). Tiered cbdc and the financial system.

Bindseil, U., Panetta, F. and Terol, I. (2021). Central bank digital currency: functional scope, pricing and controls, *ECB Occasional Paper* (2021/286).

Broeders, D. and Prenio, J. (2018). Innovative technology in financial supervision (suptech): The experience of early users.

Brown, M., Trautmann, S. T. and Vlahu, R. (2016). Understanding bank-run contagion, *Management Science* **63**(7): 2272–2282.

Brunetti, C., Di Filippo, M. and Harris, J. H. (2010). Effects of central bank intervention on the interbank market during the subprime crisis, *The review of financial studies* **24**(6): 2053–2083.

Brunnermeier, M. K. and Niepelt, D. (2019). On the equivalence of private and public money, *Journal of Monetary Economics* **106**: 27–41.

Bryant, J. (1980). A model of reserves, bank runs, and deposit insurance, *Journal of banking & finance* **4**(4): 335–344.

Cañón, C. and Margaretic, P. (2014). Correlated bank runs, interbank markets and reserve requirements, *Journal of Banking & Finance* **49**: 515–533.

Chen, Y. and Hasan, I. (2006). The transparency of the banking system and the efficiency of information-based bank runs, *Journal of Financial Intermediation* **15**(3): 307–331.

Choi, H. and Varian, H. (2012). Predicting the present with google trends, *Economic Record* **88**(s1): 2–9.

Chow, G. C. (1960). Tests of equality between sets of coefficients in two linear regressions, *Econometrica* pp. 591–605.

Cocco, J. F., Gomes, F. J. and Martins, N. C. (2009). Lending relationships in the interbank market, *Journal of Financial Intermediation* **18**(1): 24–48.

Cooper, R. and Ross, T. (1998). Bank runs: Liquidity costs and investment distortions, *Journal of Monetary Economics* **41**.

Craig, B. R., Salakhova, D. and Saldias, M. (2018). Payments delay: propagation and punishment.

D'Amuri, F. and Marcucci, J. (2017). The predictive power of google searches in forecasting us unemployment, *International Journal of Forecasting* **33**(4): 801–816.

Dávila, E. and Goldstein, I. (2020). Optimal deposit insurance, *Working Paper* .

Denbee, E., Garratt, R. and Zimmerman, P. (2014). Variations in liquidity provision in real-time payment systems.

Denbee, E., Julliard, C., Li, Y. and Yuan, K. (2021). Network risk and key players: A structural analysis of interbank liquidity, *Journal of Financial Economics* .

Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity, *Journal of political economy* **91**(3): 401–419.

Engineer, M. (1989). Bank runs and the suspension of deposit convertibility, *Journal of monetary Economics* **24**(3): 443–454.

Ennis, H. M. and Keister, T. (2006). Bank runs and investment decisions revisited, *Journal of monetary Economics* **53**(2): 217–232.

Ennis, H. M. and Keister, T. (2009). Bank runs and institutions: The perils of intervention, *American Economic Review* **99**(4): 1588–1607.

Ferrante, F. (2018). A model of endogenous loan quality and the collapse of the shadow banking system, *American Economic Journal: Macroeconomics* **10**(4): 152–201.

Freixas, X., Parigi, B. M. and Rochet, J.-C. (2000). Systemic risk, interbank relations, and liquidity provision by the central bank, *Journal of money, credit and banking* pp. 611–638.

Furfine, C. (2003). Interbank exposures: Quantifying the risk of contagion, *Journal of Money, Credit, and Banking* **35**(1): 111–128.

Garcia-de Andoain, C., Heider, F., Hoerova, M. and Manganelli, S. (2016). Lending-of-last-resort is as lending-of-last-resort does: Central bank liquidity provision and interbank market functioning in the euro area, *Journal of Financial Intermediation* .

Gertler, M. and Kiyotaki, N. (2015). Banking, liquidity, and bank runs in an infinite horizon economy, *American Economic Review* **105**(7): 2011–43.

Gertler, M., Kiyotaki, N. and Prestipino, A. (2016). Wholesale banking and bank runs in macroeconomic modeling of financial crises, *Handbook of Macroeconomics*, Vol. 2, Elsevier, pp. 1345–1425.

Goldsmith-Pinkham, P. and Yorulmazer, T. (2010). Liquidity, bank runs, and bailouts: spillover effects during the northern rock episode, *Journal of Financial Services Research* **37**(2-3): 83–98.

Gropp, R., Hakenes, H. and Schnabel, I. (2010). Competition, risk-shifting, and public bail-out policies, *The Review of Financial Studies* **24**(6): 2084–2120.

Haldane, A. G., Millard, S. and Saporta, V. (2008). *The future of payment systems*, Routledge.

Hartmann, P., Manna, M. and Manzanares, A. (2001). The microstructure of the euro money market, *Journal of International Money and Finance* **20**(6): 895–948.

Heider, F., Hoerova, M. and Holthausen, C. (2015). Liquidity hoarding and interbank market rates: The role of counterparty risk, *Journal of Financial Economics* **118**(2): 336–354.

Heijmans, R. and Heuver, R. (2014). Is this bank ill? the diagnosis of doctor target2, **1(2)**: 3–8.

Heijmans, R., Heuver, R., Levallois, C. and van Lelyveld, I. (2014). Dynamic visualization of large transaction networks: the daily dutch overnight money market.

Iacobini, A. (1994). Filter effect in the use of control charts for productive process control, *Statistica* **54**(1): 113–124.

Iori, G., De Masi, G., Precup, O. V., Gabbi, G. and Caldarelli, G. (2008). A network analysis of the italian overnight money market, *Journal of Economic Dynamics and Control* **32**(1): 259–278.

Ippolito, F., Peydró, J.-L., Polo, A. and Sette, E. (2016). Double bank runs and liquidity risk management, *Journal of Financial Economics* **122**(1): 135–154.

Ivashina, V. and Scharfstein, D. (2010). Bank lending during the financial crisis of 2008, *Journal of Financial economics* **97**(3): 319–338.

Iyer, R. and Puri, M. (2012). Understanding bank runs: The importance of depositor-bank relationships and networks, *American Economic Review* **102**(4): 1414–45.

Iyer, R., Puri, M. and Ryan, N. (2016). A tale of two runs: Depositor responses to bank solvency risk, *The Journal of Finance* **71**(6): 2687–2726.

Kahn, C. M. and Roberds, W. (1998). Payment system settlement and bank incentives, *The Review of Financial Studies* **11**(4): 845–870.

Kaminsky, G. L. and Reinhart, C. M. (1999). The twin crises: the causes of banking and balance-of-payments problems, *American economic review* **89**(3): 473–500.

Keister, T. and Narasiman, V. (2016). Expectations vs. fundamentals-driven bank runs: When should bailouts be permitted?, *Review of Economic Dynamics* **21**: 89–104.

Kiss, H. J., Rodriguez-Lara, I. and Rosa-García, A. (2014). Do social networks prevent or promote bank runs?, *Journal of Economic Behavior & Organization* **101**: 87–99.

Kokkola, T. (2010). *The payment system: Payments, securities and derivatives, and the role of the Eurosystem*, European Central Bank.

Lorenzen, T. J. and Vance, L. C. (1986). The economic design of control charts: a unified approach, *Technometrics* **28**(1): 3–10.

Lowry, C. A. and Montgomery, D. C. (1995). A review of multivariate control charts, *IIE transactions* **27**(6): 800–810.

Lowry, C. A., Woodall, W. H., Champ, C. W. and Rigdon, S. E. (1992). A multivariate exponentially weighted moving average control chart, *Technometrics* **34**(1): 46–53.

Mancini, L., Ranaldo, A. and Wrampelmeyer, J. (2015). The euro interbank repo market, *The Review of Financial Studies* **29**(7): 1747–1779.

Martin, A. (2006). Liquidity provision vs. deposit insurance: preventing bank panics without moral hazard, *Economic Theory* **28**(1): 197–211.

Martin, A., Skeie, D. and Thadden, E.-L. v. (2014). Repo runs, *The Review of Financial Studies* **27**(4): 957–989.

Martin, C., Puri, M. and Ufier, A. (2017). Deposit inflows and outflows in failing banks: The role of deposit insurance, *Technical report*, mimeo.

McAndrews, J., Sarkar, A. and Wang, Z. (2017). The effect of the term auction facility on the london interbank offered rate, *Journal of Banking & Finance* **83**: 135–152.

Montgomery, D. C. (1980). The economic design of control charts: a review and literature survey, *Journal of Quality Technology* **12**(2): 75–87.

Nelson, L. S. (1984). The shewhart control chart-tests for special causes, *Journal of Quality Technology* **16**(4): 237–239.

Nelson, L. S. (1989). Standardization of shewhart control charts, *Journal of Quality Technology* **21**(4): 287–289.

Oliveira, R. d. F., Schiozer, R. F. and Barros, L. A. d. C. (2014). Depositors' perception of 'too-big-to-fail', *Review of Finance* **19**(1): 191–227.

Panetta, F. (2018a). 21st century cash: Central banking, technological innovation and digital currencies, *Do We Need Central Bank Digital Currencies* .

Panetta, F. (2018b). 21st century cash: Central banking, technological innovation and digital currencies, milano.
URL: *https://www.bancaditalia.it/pubblicazioni/interventi-direttorio/int-dir-2018/en-panetta-07062017.pdf?language$_i$d = 1*

Piquard, T. and Salakhova, D. (2019). Substitution between secured and unsecured interbank markets, counterparty risk and the opportunity cost of collateral, *Working Paper* .

Rainone, E. (2017). Pairwise trading in the money market during the european sovereign debt crisis, *Bank of Italy Temi di Discussione (Working Paper) No* **1160**.

Rainone, E. (2020). The network nature of over-the-counter interest rates, *Journal of Financial Markets* **47**: 100525.

Roberts, S. (1959). Control chart tests based on geometric moving averages, *Technometrics* **1**(3): 239–250.

Rosati, S. and Secola, S. (2006). Explaining cross-border large-value payment flows: Evidence from target and euro1 data, *Journal of Banking & Finance* **30**(6): 1753–1782.

Scheffe, H. (1947). The relation of control charts to analysis of variance and chi-square tests, *Journal of the American Statistical Association* **42**(239): 425–431.

Shewhart, W. A. (1926). Quality control charts, *Bell Labs Technical Journal* **5**(4): 593–603.

Soramäki, K., Bech, M. L., Arnold, J., Glass, R. J. and Beyeler, W. E. (2007). The topology of interbank payment flows, *Physica A: Statistical Mechanics and its Applications* **379**(1): 317–333.

Taylor, J. B. and Williams, J. C. (2009). A black swan in the money market, *American Economic Journal: Macroeconomics* **1**(1): 58–83.

Waldo, D. G. (1985). Bank runs, the deposit-currency ratio and the interest rate, *Journal of Monetary Economics* **15**(3): 269–277.

Weidmann, J. (2018). Opening speech of the fourth cash symposium of the deutsche bundesbank, frankfurt am main.
URL: *https://www.bundesbank.de/Redaktion/EN/Reden/2018/2018*$_0$2$_1$4$_w$*eidmann.html*

Weyman, J. (2016). Risks in faster payments, Retail Payments Risk Forum Working Paper, Federal Reserve Bank of Atlanta.

Wu, T. (2011). The us money market and the term auction facility in the financial crisis of 2007–2009, *Review of Economics and Statistics* **93**(2): 617–631.

Zeileis, A., Leisch, F., Hornik, K. and Kleiber, C. (2001). strucchange. an r package for testing for structural change in linear regression models.

Supplementary Online Appendix

REAL-TIME IDENTIFICATION AND HIGH FREQUENCY ANALYSIS OF BANK RUNS

Edoardo Rainone

Bank of Italy

edoardo.rainone@bancaditalia.it.

A Additional Figures and Tables

Figure A.1: A Simplified Schema of How Customers Can Move Deposits in Payment Systems

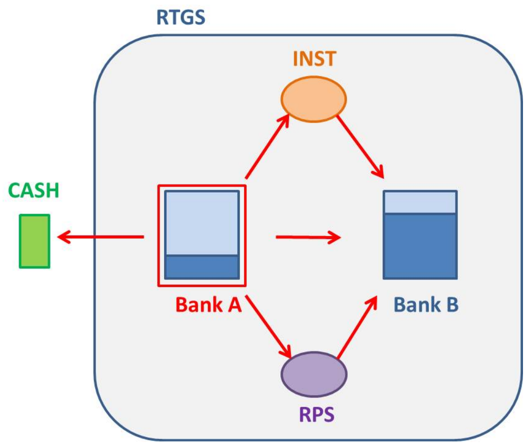

Notes. The big box is a stylized representation of a RTGS, where the central bank does its operations with commercial banks and where the latter have their reserve accounts (soft blue bins), all is settled in central bank money (dark blue in soft blue bins). The small violet circle represents a retail payment system (RPS). The small orange circle represents an instant payment system (INST). The red arrows represent outflows of central bank money from bank A reserve account, to bank B account via a RPS, INST, directly through the RTGS, or to banknotes.

Figure A.2: Parametric vs Nonparametric Densities with Real Data

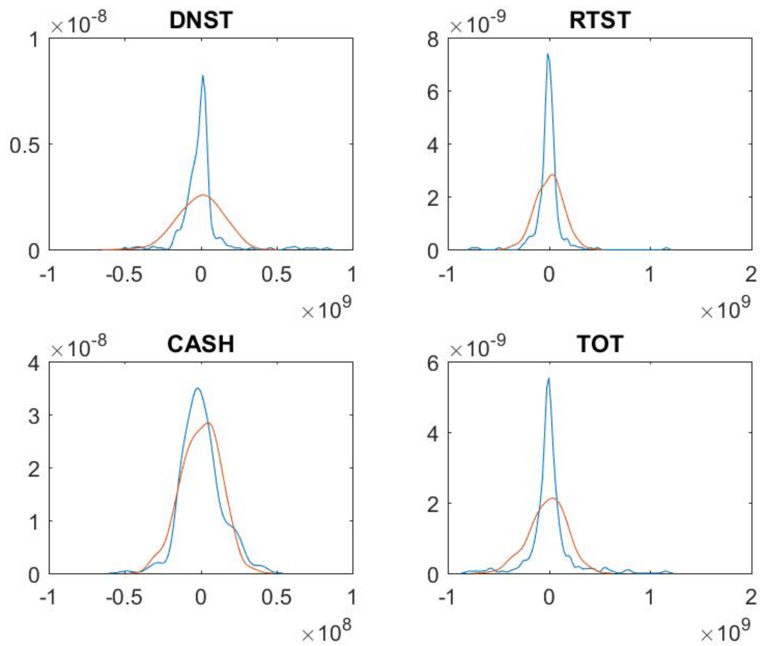

Notes. Empirical kernel densities in blue, theoretical normal distributions in orange. Densities of daily changes of each bank net position in each channel. *Deferred net settlement transfers* ($DNST$) are the bank's multilateral position settled by domestic and international RPSs. *Real-time settlement transfers* ($RTST$) are bilateral gross interbank transfers settled in real-time in central bank money for customer payments. *Cash withdrawals* ($CASH$) are cash operations by commercial banks exchanged with reserves in TARGET2. TOT is the sum of all positions.

Figure A.3: Sigma vs Nonparametric Approach

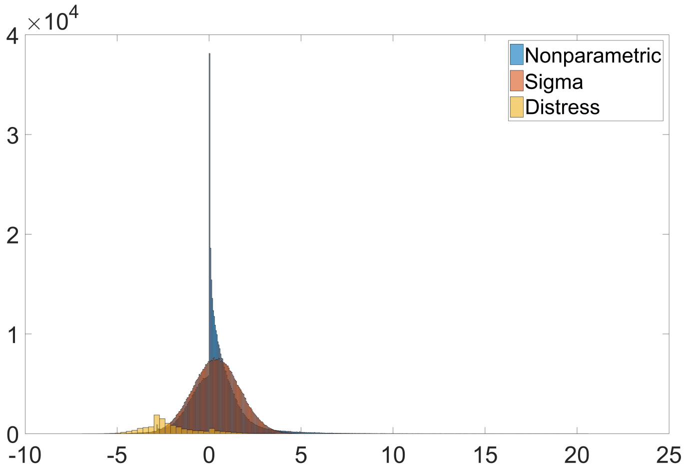

(a) $\zeta = 1$

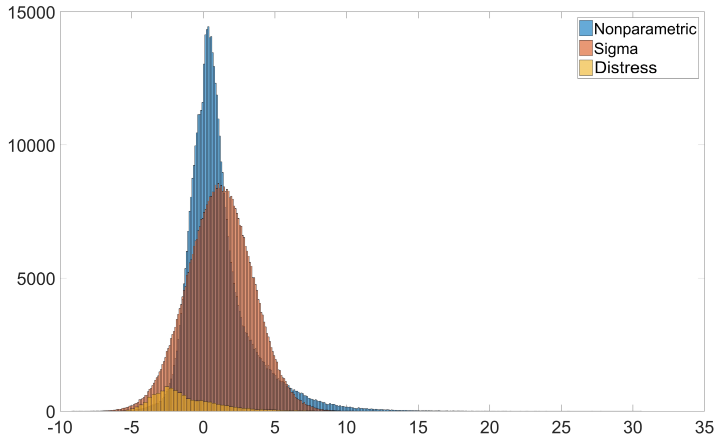

(b) $\zeta = 3$

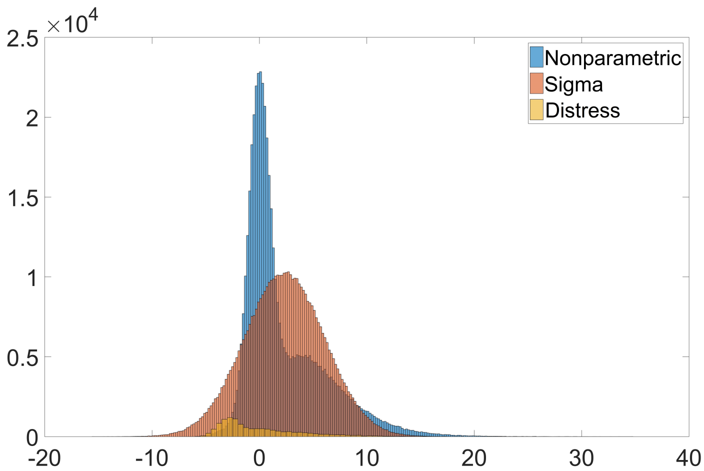

(c) $\zeta = 6$

Notes. Histograms computed on data from 1000 replications, with 500 processes with distress episodes and 500 without. The red histogram reports the inferred distribution using the sigma approach. The blue histogram reports the distribution inferred with the nonparametric approach. The yellow histograms reports the distribution of observations under distress. The DGP is described in Section 3.1. ζ are the degrees of freedom.

Figure A.4: Exogenous Break Points - Graphical Analysis.

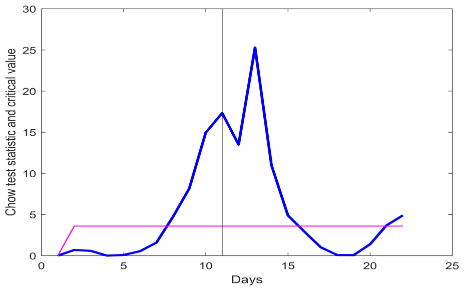

(a) All outflows

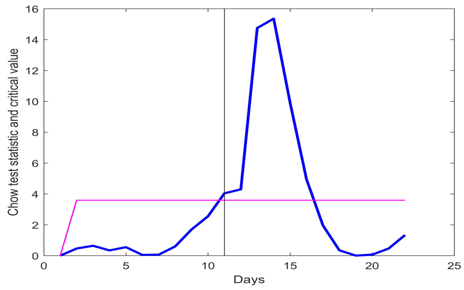

(b) Deferred

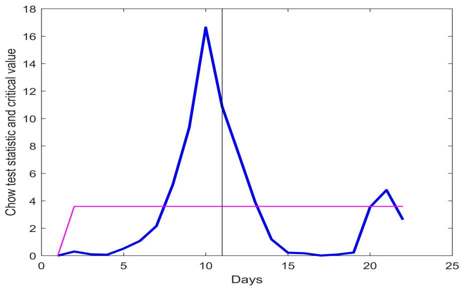

(c) Real-time

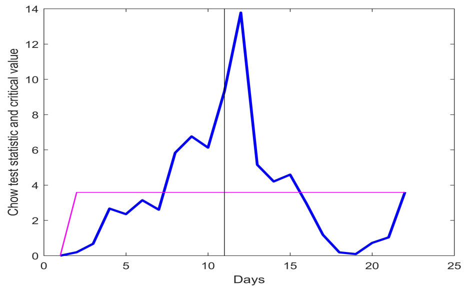

(d) Cash

Notes. Day 11 is the first estimated day of distress. Blue line: Chow test on every day on changes of net positions for each channel summed for all the cases considered. Time range for the test is 20 days around the pivotal day for which the test is computed. The position is centered to the estimated beginning (black vertical line). The violet line is the critical value.

Figure A.5: Endogenous Break Points - Graphical Analysis.

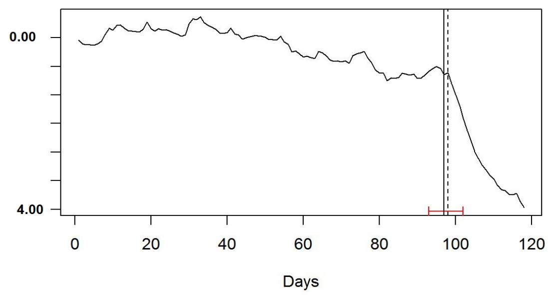

(a) All outflows

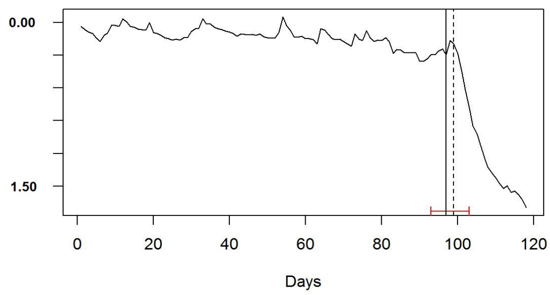

(b) Deferred

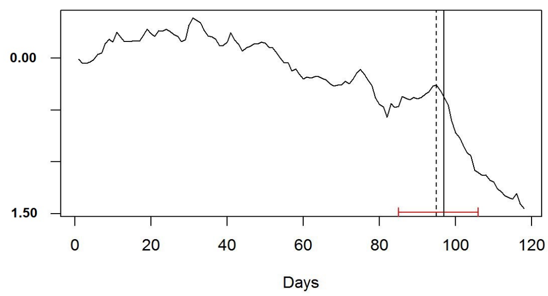

(c) Real-time

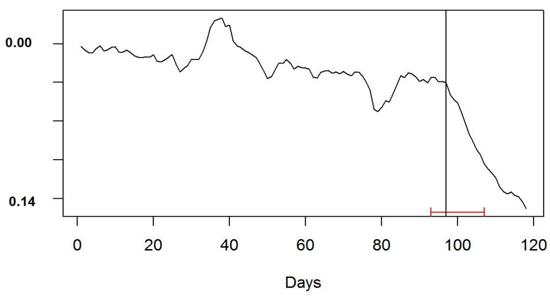

(d) Cash

Notes. Notes. x-axis: days. 100 days before and 20 after the beginning of the episode. y-axis: cumulated outflow expressed as a percentage of deposits. Day 100 is the first day of deposits outflows estimated by ReNoSCh, a black vertical line keeps track of it. Black line: cumulated net positions in euro for each channel averaged across all the cases considered expressed as a percentage of deposits. The test proposed by Bai and Perron (2003) is used to endogenously estimate the break point on changes in the net positions. The endogenous break point estimate is represented with a vertical dotted line. Its 95% confidence interval is in red.

Figure A.6: Routes to Safety

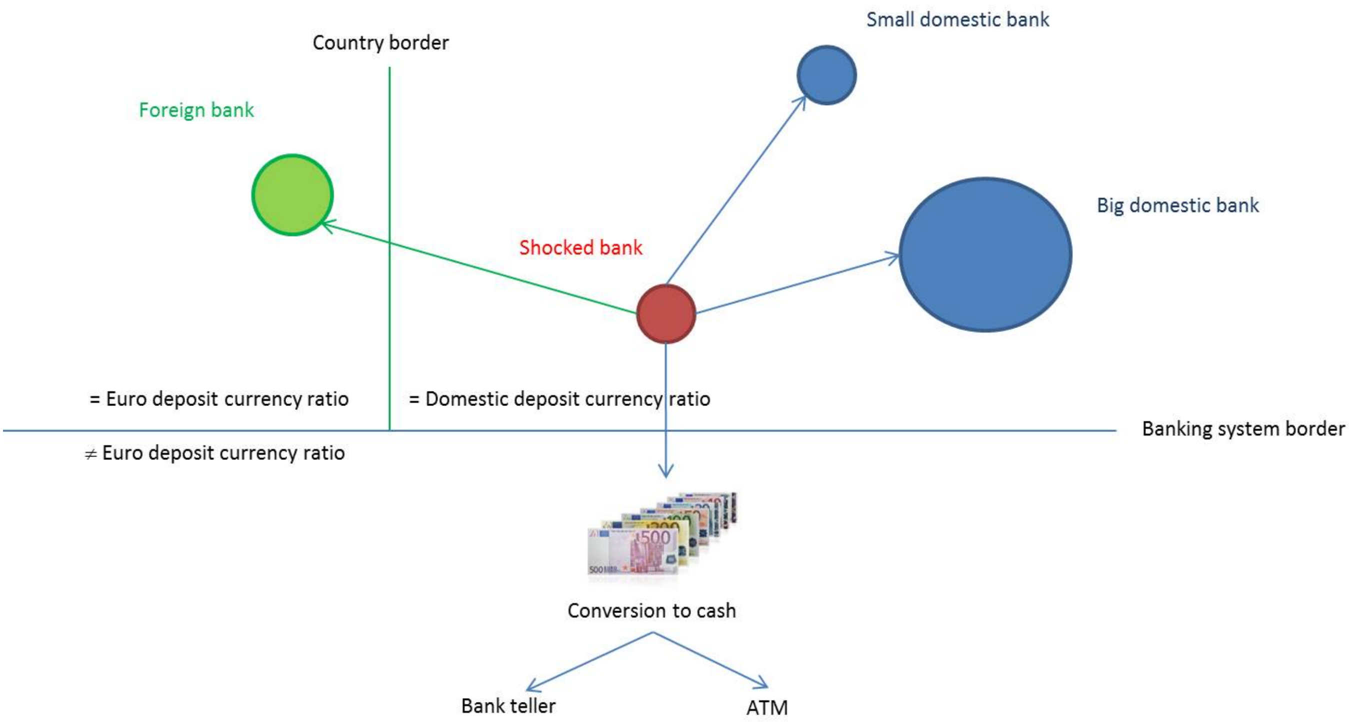

Figure A.7: Interest Rates Variation around Distress Episodes

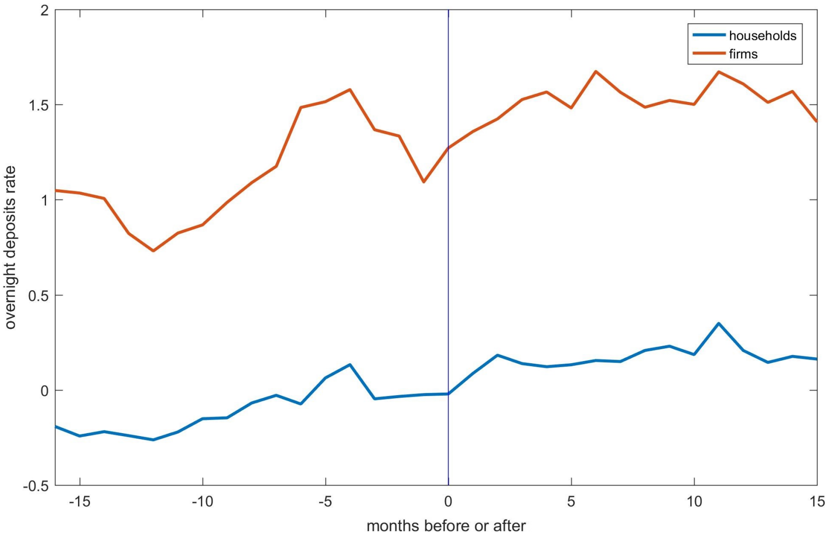

Notes. Y-axis: Rates on overnight deposit accounts in percentage points for household and firms obtain from the MIR database. X-axis: months around a distress episode. We consider 15 months before the episode and 15 months after. The vertical line represent the month in which the shock is detected. We take the spread from the average rate in the market and divide it by the standard deviation of the rate distribution for each month, in order to compare periods with high dispersion versus low dispersion periods and variation of average interest rates.

Table A.1: Number of Warning Days and Distress Detection

Number of warning days (s)	Sequence with distress																
	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0		
2					N	N	W	W	W	W	W	W	W	N	N	DISTRESS	OK
3					N	N	N	W	W	W	W	W	N	N	N	DISTRESS	OK
4					N	N	N	N	W	W	W	N	N	N	N	NO DISTRESS	NOK
	0	1	0	1	0	1	1	1	1	1	0	0	0	0	0		
2					W	W	W	W	W	W	W	W	W	N	N	DISTRESS	OK
3					N	W	W	W	W	W	W	W	N	N	N	DISTRESS	OK
4					N	N	N	W	W	W	W	N	N	N	N	NO DISTRESS	NOK
	1	0	1	0	0	1	1	1	1	1	0	0	0	0	0		
2					W	W	W	W	W	W	W	W	W	N	N	DISTRESS	OK
3					N	N	W	W	W	W	W	W	N	N	N	DISTRESS	OK
4					N	N	N	N	W	W	W	N	N	N	N	NO DISTRESS	NOK
	Sequence without a distress																
	1	0	1	0	0	1	0	1	0	1	0	0	0	0	0		
2					W	W	W	W	W	W	W	W	N	N	N	DISTRESS	NOK
3					N	N	N	N	N	W	N	N	N	N	N	NO DISTRESS	OK
4					N	N	N	N	N	N	N	N	N	N	N	NO DISTRESS	OK
	1	0	0	1	0	0	1	0	0	1	0	0	0	0	0		
2					W	N	W	W	N	W	W	N	N	N	N	DISTRESS	NOK
3					N	N	N	N	N	N	N	N	N	N	N	NO DISTRESS	OK
4					N	N	N	N	N	N	N	N	N	N	N	NO DISTRESS	OK
	1	0	0	1	0	1	0	0	1	0	1	0	0	1	0		
2					W	W	W	W	W	W	W	W	W	W	W	DISTRESS	NOK
3					N	N	N	N	N	N	N	N	N	N	N	NO DISTRESS	OK
4					N	N	N	N	N	N	N	N	N	N	N	NO DISTRESS	OK

Notes. 0 is a day not below the LCL, 1 is a day below th LCL. In this numerical example 5 consecutive days constitute distress. The first three sequences are characterized by distress, the last three are not. The first column reports the number of warning days in the last 5 needed to have a warning day. N means that the number of warning days in the last 5 days does not exceed the threshold, W means that it does and thus that day is a warning day. If we have a sequence of 5 consecutive W the second to last column reports DISTRESS. If it is a true distress episode the last column reports OK.

Table A.2: Stationarity before and after the Distress Episodes - ADF test

Dependent: daily bank-specific net position in			Raw	Regularized
Digital transfers				
	Pre			
		ADF stat	-4.7089	-7.5868
		p-value	0.0010	0.0010
	Post			
		ADF stat	-4.5339	-3.0197
		p-value	0.0010	0.0037
Cash				
	Pre			
		ADF stat	-2.4966	-6.3786
		p-value	0.0134	0.0010
	Post			
		ADF stat	-3.4448	-4.4594
		p-value	0.0010	0.0010

Notes. The dependent is the daily net position in euro respectively for digital transfers and cash withdrawals computed as the sum of credits minus the sum of debits for each bank hit by a shock. The augmented Dickey-Fuller test (ADF) test on different lines refer to separate regressions. The pre and post periods are the 20 settlement days before and after the shock.

Table A.3: Autoregression terms - Raw and Regularized Data

Dependent: daily bank-specific net position in	Raw	Residualized
Interbank digital transfers	0.0024	0.0061
	(0.0102)	(0.0102)
Cash	0.0402 ***	0.0167
	(0.0102)	(0.0102)
OMA	-0.0008	
	(0.0102)	
Collateral	-1.74E-05	
	(0.0102)	
Unsecured	0.0176 *	
	(0.0102	

Notes. * : $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$. OLS estimates of the lag term of a AR(1) model with constant and trend. We pooled together all the identified episodes. The coefficients on different lines refer to separate regressions. All have 9,804 observations. The dependent is the daily net position in euro computed as the sum of credits minus the sum of debits for each bank. The pre and post periods are the 20 settlement days before and after the shock. Collateral is computed by taking the net cash position of each bank on transactions settled through the CSDs, it includes the the transactions with the CCPs. The net cash position in the unsecured channel is computed by using all the e-MID transactions settled by each bank.

B Deposits and Reserves

When deposits move from a bank to another one, the operation involves a change in both assets and liabilities for both banks. In Figure B.1 a simple example is given. If a depositor of bank i transfers p to bank $-i$ (which can be thought as the rest of the banking sector), d_i decreases by p and d_{-i} increases by p on the liability side. On the assets side, the interbank transaction is settled using reserves held at the central bank, thus r_i decreases by p and r_{-i} increases by p. Also deposits converted in cash are traceable, in this case $-i$ is directly the central bank.

This mechanism is at the crux of banking intermediation in payments. The exchange of reserves guarantees that there is no counterparty risk left in the payment operation. Without access to and settlement in central bank money bank i and $-i$ should have bilateral accounts or use an asset not free from risk to clear the obligation. Both are undesirable from a systemic risk perspective and from the bank business perspective, as the latter is just selling a payment service and would avoid engaging in credit risk. See Directive 98/26/EC for more details on the obligations related to the settlement finality of payments in Europe.

Both deposits outflows related to distrust and regular outgoing payments imply such movements of reserves; we describe in Section 2.2 how to identify the first cause. The settlement in reserves is operated in the real-time gross settlement (RTGS) system, where banks' reserve accounts are digitally managed and updated. It follows that if we can track p from the reserve leg of the transaction, we obtain information on the deposit leg as well.

Figure B.1: Deposits and Reserves

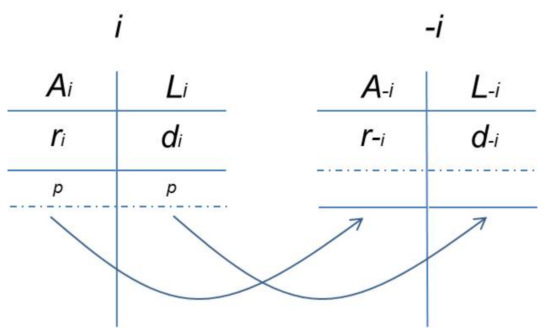

Notes. A simplified bank balance sheet is represented. A represents assets, L liabilities. i is a bank, $-i$ is the rest of the banking system. r_i are i's reserves, d_i are i's deposits.

C Source of Information in TARGET2

Here we describe in detail how to collect this type of data in TARGET2, the euro RTGS.[35] TARGET2 is based on an integrated central technical infrastructure, called the Single Shared Platform, it is operated by three national central banks: Bank of Italy, Banque de France and Deutsche Bundesbank. The implementation of TARGET2 was based on a decision of the ECB Council of autumn 2002. TARGET2 started operations on 19 November 2007, replacing TARGET. Central banks, commercial banks and other financial market infrastructures can submit payment in euro to TARGET2, where they are processed and settled in central bank money. European banks hold and manage their reserve accounts on TARGET2. Configurations in other RTGS systems, like FEDWIRE for the dollar, BOJ-NET for the yen or CHAPS for the sterling are not extremely different. We take advantage of transaction-level data for each participating bank, which allows us to reconstruct the banks' customers behavior in a detailed way. More specifically, with TARGET2 granular data we can track bank i's customers transactions in four ways.

The first is the bank's multilateral position settled by domestic and international RPSs. Given its time lagged nature we call it *deferred net settlement transfer* channel ($DNST$). We consider the national RPSs and STEP2, an international RPS owned and managed by the EBA. STEP2 is a Pan-European ACH processing payments in euro. The platform is one of the key clearing and settlement mechanisms in the Single Euro Payments Area (SEPA).[36] Together the payments settled through these systems represent the vast majority of interbank retail transfers, including transactions to merchants, deposit transfers, card payments, and so on.

The second source of information is from the settlement process of instant payments. For euro payments, recently the European Banking Association (EBA) launched a service that uses pre-funding to settle these payments.[37] The Eurosystem recently launched TIPS, a new platform that allows the settlement of instant payments directly in central bank money.[38] Since instant payments settle in real-time, banks have to dedicate part of their reserves or some collateral to pre-fund them.[39] We label changes in these balances as instant settlement transfers ($INST$). These systems can settle basically

[35] For more information about TARGET2 see http://www.ecb.europa.eu/paym/t2/html/index.en.html.

[36] STEP2 has been conceived from the outset as a Pan-European ACH for the single currency and eventually for a Single Euro Payments Area (SEPA), where banks from the different SEPA countries connect directly to exchange payment files and where appropriate routing tables enable reach to all other banks offering SEPA payments.For more information about STEP2 see https://www.ebaclearing.eu/services/step2-t-platform/overview/.

[37] See the RT1 webpage for more information https://www.ebaclearing.eu/services/instant-payments/introduction/.

[38] See the TIPS (TARGET instant payment settlement) webpage for more information https://www.ecb.europa.eu/paym/intro/news/articles_2017/html/201706_article_tips.en.html.

[39] In these sub-accounts payments can be settled in central bank money one by one, like in TIPS, or not, like in EBA RT1. The latter use pre-funding and settle in central bank money only after, while TIPS offers final and irrevocable settlement for instant payments in central bank money on a 24/7/365 basis. It allows participating banks to set aside part of their liquidity on a dedicated account opened with their central bank, from which instant payments could be settled around the clock. The balance on these accounts counts towards their required minimum reserve. These infrastructures process instant SEPA credit transfers and operate around the clock on any day of the year and support payment service providers in transferring euro transactions between payment accounts in less than 10 seconds end to end, with immediate availability of the payment amount to the beneficiary.

the same type of transactions of classic RPSs plus mobile and peer-to-peer instant transfers, which can be increasingly used by customers. With a classic RPS it takes up to one business day for a payment in euro to reach the beneficiary. With instant payments, the funds are available immediately (in the order of seconds) for use by the recipient, 24/7/365.

Much before the introduction of instant payments TARGET2 offered real-time settlement in central bank money for customer payments to participating banks, this is the third source of information available. These are gross bilateral interbank money transfers settled directly in the RTGS on behalf of customers. They are more likely to be used for high value transactions and B2B transfers. Nevertheless, banks use these payments intensively also for small transactions, probably because instant payment solutions started to appear just by the end of 2017. We label this type of transfers as *real-time settlement transfers* channel ($RTST$). As bilateral gross interbank transfers, they also provide information on the counterparty -i.e. the bank that receives the funds-. $RTST$ settle within the same day of the payment instruction, usually after few minutes, $INST$ are even faster, as they settle in few seconds.

Finally, as both reserve accounts and banknotes are the only forms of central bank money (so far), cash operations by commercial banks have to be exchanged with funds in TARGET2 ($CASH$). If the payment message is rich enough, we also have data on the denomination of banknotes.

In practice, we can construct bank A's net position on each of these channels during a time interval t by subtracting the outgoing payments to the ingoing payments related to that specific payment category.[40] The first three variables capture deposits of customers flying from a bank to another one, leaving the deposit currency ratio unaltered, while the last captures the conversion from commercial to central bank money by depositors. A nice feature of payment system data is that it allows us to identify which of these cases materialize in real-time, if signals are properly extracted. There is no need to stress the salience of this information from a financial stability perspective. In the next section, we outline the method proposed to immediately extract signals from this data.

D Tempestivity w.r.t. Supervisory Reports

An important feature of payment data is that it can provide more timely signals of shocks to funding than supervisory reports, which have usually a low frequency. If some type of distress manifests in between two reports, it may be overlooked or captured with delay. Figure D.1 gives a simple graphical example. In the blue line, a temporary shock hits d_t in between supervisory reports SR_2 and SR_3. In this case, given that the shock is absorbed in the time interval and $d_{SR_2} \approxeq d_{SR_3}$ the event is not recognized. In the red line, a permanent shock hits d_t in between supervisory reports SR_2 and SR_3. In this case, the shock is recognized with delay. Both issues amplify with the interval $SR_t - SR_{t-1}$. Important social and private costs can be saved if the shock is recognized timely. The banks itself can

[40] Other transformations of payments can be used, but the net position is the best proxy for daily variation of deposits.

reduce the cost of substituting the funds. The resolution authority has additional time to collect an organize resources to manage the shock.

Figure D.1: RTGS vs Supervisory Reports

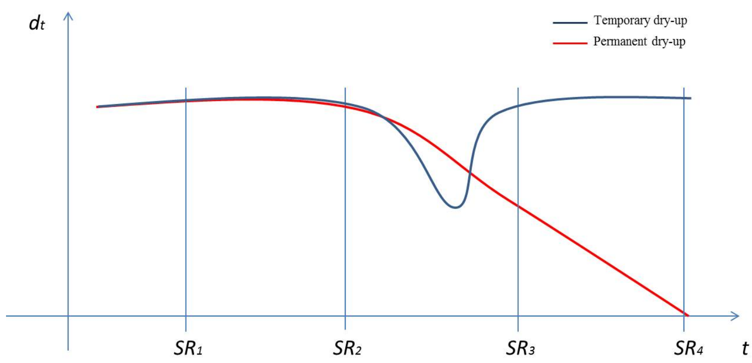

Notes. time (t) on the x-axis, deposits (d_t) on the y-axis. SR_t represents the time in which the supervisory report is delivered. The blue line represents a temporary shock, the red one a permanent shock.

E Alternative Information

Given that there are many viewpoints from which the health status of a bank can be observed, it seems that we need at least to discuss why payment systems are better than other information sources to timely identify true distress.

If we restrict our comparison to central bank-internal information, we see at least two potential candidates for monitoring banks' deposits. If the central bank is also the supervisory authority, it is likely to receive reports from the commercial bank with a certain frequency. In addition, inspections and other forms of control can be implemented. The quality and the frequency of this information depends on many factors. Even if supervisory reports can give a deeper view on the balance sheet of the bank, they have two major drawbacks. First, the frequency is usually low, and thus it can happen that distress occurs (and even ends) in between two observations (see the online appendix). Second, the bank may temporarily misreport some items. As an alternative, the central bank can monitor other aggregates that are under its direct control at a high frequency. The reserve account is an example, given that it is recorded every day for every bank to manage the reserve requirements. Nevertheless, we show in Section 4 how the drain is offset immediately by the banks, making the balance of the reserve account flat and uninformative.

Alternatively, one can look at central bank-external data. A straight source of information is market data. One can follow several indexes computed for a bank for example on Thomson Reuters, Bloomberg and so on. The problem is that market data is by construction informative for marketable debt, it can tell us whether people is selling bank's bonds or other type liabilities, but it is uninformative about deposits. A market-based measure to assess trust and credit risk is the CDS spread, unfortunately such indicator is available only for few big banks. Ratings instead are available for more banks, but updated with low frequency.

Another popular source of information is Google. Google trends has been used in several economic research papers to nowcast economic aggregates (see Choi and Varian, 2012; D'Amuri and Marcucci, 2017, for example). To test this possibility, we took the sample of episodes identified in the empirical application (see Section 3.2) and generated several time series combining the name of the bank with the words 'crisis', 'failure', 'distress' and summed them up. The indicator works pretty well in tracking these bad news episodes. Unfortunately, there is always an increase of these searches when bad news pop up, even without a drop of deposits, so the type I error is quite high. Probably people start to search these words even if they are not that bank's depositors and, even if they are, it is not automatic that they then withdraw their money.[41] It then seems a much more noisy measurement than the punctual outflows detected in payment systems.

F Structural Break Tests

Our analysis is based on the ability of ReNoSCh to identify distress episodes. ReNoSCh works on real-time, when we are at time t it says to us whether a shock is occurring exactly at time t. This is not only a desirable feature, it is partly the scope of this work. In this section, we want assess the capacity of our algorithm to detect distress episodes formally.

A straight way to assess the quality of ReNoSCh's output is to backtest it. We can check whether structural break tests, which use a larger information set (having post-shock observations available) identify the beginning of distress at the same date. Clearly these methods are not substitutes of ReNoSCh, because they can not identify breaks in real-time, but they can tell whether they would ex-post spot a break at the date identified by ReNoSCh in real-time.

A classic methodology to test for structural breaks is the Chow test. As ReNoSCh works directly on residuals of a regularization step and we are not interested on any particular structural relationship of customer payments with other variables used in the regularization process, we can just test whether the average daily net position on customer payments changes its expected value exactly when ReNoSCh says. We took all the distress episodes identified by ReNoSCh and average the net position of the bank separately for cash, real-time, deferred settlement payments and the sum of them across the episodes from 100 days before the first distress day and the following 20 days.[42] We run the test giving the identified first day of distress as a candidate for the breakpoint. The test statistics are reported in the upper panel of Table F.1, the p-values are always very small. There is significant evidence to reject the null hypothesis that the coefficients are stable after the ReNoSCh-break points occurred for all the

[41] Social networks, like twitter, also provide signals about depositors' sentiment (see Accornero and Moscatelli, 2018), but they may suffer from the same problem.

[42] Results are robust to a shortening of the pre-shock period.

time series considered.

The Chow test takes the breakpoint as given, thus the method is good for backtesting our model but is not fully comparable to it, as it does not estimate the break date. To check its sensitivity we can change the breakpoint around the one estimated by ReNoSCh and see how well the test performs. We took the 10 days before and the 10 days after the estimated day. In Figure A.4 in the online appendix we report the Chow test statistic in blue and its critical value in violet. We can see that the statistic reaches its maximum around the estimated break point, tracked by a black vertical line. Nevertheless, the blue line is above the violet in a small interval around the estimated break point, between 5 and 9 days wide. To get a more reliable assessment of the quality of our algorithm, we can consider a structural change test that provides also an estimate for the breakpoint. The test proposed by Bai and Perron (1998) is suited for this task and very popular among scholars doing this type of econometric exercises.[43] In their test there is no input regarding the breakpoint, and the method is free to estimate the optimal one. A useful feature of their estimator is also that it constructs confidence intervals for the break dates.

In the lower panel of Table F.1, we report the expected break date estimated with the algorithm of Bai and Perron (2003) and the 95 percent confidence intervals.[44] All the days are expressed as their relative distance to the first day of distress estimated by ReNoSCh. The estimated break date for the aggregate time series of customer payments, in the first column, is the day after. For the deferred electronic transfers the estimated day is two days after, for the real-time transfers it is two days before, while for the cash withdrawals it is exactly the same day. The confidence intervals vary between 9 and 21 days and are narrower for DNST and the aggregate time series. Figure A.5 in the online appendix depicts the time series of the cumulated daily net position on customer payments, the ReNoSCh break date (the solid vertical line), the estimated endogenous break point (the dotted vertical line) and its confidence interval (the red segment). ReNoSCh break dates are always included in the confidence interval, and very close to the estimated expected date. As a whole it seems that the ReNoSCh provides a good outcome when compared with structural break methods. For the aggregate time series of customer payments, the beginning of distress is estimated even one day before. On average, it slightly anticipates the break date, which in our case is better than postpone it.

G Cash Withdrawals

Let us now focus on cash. Depositors can withdraw banknotes in two ways, via ATM and directly at the bank teller.

When we think about intense distress periods, the image of long lines at the bank tellers is the first

[43] See also Bai and Perron (2003) for a description af the relative algorithm and applications.

[44] We use the R function 'breakpoints' in the package 'strucchange' (Zeileis et al., 2001). We set the break points to be up to one.

Table F.1: Structural Break Tests - Exogenous and Endogenous Break Points

		All	DNST	RTST	Cash
Exogenous break point					
	Chow test statistic	22.559	14.780	8.345	10.982
	Critical value	3.076	3.076	3.076	3.076
	p-value	0.000	0.000	0.000	0.000
Endogenous break point					
	2.5% bound	-4 days	-4 days	-12 days	-4 days
	Expected	+1 day	+2 days	-2 days	0 days
	97.5% bound	+5 days	+6 days	+9 days	10 days

Notes. Dependent variable: daily change in the net position for each channel averaged across all the cases identified. The first column reports all the customer payments, the second reports the deferred net settlement transfers, the third reports the real-time settlement payments, the last reports the cash withdrawals. We considered 120 days, 100 before the distress starts. We use a standard Chow test for the exogenous break point. The test proposed by Bai and Perron (1998) and the algorithm in Bai and Perron (2003) are used to endogenously estimate the break date. 'Expected' reports the relative position of the estimated break point w.r.t. the estimated first day of distress. The bounds of the 95% confidence interval are also reported in terms of distance from the first estimated day of distress.

to come into our minds. That is easily recognizable by everybody and translates in the ultimate form of panic. We showed that in terms of magnitude the identified distress episodes are mostly digital and thus invisible. Nevertheless, cash can still have an important role. So far, if a bank's client does not want to move her deposit to another bank, the only way she has to get central bank money instead of commercial bank money is withdrawing banknotes.

The advantage of the ATM is that it is faster to withdraw cash from almost everywhere and there is no face-to-face interaction. The latter is a desirable feature for depositors moving away their deposits. In some European jurisdictions, if a depositor withdraw a large amount of money in large denominations suddenly at the bank teller, the bank can ask a series of questions about the nature of the operation. This process is due to the prevention of money-laundering and terrorism financing which obliges banks to report 'suspect operations' to the financial intelligence unit (FIU). In this case the depositor has to declare that she is moving away her deposits or lie and risk to be filed by the bank.

Nevertheless, depositors can not withdraw large amounts of money at the ATM, there are pretty tight limits on daily operations, and only denominations up to 50 euro are often available in most of the euro countries, which are not the best to store big amounts of cash (especially because the ATM denomination is not known ex ante by the depositor).

Depending on the depositor's preferences, she may get a big amount of cash in large denominations at the bank teller, facing the risk of being filed, or repeatedly get smaller amounts at the ATM, taking the risk of not getting the last tranches because the bank is failed in the meanwhile. The smaller the deposits the fewer ATM operations needed.

Assuming a constant inventory of banknotes, if the depositors start to go to the bank teller, the bank will need large denomination banknotes. On the contrary, if the demand of cash is via ATMs it will need mostly small denominations ($\leq$ 50 euro). Given that we have granular data on which denomination is taken by each commercial bank from the central bank on a daily basis, this information is used to get more insights about depositors' preferences. We pooled all the episodes and regressed the daily net position of banks on each denomination on a dummy taking value one when the shock kicks in and zero before, the control period is 100 days before the first day estimated by ReNoSCh. Our regression model takes the following form,

$$c_{i,t} = \delta r_{i,t} + \alpha_i + \mu_t + v_{i,t}, \tag{10}$$

where $c_{i,t}$ are the daily (t) net withdrawals (withdrawals minus deposits) for the banknote denomination at day t by bank i, α_i is a bank fixed effect for denomination, μ_t are day fixed effects, $v_{i,t}$ is the error term. Table G.1 reports our results. During distress days about 0.01 percent of deposits were withdrawn (or not deposited) in ATM denominations and almost a tenth of it in 100 euro banknotes. The largest denominations, 200 and 500 euro, show an increase but smaller in value and not significant. From

Table G.1: Cash Withdrawals - Demand for Different Denominations

Dependent: daily bank-specific withdrawals in		
ATM ($\leq$ 50)	Distress	0.0064 ***
		(0.0010)
	Bank FE	Yes
	Day FE	Yes
	DW	2.0350
	Observations	9,804
100	Distress	0.0004 ***
		(0.0002)
	Bank FE	Yes
	Day FE	Yes
	DW	1.9468
	Observations	9,804
200	Distress	0.00000
		(0.00001)
	Bank FE	Yes
	Day FE	Yes
	DW	2.2670
	Observations	9,804
500	Distress	0.0003
		(0.0003)
	Bank FE	Yes
	Day FE	Yes
	DW	2.0153
	Observations	9,804

Notes. * : $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$. OLS estimates of δ. Heteroscedasticity and autocorrelation consistent estimates for all the standard errors. The dependent is the daily net position in euro for each banknote denomination and for each bank. The coefficients on different lines refer to separate regressions and are reported in percentage point of deposits of the shocked bank. R^2 and other coefficients of each regression are not reported for brevity. The pre-shock and the post-shock period are 20 settlement days before and after the shock. We pooled together all the identified episodes.

our evidences we can see that even if deposits are moved mainly digitally, cash still plays a role. Surprisingly, smaller denominations are the most affected, signaling a higher incidence of preference towards withdrawing small amounts repeatedly instead of big amounts at the bank teller.[45] This is probably more rationale for small depositors, who may also not be aware of the deposit insurance (Bartiloro, 2011) or prefer to hold a small amount of cash instead of a credit issued from the deposit insurance. Big depositors may not want to store huge amount of money in banknotes, they probably have accounts with other financial institutions and prefer to use electronic transfers. This finding has

[45] It is not possible to exclude that the bank also imposes the denomination to the customer, trying to disincentivize big withdrawals.

significant implications for financial stability, since even some insured funds are likely to flee banks in response to stress, and can serve to inform banking theory models (Dávila and Goldstein, 2020).