

Performance and Resilience to liquidity disruptions In interdependent RTGS payment systems

By

Fabien Renault*, Walter E. Beyeler♦, Kimmo Soramäki*,
Robert J. Glass♦ and Morten L. Bech*♦

2 June 2008 preliminary version (comments welcome)

Presented at the Joint Banque de France / European Central Bank conference on “Liquidity in interdependent transfer systems”, Paris, 9-10 June 2008

Abstract:

This paper analyses liquidity and credit risks in the context of interdependent interbank payment systems. A simple model is developed to investigate the operation of two real time gross settlement systems interlinked through FX transactions conducted by a set of global banks that participate in both systems. In addition, further interdependence is created by imposing a Payment versus Payment (PvP) constraint. The model illustrates under which conditions settlement of payments in the two systems becomes correlated and how credit exposures are created as the result of liquidity pressures in one of the two systems. PvP can eliminate this credit risk but will make the settlement process of each system dependent on the level of liquidity available in the other system. The consequences of an operational disruption affecting a bank participating in only one system are investigated. The crisis is shown to spread to the other system with different consequences depending on the level of liquidity available and the settlement mechanism used for FX transactions. Different channels of cross-currency channels of crisis propagation are identified.

* Banque de France, fabien.renault@banque-france.fr.

♦ Sandia National Laboratories, webeyel@sandia.gov and rjglass@sandia.gov. Messrs. Beyeler and Glass acknowledge the financial support of the National Infrastructure Simulation and Analysis Center (NISAC), a program under the Department of Homeland Security's (DHS) Preparedness Directorate. Sandia National Laboratories (SNL) and Los Alamos National Laboratory (LANL) are the prime contractors for NISAC under the programmatic direction of DHS's Infrastructure Protection/Risk Management Division. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

♣ Helsinki University of Technology, www.soramaki.net.

• Federal Reserve Bank of New York, morten.bech@ny.frb.org (corresponding author).

♦ The views expressed in this paper do not necessarily reflect those of the Banque de France, Sandia National Laboratories, Federal Reserve Bank of New York or the Federal Reserve System. The authors would like to thank Jordan Parks for excellent research assistance. We thank Denis Beau, Larry Radecki, Rob Close, participants at the Joint Bank of England/European Central Bank Conference on 'Payments and monetary and financial stability' and the annual Bank of Finland Payment System Simulation conference.

| | |
|--|-----------|
| 1. INTRODUCTION | 3 |
| 2. MODEL DESCRIPTION | 5 |
| 2.1. OVERVIEW | 5 |
| 2.2. BANKS, PAYMENTS AND TOPOLOGY..... | 6 |
| 2.3. GLOBAL BANKS AND FX TRADES | 6 |
| 2.4. CALCULATION OF THE FX EXPOSURES | 7 |
| 2.5. SIMULATION CONFIGURATION | 8 |
| 3. PERFORMANCE OF THE SYSTEM IN NORMAL OPERATION | 9 |
| 3.1. CORRELATION OF THE SETTLEMENT ACTIVITY OF TWO SYSTEMS IN NORMAL OPERATION | 9 |
| 3.2. FX SETTLEMENT RISK ARISING FROM NON-PVP TRADES IN NORMAL OPERATION..... | 11 |
| 3.2.1. FX exposures with the same level of liquidity in both systems | 11 |
| 3.2.2. FX exposures with different levels of liquidity in the two systems..... | 11 |
| 3.3. QUEUING UNDER NON-PVP AND PVP | 12 |
| 3.3.1. Queuing with the same level of liquidity in both systems | 12 |
| 3.3.2. Queuing with different levels of liquidity in the two systems | 12 |
| 4. RESILIENCE OF THE SYSTEM TO OPERATIONAL DISRUPTIONS..... | 13 |
| 4.1. EFFECT OF AN OPERATIONAL DISRUPTION ON THE SETTLEMENT ACTIVITY OF THE TWO RTGSS..... | 14 |
| 4.2. EFFECT OF AN OPERATIONAL DISRUPTION ON THE FX EXPOSURES IN THE TWO RTGSS..... | 15 |
| 4.3. DISCUSSION OF THE CROSS-CURRENCY CHANNELS OF CRISIS PROPAGATION | 16 |
| 5. CONCLUSION | 17 |
| REFERENCES | 19 |
| ANNEXES | 21 |

1. Introduction

Technological innovations and structural changes in banking are creating greater interdependence among the world's large value payment and settlement systems. Developments in technology have facilitated the emergence of systems that settle across national borders, at off-shore locations or via remote access. Given the importance of payment and settlement systems for the smooth operation of the financial system as well as its resiliency, stakeholders need to understand and assess the potential consequences of such an evolution and to understand how interdependencies influence liquidity, credit and even operational risks. In particular, the increasing scope for liquidity interdependence across systems serves to further accentuate the potential role of payment and settlement systems in the transmission of contagion effects.

The study of large value payment and settlement systems is at the intersection of several disciplines including economics, operations research and statistical mechanics of complex systems. It is a fast growing field, with a burgeoning theoretical and empirical literature on payment economics. Recent examples include McAndrews and Martin (2007) who present a theoretical model of liquidity savings mechanisms and Mills and Nesmith (forthcoming) who analyze banks' strategic decisions on when to submit transactions.

In addition, the recent development of simulation methodologies and tools able to replicate the operation of payment and settlement systems using real data has facilitated detailed stress-testing and systemic risk studies (see Mazars and Woelfel (2005), Bedford et al (2005), Schmitz et al (2006), and McVanel (2005) among others). Moreover, interesting insights are being gained from a growing literature viewing payment and settlements systems as (complex) networks where participants are represented as nodes and weighted and directed links represent either flows or bilateral risk controls.¹

Nevertheless, the study of payment and settlement systems still has several unexplored areas. A key challenge for the simulations approach is to incorporate behavior. As noted by Bech and Garratt (2006) the actions of participants have both the potential to mitigate, but also to augment the adverse effects of (wide-scale) disruptions. In fact, most simulations studies produced so far are subject to their own version of the Lucas critique.² Encouraging, preliminary progress in terms of modeling behavior has been made by Galbiati and Soramäki (forthcoming) in the context of an interbank payment system by using an agent based modeling approach. An overview of the literature up to date is available in Ledrut (2006).

A deficiency in the current literature has been the tendency to look at systems in isolation. Most studies have been dedicated to the study of single systems which precludes any rigorous

¹ The topology of payment flows between participants varies substantially across different large value payment systems. Generally, payment systems with fewer participants tend to form complete networks (see Lubloy-Szenes (2006) and Becher et al (forthcoming) whereas larger systems are often characterized by a limited core of large participants that exchange payments with many counterparties, and a large set of smaller participants that exchange a payments with only a few counterparties (see Inaoka et al. (2004), Soramäki et al. (2007)). In many instances the (degree) distribution of counterparties follows a power law. In modeling payment systems, the topology of interactions between banks can be important in determining the response dynamics to a perturbation (Bech and Garratt (2006)).

² It is naive to try to predict the effect of policy changes in a payment and securities settlement system based purely on relationships between participants and timing of instructions obtained from historical data as these both the timing and relationship would necessarily change whenever policy – the rules of the game – is changed. In order to overcome the critique one must identify the “deep parameters” such as preferences, technology and resource constraints that govern individual participant’s behavior.

treatment of system interdependencies.³ System interdependencies are, however, high on the agenda of policy makers. In 2001, the Group of Ten “Report on Consolidation in the Financial Sector” (the Ferguson report) reported that “the emergence of multinational institutions and specialized service providers with involvement in several payment and securities settlement systems in different countries, as well as the increasing liquidity interdependence of different systems, further serve to accentuate the potential role of payment and settlement systems in the transmission of contagion effects.” (Bank for International Settlements 2001)

A first step in understanding system interdependencies is being taken by a forthcoming report of the Committee Payment and Settlement Systems under the auspices of the Bank of International Settlements. It provides a qualitative framework for analyzing system interdependencies. Among other things the report identifies three different types of interdependencies. The first type is *system-based* interdependency, which includes payment versus payment (PvP) or delivery versus payment arrangements (DvP)⁴ as well as liquidity bridges between systems. The second type is *institution-based* interdependency which arises when e.g. a single institution participates or provides settlement services to several systems. The final type is *environmental-based* interdependency which can emerge if multiple systems depend on a common service provider, for example the messaging service provider SWIFT.

In this context, a stylized numerical model of interdependent interbank payment systems was developed with the objective to add to the understanding of system interdependencies. The proposed model for system interdependencies is obtained by extending the congestion and cascades model of Beyeler et al. (2007) from single real time gross settlement (RTGS) system to multiple systems operating in distinct currencies. First, we let a subset of participants settle payments in both of the systems by participating in the two RTGS. These global banks can also make FX transactions with each other, exchanging one currency for the other. This creates an institution-based interdependency between the two systems. The FX transactions can be settled via a payment-versus-payment (PvP) mechanism, which represents a system-based interdependency. In a previous paper, (Renault, Beyeler, Glass, Soramäki and Bech, 2007), we investigated how the functioning of one system depends on the other. In particular, we looked at how differences in the level of liquidity, the liquidity management practices, and the settlement methodology in the two systems could affect their respective performance and risk profiles. Our main results are recalled on chapter 3.

In this paper we also study the effect of an operational outage affecting a local bank participating in only one of the two systems. At a given time, the local bank becomes unable to transmit its payments to the system, while it still can receive payments from its counterparties, thus turning into a liquidity sink for the system. We monitor the degradation in settlement rate in both systems as a function of the total liquidity available and of the settlement mechanism used to settle FX transactions. We find that the FX linkage creates unexpected responses in both systems, during the recovery period as well as during the disruption itself.

The paper is structured as follows. Section 2 presents the model and section 3 summarizes the behavior of the unperturbed system. Section 4 describes the consequences of an operational disruptions to a large bank, which processes only local payments, on the behavior of both systems. Section 5 concludes and summarizes the paper.

³ An exception to this rule is Hellqvist and Snellman (2007) who study the interaction between the interbank payment and securities settlement systems in Finland.

⁴ These technologies ensure that each leg of a transaction is settled simultaneously and hence have the ability to eliminate counterparty credit risk.

2. Model Description

2.1. Overview

We consider a simple economy with two countries using distinct currencies. For the ease of exposition, the currencies are referred to dollar and euro, respectively. The countries are populated with *economic agents, banks, and a central bank*. Each central bank is a benevolent provider of an interbank payment system. The agents hold deposits at banks to settle obligations arising from trades with each other. In turn, banks maintain reserve balances at the central bank to settle payment instructions received from their customers and destined to agents banking at other banks. A graphical representation of the model is provided in chart 1 (annex).

In the model there are two types of banks: local and global. Local banks only participate to their domestic payment system whereas the global banks are direct participants of both systems and engage in foreign exchange trading with other global banks. Consequently, each system processes both non FX related payment instructions on behalf of local and global banks and FX related payment instructions on behalf of the global banks.

Payments are settled individually and irrevocably in continuous time. This mode of settlement is commonly referred to as real time gross settlement. To simplify matters, the systems are assumed to operate 24 hours a day and seven days a week. Consequently, any issues related to end-of-day management of settlement positions or issues related to overnight lending are ignored.

In the proposed model, the two RTGS systems are linked through two sources of interdependencies. A first interdependency is due to the presence of the “global banks”, which are direct participants in both systems and make foreign exchange trades with each other. The two legs of these trades are settled as payments in the respective payment systems. This form of interdependency can be described as an institution-based interdependency, following the taxonomy developed by the CPSS Working Group on System Interdependencies.

Second, the two systems can be linked through a payment versus payment (PvP) mechanism that ensures the simultaneous settlement of both legs of FX transactions. In the model, the PvP mechanism can be turned on (PvP) or off (non-PvP), in which case the two legs of the FX trades are settled independently. The PvP settlement represents a system-based interdependency between the two RTGS systems. Figure 1 provides a graphical illustration of the model and highlights the two sources of interdependencies between the two RTGS systems.

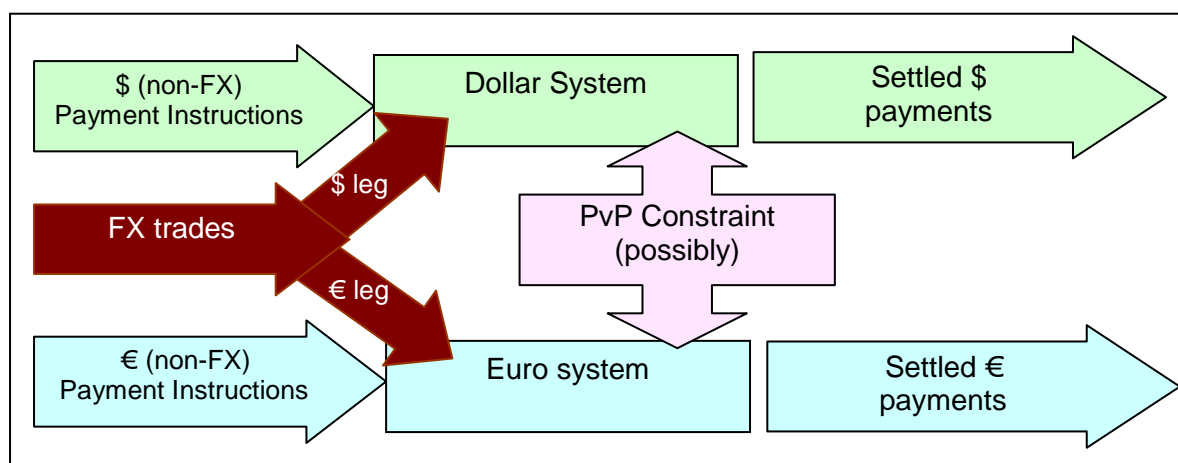


Figure 1: Interdependencies - FX trading and PvP constraint

2.2. Banks, payments and topology

A bank's ability to settle payment instructions depends on the availability of funds in its account at the central bank. We assume that banks choose to issue a payment order in the RTGS as soon as they receive a payment instruction from one of their customers. If a bank does not have the necessary liquidity to settle a payment, then the payment instructions are placed in queue by the payment system. Whenever new funds are received, they are immediately used to settle queued instructions.

It follows that a bank $i = \{1, 2, \dots, n\}$ that participates in say the dollar system can be characterized by its level of customer deposits in dollars, $D_i^{\$}(t)$, its payment instruction arrival rate, $\lambda_i^{\$}(t)$, its queue of payments awaiting settlement, $Q_i^{\$}(t)$, and its balance of reserves at the central bank, $B_i^{\$}(t)$. In order to fully characterize a global bank we also need to specify deposits, arrival process, queue and reserves in euros well as the FX trading among global banks. This is explained in section 2.3.

Payments come in many sizes but for simplicity we assume that all payments are of equal size and normalized to one. The arrival of payment instructions to a bank, is modeled as a (non homogenous) Poisson process with time varying intensity, $\lambda_i^{\$}(t)$. We assume that the arrival of payment instructions to a bank is driven by the level of deposits $D_i^{\$}(t)$ held by its customers in such a way that a higher level of deposits makes it more likely that the bank will receive requests for outgoing payments. Specifically, we assume that the level of deposits is converted into a payment instruction with a constant probability per unit time, $p^{\$}$. Consequently, the expected rate of instruction arrival to bank i per unit of time (the intensity), is given by:

$$\lambda_i^{\$}(t) = p^{\$} D_i^{\$}(t) \quad (1)$$

Accordingly, instruction arrival rate increases as incoming payments add to deposits and decreases as outgoing payments deplete deposits. To be clear, the equation above only describes the expected arrival rate. The actual number of payment instructions arriving to bank i will be the random outcome of the Poisson process.

We assign a recipient bank j to any incoming payment instruction to bank i by drawing from a discrete probability distribution defined over the set of banks. Let $\omega_{ij} \in [0;1]$ denote the probability that bank i will send a payment to bank j to the condition that bank i has received a payment instruction. We impose little structure on the matrix of conditional probabilities $\Omega \equiv \{\omega_{ij}\}_{n \times n}$ except obviously that for every bank i , $\sum_j \omega_{ij} = 1$. In fact, we allow for “on us” or book transfers ($\omega_{ii} > 0$) and we do not impose any symmetry in bilateral relationships between banks ($\omega_{ij} \neq \omega_{ji}$). As a result, we can, in principle, generate any topology of payment flows desired for the system as a whole. For example, for the payment flows to form a complete network over time a strictly positive probability mass has to be assigned to all elements of Ω .

2.3. Global banks and FX trades

In addition to receiving payment instructions to be settled in the respective payment systems, the global banks settle FX trades between each other. The local banks do not settle sides

arising from FX trades. We do not model the FX market and we simply assume throughout the paper that the exchange rate between euro and dollar is constantly maintained at a level of one to one.

We assume that the arrival of FX trades is driven by the level of deposits of the two currencies at the respective banks. The average number of euro for dollar trades that bank j receives from its clients in a given unit of time is assumed to be proportional to D_j^ϵ . Similarly the average number of dollar for euro trades that bank i receives from its clients in a given unit of time is taken proportional to $D_i^\$$. We can thus assume that the probability of one of bank i 's clients engaging in a dollar for euro trade with one of bank j 's clients is proportional to the product $D_i^\$ D_j^\epsilon$.⁵

For every pair (i,j) of global banks, the average rate at which a bank i has a dollar for euro transaction with bank j is given by:

$$\lambda_{ij}^{\$ \epsilon}(t) = p^{FX} \sqrt{\frac{D_i^\epsilon(0)}{D_i^\$(0)}} \sqrt{\frac{D_j^\$(0)}{D_j^\epsilon(0)}} D_i^\$(t) D_j^\epsilon(t) \quad (3)$$

where p^{FX} is a constant parameter describing the level of FX trading activity between the two RTGS systems. The use of the $\sqrt{\frac{D_i^\epsilon(0)}{D_i^\$(0)}} \sqrt{\frac{D_j^\$(0)}{D_j^\epsilon(0)}}$ proportionality coefficient guarantees that

$\lambda_{ij}^{\$ \epsilon}(0) = \lambda_{ji}^{\epsilon \$}(0)$ as well as a finite return time towards the initial steady state. The retained proportionality coefficient simply translates the fact that we expect certain stability regarding the currency holdings of the banks during a simulation. As in reality, we do not expect the largest participant in the euro system to be selling off all its euro holdings in order to become the largest participant in the dollar system. The FX trading activities of the global players will thus only let them oscillate around their starting position.

2.4. Calculation of the FX exposures

When FX trades are settled non-PvP, the bank that pays the first leg of the transaction bears a FX credit risk until the other leg of the transaction is settled in the other payment system. We define the time-averaged cumulated exposure of the \$ selling banks towards the € selling banks (ie the amount of € owed) and the time-averaged cumulated exposure of the € selling banks towards the \$ selling banks (ie the amount of \$ owed) as respectively:

$$Exposure_{\epsilon owed} = \sum_k \text{Value}_k \cdot \max(0; t_k^\epsilon - t_k^\$) \cdot \frac{1}{T} \quad (6)$$

and

$$Exposure_{\$ owed} = \sum_k \text{Value}_k \cdot \max(0; t_k^\$ - t_k^\epsilon) \cdot \frac{1}{T} \quad (7)$$

respectively. Here k indexes the FX transactions settled, T is the time duration of the calculation period, t_k^ϵ is the settlement time of the euro leg of the k^{th} transaction, and $t_k^\$$ is the settlement time of the dollar leg of the k^{th} transaction.

Figure 2 visualizes the concept of time-averaged exposures which correspond to the area of the colored rectangles in figure 2. The equations above simply reflect the fact that, in a FX

⁵ We assume that the customers do not have any preference regarding their FX trade counterparty,

transaction, the dollar selling bank will be facing an exposure towards the euro selling bank, if the euro leg of the transaction settles after the dollar leg of the transaction (i.e., if $t_k^\epsilon > t_k^\$$).

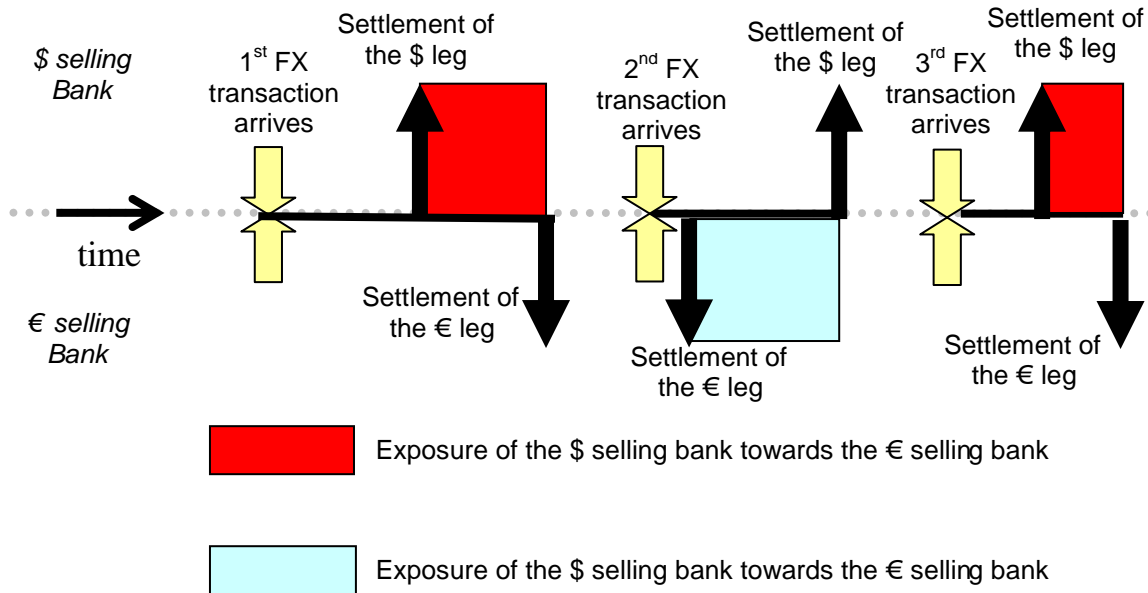


Figure 2: Exposures created by the non-PvP settlement of FX transactions

2.5. Simulation Configuration

For purposes of this exercise, we configure the model in the following manner. We assume that there are 100 participants in each payment system. The “global banks”, which participate in both systems, are taken to be the three largest participants in RTGS € together with the three largest participants in RTGS \$. Consequently, the number of local banks is 94 in each system.⁶

Furthermore, we assume that the number of counterparts each bank i will send payment to, $k_i^\$ = \sum_j 1(\omega_{ij} > 0)$, follows a power law distribution. That is bank i sends payments to $k_i^\$$ counterparties with probability

$$P(k_i^\$ = x) \sim \frac{1}{x^\gamma} \quad (4)$$

where γ is the power law coefficient be fitted to produce an average network degree of twelve. Given the set of counterparts for a particular bank i , the conditional probabilities ω_{ij} of a payment instruction being direct to a specific counterpart bank j were drawn randomly using an exponential distribution. In contrast we choose to describe the FX market as a complete network, i.e., a system where each participant trades with every other participant.

⁶ This is likely a conservative estimate for the US. Foreign banking organizations by themselves account for almost 10% of the value transferred over Fedwire. In addition, the value transferred by internationally active US banks is sizeable.

Each payment system is assumed to be in equilibrium initially in that the intensity of outgoing payments is equal to the intensity of incoming payments at time zero for each bank. We have:

$$\lambda_i^{\$}(0) = \sum_j \omega_{ji} \lambda_j^{\$}(0) \quad (5)$$

By equation (1) it follows that the initial equilibrium condition can be written as the following system of equations.

$$p^{\$}(\mathbf{I} - \mathbf{\Omega}')\mathbf{D}(0) = \mathbf{0} \quad (6)$$

where \mathbf{I} is the identity matrix, $\mathbf{D}(0)$ is the vector of initial deposits, $\mathbf{\Omega}$ is the matrix composed of the ω_{ij} , $\mathbf{\Omega}'$ is its transposed matrix, and $\mathbf{0}$ is a vector of zeros. The system of equations can be solved for the equilibrating vector of initial deposits, $\mathbf{D}(0)$, given the aggregate amount of initial deposit \bar{D} available. In other words, by allocating the initial deposits appropriately across banks, we can ensure the initial equilibrium condition of equation (5).

We follow (Beyeler et al 2007) on the initial allocation of the bank's reserve balances. Each participant sets its initial central bank balance $B_i^{\$}(0)$ (respectively $B_i^{\epsilon}(0)$) in order to control its liquidity risk (the risk of being unable to process the orders of its customers due to an insufficient balance) at the lowest possible cost (as maintaining large balances at the Central Bank entails an opportunity cost for the banks). The initial reserve balances of the banks are proportional to the square root of their initial level of deposits:

$$B_i^{\$}(0) = l^{\$} \sqrt{D_i^{\$}(0)} \quad \text{and} \quad B_i^{\epsilon}(0) = l^{\epsilon} \sqrt{D_i^{\epsilon}(0)} \quad (7)$$

where $l^{\$}$ and l^{ϵ} are parameters that characterize the total level of liquidity respectively in RTGS \$ and in RTGS €⁷

A graphical representation of the modeled system is provided in chart 2.

3. Performance of the system in normal operation

3.1. Correlation of the settlement activity of two systems in normal operation

In this section we investigate whether the settlement activity of the two payment systems becomes correlated in normal operation, because of the two interdependencies introduced in the model (the PvP mechanism and the dual participation of the global players). We consider that the settlement activity of the two systems is (positively) correlated when a period of high settlement activity (respectively a period of low settlement activity) within one system corresponds to a period of high settlement activity in the other system (respectively a period of low settlement activity).

⁷ The importance of the initial allocation of bank balances was assessed in a sensitivity study, in which different models of initial allocation were tried. It appeared that the initial allocation of liquidity between the banks does not change qualitatively the results obtained. It was also shown that for a total amount of liquidity within a RTGS, the "square root allocation" used in this paper, led to a significantly lower level of queuing than a "proportional allocation", for high levels of liquidity. This result can be intuitively related to the random walk nature of the evolution of a bank's balance (Beyeler et al 2007).

We can observe the degree of correlation between the two systems visually by using settlement rate scatter plots such as the ones presented in Chart 3 and Chart 4. Two simulations were performed to make each of these two figures. One simulation was run with a low level of liquidity (blue dots), and one simulation was run with a high level of liquidity (red dots). Each dot of the scatter plot corresponds to a certain time window of the simulation (the duration of the simulation was divided into one thousand time windows of constant duration). The abscissa of the dot corresponds to the settlement rate observed in dollar system during the considered time window (i.e., the number of local payments and FX legs settled in dollar system divided by the duration of the time window). The ordinate of the dot corresponds to the settlement rate observed in euro system during the same time window.

In both Chart 3 (non-PvP settlement of FX trades) and Chart 4 (PvP settlement), we can observe that the amplitude of the variations of the settlement rates is much higher at low liquidity. Indeed, at high liquidity, the payments are settled nearly immediately. As a consequence, the queues are almost empty and the settlement rate remains very close to the arrival rate of the payment orders. At low liquidity however, the size of the queues varies greatly over time. Periods of congestion, characterized by a low settlement rate and the building up of the queues, alternate with periods of cascades, characterized by a high settlement rate and a massive release of queued payments.

With regard to the observed degree of correlation of the two systems, Table 1 summarizes the main findings of Chart 3 and Chart 4.

| Degree of correlation between the settlement rates of the two systems | | Settlement mechanism for FX transactions | |
|---|------|--|------|
| | | non-PvP | PvP |
| Level of liquidity (the same in both systems) | Low | -0.02 | 0.83 |
| | High | 0.22 | 0.22 |

Table 1: Correlation coefficient between settlement rates of the two systems

At high liquidity, there is a slight degree of correlation between the two systems, caused by the common drive of FX trading. This was expected since a period of high FX trading will tend to increase simultaneously the throughput in both systems. The settlement mechanism (PvP or non-PvP) does not have any impact on the results at high liquidity, as all payments settle nearly immediately, irrespective of the settlement mechanism in place. The degree of correlation between the outputs of the two systems is 0.22, both in the PvP case and in the non-PvP case. This value tends to increase when the level of FX activity (the relative share of FX trades compared to the total amount of payments processed) increases. The top sketch of Chart 5 illustrates the coupling induced by the FX trading activity at high liquidity.

At low liquidity, the two systems are no longer governed by the arrival of payment orders, but rather by their internal physics of congestion (the payment orders are queued due to a lack of liquidity) and cascades (as the settlement of a newly arrived payment order can trigger the release of several queued payments). The two systems appear completely uncorrelated in the non-PvP case, as the correlation caused by the common FX input has disappeared in the internal process of congestion and cascades. The scatter plot shown in Chart 3 has thus a nearly perfect circular shape. The middle sketch of chart 5 illustrates the decoupling of the two systems.

At low liquidity in the PvP case, the settlement rate of the two systems appear highly correlated, as shown by the “comet shape” of the scatter plot presented in Chart 4. The correlation caused by the common FX input in the high liquidity case has been replaced by a mechanical PvP release correlation between the two systems. The degree of correlation of the settlement rates of the two systems is very high at 0.83. The bottom sketch of chart 5 illustrates how the PvP mechanism creates a coupling between the two systems at low liquidity.

Another way to observe the correlation between the two systems is to look at the time series of the settlement rates. Charts 6 and 7 present the variation of the settlement rate in the two systems over time, respectively in the non-PvP and PvP case, for a low level of liquidity. The two time series appear clearly uncorrelated in the non-PvP case and correlated in the PvP case.

3.2. FX settlement risk arising from non-PvP trades in normal operation

3.2.1. FX exposures with the same level of liquidity in both systems

We first investigate the magnitude of exposures when both systems are operating on the same level of liquidity. The main results are summarized in Table 2. It is not surprising to observe that the credit exposures increase sharply when liquidity is decreased. At high levels of liquidity, both legs of the FX transactions settle nearly instantly and thus the related credit exposures remain very limited. At lower levels of liquidity payments are queued and the settlement of the two legs becomes more asynchronous, thereby increasing the level of exposures. The differences in exposures between the two systems are caused by random variation.

| | | Euros owed | Dollars owed | Total exposures |
|---|---------|------------|--------------|-----------------|
| Level of liquidity (the same in both systems) | Lowest | 734 | 676 | 1410 |
| | Low | 376 | 381 | 757 |
| | High | 221 | 231 | 452 |
| | Highest | 15.3 | 13.7 | 29 |

Table 2: Exposures with same level of liquidity in both systems (non-PvP case, normal priority for FX payments, high level of FX activity).

3.2.2. FX exposures with different levels of liquidity in the two systems

It is well known that time zone differences between payment systems result in systematic exposures for non-PvP FX trades. In a similar way, when one system (for example the euro system) has a significantly higher level of liquidity than the dollar system, the euro leg of the FX transactions will settle significantly faster than the dollar leg. As a consequence, the banks that are selling euro for dollar can expect to face a higher FX settlement risk.

Table 3 sums up the results obtained when the euro system is more liquid than the dollar system.

| Level of liquidity in euro system: Highest | | Euros owed | Dollars owed | Total exposures |
|--|--------|------------|--------------|-----------------|
| Level of liquidity in dollar system | Lowest | 0.0377 | 3150 | 3150 |
| | Low | 0.413 | 1400 | 1400 |
| | High | 8.53 | 365 | 374 |

Table 3: Exposures with highest level of liquidity in the euro system (non-PvP case, normal priority for FX payments, high level of FX activity).

A comparison of Table 2 with Table 3 shows us that when the liquidity in dollar system is maintained constant at the lowest level, and the liquidity in euro system is increased from the lowest level to a very high level, the total exposures increase substantially (from 1410 to 3150).

3.3. Queuing under non-PvP and PvP

3.3.1. Queuing with the same level of liquidity in both systems

We first investigate the case where both systems have the same level of liquidity. Table 4 shows the average number of queued payments in the two RTGS systems, depending on the liquidity available and the settlement mechanism used. as a function of the level of liquidity in the two systems. The first obvious observation is that the level of queuing increases as the liquidity decreases, whether PvP is used or not. It also appears that the use of PvP settlement increases the average level of queuing (and therefore increases the average settlement delay) in both systems by 5 to 10%, depending on the liquidity level. It seems that the larger increase in relative terms occur for intermediate values of liquidity.

| Average queue in dollar system (left) and in euro system (right) | | Settlement mechanism for FX transactions | | | |
|--|---------|--|--------|--------|--------|
| | | non-PvP | | PvP | |
| Level of liquidity (same in both systems) | Lowest | 33 100 | 33 400 | 35 300 | 35 300 |
| | Low | 14 500 | 14 600 | 15 700 | 15 400 |
| | High | 4 510 | 4 480 | 4 890 | 4 900 |
| | Highest | 240 | 241 | 255 | 253 |

Table 4: Average number of queued payments with same level of liquidity in both systems (normal priority for FX payments, high level of FX activity).

3.3.2. Queuing with different levels of liquidity in the two systems

Without a PvP mechanism

This time, we investigate the consequences of a structural liquidity imbalance between the two systems. As a convention, we set the liquidity of dollar system to a lower level than the liquidity of euro system, and we observe how the level of queuing in the two systems evolves as we let the liquidity level within the two systems vary.

Table 5 shows the average queue in both systems, for various liquidity levels. As expected, the average size of the queue increases sharply for a given system when liquidity within this system is decreased. We can also note that the average size of the queue within a system does not depend on the level of liquidity available in the other system when no PvP mechanism is employed.

| Average queue in dollar system (left) and in euro system (right). | | Level of liquidity in dollar system | | | | | | | |
|---|---------|-------------------------------------|--------|--------|--------|-------|-------|---------|-----|
| | | Lowest | | Low | | High | | Highest | |
| Level of liquidity in euro system | Lowest | 33 100 | 33 400 | | | | | | |
| | Low | 33 400 | 14 600 | 14 500 | 14 600 | | | | |
| | High | 32 600 | 4 440 | 14 600 | 4 460 | 4 510 | 4 480 | | |
| | Highest | 32 900 | 235 | 14 800 | 241 | 4 500 | 238 | 240 | 241 |

Table 5: Average number of queued payments in both systems with unequal liquidity (non-PvP case, normal priority of FX payments. high level of FX activity)

With a PvP mechanism

The simulations conducted in the previous section were re-run, this time assuming that the FX transactions are settled using a PvP mechanism. For each level of liquidity in the two systems, Table 6 presents the average number of queued payments in each RTGS. Of course, the average size of the queue in a given system increases sharply when liquidity within this system is decreased, as in the non-PvP case. Contrary to the non-PvP case however, the average size of the queue in one system also depends on the liquidity available in the other system.

| Average queue in dollar system (left) and in euro system (right). | | Level of liquidity in the dollar system | | | | | | | |
|---|---------|---|--------|--------|--------|-------|-------|---------|-----|
| | | Lowest | | Low | | High | | Highest | |
| Level of liquidity in euro system | Lowest | 35 300 | 35 300 | | | | | | |
| | Low | 33 400 | 18 100 | 15 700 | 15 400 | | | | |
| | High | 33 400 | 10 700 | 14 800 | 5 890 | 4 890 | 4 900 | | |
| | Highest | 32 400 | 3 600 | 14 500 | 1 670 | 4 580 | 618 | 255 | 253 |

Table 6: Average number of queued payments in both systems with unequal liquidity (PvP case, normal priority of FX payments, high level of FX activity)

More specifically, when liquidity is decreased within the “less liquid” system, the level of queuing increases significantly within the “more liquid” system. This effect appears especially strong for intermediate levels of liquidity in the “more liquid” system. In addition, we also observe that the level of queuing in the “less liquid” system decreases when the liquidity is increased in the “more liquid” system.

We can therefore conclude that in the PvP case, the average level of queuing in one system depends not only on the level of liquidity available in that system but also on the level of liquidity present in the other system. The two systems therefore appear interlinked as an increase in the level of liquidity in one system either through a change in its participant’s behavior or through a change in the Central Bank policy will have a beneficial impact on the other system.

4. Resilience of the system to operational disruptions

In this chapter we investigate the consequences of the operational default of one participant. Due to a technical problem affecting its IT infrastructure, one participant is unable to send any payment order to the Central Bank operated RTGS system but can still receive payments from its counterparties. Consequently, the affected bank turns into a “liquidity sink”, accumulating a large balance at the Central Bank and depriving the system of its liquidity.

In this paper, we consider the case of an operational default affecting a large “local” € bank, that represents a significant share of the turnover in RTGS € but does not participate in RTGS \$. The level of liquidity is the same in the two RTGS before the operational outage.

Four cases are investigated:

- PvP, high liquidity level in the two RTGS (chart 8)
- PvP, low liquidity level in the two RTGS (chart 9)
- non-PvP, high liquidity level in the two RTGS (chart 10)
- non-PvP, low liquidity level in the two RTGS (chart 11)

Section 4.1 discusses the consequences of the outage on the settlement activity in the two RTGSs while Section 4.2 addresses the impact of the crisis on the FX exposures in the non-PvP case. Section 4.3 proposes an analysis of the cross-currency channels of crisis propagation.

4.1. Effect of an operational disruption on the settlement activity of the two RTGSs

Chart 8, 9, 10 and 11 respectively present the evolution of the settlement rate (ie the number of transactions, whether local payments or legs of FX trades, settled in a given time period) over time in each RTGS, in the four considered cases (PvP High Liquidity, PvP Low Liquidity, non-PvP High Liquidity, non-PvP Low Liquidity). At time point 10 000, a significantly large € local bank faces an operational outage and is unable to send any payment to any of its counterparties in RTGS € until time point 12 000 when the affected bank recovers. The payments that were held within the affected bank's internal IT system are then finally submitted within RTGS €.

The effects of the disruption on the two systems appear complex and manifold. In order to discuss the observed consequences of the outage, we divided the event in three parts:

- the crisis itself, defined as the time period during which the affected bank is unable to send any payments,
- the recovery, defined as the time period immediately after the end of the outage, during which the bulk of payments queued during the crisis is worked off
- the aftermath, defined as the time period comprised between the end of the recovery and the moment when the system has returned to its steady-state equilibrium of before the outage.

Table 7 presents a sum-up of the consequences of the operational disruption on the settlement rate in the two RTGS systems, in each of the considered cases.

| Case | Crisis | Recovery | Aftermath |
|--------------------------------------|---|---|---|
| PvP high liquidity (chart 8) | <ul style="list-style-type: none"> • Settlement rate in RTGS € reaches zero during outage • Settlement rate in RTGS \$ decreases during outage (-50%) | <ul style="list-style-type: none"> • Settlement rate in RTGS € increases at recovery (reaches max) • Settlement rate in RTGS \$ increases at recovery | <ul style="list-style-type: none"> • Settlement rate in RTGS € increases in a long term transient after recovery (+10%) • Settlement rate in RTGS \$ increases in a long term transient after recovery (+10%) |
| PvP low liquidity (chart 9) | <ul style="list-style-type: none"> • Settlement rate in RTGS € reaches zero during outage • Settlement rate in RTGS \$ decreases during outage (-65%) | <ul style="list-style-type: none"> • Settlement rate in RTGS € increases at recovery (reaches max) • Settlement rate in RTGS \$ increases at recovery (reaches max) | <ul style="list-style-type: none"> • marginal effect on the settlement rate of RTGS € • marginal effect on the settlement rate of RTGS \$ |
| non-PvP high liquidity (chart 10) | <ul style="list-style-type: none"> • Settlement rate in RTGS € reaches zero during outage • Settlement rate in RTGS \$ decreases during outage (-17%) | <ul style="list-style-type: none"> • Settlement rate in RTGS € increases at recovery (reaches max) • no overshoot in RTGS \$ | <ul style="list-style-type: none"> • Settlement rate in RTGS € increases in a long term transient after recovery (+10%) • Settlement rate in RTGS \$ slightly increases in a long term transient after recovery (+5%) |
| non-PvP low liquidity (chart 11) | <ul style="list-style-type: none"> • Settlement rate in RTGS € reaches zero during outage • Settlement rate in RTGS \$ decreases during outage (-25%) | <ul style="list-style-type: none"> • Settlement rate in RTGS € increases at recovery (reaches max) • Settlement rate in RTGS \$ increases at recovery | <ul style="list-style-type: none"> • no effect on the settlement rate of RTGS € • no effect on the settlement rate of RTGS \$ |

Table 7: Sum-up of the consequences of the operational disruption of a significant € local bank on the settlement activity in the two RTGS systems, depending on the level of liquidity available (the same in the two systems) and on the settlement mechanism used for FX transactions

With regard to the crisis phase, as expected based on the large size of the bank affected by the disruption, in all considered cases the outage quickly leads to RTGS € being totally deprived of liquidity, and the settlement rate within RTGS € reaching zero. It also appears from this investigation, that in all cases, the outage of the local € bank has also an impact on the settlement rate of RTGS \$. In the non-PvP high liquidity case, the exchange of local \$ payments is however very little affected by the outage and the decrease in the settlement rate of RTGS \$ during the crisis is only due to the stop of FX trading, depriving RTGS \$ from the activity related to the settlement of the \$ leg of the FX trades.

Regarding the recovery phase, we observe in all cases a dramatic increase of the settlement rate of RTGS € when the affected bank recovers the ability to send payments. The settlement activity of RTGS \$ also overshoots in the PvP case (both at high liquidity and at low liquidity), mainly due to the settlement of the \$ leg of the FX transactions that could not settle due to the lack of € balances; and in the non-PvP low liquidity case, due to the settlement of queued local \$ payments.

During the aftermath phase, we detect that the settlement rate of the two RTGS systems remain slightly larger than their steady-state value, when the level of liquidity is high. This effect is due to a slow process of return to equilibrium, from the unbalanced distribution of deposits created by the crisis. We observe no such clear increase in the low liquidity cases, although in the PvP low liquidity case, the small deposit imbalances resulting from the crisis are slowly relaxed, marginally affecting the settlement rate of the system.

A more detailed description of the unfolding of events in the different cases is provided in the Annex, beneath charts 8, 9, 10 and 11. An analysis of the cross-currency channels of crisis propagation is proposed in section 4.3.

4.2. Effect of an operational disruption on the FX exposures in the two RTGSs

Chart 12 and 13 respectively present the evolution of the FX exposures, both in terms of euros owed (corresponding to the FX transactions whose dollar leg has settled and whose euro leg is pending) and dollars owed over time in each RTGS, in the two considered cases (non-PvP High Liquidity, non-PvP Low Liquidity).

Table 8 presents a sum-up of the consequences of the operational disruption on the FX exposures.

| Case | Crisis | Recovery |
|---|--|--|
| non-PvP high liquidity (chart 12) | <ul style="list-style-type: none"> • The amount of euros owed increases dramatically • The amount of dollars owed reaches zero | <ul style="list-style-type: none"> • The amount of euros owed drops and reaches zero • The amount of dollars increases significantly |
| non-PvP low liquidity (chart 13) | <ul style="list-style-type: none"> • The amount of euros owed increases dramatically • The amount of dollars owed reaches zero | <ul style="list-style-type: none"> • The amount of euros owed drops • The amount of dollars increases significantly |

Table 8: Sum-up of the consequences of the operational disruption of a significant € local bank on the FX exposures faced by the global banks, depending on the level of liquidity available (the same in the two systems) and on the settlement mechanism used for FX transactions

As expected, we observe an extremely high level of euros owed during the crisis. This is simply due to the large number of € leg of FX transactions queued during the outage. Indeed, as all € balances quickly reach zero after the outage, the € leg of the FX trades can not settle, hence the exposure.

A less expected result is the significant increase in the amount of dollars owed during the recovery. This can be explained by the combination of two effects. First, the recovery of the affected € leads to the settlement of a very large number of queued € leg of FX transactions. At the end of the recovery phase, the FX banks therefore have very large € balances. The € leg of all new incoming FX trades is thus settled immediately. Second, the crisis has strongly modified the distribution of \$ deposits among banks. As a consequence, some FX banks have ended up with large \$ balances while other FX banks lack \$ and will accumulate queues until the system returns to its normal steady-state distribution. Many \$ leg of FX transactions thus remain queued in the meanwhile, creating FX exposures.

A more detailed description of the unfolding of events in the different cases is provided in the Annex, beneath charts 12 and 13. An analysis of the cross-currency channels of crisis propagation is proposed in section 4.3.

4.3. Discussion of the cross-currency channels of crisis propagation

Chart 14 provides a graphical explanation of the reasons why the effects of the outage on the two systems are different in each of the considered cases, highlighting the disruption path. Based on our observations presented in section 4.1 and 4.2, we can conclude that the disruption spreads from RTGS € to RTGS \$ through three different channels:

Channel 1: Low € balances at the Central bank prevents settlement of \$ leg of FX transactions in RTGS \$ in the PvP case, and generate FX exposures in the non-PvP case

In the PvP case, the lack of liquidity in RTGS € prevents the emitted FX transactions from settling. This has two main effects.

First, the settlement activity within RTGS \$ is quickly reduced, as the system is deprived from the activity related to the settlement of the \$ leg of FX transactions.

Second, the queuing of a very large number of FX transactions during the crisis (due to the lack of € balances) decreases the level of the \$ customer deposits in the FX banks, since the \$ leg of the FX transaction is also queued. The \$ holdings of the FX banks' customers are thus made unavailable for the FX banks' customers to make new FX trades or new local \$ payments. Note that the Central Bank balance of the FX banks is not decreased by the queuing of the FX transactions, but only the customer deposits at the FX banks. This effect leads to a slow decrease of the emission of new FX transactions and new \$ local payments.

In the non-PvP case, the lack of liquidity in RTGS € does not prevent the \$ leg of the emitted FX transactions from settling, but the amount of resulting FX exposures become extremely high.

Channel 2: Low € customer deposits lead to fewer emitted FX trades

In all cases, the low levels of customer deposits in RTGS € prevent the emission of new FX transactions. Basically, the FX banks' customers (who can happen to be the bank itself) stop making new FX trades simply because the level of their account at their bank has reached zero. Another possible interpretation would be that the productive agents stop trading because they face a very high uncertainty on their real € position. The end of the submission of new FX trades also deprives RTGS \$ from activity related to the settlement of the \$ leg of FX transactions, but this effect is slower than the one of channel 1) and is limited to current (rather than cumulative) rate of FX trades.

Channel 3: As not all banks are similarly affected, the system becomes unbalanced

The FX banks are affected by the crisis at different speeds. The FX banks for which the operationally disrupted bank is an important counterparty are hit sooner and see their level of € deposits decrease more rapidly. These most affected FX banks therefore start making fewer € for \$ trades and become net \$ sellers. This leads to the redistribution of \$ deposits within the global banks.

At high liquidity, this change in the \$ deposits of the FX banks spreads and leads to a redistribution of the customer deposits within the whole RTGS \$. After the recovery, the system is therefore unbalanced and slowly returns to its normal steady-state. This phenomenon generates some extra local and FX payments, hence the higher activity observed in the aftermath of the crisis at high liquidity.

At low liquidity, this change in the distribution of \$ deposits remains limited to the circle of FX banks since there is not enough liquidity available to allow for the spread of the disruption to \$ local banks. It does however cause \$ liquidity to be concentrated in the FX banks, depriving the local \$ banks of liquidity. Many \$ local payments are thus put into the queue, explaining both in the non-PvP low liquidity case the decrease of the settlement rate in RTGS \$ during the crisis and its increase during the recovery.

In the non-PvP case, the redistribution of \$ deposits and balances among the FX banks in the crisis phase, explains the sudden increase of the amount of dollars owed. Indeed, some FX banks have ended up with large \$ balances while other FX banks lack \$ and will accumulate queues until the system returns to its normal steady-state distribution. Many \$ leg of FX transactions thus remain queued in the meanwhile, creating FX exposures.

5. Conclusion

In this paper we present a parsimonious model of linked gross payment systems, each operating in a distinct currency. The model incorporates two forms of interdependencies between the two systems: dual participation by large banks in both systems and a PvP mechanism that synchronizes the settlement of the two legs of FX transactions. The dual participation, and the resulting common inflow of FX trades, creates an institution-based interdependency between the two systems. As a result, the activity of the two systems is shown to become correlated at high levels of liquidity, in the sense that a period of high settlement rate within one RTGS is statistically likely to correspond to a period of high settlement rate within the other RTGS. In the model, FX trades are settled on a gross basis, either PvP (both legs of the FX transactions can

only be settled simultaneously) or non-PvP (both legs of the FX transactions are settled independently). The use of a PvP mechanism to settle FX trades results in a system-based interdependency between the two systems. Consequently, the activity of the two systems is shown to become correlated at low levels of liquidity.

During normal operation, when the FX trades are settled non-PvP, some credit exposures are created between the global banks that engage in FX trading. Those exposures are shown to be dependent on the level of liquidity present in each RTGS. Moreover, it appears that a structural liquidity imbalance between the two systems leads to very high exposures, by acting in a similar way as a time zone difference between the two systems. When the FX trades are settled PvP, the credit exposures between the global banks vanish. However, the PvP mechanism creates another kind of interdependency between the two systems. The model shows that in the PvP case, the average level of queuing within one RTGS does not depend only on its own level of liquidity like in the non-PvP case, but also on the level of liquidity in the other RTGS. More specifically, when liquidity is decreased within the “less liquid” system, the level of queuing increases significantly within the “more liquid” system. This effect appears especially strong for intermediate levels of liquidity in the “more liquid” system. In addition, we also observe that the level of queuing in the “less liquid” system decreases when the liquidity is increased in the “more liquid” RTGS.

The consequences of an operational disruption that would prevent a local euro bank from transmitting payment orders to the Central Bank are investigated with the model. As expected, the outage eventually leads to the halt of all settlements in the euro RTGS, this system becoming totally deprived of liquidity. The outage is also shown to have an impact on the activity of the dollar RTGS in all the investigated cases. The crisis propagates to the dollar RTGS through three different channels.

The first identified channel of transmission is the settlement of FX transactions. In the PvP case, the lack of liquidity in RTGS € prevents the dollar leg of the emitted FX transmissions from settling in RTGS \$. This both immediately decreases the settlement rate of RTGS \$, deprived from the activity related to FX transactions, and also eventually leads to fewer \$ local payments being emitted, as the dollars “trapped” within the queued FX transactions, are lost for the customers of the \$ banks. In the non-PvP case, the lack of liquidity in RTGS € results in a dramatic increase in the level of FX exposures.

The second identified channel is the emission of new FX transactions. As the global banks’ customers see the level of their euro deposits decrease, all deposits becoming trapped at the defaulting bank, they engage in fewer new FX trades. As a result, even in the non-PvP case, all activity resulting from the settlement of FX transactions stops in RTGS \$.

The third identified channel is more complex. The FX banks are affected by the crisis at different speeds, depending on their relation to the defaulting bank. The most affected FX banks therefore start making fewer € for \$ trades and become net \$ sellers. This leads to the redistribution of \$ deposits within the global banks. At low liquidity, this disruption in the distribution of \$ deposits eventually leads to the queuing of several local \$ payments during the crisis. At high liquidity, the redistribution of \$ deposits affects the whole RTGS \$. This results in a long term increase of the settlement rate in RTGS \$ in the aftermath of the crisis, as the system returns to equilibrium.

The results obtained so far by the model can already be used to qualitatively describe and document the effect of the interdependency created by the FX transactions and the possible PvP mechanism on the activity of the two systems during normal operation. The model also allows us to identify the channels of propagation of the crisis across currency zones.

References

Bank for International Settlements (2001). Group of Ten Report on Consolidation in the Financial System, January, p 29

Bech, Morten and Rod Garratt (2006). Illiquidity in the Interbank, Payment System following Wide-Scale Disruptions, FRB of New York Staff Report No. 239

Becher Chris, Stephen Millard and Kimmo Soramäki (forthcoming). Network Topology of CHAPS Sterling, Bank of England Working Paper.

Bedford, Paul, Stephen Millard and Jin Yang (2005). Analysing the impact of operational incidents in large-value payment systems: a simulation approach. In Leinonen (ed). Liquidity, risks and speed in payment and settlement systems –a simulation approach, Bank of Finland Studies in Economics and Finance E:31.

Beyeler, Walter, Robert J. Glass, Morten L. Bech and Kimmo Soramäki (2007). Congestion and cascades in payment system, Physica A, Volume 384, Issue 2, pp 693-718

Galbiati, Marco and Kimmo Soramäki (forthcoming). Agent based model of payment systems, Bank of England Discussion Paper.

Hellqvist, Matti and Heli Snellman (2007). Simulation of operational failures in equities settlement. M. In Leinonen (ed). Simulation studies of liquidity needs, risks and efficiency in payment networks - Proceedings from the Bank of Finland Payment and Settlement System Seminars 2005-2006, Bank of Finland Studies in Economics and Finance, E:39.

Inaoka, H, T. Ninomiya, K. Taniguchi, T. Shimizu, and H. Takayasu (2004). Fractal Network derived from banking transaction - An analysis of network structures formed by financial institutions.”, Bank of Japan Working Paper No. 04-E-04

Ledrut, Elisabeth (2006), A tale of the water-supplying plumber: intraday liquidity provision in payment systems, DNB Working Papers No. 099

Lubloy, Agnes (2006). Topology of the Hungarian large-value transfer system. Magyar Nemzeti Bank Occasional Paper 57.

Martin, Antoine and James McAndrews (2007). Liquidity-Saving Mechanisms, FRB of New York Staff Report No. 282

Mazars, Emmanuel and Guy Woelfel (2005). Analysis, by simulation, of the impact of a technical default of a payment system participant. In Leinonen (ed). Liquidity, risks and speed in payment and settlement systems –a simulation approach, Bank of Finland Studies in Economics and Finance E:31.

McVanel, Darcy (2005). The impacts of unanticipated failures in Canada's Large Value Transfer System. . Bank of Canada Working Paper, No. 25.

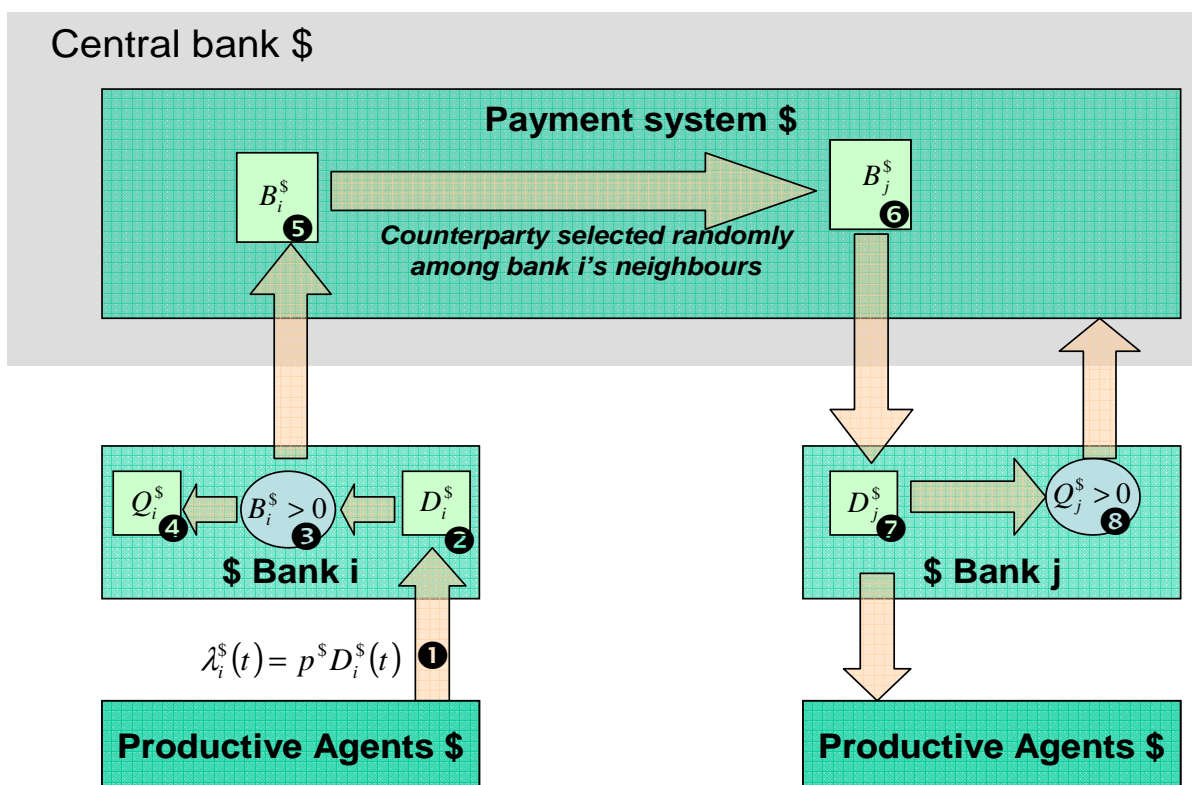
Mills, David and Travis Nesmith (forthcoming), 'Risk and Concentration in Payments and Securities Settlement Systems, Journal of Monetary Economics.

Renault, Fabien, Beyeler, Walter, Robert J. Glass, Kimmo Soramäki and Morten L. Bech (2007). Congestion and cascades in coupled payment systems, Joint BoE/ECB conference on “Payments and monetary and financial stability”, November 2007.

Schmitz, Stefan, Claus Pühr, Hannes Moshhammer, Marin Hausmann, Ulrike Elsenhuber (2006). Operational risk and contagion in the Austrian large-value payment system Artis. Österreichische Nationalbank Financial stability report, Iss. 11, pp. 96-113.

Soramäki, Kimmo, Morten L. Bech, Jeffrey Arnold, Robert J. Glass, Walter Beyeler (2007). The topology of interbank payment flows, *Physica A*, Vol. 379, pp 317-333.

Annexes



- ❶ Bank i receives a continuous stream of payment orders from its depositors. The average volume of payment orders received by a bank is taken as proportional to the current level of deposits at this bank.
- ❷ Depositor account of bank i , D_i^s is debited.
- ❸ The RTGS account balance of bank i , B_i^s , is checked.
- ❹ If Bank i does not have sufficient liquidity at the Central Bank to settle the payment, (since we consider only payments of unit size, we just check if B_i^s is greater than zero), the payment is queued.
- ❺ Otherwise, the payment is settled and B_i^s is decremented.
- ❻ The receiving bank, bank j is chosen randomly among Bank i 's counterparties, proportionally to ω_{ij} . The RTGS account of bank j is incremented.
- ❼ The depositor account of bank j is incremented. The probability of bank j to receive a payment order from one of its depositors is thus mechanically increased.
- ❽ If bank j has some outgoing queued payments waiting, the payment with the earliest submission time is released (FIFO order).

Chart 1: Processing of local payments




Chart 2: Correlation of the settlement rates in the two systems, non-PvP case




Chart 3: Correlation of the settlement rates in the two systems, non-PvP case

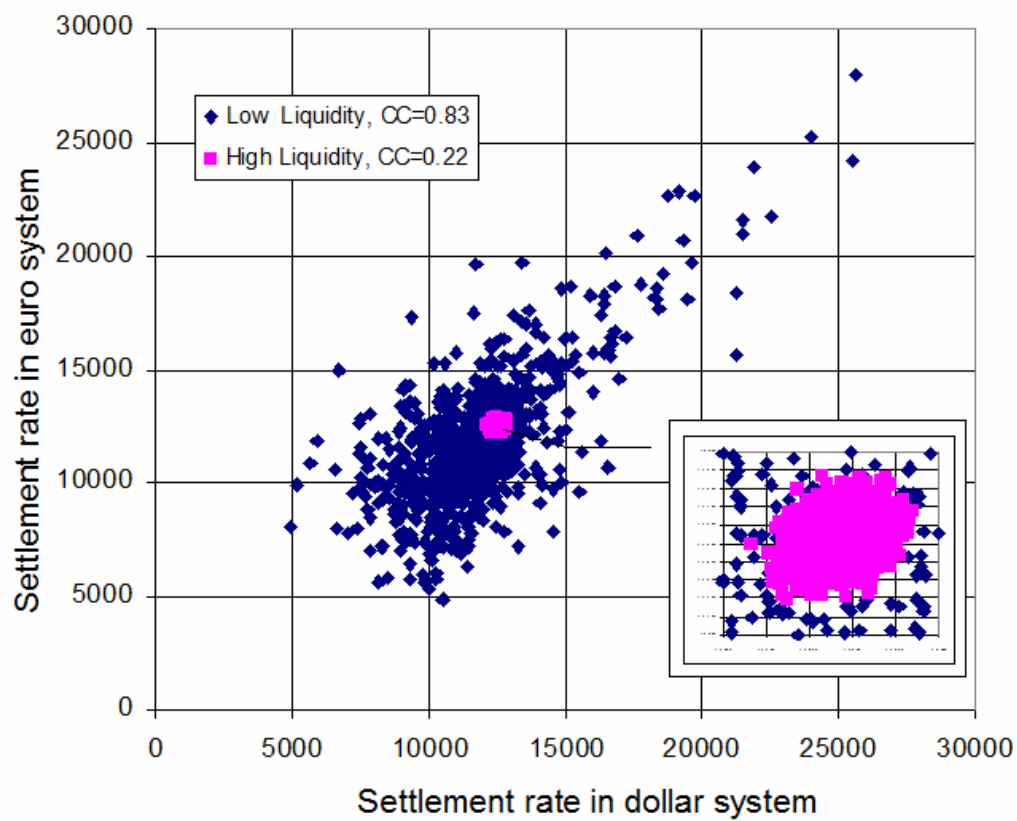


Chart 4: Correlation of the settlement rates in the two systems, PvP case

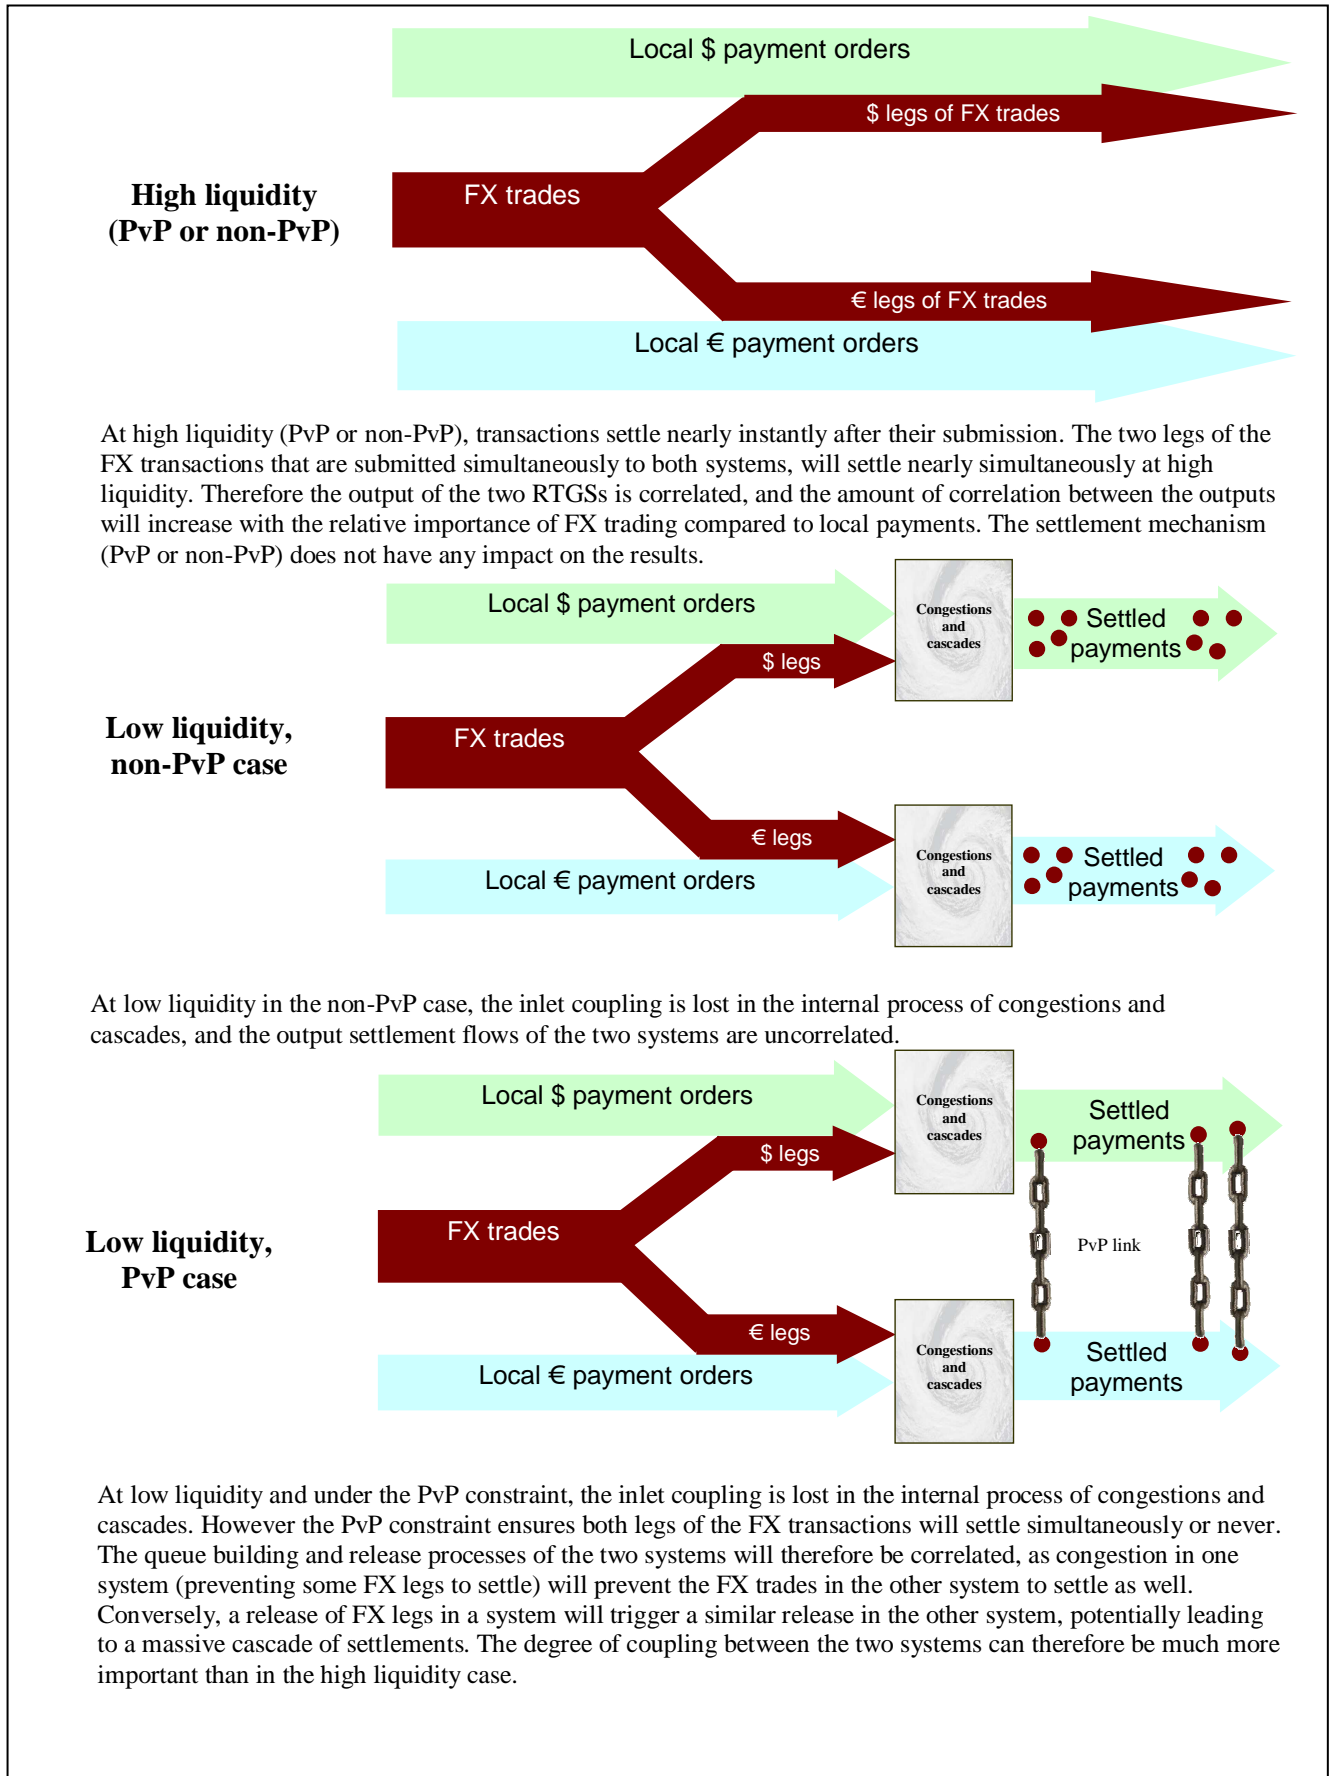


Chart 5: System interdependencies




Chart 6: Time series of the settlement rate in the two systems. non-PvP case




Chart 7: Time series of the settlement rate in the two systems. PvP case

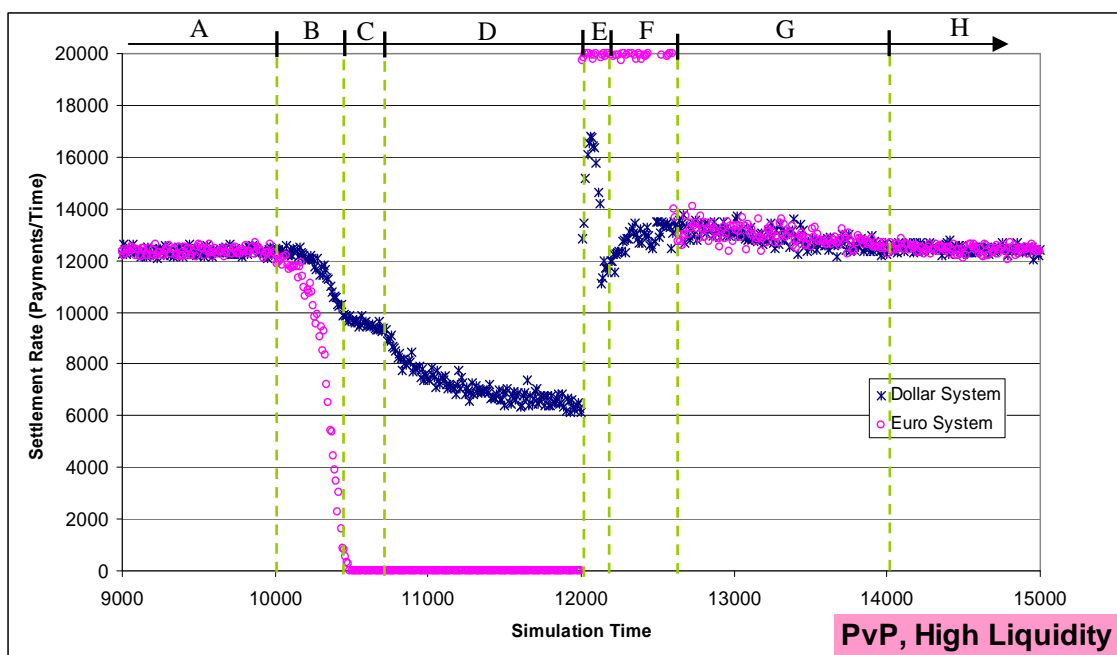


Chart 8: Time series of the settlement rate in the two RTGS systems, in the PvP case, with a high level of liquidity in both systems. A large local € bank faces an operational outage between time 10000 and 12000.

- **Period A:** The two RTGS systems are operating in steady state, each settling around 12 300 payments per unit of time. FX transactions account for 2 300 settlements per unit of time and local payments represent around 10 000 payments per unit of time in both systems.
- **Period B:** At time 10 000, a large € local bank becomes unable to send payment orders to the Central Bank. RTGS € quickly finds itself totally deprived of liquidity (i.e. all banks but the bank affected by the operational problem end up with a zero balance at the Central Bank €). As a consequence the settlement rate in RTGS € reaches zero. As FX transactions are settled PvP, the lack of liquidity within RTGS € also prevents the \$ leg of the transactions from settling in RTGS \$. The settlement activity within RTGS \$ thus decreases to a value of around 10 000 settlements per unit of time, corresponding to the domestic activity within RTGS \$.
- **Period C:** The situation remains fairly stable. The “local” activity of RTGS \$ is not yet affected while all local settlement activity within RTGS € as well as the whole FX settlement activity have come to a stop.
- **Period D:** While the € balance of the global banks at the € Central Bank is now zero, there are still € customer deposits in the global banks that continue to generate FX trades, although these trades are not settled (because all € balances are zero). As a consequence, all FX trades are queued, decreasing the level of \$ customer deposits within the global banks (the customer accounts are decreased as soon as they submit an order, even if the trade does not settle immediately). The decrease of the level of \$ customer deposits leads to a reduction of the local activity within RTGS \$, as less deposits are now available for local \$ clients to initiate payments. The settlement rate in RTGS \$ thus decreases until all global banks finally have a zero level of € customer deposits. FX trades are thus no longer generated which halts the decrease of \$ customer deposits, and allows RTGS \$ to reach a new steady state with a large fraction of \$ deposits trapped in FX queues.
- **Period E and F:** At time 12 000, the large € local bank that had been affected by an operational disruption recovers and finally transmits to the € Central Bank the payment orders of its clients.

By re-injecting liquidity within RTGS €, the recovery triggers a massive cascade of settlements of local € payments and of FX transactions. The settlement rate within RTGS € thus reaches its maximal rate, 20 000 settlements per unit of time, corresponding to the maximal system processing capacity. The settlement activity within RTGS \$ follows a more complex pattern. As RTGS € recovers, the settlement of the \$ leg of the bulk of FX transactions that were previously queued tends to increase the settlement rate in RTGS \$ (this effect dominates during period E). However, as the level of \$ customer deposits is still lower than normal (until the end of period F when all queued FX transactions have been settled), the \$ local activity remains lower than it was during period A. This effect dominates during period F.

- **Period G:** All previously queued FX transactions and local € payments have now been settled. However, as a result of the turmoil, the distribution of customer deposits within the different banks has moved away from its period A steady state. In particular the global banks that were most affected by the operational disruption (those for which the disrupted bank is an important counterparty) quickly ended up having a zero level of € customer deposits at the beginning of period D. They therefore stopped to engage in € for \$ FX trades while continuing to make \$ for € trades (these trades remained queued until the recovery, due to the lack of € balances). At the end of the crisis (at the end of period F), these global banks therefore have a higher than usual level of € customer deposits and a lower than usual level of \$ customer deposits. Due to the intrinsic stability of the model, a progressive return to the steady state is observed. This return to normal generates a surplus of FX trades and both \$ and € local payments, hence explaining the higher than normal level of settlement rate during period G.
- **Period H:** The system has returned to steady state.

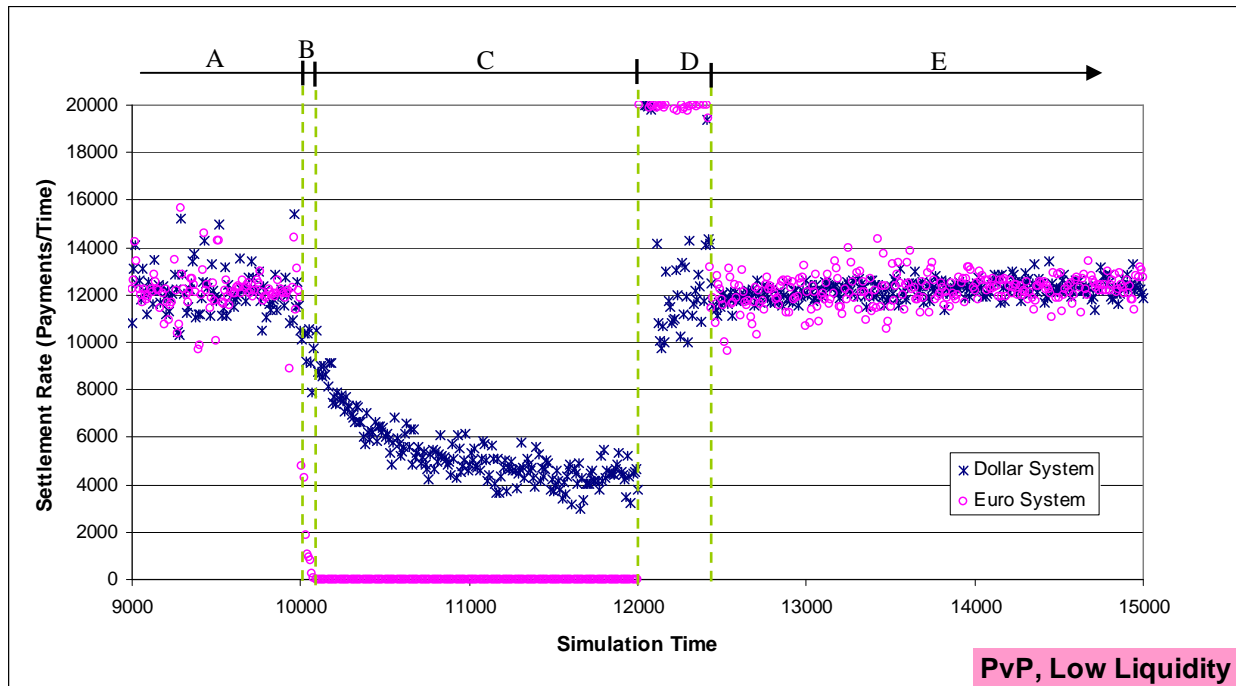


Chart 9: Time series of the settlement rate in the two RTGS systems, in the PvP case, with a low level of liquidity in both systems. A large local €bank faces an operational outage between time 10000 and 12000.

- Period A:** The two RTGS systems are operating in steady state, each settling around 12 200 payments per unit of time. FX transactions account for 2 300 settlements per unit of time and local payments represent around 9 900 payments per unit of time in both systems. Due to the low level of liquidity available in both RTGS systems, we observe a high variability in the two settlement rates, corresponding to a regime of congestions and cascades, and a slightly smaller settlement rate.
- Period B:** At time 10 000, a large € local bank becomes unable to send payment orders to the Central Bank. RTGS € quickly finds itself totally deprived of liquidity (i.e. all banks but the bank affected by the operational problem end up with a zero balance at the Central Bank €). As a consequence the settlement activity in RTGS € reaches zero. As FX transactions are settled PvP, the lack of liquidity within RTGS € also prevents the \$ leg of the transactions from settling in RTGS \$. The settlement activity within RTGS \$ thus decreases to a value of around 10 000 settlements per unit of time, corresponding to the domestic activity within RTGS \$.
- Period C:** While the € balance of the global banks at the € Central Bank is now zero, there are still € customer deposits in the global banks that continue to generate FX trades, although these trades are not settled (because all € balances are zero). As a consequence, all FX trades are queued, decreasing the level of \$ customer deposits within the global banks (the customer accounts are decreased as soon as they submit an order, even if the trade does not settle immediately). The decrease of the level of \$ customer deposits leads to a reduction of the local activity within RTGS \$, as less deposits are now available to local \$ clients to initiate payments. The settlement rates in RTGS \$ thus decreases until all global banks finally have a zero level of € customer deposits. Compared with the PvP, high liquidity case, the situation is worsened by the fact that the low liquidity available in RTGS \$ substantially increases queueing in the local banks in the following way: global banks, whose level of \$ customer deposits decreased significantly because of the FX trades are not emitting as many payments as before, become net receivers of liquidity. However the liquidity available on the \$ balances of the other banks does not allow for this net transfer of liquidity. As a result, many local \$ payments end up queued, reducing further the level of \$ customer deposits, and thus the settlement activity within RTGS \$.

- **Period D:** At time 12 000, the large € local bank that had been affected by an operational disruption recovers and finally transmits to the € Central Bank the payment orders of its clients. By re-injecting liquidity within RTGS €, the recovery triggers a massive cascade of settlements of local € payments and of FX transactions. The settlement rate within RTGS € thus reaches its maximal rate, 20 000 settlements per unit of time, corresponding to the maximal system processing capacity. The settlement activity within RTGS \$ follows a more complex pattern. As RTGS € recovers, the settlement of the \$ leg of the bulk of FX transactions that were previously queued and of the many queued \$ local payments drastically increases the settlement rate in RTGS \$. Because there are more queued payments than in the high liquidity case, the settlement rate within RTGS \$ even reaches its maximum rate of 20 000 settlements per unit of time whereas this did not occur with high liquidity. The PvP constraint, and low liquidity level, cause the settlement rate to vary widely from point to point. Because the level of \$ customer deposits is still lower than normal (until the end of period E when all excess queued transactions have been settled), the arrival rate of new \$ local instructions is lower than what it was during period A, which explains why the \$ settlement rate scatter plots exhibits a slowing increasing trend during periods D and E.
- **Period E:** Most previously queued FX transactions and local € payments have now been settled by the beginning of period E, and the remainder are slowly worked down during the rest of this period. This relaxation is slowed by the scarcity of liquidity, which tends to be concentrated in the global banks during this period. The average activity within the two systems is now comparable to its period A characteristics. Compared with the high liquidity PvP case, we do not observe a higher than normal activity. This is due to the fact that the distribution of customer deposits has not been really moved out of equilibrium since not enough liquidity was available to ensure this displacement. Once all queued payments are settled, the system is therefore in the same state as it was during period A.

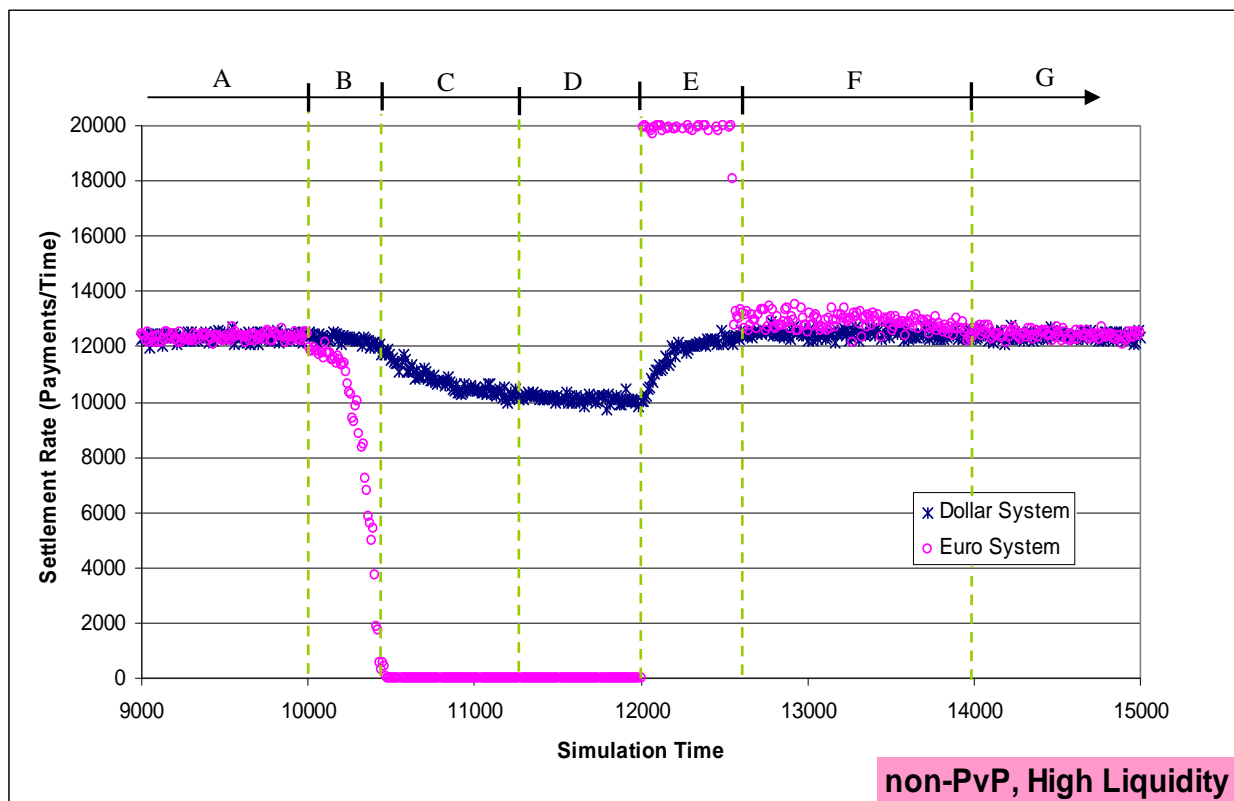


Chart 10: Time series of the settlement rate in the two RTGS systems, in the non-PvP case, with a high level of liquidity in both systems. A large local € bank faces an operational outage between time 10000 and 12000.

- **Period A:** The two RTGS systems are operating in steady state, each settling around 12 300 payments per unit of time. FX transactions account for 2 300 settlements per unit of time and local payments represent around 10 000 payments per unit of time in both systems.
- **Period B:** At time 10 000, a large € local bank becomes unable to send payment orders to the Central Bank. RTGS € quickly finds itself totally deprived of liquidity (i.e. all banks but the bank affected by the operational problem end up with a zero balance at the Central Bank €). As a consequence the settlement activity in RTGS € reaches zero. As FX transactions are settled non-PvP, the lack of liquidity within RTGS € does not prevent the \$ leg of the transactions from settling in RTGS \$. The settlement activity within RTGS \$ thus remains fairly constant at its steady state value.
- **Period C:** While the € balance of the global banks at the € Central Bank is now zero, there are still € customer deposits in the global banks that continue to generate FX trades, although only the \$ leg of these trades are settled (because all € balances are zero). However, the level of € customer deposits of the global banks decreases, because these banks issue both € local payment orders and €-for-\$ FX trades which decrement their level of € customer deposits. As a result the level of FX trading decreases and reaches zero at the end of period C when the level of € customer deposits of all global banks reaches zero.
- **Period D:** The situation remains fairly stable. The “local” activity of RTGS \$ is not affected while all local settlement activity within RTGS € as well as the whole FX settlement activity has come to a stop.
- **Period E:** At time 12 000, the large € local bank that had been affected by an operational disruption recovers and finally transmits to the € Central Bank the payment orders of its clients. By re-injecting liquidity within RTGS €, the recovery triggers a massive cascade of settlements of local € payments and of € leg of FX transactions. The settlement rate within RTGS € thus

reaches its maximal rate, 20 000 settlements per unit of time, corresponding to the maximal system processing capacity. As the queued local € payments and € legs are settled, the level of € customer deposits of the global banks increases, and the global banks resume their FX trading operations. The activity within RTGS \$ thus increases, as the \$ leg of these new FX transactions settle.

- **Period F:** All previously queued € leg of FX transactions and local € payments have now been settled. However, as a result of the turmoil, the distribution of customer deposits within the different banks is not at its steady state. Indeed the global banks that were most affected by the operational disruption (those for which the disrupted bank is an important counterparty) quickly ended up having a zero level of € customer deposits during period C. They therefore stopped to engage in €-for-\$ FX trades while continuing to make \$-for-€ trades (the € leg of these trades remained queued until the recovery, due to the lack of € balances). At the end of the crisis (at the end of period E), these global banks therefore have a higher than usual level of € customer deposits. Due to the intrinsic stability of the model, a progressive return to the steady state is observed. This return to normal generates a surplus of FX trades and € local payments, hence explaining the higher than normal level of settlement rate during period F. Contrary to period G in the PvP high liquidity case, we observe that, although both RTGS have a higher than normal settlement rate, the level of activity in RTGS € is higher than in RTGS \$. This is due to the fact that in the PvP case, the FX trades made by the FX banks are systematically put in the queue. This “saves” the \$ liquidity of the FX banks, allowing the system to wander off farther from its steady state equilibrium.
- **Period G:** The system has returned to steady state




Chart 11: Time series of the settlement rate in the two RTGS systems, in the non-PvP case, with a high level of liquidity in both systems. A large local € bank faces an operational outage between time 10000 and 12000.

- Period A:** The two RTGS systems are operating in steady state, each settling around 12 200 payments per unit of time. FX transactions account for 2 300 settlements per unit of time and local payments represent around 9 900 payments per unit of time in both systems. Due to the low level of liquidity in both systems, a high degree of variability can be observed.
- Period B:** At time 10 000, a large € local bank becomes unable to send payment orders to the Central Bank. RTGS € quickly finds itself totally deprived of liquidity (i.e. all banks but the bank affected by the operational problem end up with a zero balance at the Central Bank €). As a consequence the settlement activity in RTGS € reaches zero. As FX transactions are settled non-PvP, the lack of liquidity within RTGS € does not prevent the \$ leg of the transactions from settling in RTGS \$. The settlement activity within RTGS \$ thus remains fairly constant at its steady state value. Compared to the non-PvP high liquidity case, period B is much shorter, since as the € balances are much lower in steady state, they are faster depleted when the disruption occurs.
- Period C:** While the € balance of the global banks at the € Central Bank is now zero, there are still € customer deposits in the global banks that continue to generate FX trades, although only the \$ leg of these trades are settled (because all € balances are zero). However, the level of € customer deposits of the global banks decreases, because these banks issue both € local payment orders and €-for-\$ FX trades which decrements their level of € customer deposits but never receive € from their counterparties. As a result the level of FX trading decreases and reaches zero at the end of period C when the level of € customer deposits of all global banks reaches zero.
- Period D:** The situation remains fairly stable. All local settlement activity within RTGS € as well as the whole FX settlement activity have come to a stop. Contrary to the non-PvP high liquidity case, the “local” activity of RTGS \$ is lower than its steady state value of 9 900 settlements per

unit of time. This is due to the fact that in the low liquidity case, large numbers of \$ local payments are queued during period D, leading to a decrease of the level of \$ deposits, and thus to a decrease of the settlement activity in RTGS \$. This large queuing of \$ local payments is a direct consequence of the destabilization of the distribution of \$ deposits created by the outage. Indeed, at the beginning of period D, the distribution of \$ deposits is not at its steady state value. In particular the global banks that were most affected by the operational disruption (those for which the disrupted bank is an important counterparty) quickly ended up having a zero level of € customer deposits. They therefore stopped to engage in €-for-\$ FX trades while continuing to make \$-for-€ trades (the € leg of these trades remained queued until the recovery, due to the lack of € balances) hence decreasing their level of \$ deposits. While the distribution of \$ deposits is not at equilibrium, its return to steady state is very slow because it is controlled by the very low level of liquidity available in RTGS \$, contrary to the high liquidity non-PvP case where RTGS \$ is liquid enough to quickly find its way back to equilibrium. As a result, large numbers of \$ local payments are queued, leading to a decrease of the level of \$ deposits, and thus to a decrease of the settlement activity in RTGS \$.

- **Period E:** At time 12 000, the large € local bank that had been affected by an operational disruption recovers and finally transmits to the € Central Bank the payment orders of its clients. By re-injecting liquidity within RTGS €, the recovery triggers a massive cascade of settlements of local € payments and of € leg of FX transactions. The settlement rate within RTGS € thus reaches its maximal rate, 20 000 settlements per unit of time, corresponding to the maximal system processing capacity. As the queued local € payments and € legs are settled, the level of € customer deposits of the global banks increases, and the global banks resume their FX trading operations. The activity within RTGS \$ thus increases, as the \$ leg of these new FX transactions settle. The re-start of the FX trading also allows the distribution of \$ customer deposits to return to its initial steady state, consequently leading to the settlement of the bulk of the “excess” queued \$ local payments accumulated during period F. We therefore observe a significant peak in activity in RTGS \$ following the recovery of the € local bank affected by the operational problem, which does not occur in the case of high liquidity.
- **Period F:** The system has returned to steady state

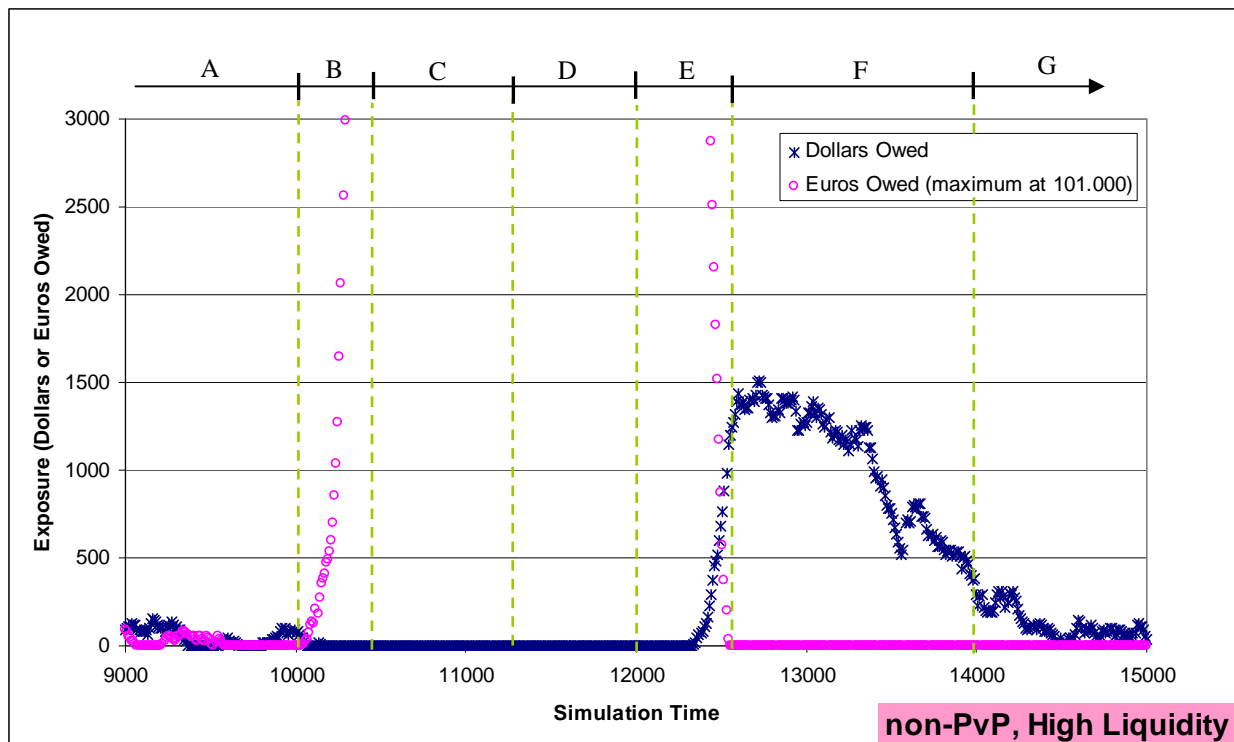


Chart 12: Time series of the FX exposures (Dollars Owed or Euros Owed), in the non-PvP case, with a high level of liquidity in both systems. A large local € bank faces an operational outage between time 10000 and 12000.

The presented time series correspond to the same simulation as chart 10. The periods A to F, defined on chart 10, have been reported on chart 12 as well.

- **Period A:** The two RTGS systems are operating in steady state. As the level of liquidity is high in the two systems, most payments are settled immediately. Few FX legs are queued, and the level of exposure remains therefore low.
- **Periods B-C-D-E:** At time 10 000, a large € local bank becomes unable to send payment orders to the Central Bank. RTGS € quickly finds itself totally deprived of liquidity, and the settlement activity within RTGS € quickly stops (cf chart 10). As the dollar leg of the FX transactions still settle, the amount of euros owed due to FX transactions (ie the number of unsettled euro leg of FX transactions for which the dollar leg has settled) reaches an extremely high level (around 100 000).
- **Period F:** The peak in dollars owed right after the recovery means that during this period the euro legs of FX transactions can settle faster, on average, than the dollar legs. This can be explained by the combination of two effects:
 - the recovery of the affected bank leads to the settlement of a very large number of queued € leg of FX transactions. This keeps liquidity in the € system within the global banks. At the end of the recovery phase, the FX banks therefore have very large € balances. The € leg of all new incoming FX trades are thus settled immediately.
 - the crisis has strongly modified the distribution of \$ deposits among banks. As a consequence, some FX banks have ended up with large \$ balances while other FX banks lack \$ and will accumulate queues until the system returns to its normal steady-state distribution. Many \$ leg of FX transactions remain queued in the meanwhile, creating FX exposures.
- **Period G:** The system has returned to steady state

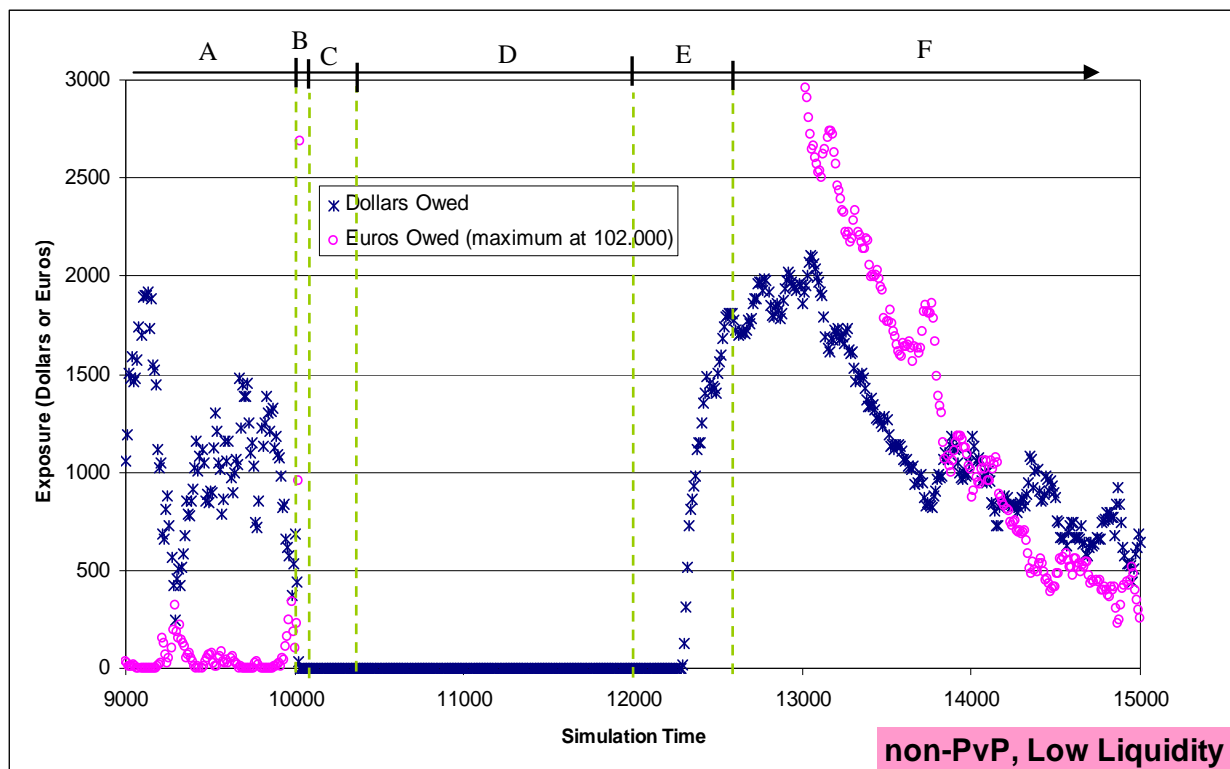


Chart 13: Time series of the FX exposures (Dollars Owed or Euros Owed), in the non-PvP case, with a low level of liquidity in both systems. A large local € bank faces an operational outage between time 10000 and 12000.

The presented time series correspond to the same simulation as chart 11. The periods A to F, defined on chart 11, have been reported on chart 13 as well.

- **Period A:** The two RTGS systems are operating in steady state. As the level of liquidity is low in both systems, some payments are queued. In average over a long period the amount of dollars owed is very close to the amount of euros owed. However, due to the internal dynamics of queue building and release of each RTGS, periods during which the amount of euros owed exceed the amount of dollars owed alternate with periods during which the amount of dollars owed exceed the amount of euros owed (this is the case during period A).
- **Periods B-C-D-E:** At time 10 000, a large € local bank becomes unable to send payment orders to the Central Bank. RTGS € quickly finds itself totally deprived of liquidity, and the settlement activity within RTGS € quickly stops (cf chart 11). As the dollar leg of the FX transactions still settle, the amount of euros owed due to FX transactions (ie the number of unsettled euro leg of FX transactions for which the dollar leg has settled) reaches an extremely high level (around 100 000).
- **Period F:** The peak in dollars owed right after the recovery means that during this period the euro legs of FX transactions can settle faster, on average, than the dollar legs. This can be explained by the combination of the same two effects presented for chart 12 in the high liquidity case (the massive release of € to the global banks due to the settlement of queued FX transactions) and the destabilization of the distribution of \$ deposits among the global banks. In contrast to the high liquidity case of chart 12, we see larger episodic exposure before and after the disruption due to the longer settlement delays caused by scarce liquidity, and a more protracted recovery period as the rate of restoration of equilibrium is throttled by liquidity limits.
- **End of period F:** The system has returned to steady state

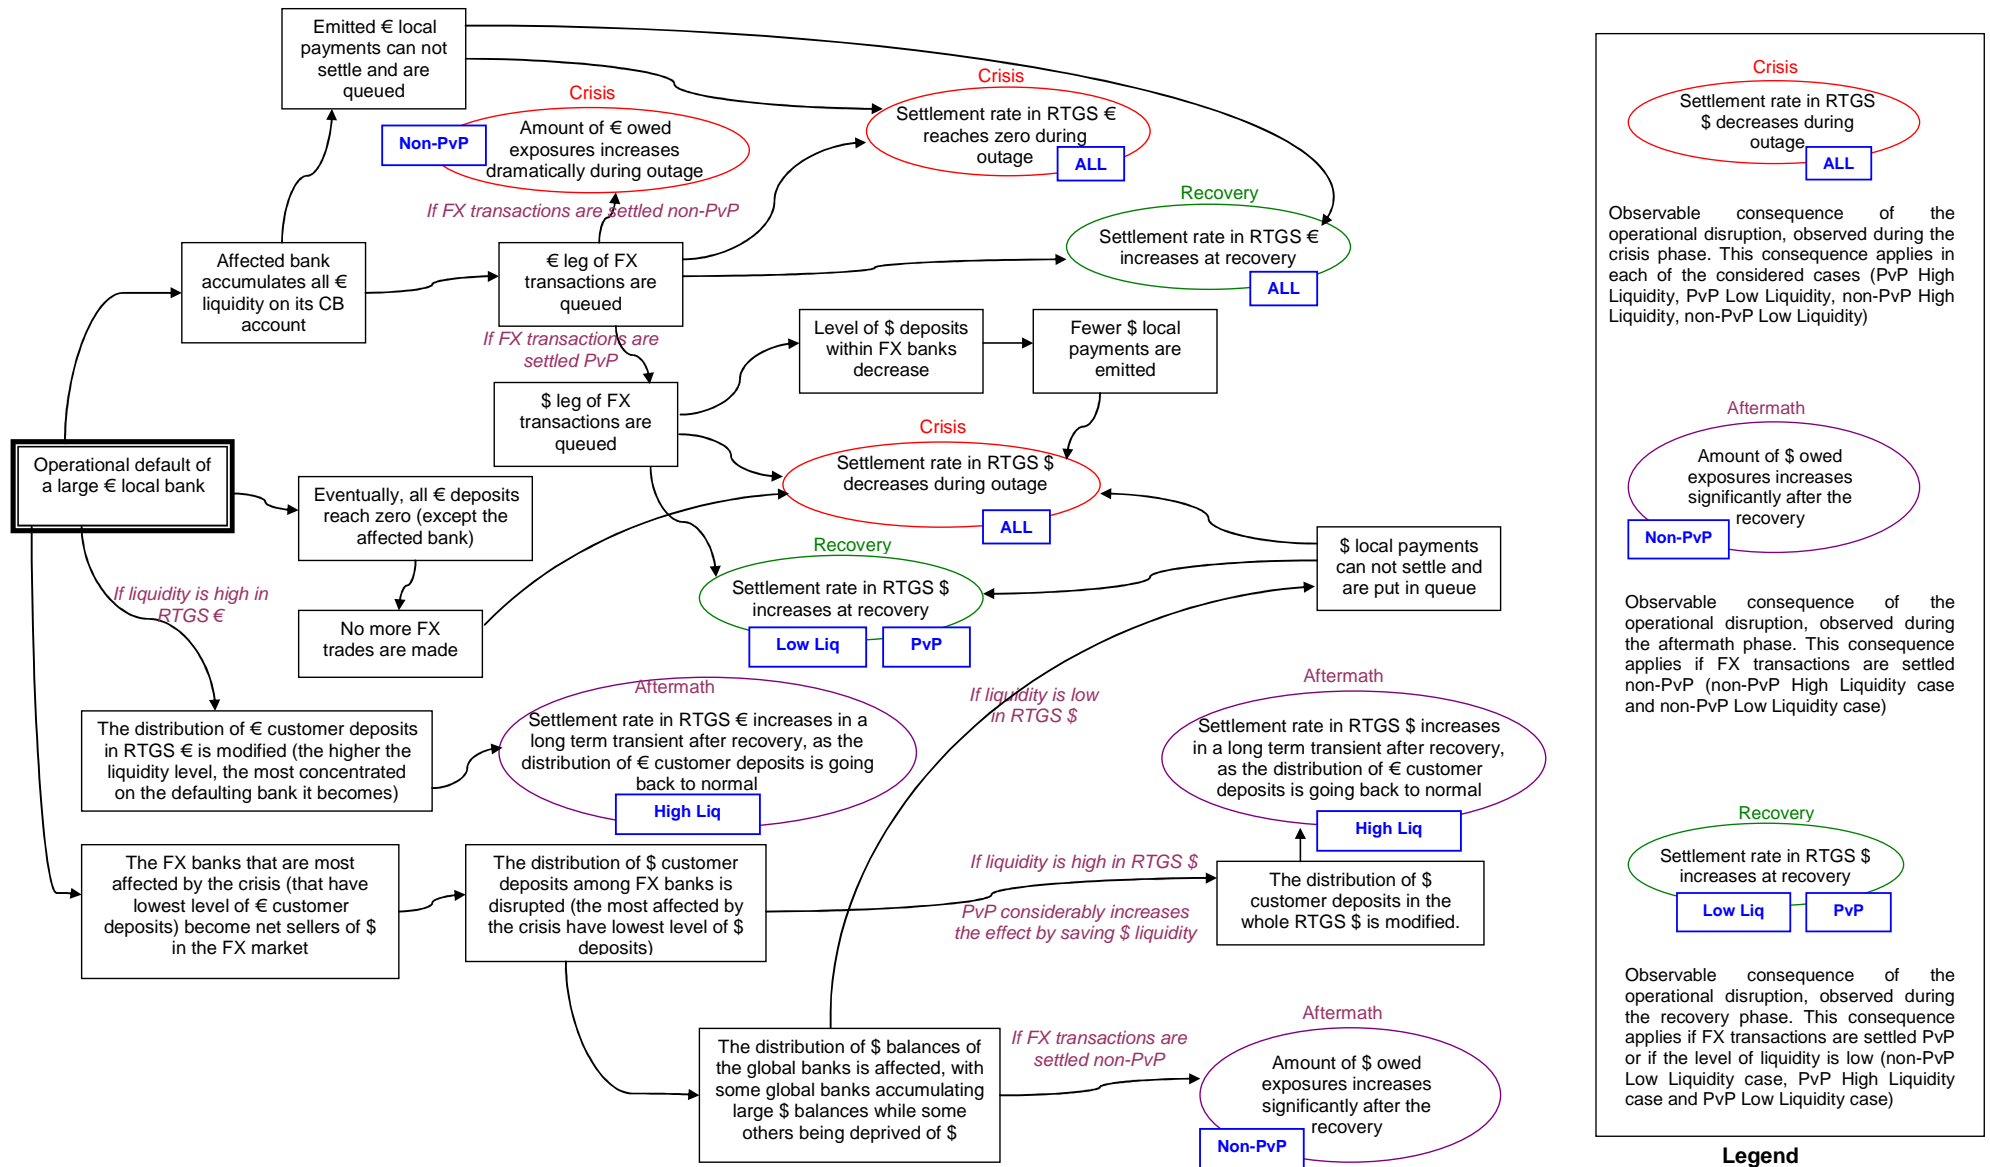


Chart 14: Unfolding of the events following the operational disruption and consequences for the two RTGS systems