Department of Economics
School of Social Sciences

An Analysis of Settlement Risk Contagion in Alternative Securities Settlement Architecture

Giulia Iori[1]
City University

Christophe Deissenberg[2]
Université de la Méditerranée and GREQAM

Department of Economics

Discussion Paper Series

No. 08/03

[1] Department of Economics, City University, Northampton Square, London, EC1V 0HB, UK. *g.iori@city.ac.uk*

[2] Université de la Méditerranée and GREQAM, Château Lafarge, Route des Milles, 13290 Les Milles, France *christophe.deissenberg@univmed.fr*

An Analysis of Settlement Risk Contagion in Alternative Securities Settlement Architectures

Giulia Iori[1] and Christophe Deissenberg[2]

[1] Department of Economics, City University, Northampton Square, London, EC1V 0HB, United Kingdom. giulia.iori@kcl.ac.uk.
[2] Université de la Méditerranée and GREQAM, Château Lafarge, Route des Milles, 13290 Les Milles, France. christophe.deissenberg@univmed.fr.

Summary. This paper compares the so-called gross and net architectures for securities settlement. It studies the settlement risk arising from exogenous operational delays and compares the importance of settlement failures under the two architectures, as a function of the length of the settlement cycle and of different market conditions. Under both architectures, settlement failures are non-monotonically related to the length of settlement cycle. There is no evidence that continuous time settlement provides always higher stability. Gross systems appear to be more stable than net systems.

Key words: Security clearing and settlement, gross and net systems, contagion.

Preamble

Manfred Gilli is a fine man. Not only is he an excellent scientist and a person of striking culture. He also inherited the physical leanness, the moral robustness, the humbleness of the mountaineers of Südtirol, his birthplace. His unobtrusive sense of humor is second to none. Most importantly perhaps, he is a profoundly gentle and kind person.

It is with a profound respect and friendship that we dedicate this paper to Manfred. It has been a pleasure and a privilege to know him for many years. May life bring us together for many more years to come.

1 Introduction

Securities Settlement Systems (SSSs) are institutional arrangements for the confirmation, clearance and settlement of securities trades and for the safekeeping of securities. They involve three steps. The first one, *trade confirmation*, aims at ensuring that the buyer and the seller (typically both banks) agree on the terms of the trade. To that purpose, following a trade each party sends an advisory message identifying the counterpart, the security, the number of shares, the invoice price, and the settlement date. After confirmation comes *clearance*, i.e., the computation of the obligations of the counterparts to make deliveries or to make payments at the settlement date. Finally, *settlement* consists of the operations by which the shares are transferred from the seller to the buyer and the payments from the buyer to the seller.

Settlement systems may operate in one tier or two tiers. In a one-tier securities settlement models all end-investor security accounts are within the Central Security Depository (CSD). This model can be found for example in the Nordic countries, Slovenia and Greece. In a one-tier system participants settle mainly their customers' transactions. In a two-tier system the CSD keeps accounts only for the participating banks (or custodians) and only the inter-participant transactions are booked on these accounts while the end-investor accounts are with the participants/custodians. See Holthausen and Tapking (2007) for an analysis of competition between CSD and custodians. The participanting banks can settle the transactions of the end-investor internally.

The banks face a variety of risks, see Committee on Payment and Settlement Systems (2001b). Among them, there is the risk that creditors do not pay back a loan (*credit risk*) or that settlement is delayed because of shortage of cash and securities (*liquidity risk*). There is the risk that securities are delivered but payment not received, and vice-versa (*principal risk*). Other risks arise from mistakes and deficiencies in information and control (*operational risk*), from the safekeeping of securities by third parties (*custody risk*), and from potential failures of the legal system that supports the rules and procedures of the settlement system (*legal risk*).

A financial or operational problem during the settlement process may make a clearing bank unable to meet its obligations. Default by a bank, in turn, may render other banks unable to meet their own obligations, triggering a chain of defaults within the SSS. In that sense, a SSS is susceptible to *systemic risk.* In addition, SSSs are critical components of the infrastructure of global financial markets. Serious dysfunction at their level have the potential to propagate to other payment systems used by or using the SSS to transfer collateral. Thus, problems in the settlement process may induce systemic risk not only for the SSS but also for the financial system and the economy as a whole. See de Bandt and Hartmann (2000) for a review on systemic risk.

Crucial characteristics of a SSS the timing and modalities of settlement. The transactions can be settled in real time or in batches. In the case of batches, the settlement can be conducted gross or net. Under gross settlement, the clearing house settles the trades in the order they have been inputed in the system by the participants. Real time settlements can only be conducted gross as *Real Time Gross Settlements* (RTGS), where payments are executed continuously via transfers of central bank funds from the account of the paying bank to the account of the receiving bank. Under net settlement, each party delivers at batch time the net amount it sold (or receives the net amount it purchased) since the last batch. Netting is appealing because it results in a very significant reduction of the amount of cash and security that needs to be available to the banks during the batch. However, a failure to settle a trade leads to an *unwind*, i.e., to the deletion of some or all of the trades in which the defaulting bank are involved, and in the re-calculation of the settlement obligations of the other banks. An unwind imposes liquidity pressures and replacement costs on the non-defaulting banks that have traded with the defaulting one. Under these conditions, the system will almost surely fail to settle if one or more of the initially non-defaulting banks proves unable to cover the shortfalls and default. It is then likely that both the securities markets and the payment system will be disrupted. To mitigate this, a partial net settlement is implemented in some markets. In that case, only the transactions that cannot be settled are deleted in such a way as to reduce the overall disruption. Likewise, real markets may implement a hybrid RTGS with partial netting queuing. Such hybrid mechanisms make use of the so called upper bound and lower bound on liquidity needs. The lower bound liquidity is the net amount of sent and received transactions. The upper bound liquidity is the amount of liquidity needed to settle immediately without queuing. The difference between upper bound and lower bound liquidity is the difference in liquidity needs between (1) a RTGS system without queuing; and (2) a deferred net settlement system with queuing until the end of the day. If the RTGS system includes a queuing facility and a netting facility at the end of the day, it can use the lower bound liquidity efficiently during the day and thereby settle some of the transactions with the same amount of liquidity as the net system, but possibly earlier during the day. See, e.g., Leinonen and Soramäki) (1999) and Committee on Payment and Settlement Systems (2001a). The usual wisdom is that reducing as much as possible the delay between trade and settlement improves the system's stability. The rational is the following. The longer is the lag between the date of trade and the date of settlement, the greater is the risk that one of the parties will default on the trade. The greater is also the possibility that the security current price will move away from the contract price, i.e., the greater is the replacement costs risk. Both the default and the replacement costs risks can be lowered by reducing the lag between trade and settlement. Thus, the G30 recommended in 1989 that the final settlement of cash transactions should occur on $T + 3$, i.e., three business days after trade date, and noted that same day settlement, $T + 0$, should be the final goal to minimize counterpart risk and market exposure. See also Leinonen (2003). Similarly, the International *Organization of Securities Commissions* (IOSCO) created, in December 1999, the *Task Force on Securities Settlement Systems*. Amongst others, the Task Force has recommended that

$T + 3$ settlement be retained as a *minimum* standard, but strongly suggests that each market assesses whether a shorter interval than $T + 3$ is needed as a function of the transaction volume, the price volatility, and the financial strength of the banks, among others.

Thus, there is a need to understand how the various type of risks are affected by a shortening of the time between trade and settlement. Consider first the principal risk. It is typically taken care of by the *Delivery versus payment* (DVP) practice that links securities transfers to funds transfers. The settlement of securities transactions on a DVP basis reduces the likelihood that the failure of a participant bank could result in systemic disruptions, but does not eliminate it. If one party fails to deliver, the counterpart still needs to replace the transaction at the current market price. The magnitude of the *replacement cost risk* depends on the volatility of the security price and on the time that elapses between the trade and the settlement dates. As this time becomes shorter, the replacement cost risk becomes less and less important. Moreover, the replacement cost risk has little systemic implication.

The credit and liquidity risks are mitigated in some markets by using a *central counterpart* (CCP) that acts as the buyer to the seller and the seller to the buyer. As previously discussed, most markets have also established central securities depositories that immobilize physical securities and transfer ownership by means of book entries to electronic accounting systems. Because of this mechanism, liquidity is usually not a problem on the security leg of the transaction if short selling is not allowed. The cash leg of the transactions is typically settled through the central bank payment system, as this has the advantage of eliminating the credit risk to the seller. Furthermore, as mentioned earlier, intraday credit is typically available (possibly against provision of collateral) to the banks.

In moving from $T+n$ to $T+0$, which is the current policy target, the liquidity risk becomes particularly important on the payments side because the in- and out-coming flows of payments are not known long in advance by the cash managers. Gridlock may occur if the flow of payments is disrupted because banks are waiting to receive payments before sending them.[3] By contrast, liquidity is not a problem on the securities side because the custodians already have the securities at the time the trade is conducted. Nonetheless, in some markets the rate of settlement falls significantly short of 100% because of human errors or operational problems on the security side. Errors or delays may result from

[3] Angelini (1998) studies RTGS systems under payment flow uncertainty. He shows that uncertainty, together with a costly daylight liquidity, may induce participants to postpone payment. This affects the quality of information available to the counterpart for cash management purposes and may induce higher than socially optimal levels of end-of-day reserve holding.

the incomplete or inaccurate transmission of information or documentation, and from system deficiencies or interruptions. Thus, a move to $T + 0$ and to real time settlement could increase settlement failures on the security side.

The previous discussion hints that there are complex interactions between the diverse characteristics of a SSS that have a number of implications at the bank as well as, arguably more importantly, at the overall settlement level. The Bank of Finland has recently developed a powerful tool for simulating SSSs in encompassing, realistic settings, see e.g. Koponen and Soramäki (1998), Leinonen (2005), Leinonen (2007). Nonetheless, additional insight on specific questions can be gained by investigating highly stylized, simple models. In particular, such simple models may help recognize salient aspects of settlement contagion in Securities Settlement Systems.

The only attempt to study contagion in a simple model of a SSS we are aware of is that of Devriese and Mitchell (2005). The authors focus on liquidity risk on the payment side and study contagion effects triggered by the default of the largest player in a gross settlement system. They show that large settlement failures may occur even if ample liquidity is provided.

By contrast, we investigate in this paper the implications of operational risks on the security side which, as noted above, may become particularly important as one goes towards $T + 0$. We are not interested in the replacement risk which, as noted previously, is not very important if the settlement interval is short. Likewise, we do not consider the credit and liquidity risks, and do not allow short selling. We use numerical simulations to compare, in a one tier framework, the performance of pure net and gross settlement architectures as a function of the length of the settlement batches. Specifically, we study the effects of increasing the number of intraday settlement batches when exogenous random delays affect the settling process. The delays are intrinsic to the system and do not depend on the length of the batches or on the gross/net arrangement. Then, a decrease in the length of the batches increases the likelihood that delays will lead to settlement failures. We compare the implications of these failures under net and gross architectures.

While the results presented here do not amount to an exhaustive study of the behavior of a SSS under alternative architectures they give first, and deep, insights on the forces at work and on the complexity of their interaction. In particular, they make clear that there is no simple monotonic relation between the length of settlement cycles and failures, and that shorter batch lengths do not necessarily improve the performance of a SSS. They suggest that the gross architecture is almost always more stable than the net one.

2 The Basic Framework

The stylized situation we consider is the following. A system of M tier one banks trade S shares of a security among themselves. Short selling is not allowed. Let's T be the length of a trading day. Settlement operates at $T+0$. Trades are registered in a queueing system and settled for by a clearing house in N intraday batches that occur at regular intervals. The N batches define N *settlement cycles* of length $T_n = T/N$. Note that, for N large, the model approximates real time settlement. The number of possible trades over a settlement cycle is $N_T = MT_n$ (M banks can trade on each time step).

Our analysis concentrates on what might happen during a settlement cycle. The settlement cycle is sub-divided in time steps, smaller than T_n, that represent the shortest time necessary for concluding a trade. At any time step $t = 1, ..., N_T$ we randomly select a buyer and a seller. The seller transfers an uniformly distributed random fraction of the shares in its virtual position, that is, of the shares it would have if all previous transactions (since the last settlement batch) had settled properly.

We assume that liquidity is always available. Thus, we consider only one of two possible reasons for default, the incapacity of the selling bank to deliver the shares on time, i.e., the security leg. The other possible cause, the inability of the buying bank to pay, or cash leg, cannot arise. Within this framework we investigate, out of all possible, one particular mechanism that could lead to a default: The existence of an unpredictable delay τ between the conclusion of a trade and the moment the trade is confirmed and cleared. This delay may be due e.g. to human or technical failure. Such a delay is always possible, even if the all actors involved acts diligently and follows best business practice. A delay τ such that the trade cannot be settled within the current cycle generates a *triggering default* as it can trigger a chain of subsequent defaults. Indeed, assume that the trade that defaults was between seller A and buyer B. The buyer B is not yet aware that the trade with A will default and may sell the security to another bank C. Bank C, in turn, may sell the security to another bank D, and so on, possibly generating a chain of defaults. These induced defaults will be called *contagious defaults*.

Note that in our scenario a default results, directly or indirectly, from involuntary causes only. It is not the consequence of a strategic behavior, such as short selling for speculative reasons.

We use simulations to compare the two settlement architectures, net and gross, previously introduced. In the case of gross settlement each trade is settled at batch time in the order it has occurred. If a trade concluded at time t does not settle, the buyer may be unable to settle a trade concluded in some $t' > t$, since the later was made assuming that the trade in t would be properly

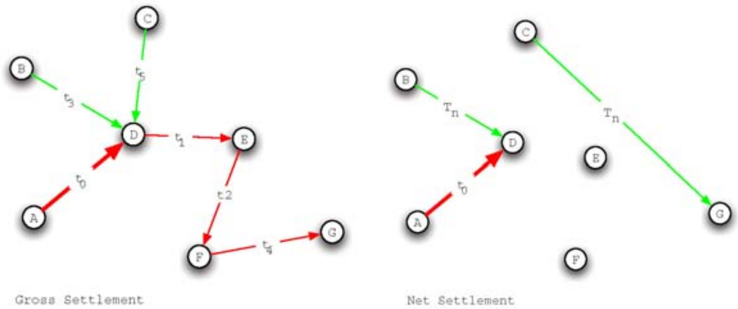

Fig. 1. Comparison of contagious defaults, triggered by an initial operational problem at time t_0, in the gross and net systems.

settled. In the case of net settlement, all trades concluded during the settlement period are settled together at batch time by netting the banks' positions. Only the bank's net position at the end of the settlement period needs to be actually settled. Thus, a default that under the gross system would induce the buyer to default on another trade may not do so under the net system.

The functioning of these two alternative systems, gross on the left and net on the right, is illustrated in Figure 1 with the help of a simple example. The dots represent the banks, the links the trades, with the arrows indicating the trade direction (from the seller to the buyer). The symbols t_j indicate the time when the trade took place, with $t_j < t_k$ for $j < k$. Originally, the banks A, B, and C have one unit of the security and the other banks have none. Under both systems the trade in t_0, plotted in both cases as a thick red link[4] *initially* defaults. In the case of gross settlement, this initial default triggers further defaults touching the trades concluded at times t_1, t_2, and t_4, for a total of four defaults (red arrows). In the netting case, by contrast, the initial default at time t_0 does not propagate through the system. Indeed, thanks to the trades at times t_3 and t_5, the overall netting position of bank D remains positive in spite of the default at time t_0. Only one default, the initial one, occurs.

In the case of net settlement, however, default is followed by an unwind. The unwind causes (a) the deletion of all the trades involving the bank that could not fulfill its commitments to sell, that is, the defaulting bank; and (b) the recalculation of the settlement obligations of the non-defaulting banks. Because the trades of the defaulting bank are deleted, it is possible that other traders will find themselves unable to settle after (b). This may trigger more failures and thus more unwinding. The settlement process is completed only

[4] All figures can be downloaded in color and high definition under http://www.giuliaiori.com/SettlementFigures

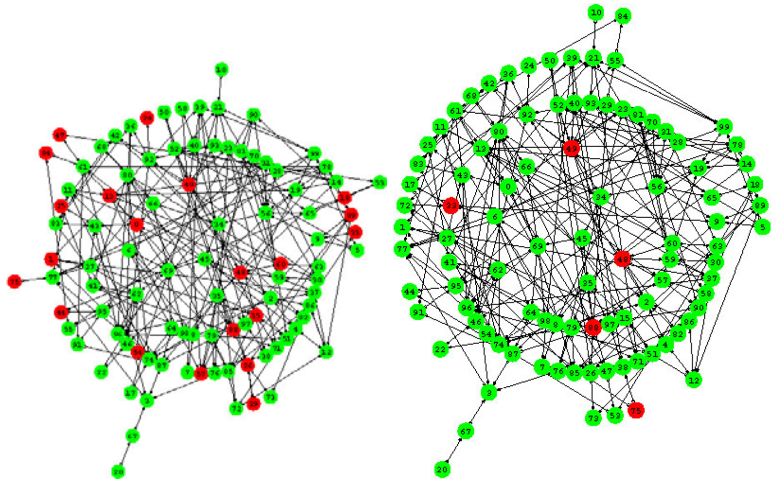

Fig. 2. Defaulting banks under gross arrangements (left) and under netting arrangements (right) as the unwinding process develops.

when all remaining banks, if any, can settle, possibly after many unwinding cycles. The effects of the unwinding process are illustrated in the Figures 2 and 3, where the red dots indicate the banks that *initially* default, the green dots those who can *initially* settle[5]. As could be expected, the number of banks that initially default (Figure 2) is lower under the net than under the gross system. Nonetheless, as the unwinding process takes place (Figure 3) and defaulting banks are removed from the system, more and more banks default under the netting arrangements. *In fine*, the number defaults is lower under the gross system (75 out of 1319) than under the net system (356 out of 1319).

The advantage of the gross settlement system is that it is always possible to identify and delete the exact trade that defaults. The disadvantage is that the order of the trades matter. The advantage of net settlement is that the order of the trades up to the batch time does not matter. This may reduce the initial number of contagious defaults. Nonetheless, via the unwinding mechanism, even trades that could otherwise settle are cancelled. This may generate new rounds of contagious defaults. We study in the next section the interplay between these two effects.

[5] The graphs were generated by simulating trading accordingly to the rules described in section 3.2.

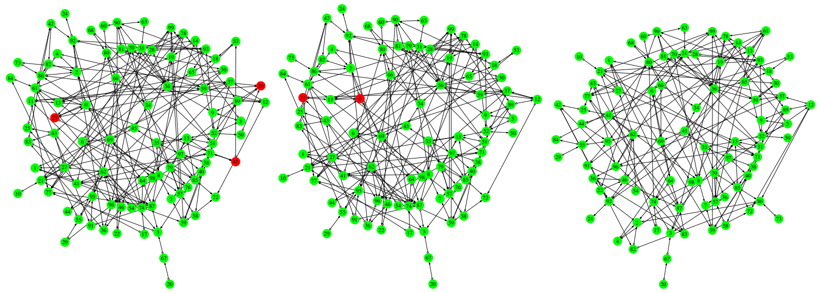

Fig. 3. Defaults in the net system as the unwinding progresses (from left to right). The final graph shows only the trades that can ultimately settle.

3 Numerical analysis

In the numerical analysis presented here, $N = 100$ banks and $S = 1000$ shares. The trading day has length $T = 65536$ time steps. Initially, the S shares are distributed randomly among the N firms. Let $\Pi_i(t)$ be the actual position of bank i at time t, and $\tilde{\Pi}_i(t)$ its virtual position. Initially, the actual and the virtual positions coincide, $\Pi_i(0) = \tilde{\Pi}_i(0)$.

At each time step t, a trade occurs with a probability λ. That is, a high value of λ indicates a very liquid market. The trade is defined in the following way. A bank i sells a integer random number of shares $s_i(t) \sim U\left[1, \tilde{\Pi}_i(t)\right]$ to buyer j. The buyer is randomly chosen. However, the seller is with probability $p \geq 0$ the buyer at time $t-1$. This introduces a structure in the sequence of trades, unless $p = 0$, in which case the choice of the seller is purely random and does not depend on past activity.

Note that $s_i(t) \leq \tilde{\Pi}_i(t)$ does not warrant $s_i(t) \leq \Pi_i(t)$. In other words, the bank sells shares that belongs to its virtual portfolio. However, should a default occur, its actual portfolio may be smaller than the virtual one and it may be unable to settle the trade it realized.

A random delay τ between the conclusion of the trade and its confirmation and clearing occurs with probability $\mu \geq 0$. If $t + \tau > T_n$, the trade is not communicated on time to the clearing house and thus automatically defaults. Otherwise it is communicated and the clearing house will *try* to settle it.

Following a trade, the positions of buyer and seller are updated in the following, self-explanatory way:

1a. If $t + \tau \leq T_n$ and $s_i(t) \leq \Pi_i(t)$,

$$\Pi_i(t) \longrightarrow \Pi_i(t) - s_i(t),\ \tilde{\Pi}_i(t) \longrightarrow \tilde{\Pi}_i(t) - s_i(t),$$
$$\Pi_j(t) \longrightarrow \Pi_j(t) + s_i(t),\ \tilde{\Pi}_j(t) \longrightarrow \tilde{\Pi}_j(t) + s_i(t).$$

1b. If $t + \tau \leq T_n$ and $s_i(t) > \Pi_i(t)$,

$$\Pi_i(t) \longrightarrow \Pi_i(t),\ \tilde{\Pi}_i(t) \longrightarrow \tilde{\Pi}_i(t) - s_i(t),$$
$$\Pi_j(t) \longrightarrow \Pi_j(t),\ \tilde{\Pi}_j(t) \longrightarrow \tilde{\Pi}_j(t) + s_i(t).$$

2. If $t + \tau > T_n$,

$$\Pi_i(t) \longrightarrow \Pi_i(t),\ \tilde{\Pi}_i(t) \longrightarrow \tilde{\Pi}_i(t) - s_i(t),$$
$$\Pi_j(t) \longrightarrow \Pi_j(t),\ \tilde{\Pi}_j(t) \longrightarrow \tilde{\Pi}_j(t) + s_i(t).$$

In the cases 1b. and 2. $\Pi_i(t)$ respectively $\Pi_j(t)$ are not updated because the trade will not take place due to a delay or to an insufficient portfolio.

The trades concluded up to time step t for which $t + \tau \leq T_n$ are stored in a matrix $J(t)$. The element $J_{i,j}(t)$ gives the net position between agent i and j at time t, as it will be known to the clearing house. If $J_{i,j}(t) > 0$ then i is a net buyer from j. If $J_{i,j}(t) < 0$ then i is a net seller to j. Obviously $J_{i,j}(t) = -J_{j,i}(t)$.

We simulate the system, under both gross and net architectures, for different values of N, and average the resulting default rates over 10000 simulations[6].

Gross Settlement: Under gross settlement, the clearing house checks all the trades in order of arrival. If at any t a bank i has committed to sell a security it did not have in its portfolio at that time, i.e. if $\Pi_i(t) < 0$, the corresponding trade cannot be settled. Accordingly, any occurrence of a negative position $\Pi_i(t) < 0$ is a contagious default. In addition, a triggering default occurs every time a trade is not communicated to the clearing house, $t + \tau > T_n$, even if $\Pi_i(t) \geq 0$.

Net Settlement: Under net settlement, the clearing house is only concerned with the final net position $n_i(T_n)$ of trader i at the batch time. Not counting the trades that were not communicated because of delays, the net position of bank i is given by

$$n_i(T_n) := \sum_j J_{i,j}(T_n).$$

If $n_i(T_n)$ is positive (negative), trader i has to receive (transfer) $n_i(T_n)$ stocks. If $n_i(T_n) < 0$ and $\Pi_i(0) > |n_i(T_n)|$ the trade can settle. However, the trade cannot settle if $n_i(T_n) < 0$ and $\Pi_i(0) < |n_i(T_n)|$. There is a contagious default. Banks who cannot settle are eliminated from the system and all their trades are canceled (unwinding mechanism). Accordingly, if bank i e.g. cannot settle, we set $J_{ij} = J_{ji} = 0$ for all j. Settlement is then attempted without the defaulting banks by recalculating the new net position of surviving banks.

[6] The C code for the simulations is available on request.

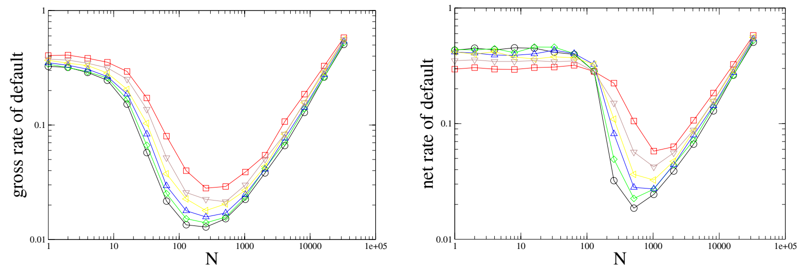

Fig. 4. Default rate r_d as a function of N for $\lambda = 1.0$, $N_a = 100$, $S = 1000$, and $p = 0$ (circle, black), $p = 0.2$ (diamond, green), $p = 0.4$ (triangle up, blue), $p = 0.6$ (triangle left, yellow), $p = 0.8$ (triangle down, gray) $p = 1$ (square, red) under gross (left) and net (right) arrangements.

3.1 Experiments I

In a first batch of experiments, we set the probability μ of a delay equal to 0, construct a single default at the beginning of the simulation, and study the induced contagious effects as trading goes on. We slowly increase the value of p (the probability that the buyer of one period is the seller of the next period), starting with $p = 0$. For each value of p, we compute the average rate of default r_d, defined as the ratio of the number of trades that fail to settle, n_f, over the total number of trades that occurred during the trading day, n_T,

$$r_d := \frac{n_f}{n_T}.$$

Figure 4 shows the outcome of these experiments as p and N increase in the case of gross (left side of the figure) and net settlement (right side). One recognizes that under gross settlement the rate of default increases with p for

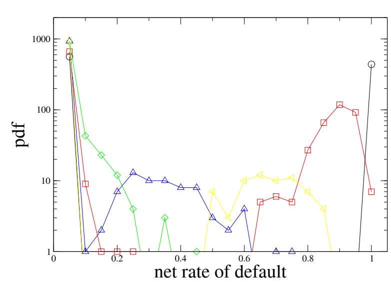

Fig. 5. Probability distribution of defaults in the netting system for $p = 0.2$, $\lambda = 1.0$, $N_a = 100$ and $S = 1000$ and $N = 1$ (circle, black), $N = 128$ (square, red), $N = 1024$ (diamond, green), $N = 8192$ (triangle up, blue), $N = 32768$ (triangle left, yellow).

any value of N. This reflects the fact that a larger p augments the likelihood of long chains of trades. Under the gross architecture, this in turn raises the likelihood of default contagion. Under the net architecture, by contrast, one might think that the trade structure and thus the value of p are irrelevant. Nonetheless, we find that when N is large an increase in p leads to a higher rate of default. Indeed, when N is large, the interval between batches is short so that few trades occur between two batches, leaving little scope for netting. Netting and gross architecture then behave almost identically. Note that for $N = 32768$ only one trade can follow the initial one within a batch, for a total of two trades. In this case the rate of default approaches 50%, as should be expected given that the initial trade defaults by construction. When N is small, a higher level of p implies a lower rate of default. That is, netting seems to perform better when the trading bank acts both as a buyer and a seller – a point that needs further study.

Under both architectures, the rate of default depends non-monotonically on N. It reaches a clear minimum in the range $N = 100$ to 1000. In the net case this can be explained in the following way. Since many trades take place between two batches when N is small, the netting mechanism is very effective. However, if a bank defaults in spite of netting, the number of transactions that are deleted via the unwinding mechanism may be very high. This is illustrated in Figure 5 where the distribution of defaults is plotted for various N given $p = 0.2$. For $N = 1$ (circle, black) the distribution is bimodal with two peaks at $r_d = 0$ and $r_d = 1$. That is, either defaults do not occur or, if they do occur, contagious effects may affect the whole settlement process. This explains the abrupt increase of r_d to about 44% as N decreases below $N = 100$. For N large, on the other hand, netting is not effective anymore. The default rate r_d increases with N since there is always at least one default, the initial one, although the number of trades becomes very small. A similar effect is found in gross systems. In that case, the increase of r_d as N decreases is due to an increase in the length of the chain of trading. For N large, netting and gross arrangements produce very similar results.

The average rate of default is always higher for netting system than for gross systems except for $p \geq 0.8$ and N small. However, this does not tell the whole story. Consider for example the case $N = 1$ and $p = 0.2$. In the netting system, no defaults occur in about 60% of the simulations. But $r_d = 100\%$ in the remaining 40% of the simulations. In the gross system, no defaults occur in only about 17% of the simulations. In more than 80% of the simulations about half of the trades cannot be settled. Hence there is no clear cut ranking of which architecture delivers greater stability.

In Figure 6, we show the transaction networks for $N = 1024$ and different values of p. For small values of p the network is disconnected in several smaller components. At the same time, the rate of default is lower in the gross case. This is in line with the predictions of Allen and Gale (2000) who suggest that systemic risk increases with the network connectivity. Nonetheless the

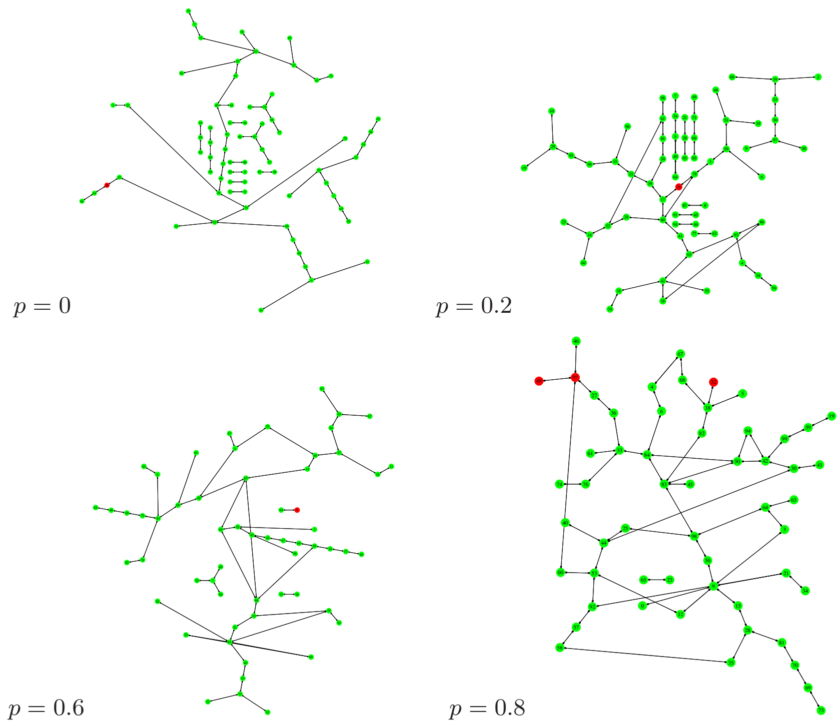

Fig. 6. Transaction network for $\lambda = 1.0$, $N = 1024$, $S = 1000$, $N_a = 100$ and $p = 0$ (top, left), $p = 0.2$ (top, right), $p = 0.6$ (bottom, left), $p = 0.8$ (bottom, right).

network structure is lost when netting multilaterally merges the positions of each bank. In fact, in the netting case we find that contagion decreases with p at low N and increases with p at high N.

3.2 Experiments II

In this second batch of simulations we set the probability p that a buyer becomes seller equal to 0, and assume that a random delay τ can occur with probability μ. We take $\tau = \tau_M \epsilon$ where $\epsilon \sim |N(0,1)|$ and τ_M is a positive constant (qualitatively similar results were obtained with a uniform distribution). A default occurs at time t if $t + \tau > T_n$.

We study the dependence of the default rate on the number N of intraday batches. Reducing the batch length has the advantage of reducing the number of parties exchanging any given security between two settlement cycles, and hence, of lowering contagion. However, a higher N decreases T_n without affecting τ. Thus, a higher N increases the rate of triggering defaults generated by random delays.

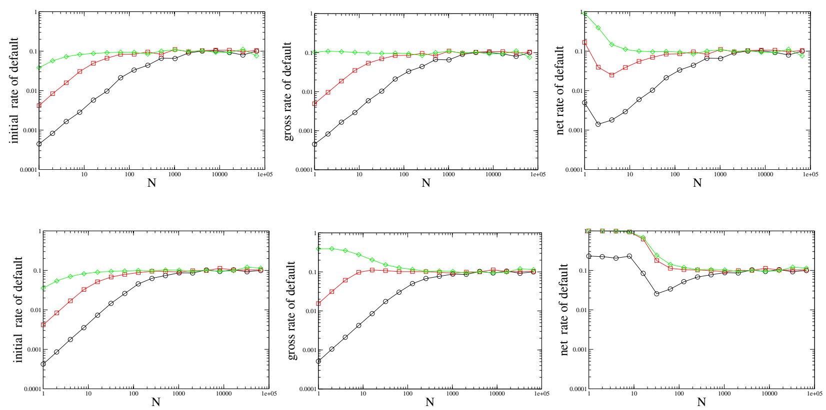

Fig. 7. Initial default rate (left), total default rate in gross systems (center) and total default rate in net systems (right), as a function of N, for $M = 100$ $S = 1000$, $\mu = 0.1$, and for $\tau_M = T$ (diamond, green), $T/10$ (square, red), $T/100$ (circle, black). Top: $\lambda = 0.01$. Bottom: $\lambda = 1$.

We compare the average rate of default r_d in the gross and net cases under different market conditions. Specifically, we let vary λ, which is a proxy for liquidity; μ and τ_M, which measure the likelihood of operational problems. Figure 7 shows the rate of triggering events (defined as the total number of triggering events divided by the total number of trades, left), the rate of default under the gross architecture (center), and the rate of default under the net architecture (right) for $\mu = 0.1$, $M = 100$, $S = 1000$ and $\lambda = 0.01$ (top) or $\lambda = 1$ (bottom). The curves correspond to $\tau_M = T$ (diamond, green), $\tau_M = T/10$ (square, red), and $\tau_M = T/100$ (circle, black). One observes that the number of triggering events increases with N both for $\lambda = 0.01$ and $\lambda = 1$. The *gross* rate of default increases with N when λ is small. For large λ, it decreases with N for $\tau_M = T$, increases with N for $\tau_M = T/100$, and for $\tau_M = T/10$ it is first increasing then decreasing in N with a maximum around $N \approx 16$. The *net* rate of default decreases with N when λ is large. For small λ it decreases with N for $\tau_M = T$, but reaches a minimum and then increases with N for $\tau_M = T/10$ and $\tau_M = T/100$.

Figure 7 reveals that gross systems are little sensitive to contagion when markets are illiquid. In this case the average rate of default is almost identical to the initial rate of default (top, left and center). However, contagious effects can be significative in the gross case when markets are liquid, due to the long chains of trade that can arise in this context (center, bottom). By contrast, in net systems contagion can be very important even when there is limited trading activity.

Further investigation is required to explain this rich variety of behavior.

4 Conclusions

In this paper we examined some issues related to the performance of different securities settlement architectures under the assumption of exogenous random delays in settlement. In particular, we focused on the impacts of the length of settlement cycles on default under different market conditions. Factors such as the market liquidity, the trading volume, the frequency and length of delays, and to some extend the trade structure, were taken into account.

We found that the length of settlement cycles has a non-monotonic effect on failures under both gross and net architectures. This reflects to a large extend the interplay between (a) the stabilization resulting from a decrease in system defaults due to a shorter settlement cycle involving fewer parties; and (b) the destabilization resulting from an increase in triggering defaults due to the greater likelihood that a delay will impair a transfer. We also showed that, contrary to a common wisdom, real time settlement does not improve the performance of the settlement process under all market conditions. Finally, under the scenarios we studied, the gross architecture appears to be more stable than the net one.

The susceptibility of a SSS to contagion depends of many factors in addition to the institutional arrangements governing the exchanges. These factors, that were not taken into account in this paper, include among others the topology of the transaction network and the distribution of portfolio among the banks. The impact of factors of this kind have been in part investigated in the financial literature using techniques from graph theory, for example in the analysis of the interbank market. We are considering to do a similar analysis for SSSs. First steps in this direction would be introduce heterogeneity on the bank size, to simulate a two tier system, and to link the trading activity of each bank to its size. In that way, bigger banks would possibly become hubs of the transaction network, with important implications for contagion.

Another interesting extension would to take into account the possibility of strategic default, i.e., of rational decisions by some participants not to settle in response to movements in securities prices. Doing so may be particularly challenging since, ideally, one should also take into account the behavior of the operator of the SSS. Indeed, the operator typically attempts to discourage strategic default by imposing a fine aiming at taxing away any potential gain from defaulting.

Albeit further investigation is needed to fully understand the complex mechanisms underlying settlement risk in SSSs, we trust that the existing results convincingly show that the approach pursued in this paper sheds light on the behavior of SSSs that could not be gained otherwise, and may potentially help improve their architecture.

Acknowledgements

This paper has crucially benefited from the Lamfalussy Fellowship Program sponsored by the European Central Bank. It expresses the views of the authors and not necessarily represents those of the ECB or the Eurosystem.

The authors are grateful to Mark Bayle, Philipp Hartmann, Cornelia Holthausen, Cyril Monnet, Thorsten Koeppl, Jens Tapking and Saqib Jafarey for valuable comments and interesting discussions. The paper was written within the ESF–Cost Action P10 "Physics of Risk". It benefited from the valuable comments of three anonymous referees. All errors and omissions are the authors sole responsibility.

References

Allen, F. and Gale, D.: 2000, Financila contagion, *Journal of Political Economy* **108**, 1–33.

Angelini, P.: 1998, An analysis of competitive externalities in gross payment systems, *Journal of Banking and Finance* **21**(1), 1–18.

Committee on Payment and Settlement Systems: 2001a, Core principles for systemically important payment systems, CPSS publication No. 43, Basel.

Committee on Payment and Settlement Systems: 2001b, Recommendations for securities settlement systems, CPSS publication No. 46, Basel.

de Bandt, O. and Hartmann, P.: 2000, Systemic risk: A survey, *Working Paper 35*, European Central Bank, Frankfurt.

Devriese, J. and Mitchell, J.: 2005, Liquidity risk in securities settlement, *CEPR Discussion Paper No. 5123* .

Holthausen, C. and Tapking, J.: 2007, Raising rival's costs in the securities settlement industry, *Journal of Financial Internediation* **16**(1), 91–116.

Koponen, R. and Soramäki, K.: 1998, Intraday liquidity needs in a modern interbank payment system – a simulation approach, *Expository Studies 14*, Bank of Finland, Helsinki.

Leinonen, H.: 2003, Restructuring securities systems processing – a blue print proposal for real-time/$t + 0$ processing, *Discussion Papers 7*, Bank of Finland.

Leinonen, H.: 2005, Liquidity, risks and speed in payment and settlement systems – a simulation approach, *Expository Studies 31*, Bank of Finland, Helsinki. Editor.

Leinonen, H.: 2007, Simulation studies of liquidity needs, risks and efficiency in payment networks, *Expository Studies 39*, Bank of Finland, Helsinki. Editor.

Leinonen, H. and Soramäki), K.: 1999, Optimizing liquidity usage and settlement speed in payment systems, *Discussion Paper 16*, Helsinki.