Journal of Financial Market Infrastructures Volume 1/Number 1, Fall 2012 (3–27)

Auto-collateralization as a liquidity-saving mechanism

Søren Korsgaard
Danmarks Nationalbank, Havnegade 5, 1093 Copenhagen K, Denmark;
email: sk@nationalbanken.dk

This paper examines the effects of introducing an auto-collateralization scheme in a simulated securities settlement system. Using artificially generated data, it shows that auto-collateralization can substantially lower liquidity needs without delaying settlement. By allowing for a more efficient usage of collateral, it also reduces intraday credit exposures. The paper further demonstrates how these efficiency improvements in usage of liquidity depend on the structure of the settlement cycle. The benefits of auto-collateralization are particularly large when system participants' incoming and outgoing payments are negatively correlated over time. Finally, the impact of participant defaults is assessed, and it is shown that auto-collateralization protects against liquidity risk in certain stress scenarios.

1 INTRODUCTION

Over time there has been a shift in the structure of payment and securities settlement systems (SSSs) away from designs known as deferred net settlement (DNS) systems to real-time gross settlement (RTGS) systems. In a DNS system payment messages are transmitted continuously, but accumulated over a given period of time. Transactions are settled at periodic intervals such as at the end of a business day, and are netted multilaterally such that each participant has a single payment obligation. This design economizes on participants' liquidity needs. In contrast, transactions are settled immediately and irrevocably upon receipt of payment messages in RTGS systems. This immediacy comes at the cost of an increased need for liquidity.

The proliferation of real-time gross settlement has largely come about at the behest of central banks. In 1985 three central banks had implemented these systems, but by the mid 2000s that figure had risen to ninety (Bech and Hobijn (2007)). One reason for central banks' keenness on RTGS systems is that these eliminate the credit risks that financial institutions otherwise face in DNS systems. The relative costs and benefits of

---

The author would like to thank Marco Massarenti and seminar participants at the De Nederlandsche Bank and University of Tilburg Financial Infrastructure Research Conference 2012 for useful comments and feedback.

gross and net settlement systems are discussed extensively in, for example, Angelini (1998), Kahn *et al* (2003) and Freixas and Parigi (1998).

A problem with gross settlement is that the associated increase in liquidity needs may induce strategic behavior that leads to settlement delays (Bech and Garratt (2003)), defeating the purpose of ensuring immediacy of settlement. Participants will rationally piggyback on others rather than provide costly liquidity themselves. They hesitate to send payments until they have received payment themselves. This results in inefficient equilibriums, a sort of waiting game, in which all transactions are postponed.

In order to counter this tendency, operators of payment and settlement systems have introduced various liquidity-saving mechanisms (LSMs) in the design of payment and settlement systems. An LSM is essentially a design rule intended to "grease" the liquidity flow. Examples of LSMs include queuing facilities, which release payments when some predefined event takes place, and netting algorithms designed to resolve gridlock situations in a payment system. Norman (2010) provides a succinct introduction to LSMs.

This paper examines a particular LSM known as auto-collateralization (sometimes called "self-collateralization"). Auto-collateralization schemes have been implemented in a number of SSSs, for example in the Danish SSS operated by VP Securities. A description of the Danish scheme and its effects on financial stability can be found in Danmarks Nationalbank (2005, 2008).

The defining characteristic of auto-collateralization schemes is that they permit participants in an SSS to use purchased securities immediately as collateral. In order to gain access to liquidity, a requirement is of course that the securities in question are eligible as collateral at the relevant liquidity provider, in practice typically a central bank. The mechanics of auto-collateralization can be illustrated by example. Consider the following. A bank purchases and immediately sells on a security worth 100. Both transactions will be processed in the same settlement cycle, but the bank's purchase is to be settled before the sale. In many SSSs participants need to reserve liquidity for settlement prior to the commencement of the settlement cycle. The bank would therefore have to fund the amount of 100, for example by borrowing the amount against some collateral. Often the cash leg of securities transactions are settled in central banks, which also provide intraday credit for settlement purposes. Thus, the bank might pledge some eligible securities as collateral to the central bank, obtain an intraday loan of 100, and repay the loan by the end of the settlement cycle. Note that the bank would need to hold collateral worth more than 100, as the central bank will also demand a haircut reflecting the quality of the collateral. If auto-collateralization were to be implemented, this situation would play out differently. Assuming that the purchased security is eligible for credit at the central bank, the bank would be able to use the security as collateral for the purchase of that security. The bank's initial liquidity need would therefore drop to an amount equal to the haircut demanded by the central bank, and its collateral requirement would fall correspondingly. Consequently, the effect of auto-collateralization is to economize on collateral.

In the simple example above the bank had a net liquidity need of zero, and it turns out that features such as the degree of netting matter for the efficacy of auto-collateralization. If a participant in an SSS is only a buyer of securities on a given day, the benefits of auto-collateralization are less pronounced than if the participant is both a buyer and a seller. In the first case the participant will still have to fund the entire purchase at the end of the day when repaying the intraday credit. The provision of intraday credit to effect the settlement only delays the timing of the payment, and risk is shifted to the central bank for a period during the day. In the second case the participant may have no net funding requirement at all, but still have a large liquidity need if its securities purchases precede its securities sales. Implementing auto-collateralization reduces this temporary liquidity need substantially.

The above discussion sheds light on two features of auto-collateralization. It reduces immediate liquidity needs, and the benefits of auto-collateralization hinge on the structure of the settlement cycle. It may also offer protection against the liquidity risks that arise in stress scenarios such as when a large participant in settlement system defaults. Though the extent to which this is the case depends on the structure of transactions. The intuition is that while auto-collateralization generally protects against unexpected liquidity risks, those are not necessarily the risks that materialize in typical stress scenarios. To elaborate, if a participant sells securities and expects to receive liquidity, but does not do so because the buyer defaults, then the selling participant can still use those securities to generate liquidity via the auto-collateralization scheme. Thus, auto-collateralization does offer protection against the liquidity risks that arise in a default situation. However, the days when the consequences of settlement failures are greatest are typically those in which large net positions are transacted, which happens to be when certain participants are mainly buyers and others sellers (say, at large bond auctions). In that case the extra liquidity may not be of much use to the sellers since they only have incoming payments anyway.

This paper explores each of the above-mentioned features of auto-collateralization in greater detail using a simulation approach. In recent years a number of studies have investigated the mechanics of payment and SSSs by use of simulations. This development is partly attributable to the development of the BoF–PSS2 simulator by the Bank of Finland (see, for example, Leinonen and Soromäki (2003)). More generally, simulation is a fruitful way of studying payment systems that are characterized by multiple participants and transactions, algorithmic rules, etc, that are difficult to model theoretically. For the same reason, studies that look at the introduction of particular system design rules or LSMs such as splitting of payments and the use of queuing and netting facilities have typically used simulations.

2 LIQUIDITY-SAVING MECHANISMS AND THE PROVISION OF INTRADAY CREDIT

Auto-collateralization can be viewed as an LSM. By allowing participants in an SSS to use purchased securities as collateral, it reduces their total collateral needs or, put differently, improves their access to liquidity for a given amount of collateral held. A number of actual SSSs have implemented auto-collateralization schemes, generally in cooperation with central banks in their respective markets. These include the SSSs operated by VP Securities (Denmark), Euroclear UK and Ireland (UK) and Clearstream Banking Frankfurt (Germany). Auto-collateralization will also be implemented in Target2-Securities.

A number of studies have examined the effects of introducing LSMs. Many of these analyze the effects of LSMs by simulating the flow of actual payment data with and without those LSMs. The use of actual data and simulations, which replicate actual payment or SSSs, holds the advantage of descriptive realism, but such studies are subject to a sort of Lucas critique. That is, if the rules of the games in a system are altered, so the behavior of participants will change.

Some (nonsimulation) studies have also explored the mechanics of payment systems in a theoretical setting in which behavioral or strategic effects are taken into account. An example is Martin and McAndrews (2008), who highlight the welfare-improving effects of properly designed LSMs. They show how LSMs have the potential to improve the functioning of payment systems and welfare, provided that the LSMs offer the right incentives to participants.

Simulation studies and the theoretical studies complement each other. The former showcase the real-life complexity of payment and settlement systems, but cannot easily incorporate behavioral changes, while the converse is true of the latter. Being in the mould of the simulation studies, this paper also potentially suffers from the same problems as laid out in the common critique of simulation studies. However, this critique is likely to have minimal relevance in the case of auto-collateralization for at least two reasons. First, the effect of introducing auto-collateralization is to reduce the opportunity costs associated with holding collateral. This gives an incentive to settle earlier and so only has a positive effect. Second, this study considers an SSS and not a payment system, and the strategic logic in these systems is not the same. In a payment system the payer decides when to submit a payment, but in a securities trade there are two parties who, typically two or three days in advance, must submit matching settlement instructions. Moreover, the SSS may have entirely its own logic (eg, algorithms) for the order in which transactions are settled, so "gaming" the system is more difficult.

The studies most similar to this are those that examine the liquidity effects of some specific intervention in a system. An early example is Leinonen and Soromäki

(1999) who describe the general trade-off between settlement delays and efficient liquidity usage. A later example is Arjani (2006), in which the effect of introducing a complex queue-release algorithm in the Canadian large-value payment system is analyzed. A number of other studies have considered different types of LSMs, for example splitting of payments (Denbee and Norman (2010)) or using a "receipt-reactive queue" (Jackson and Ercevik (2009)).

The relationship between settlement delays and liquidity needs is often depicted in the form of a frontier, such as the stylized frontier depicted in Figure 1 on the next page.

Point A corresponds to an RTGS system in which all transactions are settled immediately upon submission. There are no settlement delays, but it is the structure that requires participants to hold the most liquidity. At the other extreme is a system in which all transactions are netted at the end of the day. Full multilateral netting minimizes liquidity needs, but maximizes settlement delay.

As an LSM, auto-collateralization is worthy of note because its effects are different from those of most other LSMs. An efficient queuing facility, for example, might lower the "inner points" of the frontier, but the end points, A and B, will remain the same. Auto-collateralization is a different story. Strictly speaking, the frontier is unchanged as the same amount of liquidity is needed to settle the transactions. The amount of collateral needed to obtain this liquidity, however, is much smaller, and graphically this can be thought of as a lowering of point A.

To explain this more fully, note that the minimum or net amount of liquidity required by a participant is given by:

$$\min\left(0, \sum_{j=1}^{N} V_j\right) \tag{2.1}$$

where V_j denotes the value of the jth transaction involving that participant, with transactions ordered by their time of occurrence, and V_N is the final transaction of the settlement period.

The above would be the liquidity required in a system with end-of-day multilateral netting. The maximum amount of liquidity needed is given by:

$$\min_k\left(0, \sum_{j=1}^{k} V_j\right) \tag{2.2}$$

Here the k reflects the fact that maximal liquidity need may occur at some time before N. In order to borrow intraday credit corresponding to the amount in (2.2), a participant would need collateral worth:

$$V_{\text{collateral}} = \frac{1}{1-h}\left(-\min_k\left(0, \sum_{j=1}^{k} V_j\right)\right) \tag{2.3}$$

FIGURE 1 Stylized liquidity frontier.

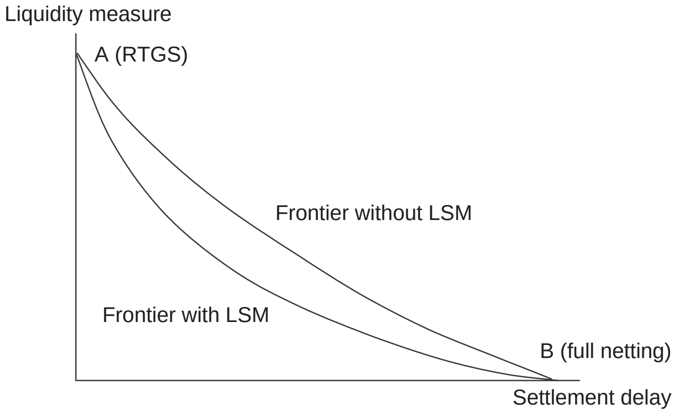

where h denotes the required haircut. (For simplicity of notation, a single haircut for all assets is assumed. In practice, the haircut demanded depends on the characteristics of the collateral posted.) With auto-collateralization, participants can use purchased securities as collateral, and only need to put up initial collateral worth:

$$V'_{\text{collateral}} = \frac{h}{1-h}\left(-\min_k\left(0, \sum_{j=1}^{k} V_j\right)\right) \tag{2.4}$$

At the end of the day the participant will either have a net debit or net credit position. The participant can use incoming flows to repay the intraday credit granted to it, and by the end of the day it will have a cash outflow equal to (2.1). This pattern is exemplified in Figure 2 on the facing page, and further practical illustrations are given in Section 4, where the simulation results are given.

Figure 2 on the facing page illustrates the hypothetical situation of a participant with time-varying liquidity needs. At its lowest, the liquidity position of the participant is -10. In a gross settlement system without auto-collateralization a participant relying entirely on borrowing would therefore have to borrow 10, pledging $10/(1-h)$ as collateral. (In a DNS system the position that needs to be financed is equivalent to the end-of-day liquidity position, which is 4 in this case.) With auto-collateralization, $(1-h)$ multiplied by the liquidity position can be borrowed by using purchased securities as collateral, when the liquidity position is negative. When the liquidity position is at its minimum, an amount of 9 can be borrowed, assuming a haircut of 10%. The liquidity, which the participants need to hold in advance, is therefore only 1, a reduction of $100\% - h$.

FIGURE 2 Hypothetical liquidity pattern.

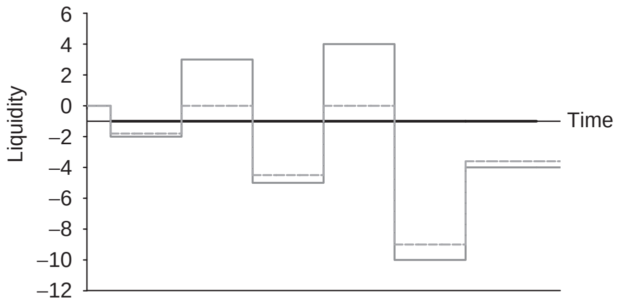

Gray solid line: liquidity position. Black line: initial credit. Gray dashed line: auto-collateralization credit.

This illustrates how auto-collateralization reduces immediate liquidity needs. Since the participant's end-of-day position is a net debit position, it will still have to repay the intraday credit obtained during the day. It is worth noting that this amount is much lower than the amount that would have been extended as intraday credit under gross settlement, thereby also reducing intraday credit exposures. This is only true, though, because of the presence of offsetting transactions during the settlement cycle. This highlights a more general feature of LSMs, namely, that their ability to improve or "grease" the settlement depends on offsetting. A queuing facility, for instance, is of no use if participants either only submit or receive payments. What separates auto-collateralization from many other LSMs is that it does not rely on reordering of transactions (thereby producing settlement delays), but it does add liquidity via more efficient collateral use.

While it is clear that auto-collateralization is beneficial to participants in a settlement system since it reduces their collateral needs,[1] it is less obvious that auto-collateralization also improves settlement efficiency itself since the liquidity frontier remains unchanged. However, the literature on intraday credit, closely related to the literature on LSMs, offers some very interesting insights that suggest that this is indeed the case.

Much of the literature on intraday credit starts from the same observation as that of Bech and Garratt (2003), namely, that participants in payment and settlement systems have an incentive to delay payments if liquidity is costly. Liquidity can have different types of costs. The main direct cost is interest, but most central banks lend

---

[1] Note also that lenders such as central banks are no worse off with auto-collateralization. In fact, their credit exposures can decrease.

intraday at low or zero rates. As Kahn and Roberds (2001) note, a common argument supporting this is that RTGS would otherwise impose undue costs on banks and their customers. Still, the central banks who lend intraday at zero interest only do so against collateral, and holding collateral may be associated with an opportunity cost. Bech and Garratt (2003) show that both uncollateralized, interest-bearing intraday credit and collateralized, interest-free intraday credit can result in suboptimal equilibriums.

In practice, there seems to have been a move toward central banks providing collateralized but interest-free intraday credit. This is described, for example, by Martin (2004), who offers a theoretical analysis to support this move. A concrete reason why this type of liquidity provision may have been favored is that the opportunity costs of holding collateral have been low for financial institutions. Since they have been required to hold quality assets as part of their prudential asset requirements anyway, the opportunity costs of using these assets as collateral for payment or settlement purposes are likely to have been minimal (see, for example, Ball *et al* (2011)). As noted earlier, the effect of auto-collateralization is to further reduce collateral needs and thereby also the opportunity costs of collateral. Since costly collateral is otherwise a cause of settlement delays (due to strategic behavior), it follows that mechanisms that reduce collateral costs also contribute to earlier settlement.

Apart from liquidity savings, auto-collateralization schemes may also entail other, less easily quantifiable, benefits. They can potentially reduce operational risk, for instance. By not having to transfer money to dedicated settlement accounts each day, there is a reduced likelihood of manual errors creeping in. Moreover, auto-collateralization can bring about greater resilience against liquidity shocks. If banks economize on liquidity and reserve just enough liquidity for settlement, they are vulnerable to not receiving liquidity from others. With auto-collateralization this type of settlement risk only occurs if the participant does not have enough collateral.

3 SIMULATION, DATA GENERATION AND MODEL SETUP

Numerous studies on liquidity issues in payment and settlement systems have been produced using simulation techniques and especially the BoF–PSS2 simulator.[2] Many of these studies rely on actual data, mostly from payment systems (more specifically, central-bank-operated RTGS systems). Few simulation studies have been produced on SSSs, presumably reflecting the fact that data from SSSs is harder to come by. This paper relies on artificially generated data for the same reason, but there are other powerful reasons in favor of using artificial data when answering research questions.

---

[2] The Bank of Finland has a website that describes the BoF–PSS2 simulator in greater detail. The site also offers an overview of the publications, which are based on use of simulation tools. See http://pss.bof.fi/Pages/Publications.aspx for details.

Artificial data allows us to easily tweak the structure of the data and thereby isolate effects. Using *ceteris paribus* assumptions is more suitable when we control the data generating process. Moreover, the answers to the questions of interest in simulation studies are often qualitative rather than quantitative, and so it matters little whether the figures obtained are quantitatively realistic. In principle, it would also be possible to circumvent the Lucas critique with artificially generated data. Behavior, in the form of decision-making rules, could be incorporated since one controls the data generating process.

The generated data must satisfy certain characteristics and logic to be used in the BoF–PSS2 simulator. Schulz (2011), who also uses artificially generated data, offers an overview of how a data set must be structured and how data can be generated.

In order to perform the BoF–PSS2 simulations done in this paper, three data sets must be created. One is a list of participants, one a set of credit limits for each participant, and one a set of transactions. The participant data set consists of fifty participants. The amount of liquidity (credit) available to each participant is varied from simulation to simulation (see Section 4 for more details).

The paper uses three different transactions data sets for the various simulation analyses performed. In the following we explain how the first data set is constructed. The second and third data sets are built in a nearly identical fashion, and the differences are explained when the data sets are introduced in Section 4. The first transactions data set consists of 10 000 transactions, which are generated using the following rules.

(1) Each transaction has value 1. The transaction value itself is irrelevant, and there is therefore no reason to vary this parameter. It would be relevant if other LSMs, such as payment splitting, were investigated.

(2) For each transaction there is a buyer and a seller (with the constraint imposed that these must not be identical).

(3) Each participant is characterized by a "size parameter", which follows a log-normal distribution. This is constructed by drawing fifty random numbers from a standard normal distribution and raising them to the power of e.

(4) The probability of a participant being party to a trade is given by its size parameter divided by the sum of all size parameters.

(5) Participants are then assigned as sellers/buyers to transactions by first constructing a cumulative probability distribution and then, for each transaction and buyer/seller, drawing a uniformly distributed random number in the range $[0, 1]$. Participants are then assigned to transactions as buyers or sellers by mapping the random numbers to the cumulative size distribution.

FIGURE 3 Method of analysis.

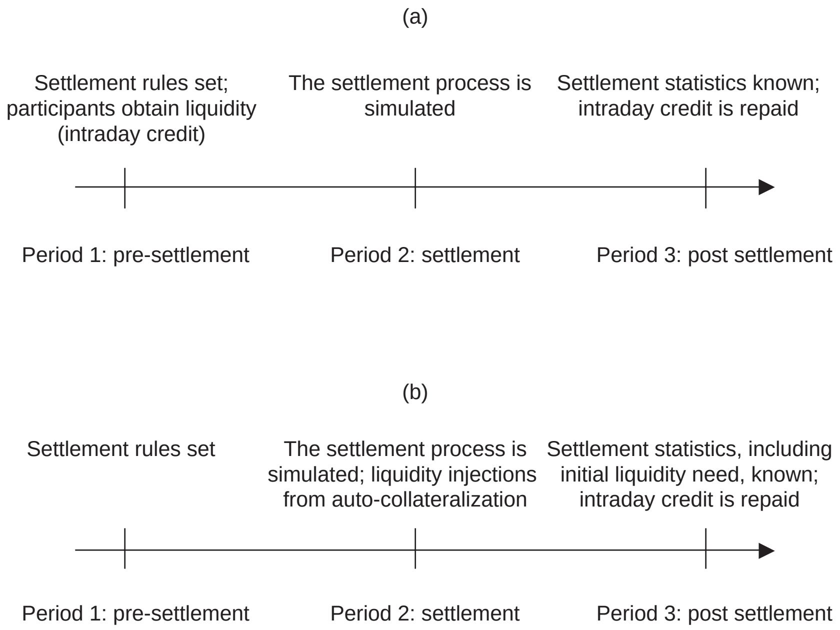

(a) BoF–PSS2 simulator (no auto-collateralization). (b) Simulation with auto-collateralization (author's own calculations).

The analysis proceeds by first running this data in the simulator in different settings: liquidity levels are varied, settlement algorithms are changed, etc. The BoF–PSS2 simulator does not include an algorithm that contains an auto-collateralization feature. If it existed, it would essentially be an algorithm continuously linking the intraday credit limit and transactions data sets. Since it does not exist, however, the analyses using auto-collateralization are done outside the simulator and are then compared with the simulated results. This restricts the analyses somewhat. For instance, it hinders the combined use of auto-collateralization and other LSMs that are built into the simulator. But there are still many issues, which can be satisfactorily addressed nonetheless.

The analyses mostly proceed by comparing the same situation with and without auto-collateralization. The two analyses work in slightly different ways (see Figure 3).

When performing the outside-simulator analyses (including auto-collateralization), the logic is in a sense "backward" as compared with the BoF–PSS2 simulations. First, some benchmark, such as the amount of liquidity required to ensure settlement without any delay, is calculated using the simulator. Subsequently, a question such as the following is asked. If auto-collateralization were implemented, how much initial liquidity would be required to effect settlement and reach that same benchmark?

4 SIMULATION RESULTS

In the absence of an auto-collateralization scheme or other LSMs such as queuing facilities, only a small fraction of transactions can be settled in time when liquidity is scarce. A useful statistic for evaluating the timeliness of settlement is the settlement delay (SD) indicator,[3] which is calculated as follows:

$$\mathrm{SD} = \frac{\sum_{i=1}^{n}\sum_{k=1}^{d} q_{i,k}a_{i,k}}{\sum_{i=1}^{n}\sum_{k=1}^{d} s_{i,k}a_{i,k}} \tag{4.1}$$

where q is the queuing time for each payment, s is the time difference between the submission and end of day (ie, the maximum possible delay), $i_{j,k}$ is the value of payment k of participant i and d is the total number of payments. This indicator essentially measures the proportion of actual delays to possible delays.

Some other interesting measures are the lower and upper liquidity bounds. The lower liquidity bound corresponds to the level of liquidity needed to settle all transactions if they are fully netted. The upper liquidity bound is the level of liquidity required to settle all transactions on time in an RTGS system. For the first of the artificially generated data sets the lower bound is 333 while the upper bound is 606. Nearly twice the amount of liquidity is needed to settle with gross rather than net settlement. This is a fairly small difference, which reflects the purely random nature of the data. In actual payments data, the largest net outflow during a day, ie, the upper bound, is typically a larger multiple of the total net outflow during the day, ie, the lower bound.

How much does settlement performance improve as participants' credit limits (available liquidity) are increased? In the first simulation no LSMs are included, and unsettled transactions are assumed to be settled at the end of the day. There will be enough liquidity for that, given our assumptions. A type of liquidity frontier is generated as follows. For each participant i the lower bound ($l_{\text{lower},i}$) and the upper bound ($l_{\text{upper},i}$) of liquidity are identified, and a number of intermediate cases are simulated and then calculated ($l_i = \{l_{\text{lower},i} + x(l_{\text{upper},i} - l_{\text{lower},i})\}$, where $x = 0.1, 0.2, \ldots, 0.9$). It turns out that adding extra liquidity has very little effect. At $x = 0.1$, only 471 of the 10 000 transactions are settled in time, and this rises gradually to (the still low) 946 at $x = 0.9$, only to jump to all 10 000 transactions when each participant is at

---

[3] The settlement delay indicator is one of several metrics frequently used in simulation studies. For an overview of such metrics, see the technical documentation (specifically the "Databases and files description") available at the BoF–PSS2 website.

FIGURE 4 Settlement delay versus available liquidity.

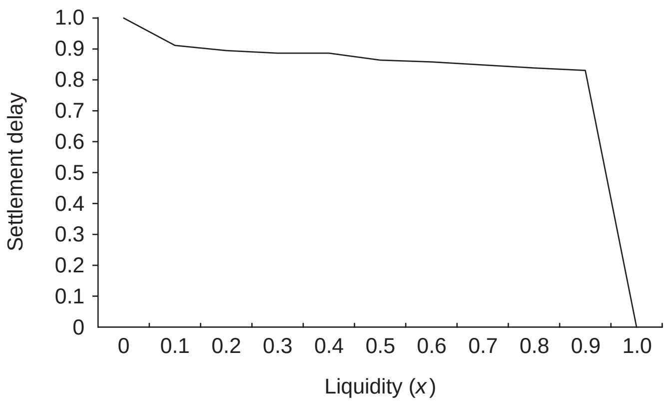

The x on the horizontal axis represents the x in the following calculation for each participants liquidity position $l_i = \{l_{\text{lower},i} + x(l_{\text{upper},i} - l_{\text{lower},i})\}$, $x = 0.1, 0.2, \dots, 0.9$.

the upper bound. This is shown in a slightly different fashion in Figure 4, where the settlement delay is plotted against the amount of liquidity available. The main reason for the diminutive effect of adding more liquidity is the means by which the data has been generated. Due to the randomness of the data, the difference between the lower and upper bonds is relatively small, which results in a small effect of adding small increments of extra liquidity. Later, when more structure is imposed on the data, extra liquidity produces much greater reductions in settlement delays.

Adding some LSM, even a simple first-in–first-out (FIFO) queue, improves settlement performance dramatically. In fact, if a FIFO queue is added, the settlement delay drops to 0.01 even though nearly 3000 transactions are queued at some point during the settlement day. This is an indicative result. The problem facing participants is not so much that they lack a little liquidity; the settlement process largely remains stuck when small increments of liquidity are added (see Figure 4). On the other hand, introducing an LSM such as reordering transactions has a large effect in this case.

The effects of introducing auto-collateralization are more interesting for our purposes. As previously explained, this cannot be done directly in the BoF–PSS2 simulator as there is no algorithm that replicates auto-collateralization. However, some ad hoc analysis of the data set can be performed and then compared with the simulator results.

First, we ask the following question. Given that an auto-collateralization scheme is implemented, how much liquidity is needed to settle all transactions without delay? This can be calculated outside the simulator if a few assumptions are made. It is first worth noting that many of the participants must have held at least a given number of securities prior to the commencement of the settlement (since otherwise they would not have been able to sell them). How many securities? At minimum a number of securities equal to their maximum credit position during the day, since at that point in time the participants must on a net basis have sold that many securities. We call this n_i. Note that these securities could not have been used for liquidity purposes in the absence of auto-collateralization. In the presence of auto-collateralization these are a source of liquidity. The initial liquidity position of participant i is thus given by:

$$\text{liquidity}_i = (1 - h)n_i + l_i \tag{4.2}$$

where l_i is the amount of liquidity from sources other than auto-collateralization and h is the haircut. With each trade both the number of securities and the liquidity position is changed. We are seeking, for each participant, to find the minimum initial l_i that permits uninterrupted settlement. A haircut of 10% is used in the calculations.

Calculating the minimum l_i for each participant and summing over i, we find that this figure is 67.9. This is a low figure when compared with the 606 required to ensure the settlement without delays in a traditional RTGS setup. The figures are equivalent to participants having to hold collateral worth 31 in the auto-collateralization case rather than 673.3 in the ordinary RTGS case. There is also the difference that the build-up of intraday credit is different. In the regular RTGS case participants borrow the required 606 in period 0, whereas the build-up of intraday credit is gradual in the auto-collateralization case. This is illustrated in Figure 5 on the next page.

Another advantage from a lender's perspective, at least if that lender is a central bank that provides intraday credit at zero interest, is that total exposures are much smaller than in the regular RTGS case. In fact, at its highest the exposure is 355, which is not much greater than the lower liquidity bound of 333. By using collateral more efficiently, the intraday credit exposures of the central bank are thus much lower with auto-collateralization than without it. For the participants there is also a direct saving since the "saved" collateral is now free to be used for other purposes.

4.1 Changing the structure of the settlement cycle

In Section 2 it was noted that the benefits of auto-collateralization would depend on the structure of the settlement cycle. Evidently, the benefits in terms of smaller collateral needs are substantial in an RTGS setup, but what if the system nets transactions in batches? The BoF–PSS2 also includes algorithms for netting systems and thereby permits analysis of such cases as well. There are various algorithms and design rules that can be applied in a delayed net settlement (DNS).

FIGURE 5 Credit build-up through settlement period.

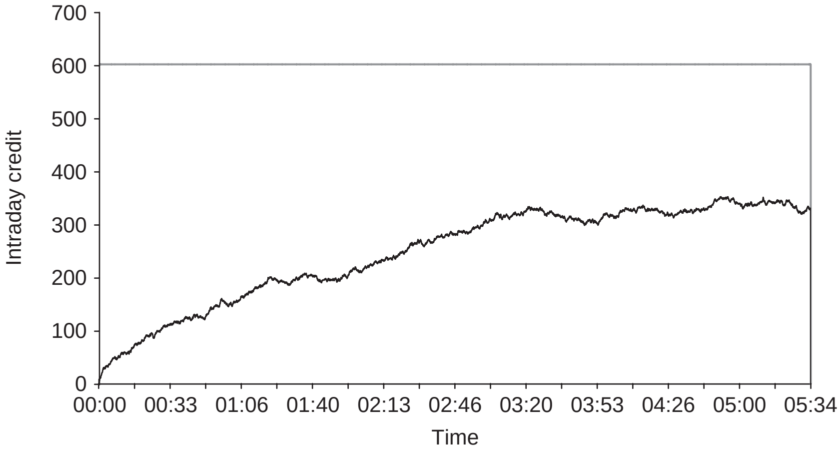

In both cases the outstanding amount of intraday credit is assumed to be repaid at the end of the settlement cycle. Gray line: RTGS. Black line: auto-collateralization.

4.1.1 Multilateral net settlement

If we perform an analysis similar to that in the gross settlement case above, only with multilateral netting batches instead of RTGS, most of the settlement will be forced to take place at the end of the day. In order to do so, the timing of the netting batches must be determined first. The system is initially set up such that it nets transactions every hour and immediately after the final transaction. The multilateral net settlement algorithm of the simulator is of the "all-or-nothing" kind. If just one participant is short of liquidity, all of the transactions are postponed to the following netting cycle. We later relax this assumption by allowing for the possibility of partial settlement.

Recall that the two benchmark liquidity levels are the upper and lower bounds of liquidity. The upper bound is the amount of liquidity required to perform RTGS, while the lower bound is the amount required to settle all transactions at the end of the day. The intermediate cases are liquidity levels in between these two points $[l_i = \{l_{\text{lower},i} + x(l_{\text{upper},i} - l_{\text{lower},i})\}$, where $x = 0.1, 0.2, \dots, 0.9]$. It turns out that nothing can be settled before the end of the day in most intermediate cases. Everything will be settled at the end of the day for liquidity levels ranging from $x = 0.1$ to $x = 0.7$. Some early settlement batches can be completed for $x = 0.8$ and $x = 0.9$ and, with $x = 1$, all batches can be completed with a settlement delay of 0.15.

Further cases are analyzed in Table 1 on the next page, in which three parameters are varied:

(1) the amount of liquidity ($x = 0.1$, $x = 0.5$ and $x = 0.9$);

(2) the use of full versus partial settlement; and

(3) the number of netting batches.

Adding extra liquidity has limited effect on settlement delays and the efficient use of liquidity, except perhaps in the case of six netting cycles and full settlement. The other indicators, the consumed liquidity indicator and the average number of queued transactions tell a similar story. The consumed liquidity indicator is a measure of how much liquidity (in addition to the participant's start-of-delay balances) has been used to effect the settlement. It is lowest in the case of few netting cycles and full settlement since that situation requires the least liquidity. On the other hand, it tends to increase as participants start the settlement cycle with more liquidity. This is somewhat counterintuitive since that would seem to reduce the need for additional liquidity. The reason why the opposite is the case is that, without much liquidity to begin with, essentially no liquidity is consumed because practically all payments end up being postponed for end-of-day settlement anyway.

An intuitive explanation for this is that netting itself can be thought of as an LSM, and thus adding liquidity in a system designed to work with a minimal amount of liquidity is of little use. A greater improvement results from introducing more settlement cycles, especially when combined with partial settlement. Settlement delays drop markedly without reducing efficient liquidity usage much. The same is true of using a partial settlement algorithm rather than the full ("all-or-nothing") algorithm, as settlement delays are much lower where liquidity usage is about the same. The average queue values, a measure of the number of transactions waiting to be settled, tell a similar story.

The performance of an auto-collateralization scheme can also be analyzed in the context of netting schemes. If one introduces auto-collateralization in a setting with six netting cycles, the initial amount of liquidity required to perform with each netting cycle running on time is 24.1. This corresponds to a settlement delay of 0.15. Considering that this cannot even be achieved with $x = 0.9$, ie, nearly the same amount of liquidity required to perform an RTGS settlement with no delay, auto-collateralization is also highly effective in the netting case. In fact, in the Danish case auto-collateralization is only used in systems based on net settlement.

With just two netting cycles, participants' liquidity needs are slightly lower at 22.0, which shows that reducing the number of netting batches only results in minor liquidity savings. Put differently, the liquidity advantages of increasing the degree

TABLE 1 Deferred net settlement simulations.

Setup	Settlement delay [0 to 1]	Consumed liquidity indicator [0 to 1]	Average number of transactions in queue
0.1, full, 6	0.87	0.04	107.16
0.5, full, 6	0.87	0.05	107.16
0.9, full, 6	0.69	0.06	36.05
0.1, partial, 6	0.18	0.03	19.07
0.5, partial, 6	0.17	0.04	18.18
0.9, partial, 6	0.17	0.05	18.06
0.1, full, 2	0.99	0.03	100.61
0.5, full, 2	0.99	0.03	100.61
0.9, full, 2	0.99	0.03	100.61
0.1, partial, 2	0.51	0.03	51.83
0.5, partial, 2	0.50	0.04	51.06
0.9, partial, 2	0.50	0.04	50.80

Setup is x (liquidity), settlement algorithm, number of netting cycles. For details on the computation of the consumed liquidity indicator, see Leinonen and Soromäki (1999).

of netting are much smaller when an auto-collateralization scheme is present, since auto-collateralization already substantially reduces liquidity needs.

4.1.2 Ordered data

Instead of tweaking the setup of the settlement cycle, we can also impose structure by changing the order of the data and making it less random. This is done by letting the probability of a participant being either a buyer or seller depend on the time of day. If the probability of participant being a buyer is p_i for the first 5000 transactions, it is set at $q_i = 1 - p_i$ for the following 5000. To preserve the property that the expected net liquidity position of a participant is zero, thus making the following analysis more comparable to the one above, the probabilities are chosen such that their expected value is 0.5. The p_i are generated as follows. First, fifty numbers ranging from 0 to 1 are drawn from a uniform distribution, and one is assigned to each participant. They are then multiplied by the size parameter (see Section 3 on the construction of the first data set), and an “adjusted” size parameter results. The probability p_i is then calculated by dividing the adjusted size parameter by the sum of all adjusted size parameters. The probability of participant i being a seller is 1 minus the probability of being a buyer.

Imposing such structure on the data can be justified from a practical perspective. Actual SSSs are typically organized such that various events, such as the processing

FIGURE 6 Settlement and liquidity usage.

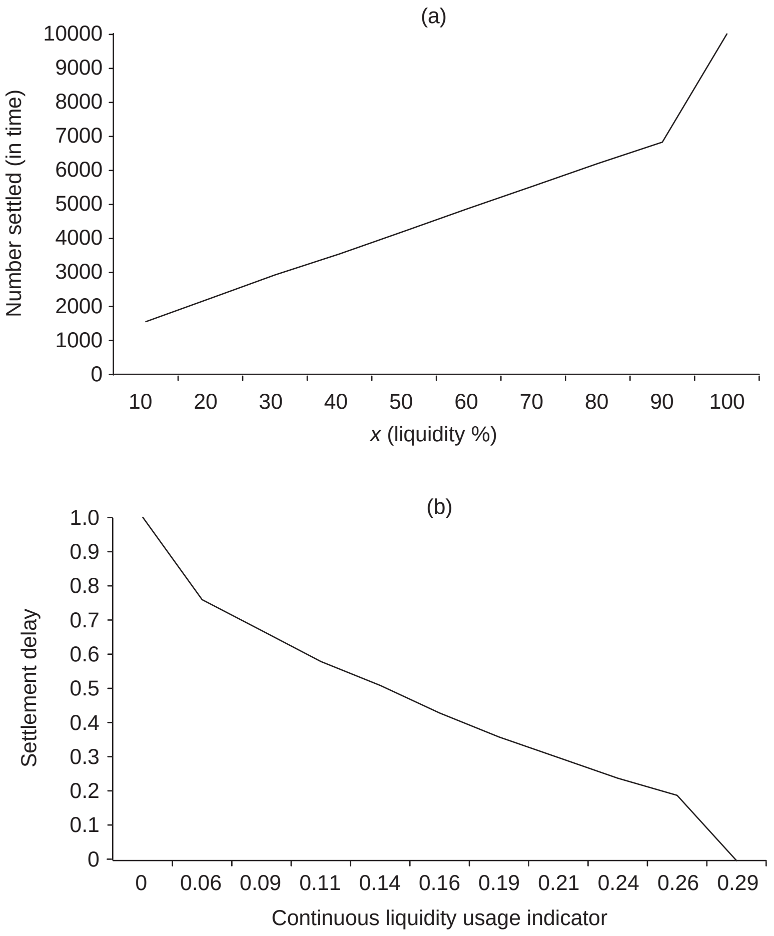

(a) Number of transactions, settled immediately upon submission for various levels of liquidity. (b) Liquidity frontier: the relationship between efficient liquidity usage and settlement delays.

of corporate actions or issuance of new securities and the settlement of securities transactions, take place at different times in the settlement cycle. In such cases it will often be true that some agents receive liquidity at one time and make payments at others. The section on default risk describes a practical example from Danish financial markets.

The new data set is markedly different from the former one, from a liquidity and settlement perspective. The lower bound of liquidity, now 374, has not changed much, but the upper bound has changed to 2859, reflecting the greater liquidity needs induced by about half the participants first being mainly buyers and subsequently sellers of securities. Moreover, adding more liquidity to the system now produces a much greater effect (part (a) of Figure 6 on the preceding page); the same is shown in the liquidity frontier graph (part (b)). Unlike Figure 4 on page 14, these figures are more akin to the stylized liquidity and delay frontiers that one would find in actual payments data.

Again, if one starts from the assumption that participants have just enough liquidity to cover their net settlement needs, adding simple LSMs such as a FIFO queue improves the settlement performance since there is enough liquidity in the system as a whole. However, the effect is much smaller in the case of the current data set as the settlement delay drops to 0.23 against 0.01 using the earlier, unstructured data, in both cases with $x = 0.1$. The intuition behind this change is that LSMs typically work best when there is enough liquidity, but the transactions simply need to be reordered slightly or perhaps netted. This was the case when the transactions were purely random; with sets of participants systematically buying and selling at different points in time, simply reordering transactions slightly will not suffice. In this case more liquidity has a greater effect since precisely what is needed to secure immediate settlement is more liquidity.

Since auto-collateralization provides more liquidity for a given level of collateral, we should expect it to be a particularly useful LSM in the instance where buyers and sellers tend to be displaced in time. Indeed, the initial liquidity needs of participants can be calculated to be 174.5, implying that participants collectively need collateral worth 193.9 if they are to lend the whole amount to be settled. This is to be compared with a need for collateral of 3176.7 in the RTGS case where everything is settled without delays.

It is also instructive to compare the borrowing patterns in this case and in the former case. Now, credit usage peaks exactly midway through the settlement cycle as this is the point where those more likely to be buyers in the meantime become more likely to be sellers. In effect, the netting starts kicking in at that point. This means that the participants who were more likely to be buyers in the first period can start repaying their auto-collateralized intraday credit as time proceeds, and they begin receiving funds from securities sales.

From the lender’s perspective, the advantages of auto-collateralization are not as great in terms of reducing the amount of intraday credit granted, as in the case with purely random data (see Figure 5 on page 16), at least, not when the provision of intraday credit peaks in the middle of the day. However, auto-collateralization still reduces lending needs and thereby credit exposures substantially, since the participants only borrow the amount they actually need to effect settlement and not some larger amount due to, say, precautionary motives.

FIGURE 7 Credit build-up throughout settlement period (structured data).

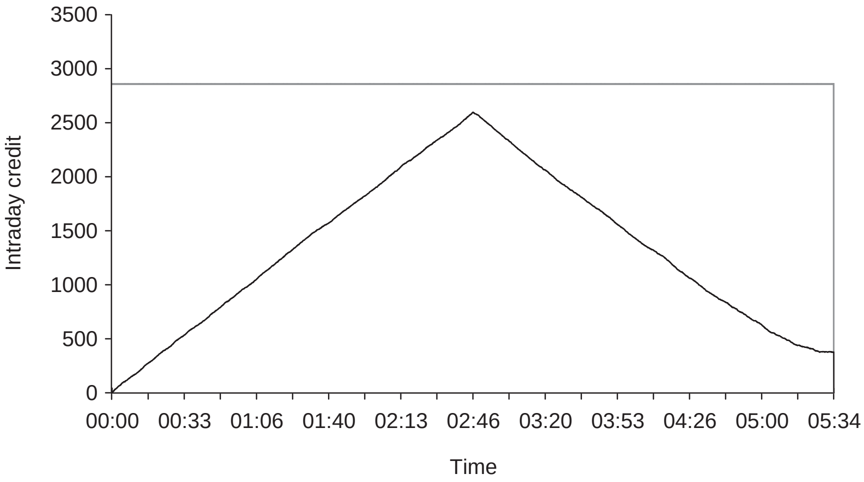

Gray line: RTGS. Black line: auto-collateralization.

4.2 Default risk

According to the international standards for financial market infrastructures,[4] the operators of SSSs must evaluate the possibility of failures, and consider means to ensure continued settlement even if, for example, participants with large net debit positions fail. An interesting question is whether auto-collateralization schemes offer added protection in such scenarios. As a general rule the answer to that question ought to be in the positive. Consider the situation in which a participant in an SSS expects to sell securities, but its counterparty defaults. Without auto-collateralization this would result in a negative liquidity shock of an amount equal to the value of the securities, which we denote v. With auto-collateralization the securities can immediately be mobilized as collateral, and the seller can borrow $v(1-h)$, reducing the magnitude of the liquidity shock.

In practice, though, the days when stress scenarios are likely to have the greater consequences are those where large net transfers are made from some participants to others. In such situations it is less clear that auto-collateralization is necessarily helpful. To illustrate this, assume that participants are either buyers or sellers. If a large buyer defaults, a number of sellers will receive less liquidity, but their liquidity

---

[4] Examples of such standards include the new CPSS/IOSCO principles for financial market infrastructures as well as earlier standards such as the ESCB/CESR-recommendations for SSSs.

FIGURE 8 Liquidity exchanged in the Danish SSS.

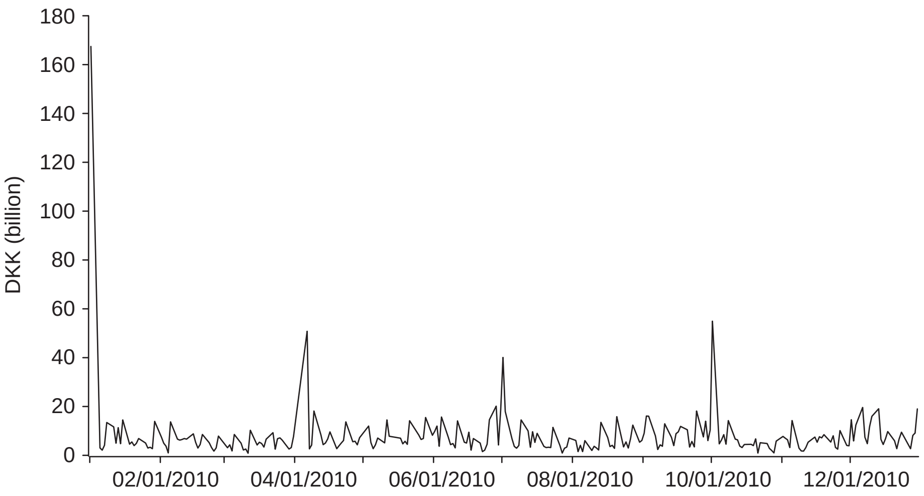

The figure shows the daily amount of liquidity exchanged in the nightly settlement batches in the SSS operated by VP Securities.

position is positive in any case, and access to more credit via auto-collateralization will not be used. It is, of course, likely that the sellers had other plans to dispose of the liquidity they expected to receive, but they will need to get hold of that extra liquidity whether auto-collateralization is present or not.

A concrete example from the Danish settlement infrastructure sheds light on an example of the above. Each year, by far the largest settlement occasion in Denmark is the first banking day of January. As shown in Figure 8, the amount of liquidity exchanged in the Danish SSS is much greater at year-end than on any other day during the year. This is due to the fact that many Danish thirty-year mortgage loans are financed by annual issuances of new mortgage bonds. The settlement of the sale of these new bonds takes place on that day, which means that large net sums are transacted with banks buying the mortgage bonds while the mortgage issuers are the net recipients of liquidity. Moreover, at about the same time as new mortgages bonds are issued, mortgage bonds issued in previous years expire.

In practice, settlement takes place in separate batches that occur at different times of the day. During the night time the newly issued mortgage bonds are transferred to the buyers. This means that banks must make a substantial payment to mortgage issuers at that point in time. However, the existing bonds are only redeemed in the morning, at which time the banks will receive a substantial amount of liquidity. The upshot is that there is a short window of time between the issuance of new bonds and the redemption of old bonds in which banks face a large liquidity need. Using the auto-collateralization scheme, they can cover a large part of this need by immediately using the newly purchased mortgage bonds as collateral for the very purchase of those bonds.

How does this relate to default risk? Suppose a large bank were for some reason unable to participate in the nightly settlement. Then the mortgage issuers would be receiving less liquidity than expected, but this would not immediately affect them as they are nonetheless still net recipients of liquidity. It would, of course, cause difficulties with the redemption of existing bonds. The mortgage issuers could obtain the needed liquidity to repay existing bondholders by using the auto-collateralization scheme, but they might as well use those bonds as collateral via traditional pledge schemes. Indeed, they have to do so if they needed liquidity for an extended period, as auto-collateralization can only be used for intraday credit.

The question of default risk can also be addressed using the simulator approach. Continuing with the existing data set, we now examine the effects of removing either the participant or the two participants with the largest net debit positions from the data sets. By removing all transactions involving the participant with the largest net debit position, the number of transactions drops to 9476 (from 10 000). If two participants are removed, the number of transactions drops to 7738.

In order to calculate the extra liquidity needed, these two transaction data sets are examined using the simulator, and the upper and lower liquidity bounds are recalculated for each participant. The extra liquidity need or the liquidity shock can then be calculated as follows:

$$\mathrm{LS1} = \sum_{i \varepsilon I} \max(0, \mathrm{LB}_{i,\mathrm{shock}} - \mathrm{LB}_i) \tag{4.3}$$

where i denotes the participants that belong to the set of participants that have not been removed, I. LB denotes the lower bound of liquidity. The calculation is the same for the upper bound.

The max function is used to reflect the fact that a positive liquidity shock, ie, a participant unexpectedly getting a smaller liquidity need, does not impede that participant's ability to settle due to insufficient liquidity.

4.2.1 Measuring liquidity shocks

In order to gauge the effect of a liquidity shock, one must construct a measure of that shock. A reasonable candidate is to compare the size of (4.3) with the total liquidity need in the nonstress situation (for comparability, the liquidity needs of the removed participants are excluded from this equation). This measure of liquidity shock can be

TABLE 2 Stress scenarios.

	S1	**S2**	**S3**	**S4**
Number of remaining trades	9476	7738	8431	7276
Total liquidity removed from system	40	75	1019	1888
LS1 (lower bound)	55	94	144	294
LS1 (upper bound)	32	113	156	307
LS2 (lower bound)	0.165	0.314	0.038	0.101
LS2 (upper bound)	0.012	0.052	0.041	0.104

S1: stress scenario in which the participant with the largest net debit position is removed (data set 2). S2: stress scenario in which the participants with the two largest net debit positions are removed (data set 2). S3: stress scenario in which the participant with the largest net debit position is removed (data set 3). S4: stress scenario in which the participants with the two largest net debit positions are removed (data set 3).

computed as:

$$\mathrm{LS2} = \frac{\mathrm{LS1}}{\sum_{i \varepsilon I} \mathrm{LB}_i} \tag{4.4}$$

In order to show that liquidity risks can be lower, even if more liquidity is removed from the system, a third transactions data set is constructed. The method of construction is identical to that used to create the second data set, except that no buyer–seller split is made at 5000 transactions. That is, some participants are more likely to be sellers and others more likely to be buyers throughout the entire range of transactions. The upshot is that LSMs such as netting have limited effect, and the upper and lower bounds are very close to each other. The results of these exercises are reported in Table 2.

Table 2 compares a total of four scenarios. In the first two, S1 and S2, the ordered data is used, ie, the data in which some participants are first more likely to be sellers, and subsequently more likely to be buyers, or vice versa, and in the last two, S3 and S4, the data for which participants are more likely to be either buyers or sellers throughout the entire data set is used. The difference between S1 and S2 is that the default of the two largest participants is simulated in S2 compared with just the largest in S1. The same difference applies to S3 and S4.

Table 2 highlights a number of interesting facts. The amount of liquidity removed from the system is much greater when the data set is structured in such a way that there is limited netting, as in scenarios 3 and 4. Yet the liquidity shocks are not correspondingly greater. The absolute increase in liquidity needs (LS1) is comparably small, and in relative terms the shock is greater in scenarios 1 and 2, at least for the lower bound of liquidity.

The question is what this reveals about the utility of auto-collateralization in stress scenarios. It has been argued that auto-collateralization protects against liquidity shocks, and liquidity shocks tend to be greater in situations where participants have both incoming and outgoing payments. One way of thinking about this is to consider the following. If the same amount of liquidity were removed in scenarios 3 and 4 as in scenarios 1 and 2, the liquidity shocks would be much greater in scenarios 1 and 2. It follows that having auto-collateralization would also be more valuable in scenarios 1 and 2.

A corollary is that auto-collateralization may offer relatively little protection against liquidity risk in the cases envisaged in international standards. Large net debit positions occur exactly when some participants are “systematic” buyers and others “systematic” sellers of securities, days such as year-end in Denmark. The benefits of auto-collateralization might, in fact, be greater at times when net debit positions are near zero while there are large but off-setting gross transactions. That, of course, does not imply that an auto-collateralization scheme does not protect against liquidity shocks. It does. It is simply an observation that days with particularly large liquidity exchanges between participants are likely to be days on which there is little scope for netting and therefore little scope for LSMs to work efficiently.

5 CONCLUDING REMARKS

This paper has argued that auto-collateralization is a useful LSM. While auto-collateralization does not strictly reduce the amount of liquidity needed to complete a settlement, it economizes on the collateral (typically) required to obtain liquidity, and thereby increases access to, and reduces the cost of, that liquidity. And by utilising collateral more flexibly, settlement can be completed with less delay and less liquidity and intraday credit.

The benefits of auto-collateralization depend on the design of the settlement system. More liquidity is needed with gross than net settlement, and so the advantages of reducing liquidity needs via auto-collateralization are more pronounced in the case of gross settlement. Still, auto-collateralization is a complement and not a substitute for other LSMs such as netting, and will be effective in their presence. Some of the cases explored in this paper even suggest that auto-collateralization may be a much more effective mechanism than, for example, netting. Auto-collateralization can also offer protection against liquidity shocks, though it may not be a particularly effective means of protection in the cases when some constituency either tends to act as buyer or seller.

In general, it seems that implementing auto-collateralization is an efficient means of improving settlement efficiency at low cost. This appears to be true for LSMs in general. In a study of the benefits of implementing an LSM in Fedwire, Atalay *et al* (2010) find that the total gains are likely to be substantial compared with the costs. Are there then no disadvantages to introducing LSMs such as auto-collateralization?

Apart from the direct costs of implementing LSMs, which are likely to be very low, there may be some less obvious side effects. As an example, consider the case of the Danish SSS in which settlement risk is highly concentrated at year-end (see Figure 8 on page 22). This concentration of risk has been criticized by, among others, the Danish central bank as a possible financial stability risk but, in a way, auto-collateralization has helped make it possible since banks, in the absence of auto-collateralization, would experience dramatically increased liquidity needs that might render such a concentration much less attractive. To be sure, this is not a particularly persuasive argument against auto-collateralization since, for example, spreading the auctions and settlement of mortgage bonds while preserving auto-collateralization would clearly be a first-best option.

Moreover, there are other tendencies that strengthen the need for LSMs. As described by Ball *et al* (2011), regulators are increasingly concerned that assets included in and intended to be part of banks' prudential asset buffers are tied up as collateral in payment and settlement systems. If banks were required to clearly separate assets for those purposes, the opportunity cost of holding collateral and thereby the incentive to delay payment would increase. This raises the value of having countervailing measures such as LSMs, which ensure a smooth settlement while reducing liquidity needs.

REFERENCES

Angelini, P. (1998). An analysis of competitive externalities in gross settlement systems. *Journal of Banking and Finance* **22**(1), 1–18.

Arjani, N. (2006). Examining the trade-off between settlement delay and intraday liquidity in Canada's LVTS: a simulation approach. Working Paper, Bank of Canada.

Atalay, E., Martin A., and McAndrews, J. (2010). Quantifying the benefits of a liquidity-saving mechanism. Federal Reserve Bank of New York Staff Reports 447.

Ball, A., Denbee, E., Manning M., and Wetherilt, A. (2011). Intraday liquidity: risk and regulation. Bank of England Financial Stability Paper 11 (June).

Bech, M. L., and Garratt, R. (2003). The intraday liquidity management game. *International Journal of Central Banking* **3**(3), 147–181.

Bech, M. L., and Hobijn, B. (2007). Technology diffusion within central banking: the case of real-time gross settlement. *International Journal of Central Banking* **3**(3), 147–181.

Danmarks Nationalbank (2005). Payment systems in Denmark. URL: www.nationalbanken.dk/DNUK/Publications.nsf/side/1E8FEC8F259E61FFC125706C003D4409/$file/payment_systems05.pdf.

Danmarks Nationalbank (2008). Financial stability, 2nd half 2008. URL: www.nationalbanken.dk/C1256BE9004F6416/side/9D5E7E2246BE3E5FC1257535003FA1E1/$file/fin_stab_uk_2-2008_web.pdf.

Denbee, E., and Norman, B. (2010). The impact of payment splitting requirements in RTGS. Working Paper, Bank of England (October).

Freixas, X., and Parigi, B. (1998). Contagion and efficiency in gross and net interbank payments systems. *Journal of Financial Intermediation* **7**, 3–31.

Jackson, J., and Ercevik, K. (2009). Simulating the impact of hybrid design on the efficiency of large-value payment systems. In *Simulation Analyses and Stress Testing of Payment Networks: Proceedings from the Bank of Finland Payment and Settlement System Seminars 2007–2008*, Leinonen, H. (ed), pp. 189–228. Bank of Finland, Helsinki.

Kahn, C. M., and Roberds, W. (2001). Real-time gross settlement and the costs of immediacy. *Journal of Monetary Economics* **47**(6), 299–319.

Kahn, C. M., McAndrews J., and Roberds, W. (2003). Settlement risk under gross and net settlement. *Journal of Money, Credit and Banking* **35**(4), 591–608.

Leinonen, H., and Soromäki, K. (1999). Optimizing liquidity usage and settlement speed in payment systems. Discussion Paper, Bank of Finland.

Leinonen, H., and Soromäki, K. (2003). Simulating interbank payment and securities settlement mechanisms with the BoF–PSS2 simulator. Discussion Paper, Bank of Finland.

Martin, A. (2004). Optimal pricing of intraday credit. *Journal of Monetary Economics* **51**(6), 401–464.

Martin, A., and McAndrews, J. (2008). Liquidity saving mechanisms. *Journal of Monetary Economics* **55**(3), 554–567.

Norman, B. (2010). Liquidity saving in real-time gross settlement systems: an overview. Financial Stability Paper 7, Bank of England (May).

Schulz, C. (2011). Liquidity requirements and payment delays: participant type dependent preferences. Working Paper, European Central Bank (February).