Liquidity-Saving Mechanisms: How they work, and their effect on Continuous Linked Settlement

by

Daniele Costanzo

Written under the supervision of:

Professor Thorsten Koeppl

An essay submitted to the Department of Economics
in partial fulfillment of the requirements
for the degree of Master of Arts

Queen's University
Kingston, Ontario, Canada

August 2012

Acknowledgements

This essay was completed with the help and guidance of several individuals who, in some way, extended their valuable insight and assistance. Among others, I am grateful to:

First and foremost, Professor Thorsten Koeppl for his patient and unfailing support as my supervisor;

Paul Miller and James Chapman at The Bank of Canada for inviting me to the Economics of Payments VI conference in Ottawa, Ontario, Canada;

My friends and classmates at Queen's University, especially Maxime Leboeuf and Nick Andexer, for helpful conversation and collaboration throughout the research and writing process, and throughout the year;

Rachel Ferreira of the Queen's Economics Department for always having the answers and making life easier;

Lastly, my family and girlfriend, Lisa Bianco, for their unrelenting urging and for listening to me talk about a topic that is utterly uninteresting to them.

Abstract

The issue of how to make real-time gross settlement (RTGS) more efficient has been of much interest to economists. Many systems have adopted liquidity-saving mechanisms (LSMs) to attempt to reduce the burden of liquidity costs by decreasing the amount of liquidity required to settle transactions. The goal of this paper is to act as a point of reference to familiarize one's self with the concepts of settlement risk and liquidity-saving. It aims to do so by providing a source that describes how these mechanisms work, using examples of some that are currently employed by different systems. Furthermore, it provides a survey of some of the more influential research in the area of payments economics which focuses on liquidity management. Lastly, it applies some of the main findings of these works to a relatively new system, continuous linked settlement (CLS), which focuses on the mitigation of foreign exchange settlement risk. I discuss how CLS uses LSMs to mitigate the burden of liquidity in the area of foreign exchange settlement, and how the use of LSMs by large value payment systems can affect settlement in CLS. The literature covered in this paper is unanimous in one contention: Liquidity-saving mechanisms are not always welfare-improving. Indeed, the effectiveness of a LSM heavily depends on the surrounding financial system framework.

Contents

1 Introduction

A great deal of attention has been directed toward the design and operation of large-value payment systems (LVPSs), and, as Martin and McAndrews (2008) point out: "...designs known as 'real-time gross settlement' (RTGS) systems have proliferated widely in the last 15 years." RTGS services include settlement account services and settlement credit services, the settlement being an "ultimate" settlement in central bank money with finality. In this sense, RTGS services are a safety net as they provide certainty that those participants that expect payments from others whose creditworthiness is questionable will receive the payments from the central bank. In RTGS, funds are transferred between participants on a real-time and gross basis. While default risk is null, there are high liquidity costs and levels of payment postponement associated with RTGS. Because funds are transferred on a gross and real-time basis, participants must hold large amounts of liquidity, leading to high interest and opportunity costs.

At the other extreme of payments systems lies uncollateralized deferred net settlement (DNS). In this environment, funds are transferred with a delay and gross payments are netted against each other - only net balances have to be settled. Of course, liquidity costs in this system are much lower, as liquidity is only required at the end of the day after netting has taken place. Because transactions can continue to be made throughout the day without liquidity on hand, however, the probability of default and contagion is much higher under DNS. This is because there is no guarantee that those funds will be available. Clearly, then, LVPSs face a tradeoff between liquidity costs and default risk. As Chiu and Lai (Bank of Canada FSR), referring to RTGS and DNS, succinctly put it: "...the optimal design of settlement systems requires joint consideration of these policy instruments."

And so, the issue of how to make RTGS more efficient has been of much interest to economists. Many systems have adopted what are called liquidity-saving mechanisms (LSMs) to attempt to reduce the burden of liquidity costs by decreasing the amount of liquidity required to settle transactions. The goal of this paper is to be used as a point of reference to familiarize one's self with the concepts of settlement risk and liquidity-saving. It aims to do so by describing how these mechanisms work, using examples of some that are currently employed by different systems. Furthermore, it provides a survey of some of the more influential research in the area of payments economics which focuses on liquidity management. Lastly, it applies some of the main findings of these works to a relatively new system, continuous linked settlement (CLS), which focuses on the mitigation of foreign exchange settlement risk, a topic of increasing importance in the age of globalisation. I also discuss how CLS uses LSMs to mitigate the burden of liquidity in the area of foreign exchange settlement, and how the use of LSMs by large value payment systems affect settlement in CLS.

The literature covered in this paper is unanimous in one contention: Liquidity-saving mechanisms are not always welfare-improving. Indeed, the effectiveness of a LSM depends heavily on the financial system framework. For instance, certain LSMs are more effective when delay costs are very high, but welfare-worsening with lower delay costs. Furthermore, the impact of LSMs may not be isolated to the system in which they are employed. Instead, it is conceivable that the success or failure of a LSM can have far-reaching consequences for the international financial system. CLS, for example, can be affected both positively and negatively by certain systems' LSMs. In this sense, the success of CLS in removing foreign exchange settlement risk may depend, at least in part, on the LSMs used by CLS and associated systems.

2 Background: The Basic Tradeoff in a LVPS

In order to properly understand liquidity-saving mechanisms, it is necessary to first know of the risks and costs that banks in a LVPS face in their everyday transactions. The main tradeoff in payment systems is between the costs associated with the function of the system and the amount of risk it allows. As Berger *et al* (1996) note, basically any form of payment involves some degree of risk to the parties involved in the transaction. In some form or another, whether explicitly or implicitly, debt is involved. Multiple parties may, at one point or another, handle the payment instruments to a transaction, and therefore, failure of any of these parties to perform their function prior to the funds being completely delivered from a payer to a payee could result in the break-down of the transaction.[1] Credit risk and liquidity risk are two important drivers of such settlement risk. The former is the risk that the amount agreed upon in a transaction is not transferred to the payee on time because the payor (or its bank or other intermediary) does not have sufficient funds. According to Berger *et al*, the latter refers to the situation in which funds cannot be transferred on time due to (1) the illiquidity of assets, (2) a market breakdown in which normally liquid assets are illiquid and cannot be sold, or (3) a transmission problem in the payments system (eg. power or computer outage) preventing funds transfer. This last liquidity risk driver is what I henceforth refer to as operational risk.

2.1 Credit Risk

A payment system may face credit risk from its payment and settlement processes, or both.[2] This risk is largely driven by current exposures from extending intraday

[1] See Berger, Hancock and Marquardt (1996)
[2] CPSS (2012)

credit to participants. In this sense, central banks that operate payment systems and provide intraday liquidity face this exposure. Of course, collateral can be required in order to extend credit; however, this does not guarantee the removal of credit risk. For instance, a system could face potential future exposure if the value of collateral posted by a participant fell below the amount of credit extended to it, thus leaving a residual exposure.[3]

Systems that employ deferred net settlement (DNS) tend to attract credit risk on two fronts. First, for a system that explicitly guarantees settlement, the guarantor of the arrangement faces current exposure in the situation when a participant does not meet its payment or settlement obligations. Second, even when such arrangements have not been made, participants still face this risk with each other. Consider, for example, the possibility of a bank in a DNS system that is incapable of meeting its end-of-day net obligation. In this case, the bank's inability to pay represents credit risk, and the potential spillover from this event (i.e. the potential subsequent insolvency of other participants) could lead to financial contagion. A pure RTGS system, on the other hand, does not face credit risk, as payments are settled on a real-time and gross basis.

2.2 Liquidity Risk/Cost

Indeed, new developments in large-value payments systems, along with the increased use of central counterparties and the introduction of CLS for foreign exchange settlement, have helped to reduce the level of credit risk associated with settlement. They have also, however, made settlement more challenging for banks in terms of liquidity usage. With RTGS, for example, participants must have ample amounts of liquidity on hand in order to meet funding requirements in real-

[3]CPSS (2012)

time. To help with this, central banks offer intraday liquidity to banks, or allow for intraday overdrafts. In order to obtain this liquidity, however, collateral is usually required. For this reason, liquidity costs are increased because, in a system such as RTGS, if a firm is unable to sell assets, it must borrow intraday liquidity from the central bank, at the same time posting collateral.

2.3 Deferred Net Settlement

The two systems discussed above (DNS and RTGS) are at opposite extremes of the liquidity-risk tradeoff. The former is associated with low intraday liquidity costs, and higher settlement risk. This is because settlement does not occur regularly throughout the day, but at longer intervals or even simply once at the end of the day. The idea behind this system is to create efficiency in terms of liquidity usage as participants are not required to hold large amounts of liquidity during the day, and need not borrow from the central bank. Instead, by netting the outgoing and incoming payments throughout the day, liquidity requirements are vastly reduced, especially with multilateral netting, and participants need only hold the amount of liquidity calculated according to this netting at the prescribed time(s) of the day. This essentially eliminates more of the interest costs of liquidity as participants need not take part in intraday borrowing since settlement occurs relatively infrequently. At the same time, DNS promotes smaller net positions which can mitigate the risk of gridlock. Consider the following example:

There are 4 banks: A, B, C, and D. Table 1 has the system's payment schedule. In this example, each bank's end-of-day net position is zero. In other words, banks in this example (in a DNS system) need no liquidity at the end of the day to settle their payments. In this very simple case, therefore, there are no liquidity costs *and* no end-of-day credit risk because there are no outstanding obligations.[4]

Table 1: Payment Schedule

Time	
9:00am	Bank A pays Bank B $1B
	Bank C pays Bank D $2B
	Bank D pays Bank B $1B
Noon	Bank A pays Bank B $1B
	Bank B pays Bank C $3B
3:00pm	Bank C pays Bank A $3B
	Bank A pays Bank C $1B
	Bank D pays Bank C $1B

2.4 Real-time gross settlement

At the other extreme, real-time gross settlement offers virtually immediate finality of payments which limits the potential for systemic risk from unsettled claims since it does not allow for the accumulation of gross exposures between settlement times. RTGS, however, is relatively costly due to the need for liquidity management and the costs of acquiring intraday liquidity. Consider the prior example. In this simple case, the total liquidity requirement in a pure RTGS system is $13B compared to $0 in the DNS system. Furthermore, in RTGS, the timeliness of incoming payments is of greater importance because outgoing payments may be conditioned upon the receipt of other funds. In the situation where firms have incentive to delay outgoing payments, queuing of payments, widespread liquidity problems, and even gridlock, could occur. The trade-off between DNS and RTGS, therefore, can be described as the risk of default and contagion associated with the former, and the cost of intraday liquidity and system gridlock associated with the latter. Of course, these costs depend on the availability of intraday credit

[4]Of course, to have a net position of zero is highly unlikely; however, this simplified example illustrates the reduction in liquidity cost when compared to real-time gross settlement.

and other financial system requirements such as risk management and collateral requirements.

Indeed, these policies can combine to make DNS and RTGS extremely similar. For instance, lower delay costs incentivise more delay and therefore increase settlement risk in RTGS, while decreasing the intervals between settlement in DNS would increase the relative cost of liquidity associated with that system. Clearly, policies can act to provide some movement from the extremes in terms of the cost-risk tradeoff. For example, the United States' Fedwire system is a RTGS system where banks are allowed unsecured intraday overdrafts. These overdrafts have a limit and must be repaid by the end of the day, which effectively represents a shift from RTGS toward DNS as it changes the timing of the liquidity requirement to the end of the day while still satisfying a transaction obligation. Another policy that represents a movement along the cost-risk tradeoff is the provision of intraday liquidity which must be repaid by the end of the day. On the other hand, to reduce some of the costs associated with acquiring and holding liquidity altogether, systems can employ what are called liquidity-saving mechanisms (LSMs).

3 Liquidity-saving mechanisms

With a better understanding of the risks and costs faced by payment systems, the reasons to employ liquidity-saving mechanisms become more obvious. Clearly, liquidity costs, in the form of explicit interest rates for borrowing intraday liquidity and implicit opportunity costs, can be quite large and can thus harm a system participant's ability to earn profits. On the other hand, delay costs can also have a negative impact on a participant's day-to-day positions. The tradeoff between the two is that the higher the relative price of liquidity, the more participants delay their payments. The potential for gridlock in a system with a shortage of liquidity therefore poses increased operational risk as much higher volumes of payments may be processed all at once at the end of the day. For these reasons, payment and settlement systems have attempted to mitigate these costs to allow for increased efficiency.

LSMs are often used to optimize real-time systems by mitigating the effects of both liquidity and delay costs, and relieving, or indeed preventing, gridlock. There are a variety of LSMs available to payment systems and they may be used in different ways to tackle liquidity inefficiencies. What is very important to note, as Galbiati and Soramäki (2010) do, is that the use of a LSM need not reintroduce settlement risk. Indeed, they claim that to ensure this, it is sufficient to establish that a payment placed in a LSM creates no presumption of settlement and that its legal status remains identical to that of a non-submitted payment.[5] Under this assertion, settlement can only then occur when an offsetting cycle forms. At this point, payments instantly settle according to the real-time, gross, risk-free modality. This is not to suggest that LSMs *never* reintroduce settlement risk; but, that the reintroduction of risk is not necessary. If, therefore, the system's main objective is to reduce risk, it may do so while still reducing liquidity costs by using an appropriate LSM. Norman (2009) provides a wonderful survey of the different LSMs available. In this section, I provide a brief summary of his survey along with further relevant comments.

Norman notes that there are four main types of drivers of liquidity in an RTGS system: (i) need for cash in advance; (ii) the timing and urgency of payments; (iii) the size of payments; and, (iv) the aggregate end-of-day balance and its predictability. To begin, the more cash needed in advance, the higher the associated liquidity costs, especially the opportunity costs associated with holding liquidity. Next, the more urgent the payment, the higher the relative cost of delay and there-

[5] An example would be a payment placed in an internal queue.

fore the need for liquidity is again increased. Third, the smaller the payment, the easier it is to settle early, given a level of available liquidity and the receipt of outside payments. The opposite is the case for large payments. And lastly, the less predictable the end-of-day balance, the higher the possibility of holding costly idle liquidity.[6] The role of a LSM is to mitigate the effects of these drivers and attempts to do so can be decentralised and/or centralised.[7]

As a decetralised approach, banks can adopt monitoring and limiting whereby they set limits on the difference between received and sent payments from other participants, thus rewarding those that pay, and punishing those that do not. Without this, banks might always wait for incoming payments which could lead to gridlock. Even if gridlock does not occur, long delays of payments can lead to later settlement exposing the system to greater operational risk which leads to less predictability and thus higher (and inefficient) precautionary balances of liquidity. In terms of centralised approaches, artificial deadlines called throughput guidelines can be put in place to ensure gradual settlement of payments throughout the day.[8] An alternative centralised approach is to provide a natural incentive to pay early by setting banks a tariff for settling payments that become more expensive throughout the day.[9] Both of these methods encourage liquidity-saving by creating opportunities early in the day for payments between banks that broadly offset to be settled with relatively little liquidity.

Offsetting algorithms (OAs) are by far the most widely used centralised method of liquidity saving.[10] An OA works by matching individual (broadly) offsetting

[6] Readers are encouraged to see Norman (2009) for further explanation of these liquidity drivers.

[7] Balance-reactive and receipt-reactive gross settlement can be implemented with either approach.

[8] Norman (2009) notes the drawbacks of throughput guidelines, namely that it is difficult to distinguish between banks that receive instructions only late in the day and those that deliberately delay their payments to free-ride on incoming payments.

[9] Ota (2010)

[10] Norman (2009) points out that only OAs can cut the need for cash in advance for settling individual payments. This may help to explain their prevalence in payment systems.

payments (bilaterally or multilaterally), and settles them simultaneously thereby reducing the liquidity requirement to the net difference between the values of the payments that have been matched. Norman (2009) provides the following example: There are four banks: A, B, C, and D.

9 payments are effected simultaneously (i.e. without relying on incoming payments). Figure 1 illustrates the flow of payments when there is no OA.

Figure 1: No Offsetting Algorithm

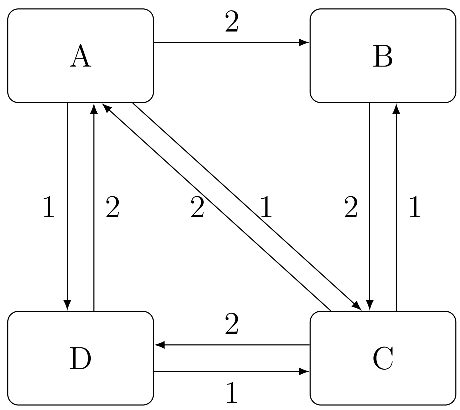

Note the individual banks' liquidity requirements:

A: $2 + 1 + 1 = 4$

B: 2

C: $2 + 1 + 2 = 5$

D: $2 + 1 = 3$

Total liquidity requirement: 14

Figure 2 illustrates the effect of bilateral offsetting. The system's liquidity requirements are dramatically reduced, totaling 6.

Figure 2: Bilaterally Offsetting Algorithm

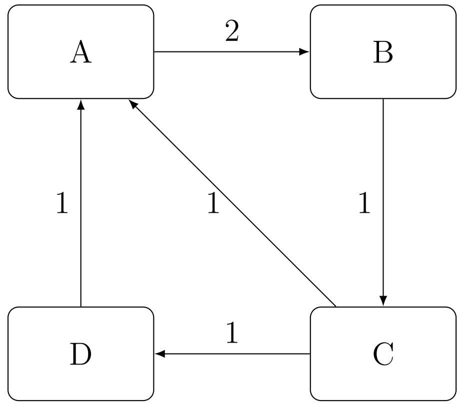

And finally, Figure 3 demonstrates the effect of multilateral offsetting. The liquidity requirement is now reduced to 1.

Figure 3: Multilaterally Offsetting Algorithm

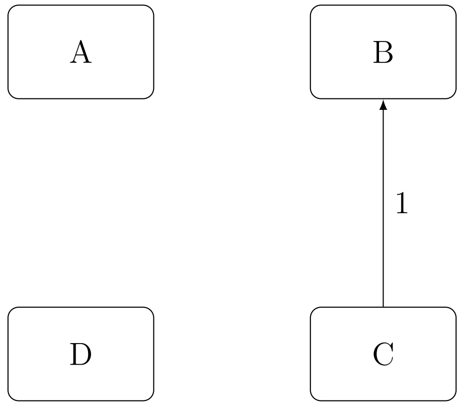

In this example, multilateral offsetting is the most efficient use of liquidity. It is obvious, therefore, that OAs are good options for dealing with type (i) liquidity drivers; however, they pose a problem of timing, particularly in real-time systems. Should netting calculations take a long time, the decision to continue processing payments becomes more difficult because it hinders the real-time risk-free modality of RTGS. In fact, in some situations, offsetting algorithms do not achieve any liquidity saving. This is particularly true when there are large transactions to be settled because it can be difficult to accumulate enough small payments in the opposite direction. Furthermore, they do nothing to deter the other types of liquidity drivers identified by Norman (2009).

For these reasons, some systems have complemented offsetting by adopting payment splitting.[11] The splitting of large payments encourages partial intraday repayment of loans which could also help reduce credit risk outside of the payment system.[12] Indeed, the splitting of payments can also help to alleviate gridlock, as smaller payments are less likely to be delayed.[13] Furthermore, to deal with the aforementioned issue of urgency, most systems have adopted a "hybrid" framework by which participants are able to reserve liquidity outside of queues in order to make normal RTGS payments. This way, urgent payments need not wait to be offset in the queue, nor need they depend on account balances. Finally, LSMs cannot be used to affect end-of-day balances; however, they can reduce banks' precautionary liquidity needs by increasing predictability. This helps to decrease the amount of idle liquidity in the system, thus promoting increased efficiency. Norman provides an elegant summary of LSMs and liquidity drivers which I summarize in Table 2.

---

[11] Norman (2009) identifies CLS, SIC, and BOJ-Net as using payment splitting.

[12] Millard and Polenghi (2004)

[13] The problem with payment splitting is the legal ramification of having a payment halted after a portion of it has been made, but prior to its completion. Obviously, the legal framework needs to take this into consideration.

Table 2: LSM and Liquidity Driver Summary

Approach to Liquidity Saving	Need for liquidity in advance	Size of Payments
(Bilateral) limits/reciprocal payments strategy	n/a	n/a
Central queuing/liquidity reservation	n/a	n/a
Throughput guidelines	n/a	n/a
Graduated intraday tariff	n/a	n/a
Offsetting Algorithm	Liquidity need reduces to net difference of offsetting payments	n/a
Payment Splitting	n/a	Liquidity need reduces to size of splitting threshold

Note: All approaches to liquidity saving described above could potentially lead to (non-urgent) payments being submitted/settled earlier that would be otherwise the case, thereby increasing the scope for co-ordination with (urgent) payments in the opposite direction, thereby helping to reduce liquidity need.
Note: The liquidity need driven by aggregate end-of-day balance *cannot* be reduced; however, any liquidity saving approach that leads to earlier submission/settlement may improve predictability of the end-of-day balance thereby reducing precautionary liquidity need.

4 The Literature

A great deal of attention has been directed toward the design and operation of large-value payment systems, and much of it is concerned with the operation and effectiveness of liquidity-saving mechanisms. The goal of this section is to provide some literature concerning LSMs to assist in understanding the challenges that are faced by settlement systems attempting to mitigate the costs of liquidity. The papers discussed below outline several parameters to be considered when constructing liquidity management policies including delay costs, liquidity costs, size of the system, and certain systemic shocks. They consist of a simulation study, an interbank model comparing welfare with and without LSMs, and a game-theoretic approach to modelling a payments system. These three papers are particularly important as they stress the importance of parameters to consider when deciding on a type of LSM, the possible welfare implications of a LSM, and the importance of the efficiency tradeoff between RTGS and a LSM. A better understanding of these concepts is integral for any *a priori* insight into the effects a LSM will have on a system. These insights will become useful later in this paper when a discussion of LSMs and continuous linked settlement is presented.

4.1 Leinonen and Soramäki (1999)

Leinonen and Soramäki (1999) (henceforth referred to as LS99) quantify the relationship between liquidity usage and settlement delay in net settlement, real-time gross settlement, and hybrid systems. They also quantify the combined costs of liquidity and delay in these systems.[14] The purpose of their paper is to analyse ways of reducing costs via optimization features using LSMs. To do so, they use a payment simulator developed at the Bank of Finland[15] and data from the Finnish payments system. In their framework, the tradeoff is between liquidity costs and delay costs; however, the latter can be viewed as a risk in that a cost of long delay is the higher potential for default. As Berger *et al* (1996) point out, this tradeoff can be affected by technological, financial, and regulatory innovations.

Two important definitions provided by LS99 are as follows: (1) a 'rigid' liquidity regime is one in which participants cannot adjust the amount of liquidity they hold after they receive the (pre-specified) amount of funds from the central bank; and (2) a 'flexible' regime is one in which participants can increase or reduce the external funds in the system. The advantage of the latter is that changes in liquidity needs can be more easily met when participants are allowed to alter their

[14] These costs are actually quite difficult to measure. A shortcoming of this paper, which is noted by the authors, is that they assume linear cost functions when, in reality, this is probably not the case.

[15] Another shortcoming noted by the authors is that the pattern of payment flows affects liquidity usage and the efficiency of optimization measures. Therefore, the exogenous nature of payments flows is a shortcoming of the the simulations. If the features of the settlement system changes, the simulation cannot account for the potentially altered behaviour of customers and participants.

holdings throughout the day. LS99 use these regimes to analyse the relationship between liquidity usage and settlement delay. The inverse relationship between the two can be explained as follows: participants with high levels of liquidity will delay very little by making payments early in the day to avoid delay costs, whereas participants with low levels of liquidity will delay payments in hopes to fund them with incoming payments. LS99 provide Figure 4 to illustrate the relationship between liquidity usage and settlement delay for rigid liquidity regimes. Curve AC shows that, if the number of net settlements during the day is increased, payment delay can be reduced in exchange for greater liquidity. LS99 note that, if the number of settlements is increased to the point where net settlement is executed after each transaction (against earlier netted payments), the liquidity usage and settlement delay correspond to those of a RTGS system without queuing as well as to a continuous net settlement (CNS) system with fully collateralised debt limits.

Figure 4: Relationship between liquidity usage and settlement delay, rigid liquidity regime

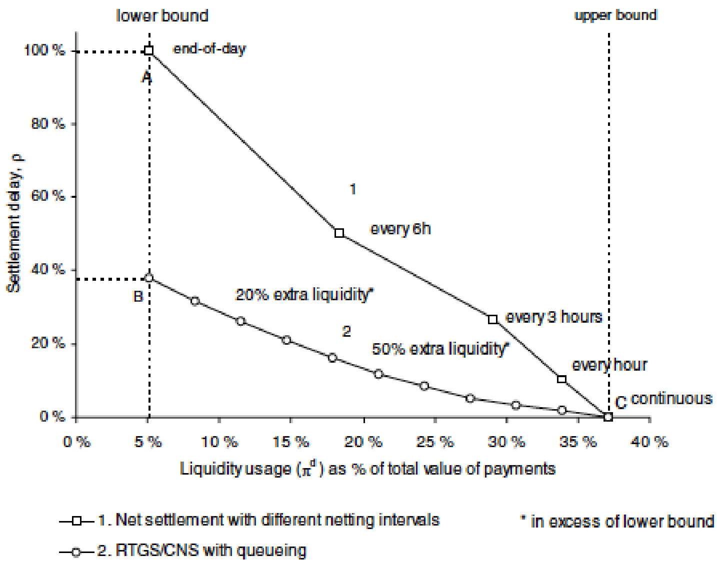

[a] Source: Leinonen and Soramäki (1999)

The system level tradeoff curve (line segment CB) shows that a RTGS or CNS system can reduce liquidity by moving northwest on the tradeoff curve. The curve is convex, indicating the diminishing returns of liquidity usage in terms of settlement delay, and vice-versa. Obviously, then, a participant's behaviour depends on their weights of the cost of liquidity relative to that of settlement delay. For example, a bank that weights the cost of liquidity high relative to that of settlement delay will choose a point near B. Furthermore, the illustration shows that a RTGS (CNS) system can operate on the same amount of liquidity as any of the netting systems, with reduced payment (settlement) delay. In fact, the authors find that these delays were reduced by 62%.

They also note that, at equal costs for liquidity and delay, the combined cost is at its maximum with end-of-day settlement. This is the case when all payments are delayed until the end of the day, and the end-of-day liquidity usage equals the lower bound of liquidity.[16] Note then that reducing the interval of net settlement reduces the combined cost since the increased cost due to rising liquidity usage is less than the cost reduction on settlement delay. Therefore, in this case, it can be concluded that costs are minimized with CNS. For RTGS with queuing in a rigid liquidity regime, liquidity costs are minimized and delay costs are maximized at the lower bound of liquidity. This is intuitive: the less liquidity, the greater the delays. At the upper bound, costs stem solely from liquidity usage as participants no longer delay but rather settle on a real-time gross basis. Indeed, theoretically, the only time the costs for net and gross settlement are equal is when payment delay is free. In this case, end-of-day settlement minimizes the costs for netting systems and the use of lower bound liquidity minimizes costs for RTGS systems.

[16] The lower bound of liquidity is defined as the minimum external liquidity needed in all settlement systems where payments can be delayed (eg by queuing) whether operating on a net or gross basis. The upper bound of liquidity equals the amount of liquidity that must be available to the participants for immediate settlement throughout the day. Any liquidity above this remains idle.

Therefore, the efficiency of a settlement arrangement hinges heavily on the parameters of interest, namely the price of liquidity and the cost associated with delaying payments. A key insight of this paper is that the type of liquidity regime is important. LS99 point out that the relative cost advantage of the flexible regime ranges from 1% at the lower bound of liquidity, to 25% at the upper bound. The reason for this is quite obvious: at the lower bound, there is little idle liquidity and therefore flexibility of the regime has little effect on the costs associated with settlement. Flexibility, they note, is a feature that helps banks overcome wider variations in liquidity demand at certain times of the day. On the other hand, for a rigid regime, participants will operate at the lower bound of liquidity if the interest cost of delay is less than 52% of the interest cost of liquidity. As delay costs change, participants move along the tradeoff curve as illustrated above. They find that if delay costs are about twice as high as liquidity costs, the system-level cost-optimum is for each participant to settle its payments immediately. In a rigid regime, however, participants will tend to delay payments and thus incur higher delay costs because part of this cost is absorbed by the use of idle liquidity.

LS99 also consider the potential for gridlock and identify two methods by which it can be alleviated and/or avoided. For instance, queued payments can be netted (either fully or partially) to reduce the maximum daily time a system is gridlocked. In their simulation, the resulting duration of gridlocks was reduced to approximately the time interval between the nettings. The second method suggested is payment splitting. Again, this helps to prevent gridlock since small payments are more easily offset than large payments. When the size of payments is smaller, the amount of idle liquidity is lower and a greater share of the available liquidity is employed for settlement. Essentially, these two methods act to shift the tradeoff curve toward the origin and each participant can effect faster settlements with a given amount of liquidity. The simulations suggest that both are effective methods, but that payment splitting is technically very efficient at alleviating the gridlock problem. LS99 find that in many cases it was more effective than the netting algorithms, especially when over 5% of transactions were split and when more than the lower bound of liquidity was available to participants.

It would appear that Central Banks can also support risk reduction and rapid payments targets by providing low-cost intraday liquidity as well as more flexible ways for participants to add or withdraw liquidity from the system. Again, at higher levels of liquidity, LS99 show that flexible regimes reap much higher rewards in terms of efficiency. In the cases when liquidity is scarce, optimizing gridlock features are required and effective. In this study, payment splitting and the partial netting of queues were most efficient. Of course, this may not be the case for all payments systems and much attention needs to be made to the many schemes available. Indeed, the efficiency of these methods depends on the flow of payments processed. The important thing to note (which LS99 does) is that the costs of the settlement process can be reduced by applying optimization methods in situations with limited amounts of available liquidity. At very low levels of liquidity, these methods can be effective.

4.2 Martin and McAndrews (2008), and Jurgilas and Martin (2010)

From an economic theory standpoint, it is naturally important for us to consider the welfare effects of a LSM. To do this, Martin and McAndrews (2008) (henceforth referred to as MM08) study the incentives of participants in a RTGS with and without the addition of a LSM. In the model, participants face a liquidity shock and different costs of delaying payments. They trade-off the latter with the cost of borrowing liquidity, much as in the previous analysis by LS99. Unfortunately, the risk of default is not taken into consideration in this paper; however, the results are still interesting and worth noting. The authors show that the design of a liquidity-saving mechanism has important implications for welfare. They illustrate the decisions of participants under different (positive, negative, or no) liquidity shocks and ask whether a LSM can improve these outcomes and find that adding a LSM is not necessarily welfare-improving.

MM08 model a LSM arrangement wherein a payment may be made only if they are offset by an incoming payment. After participants observe their liquidity shock and the time criticality of their payment, they have the choice to put the payment they must make to another participant into a queue. These payments can offset multilaterally as described in Section 3. It is assumed that there is a non-strategic agent (such as CLS) that does not use a queue but that receives time-critical payments and sends payments to participants. In a situation without liquidity shocks, then, it is obvious that queuing a payment weakly dominates delaying it. Therefore, when there are no liquidity shocks, all payments are released in the morning and all participants achieve the highest possible payoff. In the case with a liquidity shock, however, this is no longer the case. In short, depending on the length of the queue, several equilibria can occur. As the cost of delay increases, participants queue fewer payments. If this cost is small enough, it remains an equilibrium for all participants to put their payments in the queue.

It is important to note that when the fraction of participants who make their payment early decreases, so does welfare. As the authors note, the queue can eliminate some beneficial coordination between agents by giving them the option to send payments conditional on receiving them. They provide an example showing that, for some parameters, welfare can indeed be lower with a LSM. MM08 present two general mechanisms: balance reactive (BRLSM) and receipt reactive

(RRLSM) liquidity-saving mechanisms. Under the former, participants can condition the release of a queued payment on their level of balance. Under a RRLSM, offsetting payments are used to trigger the release of payments in the queue regardless of the participant's balance. MM08 show that, for BRLSMs, welfare can be higher than regular RTGS when delay costs are high and the proportion of time-critical payments are relatively low; but, welfare can be lower than RTGS when the costs of delay are small and the payments to the outside settlement system are large. When compared to an RRLSM, the BRLSM yields a better outcome when delay costs are high and the settlement system is not too large. Welfare under the RRLSM, however, is always at least as high as in the regular RTGS.

Interestingly, however, Jurgilas and Martin (2010) use the same basic model to show that, unlike in the fee-based system above, a LSM *always* improves welfare in a collateral-based RTGS system. Indeed, in a collateral-based system with a LSM, they find that the planner's allocations are the same for all parameter values. The LSM aligns the bank's incentives with those of the planner by allowing offsetting payments to settle with less collateral. Because the cost of collateral is sunk at the time banks learn their liquidity shock, the planner does not face the incentive to have banks with large negative liquidity shocks delay payment. This is interesting as many developed LVPSs are collateral-based.

These papers are important as they make it very clear that the introduction of a LSM is not always welfare-improving. On the contrary, payment systems need to take great care to ensure that the introduction of such mechanisms fit the parameters of the current economic framework. Furthermore, while simulations provide good insight into the effects of such changes on a particular system, as noted earlier such simulations cannot take into account the exogenous character of payment flows and the potential for changes in the behaviour of economic agents. This is why models such as the one presented in these two papers can be useful when considering implementation of such mechanisms. The difficulty for operators resides in determining the size of the parameters of these models, and evaluating the importance of the underlying assumptions and abstractions.

4.3 Galbiati and Soramäki (2010)

Galbiati and Soramäki (2010) (henceforth, GS10) also delve into the question of the benefits of a LSM. They provide a single simultaneous-move game, for which they find the Nash equilibria, thus providing a game-theoretic representation of a payment system. The payoffs are numerically generated by an algorithm which mimics a payments system. The authors allow banks to exchange many payments over many time-intervals, generating complex liquidity flows with queues, gridlocks, and cascades.[17] This is important because it provides a very similar landscape to that of the reality faced by banks. The paper focuses exclusively on the symmetric equilibria of the game where all banks choose the same urgency threshold, λ, and beginning liquidity, τ, and each action is a best reply to others' choices. For simplicity, a strategy profile is called 'equilibrium' only if any unilateral deviation from it is not beneficial to the deviator - even if it would lead to a non-symmetric outcome.[18]

GS10 concentrate on an extreme form of internal queuing that delays payments until the end of the day because it is far more difficult to assess *a priori* what the effect would be on equilibrium choices of "smarter" internal queues. They propose the following key question on the mechanical gains of a LSM: *"Given a level of delays, what is the minimum amount of liquidity needed not to exceed it (allowing*

[17]Readers are encouraged to see Galbiati and Soramäki for the specifics of the model.
[18]Obviously, this is subject to the possibility of missing equilibria where banks adopt different, albeit mutually optimal, choices.

any choice of beginning liquidity)?" They find that a LSM reduces liquidity needs only if the target delay level does not fall below a threshold. If delays were forced to remain below this, the planner would choose to use the RTGS stream for all payments by setting $\tau = 0$ and providing large amounts of liquidity to the system. If delays are allowed to exceed this level, then the opposite is true. That is, $\tau = 1$ and the planner uses the LSM exclusively. This is, of course, very similar to what LS99 asserts. So, if the system is suffering from a shortage of liquidity, or if delay costs are substantially low and therefore participants are willing to accept longer delays, the LSM should be used exclusively. GS10, however, correctly points out that the real gains from a LSM emerge when choices are made by independent banks in a strategic context. Indeed, in this scenario, both streams are typically used. Consequently, the efficiency of each stream depends on the share of payments routed to it by the 'others'. Because 'others" liquidity and 'my' liquidity are complements, the more liquidity posted by 'others', the more attractive RTGS becomes. Therefore, the efficiency of RTGS depends further on the amount of liquidity committed by 'others' and, consequently, there exists an efficiency tradeoff between RTGS and LSM: the higher one's share of payments, the more efficient it becomes and the lower the efficiency of the other.

For an equilibria analysis, GS10 use a wide range of liquidity prices and keep the price of delays fixed. As noted, the equilibrium choices of banks are different from those of the planner because they are not dichotomous. This is key considering, in reality, banks have several options at their disposal. Again, there is a tradeoff between liquidity and delays and therefore, as the relative price of liquidity rises, banks post less liquidity and resort more to internal queues. Since this differs from the strategy of the cost-minimizing planner, we know that this is inefficient. In this case, the planner would provide more liquidity and would never delay payments internally. GS10 find that the equilibrium costs are always more than 15% higher than the social optimum, reaching multiples thereof at higher liquidity prices.[19] According to GS, two externalities explain this inefficiency:

- A possible positive externality in liquidity provision is that incoming payments to a bank can be recycled. Therefore, liquidity can be thought of as a common good.[20] For this reason, equilibrium liquidity provision may fall short of the social optimum.

- A possible negative externality of internal queues is that banks have incentive to delay less urgent payments and use liquidity for more urgent ones (such as those to CLS); but, by doing so, they slow down the liquidity recycling in RTGS, which in turn affects other banks. Therefore, banks queue more than they should.

Because LSMs offer a more efficient way of queuing payments, banks use them more intensely, which, as mentioned above, leads to less efficiency in RTGS. GS10 assert that the "main novelty" with a LSM compared to internal queues is that, for an intermediate range of liquidity costs, multiple equilibria emerge:

- A corner equilibrium where all payments are settled via the LSM.

- Equilibria with moderate use of both ('good' equilibria)

- Equilibria with high usage of both liquidity (RTGS) and the LSM ('bad' equilibria)

The first equilibrium is socially optimal and occurs when liquidity prices are very high. The second set of equilibria are typically those with the lowest costs.

[19] At *extremely* high prices, the equilibrium coincides with the planners optimal choice, namely exclusive use of a LSM.

[20] See Angelini (1998)

The last, however, are the most costly and can be explained by the the fact that a LSM features economies of scale. Because of this, its high usage may be self-sustaining; however, overuse of a LSM is detrimental to the RTGS stream. Because of the efficiency tradeoff mentioned above, the decreased efficiency of RTGS caused by the subsequent reduction in recycling leads to the RTGS stream requiring inefficiently high amounts of liquidity.

Their conclusion is straight-forward and consistent with the two studies above: a system with a LSM is capable of delivering better outcomes than those resulting in the absence of a LSM. For a range of liquidity prices, there are 'good' equilibria which feature lower total costs than their non-LSM counterparts due to faster settlement and often also due to lower liquidity usage. There is, however, a range of intermediate liquidity prices for which a system with a LSM may generate 'bad' equilibria with high liquidity usage and intense use of the LSM, therefore requiring higher costs than a system without a LSM. They point out that this suggests that LSMs can be useful, but that they may need some co-ordination device to ensure that banks arrive at a 'good' equilibrium. This is consistent with the MM08 assertion that the introduction of a LSM can remove some beneficial co-ordination between participants.

5 Continuous Linked Settlement

The discussion and literature review above lay a solid foundation for thinking about the effects of LSMs in different payment systems. In this section, I consider their effects on a delivery system called continuous linked settlement (CLS). I do this because the modern global economy could well be defined by its interconnectedness and interdependence. The recent financial crisis and the current European sovereign debt crisis have exposed the fact that the economies of the world are more sensitive than ever to the happenings across borders and waters. As a consequence of this, foreign exchange policies are of utmost importance. In the past, foreign exchange transactions would take place between two or more parties in private. This meant that the party to fulfill its obligation first was exposing itself to foreign exchange settlement risk. It exposed itself to the possibility of receiving payments late, or not at all. Indeed, depending on the size of the transaction, liquidity and solvency problems can arise when a party to a transaction fails to fulfill its obligation. If significant enough, a chain reaction could ensue and a systemic breakdown would be a possibility. The most documented example of risks associated with foreign exchange is Herstatt risk, named after a small German bank, Bankhaus Herstatt that failed in 1974.

Herstatt risk refers to risks which span more than one market or time zone, usually involving foreign exchange or securities instruments (CLS 2009). After being ordered to liquidate, Herstatt's New York correspondent bank suspended all outgoing dollar payments which resulted in their counterparties being fully exposed to the value of the deutschmarks they had paid them earlier in the day. When such foreign transactions fail to settle, significant losses can be incurred. In the case of this particular incident, the international banking system nearly collapsed. To help reduce these risks, CLS was introduced in 2002 by the world's largest foreign exchange banks in response to central bank concerns about the impact of potential foreign exchange settlement failures on the global financial system (CLS 2009). Indeed, the losses caused by the failure of Lehman Brothers in 2008 were limited by the use of CLS (CLS 2009).

CLS acts as a common counterparty (not a central counterpary) between participating banks, holding separate RTGS settlement accounts with each of the participating central banks. Each member also has multicurrency accounts with

CLS Bank which communicates to members, RTGS systems, and third parties via SwiftNet. Using this infrastructure, CLS settles transaction simultaneously in real time. Participants, however, fund their settlement accounts on a net basis. Settlement takes place during a five-hour overlapping window when each of the banks' respective RTGS systems are open and able to receive and make payments. At the start of the day, CLS Bank calculates each member's pay-in schedule based on the instructions due to settle that day. This schedule includes the minimum amount of each currency to pay and the time at which the payments must be made. Members can check their net pay-in totals at any time before the start of the settlement day. Settlement instructions can be submitted directly to CLS Bank for processing until 06:30 CET, at which time members receive their final pay-in schedule for that settlement day. Once the first funding is paid in, the settlement cycle starts at 07:00 CET. CLS ensures that both settlement member accounts are in credit and that settlement of an instruction will not cause either settlement member to exceed its limits. It then settles the transactions that are validated and matched. If an instruction fails the checks, it is sent back into the cycle which repeats every few minutes. If an instruction fails at the end of the day, it is not executed and no funds are exchanged. The completion time target is 09:00 CET, after which all short and long positions are paid out to ensure no outstanding funds are left in the settlement accounts.

5.1 Liquidity Saving and CLS

As with payment systems, CLS (which is a delivery system) faces the tradeoff between liquidity costs and risk discussed above. To deal with the settlement risk associated with foreign exchange, it settles transactions payment-versus-payment, using a real-time settlement system. Payment-versus-payment (PVP) is a mechanism that ensures that a final transfer of one currency occurs only if a final transfer of the other currency takes place. Therefore, there still exists the risk that a participant does not receive the currency it demands; however, exposure is limited because there is no risk of losing a sent payment. Because PVP helps guarantee that both parties meet their obligations, these foreign exchange transactions become self-collateralising. PVP and real-time settlement, therefore, eliminate a great amount of risk to participants; however, if settlement was achieved on a gross basis, the liquidity requirements would be significant. To mitigate these costs, CLS has incorporated some liquidity-saving mechanisms. For instance, participants fund their settlement accounts on a net basis which drastically decreases their liquidity requirements.[21] As described earlier, netting of obligations can drastically limit the costs of liquidity for a participant by cutting obligations significantly. CLS also uses payment splitting to ease the pressure that larger transactions can have on the system.

Another mechanism that CLS uses is called an in/out swap. An in/out swap is a transaction by which time-critical liquidity requirements are reduced. A participant's net short positions and net long positions can be substantial in which case it may be helpful for two participants with opposite positions to trade currencies within CLS to reduce these positions. An in/out swap offers the possibility of doing this with large mutually offsetting balances. For example, a participant with a large net short position in CAD and net long position in USD can agree to perform a currency swap with another participant with reverse balances. The first leg settles inside CLS, and the second leg settles outside CLS. The effect is to reduce the net short position that settles within CLS, thus reducing the funding requirement of the participants taking part in the swap. Due to the outside leg of the swap, the participants would then have exposure outside of CLS as they

[21] Multilateral netting in CLS has resulted in the payment flows actually transferred between participants only amounting to around 2% of the gross amounts actually settled (CLS 2009).

must now use traditional correspondent banking; however, this is small relative to the overall total position traded in CLS. As a result of the swap, the FX positions of each member have not changed but each one has a reduced pay-in to CLS. These swaps illustrate the tradeoff between risk and liquidity costs wonderfully: the banks are subjecting themselves to outside exposure in return for reduced liquidity requirements. In a sense, these swaps are contrary to the spirit of CLS (which was created to eliminate foreign exchange settlement risk) because it is reintroducing this risk on the outside leg; however, they work as LSMs do: to promote efficiency by reducing liquidity requirements in CLS itself.

What may be just as interesting as the liquidity-saving measures taken by CLS are the effects of such measures taken by large value payments systems on CLS. For instance, what is the effect on CLS of a system that employs pure RTGS compared to a system that operates under a hybrid framework? In other words, what kind of impact do policies employed by central banks have on CLS? Consider, for example, the policy of the Swiss National Bank (SNB) of providing interest-free intraday liquidity. This enables participants in the Swiss repo market to draw on Swiss franc sight deposits with the SNB against the lodging of collateral. This liquidity must be repaid by the end of the settlement day of the Swiss RTGS system. This policy helps participants deal with the increased time-critical liquidity requirements that CLS imposes on its participants by effectively removing the time-critical aspect of CLS. Instead, banks must now weigh the cost of delaying a CLS payment (the reintroduction of foreign exchange settlement risk) with the cost of acquiring liquidity by the end of the day.

This is exactly the strategy outlined by Leinonen and Soramäki (1999) above, namely that Central Banks can support risk reduction and rapid payments targets by providing low-cost liquidity as well as more flexible ways for participants to add or withdraw liquidity from the system. Recall that they found that at higher levels of liquidity, flexible regimes reap higher rewards in terms of efficiency. If we also consider that gridlock is much less likely in the situation with ample liquidity, we may note that, if all Central Banks were able to provide interest-free intraday liquidity, there may be reduced need for optimizing gridlock features such as payment splitting because large payments could be made to CLS with relative ease. Of course, while this option is available in all payments systems, it may not benefit the CLS process directly. For Asian participants, for example, the end of the settlement day is the same as the time of settlement in CLS.

Furthermore, if we consider the bank behaviour described in Martin and McAndrews (2008), we recall that, at sufficiently low delay costs, all participants, despite their liquidity shock, delay all of their payments. This includes time-critical payments to the nonstrategic agent, in this case CLS. As delay costs increase, participants begin to make their time-critical payments early. Therefore, if the relative cost of delay can be made larger, or if foreign exchange settlement risk can be made even less appealing, banks should be incentivized to make early payments to CLS. Offering interest-free intraday liquidity is one way of doing this. Another way would be to introduce the tariffs and/or throughput guidelines discussed above. Indeed, this could be applied within CLS itself within the five-hour settlement period. The introduction of such mechanisms should be done with caution, however. Because the settlement period is relatively short, throughput guidelines and tariffs could inefficiently increase the liquidity requirements of participants.

We have also seen how the choice of LSM is important for welfare maximization within a LVPS. If we consider the effect they have on CLS, we see that it can be important here as well. Consider the two LSMs described by MM08: balance reactive (BRLSM) and receipt-reactive (RRLSM). Recall that the authors show that, in some cases, and when delay costs are small and payments to the outside settlement system are large, the BRLSM can result in participants sending some time-critical payments (such as those to CLS) to the queue, rather then sending them in the early period. If this is the case, plain RTGS can achieve higher welfare and timely payments to CLS because it creates some incentives for coordination. An RRLSM, however, weakly dominates RTGS because participants either submit all of their time-critical payments early (in the case of high delay costs) or submit all their payments to the queue. Depending on the parameters, therefore, the choice of LSM in a payments system has implications for timely payment in CLS and therefore, by extension, the reintroduction of foreign exchange settlement risk.

Recall again that the introduction of a LSM does not necessarily result in increased efficiency within a LVPS. As mentioned, Galbiati and Soramäki (2010) find that there can also exist 'bad' equilibria by which it is meant that there is both intense LSM usage and high RTGS usage, resulting in high liquidity costs. Again, this is because there exists an efficiency tradeoff between RTGS and LSMs: the more heavily one is used, the less efficient the other becomes. Therefore, in some cases, high LSM usage results in RTGS inefficiency, causing reductions in payment recycling and therefore increased liquidity requirements. Taking the large liquidity requirements imposed by CLS into consideration, it becomes evident that such pressure could easily result in the delay or queuing of a time-critical (CLS) payment. Again, interest-free intraday liquidity and the potential for in/out swaps should help to alleviate such pressure.

In the future, it would be interesting to test the effects of these external LSMs empirically. This could perhaps be accomplished by simulating payment flows to CLS Bank from a set of participants in LVPSs with and without particular LSMs. In doing so, it would be best to control for some of the important parameters discussed above such as delay costs, liquidity costs, and payments system size. Doing this, one may be able to determine the effect a particular external LSM (offsetting algorithm, payment splitting, tariffs, etc.) has on CLS. This may be useful for determining which combinations of LSMs are most effective at reducing liquidity costs, given certain intensities of time-critical payments. Doing this may allow for a welfare analysis from the CLS perspective rather than from that of a LVPS. The policies of central banks seem to imply that hybrid settlement systems are best to ensure participants are better able to make such time-critical payments; however, it would be interesting to explore alternative options in terms of strictly using a combination of liquidity-saving mechanisms.

6 Concluding Remarks

Obviously, a substantial amount of work has been dedicated to the area of payments systems. Much discussion has been had over the most effective methods of dealing with the risk-cost tradeoff that all such systems face. The purpose of this paper was to provide a source of reference to help understand this tradeoff and the ways in which LVPSs deal with it. Technological, financial, and regulatory innovations can all help to move along this tradeoff curve, even shifting it toward a situation of lower risk *and* cost. A liquidity-saving mechanism, however, is only concerned with reducing the cost of liquidity. There are many ways to do this including offsetting and splitting payments. The problem with LSMs lies not in a lack of options but in their implementation. This paper offers a thorough review of three important papers in the area of liquidity management, all of which point to two key facts: LSMs are not always welfare-improving and, therefore, their implementation must take the unique economic environment surrounding the system into account. Some key parameters in this decision include the costs of delaying payments, the costs of acquiring liquidity, the size of the system, and the liquidity rigidity of the regime in question. Indeed, flexible regimes with high delay costs and inexpensive intraday liquidity may require much less in terms of LSMs than a more rigid regime with costly liquidity.

Lastly, this paper offers an overview of CLS, a relatively new and extremely important innovation in the area of foreign exchange settlement. Indeed, with the increasing interconnectedness of the world's largest economies, the functionality of CLS is of utmost importance. As a result, it represents an excellent opportunity to consider the effects of LSMs on the global economy. This paper discusses the use of netting, payment splitting, and in/out swaps in CLS and how they help to reduce the liquidity requirements imposed on participants by CLS's payment-versus-payment delivery. Future research should focus on the effects of external liquidity-saving mechanisms on CLS and its timely delivery of funds, controlling for the key parameters discussed above such as delay costs, liquidity costs, and the intensity of payments made to CLS. While hybrid RTGS systems appear to be most common at the moment, whether strict use of LSMs could ever be beneficial for paying time-critical payments should be investigated. If so, this could represent even further liquidity savings for participants in practice.

References

Berger, A, D Hancock and J Marquardt (1996). "A Framework for Analyzing Efficiency, Risks, Costs, and Innovations in the Payments System." *Journal of Money, Credit and Banking*, Vol. 28, No. 4, Part 2: Payment Systems.

Chiu, J and A Lai (2007), "Modelling Payments Systems: A Review of the Literature," Bank of Canada Working Paper 2007-28.

CLS (2009) "Continuous linked settlement: The great FX fix"

Committee on Payment and Settlement Systems (2012), "Principles for financial market infrastructures," Bank for International Settlements and International Organization of Securities Commissions.

Galbiati, M and K Soramäki, (2010), "Liquidity-saving mechanisms and bank behaviour" Bank of England Working Paper No. 400

Hampton, T (1999), "Intra-day liquidity and real-time gross settlement - 18 months on." Reserve Bank of New Zealand Bulletin Vol. 62, No. 4

Jurgilas, M and A Martin (2010), "Liquidity-saving mechanisms in collateral-based RTGS payment systems." Bank of England Working Paper No. 389.

Kahn, C and W Roberds (2000), "The CLS Bank: A Solution to the Risks of International Payments Settlement?" Federal Reserve Bank of Atlanta, Working Paper 2000-15a.

Leinonen, H and K Soramäki (1999) "Optimizing Liquidity Usage and Settlement Speed in Payment Systems" Bank of Finland Discussion Papers

Mägerle, J and D Maurer (2009), "The Continuous Linked Settlement foreign exchange settlement system (CLS)." Swiss National Bank

Martin, A and J McAndrews (2008), "Liquidity-saving mechanisms," *Journal of Monetary Economics*, Vol. 55, No. 3, April, pages 554-67.

Norman, B (2010), "Liquidity saving in real-time gross settlement systems - an overview" Bank of England Financial Stability Paper No. 7

Ota, T (2011), "Intraday two-part tariff in payment systems," Bank of England Quarterly Bulletin; 2011 2nd Quarter, Vol. 51 Issue 2