Financial Stability Paper No. 11 – June 2011

Intraday liquidity: risk and regulation

Alan Ball, Edward Denbee, Mark Manning and Anne Wetherilt

BANK OF ENGLAND

BANK OF ENGLAND

Financial Stability Paper No. 11 – June 2011

Intraday liquidity: risk and regulation

Alan Ball, Edward Denbee, Mark Manning and Anne Wetherilt

The authors are grateful to Damian Harland at the FSA, Toby Davies, John Jackson, Edwin Schooling Latter, Pippa Lowe, Lisa Newman, David Norcross, Ben Norman, Will Speller and Peter Zimmerman at the Bank of England, Antoine Martin at the Federal Reserve Bank of New York and Rodney Garratt at the University of California, Santa Barbara for useful comments and suggestions. Any remaining errors are our own.

alan.ball@bankofengland.co.uk
Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH

edward.denbee@bankofengland.co.uk
Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH

mark.manning@bankofengland.co.uk
Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH

anne.wetherilt@bankofengland.co.uk
Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH

The views expressed in this paper are those of the authors, and are not necessarily those of the Bank of England. This paper was finalised on 1 June 2011.

ISSN 1754–4262

Contents

Intraday liquidity: risk and regulation

Alan Ball, Edward Denbee, Mark Manning and Anne Wetherilt

Banks require access to liquidity intraday in order to settle obligations in payment and settlement systems. The recent financial crisis has highlighted the need for banks to improve their liquidity risk management, including the management of intraday liquidity risk. The FSA's new liquidity regime includes intraday liquidity as a key risk driver and requires that banks calibrate their liquid asset buffers considering their need for liquidity intraday, both in normal and stressed circumstances. The Bank fully supports this approach. However, this will increase the cost of intraday liquidity and so could create incentives for banks to change their behaviour as they seek to minimise costs. If this results in payment delays, it risks jeopardising the smooth functioning of payment and settlement systems. There are a number of tools that authorities could use to minimise the chance of adverse behavioural changes. Such tools include the introduction of liquidity saving mechanisms, the strengthening of throughput rules, payment tariffs that vary through the day, setting central bank collateral eligibility criteria for intraday liquidity and regulatory 'deep dive' assessments.

Introduction

The provision of payment services is one of the key functions of the financial system.(1) Therefore, as part of its broader financial stability remit, the Bank of England is concerned that banks adequately manage the intraday risks that arise as part of their obligations to provide payment services to the wider economy. To the same end, the Bank is also concerned that the United Kingdom's large-value payment and settlement systems function smoothly and are resilient.

This paper aims to outline how intraday liquidity risks arise in large-value payment systems and to explain how system participants currently manage these risks. It then describes the policies that authorities in the United Kingdom and abroad are implementing to strengthen the resilience of payment systems to intraday liquidity stress. These actions on intraday liquidity risk will likely increase the cost of system participation, which may change the incentive structure for system participants. We present a series of policy tools that authorities could use to mitigate any potential unintended behavioural changes that could arise as a result of forthcoming liquidity regulation.

The large-value payment and settlement systems, CREST and CHAPS, play a vital role in the United Kingdom's financial system. On average, in 2010, over £730 billion of transactions were settled every day across the two systems. This equates to the United Kingdom's 2010 nominal GDP settled every two days.(2)

The introduction of real-time gross settlement (RTGS) into CHAPS in 1996 eliminated settlement risk between participant settlement banks as payments are settled finally and irrevocably, individually and in real time.(3) However, participation in RTGS systems, directly or indirectly, requires that banks have access to liquidity intraday. The provision of this creates liquidity exposures between banks. Banks send payments but expect to receive the funds back later in the day to meet other outgoing payment obligations. In other words, credit risk is exchanged for liquidity risk.

Banks typically manage their payment flows so that they end the day flat. That is, the value of incoming payments is roughly equal to the value of outgoing ones. This means that the liquidity exposures that arise from payment and settlement activity are created and extinguished within the working day. While banks actively manage these exposures they can nevertheless be very large: almost all banks in CHAPS regularly have intraday liquidity exposures in excess of £1 billion to individual counterparties.(4) For larger banks these exposures are regularly greater than £3 billion. In CREST bilateral liquidity exposures can exceed £6 billion (although these are typically secured with highly liquid assets).(5)

(1) See King (2009) and Bank of England (2010).
(2) Based on 2010 UK GDP at market prices of £1.456 trillion from ONS.
(3) In the Ernest Sykes Memorial Lecture in 1989 Governor Robin Leigh-Pemberton spoke of the risk advantages of payment systems which could reduce the duration of settlement counterparty exposures. RTGS eliminates these exposures entirely. This was followed by a discussion paper that explored the risks of the payment and settlement landscape (Bank of England (1989)).
(4) RTGS has ensured that these liquidity exposures are never translated into credit exposures.
(5) Source: Bank of England payments database.

Furthermore, settlement banks may extend uncollateralised intraday credit to their customers who access the system indirectly. Such members may be heavily reliant on their settlement bank to provide intraday liquidity. These credit exposures can also be very large. The Bank estimates that some individual intraday credit lines can be in excess of 10% of a settlement bank's core tier one capital. During a crisis, settlement banks may be reluctant to provide liquidity, given that this exposes them to credit risk, putting an extra burden on the indirect member.

Some observers have pointed to the general smooth functioning of the payment and settlement systems during the financial crisis as a sign that intraday liquidity risks are well-managed. However, intraday liquidity played an important role when individual institutions were experiencing liquidity stress. The widespread allowance of double duty (prudential asset buffers being used to support intraday liquidity needs) and authorities' reluctance to impose additional liquidity costs on banks has meant individual institutions were more exposed to liquidity stress. Since the crisis, authorities have begun to focus on intraday liquidity. In September 2008 the Basel Committee on Banking Supervision (BCBS) published 'Principles for Sound Liquidity Risk Management and Supervision' which provides guidance to banks and supervisors on liquidity risk. Principle 8 explicitly tackles banks' management of intraday liquidity risk. And the Basel III liquidity framework published in December 2010 says that the Basel Committee is currently reviewing how to tackle intraday liquidity risk.[1]

In the United Kingdom, the Financial Services Authority (FSA) has included intraday liquidity as a key risk driver in its new liquidity regime. The aim of the authorities is to ensure that banks are able to meet their payment and settlement obligations on a timely basis in both normal financial conditions and under stress. This will require banks to (a) have adequate systems and controls to manage intraday risks, (b) report data on intraday liquidity usage and risks and (c) have an adequate pool of high quality liquid assets to support intraday needs in both normal and stressed conditions.

While strengthening the management of intraday liquidity risk, these requirements will raise the opportunity cost of providing liquidity in the payment system. Previously, banks' regulatory liquid asset requirements were not calibrated to include intraday liquidity risk. However, these liquid assets were being used as collateral, intraday, in the large-value payment and settlement systems. In essence, the liquid asset buffer was being used to support intraday liquidity needs, but regulatory requirements were not calibrated to take into account these needs (a practice referred to as double duty). For UK settlement banks this meant that the cost of intraday liquidity was, in effect, negligible. Amongst other things, this has led to a situation where settlement banks willingly provide liquidity to the system as a whole, thereby allowing payments to be settled throughout the day with minimal delays.

The Bank supports the FSA's approach of including intraday liquidity risk as a key risk driver in their wider liquidity regulations and so making banks hold liquid assets to support their intraday liquidity needs over and above those for balance sheet resilience. Banks are currently required to calculate their intraday liquidity needs but only to hold a proportion of the requisite assets. The intention is that the buffers will gradually increase to cover all elements of liquidity risk. This will increase the cost of providing intraday liquidity which may, in turn, affect the willingness of banks to provide liquidity in the UK payment and settlement systems. As such it warrants careful consideration.

The paper is organised as follows: Section 1 explains why intraday liquidity risk matters, and how it is currently managed. Section 2 outlines various proposals for managing intraday liquidity risk and discusses their key features: systems and controls, monitoring and buffer requirements. Section 3 discusses the risks posed by double duty and how intraday liquidity regulations can address them. Section 4 then considers how intraday liquidity regulation might change incentives to provide liquidity in the payment system. In light of this discussion, Section 5 offers some thoughts on how to ensure that implementing intraday regulation can be best managed so as to minimise the risk of disruption to the payment system. These include the use of liquidity savings mechanisms and throughput rules. Section 6 concludes.

1 Why intraday liquidity risk matters

1.1 Intraday liquidity and payment system risks

Most large value payment systems around the world settle payments using RTGS. This means that payments are settled individually on a gross basis requiring that banks have sufficient liquidity on their settlement accounts for each payment. Whilst banks can re-use liquidity from incoming payments to fund outgoing payments, timing mismatches between incoming and outgoing payments can lead to significant liquidity needs.[2] Moreover, these intraday positions can be much larger than banks' end of day net positions.

Settlement banks meet these intraday liquidity needs using their own funds. These can be in the form of reserve account balances, intraday borrowing secured with eligible collateral from the central bank or borrowing in the interbank money markets.[3] Indirect participants, who access systems via a

(1) Basel Committee on Banking Supervision (2010).
(2) See Manning *et al* (2009) for a more in-depth discussion of the economics of large-value payment systems.
(3) Settlement banks may use other sources of funding for their intraday liquidity needs such as committed and uncommitted intraday credit lines but it is not thought that these are significant.

correspondent bank, need to have cash pre-placed with their settlement bank or have access to an intraday credit line, which may or may not be secured.(1) For further explanation of how intraday liquidity needs arise and are met in an RTGS payment system see the Annex.

Typically, payments sent in large-value payment systems are not intraday time critical: they do not have to be settled by a particular time during the day. However, there are some exceptions, such as CLS pay-ins,(2) CCP margin payments(3) and retail system clearings, which have to be settled by a specific point in the day.(4) Anecdotal evidence suggests that settlement banks in the United Kingdom consider approximately 4% of total payments by value and 5% by volume to be time critical. Settlement banks typically manage their payments by first releasing small and time-critical payments, before submitting their larger payments. They report that a large number of payments are known in advance (eg as a result of securities transactions), which facilitates the timing of payments and aids liquidity management. Settlement banks may also instruct their large clients to either pre-position funds or give advance notice of large payments that need to be made on their behalf.

In other words, through careful planning of payment flows, a settlement bank can reduce the need for intraday liquidity. This, however, is not without costs. As settlement banks aim to match incoming and outgoing payments to reduce their intraday liquidity usage, they may find it necessary to delay payments to other settlement banks. In the extreme, banks may adopt a 'receipt-reactive' strategy, only making outgoing payments using liquidity received from incoming payments.(5)

Liquidity provision in a payment system has much in common with the Volunteer's Dilemma.(6) This is a game in which everyone benefits from someone else contributing. But every participant would prefer to free ride on the contributions from others. Liquidity provision is similar; once one bank has made a payment, other banks can use that liquidity to fund their outgoing payments at no cost to themselves. In this type of game the strategic solution can result in the underprovision of liquidity.(7)

However, participation in a payment system is a repeated game. It is generally believed that repeated games tend to result in co-operative solutions.(8) The relatively smooth flow of payments throughout the day in CHAPS can be viewed as evidence of long-term co-operative behaviour. If a bank unilaterally delays payments, it will indeed reduce its own intraday liquidity needs (see Box 3). But this means that other settlement banks will receive payments later, and in turn, will either have to wait for matching funds or use their own liquidity to fund outgoing payments. Hence, delay by one settlement bank can lead to delays by all other settlement banks and reduced turnover in the payment system as a whole.

There are further costs associated with delaying payments. Individual settlement banks may incur financial penalties associated with not settling time-critical payments on time; delaying customer payments may adversely affect the relationship between settlement bank and client; and making payments later may increase a bank's vulnerability to operational risk, in turn raising operational risk in the system as a whole.

Delays of payments until the end of the day increase the likely impact of an operational problem, in particular the potential for payments to remain unsettled and cause liquidity dislocation. For example, consider a bank that suffers a system failure so it is unable to make payments for the rest of the day. The impact of this operational issue will depend upon the value and volume of payments that are unsettled at the time of the failure and the net liquidity position that the bank has with respect to the other banks in the system. Both are likely to be higher if a bank has delayed its payments (ie it is holding more of other banks' liquidity). The impact of the operational stress is therefore greater.(9)

Furthermore, if payments relate to the settlement of a credit exposure, then delay can increase the amount of credit risk in the financial system.

As a result, each settlement bank faces a trade-off between the cost of liquidity on the one hand and the cost of delaying payments on the other. Moreover, the cost of delay is a cost to both the individual bank and to the system as a whole, and hence can be thought of as an externality. Economists have investigated the incentives that banks have to delay payments. Bech and Soramäki (2002) use simulations to show the existence of the trade-off in a pure RTGS system. Galbiati and Soramäki (2008) show that the demand for costly liquidity increases as the cost of delay increases. In other words, banks make payments earlier when it becomes more expensive to delay. Bech and Garratt (2003) use a game-theoretic framework to investigate how bank behaviour depends on the intraday credit policy of the central bank. They conclude that

(1) A correspondent bank is a direct participant of a payment or settlement system that provides access to the system to other banks, making payments on their behalf.
(2) CLS is a multi-currency cash settlement system. It requires participant banks to make payments to it, 'pay-ins', five times a day. It then uses this liquidity to settle transactions and make 'pay-outs' to participants. The failure to make a CLS pay-in on time is seen by banks to have serious consequences.
(3) Central counterparties (CCPs) take margin to ensure members' performance on potential obligations to it or cover market movements on unsettled transactions. If a member fails to make a margin payment on time then the CCP may declare that member to be in default and close out the member's outstanding positions.
(4) The CPSS report 'New developments in large-value payment systems' contains a discussion of time criticality.
(5) Johnson, McAndrews and Soramäki (2004) suggests receipt-reactive gross settlement as a liquidity saving device.
(6) See Diekmann (1985).
(7) Weesie and Franzen (1998) show that cost sharing can increase the probability that a public good is produced. However, there remains a non-zero probability that the public good is not produced. In the context of a payment system, we interpret this to imply that strategic behaviour by banks can lead to the underprovision of liquidity.
(8) For example, see Friedman (1971).
(9) See Bech (2008).

in collateralised and fee-based regimes both early settlement and delay are possible equilibria. However, in a fee-based system (where the central bank charges a fee for the use of intraday liquidity) the outcome also depends upon the probability of counterparties receiving a payment request. The result is that in some circumstances delay may be socially efficient.

In practice, the liquidity/delay trade-off differs from country to country. For example, in the Fedwire system banks pay a charge for unsecured borrowing intraday from the Federal Reserve based upon the value and time for which they borrowed.(1) In contrast, in the United Kingdom, intraday liquidity is provided free of charge, but is secured against eligible collateral. The cost of this intraday borrowing is the opportunity cost of holding such collateral. Prior to the FSA's new liquidity regime, the liquid asset buffers of UK banks were not calibrated to include intraday liquidity risk. Yet, it was this asset pool that is used to support intraday liquidity needs in the payment and settlement systems. Hence, as assets were not held specifically for intraday risk, for UK banks the intraday liquidity cost of participation was effectively zero. This is slowly changing as the FSA implements its new liquidity rules.

Fedwire and CHAPS exhibit very different payment timing profiles. There may be a number of reasons why this is the case. Armentier *et al* (2008) and Becher *et al* (2008) both suggest that the policies that govern the provision of daylight overdrafts, and therefore the cost of liquidity, are significant factors.(2) **Chart 1** shows that in Fedwire payments tend to be bunched towards the end of the day, whereas **Chart 2** shows that payments in CHAPS are distributed more evenly throughout the day.

Chart 1 Timing of payments in Fedwire

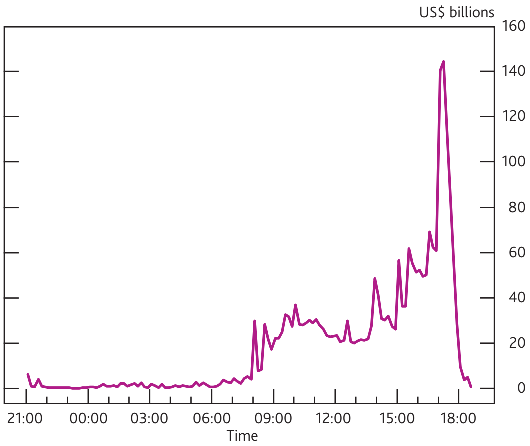

Source: Federal Reserve Bank of New York.

To conclude, for the payment system to operate smoothly, banks need to be willing and able to make payments in a timely manner throughout the entire business day. Doing so reduces operational risk in the system. But it also reduces

Chart 2 Timing of payments in CHAPS

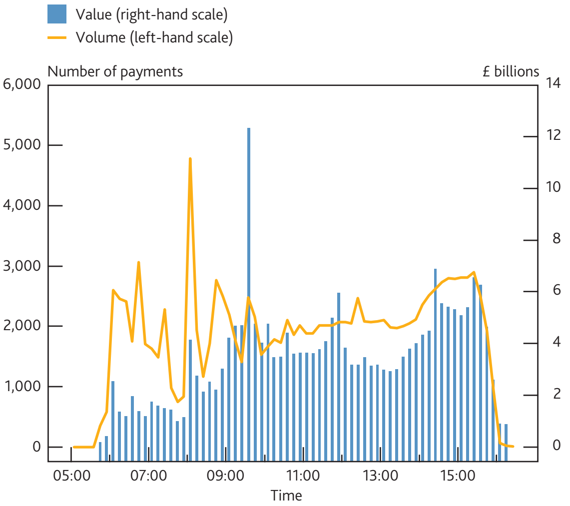

Source: Bank of England.

individual settlement banks' exposure to liquidity risk. This is explained in more detail in the next sub-section.

1.2 Intraday liquidity risk and stress scenarios

So far, we have focused on the system impact of payments being delayed. But the intraday element of the FSA's liquidity regulation is primarily concerned with ensuring that individual firms can fulfil their obligations in the payment and settlement systems. As explained in Section 1.1 intraday liquidity needs arise due to the timing mismatches between incoming and outgoing payments. Both settlement banks and their clients are vulnerable to any stress that causes this mismatch to increase, resulting in an increase in the amount of intraday liquidity required to continue making payments in a timely fashion.

There are four distinct categories of stress scenario that capture the main intraday liquidity risks in a payment or settlement system. The first three involve changes in banks' payment profiles, which in turn affect their intraday liquidity needs. The fourth involves a change in the value of assets that a bank has to support its intraday liquidity needs. The four types of stress are:

(i) a credit or liquidity shock affecting the **bank directly**, reducing counterparties' willingness to make payments to it in a timely fashion;

(ii) an operational, credit or liquidity shock affecting the ability of a **major counterparty** in the payment system to make payments to the settlement bank as expected;

(1) The Federal Reserve have announced that they intend to change their intraday liquidity policy, increasing the cost of uncollateralised daylight overdrafts and providing collateralised daylight overdrafts at zero fee. See www.federalreserve.gov/newsevents/press/other/20100930a.htm.

(2) There may also be institutional reasons why the payment profile in CHAPS and Fedwire differs. The cost of liquidity is only a partial explanation.

(iii) a credit or liquidity shock affecting a **major customer** or group of customers of the settlement bank, preventing them from receiving payments as expected; and
(iv) **market conditions change** which mean that a given pool of assets generates less intraday liquidity.

Indirect participants may experience the stresses above, but they may also be vulnerable to different manifestations of those stresses. If an indirect participant is suffering a stress, this could lead its correspondent bank to cut unsecured intraday credit lines and request prefunding or collateralisation. An indirect participant may be less vulnerable to a customer stress, but could be vulnerable to a stress afflicting its correspondent bank meaning that it is unable to submit payments in a timely fashion.

Below, we briefly explain how these stresses may manifest themselves in the payment and settlement systems.

Bank stress

Banks rely upon incoming funds from counterparties to fund outgoing payments. This recycling of funds allows intraday liquidity needs to be substantially less than the gross value of payments. If a bank is heavily reliant on incoming payments then it may be vulnerable to a liquidity stress should its counterparties decide to delay or stop making payments to it. This could occur:

- if the bank suffered, or was perceived to be suffering from, a liquidity or credit shock that made other banks in the payment system question its viability; and
- if other banks responded by demanding that they receive payments before they are prepared to send any.

The effect of this is for the stricken bank's incoming payments to be delayed and the timing mismatch between incoming and outgoing payments to increase. This would in turn lead to an increase in demand for intraday liquidity by the stricken bank.

Counterparty stress

During the day, settlement banks build up net debit and net credit positions with respect to their counterparties. These individual bilateral positions can leave settlement banks vulnerable to a stress that impacts their counterparties' ability to send payments. The following would constitute a counterparty stress scenario:

- if, at a given point during the day, a settlement bank has a large net debit position with respect to another settlement bank, ie it has sent more payments to that bank than it has received; and
- if that counterparty were unable to send payments.

The settlement bank would be short of the liquidity that it was expecting to receive. As such, even if the bank immediately stops making payments to the stricken counterparty, it will need to use more of its own liquidity to make its payments to the other banks in the system.

Customer or correspondent stress

Banks that offer correspondent banking services normally advance intraday credit to their financial institution customers, for which they may or may not charge. As a result, customer payments are typically made using the settlement bank's liquidity. This leaves the settlement bank vulnerable to any shock that hits one or more of its large customers. This could occur:

- if the settlement bank sent out payments on behalf of its customers expecting that counterparties would send payments for the customer in return;
- the customer bank suffered from, or was perceived to be suffering from, liquidity or credit issues; and
- other banks responded by delaying payments to the stricken bank.

This ultimately hits the settlement bank's liquidity position, as it does not receive payments that it was expecting.

The settlement bank then has two choices. It can either choose to stop making payments on behalf of its customer, which could potentially exacerbate the stress and increase the probability of default, or it can continue to make payments for its customer, but at the cost of using more intraday liquidity.

A bank may also find that when a customer is facing a stress there is an increase in the volume and value of payments that it needs to make. Again the settlement bank is faced with the choice of delaying payments or using more intraday liquidity to meet its customer's demands.

An indirect system participant may be vulnerable if its correspondent bank is experiencing problems. If the correspondent is suffering from liquidity problems, it may be unable to extend the same amount of unsecured intraday credit as normal. This could create an extra liquidity burden for the indirect participant or could result in payment delays.

Correspondent banking relationships also expose banks to credit risk.(1) Intraday credit lines that settlement banks extend to their customers are typically uncollateralised. This exposes the settlement bank to the risk of its customer defaulting intraday. Should a default occur after the settlement bank has sent payments on its customer's behalf, the settlement bank could find itself liable for the full value of the intraday credit line that it has extended. Conversely, some customer banks will systematically have credit balances on account with their correspondent. These credit balances are typically unsecured

(1) Harrison *et al* (2005) provide a more detailed explanation of the risk implications of correspondent banking.

and can be seen as a form of intraday lending from the customer to the correspondent bank. This exposes the customer bank to the potential default of its correspondent.

Customer stress in a securities settlement system

In a securities settlement system, such as CREST, there is also potential for customer stresses. When a settlement bank settles a securities transaction it can do so using its own liquidity or liquidity generated by one of its customers.(1) If a bank systematically relies upon the liquidity provided by a customer in order to fund its intraday liquidity needs then it could be highly vulnerable to shocks which affect that customer. For example, the customer bank could fail and so no longer participate in the system. Alternatively, the customer bank may suffer operational issues that mean it is unable to generate liquidity either by selling securities or via self-collateralising repo. Both of these shocks would require that the settlement bank raise extra liquidity in order to fulfil its normal payment and settlement obligations.

Market stress

A bank may set aside an amount of assets that it considers to be sufficient to cover its intraday liquidity needs in both normal and stressed conditions. However, during a market-wide stress the credit and liquidity quality of a particular asset could fall. This could lead to an increase in the haircut applied or to the asset becoming ineligible for central bank liquidity facilities. For indirect participants, variation in an asset's value could lead its correspondent bank to apply higher haircuts. In either case, the bank would have less intraday liquidity than it expected.

Contagion

Each of these stresses may cause the settlement bank to require more intraday liquidity in order to make its payments in a timely fashion. If the bank did not have access to this liquidity, it would have no choice but to delay payments. In addition to exacerbating perceptions of weakness, delaying payments would subsequently cause the bank's counterparties to require more liquidity to make their payments. If other banks also had insufficient liquidity to make payments, then they too would have no choice but to delay. This could result in substantial delays to payment flows and even potential gridlock. A simple illustrative example is given in Box 1.

1.3 How are these risks managed at present?

Settlement banks have developed a range of different tools to manage these intraday liquidity risks. Many settlement banks use internal schedulers with functionality that enables them to manage their payment flows smoothly throughout the day. These schedulers allow banks to limit the impact of intraday liquidity stress and respond quickly should a stress occur. For example, schedulers may allow banks to stop making payments to a particular counterparty, or take manual control of large value payments, releasing them individually.

A common feature of settlement banks' intraday liquidity management tools is the use of bilateral limits. These limits are used to manage the risk of a counterparty being unable or unwilling to make payments. To do so, banks place a cap on the amount of liquidity that they are prepared to provide to any individual counterparty, ie they limit the largest bilateral net debit position. These limits vary across counterparties and throughout the day. They are typically operational limits, and are managed dynamically during the day.

The CHAPS system rules provide further risk mitigation in the form of throughput guidelines. The rules state that a bank must send, on average over the month, at least 50% of payments by value by noon and 75% by 14:30. This helps to limit the liquidity exposures that banks have with each other and acts as a co-ordinating device which reduces the liquidity needs of the system as a whole. It also acts to limit the scope for any bank to free-ride on the liquidity provision of other banks.

Settlement banks that make payments on behalf of customers need to manage the associated liquidity and credit risk. They can do this by limiting the amount of intraday credit that they extend to their customers. Once a customer has used all of its credit, the settlement bank will stop making payments on its behalf until it receives incoming payments. These limits are typically unsecured and unadvised (the client is unaware of its intraday limit). This gives settlement banks the freedom to cut the credit that they are prepared to extend if they believe there is an increased chance of default. Some settlement banks may also ask their customers to collateralise their intraday credit lines, protecting them against the risk of an intraday default.

Beyond these risk controls there are other, less formal, mechanisms which allow banks to manage intraday liquidity risk effectively. Firstly, in the United Kingdom the small community of settlement banks helps to create a co-operative system in which banks can co-ordinate payment flows. This is especially useful in a crisis when liquidity managers and payment system operators can manage liquidity risk by direct communication with each other. Secondly, many of the system participants have built up years of experience managing payment flows. This enables them to identify unusual behaviour quickly and respond accordingly, thus limiting the impact of intraday liquidity stress should it occur.

The design of payment and settlement systems can help to reduce intraday liquidity needs and aid intraday liquidity management. For example, a number of RTGS payment

(1) In CREST, liquidity can be generated in a number of ways. The primary mechanism is the self-collateralising repo (SCR). When a security that is eligible for the Bank of England's intraday liquidity facilities is purchased it is automatically repoed to the Bank in return for intraday liquidity to the value of the security minus a haircut. The SCR mechanism allows the liquidity generated by the SCR to be simultaneously used to fund the purchase of the security.

Box 1
Stress scenarios

Intraday liquidity stresses

This box contains a stylised example of the different types of intraday liquidity stresses that a bank could be exposed to and which may cause an increase in its intraday liquidity needs. It uses a single day of real payment data taken from the RTGS system for payments between three banks, labelled A, B and C. Each bank has obligations to make payments to each of the other banks. Payments are made using a combination of incoming payments and intraday liquidity in the form of central bank reserves.

The data have been slightly adjusted so that each bank sends the same value as it receives. This box is meant to be illustrative of how these risks can materialise and not to give a quantitative assessment of how large they are.

Chart A shows how Bank A's payment profile and intraday liquidity needs evolve during the payment system day. The blue line shows the cumulative amount that Bank A has sent to Banks B and C since the start of the day. The pink line shows the cumulative value of payments that Bank A has received from Banks B and C since the start of the day. The orange line shows Bank A's net position.

Chart A Bank A's payments sent, received and net position

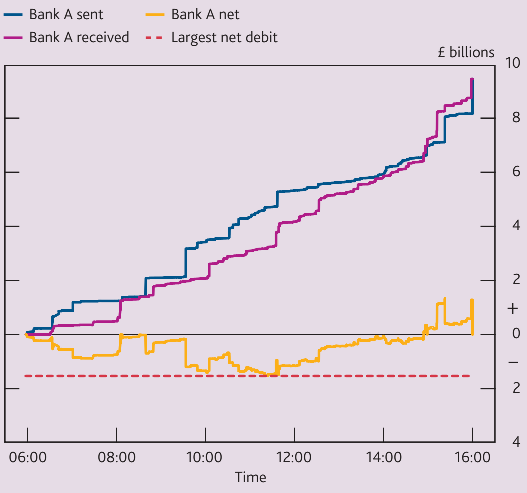

Sources: Bank of England and authors' calculations.

When the blue line is above the pink line it indicates that Bank A has sent more payments than it has received and so has used its own liquidity to fund its payments. Conversely, when the blue line is below the pink line it indicates that Bank A has received more payments than it has sent.

Bank A's intraday liquidity needs are equal to the largest net debit position that it incurred during the day. This is the minimum amount of central bank reserves needed in order to make all its payments without delay. In this example, Bank A's largest net debit position, and thus liquidity need, was £1.53 billion. This is indicated by the dashed red line.

Bank stress

In our example, Bank A is the stricken bank. Banks B and C both delay all their payments to Bank A by 30 minutes. Bank A makes payments at the same rate in order to convince the other banks that it is capable of continued participation in the system.

Chart B shows how Bank A's payment activity and liquidity needs evolve across the day. The delay in payments to Bank A is shown by the rightward shift of the pink, payments received line. These delays lead to an increase in the largest net debit position, and thus liquidity needs. As a result of this stress, Bank A's intraday liquidity needs increase by 40%, from £1.53 billion to £2.15 billion.

Chart B Bank A's payments sent, received and net position after it suffers a bank stress

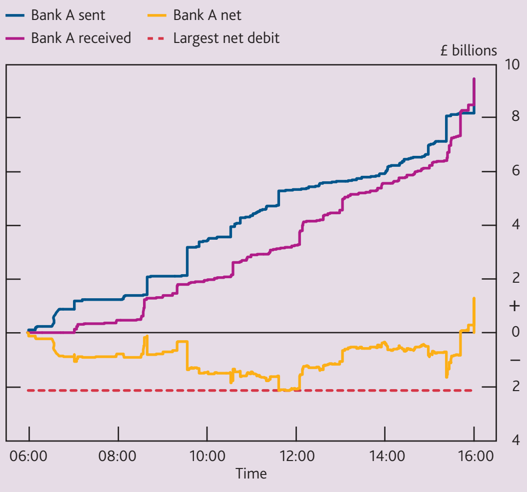

Sources: Bank of England and authors' calculations.

Counterparty stress

In our example, Bank C suffers a stress at around 9.30 am. This prevents Bank C from sending any payments to Bank A until the end of the day. Bank A responds immediately by stopping making payments to Bank C, again until the end of the day. Bank A continues to send and receive its payments to and from Bank B as normal.

Chart C shows how Bank A's liquidity needs evolve during the day. After the stress Bank A only sends and receives payments to Bank B. Bank A's largest net debit position, its intraday liquidity need, increases as a result of the liquidity stress. Bank A was relying on incoming payments from Bank C to fund some of its payments to Bank B. As these funds did not materialise, Bank A now requires more funds. This is indicated by the dashed red line which has increased 45% from £1.53 billion to £2.23 billion.

Chart C Bank A's payments sent, received and net position after Bank C suffers a stress

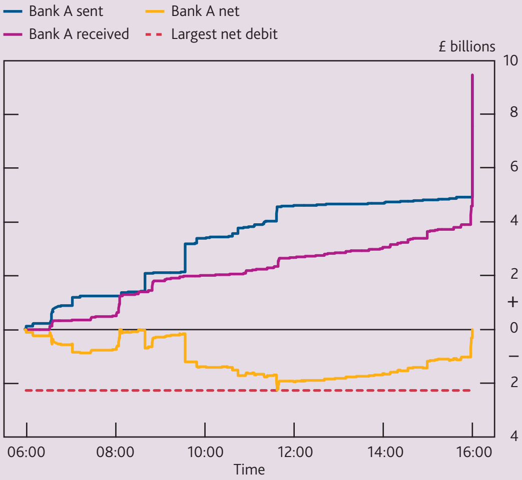

Sources: Bank of England and authors' calculations.

Customer stress

In this example one of Bank A's customers suffers a stress. This customer is responsible for 10% of the payments that Bank A sends and receives. Not having a direct relationship with Bank A's customers the other settlement banks delay all payments for 90 minutes whilst the customer continues to make payments at the normal rate. Bank A grants sufficient liquidity to allow it to do so. This stress begins at 9.30 am.

Chart D shows how Bank A's payment activity and liquidity needs evolve across the day. Compared to the base case, a proportion of the payments have been delayed for 90 minutes. This shifts the red payments received line to the right, resulting in an increase in the timing mismatch between incoming and outgoing payments. Bank A's intraday liquidity needs increase by 6% from £1.53 billion to £1.62 billion.

Chart D Bank A's payments sent, received and net position after one of its customers suffers a stress

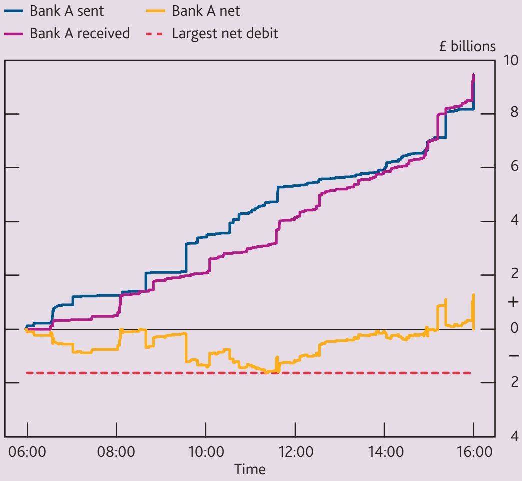

Sources: Bank of England and authors' calculations.

systems have introduced liquidity saving mechanisms such as an offsetting algorithm (see Section 5.1) to improve the liquidity efficiency of the system.

Together, these tools helped the payment and settlement systems function smoothly during the worst days of the recent financial crisis. These tools are, however, not perfect risk mitigants and could introduce another set of risks, for example reliance on individuals' experience to manage payment flows. They also do not tackle the risks posed to the customer banks. Prior to the recent financial crisis, regulators did not focus on intraday liquidity risk and there were no standard monitoring or liquidity management measures in place. In the next sections we explain how new intraday liquidity regulations in the United Kingdom and abroad are expected to meet some of these shortcomings.

2 Intraday liquidity regulation

2.1 Intraday liquidity regulation

In the 'Principles for Sound Liquidity Risk Management and Supervision', the Basel Committee on Banking Supervision has identified the issue of the management of intraday liquidity risk and collateral. In particular, principle 8 states that: 'A bank should actively manage its intraday liquidity positions and risks to meet payment and settlement obligations on a timely basis under both normal and stressed conditions and thus contribute to the smooth functioning of the payment and settlement systems.' Beyond this there are references to intraday liquidity management in a number of other principles.[(1)]

The Principles advocate an approach that places the emphasis on individual banks to have adequate systems in place to measure and manage their intraday liquidity risk. The role of the supervisor is to ensure that banks adhere to these Principles. The Principles highlight the importance of intraday liquidity management both due to its impact on the bank itself and the impact on the bank's counterparties.

In the United Kingdom, the FSA's regulatory approach is consistent with that advocated in the Principles. The responsibility for ensuring that systems and controls are adequate to measure and manage intraday liquidity risks rests with the individual bank and the FSA's role is to ensure that the systems and controls are robust. In addition, however, the FSA requires banks to calculate how much intraday liquidity they would need in both normal and stressed conditions. In time, banks will be required to hold sufficient liquid assets to cover these needs. This will be in addition to the assets that banks will be required to hold to meet their balance sheet liquidity buffer (hereafter 'prudential asset buffer') requirements. This section describes the two main approaches available to supervisors to ensure that banks measure and manage their intraday liquidity risk — monitoring tools and buffer requirements — and discusses their costs and benefits.

2.2 Systems and controls and monitoring requirements

The 'Principles for Sound Liquidity Risk Management and Supervision' state that banks should:

i) have the capacity to measure expected inflows and outflows, anticipate the intraday timing and forecast the range of potential net funding shortfalls that might arise during the day;
ii) have the capacity to monitor intraday liquidity positions against expected activities and available resources;
iii) arrange to acquire sufficient intraday funding to meet their intraday objectives;
iv) have the ability to manage and mobilise collateral as necessary to obtain intraday funds;
v) have a robust capability to manage the timing of their liquidity outflows; and
vi) be prepared to deal with unexpected disruptions to their intraday flows.

The supervisor would be expected to monitor that banks adhere to these principles. And there are plans to introduce internationally agreed monitoring metrics for intraday risk in due course.

The benefits of introducing monitoring requirements are threefold. They encourage individual banks to take a more systematic and disciplined approach to intraday liquidity management; they allow greater transparency to the supervisor and payment system overseer; and they permit the supervisor to carry out cross-sectional analyses of banks' resilience to intraday liquidity shocks as well as monitoring how a bank's resilience evolves over time.

2.3 Buffer requirements

An intraday liquidity buffer aims to ensure that banks always have sufficient intraday liquidity to fund their intraday positions, and that intraday liquidity needs do not reduce banks' resilience to other liquidity shocks.

The FSA's new liquidity regime will require banks to hold a buffer of highly liquid assets calibrated to ensure that they have sufficient intraday liquidity to make payments in a timely manner in both normal and stressed conditions. Intraday liquidity needs will be calculated in addition to the liquid assets needed by a bank for balance sheet resilience purposes.[(2)] Intraday liquidity risk is considered one of the key risk drivers that determine a bank's liquid asset buffer requirement.[(3)]

(1) For example, Principle 9 on collateral says that 'systems should be capable of monitoring shifts between intraday and overnight or term collateral usage' and Principle 10 on stress testing says that 'tests should consider the implication of the scenarios across different time horizons, including on an intraday basis'.
(2) The Basel Committee on Banking Supervision has agreed international standards for calculating the liquid asset buffers that banks will be required to hold for balance sheet resilience purposes. The standards relate to two ratios: a Liquidity Coverage Ratio and a Net Stable Funding Ratio. See Basel Committee of Banking Supervision (2010).
(3) See FSA handbook BIPRU 12.5.14 R, http://fsahandbook.info/FSA/html/handbook/BIPRU/12/5.

The principle for a settlement bank's intraday liquidity buffer calculation is explained in the FSA handbook.(1) Banks are required to calculate how much intraday liquidity they need in normal times to fund their participation in payment and settlement systems and then estimate how this could change under liquidity stress. The stresses referred to are an unforeseen, name-specific liquidity stress lasting for two weeks and an unforeseen, market-wide liquidity stress lasting three months.

The handbook also says that banks should take into consideration both their own liquidity needs and those of customers. And should consider the impact of other participants withholding some or all of their payments and customers increasing the value and volume of payments.

Banks are required to estimate the size of their own buffers as part of the Individual Liquidity Adequacy Assessment (ILAA) process. A bank's ILAA is then reviewed by the FSA who will come to a conclusion about the appropriate level of liquidity protection.

3 Intraday liquidity regulation and double duty

This section explains the practice of double duty and why it may bring risks to participants in payment and settlement systems (Section 3.1). It then considers what policy actions supervisors can take to limit these risks and examine how these actions need to be accompanied by measures to mitigate any possible adverse effects on payment flows (Section 3.2). Box 2 uses the Lehman Brothers bankruptcy as an example of the role double duty can have in a liquidity crisis.

3.1 Double duty

Central banks and supervisors encouraged the introduction of RTGS into payment systems to minimise settlement risk for high-value payments.(2) But they recognised that there was a significant liquidity cost to using it. In RTGS systems, timing mismatches between outgoing and incoming payments lead banks to use liquidity intraday. This intraday liquidity need can be many times the bank's end of day liquidity need. To reduce this cost, many authorities decided to allow banks to use their prudential asset buffers, held for wider liquidity resilience, to support payments activity intraday without making banks include intraday liquidity risk in the calibration of the buffer: a practice known as 'double duty'.

Double duty reduces the cost of making payments in RTGS. In fact, if a bank's intraday liquidity usage is less than the amount of liquid assets it is required to hold by its supervisor, then the opportunity cost is zero. This reduces the incentives for banks to delay payments and thus facilitates the smooth functioning of payment and settlement systems.

However, the practice of double duty is not without risks. Conceptually, there is one significant risk associated with using the same assets for two separate purposes: when the assets are being used for purpose A they are not available for purpose B and *vice versa*. In practice this manifests itself in two ways:

(a) *Balance sheet resilience risk*: if a bank's prudential asset buffer is serving the purpose of providing intraday liquidity, then it cannot be as effective as a buffer against a run on deposits.
(b) *Intraday liquidity risk*: if a bank suffers a prolonged balance sheet liquidity stress, this uses up the bank's prudential asset buffer meaning the bank has insufficient funds available to operate effectively in the payment systems.

Balance sheet resilience risk

If a bank suffers from a balance sheet stress it will need to liquidate its prudential asset buffer in order to pay its wholesale and retail depositors. This means the bank will need to sell or encumber some assets to generate cash, which it then pays out through the payment systems. If the assets that the bank had earmarked to support this stress have been used intraday, then the bank may not have immediate access to them and so be unable to make timely payments to its depositors. Thus, the assets are not always available for the 'run' they were designed to protect against. If the bank delays payments to its depositors then other market participants may question the bank's ability to pay. This could further exacerbate the balance sheet stress.

This is the scenario that Lehman Brothers faced in the run-up to its bankruptcy. During this period, Lehman's correspondent banks in various systems demanded that it collateralise or prefund its intraday liquidity usage. These assets were taken from the bank's liquid asset pool and so were no longer available to pay out its depositors. Despite appearing to be liquid, by 15 September 2008 Lehman had insufficient unencumbered liquid assets to support its projected cash outflow. Box 2 outlines in more detail the role intraday liquidity risk played in Lehman's collapse.

In normal times, or at the beginning of a period of liquidity stress, the probability of this risk crystallising is relatively low. Banks' prudential asset buffers are typically many times larger than their intraday liquidity needs. However, as a balance sheet stress progresses, the prudential asset buffer will shrink, increasing the probability of this risk materialising. As such, as in the case of Lehman Brothers, a bank's intraday liquidity needs may increase and take up an ever increasing proportion of the bank's liquid assets.

(1) www.fsa.gov.uk/pages/handbook.
(2) See Committee on Payment and Settlement Systems (1997).

Box 2

Intraday liquidity risk: the case of Lehman Brothers

On 15 September 2008, Lehman Brothers Holdings Inc. filed for Chapter 11 bankruptcy.(1) Evidence presented in the Valukas report suggests that a lack of liquidity turned out to be one of the key elements that precipitated the firm's collapse: on 14 September 2008, Lehman no longer had sufficient available liquidity to fund its daily operations. A significant factor in Lehman's shortage of liquidity was the provision of collateral to its clearing banks to fund its intraday liquidity positions. Lehman was using an increasing portion of the asset pool to cover intraday liquidity risk. Ultimately this meant Lehman's pool of liquid assets was unable to meet the outflows it was calibrated for. This is precisely the risk posed by double duty.

Lehman's correspondent banking relationships

Across the world Lehman used a number of different correspondent banks to provide it with payment and settlement services. In the United States, the most significant was JP Morgan who acted as their agent in tri-party repo arrangements.(2) Citibank provided Lehman with access to CLS, the multicurrency foreign exchange settlement system. HSBC provided Lehman with settlement services in the United Kingdom, acting as their settlement bank in CREST. Lehman also had correspondent banking relationships with Bank of America, Bank of New York Mellon and Standard Bank. In each case, Lehman's correspondent provided it with intraday credit to allow it to settle trades and make payments.

Lehman's liquidity pool

Lehman reported that it had a liquidity pool 'primarily intended to cover expected cash outflows for twelve months in a stressed liquidity environment.' The pool was designed to cover a range of cash outflows, while assuming that the firm could not issue any unsecured debt or generate liquidity outside that pool of assets (eg by selling less liquid assets). Lehman reported that at the end of the first quarter of 2008, the liquidity pool consisted of US$34 billion. At the end of the second quarter it was US$45 billion and at the end of the third quarter US$42 billion.(3) The size of the asset pool gave market participants comfort that Lehman had low vulnerability to liquidity risk.(4) But this asset pool increasingly contained assets that Lehman had pledged or transferred to its correspondent banks intraday. Although these assets were viewed as unencumbered overnight, a failure to pledge these assets at the start of the next day would have likely had negative implications for the willingness of these banks to clear for Lehman. In other words, the assets were not available to meet other outflows.

Lehman's correspondent banks request collateralisation and prefunding

In the run-up to Lehman's bankruptcy, its correspondent banks became less willing to advance intraday credit to allow Lehman to participate in payment and settlement systems. Throughout 2008 Lehman's correspondents started to apply a number of different conditions to reduce their exposure to Lehman. At first, haircuts were increased on tri-party repo collateral to protect the correspondent bank. This was followed by banks asking for explicit collateralisation of positions. When market conditions deteriorated, Lehman's collateral was revalued and it was required to post extra collateral. Some banks preferred to ask Lehman for cash deposits to prefund their intraday liquidity needs.

In total, by 12 September 2008, Lehman's correspondent banks had received over US$16 billion in collateralisation and prefunding.

Lehman runs out of liquidity

In Lehman's last days its liquid asset pool contained more and more assets that had been pledged intraday to its correspondent banks. According to Lehman's own post-mortem analysis, its liquid asset pool decreased from US$41 billion prior to its earnings announcement on 9 September 2008 to about US$25 billion on 12 September. The stressed market conditions of the time meant that approximately US$7 billion of these assets were not immediately liquid. Of the remaining US$18 billion of assets, US$16 billion were required by Lehman's correspondent banks. This left only US$2 billion of truly liquid assets available to meet outflows. On 15 September, Lehman was forecast to have a net cash outflow of US$4.5 billion. With insufficient liquid assets to cover this shortfall Lehman filed for Chapter 11 bankruptcy protection.(5)

The risks posed by double duty finally came to bear. Lehman had not calibrated its liquid asset pool to include intraday liquidity risk, but in a crisis intraday liquidity turned out to be a significant drain on its asset pool. Lehman found that it was significantly less liquid than both it and the market believed it to be. It was this lack of liquidity that was the final straw that broke Lehman Brothers.

(1) This box draws heavily on Volume 4 of the Chapter 11 proceedings report prepared by the Examiner, Anton R. Valukas. The full report is available at: http://lehmanreport.jenner.com.
(2) A tripartу repo is a repo in which a third party (eg a custodian bank, a clearing house or a CSD) is responsible for the management of collateral during the life of the transaction.
(3) Valukas, A (2010), page 1,409.
(4) Valukas, A (2010), page 1,415.
(5) Valukas, A (2010), page 1,454.

Intraday liquidity risk

Prudential asset buffers are typically calibrated to ensure that banks hold sufficient liquid assets to withstand a prolonged liquidity stress. For example, the BCBS' proposed Liquidity Coverage Ratio aims to ensure that a bank has 'an adequate level of unencumbered, high quality liquid assets that can be converted into cash to meet its liquidity needs for a 30-day calendar time horizon under a significantly severe liquidity stress scenario specified by supervisors'.(1)

As the 30-day balance sheet stress progresses, the bank's resilience to an intraday liquidity stress will fall with the size of the prudential asset buffer. Eventually the bank may have insufficient liquidity to be able to conduct its normal payments and settlement business in a timely manner and so be forced to delay payments to its counterparties. At the extreme the bank will have to behave 'receipt reactively' all day, relying entirely on incoming payments to fund its outgoing ones. In these circumstances, the bank's counterparties will be unlikely to be willing to send payments to it. The stricken bank may then be unable to continue to participate in the payment and settlement systems.

An alternative way to express this is to say that a bank needs to continue to operate in the payment and settlement systems to avoid default. If this requires the bank to use a fixed amount of collateral each day to generate intraday liquidity, it can only use the remainder of its collateral pool to meet outflows. Therefore, in practice, if that pool is sized on the basis of potential outflows from other sources and does not include an intraday component, then the bank will have less resilience to liquidity stresses than the supervisory authorities intended.

This argues in favour of authorities taking action to tackle the risks posed by double duty.

3.2 Intraday liquidity regulation and double duty

There are a number of ways that authorities can tackle the risks posed by double duty. The strictest way would be an outright ban on prudential assets being used intraday in the payment and settlement systems under any circumstances. However, this may not be desirable or, indeed, feasible. When a liquidity stress occurs, the payouts need to be made in the payment and settlement systems. It seems reasonable that a bank should use its prudential asset buffer to support this. After all this is why the bank has such a buffer. Furthermore, were a severe intraday liquidity stress to occur, it seems inconceivable that a supervisor would prefer that the bank protected its prudential asset buffer by failing settlement of a time critical payment rather than used those assets to fill a temporary liquidity shortfall.

One alternative approach is to allow the use of prudential assets only in stressed circumstances with the agreement of the bank's supervisor. However, it is not always straightforward to understand *ex ante* if a stress is an intraday or an overnight stress.

A further option, which is the essence of the FSA's intraday liquidity regulation, is for banks to calibrate their liquidity buffers to include intraday liquidity needs as a separate risk driver. This does not prohibit banks from using the rest of their prudential asset buffer intraday, as the portion assessed for each risk driver is fungible and available to meet liquidity requirements of any origin. However, a properly calibrated buffer would make it highly unlikely that banks would fail to meet a regulatory stress scenario that is based on liquidity outflows that are overnight or longer on account of needing to continue participation in the payment and settlement systems. Banks could still hold as much of their liquidity buffer in the payment systems as they required. If a bank habitually used its prudential asset buffer it would indicate that the buffer was insufficient and should be recalibrated. In effect, the FSA has decided to limit the risks of double duty described above without completely abolishing it.

The introduction of an intraday liquid asset buffer to deal with the risks of double duty increases the cost of intraday liquidity. The implications of this are discussed in Section 4.

4 Incentives created by intraday liquidity regulation

There are many benefits to introducing more rigorous controls around banks' monitoring and management of intraday liquidity risk. However, doing so may change the incentives for banks that participate in the payment and settlement systems. This could have a negative impact on system stability. When implementing intraday liquidity regulations, authorities need to pay particular attention to the incentives created and ensure that they have adequate tools available for dealing with any unintended systemic consequences.

Systems and monitoring

A focus on banks' internal intraday liquidity risk management systems is unlikely to create adverse incentives in themselves. If intraday liquidity monitoring extends to the use of defined metrics about liquidity usage, the risk of such adverse behavioural change is more likely.

In particular:

- if banks know exactly which dimensions the supervisor will use to assess them, then this may create incentives to alter payment flows to manipulate the metrics;
- if a bank is unsure how it compares to its peers it may attempt to manage its liquidity more tightly, even when unnecessary, for fear of attracting regulatory attention.

(1) Basel Committee on Banking Supervision (2010), page 3.

This could result in delayed payments and increased operational risk.

Buffers

Introducing intraday liquidity buffers should make banks more resilient to any potential liquidity stress. However, their introduction also makes intraday liquidity more costly. This is because of the opportunity cost of banks changing their portfolio allocations to hold more high-quality, low-return assets in the place of lower-quality, higher-return assets. Banks in a payment system may have arrived at an equilibrium that depends upon the collective trade-off between the cost of liquidity and the cost of delay. When that trade-off is changed to make liquidity more costly, banks may not be willing to provide liquidity at the same level. In particular, banks that believe they provide more than their fair share of liquidity in the near costless regime, may want to cut back their liquidity provision.

The increase in liquidity cost would be borne by most banks participating in a payment system, although not necessarily evenly. If the buffer is calibrated to take account of both a bank's normal intraday liquidity usage and its exposure to intraday liquidity risk then those banks that habitually use more liquidity or are more vulnerable to intraday liquidity stresses would be expected to hold a larger buffer.

It is reasonable to expect that banks will use historical data on payment activity to calibrate their buffers. Historical liquidity usage can give an indication of the amount of liquidity required to fulfil normal obligations. And 'what-if?' analyses of various stress scenarios can illustrate the potential impact of liquidity stresses on a bank's intraday liquidity needs.

However, basing the buffer on historical data creates incentives for banks to reduce their liquidity usage in order to benefit from a lower buffer come recalibration. A standard result in the payments economics literature is that costly intraday liquidity introduces incentives for banks to delay payments, waiting for incoming payments before sending outgoing ones.(1) It should be apparent, however, that if all banks try to delay payments then the result is that no-one benefits from reduced liquidity usage and there is an increase in systemic operational risk: an outcome that adversely affects everyone. Box 3 discusses the potential impact of banks delaying payments.

As explained in Section 1.1, there is evidence of co-operative behaviour in the United Kingdom payment and settlement systems. Indeed, this is the expected theoretical outcome in a repeated game. Therefore, it is likely that banks will find a co-operative equilibrium in a world with costly liquidity. But, it is not clear what that equilibrium will be and whether it will be optimal from a system stability perspective.

Furthermore, it is possible that the implementation of new regulation will cause disruption in the payment systems as banks delay payments in an attempt to reduce liquidity usage. Authorities have a role in both influencing the new equilibrium and ensuring that the transition happens smoothly. Section 5 discusses tools that could help guide this process.

5 Ensuring a 'good' equilibrium

A number of alternative policy and system-design measures could, in principle, be taken to mitigate the risk of a 'bad' equilibrium emerging. These fall into two broad categories:

i) measures that lower the system-wide demand for liquidity by improving coordination, and hence liquidity-recycling, in the payment system; and
ii) measures that lower the opportunity cost of generating intraday liquidity.

Some of these options may build on existing design features of large value payment and settlement systems; others may require that entirely new processes or arrangements be introduced.

The principal alternative options falling into these categories are listed in **Table A** and elaborated in the remainder of this section. Either set of measures may also be complemented by 'deep dive' examination of banks' intraday payments behaviour to inform the appropriate calibration of banks' intraday buffers and to ensure banks are not attempting to free ride on the liquidity provision of others to reduce their own liquidity needs.

Table A Policy options to mitigate the risk of payment delays

Measures that improve liquidity-recycling in the payment system	Measures that lower the opportunity cost of generating intraday liquidity
Implementation of liquidity-saving mechanisms	Setting the central-bank eligible criteria for collateral sufficiently wide: including, foreign-currency collateral, where appropriate
Implementation and enforcement of binding throughput guidelines	
Creating incentives for early submission of payments through variable tariffs	

5.1 Implementation of liquidity-saving mechanisms

Altering the design of the RTGS system to directly improve liquidity efficiency is a first way to mitigate the risk of settlement delays. In particular, the introduction of so-called liquidity-saving mechanisms (LSMs) can allow a bank to significantly reduce the amount of liquidity required to meet a given quantum of payments while minimising the delay in settlement.(2) In fact, LSMs can incentivise earlier submission of payment instructions compared to plain-RTGS. In recent

(1) See Manning *et al* (2009) and Bech and Garrett (2003).
(2) See Norman (2010) for more explanation of liquidity-saving mechanisms.

Box 3
Impact of delays on liquidity in the payment system

This box aims to illustrate the possible impact of banks delaying payments in response to an intraday liquidity buffer requirement. As in Box 1, the charts reflect payments between three banks, labelled A, B and C. The exercise is meant to be illustrative and not to reflect the quantitative potential impact of banks delaying payments.

Bank A delays all of its payments by 30 minutes

Bank A sends and receives payments to and from Banks B and C. Bank A unilaterally decides to delay payments so as to reduce its intraday liquidity requirements. This results in an increase in operational risk for Bank A as payments are more concentrated towards the end of the day.

Chart A shows that Bank A is able to reduce its intraday liquidity needs by unilaterally delaying payments to Banks B and C. The blue, payments sent line shifts to the right and consequently Bank A's largest net debit position falls. Bank A's intraday liquidity needs reduce by 16% from £1.39 billion to £1.17 billion.

Chart A Bank A's payments sent, received and net position if it delays all payments by 30 minutes

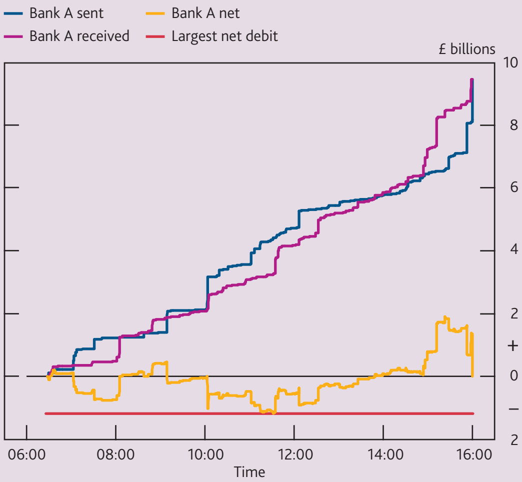

Sources: Bank of England and authors' calculations.

However, as illustrated in **Table 1**, there is a systemic cost to these delays. Firstly other banks intraday liquidity needs increase as a result: Bank B's increase by 3%, whereas Bank C's increase by 41%. The aggregate effect is to increase the amount of liquidity needed to settle all payments from £3.82 billion to £4.14 billion. Secondly, as shown in **Table 1**, noon throughput falls from 52.7% to 51.1% and the average time at which payments were settled increases by 7 minutes. These two measures indicate that the system is more vulnerable to operational risk.

Table 1 Descriptive statistics when Bank A delays

	Base case	Bank A delays
System throughput	52.7%	51.1%
Average time of settlement	11:56 am	12:03 am
Total intraday liquidity usage	£3.82 billion	£4.14 billion

All banks delay all payments by 30 minutes

In a second scenario, Banks B and C respond to Bank A's payment delays by also delaying all their payments by 30 minutes. The result is that there is no change in the relative timing of payments, but all payments are made later in the day. This is shown in **Chart B**.

Chart B Bank A's payments sent, received and net position if all banks delay all payments by 30 minutes

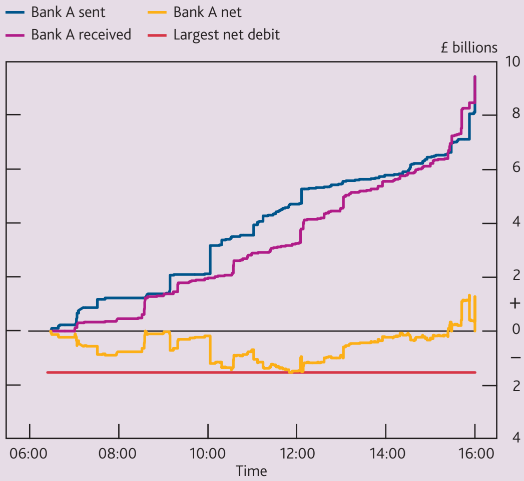

Sources: Bank of England and authors' calculations.

Table 2 shows that this is a situation in which no one benefits, but everyone suffers. The total liquidity usage across the system is unchanged, as is the distribution of that liquidity usage. However, by delaying all payments by 30 minutes the banks have increased operational risk in the system: noon throughput has fallen by 17 percentage points and the average time of settlement is now 27 minutes later.(1)

Table 2 Descriptive statistics when all banks delays

	Base case	All banks delay
System throughput	52.7%	35.7%
Average time of settlement	11:56 am	12:23 pm
Total intraday liquidity usage	£3.82 billion	£3.82 billion

(1) The delay in the average time of settlement is not quite 30 minutes as payments that were delayed until after 4 pm are assumed to have been settled at 4 pm.

years, such design features have been introduced in a number of countries' large-value payment systems (CPSS (2005) and Bech *et al* (2008)).(1)

In most cases, a system design that incorporates liquidity-saving features comprises the following key elements:

- a central queue;
- an offsetting algorithm; and
- liquidity-reservation functionality.

A central queue allows payment orders to wait for settlement where they are 'visible' to an offsetting algorithm. An offsetting algorithm finds pairs or groups of payments that can be settled simultaneously, with reduced liquidity usage. Liquidity-reservation functionality allows banks to set aside liquidity for time-critical payments to ensure that they can be settled without delay. It also acts to constrain the liquidity available to non-priority payments allowing them to queue and take advantage of the liquidity benefits of the offsetting algorithm.

In a system in which all three elements are present participants channel payments according to their priority: 'urgent' payments are submitted for immediate settlement, while non-priority payments are submitted to the central queue awaiting 'conditional release.'

In TARGET2 and similar systems, payments are released for settlement subject to an algorithm that works through the queue and identifies pairs or groups of offsetting payments that can be settled simultaneously, subject to the balance on the participant's settlement account not falling below a threshold level.(2) Consistent with the RTGS philosophy, offsetting payments are, from a legal standpoint, still settled on a gross basis with intraday finality under such arrangements. Hence there is no reintroduction of credit risk. However, the liquidity requirement reduces to the net difference between the values of the offsetting payments.

In principle, centralised queuing combined with an offsetting algorithm should encourage early submission of non-priority payment instructions. Willison (2005) finds that the incentive to delay payments declines upon the introduction of such features: if every bank submits payments early, there is a greater likelihood that the algorithm will identify offsetting payments.

Theoretical and empirical work to date generally concludes that introducing such measures would indeed deliver liquidity savings.(3) Norman (2010) quotes estimated savings of 20% and 15%, respectively, in the Korean and Japanese systems, where offsetting algorithms were recently introduced. Preliminary estimates of prospective system-wide savings in CHAPS are of the order of 30% (Ercevik and Jackson (2009) and McLafferty and Denbee (2011)), though these do not accrue evenly to all participants. McLafferty and Denbee also suggest that it is dividing payments into a priority and a non-priority channel that produces the liquidity savings. The sophistication of the offsetting algorithm serves to reduce settlement delay for a given liquidity saving.

The Bank of England has acknowledged these benefits and has commenced development of an LSM for CHAPS payments. The intention is for it to be implemented in 2013.

5.2 Throughput rules

Throughput rules establish the minimum proportion of a participant's daily settlement flow (measured in either volume or value terms) that must be settled by a particular time. Several countries' large-value payment systems have such rules in place, including those in Canada, Hong Kong and the United Kingdom. As mentioned earlier, in the United Kingdom, members of the CHAPS Clearing Company have mutually agreed to submit, on average, 50% of daily settlement flow by noon, and 75% by 2.30 pm, each calendar month. As described in Becher *et al* (2008), enforcement of the throughput guidelines in CHAPS currently relies on peer pressure rather than financial or regulatory sanctions. Persistent breaches result in a participant being called to give evidence to a committee of its peers, the so-called 'Star Chamber'.

Buckle and Campbell (2003) demonstrate that throughput rules can promote efficient liquidity recycling.(4) The authors show that, as long as the penalty for non-compliance with the throughput rule exceeds the private benefit of delaying payments, banks will comply and the system-wide liquidity requirement for a given quantum of payments will fall (ie liquidity turnover will increase). The authors also consider the optimal design of throughput rules; in particular, the optimal number of throughput requirements within the payments day. They demonstrate that there will, in theory, always be a positive (albeit diminishing) efficiency benefit to implementing an additional throughput requirement: 'in the limit, all payments could be made with an infinitesimal amount of collateral recycled infinitely frequently during the day.' This is illustrated in Box 4.

(1) Ten years ago, in a sample of a dozen of the largest payment systems in the world, approximately 3% of payments by value were settled in systems using liquidity-saving features; by 2005 this proportion had risen to 32%. See Bech *et al* (2008). This trend has continued since, most notably with the introduction of TARGET2 in the euro area.

(2) For more detail on liquidity-saving mechanisms in theory and in practice, see Manning *et al* (2009) and McAndrews and Trundle (2001).

(3) Willison (2005), Jurgilas and Martin (2010) and Martin and McAndrews (2008) examine the effectiveness of such measures from a theoretical standpoint. Galbiati and Soramaki (2010) apply agent-based modelling techniques and conclude that when the opportunity cost of liquidity is high a system that allows payments to be channelled to an offsetting stream improves efficiency, with better co-ordination of settlements.

(4) There is, however, an inherent source of inefficiency in throughput requirements. That is, they enforce the same payment timing on all banks, irrespective of possibly differing liquidity costs and differing customer preferences. For instance, if one bank's liquidity cost is higher than the others, it is efficient for that bank to submit its payments later than banks that face lower liquidity costs. The throughput guideline ignores these bank-specific factors.

Box 4
Throughput requirements: a stylised example

This box illustrates how enhanced throughput guidelines could be used to encourage banks to submit payments earlier. They could also act as a co-ordinating device which may result in a reduction in overall liquidity usage. As in the previous boxes, the data used here are adjusted real data for payments between three banks on a single day. The results are illustrative of the possible impact of more throughput guidelines and are not meant to provide quantitative estimates of the scale of the impact.

Chart A presents the banks' baseline sending profiles, where the y-axis captures the percentage of a day's payments sent.

Chart A Banks' base-line sending profiles

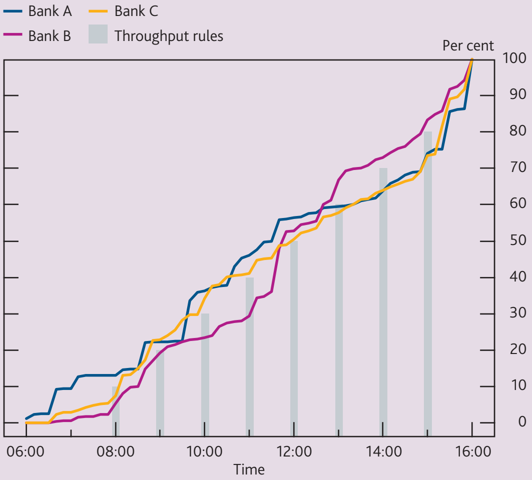

Sources: Bank of England and authors' calculations.

Let us now assume that multiple throughput requirements are introduced: 10% by 8 am; 20% by 9 am; 30% by 10 am up to 80% by 3 pm. As is clear from **Chart A**, in order to comply with the early throughput guidelines, Bank B would need to change its sending profile considerably. Banks A and C would also have to make some changes to their sending profile.

Chart B shows the profile with behavioural changes to meet the guidelines.

In this example, the introduction of throughput rules leads to a decline in liquidity needs for all three banks, reflecting the improved liquidity recycling early in the day. Bank B has to commit additional liquidity to the system to meet the earlier throughput requirements and so uses more liquidity. Banks A and C, however, enjoy savings of 3% and 30%, respectively, since they are able to recycle incoming flows from Bank B.

Across the system liquidity needs fall by 5% and earlier settlement reduces the potential impact of operational risk.

Chart B Banks' sending profiles with behavioural change to meet the throughput guidelines.

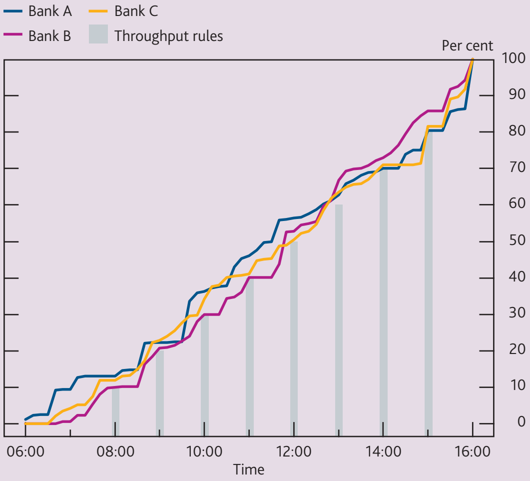

Sources: Bank of England and authors' calculations.

Table 1 Impact of throughput on banks' liquidity needs

	Liquidity needs without throughput rules	Liquidity needs with throughput rules	Percentage change
Bank A	£1.53 billion	£1.48 billion	-3%
Bank B	£1.34 billion	£1.49 billion	+11%
Bank C	£945 million	£659 million	-30%
System	£3.82 billion	£3.63 billion	-5%

The authors recognise, however, that in practice there will be a finite optimal number of throughput requirements: first, payments are not infinitely divisible; and, second, since the arrival of payment instructions is stochastic and uncertain, at some point participants will fail to comply simply because they have no payments to send within the specified time-frame.

This exposes a crucial implementation challenge. In CHAPS, while most participants tend to meet both the noon and 2.30 pm throughput guidelines and on average system throughput is above 50% at noon, some participants still persistently miss the target.(1) These banks must demonstrate to the CHAPS 'Star Chamber' that their failure to meet the guideline is not due to their withholding liquidity, but rather the absence of payment instructions to submit to the system.

This illustrates the difficulties in designing an optimal enforcement framework for throughput rules. Arguably, a stronger enforcement mechanism than peer pressure would be preferable, to ensure that the cost of non-compliance exceeds the cost of payments delay. Peer pressure is also unlikely to be credible when a bank feels that its capital is at risk, ie it is unlikely to prevent a bank from delaying payments to reduce credit risk to other participants experiencing solvency concerns.

Even if enforcement problems could be resolved, there would remain scope for banks to delay payments strategically within a throughput period, potentially leading to spikes in throughput just ahead of the target time. Indeed, as evident in **Chart 2** (Section 1.1), this pattern is currently observed in CHAPS. Any strengthening of throughput guidelines should perhaps therefore also include more extensive monitoring of banks' payments behaviour, to actively discourage such strategic behaviour (see Section 5.5, below).

5.3 Time-varying tariff

An alternative to hard submission deadlines enforced through penalties is to incentivise early submission through pricing. A progressive tariff, which increases through the day, might be expected to give banks good incentives to submit payments early. And in contrast to the blunter instrument of throughput guidelines, each bank can continue to optimise its sending profile based on its own liquidity and delay costs (incorporating the increase in the tariff).(2)

The Swiss payment system, SIC, is perhaps the leading example of a system with a progressive tariff (see **Table B**). In the Swiss case, the tariff increases exponentially over time.(3) The Swiss tariff also rises much more steeply for high-value payments. This would seem appropriate: first, banks face a sharper incentive to delay high-value payments due to the higher associated liquidity demand; and second, the system-wide liquidity impact of the delay of a high-value payment will be significantly greater than that of a low-value payment.

Table B SIC tariff structure

Time	Transaction value (CHF)	Transaction fees (CHF)[a]	
		Delivery	Settlement
End-of-day processing — 8 am		0.007	0.00
8 am–11 am	<100,000	0.01	0.02
	>100,000	0.10	0.15
11 am–2 pm	<100,000	0.10	0.20
	>100,000	0.30	0.70
2 pm — end of day	<100,000	0.40	0.60
	>100,000	1.00	2.00

Source: SIX Interbank Clearing.

(a) Fee for sending banks. The recipient bank pays a 0.02 CHF fixed-fee regardless of payment size.

Calibration of the tariff is important, but challenging. To impact effectively on banks' incentives, the tariff must be calibrated to alter the bank's optimal trade-off between liquidity cost and delay cost. James and Willison (2004) estimate the opportunity cost of posting liquidity in CHAPS as being equal to the difference between a bank's unsecured and secured costs of borrowing, which (at that time) was equal to 7 basis points per annum. This equates to a daily cost of around 0.02 basis points. With a delay cost of zero, a £1 million payment should attract a tariff of the order of £2; with a positive delay cost, the required tariff would decline.

Theoretically, there may be a case for a linear, rather than exponential tariff, since in the presence of stochastic operational shocks, the delay cost faced by a bank itself rises exponentially. As such, an exponentially increasing tariff could impose an excessive penalty (particularly where a bank may receive a material proportion of customer instructions late in the day).

5.4 Collateral eligibility criteria

The overall cost of providing intraday liquidity to the system depends on both the quantum required to settle payments on a timely basis and the unit opportunity cost of generating liquidity. The policy measures described above all consider the former, but in theory steps could also be taken that operate on the latter. In a collateralised system, one such measure might be to lower the opportunity cost of generating liquidity by modifying the eligibility criteria of collateral accepted for intraday liquidity. By doing so, central banks would enable settlement banks to post assets that they are already holding on their balance sheets, and for which they have no alternative intraday use.

(1) See Bank of England (2009).
(2) See Ota (2011).
(3) Serbia has introduced a progressive tariff, which also increases exponentially through the payments day. Brazil's STR system also has a time varying tariff, with payments settled before 9 am charged half the normal fee.

When considering the appropriate set of eligible collateral, central banks may find there are a number of constraints to providing funds against less liquid securities. First, it may provide incentives for banks to readjust their portfolio composition towards riskier assets, although on an intraday time-frame there may be few opportunities for this (perhaps only the extension of credit by settlement banks to their customer banks). In this regard, more intensive monitoring could assist.

Another, perhaps more material, concern is that it could introduce tensions between monetary policy and payment systems objectives. If the central bank has wider eligibility criteria for intraday liquidity than for access to the operational standing facilities (OSFs) this could cause difficulties at the end of the day. For example, a settlement bank that has borrowed intraday liquidity from the Bank via intraday repo may find that it is short of cash at the end of the day and so be unable to unwind the repo. Furthermore, the intraday repo may have been secured with assets that are ineligible for the OSFs. This would mean that the repo could not be rolled without the bank replacing the collateral with higher-quality assets. The bank may not have such assets available and so be unable to borrow from the central bank, interfering with the functioning of monetary policy implementation.

5.5 'Deep dive' examination of banks' payments behaviour

A possible complement to mechanisms to improve coordination of payments may be intensive analysis of banks' payments behaviour as captured in their liquidity reports to the prudential regulator.

In the United Kingdom's case, the framework for such analysis could be the FSA's Individual Liquidity Adequacy Assessment (ILAA). In the presence of an additional component of the liquidity buffer for intraday liquidity, a bank's may wish to delay payments to reduce intraday liquidity usage and so reduce the size of the intraday element of its liquid asset buffer upon recalibration.[(1)]

As such, if a bank were to submit its liquidity report to the prudential regulator, arguing for a sizable reduction in its intraday liquidity buffer, then, a 'deep dive' examination of that bank's payment system performance could be carried out, drawing on information both from transaction records within the payment system, and information sought directly from the bank on the timing of receipt of client instructions. Such analysis would enable the regulator to gauge whether the reduction in the required buffer reflected:

- a change in the bank's underlying business (ie the magnitude and timing of customer or proprietary flows); and/or
- improved payments efficiency (better co-ordination of payments with others in the system); and/or
- free-riding on other banks' liquidity (which would be evident from, on average, later submission times, and — combining information from the payment system and bank's internal scheduler — delays between the timing of the arrival of payment instructions and the timing of settlement in the system).

If such an examination revealed that the bank had indeed been engaging in free-riding, then the regulator could insist on maintaining the intraday liquidity buffer at the pre-existing level, thereby preserving system-wide liquidity and facilitating good recycling.

If the threat of such a 'deep dive' examination is credible — which it should be if the process is well articulated and the authorities have the appropriate powers to request information and the tools to conduct the analysis — the *ex ante* incentive to delay will be minimised.

5.6 Proposed changes to securities settlement systems

There are a number of proposed system enhancements that will reduce liquidity risk in securities settlement systems. EUI, the operator of the CREST service, has announced the introduction of a 'Term DBV' product in June 2011.[(2)] This will be in addition to the existing overnight delivery-by-value (DBV) mechanism in CREST. When a firm enters a DBV transaction it unwinds every morning (ie the cash is returned to the cash lender and the securities are returned to the cash borrower) and then it re-winds every evening. The Term DBV product will allow repos to be transacted and only unwound at the end of the term. This will lead to a reduction in the need for intraday liquidity for the cash borrowing bank.

There are also plans to move from a supply driven liquidity generation model, whereby intraday liquidity is automatically created from the central bank when eligible securities are purchased, to a demand driven one where liquidity is created when it is needed to settle a transaction. This should limit the amount of liquidity that is generated. To the extent that this incentivises more prudent liquidity management this will reduce the potential for intraday liquidity stresses to crystallise.

These enhancements show how the design of security settlement systems can play a significant role in reducing liquidity risk and influencing participant behaviour.

(1) Since, once in place, the incremental opportunity cost to the bank of drawing on the buffer to facilitate settlement of payments early in the day is zero.
(2) See Euroclear UK and Ireland (2010).

6 Conclusion

This paper has sought to clarify the main features of intraday liquidity usage and the intraday liquidity risks that arise in payment and settlement systems. In particular, the paper has highlighted how the practice of double duty, where prudential asset buffers are used as collateral in the payment system without being calibrated to include intraday liquidity risk, may create undesirable risks, both to individual participants and the system as a whole.

The Bank supports the FSA's approach of including intraday liquidity risk as a key risk driver in their wider liquidity regulations and so making banks hold liquid assets to support their intraday liquidity needs over and above those for balance sheet resilience. However, while critical in mitigating intraday liquidity risk, increasing the size of a bank's liquid asset buffer to include this risk might create incentives for banks to reduce the provision of intraday liquidity in large-value payment systems.

This paper has shown that such a change in behaviour could have the undesirable effect of increasing operational risk, and potentially credit risk, in these systems. The paper has therefore outlined a menu of options for mitigating this risk, some relating to changes in the payment system, others aimed at lowering the opportunity cost of providing liquidity. Together with regulatory vigilance, they might ensure that the full benefits of intraday liquidity regulation are achieved, without undermining the resilience of large-value payment systems.

Annex
Payment systems and intraday liquidity

Intraday liquidity

In an RTGS payment system, the settlement of payments requires that banks have access to intraday liquidity. To illustrate, two banks, Bank X and Bank Y, each have to make two payments of one unit to the other bank. The order of the settlement of these payments determines how much liquidity each bank needs and how much is required in aggregate.

Case 1

Bank X makes its first payment, then Bank Y makes both of its payments and finally Bank X makes its second payment.

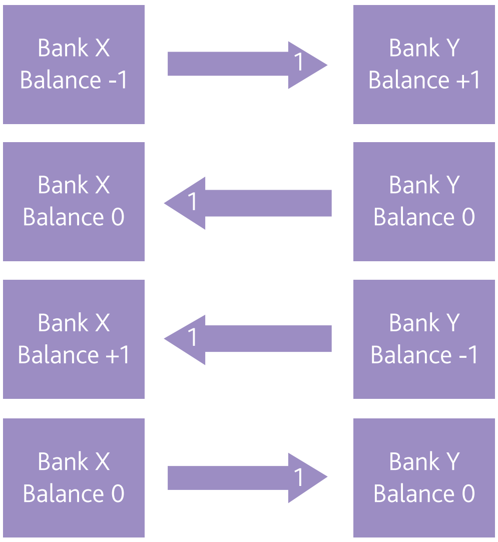

The result is that both banks use 1 unit of liquidity.

Case 2

Bank X makes both of its payments and then Bank Y makes both of its payments.

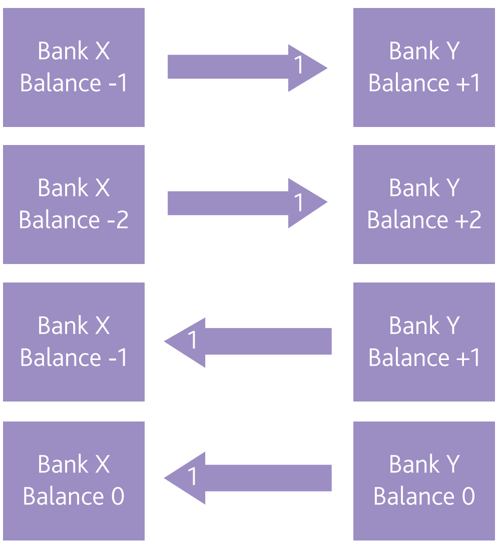

The result is that Bank X uses two units of liquidity whereas Bank Y uses none.

Intraday liquidity sources

In both cases Bank X has to have access to intraday liquidity in order to fund its payments. In case 2, Bank Y was able to make all of its payments without access to intraday liquidity beyond the payments that it received from Bank X.

Bank X can obtain intraday liquidity from a number of sources:(1)

i) Central bank reserves. It can start the day with cash on its central bank reserve account which can be used for making payments.

ii) Collateral that can be pledged/or repoed to the central bank in return for intraday liquidity. Most Central Banks offer settlement banks free collateralised intraday liquidity.

iii) Overnight money markets. Bank X could enter the overnight unsecured money market and borrow some funds. If these funds are received before it needs to make the first payment, it can use these funds for intraday liquidity.

Indirect participants

Many banks do not participant directly in large value payment and settlement systems but use the services of a correspondent bank instead. A correspondent bank is a direct participant of the system and offers to settle payments on behalf of another bank.

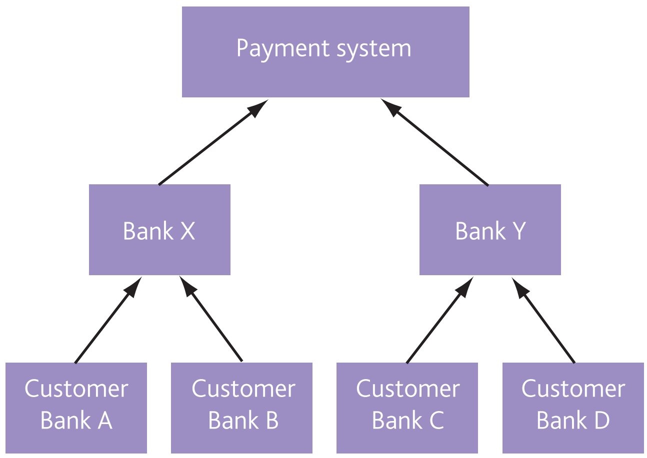

Bank A sends a payment to Bank D:
Bank A will send an order to Bank X to ask it to make a payment on its behalf. Bank X will debit Bank A's account and send the message to the payment system. This will debit Bank X's settlement account and credit Bank Y's account. When Bank Y receives the payment it will in turn credit Bank D's account.

(1) There are a number of other sources of intraday liquidity that a bank could use such as intraday credit lines or incoming payments from ancillary systems. For simplicity we have only included the most commonly used.

Bank A sends a payment to Bank B:
Bank A will send an order to Bank X to ask it to make a payment on its behalf. As Bank B is also a customer of Bank X the payment does not need to be sent to the payment system. Bank X will debit Bank A's account and credit Bank B's account. The payment will settle across the books of Bank X. This type of payment is known as an internalised payment. Bank X requires no liquidity to settle such a payment, but Bank A does.

References

Arjani, N (2006), 'Examining the trade-off between settlement delay and intraday liquidity in Canada's LVTS: a simulation approach', *Bank of Canada Working Paper no. 2006–20*.

Armentier, O, Arnold, J and McAndrews, J (2008), 'Changes in the timing distribution of Fedwire Funds transfers', *FRBNY Economic Policy Review*, September.

Bank of England (1989), *Payment and settlement systems risks and efficiency: A discussion paper*, November.

Bank of England (2009), *Payment Systems Oversight Report 2008*.

Bank of England (2011), *Payment Systems Oversight Report 2010*.

Bank of England (2010), *Financial Stability Report*, June.

Basel Committee on Banking Supervision (2008), 'Principles for sound liquidity risk management and supervision', *Bank for International Settlements*.

Basel Committee on Banking Supervision (2009), 'International framework for liquidity risk measurement, standards and monitoring', *Bank for International Settlements*.

Basel Committee on Banking Supervision (2010), 'Basel III: International framework for liquidity risk measurement, standards and monitoring', *Bank for International Settlements*.

Bech, M (2008), 'Intraday liquidity management: a tale of games that banks play', *FRBNY Economic Policy Review*, September.

Bech, M and Garratt, R (2003), 'The intraday liquidity management game', *Journal of Economic Theory*, Vol. 109, No. 2.

Bech, M and Soramäki, K (2005), 'Gridlock resolution and bank failures in interbank payment systems', in *Liquidity, Risk and Speed in Payment and Settlement Systems — a Simulation Approach*, Bank of Finland Studies E: 31–2005.

Bech, M, Preisig, C and Soramäki, K (2008), 'Global trends in large-value payments', *FRBNY Economic Policy Review*, September.

Becher, C, Galbiati, M and Tudela, M (2008), 'The timing and funding of CHAPS sterling payments', *FRBNY Economic Policy Review*, September.

Beyeler, W, Glass, R, Bech, M and Soramäki, K (2007), 'Congestion and cascades in payments systems', *Physica A: Statistical Mechanics and Its Applications*, Vol. 384, No. 2.

Buckle, S and Campbell, E (2003), 'Settlement bank behaviour and throughput rules in an RTGS payment system with collateralised intraday credit', *Bank of England Working Paper no. 209*.

Central Bank of Brazil (2008), *Brazilian Payment System*, September.

Committee on Payment and Settlement Systems (1997), 'Real-time gross settlement systems', *Bank for International Settlements*, March.

Committee on Payment and Settlement Systems (2001), 'Core principles for systemically important payment systems', *Bank for International Settlements*, January.

Committee on Payment and Settlement Systems (2005), 'New developments in large-value payment systems', *Bank for International Settlements*, May.

Committee on Payment and Settlement Systems (2006), 'Cross-border collateral arrangements', *Bank for International Settlements*, January.

Committee on Payment and Settlement Systems (2007), 'Payment systems in Serbia', *Bank for International Settlements*, June.

Diekmann, A (1985), 'Volunteer's dilemma', *Journal of Conflict Resolution*, Vol. 29, No. 4, pages 605–10.

Ercevik, K and Jackson, J (2009), 'Simulating the impact of a hybrid design on the efficiency of the CHAPS large-value payment system', in Leinonen, H (ed), *Scientific monographs E:42*.

Euroclear UK and Ireland (2010), 'Euroclear UK and Ireland's system enhancements of June 2011 due to introduction of the Term DBV service', *Technical Newsletter*, October.

Financial Services Authority (2009), 'Strengthening liquidity standards', *Policy Statement 09/16*, October.

Folkerts-Landau, D, Garber, P and Schoenmaker, D (1996), 'The reform of wholesale payment systems and its impact on financial markets', *Group of Thirty Occasional Paper no. 51*.

Friedman, J (1971), 'A non-cooperative equilibrium for supergames', *Review of Economic Studies*, Vol. 38, No. 1, pages 1-12.

Galbiati, M and Soramäki, K (2008), 'An agent-based model of payment systems', *Bank of England Working Paper no. 352*.

Galbiati, M and Soramäki, K (2010), 'Liquidity saving mechanisms and bank behaviour', *Bank of England Working Paper no. 400*.

Green, E (2007), *The role of Central Banks in payment systems*, in Haldane *et al* (2007), pages 45–56.

Haldane, A, Millard, S, and Saporta, V (eds), 'The future of payment systems', Routledge International Studies in *Money and Banking*, 43, Routledge, Abingden, United Kingdom.

Harrison, S, Lasaosa, A and Tudela, M (2005), 'Tiering in UK payment systems: credit risk implications', *Bank of England Financial Stability Review*, December.

Jackson, J and Manning, M (2007), *Central Bank intra-day collateral policy and implications for tiering in RTGS payment systems*, in Haldane *et al* (2007), pages 138–59.

James, K and Willison, M (2004), 'Collateral posting decisions in CHAPS Sterling', *Bank of England Financial Stability Review*, December.

Johnson, K, McAndrews, J and Soramäki, K (2004), 'Economizing on liquidity with deferred settlement mechanisms', *FRBNY Economic Policy Review*, December.

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

Kahn, C and Roberds, W (1998), 'Payment system settlement and bank incentives', *Review of Financial Studies*, Vol. 11(4), pages 845–70.

King, M (2009), Speech to Scottish business organisations, Edinburgh.

Manning, M and Willison, M (2006), 'Modelling the cross-border use of collateral in payment systems', *Bank of England Working Paper no. 286*.

Manning, M, Nier, E and Schanz, J (2009), *The economics of large-value payments and settlement*, Oxford.

Martin, A (2004), 'Optimal pricing of intraday liquidity', *Journal of Monetary Economics*, Vol. 51(2), pages 401–24.

Martin, A and McAndrews, J (2008), 'Liquidity-saving mechanisms', *Journal of Monetary Economics*, Vol. 55, No. 3.

McAndrews, J and Trundle, J (2001), 'New payment system designs: causes and consequences', *Bank of England Financial Stability Review*, December.

McLafferty, J and Denbee, E (2011), 'Liquidity saving in CHAPS: a simulation study', *Bank of England Working Paper*, forthcoming.

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 of payment systems', *Bank of England Working Paper no. 428*.

Leigh-Pemberton, R (1989), 'Challenges facing the Sterling wholesale payment systems: Ernest. Sykes Memorial Lecture'.

Valukas, A (2010), 'Lehman Brothers Holdings Inc. Chapter 11 proceedings Examiner's report', *Jenner & Block*, March.

Weesie, J and Franzen, A (1998), 'Cost sharing in a volunteer's dilemma', *The Journal of Conflict Resolution*, Vol. 42, No. 5, pages 600–18.

Willison, M (2005), 'Real-time gross settlement and hybrid payment systems: a comparison', *Bank of England Working Paper no. 252*.