Working Paper no. 337

# Risks and efficiency gains of a tiered structure in large-value payments: a simulation approach

Ana Lasaosa and Merxe Tudela

January 2008

Bank of England

# Risks and efficiency gains of a tiered structure in large-value payments: a simulation approach

*Ana Lasaosa**
*and*
*Merxe Tudela***

Working Paper no. 337

* Financial Resilience Division, Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH.
Email: ana.lasaosa@bankofengland.co.uk
** International Economic Analysis Division, Monetary Analysis, Bank of England, Threadneedle Street, London, EC2R 8AH.
Email: merxe.tudela@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We wish to thank members of the Systemic Risk Reduction Division, participants in 'The Economics of Payments workshop 2006' and 'The 4th Simulator Seminar for Payment and Settlement System Experts, 2006' for valuable comments and suggestions. This paper was finalised on 11 October 2007.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.html.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.


ISSN 1749-9135 (on-line)

**Contents**

**Abstract**

The large-value payment system in the United Kingdom (CHAPS) is highly tiered: a few settlement banks make payments on behalf of many customer banks. This paper makes use of a simulation approach to quantify by how much tiering affects, on the one hand, concentration and credit risk and, on the other, the liquidity needs of CHAPS. We do so by creating scenarios where current settlement banks become customer banks and thus we increase the degree of tiering. The results show that concentration risk would rise substantially in what is already a highly concentrated system. As for credit risk, the size of intraday exposures compared with settlement banks' capital is very small and therefore the likelihood of contagion remote. More importantly, the increase in credit risk brought to the system by settlement banks leaving CHAPS bears little relationship to the values settled by each individual bank. We find that increasing the degree of tiering in CHAPS leads to substantial liquidity savings – although the liquidity saved is only a fraction of the spare liquidity currently posted in the system. Most of the savings are due to liquidity pooling rather than to internalisation of payments. There is a strong relationship between changes in values settled and liquidity needs. This relationship can be used to forecast the impact on liquidity needs if more banks were to join CHAPS. The quantification of the trade-off between risk and efficiency in different scenarios provides policymakers with a useful analytical framework for analysing the effects of tiering.

Key words: Large-value payment systems, liquidity, tiering.

*JEL* classification: G21 - Banks; other depository institutions; micro finance institutions; mortgages.

**Summary**

Only a few banks are direct members of the Clearing House Automated Payment System (CHAPS), the UK large-value payment system. The vast majority of banks access the system indirectly as second-tier banks, through any of the few direct members (settlement banks). We describe a system in which a very small proportion of banks are direct members as a highly tiered system. The degree of tiering affects both how risky and how efficient the UK system is. Recent research has classified the various risks and benefits of tiering in large-value payments, but much less progress has been made in quantifying these risks and benefits. This paper seeks to fill this gap. It does not attempt to establish the relationship between normal and stressed liquidity needs, or how liquidity insurance should be regulated.

In order to gauge how the degree of tiering in CHAPS affects risks and benefits, we need to be able to vary the degree of tiering while holding other factors constant. A simulation approach allows us to do this. We create artificial versions of CHAPS where we increase the degree of tiering by reducing the number of direct members. We then use the simulation results to quantify the impact of tiering on concentration risk (a large settlement bank being a potential single point of failure), on credit risk (how exposed settlement banks are to second-tier banks) and on how much liquidity the system needs for it to operate.

The results show that, in a more tiered system, concentration risk would rise substantially. The credit risk incurred by direct members extending unsecured intraday overdrafts to their customer banks for their payments business would not be substantial under normal circumstances. The likelihood of contagion of credit problems to the broader financial system would be remote in our more tiered system. More importantly, our analysis has shown that the increase in credit risk brought to the system by settlement banks leaving CHAPS bears little relationship to the values settled by each individual bank. The key determining factor of the size of settlement banks' intraday credit exposures to second-tier banks is the timing of intraday payments of second-tier banks – a variable that central banks do not observe directly.

Increasing the degree of tiering in CHAPS leads to substantial savings in the amount of liquid assets that settlement banks need to post every day. Only a small proportion of these savings are due to settlement banks settling payments across their own books. Moreover, the clear relationship between changes in values settled and liquidity needs shown by our simulations make it possible to project what would happen if current second-tier banks joined CHAPS as direct members. We estimate that the liquidity needs could increase by £8 billion in aggregate if as many as five large banks (in terms of values of payments processed) joined CHAPS – the opposite case to the one analysed so far. While this figure is significant, it is only a fraction of the £17 billion spare liquidity posted on average in the system as a whole every day.

## 1 Introduction

Only a few banks are direct members of the UK large-value payment system. The vast majority of banks access the system indirectly, through any of the few direct members. We describe a system in which a very small proportion of banks are direct members of a highly tiered system. Tiering is therefore defined as the proportion of banks that are direct members of the payment system. In this context, higher tiering means fewer direct members. The degree of tiering affects both how risky and how efficient the UK system is. While recent theoretical research and central bank analysis[1] has classified the various risks and benefits of tiering in large-value payments, much less progress has been made in quantifying these risks and benefits. This paper seeks to fill this gap. Its scope is restricted to risks and efficiency gains under normal circumstances. Under stressed circumstances, tiering can create additional operational and liquidity risks that are not discussed in this paper. As a result, we do not assess what degree of tiering is optimal from an overall welfare perspective.

CHAPS[2] is the large-value payment system in the United Kingdom. It has been a real-time gross settlement (RTGS) system since April 1996, processing an average of £200 billion and over 100,000 payments each day. Member banks hold settlement accounts with the Bank of England and they can run fully collateralised intraday overdrafts. Only fourteen settlement banks and the Bank of England are direct CHAPS members.[3] We refer to these banks are direct members, fist-tier banks or settlement banks. All other banks (the second-tier banks) have to access CHAPS through these direct or first-tier members. They are called indirect members, second-tier banks or customer banks. A similar degree of tiering is found in the embedded payment system of CREST, the UK securities settlement system.

In order to gauge how the degree of tiering in CHAPS Sterling affects measures of risks, we need to be able to vary it while holding other factors constant. This is not possible with cross-country comparisons, where many factors other than the degree of tiering vary. A simulation approach, by contrast, is particularly useful in creating artificial versions of one payment system with different degrees of tiering. This paper makes use of a simulation approach to quantify by how much tiering affects, on the one hand, concentration and credit risk and, on the other, the liquidity needs of the UK large-value payment system. Quantifying the trade-off between risks and efficiency enables us to provide policymakers with a framework in which to discuss the effects of alternative models. The paper then goes on to analyse how much of the liquidity gains are due to pooling and how much to

---

[1] See Kahn and Roberds (2005), Jackson and Manning (2007), Harrison, Lasaosa and Tudela (2005) and Bank of England *Payment Systems Oversight Report* (2005) (www.bankofengland.co.uk/publications/psor/psor2005).

[2] There are in fact two different large-value payment systems, CHAPS Sterling and CHAPS Euro. Since the values settled in CHAPS Sterling are 20 times higher than those settled in CHAPS Euro, our analysis will be restricted to CHAPS Sterling. All references to CHAPS in the paper refer to CHAPS Sterling.

[3] CLS Bank and Abbey joined at the end 2005. We count RBS and Natwest as a single member even though they actually have two separate accounts. A complete membership list can be found in www.apacs.org.uk/uk_payment_schemes/chaps_clearing_1.html.

internalisation of payments. It does not attempt to establish the relationship between normal and stressed liquidity needs, or how liquidity insurance should be regulated.

The rest of the paper is structured as follows. The next section outlines how tiering may affect risks and liquidity efficiency in the UK large-value payment system. Section 3 introduces the methodology used in our analysis. Section 4 goes on to describe the main results. Our concluding comments are presented in Section 5.

## 2 Tiering, risks and efficiency

In their 2003 Financial System Stability Assessment of the United Kingdom,[(4)] the International Monetary Fund (IMF) highlighted the potential risks arising from the highly tiered structure of the UK large-value payment systems. Table A shows how the degree of tiering in CHAPS – defined as the proportion of banks accessing the system directly – is high compared with other countries.[(5)] The IMF drew attention to the exposures arising between the first-tier and second-tier institutions and the potential for contagion risk, that is, the risk that credit problems in a second-tier bank might spill over to first-tier banks. Since the IMF assessment was published, the Bank has analysed the risks arising from tiering in UK payment systems.[(6)]

---

(4) Available at www.imf.org/external/pubs/ft/scr/2003/cr0346.pdf.

(5) Tiering defined in this way is obviously correlated with the general degree of concentration in the banking system: if only a few banks account for a large share of the banking market (the UK case), access to payment systems is likely to be tiered due to economies of scale.

(6) 'Strengthening financial infrastructure' article of the December 2004 *Financial Stability Review*, available at www.bankofengland.co.uk/publications/fsr/2004/fsr17art4; 2004 and 2005 *Payment System Oversight Reports* available at www.bankofengland.co.uk/publications/psor/psor2004 and www.bankofengland.co.uk/publications/psor/psor2005.

**Table A: Settlement banks in large-value payment systems**

| Country | System name | No. of settlement banks[a] | No. of credit institutions |
|---|---|---|---|
| United Kingdom | CHAPS Sterling | 13 | 420 |
| | CHAPS Euro | 19 | |
| Belgium | ELLIPS | 16 | 109 |
| Canada | LVTS | 14 | 45 |
| France | TBF | 156 | 1,067 |
| | PNS | 21 | |
| Germany | RTGS Plus | 93 | 2,370 |
| Italy | BIREL | 204 | 821 |
| Japan | BOJ-NET | 371 | 506 |
| Netherlands | TOP | 106 | 95 |
| Sweden | E-RIX | 13 | 125 |
| | K-RIX | 19 | |
| Switzerland | SIC | 307 | 327 |
| United States | Fedwire | 7,736 | 8,130 |
| | CHIPS | | Not available |
| European Union | TARGET | 1,579 | _ |
| | Euro1 | | |

Sources: Committee on Payment and Settlement Systems (CPSS) Statistics on payment and settlement systems in selected countries (2005) and OECD's 'Bank Profitability' (2003).

[a] Includes central banks. Data for 2003. As for 2006, CHAPS Sterling has fifteen members, including the Bank of England.

A tiered payment system can give rise to several types of risk relative to a system in which all banks are direct participants:

*Credit risk*: Credit exposures arise when settlement banks offer their customer banks overdraft facilities when making outward payments on their behalf. Conversely, when customer banks hold positive intraday balances at their settlement bank, the customer bank is exposed to the settlement bank. As Flannery (1996) points out, a tiered system that relies heavily on private credit may not function so well in times of crises.

*Concentration risk* (also referred to as 'node risk'): Tiering increases the concentration of payments in each settlement bank. The system is therefore more sensitive to temporary outages experienced by individual settlement banks. The consequences of a temporary disruption can be either merely operational (other banks' inability to make and receive payments) or also liquidity-related. A stricken bank able to receive but unable to send payments may become a 'liquidity sink', draining the rest of the system of the liquidity needed to continue making payments. The larger the proportion of the liquidity in the system controlled at any one time by an individual bank, the greater the risk that operational problems lead to liquidity difficulties.

*Legal risk*: The finality of payments made or received on behalf of indirect participants is not as well defined as in systems designated under the European Union's Settlement Finality Directive such as CHAPS. In addition, a tiered system makes internalisation of payments possible. Internalised payments are payments made between customer banks of the same settlement bank and settled internally across the settlement bank's books without being forwarded to the payment system. For a given degree of tiering, a greater degree of internalisation may increase the legal risk of payments being unwound.

*Liquidity dependency*: Second-tier banks may view their settlement banks as lenders under all conditions. First-tier banks, conversely, may depend on the incoming transactions of their customer banks for their own liquidity needs.

So far, the Bank's analysis of the risks involved in tiering has focused on credit risk to settlement banks. Based on evidence that first-tier banks extend unsecured intraday credit to second-tier banks, Harrison, Lasaosa and Tudela (2005) analysed the credit risk exposure of settlement banks using a standard credit risk model. They examined the change in the distribution of credit losses incurred by a bank that moves from only processing payments on behalf of its own customers to carrying out correspondent business on behalf of second-tier banks. The model was calibrated to UK financial infrastructures. It concluded that, in normal market conditions, the credit risk to first-tier banks appeared to be low. Even under stressed circumstances, the assumptions had to be extreme to lead to a significant increase in the credit risk faced by the settlement bank.

As for concentration risk, research carried out by the Bank has found that operational problems at individual CHAPS settlement members would not in general prevent the remaining banks from making payments to each other due to liquidity shortages – providing that banks quickly stopped making payments to the stricken bank.[7] Bedford, Millard and Yang (2004) use a simulation approach and find that the system exhibits a high level of resilience. Nonetheless, the operational failure of a key node bank would still disrupt all of its own customers and those to whom they were making payments.

Jackson and Manning (2007) construct a model that examines the key factors affecting banks' decisions whether to become direct members of a particular system or not, and the central bank's decision to require collateralisation of intraday credit or not. Their findings suggest the existence of economies of scale in correspondent banking, which make concentration likely.

---

[7] Whether banks do so in practice is an interesting subject for further research. If they do not, the results of Bedford *et al* (2004) may no longer hold.

A tiered payment system has benefits too. It can reduce systemic risk[8] in two different ways. First, a tiered payment system depends less on the central infrastructure because some payments are 'internalised'. In the case of some operational failure of the central system, payments across the books of settlement banks can still take place. A second, potential, risk-related benefit from tiering is the increase in monitoring by first-tier banks of the financial position of second-tier banks. Kahn and Roberts (2005)[9] argue that tiering increases the level of monitoring by first-tier banks and reduces the incentive to default by second-tier banks. If a second-tier bank proves itself to be unreliable, it will be required to collateralise fully its payment activity at an additional cost. If it is reliable, it only needs to be monitored. Hence, the first-tier bank has an incentive to monitor efficiently and the second-tier bank has an incentive to behave reliably.

- So far we have compared the systemic risks involved in a tiered and non-tiered payment system. But a tiered system can also be more efficient than one where all banks are direct members of the payment system. In a competitive market, the banks that have a competitive advantage when offering correspondent banking services become settlement banks. Larger banks are normally better placed to do so due to existing economies of scale in several areas:
- Infrastructure: IT, contingency arrangements, administration.
- Direct membership tends to be relatively more expensive for banks with small volumes of payments, due to the existence of both fixed fees and a sliding scale for per-transaction (volume-based) fees.
- Dedicated staff: a well-staffed liquidity management team.
- Liquidity requirements of RTGS systems: RTGS systems, such as CHAPS, are liquidity-intensive because they require payments to be made on a gross basis and to be fully pre-funded. Liquidity needs might decrease with the degree of tiering due to two effects, liquidity pooling and internalisation of payments – Section 4.5 explores which of the two effects is more important in CHAPS.

This paper focuses on the risks to and efficiency gains from tiering that we can quantify with the data at our disposal: concentration risk, credit risk, and liquidity savings. The liquidity demands analysed here are those the system needs to function under normal circumstances, where lower needs are better. The question of which level of liquidity would be sufficient under stressed circumstances is out of the scope of this paper, as are two related questions: whether a payment system that is liquidity-efficient needs proportionally more or less liquidity insurance in the event of a crisis, and whether the payment system should be designed in such a way that forces members to hold liquidity insurance.

---

[8] We define systemic risk as a risk to the financial system that is not adequately internalised by system participants and that imposes material costs to the banking system should it materialise.
[9] Available at www.bankofengland.co.uk/financialstability/futureofpayments/kahnrobertsBOE.pdf.

## 3 Methodology

We do not have access to actual data of CHAPS under different degrees of tiering that would allow us to study the implications of a more, or less, tiered structure. We can turn, though, to simulation techniques to analyse different tiering scenarios and their implications for liquidity needs and system risks. Specifically, we use the Bank of Finland payment and settlement system simulator (BoF-PSS2) initially developed in the mid-1990s by the Bank of Finland to study the effects of the introduction of European Monetary Union (EMU) on the Finnish payment systems. Since then the simulator has been developed for and used by several central banks.(10)

The BoF-PSS2 models settlement processes according to a set of rules defined for a payment system environment, giving as outputs account balances and payments settled and received that can easily be analysed within the simulator or exported to other programs. These outputs allow us to draw conclusions about how system characteristics such as credit risk, liquidity consumption, settlement speed or gridlock resolution vary under the different scenarios.

The first step in a simulation process is to establish a 'benchmark' against which other simulations (or scenarios) are compared. Our benchmark involves a simple replication of real-life CHAPS Sterling using actual transactions for June 2005. This means replicating a total of 2.5 million payments accounting for £4.2 trillion of total value transferred. We then construct different scenarios in which we vary the level of tiering in CHAPS Sterling.

For each transaction we have information on the settlement banks sending and receiving the payment, the payment amount and the exact time at which it was sent. We do not have information on end payer or payee, or an indication of whether the payment is sent on behalf of the settlement bank itself or one of its customer banks. And, obviously, internalised transactions are not included in the data since they do not go through CHAPS.(11)

Our simulations involve an *increase* in the degree of tiering and, therefore, a reduction in the number of direct CHAPS members. Decreasing rather than increasing the number of direct members allows us to observe the payments sent and received by the banks that we will turn into customer banks in the data. We can then assign them to a settlement bank once they become customer banks and conduct the simulations. Since our data does not identify the transactions in which customer banks are involved, we could not construct a scenario where customer banks become direct CHAPS members.

We increase the degree of tiering in the system by turning the seven smallest banks by value into customer banks of the three major settlement banks in turn. All together the seven smallest banks

(10) See Leinonen (2005) for a more technical and precise description of the BoF-PSS2.

(11) See Section 4.3 below for a detailed explanation of internalisation.

account for 17% of CHAPS transactions by value and 19% by volume. We start off by assigning (or converting into customer) the smallest bank (accounting for less than 1% of CHAPS values) to one of the three major settlement banks (let us call it major settlement bank 1); we then assign this same bank to another of the three major settlement banks (major settlement bank 2); in another step we assign it to a third major settlement bank (major settlement bank 3). We continue assigning the smallest and second smallest banks to major settlement bank 1, then to major settlement bank 2 and finally to major settlement bank 3. We go on with the three smallest banks in a similar way and continue this process up to the seven smallest banks. This gives us 21 different tiering scenarios.

We construct two additional scenarios in the following way. We take the seven smallest banks as before and we assign some of them to settlement bank 1, some of them to settlement bank 2 and the rest to settlement bank 3. The choice of which of the smallest banks to assign to which of the major settlement banks is based on the relative proportion of bilateral payments between the two types of banks, and in two ways, first considering payments by volume and second by value. New customer banks are likely to choose as their settlement banks those banks with whom they have a large proportion of payments by virtue of their business, since internalised payments are usually charged at a lower rate than those going through CHAPS. We call the first scenario ‘assign by volume’, and the second one ‘assign by value’. We have 23 different scenarios in total.

There are two main assumptions embedded in this experimental design. First, we assume that the timing of payments does not change when a settlement bank becomes a customer bank. Settlement banks have more discretion than customer banks about when to send payments to CHAPS – discretion constrained by intraday deadlines prevalent in financial markets and by throughput guidelines. But, with the information at our disposal, any changes to the timing of payments in our simulations would be completely arbitrary. The second assumption is that settlement banks that become customer banks take their own customer banks with them because there is no immediate reason to believe otherwise. Any other alternative would be arbitrary as well.

## 4 Results

In what follows we analyse how our measures of concentration risk, credit risk and liquidity efficiency change in the different scenarios.

### *4.1 Concentration risk*

Following James (2003), we gauge the increase in concentration risk by looking at individual banks’ share of all payments going through CHAPS. This measure shows the proportion of payments in the system that would be affected by an operational outage in that particular bank. This measure of concentration risk does not take into account the interconnectedness of the node. Concentration risk defined in this way does not encompass the impact of payments going through other settlement banks that could be affected through liquidity sink channels in such an operational crisis. We calculate the share of payments based on the value of all incoming and outgoing payments for each bank. Concentration risk for bank *j* is defined as value settled in CHAPS by bank *j* plus value received through CHAPS by bank *j* over the totals sent and received in CHAPS by all settlement banks.

$$Node\ Risk_j = \frac{Value\ Sent_j + Value\ Received_j}{CHAPS\ Totals\ Sent\ and\ Received}$$

Current concentration figures are 26% for the largest settlement bank, 25% for the second largest and 17% for the third largest.

Table B shows how concentration in CHAPS increases when the same seven customer banks are assigned to a combination of the three major settlement banks. As described in the previous section, we decide which customer bank to assign to each settlement bank according to the most common interbank payment flows that we observe in our data.

**Table B: Increase in CHAPS Sterling payments share when new customer banks are assigned to several settlement banks**

| | Major settlement bank 1 | Major settlement bank 2 | Major settlement bank 3 |
|---|---|---|---|
| **Current share** | **26.0** | **24.5** | **17.1** |
| Assign by value[a] | 31.7 | 23.3 | 29.3 |
| Assign by volume[a] | 30.4 | 23.9 | 30.4 |

Sources: Payments database and Bank calculations, average over business days, June 2005.
[a]For a definition of these terms or scenarios see page 11.

Chart 1 shows the increase in the share of payments of each major settlement bank when all the new customer banks are assigned to it. Each line shows the results of a different simulation: when all new customer banks are assigned to major settlement banks number 1, 2 and 3. Each point represents the average over the business days in June 2005. The unfilled shapes correspond to our benchmark case of current concentration risk figures. The degree of tiering moves up as we move towards the right of the chart, with each point on the horizontal axis representing one more small settlement bank that becomes a customer bank. Each differently coloured line shows how the concentration risk increases when the new customer banks are assigned to each of the three major settlement banks. Finally, the two single brown dots show the share of payments for major settlement bank 1 in scenarios ‘assign by value’ and ‘assign by volume’.

**Chart 1: Concentration risk (average June 2005)**

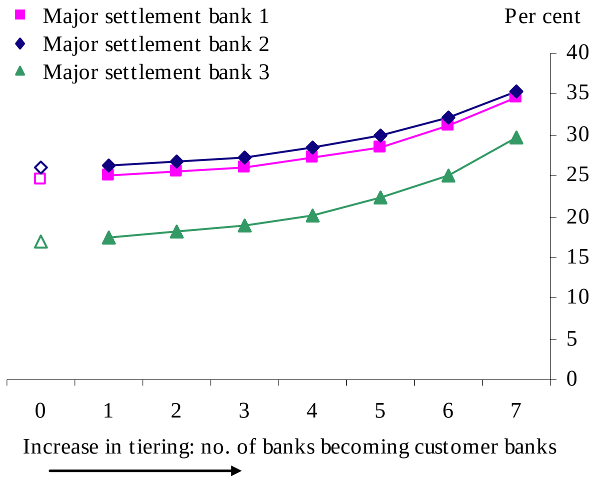


Sources: Payments database and Bank calculations.

We can see that the increase in concentration risk is greatest when the new customer banks are assigned to the major settlement bank that currently has the smallest payment share of the three. This is not surprising. Other things equal, the bigger a settlement bank is, the larger the proportion of its payments that can be internalised. Internalised payments are not sent through CHAPS, so internalisation decreases both the numerator and the denominator in our measure of concentration risk – but the decrease is stronger in the denominator. As a result, the more internalised payments a bank has, the smaller the impact of making other settlement banks their customer banks on concentration risk. The increase when either major settlement bank 1 or 2 capture the new customer banks is practically identical.[12] There is little day-to-day variation in concentration risk within our sample. The coefficient of variation is low, ranging between 0.02 and 0.03 in all our scenarios.

### *4.2 Credit risk*

We measure the increase in credit risk in the different scenarios by taking the maximum intraday liquidity that banks currently need as direct CHAPS members to be able to settle all transactions on a gross basis in real time. CHAPS settlement banks have to post collateral to access intraday credit from the Bank of England. We know that it is common practice in the United Kingdom to grant customer banks unsecured intraday overdraft facilities – especially to large customers. Thus, the maximum intraday liquidity currently used by settlement banks is likely to be a reasonable measure of the maximum unsecured intraday credit that each individual bank would need to obtain from their

---

[12] The change in our concentration risk measure as settlement banks become customer banks is made up of an increase in the numerator and a decrease in the denominator (total values settled) due to internalised transactions. This change in the denominator makes potential statistical relationships between overall values settled in CHAPS and concentration risk hard to interpret, despite obvious correlation coefficients. For this reason, we do not chart the relationship between values settled and concentration risk.

settlement banks if they became indirect CHAPS members. It is worth stressing that this would be an upper bound measure: it is reasonable to assume that settlement banks would have some degree of discretion of when to process their customers' payments and try to avoid building unnecessary credit exposures. Table C shows the monthly average of maximum intraday liquidity for each settlement bank that becomes a customer bank.

The first point to emerge from this table is that the magnitude of the risk is insignificant with respect to the amounts of Tier 1 capital held by the three large settlement banks: between £19 and £43 billion. This finding is consistent with Harrison *et al* (2005) who conclude that in normal market conditions the risk to settlement banks from their intraday credit exposures to second-tier banks appears to be low. Interpreting these results, one should bear in mind that these figures are monthly averages of intraday maxima – the exposures on which the averages are calculated may have lasted for only a few seconds each day. And the fact that the mean is in all cases higher than the median implies that the distribution is asymmetric, skewed towards a relatively small number of high exposures. The coefficients of variation displayed in Table C range between 0.46 and 0.78, reflecting a very disperse distribution of intraday maxima over the month.

**Table C: Maximum intraday credit that would have to be extended to each individual bank**

| Customer bank number | Maximum intraday credit (£ millions) | | | Share of CHAPS payments (per cent) | | Share of payments made on behalf of customers | | (a)/(mean value settled) |
|---|---|---|---|---|---|---|---|---|
| | Mean (std. dev.) (a) | Coeff. variation | Median | By value | By volume | By value | By volume | |
| 1 | 253 (135) | 0.53 | 248 | 0.5% | 1.2% | 9.7 | 31.7 | 0.23 |
| 2 | 385 (177) | 0.46 | 367 | 0.8% | 0.9% | 9.0 | 9.0 | 0.23 |
| 3 | 770 (421) | 0.55 | 615 | 0.8% | 2.2% | 0.0 | 0.3 | 0.45 |
| 4 | 718 (557) | 0.78 | 653 | 1.9% | 7.5% | 2.2 | 1.5 | 0.18 |
| 5 | 438 (320) | 0.73 | 341 | 2.6% | 0.7% | 0.0 | 0.0 | 0.08 |
| 6 | 1,042 (673) | 0.65 | 860 | 3.9% | 1.8% | 12.0 | 41.3 | 0.13 |
| 7 | 735 (486) | 0.66 | 587 | 6.6% | 5.1% | 59.9 | 54.3 | 0.05 |

Sources: Payments database (June 2005), 2005 Correspondent Banking Survey and Bank calculations.

It is also apparent from Table C that the size of the exposures is not proportional to the share of total CHAPS payment values made up by each bank.[13] Bank 3, with 0.8% of CHAPS payments, has similar mean peak liquidity usage as bank 7, with 6.6% of CHAPS payments. The correlation coefficient between the two series is only 0.52, with the ratio varying from 0.05 for bank 7 to 0.45 for bank 3. Economies of scale in liquidity usage (described in the next section) could explain why credit risk increases less than proportionally with the value of CHAPS payments. But they do not

[13] The correlation coefficient between the median maximum liquidity needed and the share of CHAPS payment values is 0.5.

explain why two banks with a similar share of CHAPS values have different maximum liquidity needs. The relative size of each bank's correspondent business (column 5 of Table C) does not emerge as a factor, either. Banks 3 and 5 have virtually no correspondent business and very different peak liquidity usage. The most plausible explanation is that these observed variations in liquidity usage may stem from differences in their customer base affecting the timing of incoming and outgoing payments – if a settlement bank needs to make a large number of payments before receiving many payments it has higher peak intraday funding needs.

The implication of the above analysis is that, unlike in the case of concentration risk, policymakers interested in reducing credit risk by encouraging more second-tier banks to join the payment systems cannot rely on the value of sterling payments processed as a reliable measure of potential intraday credit exposure. The key factor determining the size of intraday exposures is the intraday pattern of payments. Unfortunately, precise data on the intraday pattern of payments of customer banks is not usually directly available to policymakers, making their task hard.

### *4.3 Efficiency gains*

Compared with deferred net settlement (DNS) systems, RTGS systems reduce credit risk at the expense of increasing liquidity costs.(14) Not being able to net off payments increases the liquidity needed by each settlement bank and, therefore, the overall liquidity needs of the system. In the case of CHAPS, banks obtain liquidity in central bank money by posting eligible collateral with the Bank of England.

As discussed above, tiered payment systems are more concentrated. Concentration in payment systems leads to liquidity savings due to two effects that go in the same direction: liquidity pooling and internalisation of payments.

Liquidity pooling: despite the fact that payments cannot be netted off in a RTGS system, the larger the number of payments received and sent by a given settlement bank, the higher the probability that incoming payments fund (totally or partially) outgoing payments.(15) This results, on average, in smaller peaks of both liquidity needs (when the balance is negative) and liquidity surpluses (when the balance is positive). Since we are concerned with the average of *maximum* intraday liquidity needed, we expect our scenarios involving fewer but larger settlement banks to show a fall in liquidity needed with respect to current CHAPS figures. This saving involves no change in the payment values settled through CHAPS.

Internalisation: when a settlement bank becomes a customer bank of another settlement bank, all transactions between the two (either on their own behalf or on that of their customers) that used to go

(14) Selgin (2004) argues that the drive to replace DNS with RTGS systems on credit risk grounds is misguided. He claims that the only credit risk arising in DNS systems are those granted by the receiving banks to their customers. He also argues that regulators' guarantees distort the market.
(15) Except in very asymmetrical distributions for incoming and outgoing payments.

through CHAPS are internalised. Since they are no longer sent to the payment system, the settlement bank does not need to obtain intraday credit from the Bank of England to fund them. The internalisation effect leads to a decrease in liquidity needed as a result of the decrease in values settled through CHAPS. An estimation of the relative size of each effect is presented in Section 4.5.

Chart 2 shows the savings in liquidity needs as tiering increases; similarly Chart 3 shows the savings in liquidity needs relative to the decrease in overall values settled in CHAPS for each of our simulations. The figures are percentage change relative to current values. Different colours represent different scenarios. The three scenarios where all new customer banks are assigned to one large settlement bank have seven points each, with two extra data points (brown in the chart) for the simulations where all seven customer banks are assigned to a combination of settlement banks. The data points for the 'assign by value' and 'assign by volume' scenarios in Chart 2 overlap and therefore only one is visible.  The liquidity needs of the system are defined as the sum of each individual settlement bank's intraday maximum liquidity requirements. As tiering increases and the number of CHAPS settlement banks falls, the values settled in the system go down.

**Chart 2: Increase in tiering and liquidity needs**

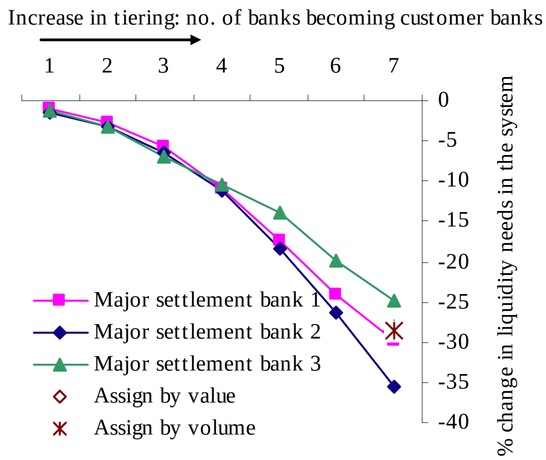


Sources: Payments database and Bank calculations.

**Chart 3: Percentage changes in liquidity needs in CHAPS and values settled in CHAPS**

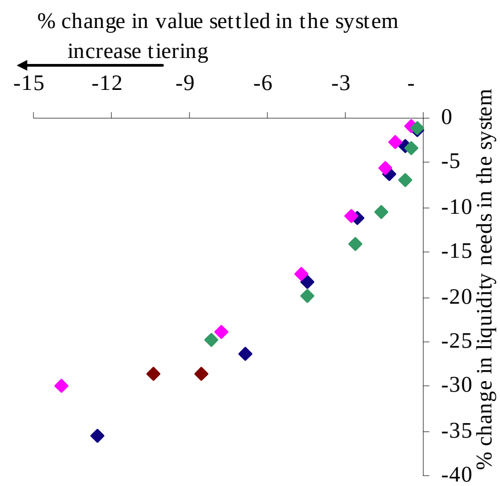


Sources: Payments database and Bank calculations.

The chart shows substantial liquidity savings associated with a reduction in the number of settlement banks (an increase in tiering). Moreover, the savings are similar across the different scenarios: dots of different colours in the chart are relatively close together. The maximum saving is 36% of current liquidity needs when all new customer banks are assigned to settlement bank 2. This is equivalent to £5.9 billion liquidity. It is apparent from Chart 3 that there is a close relationship between changes in overall values settled in CHAPS and changes in liquidity needs. For a given percentage reduction in values settled, the reduction in liquidity needs is approximately three times as big.

The variation in liquidity savings across the month is higher than in the case of concentration risk but much lower than for credit risk. The coefficients of variation in all 23 scenarios range between 0.12 and 0.15 – compared to 0.02-0.03 for concentration risk and 0.46-0.78 for credit risk.

### *4.4 Decrease in tiering: liquidity costs*

The close relationship between values settled and liquidity needs found across all scenarios allows us to attempt a forecasting exercise. Our interest is to try and gauge how much liquidity would CHAPS need if some large (in terms of values of payments processed) customer banks became settlement banks. This boils down to fitting a line to the points shown in Chart 3 and projecting it to the positive quadrant of the horizontal axis. This will give the increase in liquidity needs expected when the degree of tiering in the system decreases and the values settled in CHAPS increase.

Since no data points from our simulations lie in the positive quadrant of the horizontal axis we need to make strong assumptions about the functional form of the relationship between positive increases in values settled and positive increases in liquidity needs. The only information we have at our disposal is the relationship found in the negative quadrant. We try three regression specifications: linear, quadratic and cubic. The criteria for choosing a particular forecasting specification is twofold: goodness of fit to the observed data points and plausibility of the predicted values in the positive quadrant.

The three regressions fitted the data well, with adjusted $R^2$ of 0.95 for the linear equation, 0.99 for the quadratic and 0.99 for the cubic. But a quadratic form extended into the positive quadrant resulted in implausibly high values of liquidity needs: 100% extra for an increase in values settled of 7%. Our chosen functional forms for forecasting purposes, therefore, are a linear equation and a cubic equation. The cubic equation is used to create a concave mirror image of the convex line on the negative quadrant on the positive quadrant. We believe this is a plausible functional specification: extra liquidity needs become smaller as new banks join CHAPS, just as extra savings become smaller as more banks leave the system.

The points chosen for the forecast are not arbitrary. They correspond to the values that would settle in CHAPS if the largest, second largest and up to five largest customer banks joined CHAPS. We know the values from the 2003 CHAPS Traffic Survey, a data set that includes the values of transactions settled by current settlement banks on behalf of their (anonymised) customer banks.[(16)] Transactions originated by the largest customer bank makes up 4.2% of current CHAPS values, and those originated by the five largest customer banks make up 15.7%.

---

[(16)] The values used are likely to be biased downwards. It is plausible to assume that customer banks have to make payments to their settlement banks and *vice versa* – these payments will, by definition, be internalised. This internalisation would disappear once the customer banks became settlement banks, so the overall values settled through CHAPS are likely to be higher.

Charts 4 and 5 depict the linear and cubic equations respectively. In the linear specification, the 95% confidence interval for the coefficient is between 2.7 and 3.3 (with no constant). That implies that for any given percentage increase in values settled, the percentage of extra liquidity needed is three times as high. This assumption becomes less plausible as we predict larger increases – one would expect to find some deceleration in the rate of liquidity needed as tiering tends towards zero. The cubic specification accounts for this deceleration. The linear coefficient lies in a 95% confidence interval between 3.74 and 4.43, and the cubic one between -0.012 and -0.007. In fact, the deceleration captured by the latter negative coefficient is so strong that there is a slight percentage decrease in liquidity needs for the last two points in our forecast. Chart 6 depicts the 95% confidence intervals around the forecast, while Table D gives the values of extra liquidity needs that an increase in the number of settlement banks would bring about.

**Chart 4: Predicted changes in liquidity needs based on changes in value settled – linear prediction**

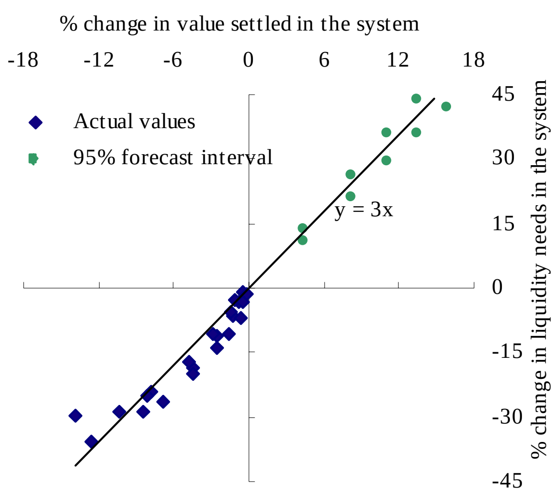


Sources: Payments database and Bank calculations.

**Chart 5: Predicted changes in liquidity needs based on changes in value settled – cubic prediction**

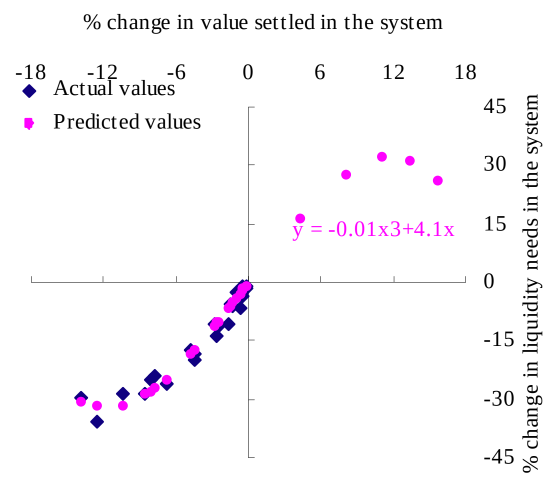


Sources: Payments database and Bank calculations.

**Chart 6: Forecast changes in liquidity needs and confidence intervals – cubic prediction**

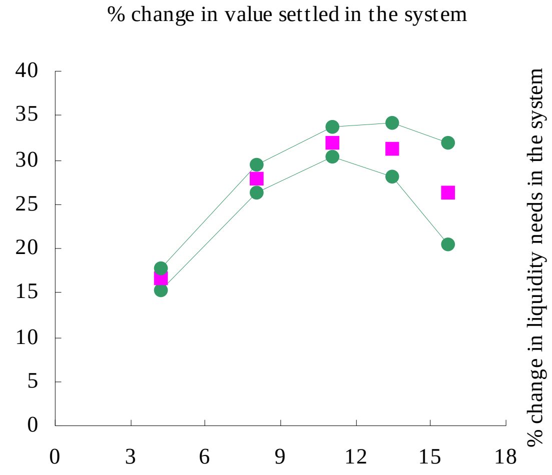


Sources: Payments database and Bank calculations.

**Table D: Changes in value settled and changes in liquidity needs**

| | Change in values settled, per cent (in £ billion) | Change in liquidity needs, per cent (in £ billion) | |
|---|---|---|---|
| **No. of customer banks joining CHAPS** | | **Linear** | **Cubic** |
| 1 | 4.2 (£8.2 bn) | 12.6 (£2.1 bn) | 16.5 (£2.7 bn) |
| 2 | 8.0 (£15.6 bn) | 23.9 (£4.0 bn) | 27.8 (£4.6 bn) |
| 3 | 11.1 (£21.4 bn) | 32.9 (£5.4 bn) | 32.0 (£5.3 bn) |
| 4 | 13.4 (£26.1 bn) | 40.0 (£6.6 bn) | 31.2 (£5.2 bn) |
| 5 | 15.7 (£30.5 bn) | 46.7 (£7.7 bn) | 26.2 (£4.3 bn) |

Sources: Payments database and Bank calculations.

How valuable these liquidity savings are depends on how scarce liquidity is. Data show that there is a lot of spare liquidity (defined as the difference between the maximum liquidity posted intraday and the maximum liquidity used intraday) in the system (see Chart 7). Banks posting more liquidity than needed is an indication that liquidity in the system is cheap. This may well be the case, especially for UK-owned banks subject to the Sterling Stock Liquidity Regime (SLR) by the regulator. Under the SLR, banks must hold a stock of eligible liquid assets overnight. The list of eligible assets broadly coincides with assets that can be used as collateral to obtain intraday credit with the Bank of England. If there is practically no opportunity cost in using the eligible assets intraday, then banks may decide to post more liquidity than needed. Foreign banks operating in the United Kingdom, on the other hand, are subject to a maturity mismatch approach, which requires that they have incoming liquidity to fund known outflows. Intraday liquidity may be more expensive for them.

**Chart 7: CHAPS spare liquidity**

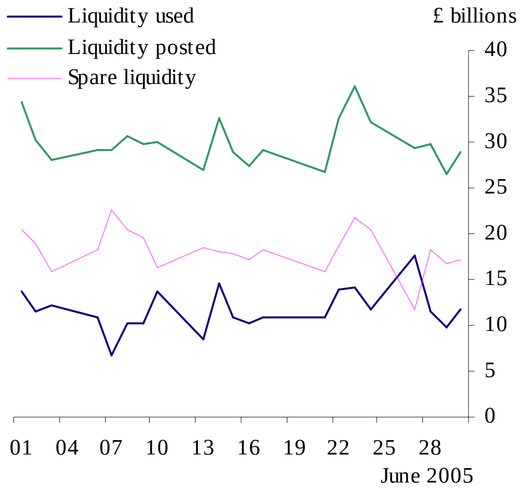


Sources: Payments database and Bank calculations.

In June 2005, the average spare liquidity was £17 billion (£18 billion median). The standard deviation over the month was £5 billion. The differences across banks are marked. The settlement bank with most spare liquidity had an average of £5 billion, that with the least just £10 million. Given these figures, the increase in liquidity needs of up to £8 billion suggested by our forecasts do not appear, at first glance, a disproportionate price to pay for the potential reductions in node and credit risk that a decrease in tiering in CHAPS would bring at the system level.

### *4.5 Liquidity pooling versus internalisation*

The liquidity gains observed in a more tiered system can stem from either an increase in the pooling of liquidity or from internalisation of payments. To our knowledge, no paper in the literature has estimated which of these two effects drives the savings.

We cannot disentangle the proportion of liquidity gains due to pooling and to internalisation with absolute precision. The reason for this ambiguity is that the amount saved at each transaction will depend crucially on the liquidity position of the settlement bank at the precise point in time when the payment takes place. But we can calculate the upper and lower bound for the proportion of savings due to each effect. This information will give a clear indication of each factor's relative size. The appendix gives a detailed explanation of how the intervals are calculated and why they are the closest we can get to quantifying the relative savings due to pooling and to internalisation.

Our results show that the vast majority of liquidity savings are caused by pooling. Table E displays the estimated intervals for the proportion of liquidity savings in CHAPS that are due to each factor. The average is calculated as the daily mean over our sample period and across all scenarios. Pooling is, on average, eight times bigger than internalisation. Table E also presents the upper and lower bound for the minimum (maximum) daily savings for the scenario with the minimum (maximum) savings. The lower bound of the minimum savings due to pooling (63%) is clearly higher than the upper bound of the maximum savings due to internalisation (37%). We can therefore be confident that liquidity pooling accounts for most of the savings observed.

**Table E: Estimated range of savings due to internalisation and to liquidity pooling**

| | Internalisation | Liquidity pooling |
|---|---|---|
| Average | 1%-22% | 78%-99% |
| Minimum | 0%-2% | 63%-96% |
| Maximum | 4%-37% | 98%-100% |

Sources: Payments database and Bank calculations.

Finally, Charts 8 and 9 show that the proportion of savings due to internalisation increases with the number of new customer banks assigned to each major settlement bank.

**Chart 8: Ranges of liquidity savings due to pooling (%) when banks are assigned to major bank 1**

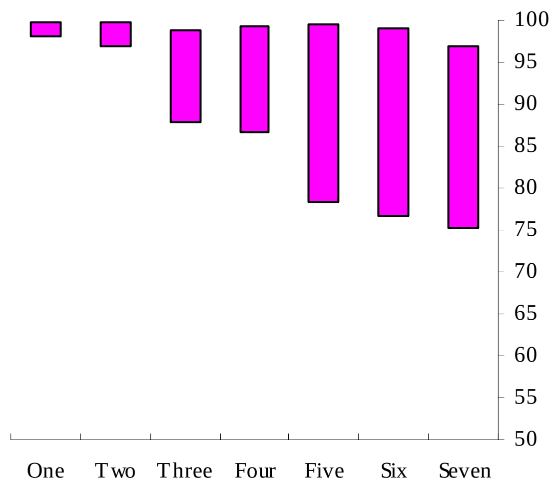


Sources: Payments database and Bank calculations.

**Chart 9: Ranges of liquidity savings due to pooling (%) when banks are assigned to major bank 2**

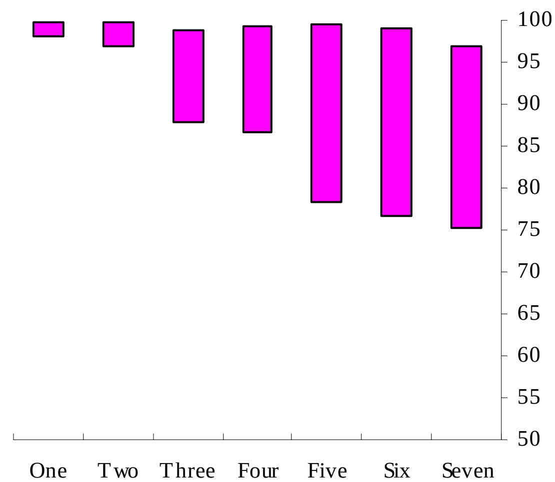


Sources: Payments database and Bank calculations.

## 5 Conclusions

This paper uses a simulation approach to quantify the impact of a change in the degree of tiering in the structure of the UK large-value payment system on concentration risk, credit risk and liquidity efficiency under normal circumstances. It does so by creating artificial scenarios where the number of direct participants in CHAPS Sterling is reduced one by one, thus increasing the degree of tiering.

The new customer banks are assigned to either one or a combination of the large (by value of payments processed) settlement banks.

The results show that concentration risk would rise substantially in what is already a highly concentrated system. The increase in risk is slightly smaller than the share of payments accounted for by the settlement banks becoming customer banks because of internalisation effects. As for credit risk, our figures have confirmed previous analysis at the Bank that, under normal circumstances, the size of intraday exposures compared with the settlement bank's capital is very small and therefore the likelihood of contagion remote. But, more importantly, our analysis has shown that the increase in credit risk brought to the system by settlement banks leaving CHAPS bears little relationship to the values settled by each individual bank. The key determining factor of the size of intraday credit exposures is the pattern of intraday flows of second-tier banks – a variable that central banks do not observe.

Increasing the degree of tiering in CHAPS leads to substantial liquidity savings. Our analysis has shown that the vast majority of the savings are due to liquidity pooling rather than internalisation. Moreover, the clear relationship between changes in values settled and liquidity needs shown by our simulations make it possible to project what would happen if current customer banks joined CHAPS as settlement banks. We estimate that liquidity needs could increase by £8 billion in aggregate if as many as five large banks (in terms of values of payments processed) joined CHAPS. While this figure is significant, it is only a fraction of the £17 billion spare liquidity posted on average in the system as a whole during the same time period. The paper has not attempted to establish the relationship between normal and stressed liquidity needs, or how liquidity insurance should be regulated. Indeed, it is worth bearing in mind that liquidity costs and efficiencies savings do not in and of themselves make a tiered system optimal from a welfare perspective. Under stressed circumstances, tiering can create additional operational and liquidity risks that are not discussed in this paper.

Two possible extensions of this paper stand out as promising. One involves simulating operational outages of individual banks in the more concentrated scenarios in order to analyse how robust CHAPS would be in terms of liquidity compared with the current situation. This exercise will complement the direct effects of an individual bank's operational outage (share of payments usually settled by that bank that are disrupted) by adding second-order disruptions through potential liquidity shortages in the system. A second possible avenue for further work is to model the relationship between settlement banks and the new customer banks as an ancillary system. Such an exercise would provide us with a more detailed picture of the sources of liquidity needs in corresponding banking.

## Appendix: Liquidity pooling versus internalisation

As stated in the main body of the paper, concentration in payment systems leads to liquidity savings because of two effects: internalisation of payments and liquidity pooling. When more payments are internalised, the settlement bank does not need to find the liquidity to fund them and the liquidity needs of the system decrease. And a settlement bank with a larger number of transactions is likelier to have incoming payments funding its outgoing transactions. As a result, the maximum liquidity need intraday is likely to be lower.

In practice, though, it is very difficult to disentangle both effects since the origin of the precise liquidity savings at any point in time during a day depends crucially on the liquidity position of the settlement bank at that point in time. A payment sent, say, from settlement bank A to settlement bank B will be internalised if settlement bank A becomes a customer of settlement bank B. This, in theory, could reduce the liquidity needs of settlement bank B. But it will not do so if bank B was already in liquidity surplus when the saving took place. We need to bear in mind that liquidity needs are defined as the intraday *maximum*, not average. Conversely, a potential saving caused by liquidity pooling may not affect the maximum intraday liquidity needs if the settlement bank was already in surplus. In what follows we run through a stylised example that attempts to disentangle these two effects.

Table A1 lists the payment structure in a very simple system with four settlement banks and eight transactions at different times within the day. We call this ‘the benchmark model’. We are interested in comparing the liquidity needs in the benchmark model with the liquidity needs in a situation when bank AA becomes a customer bank of bank MM and all payments between them are internalised. We have set up the transactions in such a way that we only need to compare the liquidity needs of bank MM before and after AA becomes its customer bank to see how the liquidity needs of the system changes. AA starts receiving money before having to send any payments. AA always has surplus liquidity, thus not affecting the maximum intraday liquidity needs of the system. The comparison is not affected by the liquidity needs of BB and CC, either – they are the same before and after AA becomes MM’s customer bank.

**Table A1: Payments structure – benchmark model**

| Time | Payer | Payee | Amount |
|---|---|---|---|
| 1 | MM | AA | 90 |
| 2 | MM | CC | 10 |
| 3 | BB | AA | 50 |
| 4 | MM | AA | 25 |
| 5 | CC | AA | 100 |
| 6 | AA | BB | 10 |
| 7 | AA | MM | 10 |
| 8 | MM | BB | 100 |

Table A2 calculates the liquidity needs under different calculations. As explained we can just focus on MM's liquidity needs for comparison purposes.

**Table A2: Liquidity needs**

| **Benchmark** **[A]** | | **AA customer of MM, payments between them internalised** **[B]** | | **As benchmark but excluding payments between AA and MM** **[C]** | |
|---|---|---|---|---|---|
| Time | Net debit | Time | Net debit | Time | Net debit |
| 1 | 90 | 1 | 0 | 1 | 0 |
| 2 | 100 | 2 | 10 | 2 | 10 |
| 3 | 100 | 3 | -40 | 3 | 10 |
| 4 | 125 | 4 | -40 | 4 | 10 |
| 5 | 125 | 5 | -140 | 5 | 10 |
| 6 | 125 | 6 | -130 | 6 | 10 |
| 7 | 115 | 7 | -130 | 7 | 10 |
| 8 | 215 | 8 | -30 | 8 | 110 |
| Maximum liquidity needs to be able to settle all transactions in RTGS | | | | | |
| | 215 | | 10 | | 110 |

The first two columns in Table A2 refer to the benchmark situation (when AA is a direct CHAPS member). Each row gives, after each transaction, the net debit position of bank MM. For example, after transaction 1, MM bank is in a net debit position of 90 (since it has to pay to AA 90 without having received any payment yet). At time 2, MM has to make a further payment of 10, increasing its net debit position to 100. No modifications occur after transaction 3 since it does not involve bank MM. After transaction 4, MM has to make another payment of 25 further increasing its net debit position to 125. Transactions 5 and 6 do not affect MM's net debit position, whereas transaction 7 decreases it to 115 due to the payment of 10 received from AA. There is a further increase in MM's net debit position of 100 to 215 after transaction 8. The maximum liquidity need is therefore 215.

It is plausible to think that the difference between calculations [A] and [C] would give us the liquidity savings due to internalised payments. The only difference between those two calculations is the payments between the settlement bank and the bank to become customer of the settlement bank; this gives a saving in liquidity of 105 (215-110). Likewise, the difference in liquidity needs between calculations [B] and [C] would show the savings in liquidity due to liquidity pooling. The difference between that pair of calculations corresponds to the transactions between any other settlement bank and the new customer bank. This gives a saving in liquidity of 100 (110-10). The total savings in liquidity are thus 205 (215-10), that is, the sum of liquidity savings due to internalisation and the liquidity savings due to liquidity pooling.

Let us consider now one more calculation [D], as in Table A2 as shown below: the liquidity needs of the system when AA becomes a customer bank of MM but the payments between them continue to be sent to RTGS. We could assume now that the difference between calculations [D] and [A] would give us the liquidity savings due to liquidity pooling, since in calculations [D] we do not allow for internalisation. Liquidity savings are now 125 (215-90).

**Table A2 – extra column: Liquidity needs**

| AA customer of MM, payments between them sent to RTGS [D] | |
|---|---|
| Time | Net debit |
| 1a | 90 |
| 1b | 0 |
| 2 | 10 |
| 3 | -40 |
| 4a | -15 |
| 4b | -40 |
| 5 | -140 |
| 6 | -130 |
| 7a | -120 |
| 7b | -130 |
| 8 | -30 |
| Maximum liquidity needs to be able to settle all transactions in RTGS | |
| | 90 |

The difference in liquidity needs between calculations [D] and [B] could be interpreted as savings in liquidity due to internalisation only. This yields a value of 80 (90-10). Adding the savings due to liquidity pooling and internalisation savings gives again total liquidity savings.

As we can see, these two methods of disentangling the liquidity savings yield different results. The reason for the difference is that when using columns [A], [B] and [C], we are assigning all savings whose cause we cannot identify to internalisation. When we use columns [A], [B] and [D], by contrast, we assigned them to liquidity pooling. We therefore conclude then that we cannot disentangle the liquidity savings due to internalisation and to liquidity pooling with absolute precision. We are able, though, to calculate upper and lower bounds for the estimates of the savings due to either internalisation or pooling in each of our scenarios. In the example presented above, the lower bound for internalisation savings is given by D-B (80), and upper bound by A-C (105). The lower bound for savings due to liquidity pooling is given by C-B (100) and the maximum by A-D (125).

## References


**Bedford, P, Millard, S and Yang, J (2004)**, 'Assessing operational risk in CHAPS Sterling: a simulation approach', *Bank of England Financial Stability Review*, December, pages 135-43.

**Flannery, M J (1996)**, 'Financial crises, payment system problems, and discount window lending' (in Controlling Risks on Large-Value Transfer Systems), *Journal of Money, Credit and Banking*, Vol. 28, No. 4, Part 2: Payment systems research and public policy risk, efficiency, and innovation (Nov. 1996), pages 804-24.

**Harrison, S, Lasaosa, A and Tudela, M (2005)**, 'Tiering in UK payment systems', *Bank of England Financial Stability Review*, June, pages 63-72.

**Jackson, J and Manning, M (2007)**, 'Central bank intra-day collateral policy and implications for tiering in RTGS payment systems', forthcoming in Haldane, A, Millard, S and Saporta, V (eds), *The future of payment systems*, Routledge.

**James, K (2003)**, 'A statistical overview of CHAPS Sterling', *Bank of England Financial Stability Review*, June, pages 115-21.

**Kahn, C M and Roberds, M (2005)**, 'Payments settlement: tiering in private and public systems', University of Illinois and Federal Reserve Bank of Atlanta, *mimeo*.

**Leinonen, H (2005)**, 'Liquidity, risks and speed in payment and settlement systems – a simulation approach', *Bank of Finland Studies*, E:31, pages 49-57.

**Selgin, G (2004)**, 'Wholesale payments: questioning the market-failure hypothesis', *International Review of Law and Economics*, Vol. 24, pages 333-50.