DRAFT – PLEASE DO NOT QUOTE

THE IMPACT OF SYSTEM DESIGN ON TIERING INCENTIVES

Robert Arculus, Jennifer Hancock and Greg Moran

Research Discussion Paper

Reserve Bank of Australia

We are grateful to Mark Manning, Andrea Brischetto and other colleagues at the Reserve Bank for their helpful comments. We would also like to thank Sammy Yousef, Rod Phillips and the Bank of Finland staff for their technical advice. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of Australia. The authors are solely responsible for any errors.

Authors: arculusr, hancockj or morang at domain rba.gov.au

Abstract

Tiering – where an institution does not participate directly in the central payments system but instead settles its payments through an agent who does – is a significant issue for payment system regulators. Indirect settlement can provide efficiency advantages, most especially in liquidity savings, but equally it can increase risk. This paper uses simulation analysis to explore the impact of payment system design on institutions' incentives to tier. We find some evidence to support our hypothesis that by reducing the liquidity benefits of tiering, the liquidity-saving mechanisms in Australia's real-time gross settlement (RTGS) system, RITS, have contributed to a low level of tiering relative to RTGS systems overseas. We also find that tiering results in only small increases to the level of concentration in RITS, and thus does little to increase concentration risk. In terms of credit risk, we find that tiering creates substantial two-way exposures between clients and their settlement banks, although the effect of system design on these exposures is mixed and not significant overall. Finally, we provide some discussion on the costs and benefits of tiering in RITS, and how a central bank might attempt to choose an optimal level of tiering.

.

JEL Classification Numbers:

Keywords:

Table of Contents

THE IMPACT OF SYSTEM DESIGN ON TIERING INCENTIVES

Robert Arculus, Jennifer Hancock and Greg Moran

1. Introduction

'Tiering' – where an institution does not participate directly in the central payments system but instead settles its payments through an agent who does – is a significant issue for payment system regulators. Indirect settlement can provide efficiency advantages, most especially in liquidity savings, but equally it can increase risk.

This paper uses simulation analysis to explore the impact of payment system design on institutions' incentives to tier. Specifically, we examine how the design of Australia's real-time gross settlement (RTGS) system – the Reserve Bank Information and Transfer System (RITS) – affects the potential liquidity savings from tiering. We also attempt to quantify the increases in concentration and credit risk caused by an increase in tiering in RITS. Finally, we offer an attempt to provide more intuitive interpretations of these costs and benefits, alongside a consideration of how a central bank might try to set a socially optimal level of tiering. Our analysis is intended to shed light on the present level of tiering in RITS, as well as inform policymakers in regard to rules that restrict tiering.

The degree of tiering varies across payment systems. The CHAPS Sterling system in the UK, for instance, is relatively highly tiered, with only 14 direct participants (not including the Bank of England) making payments on behalf of several hundred other institutions (BIS, 2009). The US Fedwire system, in comparison, has a fairly flat payments structure, with several thousand direct participants. The relatively low level of tiering in RITS is somewhat of a puzzle. Since tiering restrictions were relaxed in 2003, very few institutions have opted to settle indirectly. Several hypotheses have been proposed, including that the liquidity-saving mechanisms employed in the system reduce the incentives to tier.

The remainder of the paper is structured as follows. In Section 2, we briefly review the literature concerning the costs and benefits of tiering in payment systems. In Section 3, we provide an overview of RITS. Section 4 outlines our simulation methodology, adapted from Lasaosa and Tudela (2008). Section 5 presents estimates of liquidity savings from tiering under different system designs. Section 6 analyses the drivers of these liquidity savings. Section 7 presents estimates of the increases in concentration and credit risk associated with an increase in tiering. In Section 8, we discuss a framework for weighing the costs and benefits of tiering. Section 9 concludes.

2. The Costs and Benefits of Tiering

2.1 Liquidity Benefits

The liquidity cost of direct participation is a major factor in an institution's decision to settle directly or indirectly in a payments system. Systems that operate on an RTGS basis require participants to hold substantial liquidity in order to cover payments as they arise. In RITS, intraday liquidity is provided through interest-free reciprocal purchase agreements ('repos') with the Reserve Bank of Australia (RBA), but in using this facility participants incur the opportunity cost of the collateral posted. Insofar as different participants face different opportunity costs of collateral, there may be gains from trade in having one participant (the 'settlement bank') send and receive payments on behalf of another (the 'client').

As discussed in Jackson and Manning (2007), Adams, Galbiati and Giansante (2008) and Lasaosa and Tudela (2008), the two sources of liquidity savings from tiering are:

- 'Internalisation' savings, resulting from payments between the client and the settlement bank being settled across the settlement bank's books, rather than sent to the RTGS system; and

- 'Liquidity pooling' savings, as the combined payment flows allow more payments to be funded from receipts. Unless the client and settlement bank's peak intraday liquidity requirements occur simultaneously, tiering requires less liquidity than the sum of their individual peak requirements as payments received by one can be used to fund payments by the other.

In addition to liquidity savings, Adams *et al* (2008) identify several other factors that can encourage tiering. For instance, if there are economies of scale in liquidity costs (that is, the cost of liquidity is a declining proportion of total liquidity used), then tiering is likely to increase. Similarly, the narrower the range of acceptable collateral, the higher collateral costs are likely to be, which will also increase tiering.[1] Adams *et al* also identify that the exact pattern of tiering is a consequence of the individual choices of institutions, with multiple equilibrium outcomes possible depending on the order in which decisions are made.

From a slightly different perspective, Jackson and Manning (2007) focus on the decision by the central bank of whether or not to demand full collateralisation of intraday loans. The higher the probability of participant default, or the lower the costs of posting the collateral, the more likely demanding full collateralisation is to be optimal. Under tiering, however, requiring full collateralisation can still be optimal even if only a strict subset of participants has low collateral and monitoring costs, as those participants can then conduct payments on behalf of others. As an example, some institutions may be 'natural' holders of certain types of collateral-eligible securities – that is, the holding of such securities is optimal in terms of their business profile and preferred risk-management strategy, or is required by regulation – while other institutions may only hold such securities for the purposes of accessing central bank liquidity. In such a case, it is likely that the former institutions would have lower collateral costs than the latter.[2]

---

1 In recent years the RBA has widened the range of collateral acceptable for repos, from Commonwealth Government securities to include securities from state governments, certain supranational organisations, and certain Authorised Deposit-taking Institutions (ADIs). However, as this has been partly in response to a declining supply of Commonwealth Government securities, the net effect on tiering incentives is ambiguous.

2 Even for natural holders of such securities, however, collateral costs are not zero. Market developments while the securities are collateralised could lead to changes in the institution's preferred mix of holdings. Since it will not be possible to alter the mix until the repos have been repaid, there is some inevitable degree of market risk in posting collateral.

2.2 Monitoring Benefits

The benefits of tiering are not restricted to liquidity savings. Kahn and Roberds (2008) provide a theoretical model of payment service networks suggesting that, as a means to reduce credit risk, there is some scope to replace the posting of costly collateral with inter-agent monitoring. In a tiered network – especially one with unsecured credit and overdraft facilities – settlement banks will have an incentive to monitor the credit-worthiness of their clients. If these clients were to participate directly in the system, however, this incentive is eroded. It might be assumed that the ability of private-sector institutions to monitor each other is greater than the monitoring capacity of the central bank, which is plausible given the substantial interactions and information flows that occur between private-sector institutions in everyday business. If so, and if such monitoring is less costly than posting collateral, then – as emphasised in Kahn and Roberds – it may be optimal to even restrict access to the primary payment settlement network, with the intent of forcing monitoring activity onto participants. Chapman, Chiu and Molico (2008) make a similar point regarding the potential increases in efficiency that can be brought about by clients establishing their credit-worthiness through interactions with settlement banks.

2.3 Potential Risk Increases

Tiering in payment systems can also impose significant risks. Given default by an individual participant in a tiered network can result in significant spillover costs to the rest of the system, the incentive to monitor may not fully incorporate all relevant externalities. There is also some degree of legal risk, in that there is not the same legal certainty regarding the status of payments settled in commercial bank money. For instance, under the 'zero hour' rule a court may date the bankruptcy of an institution from the midnight before the bankruptcy order is made. In Australia a specific piece of legislation, the *Payments Systems and Netting Act*, allows the RBA to protect payments that occur in RITS from the application of this rule, but payments settled across the books of the settlement bank do not have the same protection. Tiering can also create problems of industrial organisation; Chande, Lai and O'Connor (2006) note that a settlement bank may also be a competitor with its clients in the market for retail payment services, though in their model the incentive to exploit this position is mitigated by a desire to not increase the credit risk settlement banks face from their clients. Business risk can also be an issue, in that the choice of a settlement bank to exit the market may cause a greater disruption to the payments system than would result were tiering not present.

From the point of view of the central bank in its presumed role as maximiser of total social welfare, credit risk is one of the most significant issues. In a sense, as moving to an RTGS system decreases credit risk at the expense of increased liquidity costs (see Kahn and Roberds 1998), tiering represents the possible reintroduction of credit risk. Reliance on private credit thus rises, which may complicate events in the case of a crisis (though conversely, higher reliance on public credit directly exposes taxpayers to financial system complications). Note that this credit risk is two-way. Both the settlement bank and its client are exposed to the failure of the other; the former because of the offered intraday credit and the latter due to the settlement bank's role as holder of the relevant accounts. Harrison, Lasaosa and Tudela (2005) attempt to quantify the exposure of settlement banks in the UK CHAPS system, finding that the risk is not substantial under normal operating conditions, but has the potential to rise considerably under extreme circumstances.

Finally, tiering also increases concentration risk. An operational problem at a participant, for instance, may result in that participant becoming a 'liquidity sink'.[3] The more liquidity is concentrated into fewer participants, the more severe the impact of such a problem is likely to be. On the other hand, as a tiered network depends less on the central infrastructure, it may allow some payments to still go ahead in the event of a central system failure (although not in central bank money). The net effect is ambiguous, but certainly tiering has the potential to significantly change the dynamics of system disruptions and participant failures.

---

[3] Liquidity sink describes a situation where a participant is able to receive but not send payments, and thereby drains liquidity from the system.

3. Australia's RTGS System

RITS has operated as an RTGS system since 1998. It features a centralised queue with bilateral offsetting and a liquidity reservation feature. The settlement algorithm in RITS uses a next-down loop to test each queued transaction for both individual and simultaneous settlement with up to ten offsetting transactions.[4] Participants can also reserve liquidity for priority payments using a ‘sub-limit’; payments that the participant has set to ‘priority’ will be tested for settlement against the participant’s entire balance, whereas ‘active’ payments are tested only against balances in excess of the sub-limit set by the participant. There is also a third status of ‘deferred’ available to participants, which allows them to send payments to RITS before the participant is ready to settle the payment. Deferred payments are not tested for settlement until their status is changed to either active or priority.

Initially, direct access to RITS was only available to banks, with all banks required to settle their RTGS payments using their own settlement account.[5] In 1999, following the recommendations of the Wallis Inquiry into Australia’s financial system, access was broadened to allow third-party payment providers and non-bank Authorised Deposit-taking Institutions (ADIs) to participate directly in RITS.[6] The Wallis Inquiry also resulted in the creation of the Australian Prudential Regulation Authority (APRA), which prudentially regulates all ADIs – banks, building societies, credit unions and special third-party providers of payments services.[7] While all ADIs can now become direct participants in RITS, only banks are required to hold a settlement account at the Reserve Bank.

Notwithstanding the broad scope of participation, payments through RITS are fairly highly concentrated, with the major domestic banks accounting for almost

---

4 In July 2009, the RBA added a Targeted Bilateral Offset algorithm, which allows participants to select specific payments for bilateral offset.

5 Special Service Provider accounts were set up for the building society and credit union industry associations, to allow building societies and credit unions to settle indirectly through these associations.

6 See Reserve Bank of Australia (1999) for more information.

7 APRA also has regulatory responsibilities for general insurance and reinsurance companies, life insurance, friendly societies, and most members of the superannuation industry.

60 per cent of the value of all payments made. Indeed, payments just between the four major domestic banks account for around a third of all payments. The direction of payment flows tends to have a similarly skewed distribution at the individual level; the average RITS participant, for instance, makes more than half of its payments to just three other participants (measured in value terms).

Since 2003, settlement account-holders whose RTGS payments comprise less than 0.25 per cent of the total value of RTGS transactions have been permitted to settle via an agent.[8] Prior to this, banks were prohibited from tiering. Despite the relaxation in policy, data for 2008 suggest that only six of the 34 participants eligible to settle indirectly chose to do so. Moreover, it appears that most participants in RITS which have chosen to settle indirectly have been required by their settlement banks to pre-fund their obligations (largely negating the benefits and costs of tiering discussed in section 2). It is not clear why this should be the case, or whether it should be expected to continue in the future.

4. Methodology

Our methodology is adapted from Lasaosa and Tudela (2008), who undertook a study of the effects of tiering in the UK CHAPS Sterling system using the Bank of Finland’s payment system simulator. Given the already highly-tiered nature of CHAPS Sterling, Lasaosa and Tudela focus their analysis on the potential impact of a reduction in tiering. Conversely, in the local context, we are interested in analysing why only a few participants have chosen to tier so far, and thus we investigate the potential costs and benefits from an increase in tiering.

The sample period is the month of January 2008, containing 21 business days over which 623,860 individual transactions took place with a total value of around $4.04 trillion. Excluding the four largest participants, as well as the RBA, CLS Bank and the settlement accounts of the equity and futures systems, there are 49 participants altogether that are considered as ‘tiering candidates’ in our simulations.

---

8 See Australian Prudential Regulation Authority and Reserve Bank of Australia (2003) for more information.

4.1 Tiering Order

Although there is an array of plausible methods of selecting client institutions and their respective settlement banks (see Lasaosa and Tudela 2008 for examples), we chose to allocate institutions based on the value of payments sent and received. Institutions are tiered from smallest to largest in order of their share of all payments. Our reasoning is that larger institutions should generally have a lower opportunity cost of collateral as their banking operations naturally result in them holding more eligible securities on their balance sheet, which in turn gives them a competitive advantage in the market for providing payment services. This approach is also consistent with the current formulation of RBA policy, whereby only participants whose share of RTGS payments comprise less than 0.25 per cent of the total value of RTGS transactions are eligible to tier.

The settlement bank for each individual tiering candidate is chosen as the institution with which the institution to be tiered conducts the largest share of its payments – that is, its 'largest payments partner'. This approach is likely to maximise the value of payments which are internalised, although this not a mathematical certainty.[9]

In practice such allocation decisions would be interdependent. That is, each institution's choice of settlement bank could change depending on the choices of other institutions and the subsequent sizes of different tiered networks (Adams *et al* 2008 provide an interesting model of participant tiering choice). However, it was found that attempting to account for this would have minimal material effect; for instance, when each client institution was assigned to its largest payments partner with the choices of all smaller institutions taken as given, the choice of settlement bank was only altered on four occasions.

Simulating these tiering scenarios involves replacing the client with its settlement bank in the transaction data, and deleting any payments between the settlement

---

9 A further scenario was considered, where tiering would have been conducted in the order of the share of total volumes rather than values. However, the tiering order based on volumes was not found to be especially different to the one based on values, therefore it was not pursued further.

bank and itself which result (i.e. payments that were previously between the client and the settlement bank). For multilateral batches, the settlement bank's revised position is calculated as the net of its original position and that of the clients it represents.

4.2 System Designs

To test our hypothesis that RITS' liquidity-saving features decrease participants' incentives to tier, we simulate tiering in four RTGS system designs (Table 1). Details of how RITS' bilateral offset and sub-limit features have been included in the simulations are contained in Appendix A.

Table 1: RTGS System Designs

	Central queue	Bilateral offset	Sub-limits
Pure RTGS	-	-	-
Central queue only	X	-	-
Bilateral offset only	X	X	-
RITS replica	X	X	X

4.3 Liquidity

One of the limitations of simulation is that it does not allow participants to change their behaviour in response to an external event – like increased tiering. As such, we must set the maximum available liquidity for each participant exogenously. In general, we assume that each participant begins each day with a zero balance, but has a credit limit, which varies throughout the day, equal to the actual value of its opening settlement account balance and intraday repos at each point in time during that day.[10] To prevent payments that do not settle immediately in a pure RTGS system from being rejected and remaining unsettled at the end of the day, we have assumed that all participants have access to unlimited liquidity in the pure RTGS simulations. In the tiering scenarios, we reason that the settlement bank does not

10 The RBA, CLS Bank, and the settlement accounts of the equity and futures systems are provided with unlimited funds in all system designs.

have access to collateral on its clients' balance sheets, therefore we leave the settlement bank's credit limit unchanged.[11]

We find that a small proportion (less than 1 per cent) of payments remain unsettled at the end of most simulated days. This is because payment settlement times differ across system designs while credit limits are set exogenously. To allow these payments to settle, we relax credit limits at the end of each day.

5. The Impact of Tiering on Liquidity Usage

5.1 Measuring Liquidity Usage

We use the sum of individual participants' peak intraday liquidity requirements as our measure of system liquidity usage. Given the simulations assume nil opening balances and intraday credit is provided via credit limits, a participant's peak intraday liquidity requirement is equivalent to its minimum intraday account balance. This measure seems fairly reasonable if, for instance, it is considered that funding any peak liquidity requirement for however brief a period will still require the accessing of funds to cover that position, which in turn requires holding sufficient collateral that may incur an opportunity cost.[12]

---

[11] Alternatively, we could have assumed that the settlement bank increased its credit limit (for example, by the value of its clients' credit limits). Indeed, preliminary simulations were run using additive credit limits, but this resulted in quite substantial and unrealistic increases in liquidity usage under tiering. Therefore, our preference has been to remain with fixed, non-additive credit limits.

[12] If, however, the costs of collateral are seen less as a function of foregone uses, and more as due to the fact that using securities as collateral in a repo might hamper the rebalancing of a portfolio in the event of a shift in risk preferences, our measure is less appropriate, since costs would then partially be a function of the length of time repos were accessed for.

5.2 Estimates of Liquidity Savings

5.2.1 Cumulative tiering

We first look at the case where individual institutions are tiered cumulatively, from smallest to largest, according to their share of the total value of payments. Figure 1 shows liquidity usage over the range from no tiering to tiering all candidate institutions. For all scenarios in this case, the pure RTGS system is the most liquidity intensive and the RITS replica the least intensive. Average daily liquidity usage is $14.8 billion in the pure RTGS system under the benchmark no-tiering scenario, and falls to $7.2 billion when all candidate institutions are tiered. For the RITS replica system, liquidity usage falls from $12.6 billion to $5.6 billion. Liquidity usage in the other two system designs falls in between. Of these two, the bilateral offset only system clearly uses less liquidity for approximately the first 30 tiering scenarios. For subsequent scenarios, however, the presence of bilateral offset has almost no effect. This may be due to the increasing concentration of the system; Ercevik and Jackson (2009) find that liquidity recycling increases with system concentration, thus the need for bilateral offset decreases. The share of the total value of payments settled by bilateral offset falls from 28 per cent when there is no tiering to 13 per cent when all candidate institutions are tiered.

Figure 1

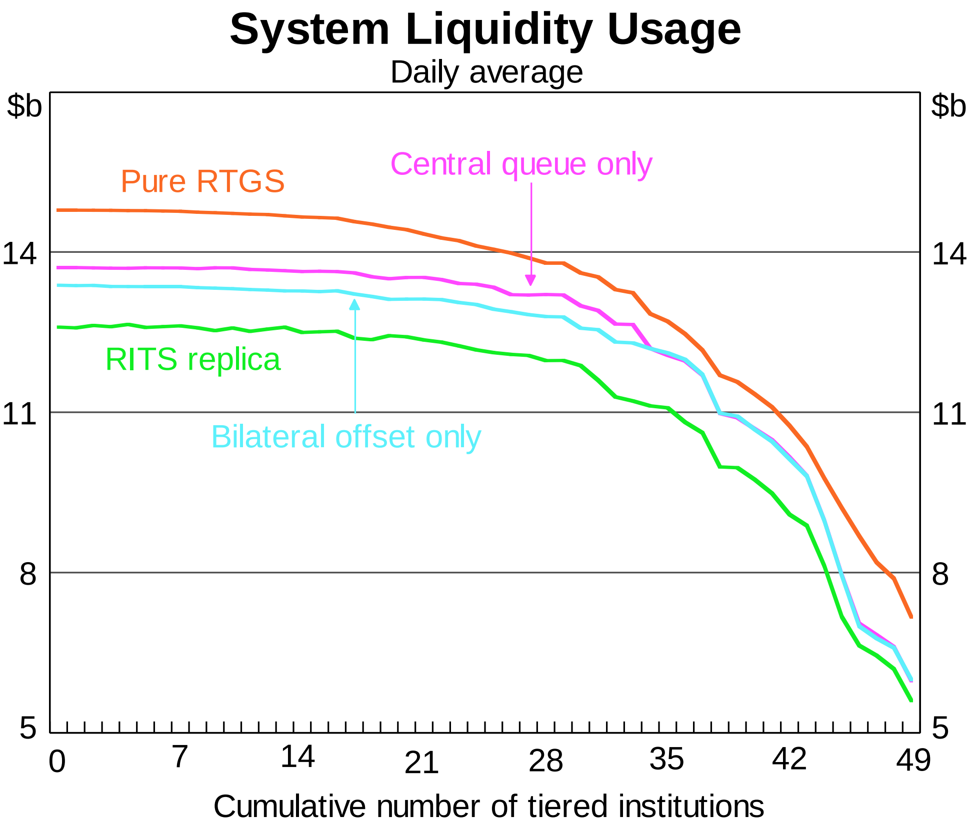

Source: RBA

Interestingly, tiering occasionally results in an increase in liquidity usage in the three systems with credit limits (Table 2). Moreover, such examples are restricted to the first 29 tiering scenarios. The most likely explanation is that, with credit limits fixed, a settlement bank is required to make more payments using the same amount of credit, which may result in settlement delays which in turn force other participants to increase their use of intraday credit. As larger institutions are tiered, this effect is likely to be outweighed by the larger liquidity savings.

Table 2: Marginal Changes in Daily Average System Liquidity Usage

Cumulative tiering

	Overall	Decreases			Increases		
	Average $m	Number	Average $m	Largest $m	Number	Average $m	Largest $m
Pure RTGS	-155.4	49	-155.4	-722.2	0	n/a	n/a
Central queue only	-157.8	43	-181.4	-1,015.9	6	11.0	24.0
Bilateral offset only	-150.7	44	-168.4	-1,020.1	5	5.4	15.3
RITS replica	-142.7	38	-193.2	-949.4	11	32.0	76.3

While a reduction in liquidity usage is not a necessary consequence of tiering, instances of increased liquidity usage are more likely a perverse consequence of the lack of behavioural response in the simulator, where agent behaviour is not optimised to the altering system structures. This is particularly true with respect to sub-limits, which in reality are likely to be a function of the size of payments a participant expects to send and receive on a particular day. As a liquidity reservation feature, sub-limits appear to exacerbate the problem. The RITS replica system experiences 11 increases in liquidity usage while the two systems without sub-limits experience around half that.

Similar to an exercise conducted by Lasaosa and Tudela (2008), we examine the relationship between the percentage changes in system liquidity usage and system value settled as tiering increases. However, in contrast to Lasaosa and Tudela, we calculate the change in value settled as the total value settled less the value of all payments to and from the most recently tiered institution, as opposed to the total value settled less payments between the most recently tiered institution and its settlement bank. We believe this alternative measure better captures both the liquidity savings from internalisation and liquidity pooling.

Lasaosa and Tudela aim to forecast the change in liquidity usage from a decrease in tiering, with the limitation that the highly tiered nature of CHAPS Sterling allows only around 20 observations over which to regress. By contrast, we are able conduct our regression over the 49 observations in our RITS replica and pure

RTGS systems (Table 3 and Figure 2). While Lasaosa & Tudela find the best fit is offered by a cubic specification, we find preferred specifications of polynomials of the fourth (in the case of the RITS replica) and fifth (for pure RTGS) order.

Table 3: Percentage change in liquidity needs regressed against percentage change in value settled in system*

System:	Constant	%Δ(settled)	%Δ(settled)^2	%Δ(settled)^3	%Δ(settled)^4	%Δ(settled)^5
RITS-replica	-0.09	2.44***	0.10**	0.00**	0.00**	
Pure RTGS	-0.24***	3.20***	0.23***	0.01***	0.00**	0.00*

* Percentage change refers to the change from a state with no tiering. *,** and *** refer to significance at the 10, 5 and 1 per cent levels, respectively, using White heteroskedasticity-consistent standard errors.

Figure 2

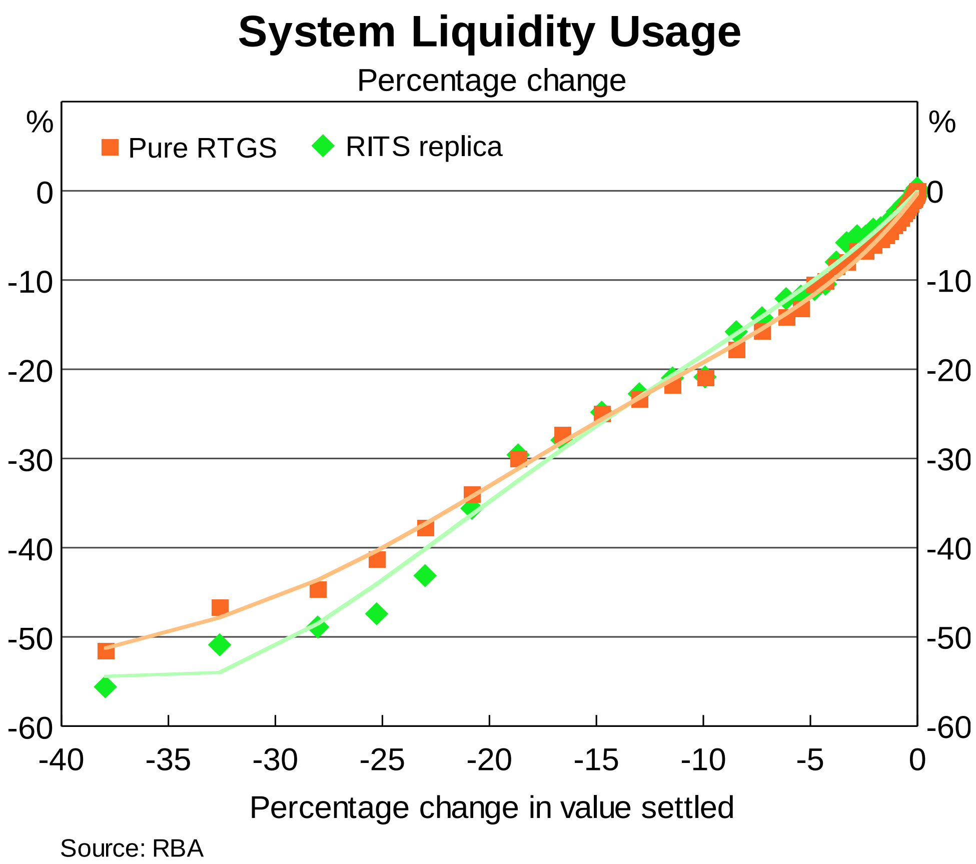

As Figure 2 shows, liquidity used in both systems tends to decline at a similar percentage rate as the value settled in the system falls (noting that the pure RTGS system is starting at a higher base of liquidity usage, as shown in Figure 1). The results suggest that, at least amongst approximately the smallest thirty participants, a 1 per cent reduction in the value settled in the system is associated with a decline in liquidity usage of around 2½ per cent in the RITS-replica and slightly more than 3 per cent in the pure RTGS system.

5.2.2 Tiering individual participants

We now look at the case where individual participants are tiered in isolation. Figure 3 shows liquidity usage with no tiering (0 on the x-axis) and with the tiering of one institution at a time in order of their share of the total value of payments. Again, the pure RTGS system is the most liquidity intensive and the RITS replica the least intensive for all scenarios in this case.

Figure 3

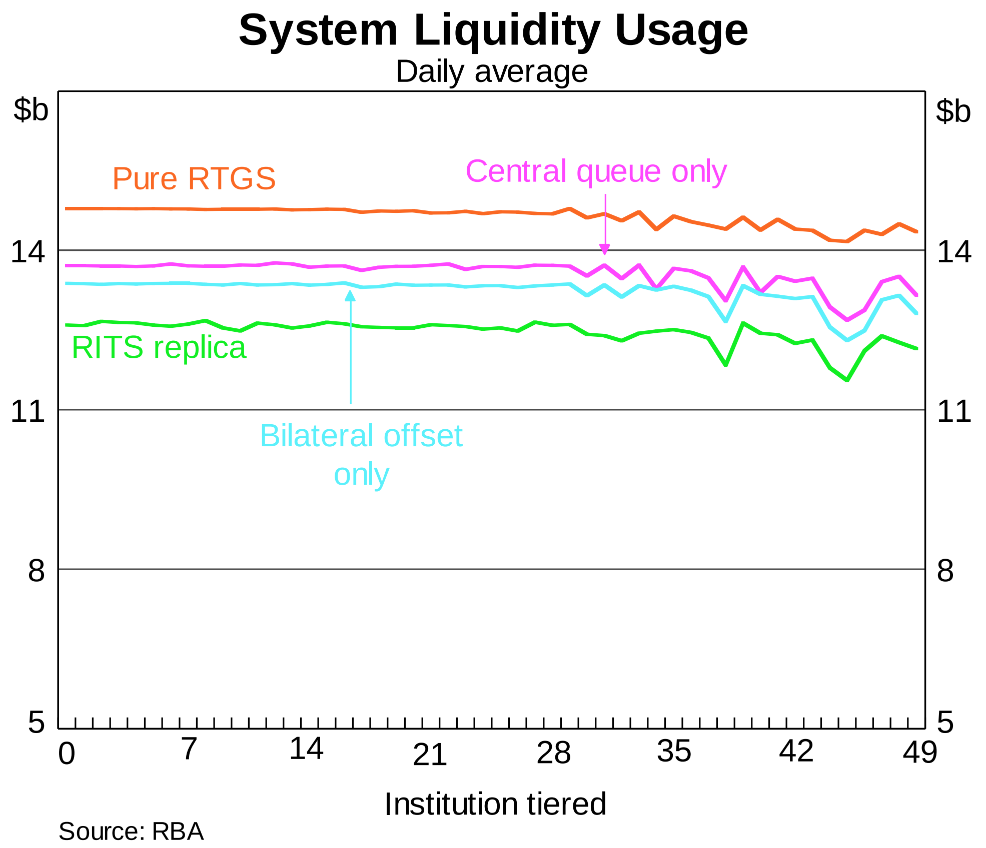

Tiering an individual participant can save up to $1 billion of liquidity, depending on the participant being tiered and the system design (Table 4). We again observe that tiering results in increases in system liquidity in some scenarios, although these instances are not restricted to the first 29 scenarios in this case.

Table 4: Marginal Changes in Daily Average System Liquidity Usage
Individual tiering

	Overall	Decreases			Increases		
	Average $m	Number	Average $m	Largest $m	Number	Average $m	Largest $m
Pure RTGS	-151.3	49	-151.3	-619.5	0	n/a	n/a
Central queue only	-142.6	38	-189.9	-1,023.8	11	20.5	53.0
Bilateral offset only	-154.7	47	-161.5	-1,075.1	2	4.6	8.3
RITS replica	-137.3	36	-200.1	-1,042.4	13	36.6	84.2

5.2.3 Summary

Overall our results provide some evidence in support of the hypothesis that the liquidity-saving mechanisms in RITS reduce the liquidity-saving incentive to tier. While our simulations clearly show that a central queue, bilateral offset and sub-limits reduce total system liquidity usage, it is not clear whether the *marginal* liquidity savings of tiering system participants are lower in RITS compared with systems with fewer or no liquidity-saving mechanisms. Across all 49 tiered participants, in both the cumulative and individual tiering cases, the average marginal liquidity saving is indeed lower in the RITS replica system than that in either the pure RTGS or central queue only systems. However, this is not always true for subsets of participants. For example, when the smallest 30 participants are considered, the average marginal liquidity saving in the RITS replica system is lower than that in the pure RTGS system (a result also borne out by the difference in the coefficients found in the regression analysis), but *higher* than that in the central queue only system.

6. Decomposing Liquidity Savings

To decompose liquidity savings into the two sources identified in the literature, namely liquidity pooling and payments internalisation, we follow Lasaosa and Tudela (2008) and run two additional sets of simulations. To isolate the impact of liquidity pooling, we run the tiered simulations including the internalised payments that were previously omitted.[13] Since these internalised payments are still being sent through the system, all the liquidity savings from tiering can be attributed to liquidity pooling. Conversely, to measure internalisation savings we omit payments between the client and the settlement bank but otherwise leave the client as a direct participant.[14] Therefore, any reduction in liquidity usage will be due to transactions between the client and the settlement bank being settled outside the RTGS system. For this exercise, we focus on the cumulative tiering scenarios alone.

Since liquidity pooling and internalisation are the only sources of liquidity savings, comparing liquidity savings in the original simulations with those in the additional simulations result in two sets of estimates for the relative importance of the sources of the liquidity savings.[15] The reason for the differences in these estimates is twofold. Firstly, the complexities of the liquidity recycling process mean that a small change in transaction data can have substantial flow-on effects on the settlement and liquidity usage profiles of a simulation. Secondly, our additional simulations do not perfectly separate out the liquidity-saving effects of tiering. Because the client still participates in the system in the internalisation simulations, the fact that it no longer receives funds through the system from the settlement bank potentially increases its liquidity needs. Note that this upward pressure on liquidity does not exist in the original tiering simulations because the client is completely removed from the system. As such, the liquidity savings yielded by the internalisation simulations are likely to be understated.

Figure 4 shows that not only have the internalisation simulations understated the liquidity savings from internalisation, they have predicted negative liquidity

---

[13] This involves transforming payments to and from the client into payments to and from the settlement bank, and not dropping payments from the settlement bank to itself.

[14] As multiple clients enter the same tiering network, all payments between them must also be tiered. For example, consider initially that participant A has tiered to participant B (i.e. participant B acts as settlement bank for participant A). To measure the internalisation effect when participant C also tiers to participant B, payments between participants C, B *and* A must all be omitted.

[15] These values should be thought of as alternative estimates, not as the upper and lower bound on a range.

savings (i.e. an increase in liquidity usage) for almost all tiering scenarios.[16] Consequently, indirect estimations of liquidity pooling's share of total liquidity savings are well over 100 per cent in most scenarios. Estimated directly, liquidity pooling typically accounts for around 95 to 100 per cent of total liquidity savings.

Figure 4

Liquidity Usage Decomposition

RITS replica

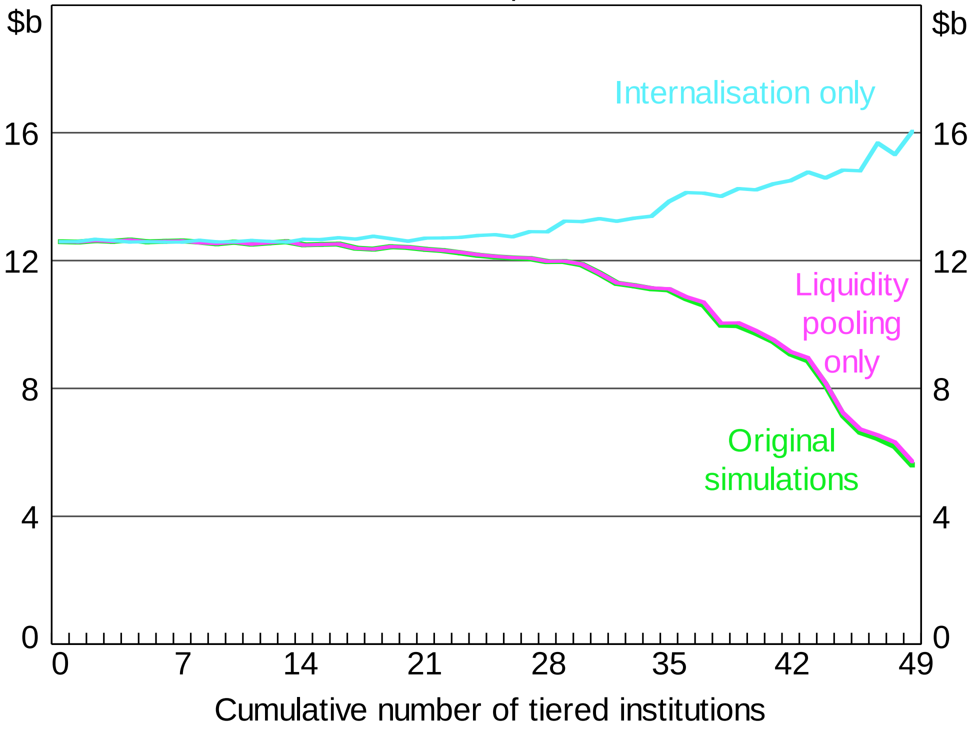

Source: RBA

Information asymmetries suggest that internalisation benefits are likely to be the primary driver in deciding to tier. In most payments networks each participant can only observe its own payments and receipts. Consequently, both it and a prospective settlement bank can estimate the liquidity savings from internalisation based on the offsetting flows between the two institutions. However, in order to estimate the liquidity pooling effect a participant would need to obtain a degree of information which commercial considerations may render infeasible.

---

16 While the results shown here are for the RITS replica system, liquidity savings in the other three system designs exhibit a similar pattern.

Our finding that the majority of liquidity savings are due to liquidity pooling is in line with that of Lasaosa and Tudela (2008). This poses an interesting question, given Lasaosa and Tudela's finding is for participants in the highly-tiered CHAPS Sterling system: if the direct participants considered for tiering by Lasaosa and Tudela are representative of all participants (both direct and indirect) in CHAPS Sterling, then this may suggest that information asymmetries do not hinder a participant's estimation of the benefits of tiering and thus its decision to tier (because a large number of participants have chosen to tier in that system despite not being able to estimate the liquidity pooling savings *a priori*). On the other hand, both our and Lasaosa and Tudela's results may be evidence of information asymmetries leading to an inefficiently low-level of tiering, because participants which have chosen not to tier are the ones that stand to gain predominantly from the unobservable liquidity pooling effect.

7. The Impact of Tiering on Systemic Risk

7.1 Concentration Risk

Indirect participants in a payments network send payment instructions to their settlement bank, which then acts on their behalf. Consequently, in choosing to tier the client becomes operationally dependent on its settlement bank. One might argue that larger institutions are better equipped to minimise the probability of an operational problem. However, by concentrating payment flows, tiering amplifies the consequences of an operational incident at the settlement bank – in particular, the size of the potential liquidity sink increases.

A general measure of this type of operational risk is the level of concentration in the system; the increase in settlement banks' share of payments as the level of tiering increases. Note that we have used the share of payments sent, rather than sent and received, as generally even when a participant suffers an operational incident they can still receive payments. While a more accurate way to model the impact of tiering on the consequences of an operational incident is to simulate operational incidents in a tiered network, this is beyond the scope of this paper.

We find that our cumulative tiering scenarios result in only a modest increase in the concentration of payments among the four largest participants. In the absence of tiering, the four largest participants account for around 57 per cent of RITS payments by value. With all other participants (apart from the RBA, CLS Bank, and the settlement accounts of the equity and futures systems) settling indirectly the combined share of the four largest participants rises by around 10 percentage points. However, since it is unlikely that an operational incident would occur at all four of the largest participants simultaneously, it is more noteworthy that the largest single share only rises 3 percentage points.

As a settlement bank's network of clients increases, more payments are settled across its books rather than settled in RITS. Figure 5 shows the increase in payments settled across the books of a settlement bank (rather than in RITS) as the level of tiering increases. While payments settled outside RITS are likely to be unaffected by an operational problem in RITS itself, this immunity is likely to be of minimal benefit due to the low likelihood of such a problem.

Figure 5

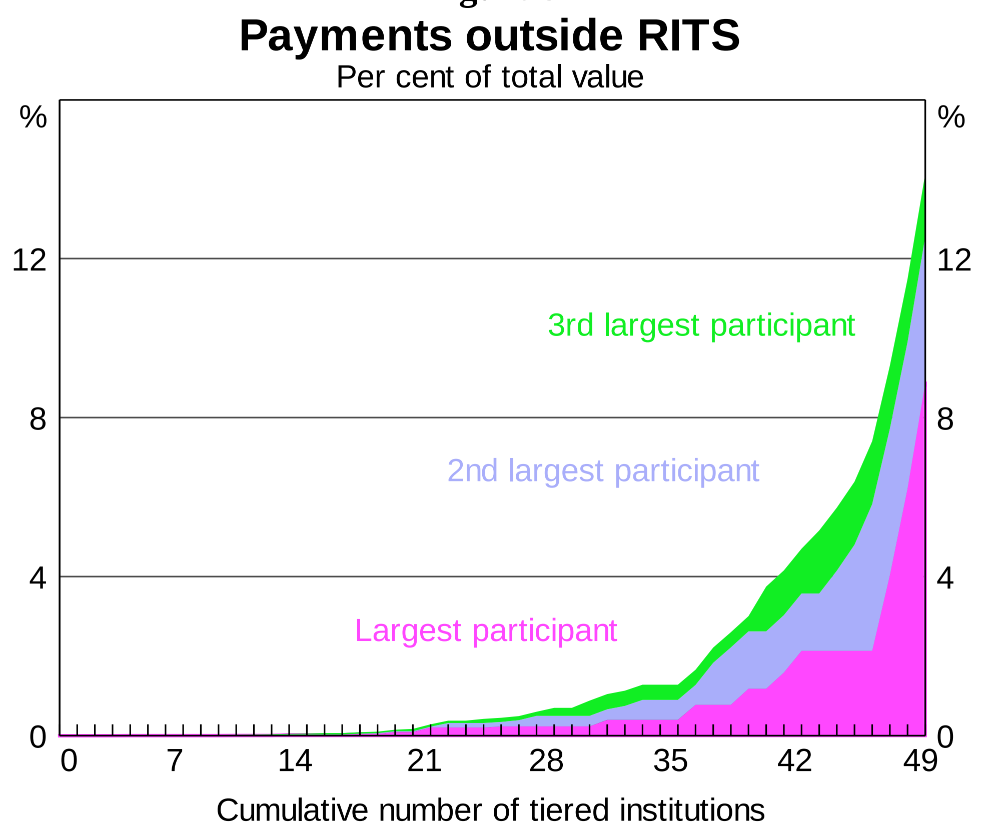

Source: RBA

7.2 Credit Risk

Tiering creates a two-way exposure between a client and its settlement bank because payments are settled across the settlement bank's books, rather than in central bank money, which is not subject to credit risk. Furthermore, these payments – unlike those in RITS – may be subject to the zero hour rule, which means that in the event of a bankruptcy their finality can be challenged. In this section we present measures of this two-way exposure for the two system designs at either end of the liquidity usage spectrum: the pure RTGS system and the RITS replica system.

7.2.1 Settlement bank exposures

A settlement bank's maximum intraday exposure to a client can be measured as the client's maximum intraday cumulative net *payment* (as opposed to receipt) position when it settles directly in the RTGS system. This measure of settlement bank exposure should be regarded as an upper bound, as settlement banks can use their discretion regarding the timing of their clients' payments to minimise their exposure and they can require clients to pre-fund settlement obligations.[17]

Nevertheless, according to this measure, a settlement bank's *average* maximum intraday exposure to any one of the smallest 29 tiering candidates over the sample period is less than $100 million – a significantly smaller amount than the tier 1 capital of the four largest settlement banks in our simulations (Figure 6). While the *largest* maximum intraday exposures over the month are roughly three times the size of the average maximum intraday exposures, they are still quite low for the smallest 29 tiering candidates (Figure 7). Unsurprisingly, maximum intraday exposures are typically much higher among the largest 20 tiering candidates. A loss equal to the largest maximum intraday exposure of around $2 billion is still, however, somewhat smaller than the tier 1 capital of the four largest settlement banks.

---

[17] Note, the timing of settlement in the tiered simulations may also vary depending on the liquidity available to the settlement agent.

Figure 6

Settlement Banks' Maximum Exposures

Average over period

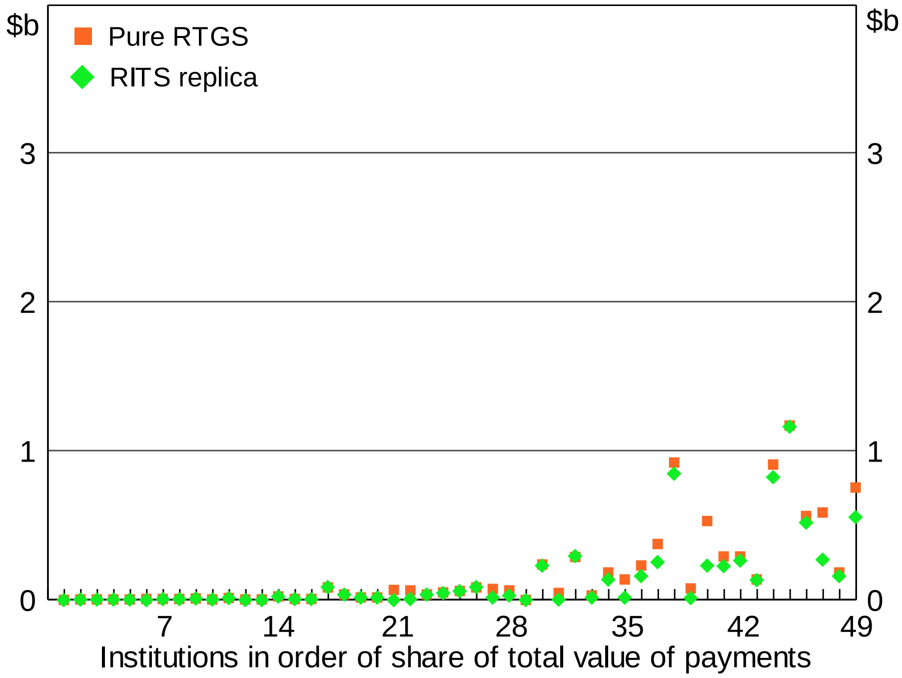

Source: RBA

Figure 7

Settlement Banks' Maximum Exposures

Maximum over period

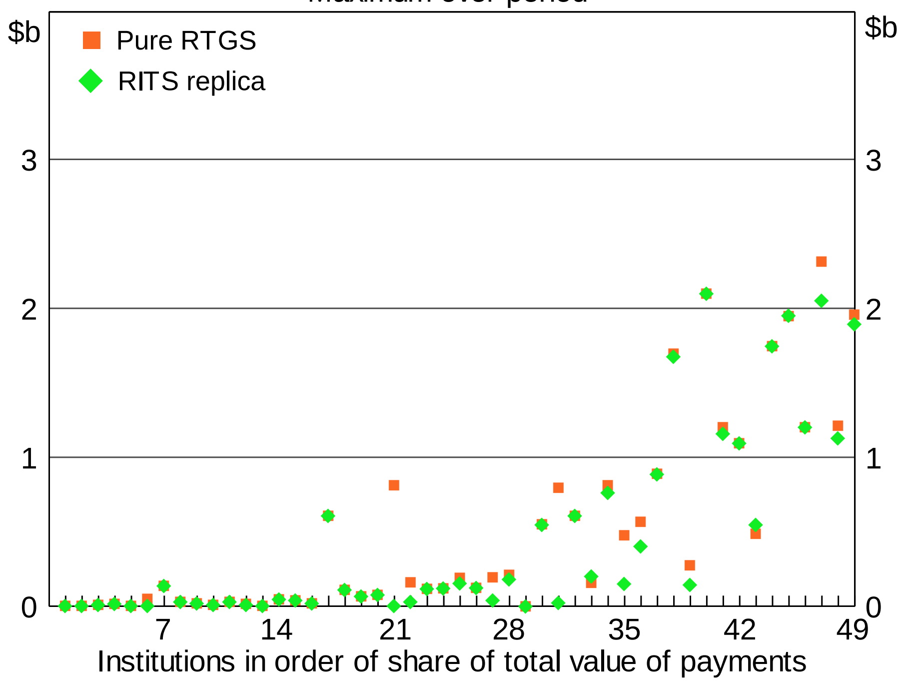

Source: RBA

Overall, system design does not appear to significantly affect a settlement bank's maximum intraday exposure. However, in the few cases where exposure is noticeably higher in one system design, this tends to be the pure RTGS system.

Our results provide more support than those of Lasaosa and Tudela (2008) for a rule that seeks to reduce credit risk arising from indirect settlement by restricting tiering based on the share of value settled. Lasaosa and Tudela found a correlation of just 0.52 between settlement banks' average maximum intraday exposures and average daily value settled; we find higher correlations of 0.60 in the RITS replica system and 0.68 in the pure RTGS system. Moreover, Figure 7 suggests a polynomial rather than linear expression might better explain the relationship between the two series.

7.2.2 Individual client exposures

A client's maximum intraday exposure to its settlement bank can be measured using the client's maximum intraday cumulative net *receipt* (as opposed to payment) position when it settles directly in the RTGS system. Because a settlement bank has discretion over the timing of payments, and because it may require pre-funding from its client, these estimates should be viewed as a lower bound.

Clients' average maximum intraday exposures are typically less than $1 billion (Figure 8). The largest maximum intraday exposures are still less than $1 billion for smaller institutions, but are as high as $3.5 billion for the largest institutions (Figure 9). Given the largest client institutions are typically global banks, their largest exposures are still somewhat smaller than their tier 1 capital.

Figure 8

Clients' Maximum Exposures

Average over period

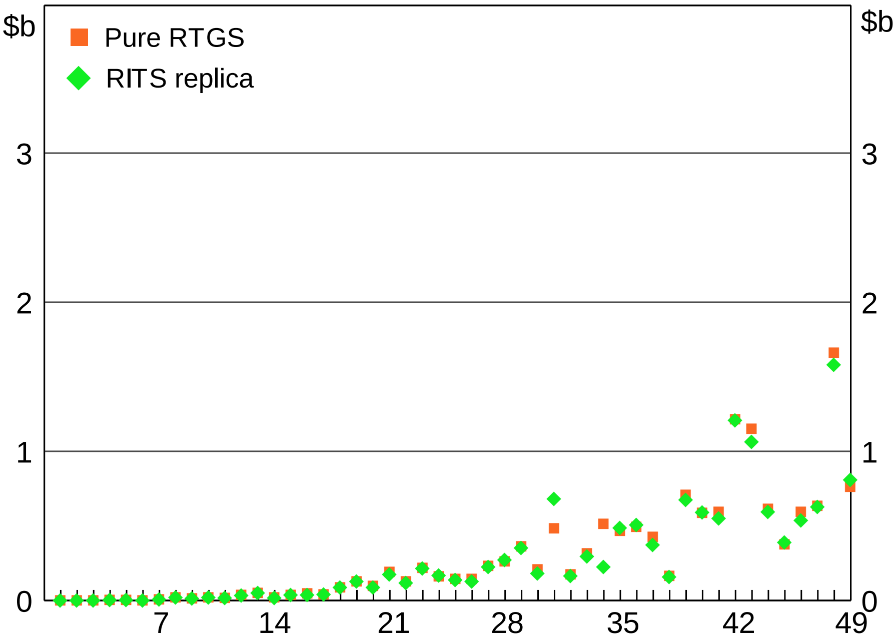

Source: RBA

Figure 9

Clients' Maximum Exposures

Maximum over period

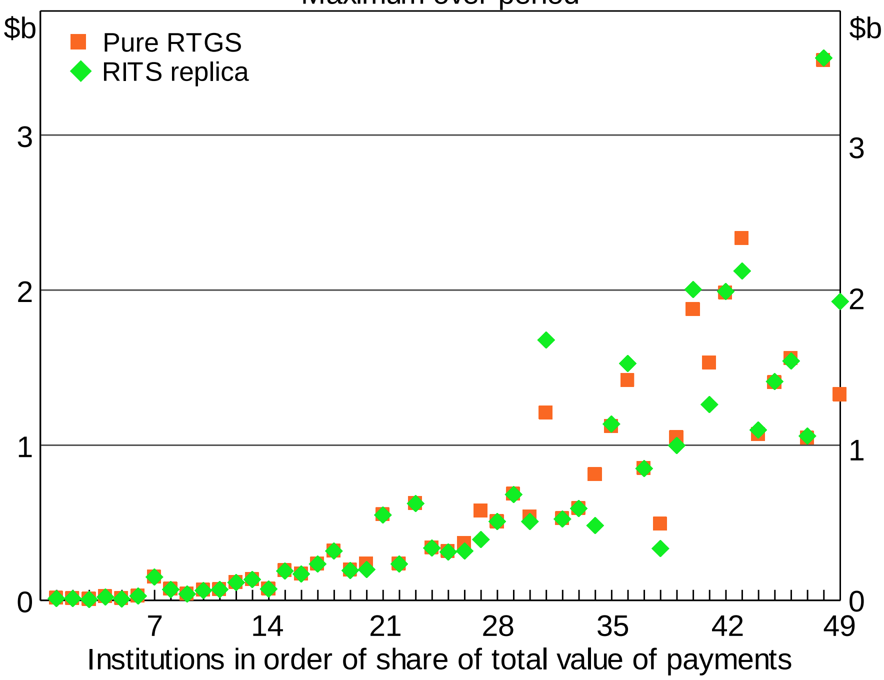

Source: RBA

Clients' largest maximum intraday exposures vary significantly across different system designs in a few cases, but do not tend to be higher for any particular system design. Unlike settlement bank exposure, where examples of higher exposures in the pure RTGS system are a likely consequence of the higher liquidity usage in that system (i.e. clients tend to have higher cumulative net *payment* positions), the effect of system design on clients' cumulative net *receipt* positions is less clear.

7.2.3 Total client exposures

While a settlement bank is unlikely to face the simultaneous default of all of its clients, if a settlement bank defaults all of its clients are exposed. To estimate the maximum total client exposure to a particular settlement bank we can sum the minute-by-minute exposures, measured using each client's cumulative net receipt position when it settled directly.[18] As noted above, these estimates of client exposures should be viewed as lower bounds.

Each observation in Figures 10 and 11 represents the maximum aggregate loss that could occur in the most recently expanded tiered network due to default by the settlement bank for that network. For example, when the 49th institution tiers in Figure 10, the average maximum intraday exposure that it and other clients tiered to the same settlement bank face in aggregate is around $4 billion in the RITS replica system.

18 Note that exposures are not multilaterally netted. Therefore, if a client has negative exposure (i.e. it owes the settlement bank), that exposure is excluded from the calculation.

Figure 10

Total Client Maximum Exposures*

Average over period

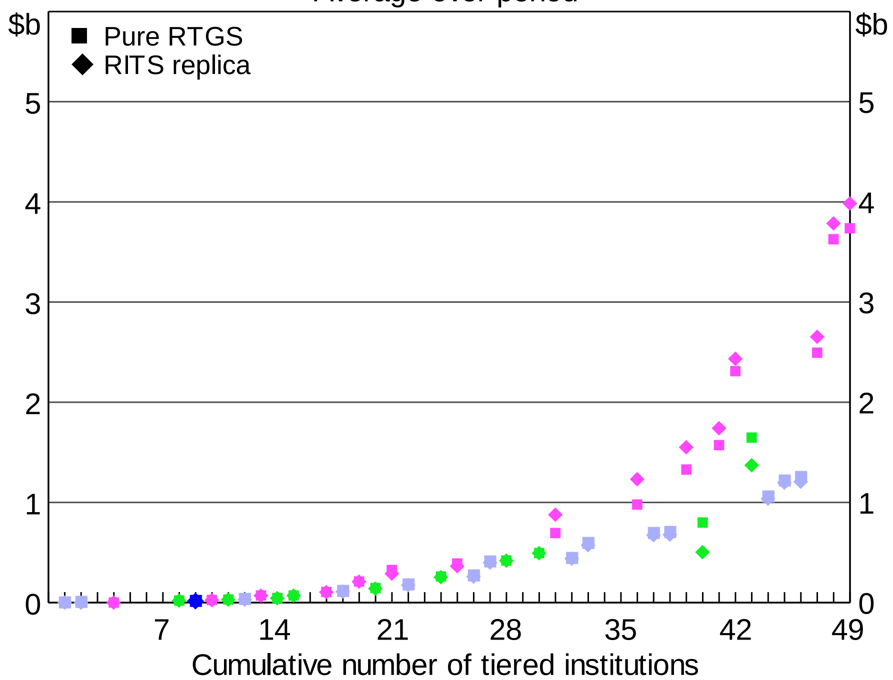

* Each colour represents one of the four largest settlement banks

Source: RBA

Figure 11

Total Client Maximum Exposures*

Maximum over period

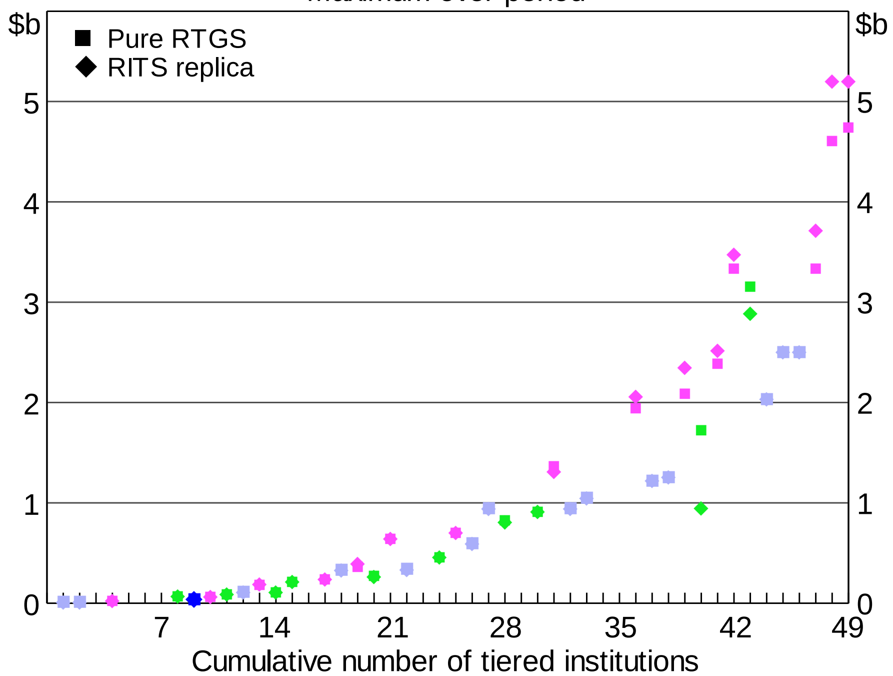

* Each colour represents one of the four largest settlement banks

Source: RBA

The effect of system design on these exposures tends to be small and inconsistent. As larger institutions tier, total client exposures are higher in the RITS replica system for the settlement bank depicted in pink, but lower for the settlement bank depicted in green. System design does not significantly affect total client exposures when only small institutions are tiered.

8. Evaluating Tiering Trade-offs

8.1 Central Bank Decision-making

In this section we provide an attempt to more firmly quantify the value of liquidity savings (accruing to the private sector) and risk associated with exposed payments (a concern both for the private sector and the central bank, insofar as the central bank's responsibilities include maintaining the stability of the financial system). The observed linear relationship between these two variables suggests the construction of a utility maximisation problem, analogous to the standard problem faced by a consumer choosing over two goods with a budget constraint. This leads to an exploration of a framework within which a central bank might choose the socially optimal level of tiering, though without identifying certain key parameters it is not possible to provide a specific solution to this issue. Throughout, our focus is on the cumulative tiering scenarios, that is, the situation where all system participants below a certain size choose to tier; this is reflective of the current formulation of RBA policy, where only participants below a particular payment-share cut-off are eligible for tiering (though, as noted, in practice it is not the case that all agents eligible to tier will choose to do so).

8.2 Valuing Liquidity Savings

Previous literature has established that it is possible to provide a quantitative estimate of the opportunity cost of liquidity usage (see James and Willison 2004). The benefit of reducing the system-wide need for liquidity depends on the marginal cost of liquidity. Since participants only pay small, flat transaction fees on intra-day repos, and as intra-day liquidity in RITS is free besides collateralisation, liquidity cost is largely a function of the opportunity cost of the collateral required to secure the intraday repos. In theory, the opportunity cost of collateral is the value of the collateral used multiplied by the spread between the unsecured and secured lending rates in the cash market. Formally, denoting the spread as i, collateral used as Y, and the opportunity cost of collateral as Q, we have

$$Q = iY$$

The intuition behind this calculation is that an institution in possession of collateral-eligible securities could use those securities to obtain funds in the secured lending market, and then lend those funds out at the unsecured lending rate. In holding (and using) the collateral to cover liquidity demands, the institution foregoes this marginal return.[19]

In practice, over the twelve months prior to the market turmoil at the end of August 2008, the average value for i was around 0.4 basis points. Based on this rate, the value of liquidity savings is less than $10,000 in all systems up until the 18th institution is tiered (Figure 12). From the tiering of the 38th institution, the value of savings is over $100,000. Corresponding to the liquidity usage in the various systems, the average value of liquidity savings is lowest in the RITS replica and bilateral offset only systems at around $61,000, and highest in the pure RTGS system at over $70,000. Average value of savings in the central queue only system is around $64,000.

---

[19] This assumes that the relevant binding constraint on collateral holdings is the necessity of using that collateral to obtain liquidity. However, financial institutions are required to hold certain levels of low-risk securities as part of their prudential regulatory requirements, and some may prefer to hold low-risk securities as part of their portfolio strategy, both factors that may provide a lower bound on the amount of collateral holdings that could be sold off. This may prevent full realisation of the gains discussed in this section.

Figure 12

Opportunity Cost of System Liquidity Savings

Daily average

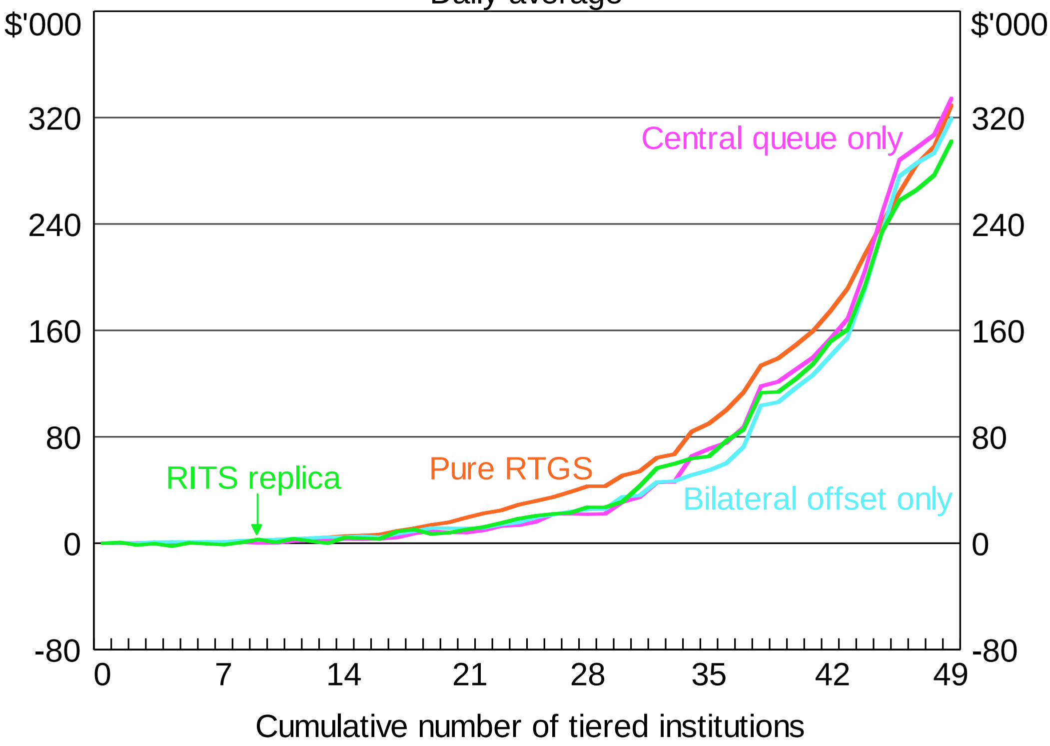

Source: RBA

A caveat here is that on several occasions over the last few years – and persistently since August 2008 – the unsecured-secured spread has been negative, suggesting our measurement of it is problematic. There may be compositional issues, as the unsecured rate is measured across inter-bank loans, while the secured rate includes loans to non-bank institutions (which are likely to have riskier credit profiles). Counterparty credit limits can also lead to negative spreads as a participant that has exhausted its unsecured limits will likely prefer to pay a premium to borrow on a secured basis rather than pay the 25 basis point penalty on overnight loans from the RBA. While this is still a valid measure of the return to the lender, in the extreme if such limits prevent on-lending of unused liquidity, the unsecured-secured spread overestimates the opportunity cost of liquidity as no revenue has been foregone. Recent Australian government guarantees on bank deposits and bond issuances have also blurred the lines between secured and unsecured lending. In any case, while the period since August 2008 has been excluded from our averaging calculation, these difficulties in measuring the spread will induce error in our estimates. (For comparison, James and Willison find a 7 basis point unsecured-secured spread in the UK money market.)

8.3 Approximating Changes in System-wide Credit Risk

Providing a monetary estimate of the 'cost' associated with having a given value of payments exposed is a much more difficult task; here we instead attempt to provide an estimate of the expected value of the size of defaulted payments that might occur as a result of participant failure. While in Section 7 above, we considered the extreme daily exposures of settlement banks to clients and vice versa, here we will instead use the average account balance of each client throughout the sample period.

In the case where a client has a positive account balance on average then in the event that they failed at a random point throughout the day it would be expected that they would be owed money by the settlement bank, and thus the failure would not be associated with payment default. If the settlement bank fails, however, then we assume they would default on the full value of the average positive balance. Conversely, in the case of a client with a negative average balance, the failure of the client would be expected to result in the default of payments by the client to the settlement bank, rather than vice versa.

The sum of the absolute value of institutions' average balances is shown in Figure 13 (absolute values are used since, as noted above, exposures arise regardless of whether balances are positive or negative). Using average balances as a measure of exposure is a significant simplification and at least three provisos should be made clear. First, that in the event of a default, we are likely to observe abnormal payment patterns immediately prior to the failure, and thus using an average is misleading. Secondly, account balances throughout the day reflect strategic behaviour by participants. In RITS, smaller participants often obtain liquidity by delaying outgoing payments until late in the day, funding them from earlier incoming payments (while larger participants are more likely to obtain liquidity through repos). Thus smaller participants may be more likely to have positive average balances, but there would be no reason to expect this behaviour to be sustained if the participant is no longer directly responsible for obtaining liquidity. Third, and perhaps most importantly, our assumption ignores the almost-certain reality that settlement banks would establish limits on the size of the credit exposures they are willing to take on with regards to their clients.

Figure 13

Total Average Value of Exposed Payments

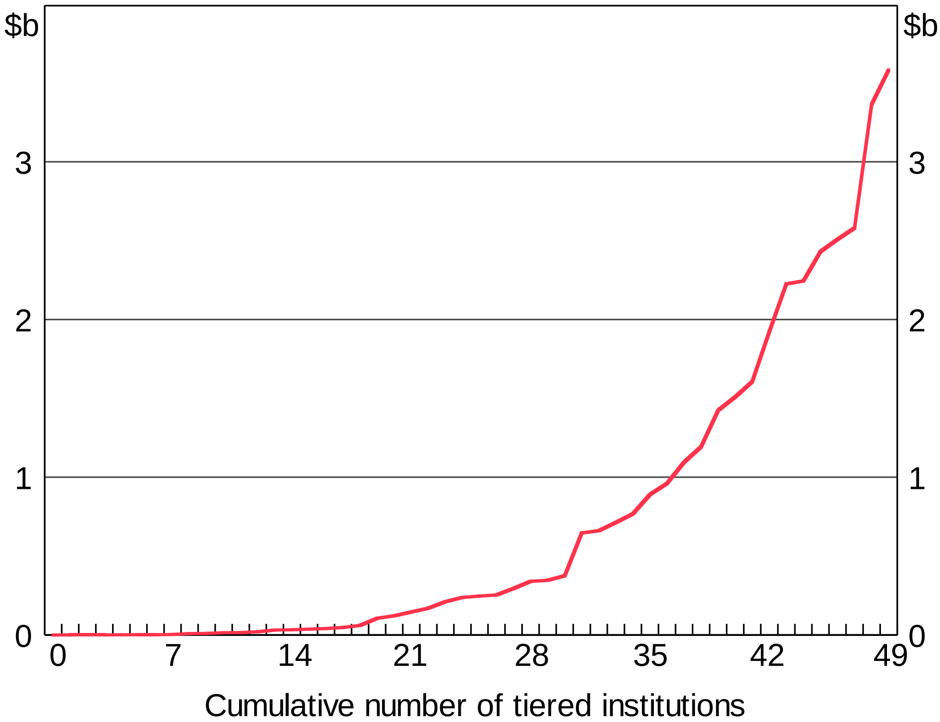

Source: RBA

Ideally, the next steps would involve estimation of the probability of default for each individual participant, and then further estimation of the likelihood of 'second-round' effects from such a default (that is, creditors to the original defaulter defaulting in turn), perhaps by comparing the size of expected defaults to the capitalisation of each creditor. Again, in another simplification, we will instead assume that the probability of any individual participant failing on any given day is k, and that the probability of any participant failing in the first round is independent of the probability of any other participant failing in the first round. This is not saying that we ignore the possibility of second-round failures; rather, we are ignoring the prospect that, if we observe one participant failing due to an exogenous event – such as a sudden decline in the value of a particular asset class – we would likely expect the probability of some other participant failing at the same time to be higher, regardless of the impact of the first participant failing to pay its debts.

Based on these assumptions, a simplistic measure of credit risk 'cost', to be denoted P, is then simply k multiplied by the size of total average exposures (denoted X) under increasing degrees of tiering. Formally,

$P = kX$

We can provide an estimate of the value of k using the credit ratings of system participants. For participants for whom ratings are readily available (approximately 80 per cent of the total), the median Standard & Poor's credit rating is A+. This is associated with a mean annual default rate since 1981 of 0.05 per cent, or around 0.0001 per cent as a daily rate.[20]

We make no attempt to model feasible second-round scenarios, but instead, in the next section, will assume that the central bank's utility is decreasing as P increases. Note, however, that this is quite problematic: it suggests the central bank would be indifferent between two separate participants each having average exposures of size z, and one individual participant having an exposure of size 2z. In practice, central bank concern over credit risk is likely to be not only a function of the total value of exposed payments, but of the distribution of those exposures amongst participants.

8.4 Central Bank Utility

As can be seen in Figure 14, the relationship between our two variables of interest – the expected average daily value of payments defaulted upon (kX or CREDLOSS); and the estimate of the opportunity cost of collateral (iY or LIQCOST) – is strikingly linear.

---

[20] Standard & Poor's (2009) itself advises against using credit ratings as a guide to default rates, as ratings may be influenced by other factors (such as the probability of a bail-out post-default).

Figure 14

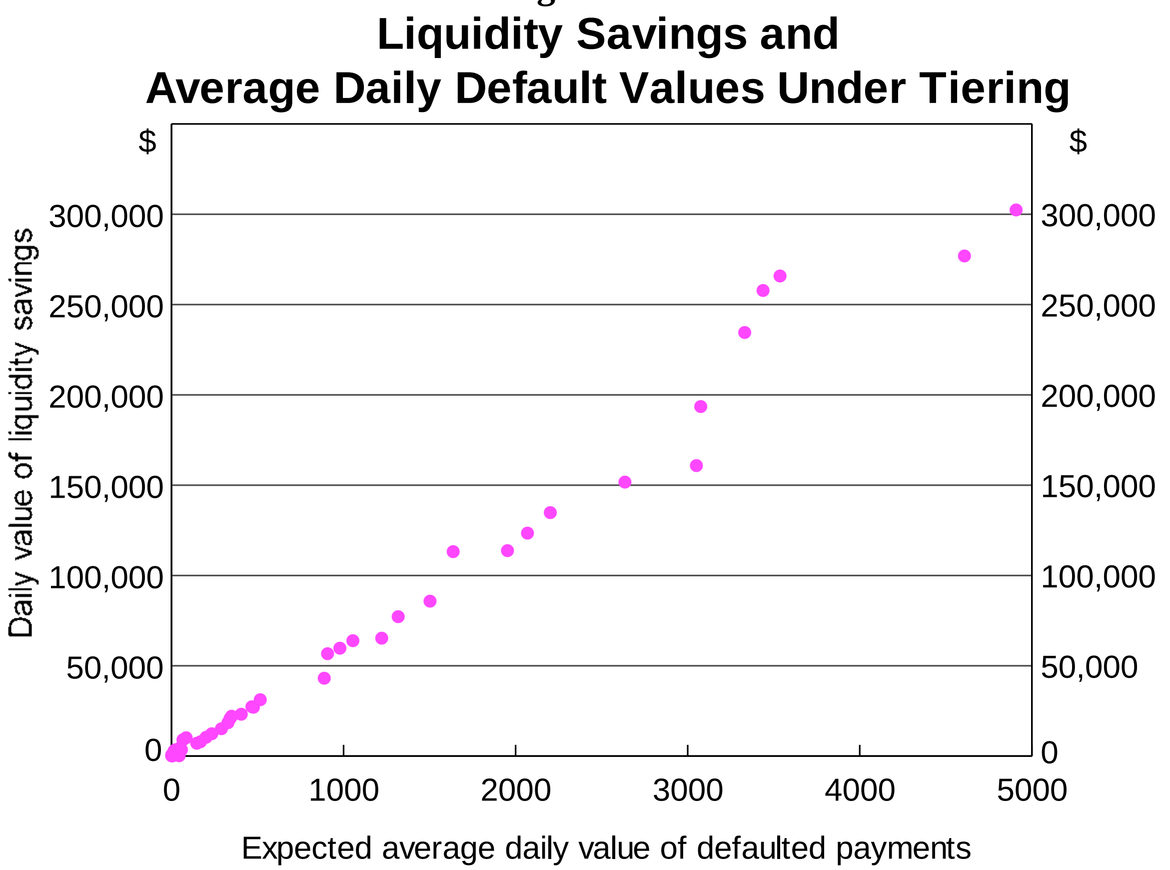

It is not mathematically necessary – nor expected – that a relationship of this type should hold, and thus this is in itself of some interest. A regression of the value of liquidity savings on the expected credit loss in levels has an adjusted R^2 value of 0.98, with the following coefficients significant at the 1 per cent level:

$$LIQCOST = -1464.55 + 64.01 \cdot CREDLOSS$$

As the regression and Figure 14 indicate, the expected value of liquidity savings is generally substantially larger than the average daily value of defaulted payments (indeed, a one dollar increase in the average daily value of defaulted payments is associated with sixty-four dollars worth of liquidity savings). Aside from the caveats already mentioned concerning how these variables have been constructed, this does not necessarily indicate that tiering is desirable from the central bank's perspective, as concern over payments defaulted upon is likely to be a function of the systemic impact of a default, not the size of the default itself.

If the expected daily value of defaulted payments is converted from a ‘bad’ to a ‘good’ – by, for instance, subtracting each measure from the maximum expected credit loss (i.e. kX when all 49 tiering candidates settle indirectly) – then the trade-off faced is one of maximising against, effectively, a linear budget constraint, analogous to standard consumer utility problems. To illustrate this, we could define central bank utility using the (relatively general) Constant Elasticity of Substitution function:

$$U = \left(aP^{r} + (1-a)Q^{r}\right)^{\frac{1}{r}}$$

where

$$r = \frac{s-1}{s},\ 0 \le a \le 1$$

and s can be interpreted as the elasticity of substitution. As s approaches 1 (r approaches zero), the function collapses to the Cobb-Douglas form; as s approaches infinity (r approaches 1), we have perfect substitutes; and as s approaches zero from above (r approaches negative infinity) we have perfect complements.

Naturally, in practice a central bank might also place weight on other issues rather than credit risk and liquidity savings. For example, as noted above, tiering impacts the concentration of payments system flows, and thus can potentially alter the level of operational risk in the system. However, credit risk and liquidity savings are likely to be the two areas of largest concern. Substituting in $P = kX$ and $Q = iY$, the problem is thus to maximise

$$U = \left(a(kX)^{r} + (1-a)(iY)^{r}\right)^{\frac{1}{r}}$$

subject to the linearity constraint

$$iY = \alpha - \beta kX$$

Standard solution techniques generate the following optimal value for Y as a function of our various parameters:

$$Y^* = \frac{\alpha}{i + \beta k \sigma}$$

where $\sigma = \left(\left(\frac{1-a}{a} \right) \beta \right)^{\frac{1}{r-1}} \left(\frac{i}{k} \right)$

If central bank's preferences can be specified in terms of a and r, it would then be straightforward to find the relevant level of tiering that generates the closest possible approximation to Y^*. For interest, in the special case of Cobb-Douglas utility (that is, assuming an elasticity of substitution of one), where the problem is instead to maximise

$$U = a \ln(kX) + (1-a)\ln(iY)$$

again subject to

$$iY = \alpha - \beta kX$$

which solves for the optimal value of Y:

$$Y^* = \frac{(1-a)\alpha}{i}$$

Having converted the expected value of defaulted payments into a 'good', we estimate α:

$$LIQCOST = 312627.00 - 64.01 \cdot CREDLOSS$$

Thus, under these highly restrictive assumptions, to support the current level of tiering in RITS where Y is effectively zero (that is the value of Y^* is closer to zero than the minimum possible liquidity savings from tiering) would require a value for a practically equivalent to one, which implies that the central bank places considerable weight on the credit risk that results from tiering.

9. Conclusions

This paper uses simulation analysis to examine the effect of system design on the incentives to participate indirectly in RTGS systems. Our methodology is adapted from Lasaosa and Tudela (2008). We are motivated primarily by explaining the low level of tiering in Australia's RTGS system, RITS, relative to RTGS systems overseas.

We find some evidence to support our *a priori* hypothesis that the design of RITS (i.e. an RTGS system with a central queue, a bilateral offset algorithm and a liquidity reservation feature) reduces the liquidity-saving incentives to tier. Compared to a pure RTGS system, for instance, regression analysis suggests the marginal liquidity savings from tiering are lower amongst the smaller system participants. However, we find that tiering under different system designs has mixed effects on liquidity usage and we attribute this to the lack of behavioural response from participants to system design changes in our simulations; specifically, the fact that participants' credit limits and use of sub-limits are unchanged across system designs. Further simulation analysis could arguably warrant the embedding of certain behavioural assumptions in the data inputted to the simulator. However, such assumptions are likely to be somewhat arbitrary and also difficult to implement given the size of the datasets that will require modification.

Similar to Lasaosa and Tudela (2008), we find that the majority of liquidity savings come from liquidity pooling rather than the internalisation of payments. In conducting this analysis, we follow Lasaosa and Tudela and run two additional sets of simulations designed to separate out the two effects. However, we note that the set of simulations designed to isolate the internalisation effect do not do so perfectly, and thus these results should be treated with caution. Conceivably, our finding that the majority of liquidity savings are sourced from the liquidity pooling effect could complicate the process of establishing tiered networks since, for reasons of confidentiality, prospective clients and settlement banks are unlikely to be able to share the payments information necessary to properly estimate the liquidity savings from liquidity pooling. However, it is unclear whether our finding helps explain the low level of tiering in RITS given Lasaosa and Tudela reach a similar conclusion for the highly-tiered CHAPS Sterling system.

We find that tiering results in only small increases to the already high level of concentration in RITS, and thus does little to increase concentration risk (as concentration risk is measured as the percentage of the value of payments sent by direct participants in the system, this finding is independent of system design). In terms of credit risk, we find that the effects of system design are limited. We note that maximum exposures are typically smaller than the participants' capitalisations. However, we do not know the level of exposure that any individual participant would be willing to accept as part of conducting its payments business, and we would expect this to be significantly lower than the participant's capital.

Finally, we offer an attempt to provide more intuitive interpretations of the main costs and benefits of tiering, leading to the construction of a simplistic framework within which a central bank might attempt to direct a system towards an optimal tiering level (though the difficult task of determining the appropriate parameters to govern such a framework is not attempted).

In sum, we do not establish any conclusive reason for the relatively low level of tiering in RITS and the seeming preference for many smaller participants to remain directly engaged in the system.

Appendix A: Sub-limits and Bilateral Offsetting

We have used bilateral limits in the simulator to replicate RITS' sub-limit feature. This has involved modifying the simulator's entry and queuing sub-algorithms so that they conduct the appropriate settlement tests (e.g. test priority payments against a participant's entire settlement account balance, and test active payments against a participant's account balance in excess of its sub-limit). In addition, data limitations mean we are unable to pinpoint when a queued payment's status is changed by the sending participant; we only know the status of the payment upon submission to the RITS queue, and the status of the payment when it was settled in RITS. Input to the simulator requires payments to have a single status which remains unchanged during queuing, thus we were required to modify our underlying transaction data. Table A1 summarises our approach.

Table A1: Payment Status And Submission Times

Status when submitted to RITS	Status when settled in RITS	Status when submitted to the simulator	Time when submitted to the simulator
D	A	A	Settlement time in RITS
	P	P	Settlement time in RITS
A	A	A	Submission time to RITS
	P	P	Submission time to RITS
P	A	P	Submission time to RITS
	P	P	Submission time to RITS

Note: In the pure RTGS system design with unlimited liquidity all payments are submitted to the simulator at the time they were settled in RITS and payment statuses are irrelevant.

Payments that were submitted to RITS as deferred are submitted to the simulator at the time that they were settled in RITS. This change is based on the assumption that the sender of a deferred payment did not intend for the payment to settle upon its submission, but rather intended to change the status of the payment at a later point in time at which it did want settlement to occur – we assume that the actual settlement time in RITS is a better approximation of this point in time. Payments that were submitted as active but later settled as priority also have their submission time to the simulator changed to their actual settlement time in RITS. A number of participants in RITS have been observed to manage liquidity by setting very high sub-limits, submitting payments to the queue as active, and subsequently changing a payment's status to priority when they want it to be settled. As such, we again assume in these cases that actual settlement time in RITS is a better approximation of the point in time at which the sending participant wished settlement to occur.

We have also designed a bilateral offset sub-algorithm for the simulator which seeks to replicate RITS' own bilateral offset algorithm. In RITS, payments which are queued for over a minute are tested for bilateral offset with up to 10 payments due from the receiving participant on a next-down looping basis.[21] By contrast, the BOBASIC bilateral offset sub-algorithm which is provided with the simulator tries to offset all queued payments between the sending and receiving participants, iteratively removing the last queued transaction to find a combination of offsetting transactions that it can settle simultaneously.

21 We have not incorporated the minute delay feature of RITS' bilateral-offset feature into our sub-algorithm, and this is not expected to affect our results significantly.

References

Adams M, M Galbiati and S Giansante (2008), 'Emergence of tiering in large value payment systems', mimeo.

Australian Prudential Regulation Authority and Reserve Bank of Australia (2003), 'Management of Exchange Settlement Accounts', Joint Media Release, March.

Bank for International Settlements (2009), *Statistics on payment and settlement systems in selected countries*, CPSS Publications No 86.

Chande N, A Lai and S O'Connor (2006), 'Credit in a Tiered Payments System', Bank of Canada Working Paper No. 2006-36.

Chapman J, J Chiu and M Molico (2008), 'A Model of Tiered Settlement Networks', Bank of Canada Working Paper No. 2008-12.

Ercevik K and J Jackson (2007), 'Simulating the impact of a hybrid design on the efficiency of large-value payment systems', in H Leinonen (ed), *Simulation analyses and stress testing of payment networks*, proceedings from the Bank of Finland Payment and Settlement System Seminars 2007–2008, Helsinki.

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

Jackson J and M Manning (2007), 'Central Bank intraday collateral policy and implications for tiering in RTGS payment systems', De Nederlandsche Bank Working Paper No. 129.

James K and M Willison (2005), 'Collateral Posting Decisions in CHAPS Sterling', *Financial Stability Review*, December 2004.

Kahn C and W Roberds (1998), 'Payment System Settlement and Bank Incentives', *The Review of Financial Studies*, 11(4), pp 845–870.

Kahn C and W Roberds (2008), 'Payments Settlement: Tiering in Private and Public Systems', *Journal of Money, Credit and Banking*, 41(5), pp 855–884.

Lasaosa A and M Tudela (2008), 'Risk and efficiency gains of a tiered structure in large-value payments: a simulation approach', Bank of England Working Paper No. 337.

Reserve Bank of Australia (1999), 'Eligibility for Exchange Settlement Accounts', Media Release No. 1999-2.

Standard and Poor's (2009), 'Understanding Standard and Poor's Rating Definitions', Standard and Poor's website, June 2009.