BANK OF ENGLAND

Working Paper No. 427

System-wide liquidity risk in the United Kingdom's large-value payment system: an empirical analysis

Marcelo Perlin and Jochen Schanz

May 2011

BANK OF ENGLAND

Working Paper No. 427

System-wide liquidity risk in the United Kingdom's large-value payment system: an empirical analysis

Marcelo Perlin[(1)] and Jochen Schanz[(2)]

Abstract

When settling their own liabilities and those of their clients, settlement banks rely on incoming payments to fund a part of their outgoing payments. We investigate their behaviour in CHAPS, the United Kingdom's large-value payment system. Our estimates suggest that in normal times, banks increase their payment outflows when their liquidity is above target and immediately following the receipt of payments. We use these estimates to determine the robustness of this payment system to two hypothetical behavioural changes. In the first, a single bank stops sending payments, perhaps because of an operational problem. In the second, it pays out exactly what it previously received, relying exclusively on the liquidity provided by other system members. Using the observed uncertainty around our estimated behavioural equations, we derive probabilistic statements about the time at which the bank's counterparties would run out of liquidity if they followed their estimated normal-time behaviour.

Key words: Payment systems, banks, network models, contagion, systemic risk, liquidity risk.

JEL classification: G21.

(1) Professor Adjunto — Escola de Administração (Federal University of Rio Grande do Sul, Brazil).
Email: msperlin@ea.ufrgs.br
(2) Bank of England. Email: jochen.schanz@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. This research was conducted when Dr Perlin was a summer intern at the Bank of England.

We are grateful to seminar participants at the Bank of England and an anonymous referee for useful comments. This paper was finalised on 31 March 2011.

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.htm

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

Summary

Intraday liquidity requirements in large-value real-time gross payment systems can substantially exceed the liquidity that its direct members hold overnight on their accounts with the central bank. As an illustration, UK banks' aggregate holdings of reserve balances with the Bank of England fluctuated around £30 billion in 2008, while the daily amount of liquidity that banks pass through the United Kingdom's large-value payment system, CHAPS, was in the order of £250 billion. To be able to process these payments, banks borrow additional liquidity intraday from the central bank, and recycle liquidity during the day: that is, they partly rely on incoming funds to settle their outgoing payments.

Banks contribute liquidity to the system by sending more payments than they received. We empirically investigate the effects that a hypothetical change in a single bank's payments behaviour has on the liquidity position of its counterparties. Our objective is to highlight the consequences for system-wide risk if these counterparties do not adapt their normal-time behaviour to the changed behaviour of this bank. To this effect, we first estimate banks' payments behaviour: that is, we attempt to find in the data a 'payments rule' that relates a bank's outgoing payments to its available liquidity and incoming payments. We then combine these rules to simulate payments behaviour in the system. In particular, we are interested in the effects that a change in a single bank's payments rule would have on the liquidity position of its counterparties.

We investigate two such hypothetical changes. First, a bank simply stops sending payments – perhaps because of an operational problem. If its counterparties continue to send payments to that bank, they transfer liquidity without receiving any in return from the bank that stops sending payments. Their liquidity buffer may shrink in response. Following our estimated payment rules, the counterparties reduce the value of payments they make, in turn causing the liquidity buffer of their counterparties to fall. We incorporate these spillovers in our simulation and compute, for each counterparty, the time and probability with which it is likely to run out of funds. Assuming that its counterparties do not deviate from their estimated rule, we find that the probability of at least one counterparty becoming liquidity constrained within the first hour is substantial. (In practice, the probability might be smaller, as banks' liquidity management is more sophisticated than we can capture with our model.)

The second change assumes that a bank stops providing additional liquidity to the system – perhaps because it finds itself short of liquidity, or because it becomes concerned about the other banks' ability or willingness to add liquidity to the system. Instead, it only sends out exactly what it has received. We show that such a tit-for-tat strategy would also reduce its counterparties' available liquidity. Again, we compute the time and probability with which the counterparties are likely to run out of funds, assuming that they continue to follow our estimated payment rules. We find that the probability of at least one counterparty becoming liquidity-constrained within the first hour is still substantial, although lower than in the previous case.

Finally, we attempt to identify factors that explain why changing some banks' payments behaviour has a greater effect on their counterparties than changing the behaviour of other

banks. A possible reason is that some banks are larger than others, or that they occupy more important positions in the interbank network. In our case, size appears to explain most of the variability of the *average* effect on the counterparties. More detailed information about the network helps to identify which counterparties are most at risk.

1 Introduction

We investigate how banks that are direct members of the United Kingdom's large-value payment system, CHAPS, manage their liquidity in CHAPS intraday. CHAPS handles nearly all large-value same-day sterling payments between banks, other than those relating specifically to the settlement of securities transactions. Every day, about £270 billion worth of payments are settled using the system. In such real-time gross settlement systems, these direct members – also referred to as settlement banks – rely to some extent on incoming payments to fund their own payments: in 2006, five settlement banks settled £5-£10 worth of payments for each pound of liquidity they had available at the start of the day, and five other banks settled even more than £10.

In CHAPS, as in other large-value payment systems, settlement banks fund their payments from two sources: first, from the central bank via a collateralised intraday loan at the start of each day; and second, from liquidity obtained from incoming payments. We empirically investigate the form these receipt-reactive strategies take.

Section 2 provides an overview of related literature. We hope to make an original contribution to the growing literature on banks' intraday liquidity management in payment systems (see Manning *et al* (2009) for an overview), and perhaps, via our method, to the wider literature on networks and prudential liquidity regulation.

In Section 3, we estimate a 'payments rule' for each bank: that is, the parameters of a (linear) function that relates a bank's outgoing payments to a measure of its liquidity position, allowing for the most recent changes to available liquidity to enter the function separately. We measure the bank's liquidity position as the deviation from what the bank would expect to have at that time of the day. The idea is that if the bank has more liquidity than it expected, it is more likely to send payments. We add the most recent changes in its liquidity position as a separate regressor because of anecdotal evidence that liquidity managers abide by liquidity limits when scheduling payments. When a bank is operating at its limit, balance changes are likely to be mean-reverting. Becher *et al* (2008) also found evidence for recent changes to play a role.

In Section 4, we combine these rules to simulate payments behaviour in the system. In particular, we are interested in the effects that a change of a single bank's payments rule would have on the liquidity position of its counterparties. In our first scenario, we assume that a bank simply stops sending payments – perhaps because of an operational problem. This reduces its counterparties' liquidity, and the value of payments they make in the following round. We incorporate these and further spillovers in our simulation and compute, for each counterparty, the time and probability with which it is likely to run out of funds.

In the second scenario, we assume that a bank stops providing additional liquidity to the system – perhaps because it is running short of liquidity. Instead, it only sends out exactly what it has received. We show that such a tit-for-tat strategy would also reduce its counterparties' available liquidity. Again, we compute the time and probability with which the counterparties are likely to run out of funds, assuming that they continue to follow our estimated payment rules.

In Section 5, we argue that size explains nearly all of the variability of the *average* effect on the counterparties. Our simulations help quantify this effect, and, because it uses detailed information about the structure of the payments network, helps identifying which counterparties are most at risk.

Section 6 contains a discussion of our method and compares it with related studies that use either a pure simulation framework with exogenous payment strategies, or an agent-based learning approach. The paper is, to out knowledge, the first attempt to estimate the robustness of an entire system of banks to changes in the behaviour of a single bank on the basis of estimated behavioural equations: how long it might take until other banks experience liquidity shortages; and which banks are most likely to experience such shortages.

2 Related literature

We first provide a brief overview of our methods and results to facilitate the comparison with the existing literature. We start by estimating a parsimonious model of banks' payments behaviour parametrically on the basis of their observed average behaviour, controlling for time-of-the-day effects. We then use the estimated equations in a simulation exercise, in which one bank's estimate behaviour is replaced by an exogenous 'payments rule' – for example, the afore-mentioned 'tit-for-tat' strategy. We explore how this change in behaviour affects the liquidity of the affected bank's counterparties and determine the time at which these counterparties are likely to run short of liquidity, assuming that they retain their normal-time behaviour.

A number of papers investigate optimal timing of payments in payment systems. In Bech and Garratt (2003), liquidity managers trade off the liquidity savings achieved when recycling incoming payments with the costs of delaying the execution of own payment instructions while waiting for incoming payments. Mills and Nesmith (2008) and Merrouche and Schanz (2010) extend their framework to study optimal timing when banks may be unable to send payments, perhaps because of an operational problem. Merrouche and Schanz (2010) also estimate the behaviour of that bank's counterparties during the operational outage non-parametrically and find that the stricken bank's counterparties reduce the payments to the stricken party in an attempt to save liquidity. Our paper estimates banks' payments behaviour parametrically on the basis of banks' average payments behaviour. This may underestimate counterparties' reaction when a bank is unable to send payments but provides a useful benchmark for hypothetical changes to a bank's payments behaviour.

McAndrews and Rajan (2000) and Becher *et al* (2008) estimate simple payments rules for the US and the UK large-value payment systems, respectively. Their studies show that payments behaviour is affected by the central bank's pricing schedule for intraday loans. If the central bank charges for these loans on a pro-rata basis, banks attempt to co-ordinate their payments to keep the duration of their loan as short as possible. If the central bank's charges are independent of the duration of the intraday loan, the incentive to co-ordinate is smaller. Nevertheless, Becher *et al* (2008) find some evidence that a bank's outgoing and incoming payments are positively correlated. They do not, however, control for time-of-the-day effects. Because they only observe when payments are settled, but not when the sending bank received its instruction to settle the payment, this correlation is may be only an artefact of an exogenously changing flow of payments instructions during the day. They also do not control for serial correlation of payments, giving rise to concerns that their parameters are not estimated consistently.

We also do not observe the arrival time of payment instructions. However, we attempt to avoid the pitfalls of Becher *et al* (2008) by normalising the value of payments made at a given time by the average value of payments that banks usually make at this time of day. This removes serial correlation of payments and should exclude the effects of potential variations in the flow of payment instructions on our estimates.

A number of simulation studies tested a payment system's ability to withstand shocks to one of its participants. Beyeler *et al* (2007) develop a stylised model in which banks employ simple rules to determine whether to make or delay a payment. They then explore the dynamics of liquidity recycling in the system with alternative initial conditions for system-level liquidity and different assumptions about the availability of a market for intraday interbank liquidity sharing. Renault *et al* (2007) extend this framework to a setting in which two large-value payment systems are linked via banks that participate in both systems, investigating how this can lead to spillover of shocks from one system to the other. We hope to add to these studies by estimating the rules that banks employ.

The rules that banks use in Beyeler *et al* (2007) and Renault *et al* (2007) are set using plausible assumptions rather than empirical estimates. Galbiati and Soramäki (2008) introduce agent-based learning into their model and allow the simulated banks to vary their behaviour over time. They show that strategies converge to a Nash equilibrium. We do not show that our estimated rules form a Nash-equilibrium of the 'game' that banks 'play' during the day. However, we allow banks' strategies to be more flexible: their form is inspired by previous studies; their parameters estimated by the data. We are also able to make probabilistic forecasts of the system's behaviour.

The paper is also related to the literature that investigates systemic liquidity risk at a longer horizon (up to several weeks). As in our paper, the question is how a single bank's decisions can lead to liquidity shortages at its counterparties. However, over a longer horizon, banks have more options to increase their liquidity: not rolling over loans; raising new funds in various markets; and selling various types of assets. Models that attempt to capture longer-term liquidity risk therefore only incorporate a very stylised description of a bank's behaviour: they rank the different options the bank has, and link the availability of these options to, for example, the bank's solvency position. See, for example, Aikman *et al* (2009) and van den End (2008).

Finally, our simulation takes the structure of payment flows in the United Kingdom's large-value payment system into account. Becher *et al* (2009) investigate this network in much more detail in a descriptive exercise; Soramäki *et al* (2007) conduct a similar exercise for Fedwire, a US large-value payment system. Wetherilt *et al* (2009) exploit the UK data to investigate changes in the sterling overnight loan market during 2006-08 using similar network concepts.

3 Bank-by-bank estimates of intraday liquidity management

We assume that three factors determine how a bank manages its liquidity intraday: its cost of obtaining liquidity from the central bank; how the level of its available liquidity compares to its target level of liquidity; and finally, how its available liquidity changed just ahead of its decision to send payments.

The first factor, cost of liquidity, seems the most obvious and is a key variable in the theoretical literature on payments behaviour. However, the marginal cost of intraday liquidity is not only low but also very similar for the banks in our sample. The Bank of England provides liquidity in the form of central bank money against collateral (mostly government bonds) during the day at no interest if this loan is repaid before the end of the day. And the marginal cost of pledging collateral with the Bank of England is small, as banks have so far been permitted to pledge securities they have to hold for prudential purposes, and prudential requirements only had to be fulfilled at the end of each day. (These arrangements are currently under review.[1]) The low cost of pledging collateral, together with the fact that banks receive funds from their counterparties during the day, caused our estimates of the influence of banks' cost of liquidity to be insignificant. We would only expect the cost of liquidity to play a role if each bank had to pledge an amount of securities sufficient to settle all its payments independently of incoming liquidity.

The second factor, the target level of liquidity, turns out to have more predictive power. Intraday liquidity targets are likely to be influenced by the desire to have sufficient liquidity to meet urgent payments to ancillary payment and settlement systems (such as CLS); by the bank's customer structure and business composition (which influences the typical profile of payment instructions and expected payment inflows); and by the time of the day (where the bank may be less willing to incur large overdrafts with the central bank at the end of the day). Furthermore, intraday liquidity targets are influenced by expected inflows, so that each bank's target is also determined by its counterparties' desire to hold a buffer and their business composition. Had we set up a game-theoretical model, targets would be well described by equilibrium payments behaviour. We make the standard assumption that banks' observed average behaviour corresponds to their equilibrium behaviour. Their (time-dependent) liquidity target can then be estimated as the average amount of central bank balances the bank had at a given point in time during the day, where the average is taken across all days in our sample.[2]

The third factor captures the receipt-reactive element in banks' payment behaviour. The more liquidity a bank has received from its counterparties, the more it is willing to pay. Liquidity managers suggest an interesting reason for this behaviour. Even though marginal costs of liquidity appear low at current liquidity levels, there is a risk that a bank might try to free-ride on the liquidity that its counterparties provide to the system. This free-riding could substantially increase the marginal cost of liquidity for its counterparties. By making its outgoing payments

[1] Financial Service Authority (2008), *Strengthening liquidity standards*, Consultation Paper 08/22.

[2] Notice that we retain variation in the target despite controlling for typical intraday patterns of payments. Incoming and outgoing payments have different normalisation bases: in each case, payments are normalised by the sending bank's average payment flows. We do not normalise liquidity balances.

conditional on incoming payments, each bank controls the amount of liquidity it contributes to the system. (There is an analogy to a game in which players contribute to a public good: see, for example, Manning *et al* (2009), Section 4.1.2.)

The following sections describe the data and the model we use to estimate each bank's behaviour. System-wide estimates are provided in Section 4.

3.1 Data

Our data set comprises payments aggregated in intervals of 10 minutes from the opening (06:00) to the closure (17:00) of the United Kingdom's large-value payment system, CHAPS. For each interval, the information is presented in a matrix, with payers in rows and payees in columns. Each cell contains the number of payments sent, their total value, and information about the distribution of the value (the number of payments pertaining to a certain value band). We only exploited information about the value of payments bank *i* receives in time interval (*t-1,t*] on day *d*, $Pay_{i,d,t}^{In}$, and the value contemporaneous payments *i* sends, $Pay_{i,d,t}^{Out}$. We exclude observations before 7:00 and after 16:00 because only comparatively few payments are made in these intervals.

Table 1: Summary statistics

	Typical large bank		Typical small bank	
Average value of payment made, £mn	£	817	£	201
Average value of payment received, £mn	£	818	£	200
Maximum value of payment made, £mn	£	15,764	£	2,219
Maximum value of payment received, £mn	£	14,490	£	3,043

Note: All values calculated on the basis of 10-minute aggregates. Large banks are those with the largest average daily payment values. Minimum values of payments were generally close to zero.

We study payments exchanged between the seven largest banks in CHAPS between 01/01/2007 and 31/12/2007. (We chose this interval to identify changes in payments behaviour during the first months of the financial crisis but could not identify any: intraday liquidity remained ample.) Banks are identified in our sample; however, for confidentiality reasons, we only provide anonymised results. Out of a total of 13 commercial settlement banks, we chose to include only the seven largest because the smaller banks make substantially fewer payments, which would lead to many 'zero' entries in our data set.[3] For obvious reasons, we also excluded the Bank of England and CLS, the former being a central bank, and the latter following a pre-determined schedule of payments. Table 1 provides summary statistics for our data set, separating the four largest banks from the three other banks. Figures 1 and 2 show the distribution of payment values and volumes during the day for a typical small/large bank.

Figure 1 shows that the value of payments made follows a similar pattern, independently of the bank's size, while the volume pattern differs. Generally, a large number of small payments are made at the start of the day: Values (Figure 1) are low at the start of the day, while volumes (Figure 2) are comparatively large, in particular for small banks (left panel).

[3] We capture about 80% of all payments made in CHAPS in our sample.

Figure 1: Intraday pattern of payments values for typical small bank (left) and typical large bank (right)

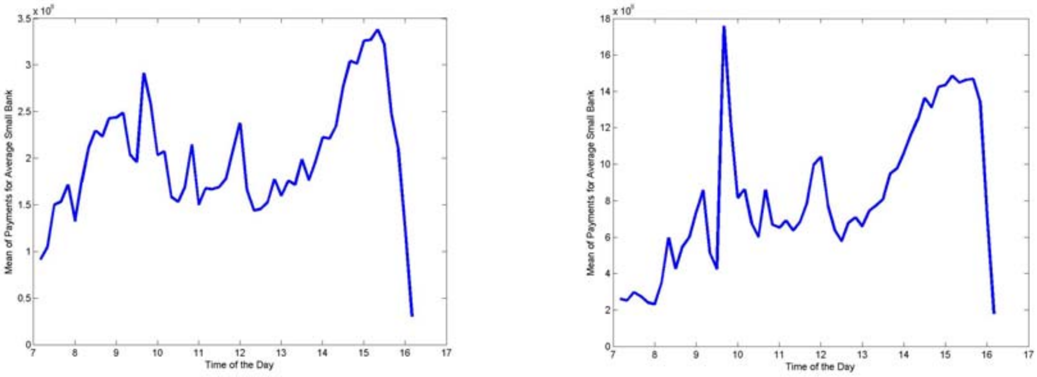

Notice: The average values of payments were calculated based on the total multilateral payments in the sample, conditional on the time of the day (7:00-16:00), over the different days of the sample.

Such small payments are settled early because they have comparatively little impact on a bank's liquidity balance but would nevertheless be costly to delay should the bank be temporarily unable to make payments because of operational problems. Total value peaks just before 10:00, before falling back and reaching its top at the end of the day. To some extent, this shape can be explained by settlement conventions in the markets where the obligations arise that are subsequently settled in the payment system, and by the settlement times in ancillary systems, such as retail payment and security settlement systems.

Figure 2: Intraday pattern of number of payments for typical small bank (left) and typical large bank (right)

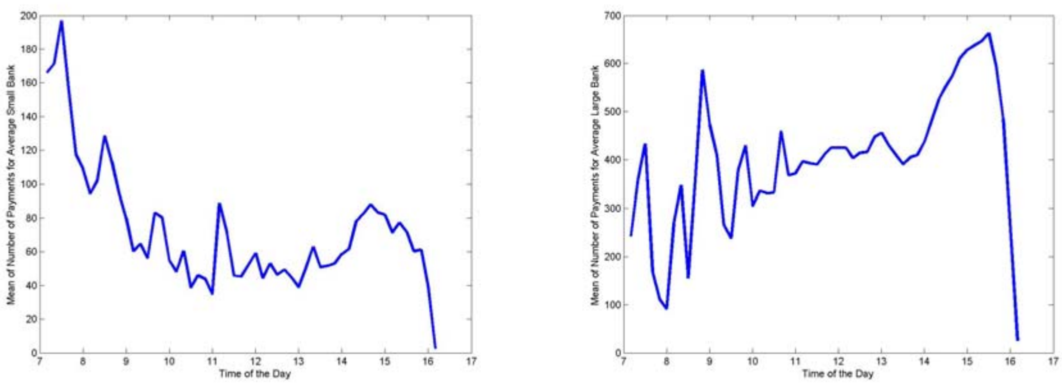

Each bank's balance – initial liquidity at the start of the day d, plus the difference between payments received and made – can then be computed as follows:

$$dayBal_{i,d,t} = dayBal_{i,d,0} + \sum_{\tau=1}^{t} Pay_{i,d,\tau}^{In} - \sum_{\tau=1}^{t} Pay_{i,d,\tau}^{Out}$$

$Pay_{i,d,t}^{In}$ - Total amount received by bank i on day d in interval t **(1)**

$Pay_{i,d,t}^{Out}$ - Total amount sent by bank i on day d in interval t

The initial value of the balance is the sum of a bank's overnight holdings of central bank money on their reserve accounts with the Bank of England, plus any central bank money borrowed from the Bank of England at the start of the day. Incoming/outgoing payments are those received from/sent to all other banks in our sample.

Our empirical models relate a bank's (normalised) outgoing payments, $Pay_{i,d,t}^{Out}$, to previous changes in the level of its balance. Section 3.2 presents first a simple model of 'short-term' liquidity management, which relates outgoing payments to balance changes that occurred in the preceding 20 minutes, and then a model of 'long-term' liquidity management, which relates outgoing payments additionally to deviations from the bank's target level of liquidity.

A technical remark before we proceed. For two reasons, our variable *dayBal* does not fully coincide with the liquidity banks have available. First, we exclude payment flows to and from the smaller banks. This should be immaterial because these flows only make up a small proportion of the value and volume of all payments. Second, we exclude inter-system transfers of liquidity. The reason is as follows. All CHAPS settlement banks for whom we estimate payment rules are also settlement banks in the UK securities settlement system, CREST. Banks can transfer liquidity between these systems. These transfers are large but infrequent and only occur mostly in the early morning and late afternoon on each day. In our estimation of payments behaviour, the corresponding large, rare changes in *dayBal* would have distorted our estimation results: effectively, they are outliers. Intra-system payments could also be treated as an additional explanatory variable: when a bank is short of liquidity in one system, it may decide to transfer liquidity from another system in which it has excess liquidity. For this paper, we decided to leave the inclusion of intra-system payments for future research.

3.2 Our hypotheses: models of intraday management of liquidity

The purpose of presenting two models is to test whether banks primarily adapt their payments behaviour to short-term changes in their balances, or whether they adopt a slightly longer-term horizon and manage their liquidity throughout the day, taking into account deviations from their target level of liquidity.

Our first model only includes a term that describes reactions to balance changes within the previous 20 minutes. This term, $shortBal_{i,d,t}$, is defined as

$$shortBal_{i,d,t} = \sum_{\tau=1}^{nLags} \left(Pay_{i,d,t-\tau}^{In} - Pay_{i,d,t-\tau}^{Out} \right) / \, nLags$$

ie, excluding payment flows in interval *t.* We set the number of lags, *nLags*, to two, after which correlations die out. (Our results turn out to be robust to the choice of lags.) Our model of 'short-term' liquidity management is given by

$$Pay_{i,d,t}^{Out*} = \alpha_i + \beta_i \cdot shortBal_{i,d,t}^{*} + \varepsilon_{i,d,t} \qquad \textbf{(2)}$$

(Stars indicate that the variables use normalised variables, see Section 3.3.) We expect α_i and β_i to be positive: α_i because some payments are likely to be made independently of incoming liquidity; and β_i because the higher the net amount of payments the bank received during the previous intervals, the more inclined it should be to make a payment.

In our second model, we investigate how the strength of the receipt-reactive behaviour depends on how much liquidity banks have available: in particular, how much more liquidity they have during the time interval under consideration compared to what they would expect to have at that time of the day. We therefore extend the specification to

$$Pay_{i,d,t}^{Out*} = \left(\alpha_i^{(2)} + \phi_{i,1} Z_{i,d,t}\right) + \left(\beta_i^{(2)} + \phi_{i,2} Z_{i,d,t}\right) \cdot shortBal_{i,d,t}^{*} + \varepsilon_{i,d,t} \qquad \textbf{(3)}$$

$Z_{i,d,t}$ is the relative deviation by which liquidity holdings as measured by *dayBal* at the start of period t exceed bank i's target for the end of period t.

$$Z_{i,d,t} = \left(dayBal_{i,d,t-1} - target_{i,t}\right) / target_{i,t}$$

Bank i's target is taken to be the bank's average end-of-period-t liquidity holding, where the average is taken across all days of our sample:

$$target_{i,t} = E_d\left(dayBal_{i,d,t}\right) = \sum_{d=1}^{nDays} dayBal_{i,d,t} / nDays$$

Figure 3 shows the typical pattern of the variable *target*, where the opening balance is normalised to zero. (Recall that our sample starts only one hour after the opening, so that the first observation in the figure is unequal to zero.) In practice, banks' opening balance is *not* zero, as they hold reserves with the central bank overnight. As mentioned above, banks also obtain a secured credit line from the central bank to settle their payments. In addition, as we explained for the construction of *dayBal*, the measure ignores intra-system transfers, and payments between banks included in the sample and those that are not. What the figures reveal, however, is that banks are willing to contribute liquidity to the system during the first hours of the day, allowing their balance to fall, and absorb liquidity towards the end of the day, ensuring that they thereby 'repay' any intraday loan the central bank provided.

Figure 3: Intraday pattern of our estimated average value of liquidity for a typical small bank (left) and typical large bank (right) in the sample.

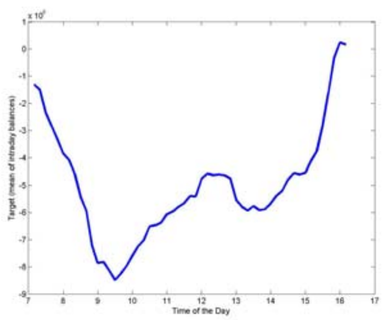

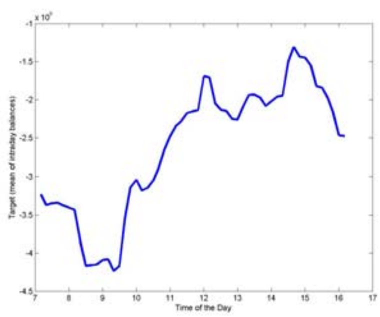

For the model in **(3)**, we expect $\alpha_i^{(2)}$ and $\beta_i^{(2)}$ to be positive for the same reasons as α_i and β_i in Model 1. We expect both $\phi_{i,1}$ and $\phi_{i,2}$ to be positive as well: a bank is more willing to make additional payments – both autonomously and in response to incoming payments – if its balance is above target. Notice that the right-hand side variables *shortBal* and *Z* predict payments behaviour in [t-1,t) and are defined with respect to payments that occurred in or before interval [t-2,t-1), helping us to avoid endogeneity problems.

3.3 Results

Before estimating Models 1 and 2, we normalise the payment values by each bank in each time interval by the average amount of payments made in this interval. Figure 4, which presents the auto-correlogram of payments made and received, illustrates the reason: payments behaviour does not change rapidly during the day, and is similar during the same interval on different days. We remove this seasonal effect to be able to identify receipt-reactive payments behaviour.

Figure 4: Auto-correlogram of the value of payments made and received for a typical bank in the sample, before and after normalisation.

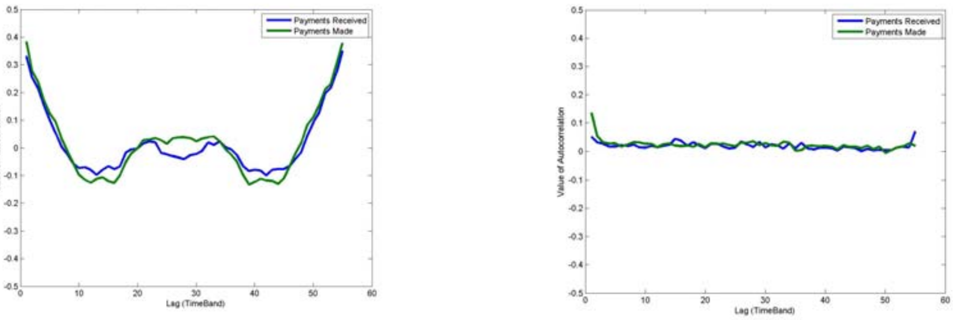

Notice: Each point in each chart measures the (auto-)correlation between payments sent during interval *t* and during interval *t-x* (green line), and correspondingly for payments received (blue line). For example, the autocorrelation between the values of payments received in two adjacent intervals is about 0.32 before normalisation (left panel), and 0.12 after normalisation (right panel).

Before normalisation, the auto-correlogram shows comparatively high correlation for the first five lags (that is, the previous hour of the same day), and for lags 50-55 (that is, for the subsequent hour on the previous day). After the normalisation, the autocorrelations are insignificant.

We estimate Models 1 and 2 bank by bank using OLS.[4] Table 2 contains the results for Model 1. The mean value of payments made, which is estimated by α_i, is virtually equal to 1 after the normalisation because the mean value of *shortBal* – the change in the payer's balance during the two preceding intervals – is virtually equal to zero. More interestingly, the β_i are significant at the 1% or 5% level for all but one bank and have the expected signs. We conclude that liquidity managers typically increase the amount of payments they make when they have received more than they sent during the preceding 20 minutes. The mean (unweighted) reaction coefficient is $\sum \beta_i / 7 = 0.13£$. This means that, for each additional pound sterling that a bank receives more than it sent during the previous 20 minutes, it sends[5] out on average an additional of £0.13.

[4] Notice that even though the regression models form a system (what one bank sends is received by another), all dependent variables are lagged. Our model therefore has a vector autoregressive (VAR) representation (see Appendix 3) and the OLS solution delivers efficient estimates.
[5] See Appendix 2 for the derivation of the parameter's interpretation.

While all parameters are significant at the 5%-level, the adjusted R^2 values in Table 2 are comparatively small. One explanation is that settlement banks appear to pass on most payment instructions immediately to the settlement system unless their liquidity balance falls below a limit, in which case they wait for incoming payments before releasing further payments. Outgoing payments are therefore likely to follow the random arrival process of payment instructions for most of the time. Only when the liquidity balance hits its lower bound, outgoing payments are dependent on incoming payments.[6] This also explains why empirically, the most recent contributions appear to have an influence on the bank's decision to make payments.

Table 2: Dependent variable is $Pay_{i,d,t}^{Out}$, the value of payments made

	$\hat{\alpha}_i$	$\hat{\beta}_i$	R^2_{adj}
Bank 1	1.00***	0.14***	1.6%
Bank 2	1.00***	0.22***	2.2%
Bank 3	1.00***	0.17***	1.9%
Bank 4	1.00***	0.05*	0.2%
Bank 5	1.00***	0.23***	2.1%
Bank 6	1.00***	0.01**	0.0%
Bank 7	1.00***	0.07***	0.2%

Note: The estimated equation is $Pay_{i,d,t}^{Out*} = \alpha_i + \beta_i \cdot shortBal_{i,d,t}^{*} + \varepsilon_{i,d,t}$ (equation **(1)**). α should be interpreted as the autonomous payments made in each interval; β as the value of payments made in response to recent changes in the liquidity balance. Newey West's covariance matrix has been used to compute standard errors. *** indicates significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 3 shows the corresponding results for Model 2. The deviation of actual liquidity from the bank's target – that is, from its typical liquidity holding at the given time of day – generally adds to the explanatory power of Model 1. The coefficients have the expected signs. Compared to Model 1, the values of the β_i remain broadly unchanged. The ϕ_i are also positive, indicating that a bank is more willing to make additional payments – both autonomously and in response to incoming payments – if its balance is above target. For example, suppose a bank makes, on average, a payment of 100£[7] in each time interval. If this bank's balance exceeds the target by 1%, the average bank sends[8] an additional $100 \cdot \sum \phi_{i,1} / 7 = 58£$ independently of incoming payments, and an additional $\sum \phi_{i,2} / 7 = 0.32£$ in response to a £1 increase in *shortBal* – the change in the payer's balance during the two preceding intervals. All but two coefficients are statistically significant at the 1% or 5% level.

[6] We have attempted to identify these limits directly in the data. On some days, they appear to be clearly identifiable. However, they appear to vary with the time of the day: a bank may decide to immediately settle an urgent payment, even if that falls below a previously set limit. Also, limits may be set on a bilateral and on a multilateral basis. We found it impossible to reliably separate limits from observed minimum balances, and opted instead for including our 'target' measure as a determinant for outgoing payments.

[7] This is equivalent to $E\left(Pay_{i,d,t}^{Out}\right) = E\left(A_{i,d,t}\right) = 100£$.

[8] See Appendix 2 for derivations.

Table 3: Dependent variable is $Pay^{Out}_{i,d,t}$, the value of payments made

	$\hat{\alpha}_i^{(2)}$	$\phi_{i,1}$	$\hat{\beta}_i^{(2)}$	$\phi_{i,2}$	R^2_{adj}
Bank 1	0.97***	0.98***	0.14***	0.78***	1.6%
Bank 2	0.99***	0.06*	0.21***	0.23***	1.8%
Bank 3	0.97***	1.50***	0.11***	0.50***	2.2%
Bank 4	0.98***	0.37*	0.05**	0.48***	0.5%
Bank 5	0.99***	0.39***	0.17***	0.07	2.6%
Bank 6	0.99***	0.25***	0.01	0.11	0.1%
Bank 7	1.00***	0.53***	0.04***	0.03	0.7%

Note: The estimated equation is $Pay^{Out*}_{i,d,t} = \left(\alpha_i^{(2)} + \phi_{i,1} Z_{i,d,t}\right) + \left(\beta_i^{(2)} + \phi_{i,2} Z_{i,d,t}\right) \cdot shortBal^*_{i,d,t} + \varepsilon_{i,d,t}$, where Z is the relative deviation of a bank's balance from its target (equation **(2)**). α should be interpreted as the autonomous payments made in each interval; β as the value of payments made in response to recent changes in the liquidity balance. $\phi_{i,1}$ measures the extent to which autonomous payments are larger when the bank's liquidity exceeds its target; $\phi_{i,2}$ the extend to which its reaction to recent changes in the liquidity balance changes when its liquidity exceeds its target. Newey West's covariance matrix has been used to compute standard errors. *** indicates significance at the 1% level, ** at the 5% level, and * at the 10% level.

4 System-wide estimates of banks' liquidity management

In this section, we use the estimated equations of Model 1 to simulate the properties of the interbank network. In particular, we are interested in the external effects of two plausible modifications of a bank's strategy compared to its estimated behaviour in normal times: in the first scenario, it stops sending payments altogether (a common result of operational problems); in the second, it adopts a non-cooperative behaviour, stops sending autonomous payments, and only sends exactly what it receives – perhaps as a result of an unanticipated liquidity shortage. This changed behaviour is assumed to start at the beginning and to remain in place until the end of the simulation.

4.1 A system-wide model of payments made and received

To simulate the behaviour of the payment system, we need

- a 'balance rule' describing the evolution of each bank's balance, and an initial value for these balances; and
- a 'payments rule' describing how much each bank pays to whom.

Equation **(1)** describes the behaviour of bank's balances and is reproduced here for easier reference (we drop the index *d* to economise on notation):

$$dayBal_{i,t} = \hat{dayBal}_{i,0} + \sum_{\tau=1}^{t} Pay^{In}_{i,\tau} - \sum_{\tau=1}^{t} Pay^{Out}_{i,\tau}$$

where the initial value of the balance is estimated using the sample average of the bank's available liquidity:

$$\hat{dayBal}_{i,0} = \frac{1}{T} \sum_{t=0}^{T} dayBal_{i,t}$$

Using the average balance rather than the initial balance provides us with a time-independent start of our simulation.[9] Notice that we do not normalise payments in the simulation: we only needed to do this in the estimation step in order to consistently estimate the parameters of the payments rule β_i.

Equation **(2)** provides the basis for the payments rule we use in the simulation. It needs to be complemented by two additional ingredients. First, we need to ensure that the payments rule is feasible: a bank's payments cannot exceed its available liquidity. Equation **(4)** contains that rule, using the estimates obtained in Section 3. Because we work with the original data in our simulation, we scale up the intercept $\hat{\alpha}_i$ by bank i's average payment value $E_\tau\left(Pay_{i,\tau}^{Out}\right)$.[10]

$$Pay_{i,t}^{Out} = \begin{cases} \hat{\alpha}_i E_\tau\left(Pay_{i,\tau}^{Out}\right) + \hat{\beta}_i \cdot shortBal_{i,t} + \eta_{i,t}^{payOut} & \text{if } \ Pay_{i,t}^{Out} \le dayBal_{i,t-1} + Pay_{i,t}^{In} \\ 0 & \text{if } \ Pay_{i,t}^{Out} > dayBal_{i,t-1} + Pay_{i,t}^{In} \end{cases} \qquad \textbf{(4)}$$

Equation **(4)** describes each bank's aggregate payment outflows, but not its inflows. Had we estimated equivalents to equation **(2)** on a bilateral basis, we could construct inflows from outflows. However, the coefficients of the bilaterally estimated payment rules had different signs and were occasionally not significant. We therefore simply assume that out of its outgoing payments, bank i sends in each interval a constant[11] share $\gamma_{i\to j}$ to each of its counterparties j. We estimate these shares by simple average payment shares over the sample period:

$$\hat{\gamma}_{i->j} = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{Pay_{i\to j,t}^{Out}}{\sum_{k\neq i}^{nBanks} Pay_{i\to k,t}^{Out}}\right) \qquad \textbf{(5)}$$

Payment flows to bank i are simply the sum of its counterparties' total payments times these payment shares:

$$Pay_{i,t}^{In} = \sum_{j\neq i}^{nBanks} \hat{\gamma}_{j->i} Pay_{j,t}^{Out} \qquad \textbf{(6)}$$

4.2 Simulation procedure

In each of the following scenarios, we substitute an exogenous payments rule for the estimates rule **(4)** for a single bank k (all else remains unchanged).
In Scenario 1, where a bank k is unable to pay (but still able to receive payments), we modify **(4)** to

$$Pay_{k,t}^{Out} = 0 \qquad \textbf{(4)'}$$

[9] An interesting extension would be to draw out the effects of a shock that hits the system at a specific time of the day.
[10] See Appendix A2 for the derivation.
[11] Allowing for time-varying shares and taking into consideration the uncertainty of their estimates in the simulation is left for future research.

for all $t>0$, and one k at a time. In Scenario 2, where a bank k pays exactly what it received in the current[12] interval, we modify **(4)** to

$$Pay_{k,t}^{Out} = Pay_{k,t}^{In} \qquad \textbf{(4)''}$$

In each case, we are interested in the effect of the change in behaviour on the other banks' available liquidity, assuming that these banks continue to follow their estimated payments rule **(4)**. In particular, we want to determine when they would run out of liquidity under these behavioural assumptions.

We undertake a deterministic and a stochastic simulation. In the deterministic simulation, we compute recursively each interval the level of liquidity each bank has available, setting the error in **(4)** to zero. In the stochastic simulation, we sample from the error distribution of $\eta_{i,t}^{payOut}$ in the system of equations given by **(4)** to derive the probability with which one, two, or even more banks would run out of liquidity if they did not change their estimated behaviour. In order to obtain a better impression for the properties of the system, we let both simulations run over 20 hours.[13] In reality, liquidity holdings are 'reset' at the start of each day when banks pledge new collateral to the Bank of England, which lasts for 10.5 hours in CHAPS. Thus, any result showing a bank running out of liquidity only after 10.5 hours effectively means that we would not expect to observe any bank running out of liquidity in reality, even if the shock hit at the start of the day.

4.2.1 Scenario 1: deterministic simulation

To give an impression of the strength of links between the banks in the system, Table 4 shows the immediate effect of bank i's (in columns) changed payment rule on the other banks' expected balances in the subsequent ten-minute interval.[14] Numbers are expressed relative to the other banks' balances at the time of the change in i's payment rule. For example, the first column shows the effect of bank a ceasing to send payments on its counterparties' liquidity: bank b's available liquidity is reduced by 0.2%, bank c's by 1%, and so on. How strong the bank is affected depends on the structure of the payment network, that is, on $\hat{\gamma}_{i->j}$ (equation **(5)**), and on the size of the bank that stops sending payments. We re-label the banks randomly in this table for confidentiality reasons.

But there are indirect effects as well, which Table 5 does not capture. Smaller receipts reduce the inclination of the bank's counterparties to send payments: in **(4)**, $shortBal_{j,t+1}$ falls compared to base case. The effect of the bank's changed payments rule therefore cascades through the system over time.

[12] That is, we assume that bank k pays 'at the end' of the interval, having observed what it received during the interval. There is no indeterminacy problem as the other banks' strategies depend on the change in their balance up to, but not including, the current interval. Alternatively, we could have assumed that bank k pays what it received during the preceding interval.
[13] That is, we have 6 ten-minute intervals over 20 = 120 iterations of the algorithm.
[14] The mathematical derivations for the expected liquidity shortfall are in Appendix 1.

Table 4: Immediate expected effect of a bank's failure to send payments

		Bank that stops sending payments						
		Bank a	Bank b	Bank c	Bank d	Bank e	Bank f	Bank g
Banks that continue to follow their payment rules	Bank a		-0.5%	-0.2%	-0.1%	-1.6%	-1.5%	-0.3%
	Bank b	-0.2%		-0.5%	-0.3%	-2.3%	-2.2%	-0.7%
	Bank c	-1.0%	-6.0%		-1.4%	-15.6%	-14.5%	-4.8%
	Bank d	-0.1%	-1.7%	-0.6%		-3.1%	-2.6%	-0.7%
	Bank e	-0.5%	-3.2%	-1.5%	-0.7%		-6.6%	-2.0%
	Bank f	-0.6%	-3.3%	-1.6%	-0.6%	-6.8%		-2.0%
	Bank g	-0.1%	-1.0%	-0.5%	-0.2%	-2.1%	-2.1%	

Despite this reduced inflow of payments, the other banks do not commensurately reduce their outflows in our simulation: the estimated payments rule showed that banks make a substantial proportion of their payments independently of preceding balance changes. As a result, the other banks only risk running out of funds some time after bank 1's failure to send payments. Figure 5 plots these results. Each line of the chart plots the expected number y of illiquid counterparties x hours after *i* stopped sending payments to them. (The y-axis ranges from 0 to 6 counterparties; the x-axis from 0 to 20 hours.)

For an example, consider bank 7. About 2.5 hours after bank 7 stopped sending payments, one of its counterparties would become illiquid if it continued to follow the payments rule that it follows in normal times. About 6 hours after it stopped sending payments, two counterparties would be illiquid, and so on.

Figure 5: Effect of failure of bank *i* to send payments on its counterparties

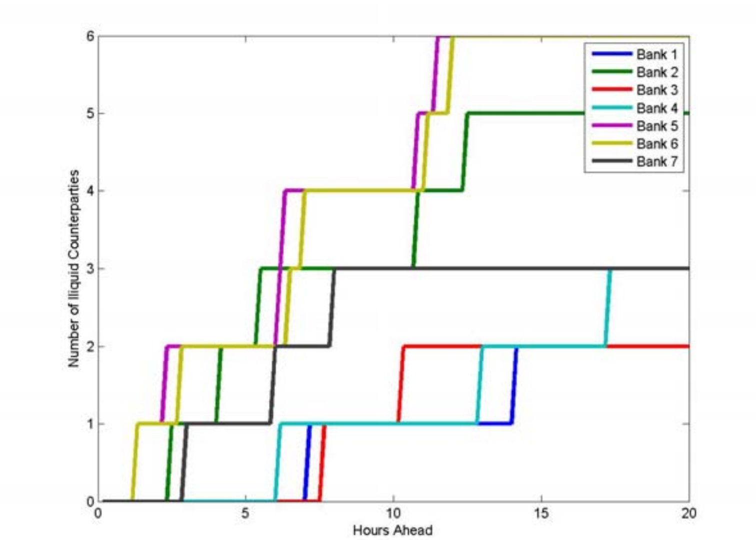

Notice: Each line of the chart plots the expected number y of illiquid counterparties x hours after *i* (see legend in figure) stopped sending payments to them. The y-axis ranges from 0 to 6 counterparties; the x-axis from 0 to 20 hours.

Clearly, the consequences of a bank's failure to send payments on its counterparties depend on that bank's properties. Broadly, there are two groups of banks: the failure of a bank in the first group (banks 2, 5, 6 and 7) to make payments causes the illiquidity of at least two other banks within the first six hours of the change in behaviour. For the remaining three banks, the external effects are smaller.

To reiterate, these expected times at which the counterparties become illiquid have been derived assuming that the counterparties do not deviate from their payments rule. In practice, banks' liquidity management is more sophisticated than we can capture with our model. Merrouche and Schanz (2010) provide evidence that banks also monitor their bilateral balances, and reduce their payments to a counterparty from which they do not receive funds. However, there are only few instances in which banks experience sufficiently long operational problems that prevent them from sending funds. (Merrouche and Schanz (2010) investigate eight instances.) We were therefore forced to estimate our payment rules on data that primarily included days in which no bank experienced operational problems, and could not find systematic evidence for *bilateral* receipt-reactive behaviour. Thus, our payments rules **(4)** only relate payments to a bank's total (multilateral) changes in liquidity, and therefore tends to underestimate a counterparty's reaction to a bank's inability to send payments. Finally, the payment system control announces operational problems of a bank to all its counterparties to support counterparties' efforts to adjust their payments behaviour.

4.2.2 Scenario 2: stochastic simulation of an operational failure

While the deterministic simulation was used to compute expected times of illiquidity, the stochastic simulation predicts the likelihood of counterparties' illiquidity. These probabilistic estimates are presumably more useful for risk-assessment purposes because they include model uncertainty: that is, the error that our estimated payments rule makes when predicting actual payments behaviour.

To construct the probabilistic estimates, we take, for each simulated ten-minute interval over a total of 50 hours, independent samples from the distribution of the sample errors of the payment rule (4):[15]

$$\hat{\eta}_{i,t}^{payOut} = Pay_{i,t}^{Out} - \left(\hat{\alpha}_i E_\tau\left(Pay_{i,\tau}^{Out}\right) + \hat{\beta}_i \cdot shortBal_{i,t}\right).$$

With these error terms we constructed one simulated path of balances for all banks – as in the deterministic case, the only difference being that in the deterministic case, the error term was set to zero. We repeated this procedure 5,000 times to construct 5,000 such paths for payments made, received and the corresponding balances for each of the seven banks in sample. For each interval of the simulated path, we then computed the proportion of simulations paths for which zero, one, two, etc… banks were illiquid. (We call a bank 'illiquid' when, according to its estimated payment rule, it should make a payment, but cannot because it lacks the necessary funds (see the case distinction in the definition of the payments rule **(4)**.) This proportion is then interpreted as the probability that zero, one, two, etc… banks would indeed become illiquid.

[15] Notice that we neglect any serial correlation of the error terms when sampling errors. We know that the cause of serial correlation is the intraday pattern (see Figure 3). However, for this paper, we are interested in explaining the part of payments behaviour that cannot simply be explained by the time of day. Therefore, we chose to ignore this pattern in our simulations. The variance of the residuals in our model is correspondingly larger. The inclusion of a time dependency in the simulation is left for future research.

This procedure is similar to the usual impulse response analysis in a VAR model (see Greene (2003), page 593).

To obtain an impression of the impact of uncertainty on these estimates, it is useful to first analyse our system for the case that all banks follow their normal payments rule **(4)**. Figure 6 shows that even in this case, there is a non-zero (about 25%) likelihood that a bank becomes illiquid within the first hours of the operations of the payment system.

Figure 6: Probability of banks becoming illiquid if their estimated payment rules are disturbed randomly

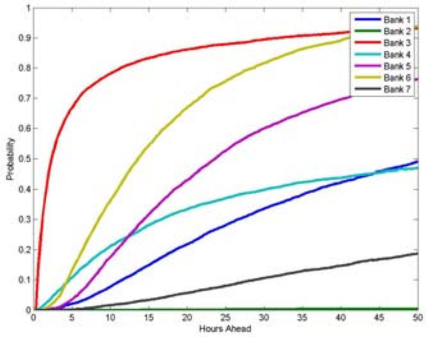

Notice: Each line of the chart plots the probability with which the corresponding bank will become illiquid when each payment rule is shocked in each interval of the iteration, x hours after the first shock.

The reason for the simulated non-zero likelihood of failure is that some banks have an apparently relatively small liquidity reserve in the payment system and manage their payments correspondingly tightly. Our payment rules, however, are only estimated with a substantial error. In the simulations, this error shows up as large random deviations from the bank's deterministic payment rule. In reality, a bank with an apparently relatively small liquidity reserve would, of course, attempt to avoid such random errors – and they would ultimately also change their payments rule to avoid illiquidity. Finally, notice that the cumulative probability distributions become successively flatter because the uncertainty in our payment rule adds up over the intervals over which we project the system into the future.

Figure 7: Probability of a number of counterparties becoming illiquid if a small (large) bank stops making payments (systemic risk)

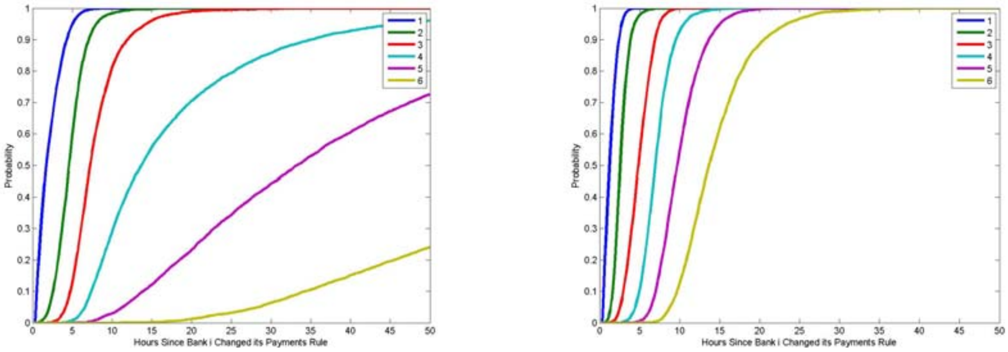

Notice: For each chart, each line plots the probability with which one, two, three, etc, of bank *i*'s counterparties become illiquid x hours after bank *i* stops sending payments. The x-axis shows the time since *i* changed its payments behaviour (from 0 to 50 hours). All counterparties are assumed to follow their normal-time payments behaviour.

Now back to Scenario 1, in which one bank does not send any payments during the entire simulation. For Figure 7, we picked a small bank (left panel) and a large bank (right panel) out of our sample and assumed that it would not send any payments during the simulation period. Figure 7 plots the likelihood that one (blue line), two (green), three (red), etc… of its counterparties become illiquid, again assuming that they follow a noisy version of their estimated payments rule.

Figure 7 shows that with near-certainty, one bank will become illiquid within the first six hours after a small bank's failure to send payments (within the first four hours after a large bank's failure to send payments); see the blue lines. For this large bank, three counterparties would become illiquid after approximately four hours with a probability of about one third (red line). Again, recall that these results were derived under the assumption that banks follow a noisy version of our estimated payment rules.

4.2.3 Scenario 2: stochastic simulation of a free-riding strategy

We now turn to the stochastic simulation of the consequences of a bank reverting to a (deterministic) tit-for-tat strategy **(4)''** at the start of the simulation: it sends out exactly the same value of payments that it received in the same interval.[16] A bank might revert to this strategy if it becomes concerned about the other banks' ability or willingness to return liquidity, or because it itself experiences a surprise shortage of liquidity. *Ex ante*, it is not clear whether the tit-for-tat strategy leads to lower or higher payments by the bank that plays according to **(4)''**: on the one hand, no autonomous payments would be made; but on the other, more payments may be made in response to incoming payments.

Figure 8: Probability of counterparties becoming illiquid if a small (large) bank reverts to a tit-for-tat strategy

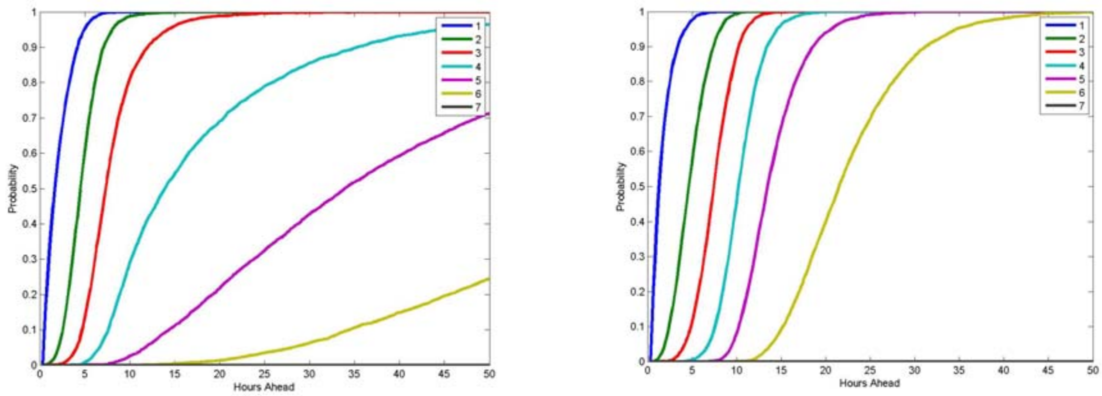

Notice: For each chart, each line plots the probability with which one, two, three, etc, of bank *i*'s counterparties become illiquid x hours after bank *i* reverts to a tit-for-tat strategy. The x-axis shows the time since *i* changed its payments behaviour (from 0 to 50 hours). All counterparties are assumed to follow their normal-time payments behaviour.

Figure 8 shows the results of the simulations, assuming that the same banks that stopped sending payments in Scenario 1 instead revert to the receipt-reactive strategy **(4)''**. It reveals that this bank indeed sends less liquidity than if it had followed the original rule **(4)**: the likelihood of a

[16] We do not report the less interesting case of the deterministic simulation.

certain number of its counterparties becoming illiquid lies in between the cases in which all play according to **(4)**, and Scenario 1, in which the bank followed **(4)'**.

For risk-assessment purposes, these results may be more important than those in Scenario 1. The reason is that receipt-reactive behaviour is less easily detectable for the bank's counterparties than complete failure to pay. And the bank that uses such a strategy would not be under the obligation to inform system control of its changed behaviour. Thus, our assumption that the other banks adhere to their standard payment rules **(4)** is more realistic than in Scenario 1. Of course, it remains reasonable to assume that the counterparties would divert from that strategy once their available liquidity has substantially fallen.

5 Size versus network attributes

In the preceding sections, we often contrasted large and small banks. 'Large' banks' failure to send payments, or their adoption of a tit-for-tat strategy, had larger external effects on the rest of the system. In theory, size is not the only network attribute that determines the degree to which shocks to the payments behaviour of a single bank are amplified throughout the system. The structure of the network matters as well.[17] Star networks depend on the operational availability of the central node. Fully connected networks are more robust unless shocks become very large. But, beyond that, a well-connected system can flip: interconnections can provide the mechanism that allows losses to cascade across all institutions.

The structure of a network can be described by a large number of variables, measuring, for example, how large the banks in the network are, and how they are linked. A recurring question in this context is whether size can more or less completely explain the influence of a bank in the network – in our case the effect of a changed payments rule. The number of banks in CHAPS is too small to investigate this hypothesis formally. In our simulation, the scaled intercepts in **(4)** essentially measure the size of the bank's average payment flows. But they are also a measure of the value of payments a bank makes independently of the payments it receives, and therefore a key component of the impact that a change in counterparty's payments rule has on the bank's available liquidity. (More precisely: the smaller the slope coefficient, the better the intercept describes the impact.) Our measures of size and impact are therefore closely correlated.

Admittedly, the small errors that such a univariate prediction would make would add up over time as banks repeatedly apply their payment rules to react to previous incoming payment flows. However, in practice, payment rules might not remain constant, and are more likely to change the more time banks have to become aware of bank *i*'s original change in behaviour. The quality of a prediction that takes additional network characteristics may not improve.

However, a more detailed knowledge of the network structure has advantages. It helps predict bilateral relationships: suppose bank *i* becomes unable to send payments, which counterparty is most likely to suffer first? And how long would its liquidity last? Figure 9 provides examples, again choosing the same pair of banks for which Figures 7 and 8 were constructed. Contrast

[17] Aikman *et al* (2009) analyse the amplification of liquidity shocks in a quantitative, simulation-based model of systemic stability based on the network of UK banks. Haldane (2009) emphasises the analogies between financial network with biological networks.

Figures 7 and 9: Figure 7 shows the system-wide impact of a small (left panel) and a large (right panel) bank failing to make payments, while Figure 9 shows the vulnerability of a specific bank to a small bank (left panel) and a large bank (right panel) failing to make payments.

Figure 9: Probability of specific counterparties becoming illiquid if a small (large) bank stops making payments (bank-specific risk)

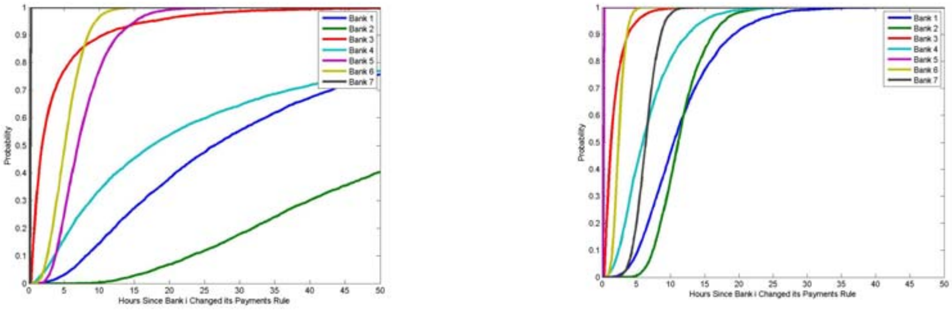

Notice: For each chart, each line plots the probability with which bank *i*'s counterparties become illiquid x hours after bank *i* stops sending payments. The x-axis shows the time since *i* changed its payments behaviour (from 0 to 50 hours).

Clearly, the size of the bank that changes its payments rule again determines the general shape of the picture. However, while the left panel could suggest that bank 2 is considerably less vulnerable to a small bank stopping to make payments than bank 1 (the green line lies considerably below the blue), the right panel shows that they are about equally vulnerable to a large bank stopping to make payments (the green line lies close to the blue). Predictions purely based on the size of the bank that originally changed its behaviour could be misleading.
We conclude that a detailed knowledge of the network of payment flows between banks can help predict the likelihood with which a bank becomes illiquid if one of its counterparties deviates from its typical payments behaviour.

6 Comparison of our approach with existing simulations

This section compares our approach with other simulation methods that have been used to study the robustness of payment systems, drawing out their relative advantages and disadvantages.

Several sources of information, and diverse sets of analytical tools, can be used to assess the robustness of a payment system to exogenous liquidity shocks (cf. Manning *et al* (2009)). The key question is how to best model banks' behaviour. Researchers using game-theoretic techniques are often required to abstract significantly from reality to be able to solve their models and are often constrained by the sometimes strong assumptions required to deliver analytically tractable solutions. Another possibility, which we chose, is to allow for a more complex environment but restrict the strategies that banks can follow.

One option is to specify banks' strategies exogenously – both their form and their parameterisation. Examples include Beyeler *et al* (2007) and Renault *et al* (2007). Their base case is that banks make payments whenever they have sufficient liquidity. The authors can then test the robustness of their results by varying the strategies. This comparatively simple treatment permits the description of the system to be relatively complex. (For example, Renault *et al*

(2007) model liquidity flows between two interlinked payment systems.) On the other hand, it is more suited to deliver qualitative rather than quantitative predictions of actual behaviour.

Another option is to specify only the form the payment rules take and leave its parameterisation open. We are aware of two variants in the payments literature: the first employing agent-based modelling techniques; the second using econometric estimates to determine the parameters.

In the first variant, which uses agent-based modelling techniques, the parameters are determined by a learning algorithm. Banks are given a simple objective function – for example, to minimise both the delay between receipt of a executing payments. The system is then simulated for an initial set of parameter values, which banks are able to adjust following an exogenous adjustment rule once an entire path has been simulated. Galbiati and Soramäki (2008) use this method. Its advantage is that once a sufficient number of paths has been simulated and the parameter values have converged, one may be able to show that the resulting strategies form a Nash-equilibrium, that is, that they are mutually optimal (subject to the form of the payments rule and the objective function). The disadvantage is that the form of the decision rule has to be simple to allow the optimal behaviour to be determined; that the objective function may not be correctly specified; and that computational requirements are substantial. In particular, the dimensionality of the problem would increase substantially if the aim was to derive probabilistic statements.

We employ the second approach, which uses econometric techniques to estimate banks' payments rules, and then use these rules to forecast the system's behaviour. Compared to the agent-based approach, the link to actual, observed behaviour is closer. However, it is less clear to what extent behaviour that we observe in normal times also describes behaviour in unusual situations, such as during operational outages, or when one bank adopts a tit-for-tat strategy.

Relatedly, we do not verify whether our estimated payments rules would be optimal in a theoretical model of payments behaviour. This, combined with the comparatively small explanatory power of our regressions, might reduce confidence in the validity of the estimated payment rules. However, because the numerical complexity of simulating the system's behaviour is smaller, we can provide probabilistic forecasts. For these, we do not rely on distributional assumptions but sample directly from the errors of the estimated payment rules. Probabilistic forecasts should be particularly useful for risk assessment purposes.

Finally, a few remarks on how our approach relates to a classic VAR model. One can think of our payment rules forming a VAR, with the auto-regressive terms given by *shortBal*, where the coefficients of the two lags are constrained to be identical. In order to obtain impulse-response functions, we could simply have traced the effect of a shock to one of the equations through the system. For example, a negative shock to one of the equations would have corresponded to a reduction the payments the respective bank makes. Shocking the system in this way would, however, not have taken an important constraint into account: that is, that a bank cannot make a payment if it lacks liquidity. In our simulation, we integrate this constraint in equation **(4)**. (Figure 5 showed the results.) In addition, our main interest lies in the effect of specific changes to a bank's behaviour – the inability to send payments, or the adoption of a tit-for-tat strategy. These cannot be satisfactorily described by a shock to the payments rule.

In summary, we have most confidence in our forecasts of the system's behaviour when we investigate small shocks to a bank's payments rule which are difficult to detect for other banks. Other banks would then be unlikely to change their behaviour from the one we estimated in normal times. Our results of Scenario 2, where a bank simply becomes stingier with liquidity but does not entirely stop paying, may therefore be more robust than those we derived in Scenario 1.

7 Conclusions and further research

When settling their own liabilities and those of their clients, settlement banks rely on incoming payments to fund a part of their outgoing payments. We investigate their behaviour in CHAPS, the United Kingdom's large-value payment system. Our estimates suggest that in normal times, banks increase their payment outflows when their liquidity is above target and immediately following the receipts of payments. We use these estimates to determine the robustness of this payment system to two hypothetical behavioural changes. In the first a single bank stops sending payments, perhaps because of an operational problem. In the second, it adopts a free-riding strategy, paying out exactly what it previously received.

Using the errors of the estimated behavioural equations, we derive probabilistic statements about the time at which the bank's counterparties would run out of liquidity if they followed their estimated normal-time behaviour. In both cases, we show that there is a considerable likelihood that at least one counterparty will become illiquid within the first hour after the change in behaviour if it continues to follow our estimated payments rule. The impact is larger the larger the bank that changes its behaviour, and larger when the bank stops sending payments.

Importantly, our payment rules are estimated across days during which such changes to payments rules were rare. (In fact, we only know about times when a bank stopped sending payments.) Because our estimates average over normal days and those were the payments rule changed, it presumably underestimates the receipt-reactive part of the payment rule and exaggerates the impact of a sustained change of that bank's behaviour. We have more confidence in our forecasts of the system's behaviour when we investigate small shocks to a bank's payments rule which are difficult to detect for other banks, such as the adoption of a free-riding strategy.

There are at least two interesting avenues for further research. The first is a less parsimonious description of banks' behaviour. For example, one could use the extended payments rule in the simulation, which also takes the deviation of a bank's actual liquidity holdings from its target into account. Alternatively, one could estimate two payment rules, one that applies once a behavioural change of a counterparty becomes known to other system participants, and the other describing behaviour in normal times. Ideally, intra-system transfers should be integrated in the estimated payments behaviour, as an additional explanatory or as a dependent variable.

The second extension could be to add further scenarios to our simulation. For example, we could investigate the effects of the time of day at which a change of payments behaviour occurs on other banks' liquidity position.[18] Or we could allow behaviour to revert to the original estimated payments rule after some time – a good description of a temporary operational problem.

[18] Merrouche and Schanz (2010) show that shocks to payments behaviour which occur in the morning have a greater effect than those that occur in the afternoon.

Appendix 1: Derivation of the expected liquidity shortfall

Given failure to pay of bank *m* for time *t*, ensuring that bank *i* is unaffected for time *t+1* implies that the balance process for *t+1* is unchanged in between the cases. Formally:

$$dayBal_{i,t+1} = dayBal^{*}_{i,t+1} + C_{i,t+1} \qquad \textbf{(A.1.1)}$$

In **(A.1.1)**, $dayBal^{*}_{i,t+1}$ is the balance process for bank *i* for time *t*+1, given that bank *m* is unable to make a payment at time *t*. The value of $C_{i,t+1}$ is the amount that would have made both process equal, that is, the impact of bank's *m* failure towards bank's *i* balance in *t*+1.
Expanding the last equation with expectations:

$$E\left(dayBal_{i,t+1}\right) = E\left(dayBal^{*}_{i,t+1}\right) + E\left(C_{i,t+1}\right)$$

$$E\left(dayBal_{i,0} + \sum_{j=1}^{t+1} Pay^{In}_{i,j} - \sum_{j=1}^{t+1} Pay^{Out}_{i,j}\right) = E\left(dayBal_{i,0} + \sum_{j=1}^{t+1} Pay^{In,*}_{i,j} - \sum_{j=1}^{t+1} Pay^{Out}_{i,j}\right) + E\left(C_{i,t+1}\right)$$

$$E\left(\sum_{j=1}^{t} Pay^{In}_{i,j} + Pay^{In}_{i,t+1}\right) = E\left(\sum_{j=1}^{t} Pay^{In}_{i,j} + Pay^{In,*}_{i,t+1}\right) + E\left(C_{i,t+1}\right)$$

$$E\left(C_{i,t+1}\right) = E\left(Pay^{In}_{i,t+1}\right) - E\left(Pay^{In,*}_{i,t+1}\right) \qquad \textbf{(A.1.2)}$$

The last result is intuitive as the difference between the scenarios is just the different amount of money bank *i* would have received in *t*+1. By expanding the last equation we get:

$$E\left(C_{i,t+1}\right) = E\left(\sum_{j=1, j\neq i}^{nBanks} \gamma_{i->j} Pay^{Out}_{j,t+1}\right) - E\left(\sum_{j=1, j\neq i, j\neq m}^{nBanks} \gamma_{i->j} Pay^{Out}_{j,t+1}\right) \qquad \textbf{(A.1.3)}$$

Now, realising the property that:

$$\sum_{j=1, j\neq i}^{nBanks} \gamma_{i->j} Pay^{Out}_{j,t+1} = \sum_{j=1, j\neq i, j\neq m}^{nBanks} \gamma_{i->j} Pay^{Out}_{j,t+1} + \gamma_{i->m} Pay^{Out}_{m,t+1}$$

we have the final result:

$$E\left(C_{i,t+1}\right) = \gamma_{i->m} E\left(Pay^{Out}_{m,t+1}\right) \qquad \textbf{(A.1.4)}$$

Appendix 2: Interpretation of coefficients in estimated payment rules

Our first regression model is:

$$Pay^{*,Out}_{i,d,t} = \alpha_{i} + \beta_{i} \cdot shortBal^{*}_{i,d,t} + \varepsilon_{i,d,t} \qquad \textbf{(A.2.1)}$$

Expanding equation **(A.2.1)**, we get:

$$\frac{Pay^{Out}_{i,d,t}}{A_{i,t}} = \alpha_i + \beta_i \cdot \left(\frac{\sum_{v=1}^{nLags} Pay^{In}_{i,d,t-v}}{B_{i,t} \cdot nLags} - \frac{\sum_{v=1}^{nLags} Pay^{Out}_{i,d,t-v}}{A_{i,t} \cdot nLags} \right) + \varepsilon_{i,d,t} \qquad \textbf{(A.2.2)}$$

In **(A.2.2)** the terms $A_{i,t}$ and $B_{i,t}$ are, respectively, the intraday means of payments sent and received. We can rewrite this as:

$$Pay^{Out}_{i,d,t} = A_{i,t}\alpha_i + \beta_i \cdot \left(\frac{A_{i,t}}{B_{i,t}} \frac{\sum_{v=1}^{nLags} Pay^{In}_{i,d,t-v}}{nLags} - \frac{\sum_{v=1}^{nLags} Pay^{Out}_{i,d,t-v}}{nLags} \right) + A_{i,t}\varepsilon_{i,d,t} \qquad \textbf{(A.2.3)}$$

So, the effect of a small change on the average past payment received, $\sum_{v=1}^{nLags} Pay^{In}_{i,d,t-v} / nLags$, on $Pay^{Out}_{i,d,t}$ is:

$$\frac{\partial Pay^{Out}_{i,d,t}}{\partial \left[\frac{\sum_{v=1}^{nLags} Pay^{In}_{i,d,t-v}}{nLags} \right]} = \beta_i \cdot \frac{A_{i,t}}{B_{i,t}} \qquad \textbf{(A.2.4)}$$

Similarly, the effect of a small change in the average payment made in the past is:

$$\frac{\partial Pay^{Out}_{i,d,t}}{\partial \left[\frac{\sum_{v=1}^{nLags} Pay^{Out}_{i,d,t-v}}{nLags} \right]} = -\beta_i \qquad \textbf{(A.2.5)}$$

For the interpretation of the overall effect of $shortBal^{*}_{i,d,t}$ in $Pay^{Out}_{i,d,t}$, one can simply assume that on expectation the ratio $\frac{A_{i,t}}{B_{i,t}}$ is equal to one, $E\left(\frac{A_{i,t}}{B_{i,t}}\right) = 1$. The intuition is that, on average, banks pay and receive the same total amount of money in different times of the day. This is not far from empirical estimates. With this simplification, we get:

$$\frac{\partial Pay^{Out}_{i,d,t}}{\partial shortBal_{i,d,t}} = \beta_i \qquad \textbf{(A.2.6)}$$

For the second regression model, the derivation is very similar, with the results:

$$\frac{\partial Pay_{i,d,t}^{Out}}{\partial\left[\frac{\sum_{v=1}^{nLags} Pay_{i,d,t-v}^{In}}{nLags}\right]} = \left(\beta_i^{(2)} + \phi_{i,2} Z_{i,d,t}\right) \cdot \frac{A_{i,t}}{B_{i,t}} \qquad \textbf{(A.2.7)}$$

And,

$$\frac{\partial Pay_{i,d,t}^{Out}}{\partial\left[\frac{\sum_{v=1}^{nLags} Pay_{i,d,t-v}^{Out}}{nLags}\right]} = \left(\beta_i^{(2)} + \phi_{i,2} Z_{i,d,t}\right) \qquad \textbf{(A.2.8)}$$

And, assuming $E\left(\frac{A_{i,t}}{B_{i,t}}\right) = 1$,

$$\frac{\partial Pay_{i,d,t}^{Out}}{\partial shortBal_{i,d,t}} = \left(\beta_i^{(2)} + \phi_{i,2} Z_{i,d,t}\right) \qquad \textbf{(A.2.9)}$$

$$\frac{\partial Pay_{i,d,t}^{Out}}{\partial Z_{i,d,t}} = A_{i,d,t}\phi_{i,1} + \phi_{i,2} shortBal_{i,d,t} \qquad \textbf{(A.2.10)}$$

Now, for the simulation of the models, we again use the simplification $E\left(\frac{A_{i,t}}{B_{i,t}}\right) = 1$. This yields:

$$Pay_{i,d,t}^{Out} = \alpha_i^* + \beta_i \cdot shortBal_{i,d,t} + \varepsilon_{i,d,t}^* \qquad \textbf{(A.2.11)}$$

Where:

$$\alpha_i^* = A_{i,t}\alpha_i$$
$$\varepsilon_{i,d,t}^* = A_{i,t}\varepsilon_{i,d,t}$$

Taking expectation of the intercept we get:

$$E\left(\alpha_i^*\right) = \alpha_i E\left(A_{i,t}\right)$$

Translating to our model, following the property $E\left(A_{i,t}\right) = E\left(nDays^{-1}\sum_{z=1}^{nDays} Pay_{i,z,t}^{Out}\right) = E\left(Pay_{i,z,t}^{Out}\right)$, last equation reduces to:

$$E\left(\alpha^{*}_{i}\right)=\alpha_i E\left(Pay^{Out}_{i,d,t}\right)$$

Therefore, the final equation that we use for the simulations is:

$$Pay^{Out}_{i,d,t}=\alpha_i E\left(Pay^{Out}_{i,d,t}\right)+\beta_{i}\cdot shortBal_{i,d,t}+\varepsilon^{*}_{i,d,t} \qquad \textbf{(A.2.12)}$$

Correspondingly, for the second model, following similar steps as Model 1, we have:

$$\begin{aligned} Pay^{Out}_{i,d,t} &= \alpha_i^{(2)^{**}}+\beta_i^{(2)}\cdot shortBal_{i,d,t}+\varepsilon^{*}_{i,d,t} \\ \alpha_i^{(2)^{**}} &= E\left(Pay^{Out}_{i,d,t}\right)\left(\alpha_i^{(2)}+\phi_{i,1}Z_{i,d,t}\right) \end{aligned} \qquad \textbf{(A.2.13)}$$

Appendix 3: How our simulation model could be written as a VAR

Defining:

$\underset{[m,n]}{X}$	- a matrix with *m* rows and *n* columns
$*$	- Matrix multiplication operator
I	- Identity matrix
α^{*}	- Scaled alpha vector
β^{*}	- Scaled beta vector

The process for payments in matrix notation, using *nLag=1*, is:

$$\begin{aligned} \underset{[1,nBank]}{Pay^{Out}_{t}} &= \underset{[1,nBank]}{\alpha^{*}}+\left(\underset{[1,nBank]}{Pay^{In}_{t-1}}-\underset{[1,nBank]}{Pay^{Out}_{t-1}}\right)*\underset{[nBank,nBank]}{\beta^{*}}+\underset{[1,nBank]}{\varepsilon_t} \\ \underset{[1,nBank]}{Pay^{In}_{t-1}} &= \underset{[1,nBank]}{Pay^{Out}_{t-1}}*\underset{[nBank,nBank]}{\lambda} \end{aligned} \qquad \textbf{(A.3.1)}$$

This implies that:

$$\begin{aligned} \underset{[1,nBank]}{Pay^{Out}_{t}} &= \underset{[1,nBank]}{\alpha^{*}}+\left(\underset{[1,nBank]}{Pay^{Out}_{t-1}}*\underset{[nBank,nBank]}{\lambda}-\underset{[1,nBank]}{Pay^{Out}_{t-1}}\right)*\underset{[nBank,nBank]}{\beta^{*}}+\underset{[1,nBank]}{\varepsilon_t} \\ \underset{[1,nBank]}{Pay^{Out}_{t}} &= \underset{[1,nBank]}{\alpha^{*}}+\left(\underset{[1,nBank]}{Pay^{Out}_{t-1}}*\left(\underset{[nBank,nBank]}{\lambda}-\underset{[nBank,nBank]}{I}\right)\right)*\underset{[nBank,nBank]}{\beta^{*}}+\underset{[1,nBank]}{\varepsilon_t} \end{aligned} \qquad \textbf{(A.3.2)}$$

Therefore, using standard VAR notation, the regression model can be specified as:

$$\underset{[1,n]}{Y_t} = \underset{[1,n]}{\alpha} + \underset{[1,n]}{Y_{t-1}} * \underset{[n,n]}{B^*} + \underset{[1,n]}{\varepsilon_t}$$

$$\underset{[n,n]}{B^*} = \left(\underset{[nBank,nBank]}{\lambda} - \underset{[nBank,nBank]}{I} \right) \underset{[n,n]}{B} \qquad \textbf{(A.3.3)}$$

References

Aikman, D, Alessandri, P, Eklund, B, Gai, P, Kapadia, S, Martin, E, Mora, N, Sterne, G and Willison, M (2009), 'Funding liquidity risk in a quantitative model of systemic stability', *Bank of England Working Paper no. 372*.

Bech, M L and Garratt, R (2003) 'The intraday liquidity management game', *Journal of Economic Theory*, Vol. 109(2), pages 198-219.

Becher, C, Galbiati, M and Tudela, M (2008), 'The timing and funding of CHAPS Sterling payments', *Federal Reserve Bank of New York Economic Policy Review,* Vol. 14(2), pages 113-33.

Becher, C, Millard, S and Soramäki, K (2009), 'The network topology of CHAPS Sterling', *Bank of England Working Paper no. 355*.

Beyeler, W, Glass, R, Bech, M and Soramäki, K (2007), 'Congestion and cascades in payment systems', *Physica A: Statistical Mechanics and its Applications*, Vol. 384(2), pages 613-718.

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

Greene, W (2003), *Econometric Analysis*, Prentice Hall, 5th edition.

Haldane, A (2009), *Rethinking the financial network*, speech delivered at the Financial Student Association in Amsterdam on 28 April.

Manning, M, Nier, E and Schanz, J (2009), *The economics of large-value payment and settlement systems: theory and policy issues for central banks*, Oxford University Press.

McAndrews, J and Rajan, S (2000), 'The timing and funding of Fedwire Funds transfers', *Federal Reserve Bank of New York Economic Policy Review*, Vol. 6(2), pages 17-32.

Merrouche, O and Schanz, J (2010), 'Banks' intraday liquidity management during operational outages: theory and evidence from the UK payment system', *Journal of Banking and Finance*, Vol. 34, pages 314-23.

Mills, D C and Nesmith, T D (2008), 'Risk and concentration in payment and securities settlement systems', *Journal of Monetary Economics*, Vol. 55(3), pages 542-53.

Renault, F, Bech, M L, Beyeler, W, Glass, R and Soramäki, K (2007), 'Congestion and cascades in coupled payment systems', *Federal Reserve Bank of New York Staff Papers no. 259*.

Soramäki, K, Bech, M L, Arnold, J, Glass, R J and Beyeler, W E (2007), 'The topology of interbank payment flows', *Physica A*, Vol. 379, pages 317-33.

van den End, J W (2008), 'Liquidity Stress-Tester: a macro model for stress-testing banks' liquidity risk', *DNB Working Paper no. 175*.

Wetherilt, A, Zimmerman, P and Sorämaki, K (2009), 'The sterling unsecured loan market during 2006-08: insights from network theory', *Bank of England Working Paper no. 398.*