BIS Working Papers
No 1318

Auction-based liquidity saving mechanisms

by Caner Ates, Morten Bech, Rodney Garratt and Marko Nanut Petrič

Monetary and Economic Department

December 2025

JEL classification: G21, C72, D44

Keywords: payment system, liquidity savings mechanism (LSM), real-time gross settlement (RTGS) system, auctions

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Auction-based liquidity saving mechanisms

Caner Ates[a] Morten Bech[a] Rodney Garratt[b]
Marko Nanut Petrič[a,c]

Abstract

This paper introduces a novel liquidity-saving mechanism (LSM) for real-time gross settlement (RTGS) systems based on an auction framework. Unlike conventional queue-based LSMs that optimise payment netting given pre-committed liquidity, the proposed mechanism reverses the process by asking participants whether they are willing to supply liquidity to facilitate the settlement of a given set of payments. These arrangements are supported through side payments from liquidity receivers to liquidity providers, determined via auctions. The mechanism is evaluated using agent-based simulations calibrated with transaction data from four RTGS systems: BCRP-RTGS (Peru), SAMOS and SADC-RTGS (South Africa), and SIC (Switzerland). Across all trials, auction success rates averaged 79% (ranging from 62 to 100%). These success rates were sufficiently high that observed failures typically did not increase liquidity usage. Overall, the findings demonstrate the importance of side payments in coordinating participant incentives, reveal the trade-off between liquidity efficiency and settlement delay, and highlight the potential of auction-based LSM designs to enhance RTGS system efficiency.

JEL classifications: G21, C72, D44

Keywords: Payment system, liquidity savings mechanism (LSM), real-time gross settlement (RTGS) system, auctions

1 Introduction and general overview

Over the past 40 years, large-value payment systems worldwide have transitioned to real-time gross settlement (RTGS) systems. In RTGS systems, payments are made on a gross basis individually and in real time, using the liquidity available to the payer at the moment of the transaction. This design ensures settlement finality and eliminates credit risk, but

[a]Bank for International Settlements, Basel, Switzerland[*]
[b]University of California, Santa Barbara, California, United States
[c]Swiss National Bank, Zürich, Switzerland[†]

[*]The views expressed are those of the authors and do not necessarily reflect the official positions of the Bank for International Settlements.

[†]The views, opinions, findings, and conclusions or recommendations expressed in this paper are strictly those of the authors. They do not necessarily reflect the views of the Swiss National Bank (SNB). The SNB takes no responsibility for any errors or omissions in, or for the correctness of, the information contained in this paper.

it also entails substantial liquidity demands. In many advanced economies, the total value of payments processed through RTGS systems can exceed annual GDP within a single week. To address these considerable liquidity demands, central banks have introduced liquidity-saving mechanisms (LSMs) to enhance the efficiency of settlement processes. (Norman (2010)).

A variety of LSM types have been developed to reduce liquidity usage in RTGS systems. The simplest forms involve operational rules that promote the timely submission of payments or fee structures that incentivise early settlement. However, the most direct approach to liquidity saving is through the use of "netting queues", in which payments are temporarily queued and netted either bilaterally or multilaterally against each other. Queue-based LSMs operate alongside many RTGS systems worldwide, including those in the United Kingdom (CHAPS), the Eurosystem (TARGET2), Japan (BOJ-NET), and Switzerland (SIC).

Queue-based LSMs aim to net or offset non-urgent payments submitted to a central queue, either continuously or at predetermined time intervals. In most implementations, participants reserve an amount of liquidity in advance for the LSM to use. Ideally, each participant's net debit position within a netting cycle is smaller than the liquidity it allocates to the queue, allowing all queued payments to be settled. However, this condition is not always met, and the LSM must then identify a subset of queued payments that can be settled with the reserved liquidity. Determining this optimal subset constitutes a constrained integer programming problem which, for large systems, is NP-hard. As a result, many LSMs use heuristic rules to winnow down or reorder the set of queued payments. These heuristic rules, while computationally efficient, are generally suboptimal, as they may not capture participants' true preferences over which payments to net or offset, nor the trade-offs they face between settlement priority and liquidity cost.

This paper presents a new concept for an LSM, building on the framework proposed by Garratt (2021). The mechanism extends the traditional centralised netting queue by inverting its logic. Instead of seeking the optimal set of payments to net, given a fixed amount of reserved liquidity, participants are asked whether they are willing to supply liquidity to facilitate the settlement of a given set of payments. To incentivise participants to do so, the new approach introduces "side payments". A side payment is a (small) reward from liquidity receivers to liquidity providers to compensate them for the costs of providing liquidity.[1] In essence, participants have the option to pay for the earlier receipt of their own funds.

The theoretical analysis presented in Garratt (2021) assumes that the LSM operator has complete information about both the benefits participants derive from settling payments and their liquidity costs. Under this strong informational assumption, it is possible to apply the Shapley value cost allocation method to achieve an optimal outcome. In practice, however, the operator does not know all the information required to ensure the settlement of the optimal subset of queued payments. The LSM developed and empirically evaluated in this paper addresses this informational constraint by eliciting participants' private valuations through an auction mechanism.[2]

[1] A side payment should not be interpreted as a bribe. This term is used in its economic sense: a transfer of resources from one party to another to incentivise or compensate for an agreement or action outside the primary terms of a deal or contract.

[2] The auction-based LSM was built and tested as part of the BIS Innovation Hub's Project Titus. For

In a live implementation of an auction-based LSM, side payments would be determined by participants' bids. They would be informed of the set of payments proposed for net settlement and the corresponding liquidity requirements, after which they would submit bids reflecting their willingness to provide or receive liquidity. A participant that ends the netting cycle in a net credit position—that is, with more incoming than outgoing value—would typically submit a positive bid, offering to pay for the liquidity benefits it receives. Conversely, a participant that ends in a net debit position would submit a negative bid, effectively requesting compensation for supplying liquidity. In summary, liquidity receivers offer compensation through positive bids, while liquidity providers seek compensation through negative bids. An auction is successful if the sum of all bids is positive.

As it was not feasible to conduct live auctions with actual treasury managers, the proposed auction-based LSM was evaluated using agent-based modelling (ABM). ABM is a simulation approach in which individual agents—here, representing system participants—follow specified behavioural rules that govern their actions within a complex system. By simulating agent interactions, the model generates system-level outcomes that emerge from individual decision-making processes. The simulations can be repeated using different data sets and preference parameters.

To evaluate the auction-based LSM, payments data from four payment systems were used: the Central Reserve Bank of Peru's (BCRP-RTGS), South Africa's domestic South African Multiple Option Settlement (SAMOS) system, the cross-border Southern African Development Community RTGS (SADC-RTGS), and Switzerland's Swiss Interbank Clearing (SIC) system.[3] For each system, the payment day was divided into 10-minute intervals and an auction was conducted for the payments submitted during each interval. If an auction was successful, the corresponding payments were settled. If it was unsuccessful, two alternative procedures were considered: settling the payments on a gross basis through the conventional RTGS stream or *pushing* them to the next interval for inclusion in the subsequent auction. The simulations were repeated using different payment data sets (payments files) and bidder preference parameters as inputs, producing corresponding bid patterns.

Pushing payments to the next interval introduces additional settlement delays but also increases the potential liquidity savings. For each payments file, the realised liquidity savings and the average delay per dollar settled are computed. In cases where payments are pushed, the incremental delay required to achieve the liquidity savings from netting all payments is also reported. Using a model of participant preferences, the cost of these additional delays is quantified and compared to the liquidity benefits obtained. Interestingly, in some systems (eg SAMOS) it is often optimal to push payments from failed auctions forward, whereas in others (eg SIC) it is not. Thus, systems considering whether to push payments can make an informed decision by weighing the social cost of settlement delays against the reduction in liquidity costs.

The auction-based LSM performs remarkably well. In simulations, auction success rates range from 62 to 100%, depending on the system, date, and data-filtering parameters (discussed below). In most instances where an auction fails, the subsequent one succeeds,

details, see https://www.bis.org/about/bisih/topics/fmis/titus.htm.

[3]These data sets were used solely to assess the feasibility of the side payment concept and should not be interpreted as a recommendation for adoption in any of the respective jurisdictions.

allowing the mechanism to capture nearly all achievable liquidity savings from payment netting.[4]

Payment system participants would not want to submit all of their payments to the LSM. It only makes sense to include payments in the queue if they expect others to submit offsetting payments. To approximate such "intelligent" queue management, the analysis is repeated using filtered data that include only payments between system participants that are included in payment cycles within each 10-minute interval. These cycles are identified using Johnson's cycle algorithm (Johnson (1977)).[5] When considering only payments from system participants that belong to cycles, liquidity savings relative to conventional RTGS processing range from 2.6 to 12.2%. This approach likely overstates the potential netting benefits, as participants do not have perfect foresight about which counterparties will be included in a payment cycle in any given interval. Nevertheless, it provides an upper bound on achievable liquidity savings through netting. The auction-based LSM is therefore evaluated using both raw ("unfiltered") data and data reflecting intelligent queue management ("filtered" data).

The remainder of the paper is structured as follows. Section 2 reviews liquidity-saving mechanisms, outlining their benefits for payment systems and the main challenges that remain. Section 3 describes the auction-based LSM design, the simulation approach used to construct participant bid functions, and the underlying data sets. Section 4 presents the simulation results for the auction-based LSM, together with statistics on potential and realised liquidity savings. The paper concludes with a summary of the key findings and broader implications of the study.

2 Background on liquidity-saving mechanisms

Before electronic ledgers and systems made finance faster and more efficient, payment systems were generally "deferred net settlement". Paper instruments such as receipts, trade slips, order forms, drafts, cheques, and promissory notes were collected throughout the day. At the end of each business day, the difference between dues to and dues from for each participant was calculated, and the corresponding net amount was debited or credited the following day. These systems were designed with efficiency and security in mind. In many jurisdictions, central banks established and managed such clearing systems, acting as both the settlement agent and the lender of last resort.(Norman et al. (2011))

Even after the transition to electronic processing removed the operational constraints associated with paper-based systems, many central bank payment systems continued to rely on deferred settlement.[6] This persisted until the risks inherent in deferred settlement — and the potential for systemic gridlock arising from a loss of confidence — became evident following the failure of Herstatt Bank in 1974 (Galati (2002)).

Advances in technology and the ever-increasing payment values during the 1980s prompted

[4]The magnitude of potential liquidity savings varies across systems: when all payments within a 10-minute interval are netted, savings relative to conventional RTGS processing range from 0.3 to 4.8%. The relatively modest savings reflect the characteristics of the underlying data, rather than any limitations of the auction-based LSM itself.

[5]See Appendix A for an explanation of Johnson's cycle algorithm.

[6]A notable exception was the Fedwire Funds Service, established in 1918 as a wire communications network for settlement among Federal Reserve Banks.

significant changes in central bank payment systems (Bech and Hobijn (2006); Bech et al. (2017)). RTGS systems were introduced, enabling payments to be settled individually and in real time, thereby eliminating settlement risk.

Eliminating settlement risk came at a cost. In RTGS systems, the absence of deferral or netting means participants must maintain sufficient liquidity to cover the full value of each payment they intend to make. For many institutions, these liquidity demands are significant. To mitigate this burden, numerous central banks introduced *intraday liquidity facilities*, which provide same-day credit to participants. While these facilities enable timely settlement, they also expose central banks to credit risk. Consequently, overdrafts are typically collateralised or subject to fees. Moreover, the design of intraday credit arrangements can encourage strategic behaviour: participants may delay payments (to avoid the opportunity cost of collateral or fees), and wait for others to settle first to reduce their own liquidity needs. Such delays, however, can hinder the smooth functioning of the payment system (Bech and Garratt (2003), Bech and Garratt (2012) and Bech (2008)).

LSMs help address many of these challenges. Although their designs vary, most allow participants to queue "non-urgent" payments and settle only the resulting net amounts. This approach introduces short settlement delays without re-introducing credit risk. It also reduces the overall liquidity required for participation, encourages prompt payment submissions, and preserves full RTGS functionality for transactions that require immediate settlement.

Both participants and the broader system benefit from LSMs. However, liquidity reserved for net settlement must still be set aside by participants, which entails a cost. When a netted payment exceeds available liquidity, the LSM must remove or postpone certain payments from the queue. Determining which payments to exclude is computationally complex and may yield sub-optimal outcomes in the absence of information about participants' payment priorities and preferences (Garratt (2021)).

During the Great Financial Crisis and its aftermath, central bank policy responses flooded the financial system with reserves, effectively eliminating the need for intraday liquidity to settle payments. Settlement times accordingly shifted earlier (Bech et al. (2012)). However, these favourable conditions from a narrow payment system perspective are unlikely to persist (Carstens (2022)). As monetary conditions tighten, the demand for intraday credit and liquidity-saving mechanisms (LSMs) is expected to re-emerge (Benos and Harper (2016)).

At the same time, emerging tokenised payment systems—such as stablecoins—are reintroducing elements of prefunding and gross settlement, both of which impose significant liquidity demands reminiscent of the early days of RTGS systems. Furthermore, ongoing modernisation initiatives are broadening RTGS access beyond traditional banks (CPMI (2021)). Non-bank participants, however, may not be eligible for intraday credit or may lack sufficient collateral, constraining their ability to settle payments promptly.

In such environments, liquidity scarcity can encourage competition rather than cooperation among participants, ultimately undermining system efficiency (Bech (2008)). As liquidity management once again rises on policymakers' agendas, innovative approaches—such as auction-based LSMs—may play an increasingly important role in addressing these emerging challenges.

3 Overview of the auction-based LSM

The auction-based LSM was designed for use in combination with an RTGS payment system. The set of payments that enter the netting queue is taken as given and auctions are used to find a mutually agreeable way to share the liquidity burden of netting these payments. Auctions are a means to elicit private information about system participants' benefits and costs.

3.1 Auction prototype

The number of bidders in each auction equals the number of system participants with either an outgoing or incoming payment in the payment queue that exists at the time when the auction is conducted. Given bidders' preferences for processing payments and their liquidity costs, bidders may have a positive or negative reservation price related to settling the payments in the queue. After computing the net liquidity obligations needed to settle all payments in a given queue, bidders with a net debit position may have a negative reservation price, as these participants will be required to provide liquidity. A bidder with a negative reservation price requires compensation to provide the liquidity needed to settle the payments in the queue. In contrast, bidders with a net credit position will certainly have a positive reservation price as they settle their payments in the queue and receive liquidity.

An interesting aspect of these auctions is that the "good" that is auctioned off is not transferred from the seller to a buyer, as in a typical auction. Rather, the objective is to implement a netting solution that benefits everyone in different ways. Variations in the benefits of a netting solution stem from the fact that each participant has distinct preferences for settling its customer payments and from the fact that, while some are obligated to provide liquidity, others are not.

What is at stake in these auctions is the provision of liquidity. The auctions determine whether those in a net debit position are willing to provide the liquidity required to settle all the payments in the queue. Their willingness to do this will depend on the urgency of the queued payments and their current liquidity position. System participants that have experienced larger outflows than inflows up to the current point in time and, hence, perceive themselves to be short of liquidity may be reluctant to provide more, while participants that have a strong liquidity position may be willing to provide liquidity cheaply. The auction is successful if the aggregate compensation required by *liquidity providers* is less than the aggregate amount that the beneficiaries of this liquidity (participants in a net credit position, ie *liquidity takers*) are willing to pay.

Unlike the conventional liquidity-saving mechanisms discussed above, the auction-based LSM elicits information from participants about their preferences for making payments and the associated costs of providing liquidity. Participant bids reflect their minimum willingness to pay/receive in exchange for providing or receiving liquidity and settling their payments at the current opportunity.

3.2 Liquidity-saving mechanism design

The auction-based LSM was implemented for our trial to run on a 10-minute "netting cycle". Each cycle follows the same process (Graph 1). The auction-based LSM receives payments from participants and stores them in a database queue, similar to a traditional LSM. The auction-based LSM then computes the multilateral net positions of all queued payments. A message is broadcast to the participants showing (i) the list of payments that they will settle in the queue and (ii) the net liquidity position between sent and received payments. If the value of sent payments exceeds the value of received payments, then the participant has a (positive) net debit position. Participants evaluate the proposal according to their preferences and liquidity costs. Each participant responds with a bid. The LSM evaluates the bids, and if the sum of payments is greater than zero (ie the total amount requested is less than the total amount offered), it calculates side payments based on a pre-determined allocation rule. It sends them, along with the underlying interbank payments, for settlement.

Graph 1: Stylised netting cycle

The specific rules for determining side payments following a successful auction in each netting cycle s are as follows: given a vector of bids $(B_i^s)_i$ with $\sum_i B_i^s \geq 0$ the transfer amount is

$$X = \frac{\sum_{i:B_i^s \geq 0} B_i^s - \sum_{i:B_i^s<0} B_i^s}{2}. \tag{1}$$

Bidder i pays $X\frac{B_i^s}{\sum_{j:B_j^s \geq 0} B_j^s}$ if $B_i^s \geq 0$ and $X\frac{B_i^s}{\sum_{j:B_j^s < 0} B_j^s}$ if $B_i^s < 0$. Note that these side payments are feasible given that $\sum_i B_i^s \geq 0$ and individually rational since each paying bidder pays less than their bid and each receiving bidder receives more than the absolute value of their negative bid (ie more than they asked for).

3.3 Simulated agents and auctions

The auction-based LSM is novel because it uses an auction mechanism to elicit participants' preferences in the payment system. Ideally, the concept would have been tested with real participants; however, this was not feasible. Instead, participants were simulated. A simulation requires a model that represents key features of the environment and a description of all necessary actions by the system's participants. Specifically, to test the auction-based LSM it was necessary to determine participants' reservation prices and then estimate a function that maps these reservation prices to bids in an auction (ie a bid function). An ABM model was developed to simulate the behaviour of commercial bank treasury managers in virtual auctions.

3.3.1 Auction-based LSM participants' utility

A participant's utility was determined by the payments made and any change in their settlement balances (Table 1).[7] Reducing a participant's settlement balance is costly because it reduces their ability to make future payments without obtaining additional liquidity from alternative sources, whereas increasing a participant's settlement balance

[7]The utility functions used in this model generalise those specified in (Garratt (2021)).

is a benefit because it gives the participant more flexibility in terms of the timing of future payments. Costs are initially assumed to be the same across participants (ie they all have the same benefits and costs for increasing/reducing their balances). However, as the payment day proceeds, these costs and benefits are adjusted to reflect changes in each participant's liquidity position (ie each participant's current settlement balance) relative to their initial settlement balance. The adjustment is defined so that a participant incurs more cost/benefit from providing/receiving liquidity when their settlement balance falls below their starting balance, and less cost/benefit from providing/receiving liquidity when their settlement balance rises above their starting balance.

Table 1: Simulated participant benefits and costs

	Benefits	**Costs**
Payments settled	Benefit parameters were drawn independently for each participant at the start of the day from a uniform distribution [.5/100,000,1/100,000] and kept the same throughout the day.	N/A
Liquidity	Liquidity received results in a benefit that increases as the participants' current settlement balance decreases relative to their starting settlement balance.	Liquidity provided results in a cost that increases as the participants' current settlement balance decreases relative to their starting settlement balance.

The costs and benefits were chosen so that the upper bound on the per-currency unit benefit of settling payments is below the per-currency unit cost.[8] This captures the idea that payments submitted to a queue are likely ones the participant would not have been willing to settle if they were required to provide the full (gross) amount of liquidity. At the start of each settlement day, participants are assigned a randomly determined, individual, per-currency unit benefit value for settling their payments (see Table 1). This assumption implies that the system participant's perceived benefit from settling each currency unit of each payment it places in the queue is the same, which is a simplification of reality. This assumption, combined with the fact that there is significant variation in the values of queued payments across participants in each netting cycle, still provides us with sufficient variation in reservation prices across participants to test the mechanism.

The formal description of how participants' costs and benefits in settling queued payments translate into reservation prices is as follows. Let $\left(p_{ij}^s\right)_{ij}$ denote the matrix of queued payments after netting cycle s, where p_{ij}^s denotes a payment from participant i to participant j, and let $d_i^s = \sum_{j \neq i} p_{ij}^s - \sum_{j \neq i} p_{ji}^s$ denote each participant i's net position that results

[8]Values in the experiment were denominated in the currency in which payments were denominated (ie Peruvian sol (PEN), South African rand (ZAR), Swiss franc (CHF) and Brazilian real (BRL)).

from netting all payments in the queue. Furthermore, let $q^s_{i,out}$ denote the total value of outgoing payments for participant i, which have a per unit benefit b_i. Let L^s_i denote participant i's liquidity position at cycle s. The cost/benefit of liquidity in each netting cycle is determined by the weight cw^s_i, where c is a constant scaling parameter and

$$w^s_i = 0.5 + \frac{1}{1 + L^s_i / L^0_i}.$$

Participant i's reservation price is given by the formula

$$r^s_i = b_i q^s_{i,out} - cw^s_i d^s_i.$$

3.3.2 Bid functions

Given their reservation prices, live participants would be able to make bids that would determine the outcomes of the auction-based LSM auctions. As mentioned above, the auction-based LSM was not tested with live bidders. Moreover, no preexisting theory can be applied to determine equilibrium bid functions in the general multilateral netting case. Consequently, an ABM was used to estimate bid functions for all participants in each payment system (see Graph 2).

Bid functions were determined by iterating a process that generates a vector of successful bids. The day was divided into 10-minute netting cycles. At the end of each netting cycle, the netting proposal is taken and each participant's reservation price is computed using the process described in subsubsection 3.3.1.

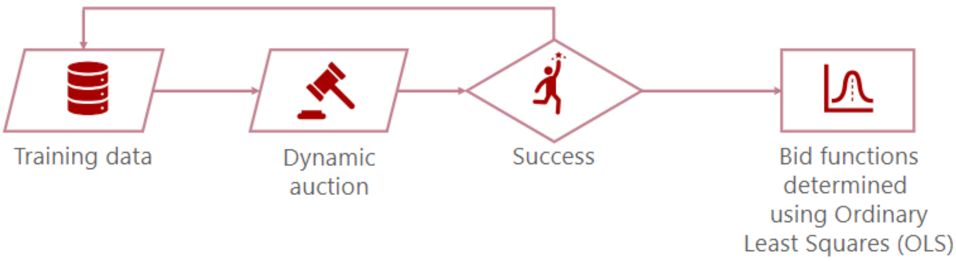

Graph 2: Establishing bid functions for auctions

Dynamic auctions were then conducted using the following process to generate bids. All bidders with a net credit position (ie those who are buyers) submit an initial bid equal to half their reservation price. All bidders with a net debit position (ie those who are sellers) submit an initial bid equal to the full amount of their liquidity cost (ie the amount of their net debit position d^s_i times their cost parameter cw^s_i). Then buyers/sellers adjust their bids upwards/downwards according to a specified random adjustment process until there is an agreement.[9] The agreement identifies final bids for each participant involved in the netting arrangement. These bids and their associated reservation prices are used to estimate bid functions for each participant using ordinary least squares.

Real payment data from four payment systems – BCRP-RTGS, SAMOS, SADC-RTGS, and SIC – were used to estimate the bid functions. Graph 3 and 4 illustrate the fit of the regressions used to determine the bid functions for participants in each system.

[9]Bid increments are determined by draws from a folded normal distribution with mean zero and standard deviation equal to $a \cdot r_i$ for buyers and $a \cdot (cw^s_j d_j + r_j)$ for the sellers, where a is an exogenously specified parameter that determines the speed of adjustment.

As mentioned above, it is impossible to check whether the estimated bid functions match those that participants would use in an equilibrium of the auction-based LSM auction. Some assurance is provided from the fact that the ABM procedure works almost perfectly in a simplified version of the problem with only two participants. Appendix B demonstrates that our procedure generates equilibrium bid functions in a setting with two participants and non-urgent payments.

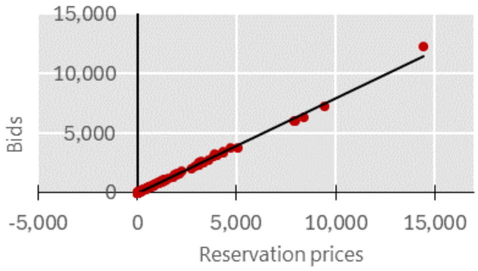

(a) BCRP-RTGS day 1: buyer bids

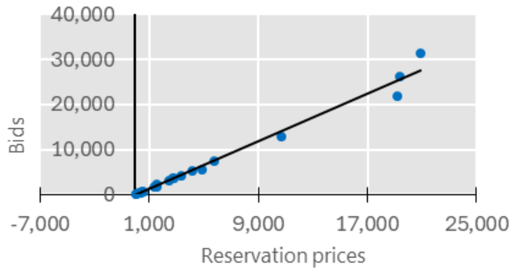

(b) BCRP-RTGS day 1: seller bids

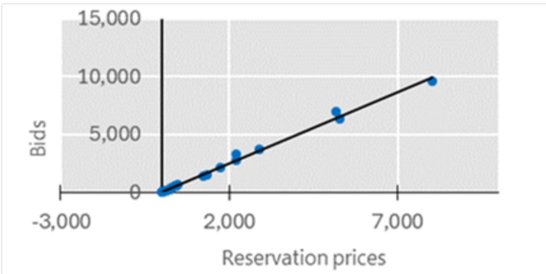

(c) BCRP-RTGS day 2: buyer bids

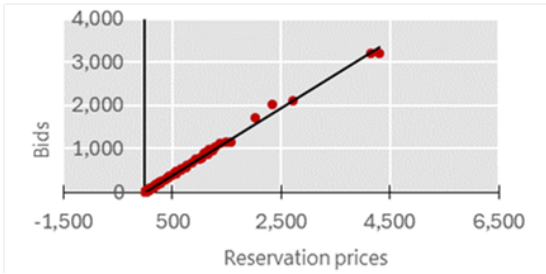

(d) BCRP-RTGS day 2: seller bids

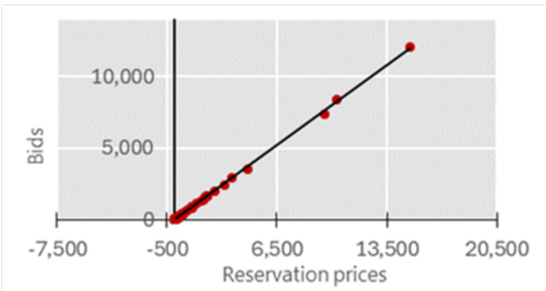

(e) SADC–RTGS day 1: buyer bids

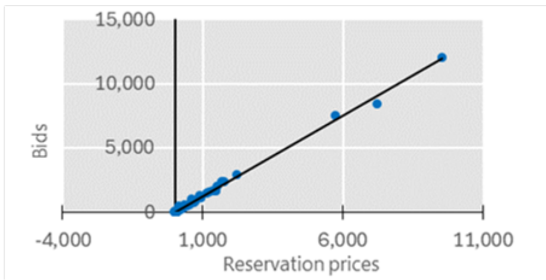

(f) SADC–RTGS day 1: seller bids

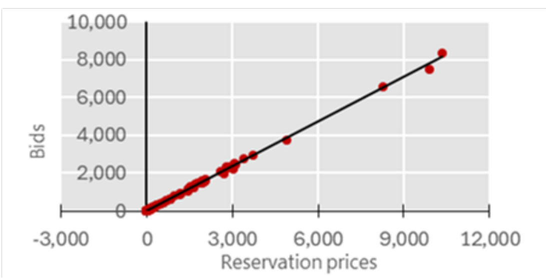

(g) SADC-RTGS day 2: buyer bids

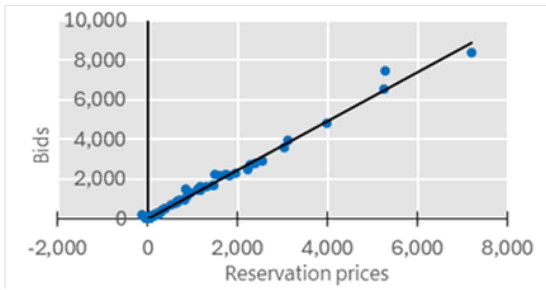

(h) SADC-RTGS day 2: seller bids

Graph 3: Bids relative to reservation prices. Each dot represents a final bid placed by a buyer (left-hand panel) or seller (right-hand panel) in a simulated auction. Bids are used to estimate the bid functions of buyers and sellers in the associated regression lines.

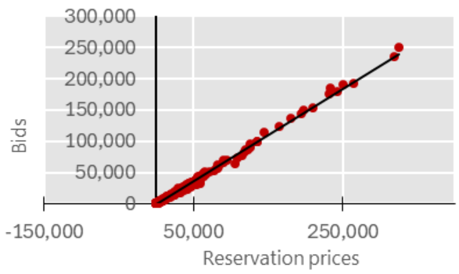

(a) SAMOS-RTGS day 1: buyer bids

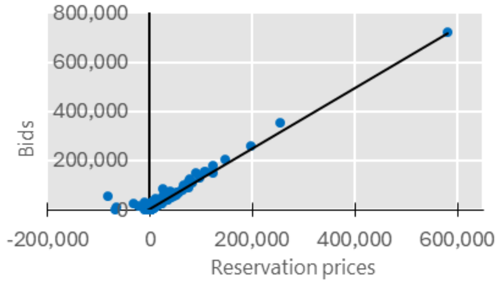

(b) SAMOS-RTGS day 1: seller bids

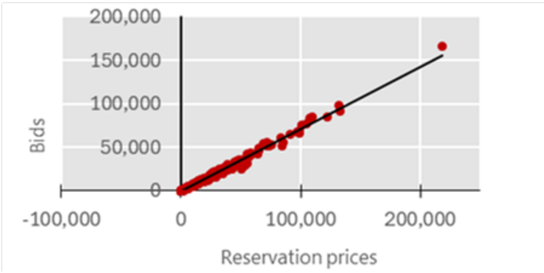

(c) SAMOS-RTGS day 2: buyer bids

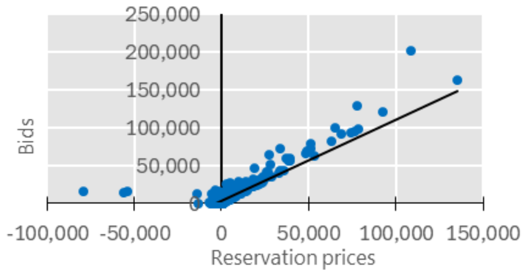

(d) SAMOS-RTGS day 2: seller bids

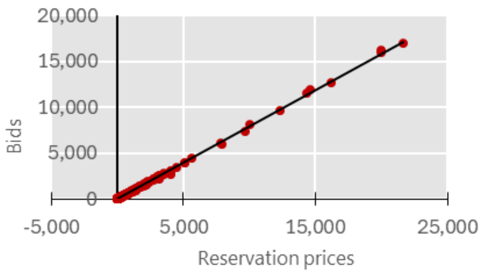

(e) SIC-RTGS day 1: buyer bids

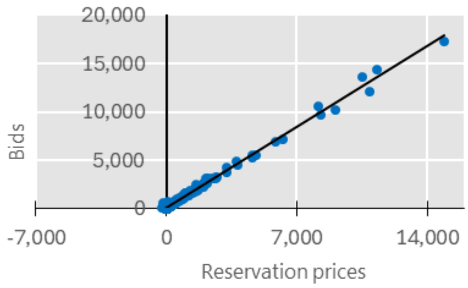

(f) SIC–RTGS day 1: seller bids

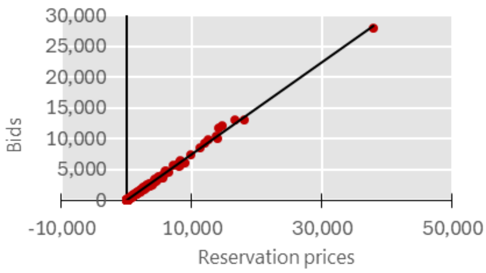

(g) SIC-RTGS day 2: buyer bids

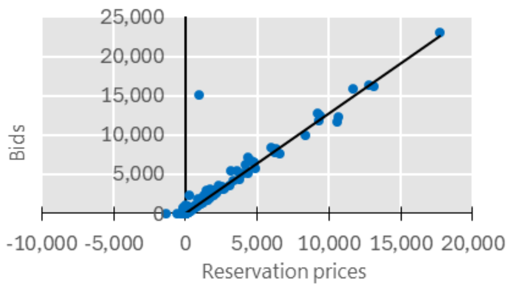

(h) SIC-RTGS day 2: seller bids

Graph 4: Bids relative to reservation prices. Each dot represents a final bid placed by a buyer (left-hand panel) or seller (right-hand panel) in a simulated auction. Bids are used to estimate the bid functions of buyers and sellers in the associated regression lines.

3.4 Payments data used for testing

Two days of anonymised data from multiple payment systems was used to train simulations and test the LSM (Table 2). Initial settlement balances were not included in the data. Instead, all participants were given a sufficiently high balance to conduct all settlements for the given day.

Table 2: Summary of key payment systems parameters for 2021; values and volumes are annual aggregates.

System	Value (USD bn)	Volume (mn)	Participants	Year founded
BCRP-RTGS	1050	1.2	53	2000
SAMOS	11 000	10.8	33	1998
SADC-RTGS	89	0.37	84	2013
SIC	45 700	893.4	319	1987

These systems differ in scale and scope. BCRP-RTGS handles wholesale payments between participants and has the smallest daily value and volume of the systems analysed. Switzerland's SIC system is at the other end of the spectrum. SIC handles retail and wholesale payments, and has the most significant daily value and volume. SAMOS is the domestic payment system for South Africa, and SADC-RTGS is a cross-border payment system for the countries in the Southern African Development Community. These systems provide intermediate cases in terms of values and volumes.

The four systems have different operating hours and different peak usage times. The BCRP-RTGS, SADC-RTGS and SIC systems have gridlock resolution mechanisms whereby payments are automatically queued if there is insufficient liquidity and then settled on a net basis. The data from the central banks include all payments processed on the selected days, regardless of whether they were ultimately queued in the home system.[10]

The systems settle in their domestic currencies. SADC and SAMOS use ZAR as the settlement currency. SIC uses CHF as the settlement currency. The BCRP-RTGS uses PEN as the main settlement currency. There is also a USD funding component to the BCRP-RTGS, however our data exclude USD-denominated payments.

4 Simulation results

The presentation of the simulation results begins by reporting the frequency with which the auction mechanism is successful. This evaluation is conducted under four scenarios reflecting the two-by-two design involving pushed versus non-pushed payments and unfiltered versus filtered data (see Table 3).

Table 3: Summary of payment types and their settlement modes.

	Pushed on failure	**Non-pushed on failure (settled using RTGS)**
Unfiltered (All payments)	Unfiltered pushed	Unfiltered non-pushed
Filtered (Johnson's algorithm)	Filtered pushed	Filtered non-pushed

Results for two versions of the data sets are presented. The unfiltered sample includes

[10] A handful of extremely large payments were deleted from each system as these would not be submitted to an LSM payments queue.

all the payments provided in the data file, except a handful of very large payments that were deleted. The filtered sample uses payments from each 10-minute time interval that are part of a cycle. Examination of the two cases likely provides reasonable benchmarks for evaluating the effectiveness of the auction-based LSM for different systems.

4.1 Auction success

An auction is successful if the sum of participant bids is non-negative. This ensures that side payments can be made and that everyone is better off if the payments in the queue are settled on a net basis and those in a net debit position provide liquidity equal to that position. If an auction is unsuccessful, the payments in the queue either roll over into the next queue and are included in the next auction (pushed) or are settled via RTGS (non-pushed). In all auction-based LSM simulations, all payments are cleared by the end of the day. In particular, if the last auction of the day fails, then all the payments in the final queue are settled on a gross basis.

Tables 4 and 7 show the probability of success of the auctions for the four cases shown in Table 3. For comparison purposes, results with and without side payments are provided. In the case of no side payments, the auction succeeds if all participants who realise a positive net debit position when the payments in the queue are netted are willing to provide that liquidity without compensation.

Table 4: Auction success probabilities (Unfiltered, non-pushed, b U[0.5,1], c=2, buyer start = $0.5r_i$)

	BCRP-RTGS		SAMOS		SADC-RTGS		SIC	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
Max number of auctions	53	52	60	59	64	63	104	108
Auction participants: Count								
Maximum	20	19	25	26	23	22	119	121
Minimum	3	4	2	2	2	2	10	7
Average	10	9	20	19	12	11	42	45
Success probability (%)								
Without side payments	1.9	1.9	1.7	1.7	1.6	1.6	1.0	0.9
With side payments	100.0	100.0	73.3	74.6	71.9	76.2	67.3	62.0

Table 5: Auction success probabilities (Unfiltered, pushed, b U[0.5,1], c=2, buyer start = $0.5r_i$)

	BCRP-RTGS		SAMOS		SADC-RTGS		SIC	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
Max number of auctions	53	52	60	59	64	63	104	108
Auction participants: Count								
Maximum	20	19	26	28	27	22	144	156
Minimum	3	4	13	2	2	2	11	8
Average	9	9	19	18	10	10	51	55
Success probability (%)								
Without side payments	1.9	1.9	1.7	1.7	1.6	1.6	1.0	0.9
With side payments	100.0	100.0	76.3	76.6	70.3	79.4	66.3	63.9

Table 6: Auction success probabilities (filtered, non-pushed, b U[0.5,1], c=2, buyer start = $0.5r_i$)

	BCRP-RTGS		**SAMOS**		**SADC-RTGS**		**SIC**	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
Max number of auctions	6	3	56	55	22	26	95	99
Auction participants: Count								
Maximum	2	2	17	16	5	7	49	61
Minimum	2	2	2	3	2	2	2	2
Average	2	2	10	10	10	10	14	14
Success probability (%)								
Without side payments	1.9	1.9	1.7	1.7	1.6	1.6	1.0	0.9
With side payments	100.0	100.0	83.9	87.3	72.7	65.4	73.7	71.7

Table 7: Auction success probabilities (filtered, pushed, b U[0.5,1], c=2, buyer start = $0.5r_i$)

	BCRP-RTGS		**SAMOS**		**SADC-RTGS**		**SIC**	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
Max number of auctions	6	3	56	55	22	26	95	99
Auction participants: Count								
Maximum	2	2	18	16	7	9	62	67
Minimum	2	2	3	4	2	2	2	2
Average	2	2	10	10	3	3	16	18
Success probability (%)								
Without side payments	1.9	1.9	1.7	1.7	1.6	1.6	1.0	0.9
With side payments	100.0	100.0	82.1	89.1	72.7	65.4	74.7	69.7

Tables 4 and 7 show that side payments are essential to encouraging liquidity provision. Without side payments, netting occurs only if the individual benefit to each liquidity provider from netting their queued payments exceeds their cost of providing liquidity. This is a rare occurrence in the auction-based LSM setup in which non-urgent payments may have very low benefits relative to the cost of liquidity provision for immediate settlement. The probability of success with side payments is consistently high across all systems and netting cycles.

The two systems with the fewest participants (BCRP-RTGS and SADC-RTGS) tend to have higher success probabilities. This may be expected, as one might have thought that it would be harder to reach an agreement with many participants. However, the auction-based LSM design reallocates individual willingness to pay through its mapping, from auction bids to payments, in a way that smooths out individual differences and reaches agreement if there is consensus on aggregate amounts offered and requested.

4.2 Liquidity savings

A common measure of the success of a liquidity-saving mechanism is the amount of liquidity saved relative to the amount required under pure RTGS. Norman (2010) summarises the benefits of specific netting queues implemented in the past and reports savings of 20% for the Bank of Korea's BOK-Wire+ payment system and 15% for the Japanese BOJ-Net system. Davey and Grays (2014) and Seaward (2016) estimate that liquidity savings from the LSM introduced to CHAPS in April 2013 amounted to around 20%. However, this number fell to near zero after the system became flush with liquidity due to quantitative easing policies.

The realised amount of liquidity savings depends on the success of the mechanism and the potential for netting within the payments data. If there are a few offsetting payments or if the sequencing of payments allows a high degree of liquidity recycling, then liquidity savings from a netting queue may be low even if the LSM functions perfectly. To isolate these two effects (auction success rates and potential liquidity savings from netting) a measure of potential liquidity savings is computed that assumes the auctions are successful 100% of the time. This measure of liquidity saving captures the netting benefits inherent in the data over 10-minute intervals, entirely independent of the mechanism's success.

4.2.1 Potential liquidity savings from netting in 10-minute intervals

The potential liquidity savings is calculated from netting in 10-minute intervals by determining the maximum net debit position for each participant that arises if the payments included in each queue are settled on a net basis. No auctions are involved in the computation of this benchmark. Payments are netted such that, in each period, participating participants have a net debit position. Liquidity is allocated accordingly, and the maximum net debit position for each participant is identified at the end of the day. The sum of these maximum net debit positions represents the liquidity requirements under netting, which can then be compared with the liquidity requirement under RTGS to determine the potential liquidity savings from netting payments in 10-minute intervals, assuming no pushed payments.

Graph 5 shows that the potential liquidity savings from netting range from 0.3 to 4.8% for the unfiltered data and from 2.6 to 12.2% for the filtered data.

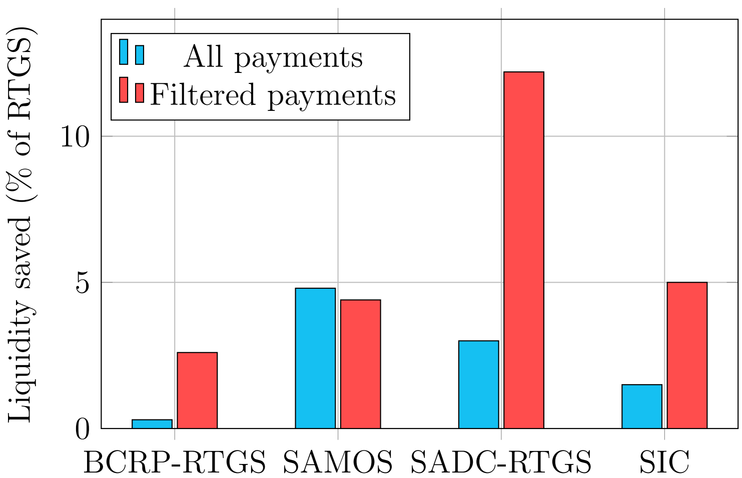

Graph 5: Potential liquidity savings from netting payments over 10-minute intervals

4.2.2 Realised savings from the auction-based LSM compared with the potential savings based on 10-minute netting intervals

The left-hand panel of Graph 6 shows the percentage of potential liquidity savings captured by the auction-based LSM in the case of non-pushed payments. In all but two cases—the unfiltered case for SAMOS and the filtered case for SIC—the auction-based

LSM achieved 100% of the potential liquidity savings from netting over 10-minute intervals. This is surprising as the auction success rates were not 100%, except in the case of Peru. However, this shows that the success rates were high enough, except in the two cases identified in the graph, so that a few missed netting opportunities did not drive participants' maximum net debit positions below the levels associated with perfect success.

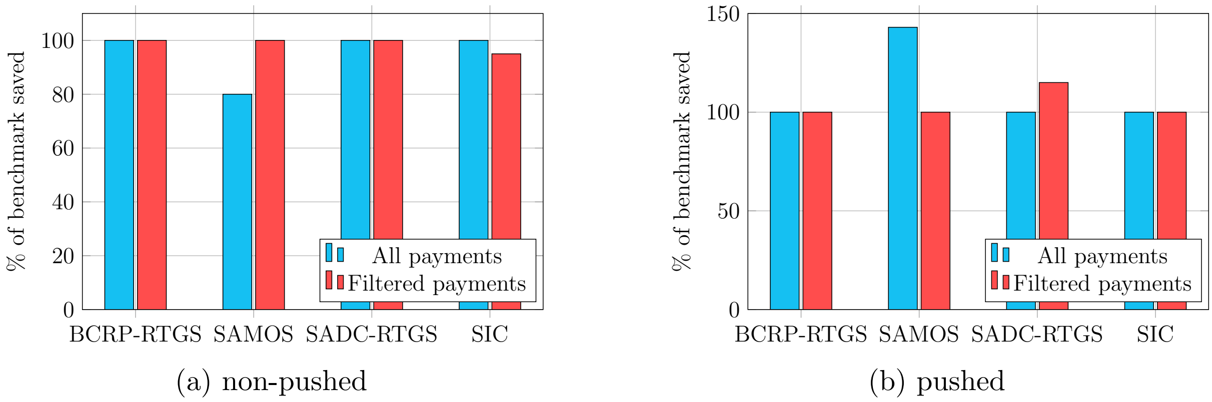

Graph 6: Realised liquidity savings as a percentage of potential savings based on netting at 10-minute intervals. "Non-pushed" means payments from failed auctions are settled immediately using RTGS. "Pushed" means payments from failed auctions are pushed forward into the next queue.

The right-hand panel considers the case in which payments associated with failed auctions are included in the following auction. Combining payment queues creates the potential for additional netting opportunities. Consequently, as our results show, the liquidity savings resulting from pushed payments start at 100% of the maximum obtainable benefit for 10-minute netting intervals and can go even higher. The bars in the left-hand panel are always at least as high as 100%. In two cases – SAMOS unfiltered payments and SADC-RTGS filtered payments - considerable additional liquidity savings result from failed auctions. This does not imply that failed auctions are a good thing because, in such cases, there is additional delay. The trade-off between liquidity savings and additional delay due to pushed payments is evaluated in the next subsection.

4.2.3 Liquidity savings and delay

Table 8 shows the average time (in seconds) that each dollar entered into the auction-based LSM is delayed. The average delay in seconds is computed as follows. For each cycle, the difference between the time a payment is submitted and the time the cycle ends is calculated. The time in seconds is then multiplied by the payment value. This is completed for all payments in the cycle and then these values are summed. The total is then divided by the total value of payments submitted during the cycle. The resulting number (in units of seconds) represents the average time each currency unit of payment is in the queue during that cycle. Finally, these numbers are averaged across all cycles. For unfiltered payments (top panel) and filtered payments (bottom panel), the first row of each panel shows the average delay in seconds in the case of non-pushed payments. For example, in the case of non-pushed payments for SIC, the average time that each Swiss franc worth of payments is delayed is 370.3 seconds. Recall that a queuing cycle is 10 minutes or 600 seconds. For context, note that if all the payments arrived at a uniform rate and were of equal value, an average delay of 300 seconds would be observed. The second row shows the delay when payments left in the queue after a failed auction are pushed into the next queue. These values are (weakly) larger than those in the first row, as pushing payments introduces additional delay. The third row calculates the difference. The difference is zero in the case of BCRP-RTGS because there are no failed auctions, and hence no pushed payments.

Table 8: The average delay in seconds per currency unit in the auction-based LSM

	BCRP-RTGS	**SAMOS**	**SADC-RTGS**	**SIC**
		Unfiltered		
Non-pushed	323.8	332.1	350.8	370.3
Pushed	323.8	446.2	352.7	434.7
Difference	0	114.1	1.9	65.4
		Filtered		
Non-pushed	14.3	251.4	60.3	169.9
Pushed	14.3	285.8	61.2	195.6
Difference	0	34.4	0.9	25.7

Table 9 compares the delay from pushing payments with the associated liquidity savings, for both the unfiltered (top panel) and filtered (bottom panel) scenarios. The first row of each panel converts the delay in seconds into a monetary value by pricing the delay in terms of each participant's benefit parameter b_i. Recall that the benefit parameters reflect the benefit, measured in the relevant currency unit, that each participant attributes to making a payment in the current cycle versus having it delayed until the next cycle. If a payment does not settle in the current cycle, then the next opportunity to settle will be at the end of the next cycle, which implies an additional 10-minute delay. The additional delay due to pushed payments can be converted into currency units by multiplying the dollar value of delayed payments for each participant in each cycle by the participant's benefit parameter and summing these values across all participants and all cycles. The resulting value is referred to as the social cost of additional delay. The second row provides the actual liquidity savings (in the relevant currency unit) from the additional delay, which results from the potential for increased netting opportunities that emerge when two queues are merged. These values are computed as the difference between the sum of all participants' maximum net debit positions that are realised when payments are not pushed and the sum of all participants' maximum net debit positions that are realised when payments are pushed. The latter value is necessarily less than the former value, as more netting opportunities result in smaller net debit positions. The third row takes the liquidity savings from the second row and computes the social benefit of not having to provide this liquidity. This social benefit is calculated by multiplying the amount of liquidity by the cost of capital (assuming an annual rate of 5% and dividing by 365 to obtain a daily rate).

According to the numbers for the unfiltered data, it would be sensible to push payments from failed auctions in both the SAMOS system and the SADC system since, in each case, the social cost of delay (ZAR 50,000 and ZAR 12, respectively) is less than the social benefit of liquidity savings (ZAR 91,781 and ZAR 55, respectively). In contrast, it would not be socially beneficial to push payments in the SIC system because the social benefit (CHF 685) is less than the social cost (CHF 5,100). The results for the filtered data

Table 9: Social cost and benefit of pushing payments from failed auctions forward to the next cycle.

	BCRP-RTGS (PEN)	SAMOS (ZAR)	SADC-RTGS (ZAR)	SIC (CHF)
	Unfiltered			
The social cost of additional delay	0	50 000	12	5100
Liquidity savings from additional delay (millions)	0	670	0.4	5
Social benefit from additional delay (5% annual rate)	0	91 781	55	685
	Filtered			
The social cost of additional delay	0	16 000	2	2300
Liquidity savings from additional delay (millions)	0	530	0	1
Social benefit from additional delay (5% annual rate)	0	87 123	0	164

are similar, although the social benefit and social cost of pushing payments are roughly equal in the case of SADC-RTGS.

4.2.4 Liquidity savings in absolute terms

The liquidity usage that result from netting in 10-minute intervals (potential and realised) are shown, together with the usages under RTGS and DNS for comparison purposes, in Tables 10 and 13. The auction-based LSM sits between the extreme scenarios of RTGS and DNS. Consistent with Graph 6b in subsubsection 4.2.3, the numbers show that with pushed payments the auction-based LSM can achieve a liquidity usage that is below that which is obtained when netting based on 10-minute netting cycles is implemented perfectly (ie no failed auctions).

Table 10: Liquidity usage (unfiltered, pushed). Potential and Realised (auction-based LSM) are based on 10-minute netting cycles. Liquidity usage is in the currency settled by each system, in millions.

Liquidity Usage	**BCRP-RTGS**		**SAMOS**		**SADC-RTGS**		**SIC**	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
RTGS	4917	4995	123 599	78 693	2642	2870	42 827	56 897
Potential	4707	4800	68 676	41 628	1738	2344	38 181	52 050
Realised (auction-based LSM)	4707	4800	96 147	60 778	1740	2346	38 186	52 051
DNS	1618	596	48 875	31 165	1614	1490	14 602	19 586

Table 11: Liquidity usage (unfiltered, unpushed). Potential and Realised (auction-based LSM) are based on 10-minute netting cycles. Liquidity usage is in the currency settled by each system, in millions.

Liquidity Usage	**BCRP-RTGS**		**SAMOS**		**SADC-RTGS**		**SIC**	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
RTGS	4917	4995	123 599	78 693	2642	2870	42 827	56 897
Potential	4707	4800	116 499	78 661	1729	2325	38 100	51 657
Realised (auction-based LSM)	4707	4800	93 199	62 929	1728	2324	38 094	51 651
DNS	1618	596	48 875	31 165	1614	1490	14 602	19 586

Table 12: Liquidity usage (filtered, pushed). Potential and Realised (auction-based LSM) are based on 10-minute netting cycles. Liquidity usage is in the currency settled by each system, in millions.

Liquidity Usage	**BCRP-RTGS**		**SAMOS**		**SADC-RTGS**		**SIC**	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
RTGS	276	105	98 265	57 854	245	371	15 380	20 598
Potential	267	104	94 910	55 931	197	270	14 454	19 961
Realised (auction-based LSM)	267	104	94 920	55 916	223	316	14 448	19 955
DNS	267	104	37 006	20 565	156	90	9849	12 754

Table 13: Liquidity usage (filtered, unpushed). Potential and Realised (auction-based LSM) are based on 10-minute netting cycles. Liquidity usage is in the currency settled by each system, in millions.

Liquidity Usage	**BCRP-RTGS**		**SAMOS**		**SADC-RTGS**		**SIC**	
	Day 1	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*	*Day 1*	*Day 2*
RTGS	276	105	98 265	57 854	245	371	15 380	20 598
Potential	267	104	84 738	55 605	223	319	15 523	20 553
Realised (auction-based LSM)	267	104	84 736	55 601	221	315	14 592	19 741
DNS	267	104	37 006	20 565	156	90	9849	12 754

5 Conclusion

The auction-based LSM demonstrates that LSMs can be designed to capture nearly all potential liquidity savings achievable from netting payments within fixed time intervals. In some systems, auction success rates reached 100%, and in instances when auctions failed, payments were typically settled in the subsequent round, resulting in minimal additional liquidity usage.

A key insight from the analysis is the importance of compensating liquidity providers for their role in facilitating settlement. The side payments proposed in this study are crucial because liquidity providers often value the immediate settlement of their own payments less than the cost of supplying liquidity. In contrast, liquidity receivers benefit from incoming funds and are therefore willing to pay. The auction mechanism effectively aligns these incentives, eliciting participants' willingness to pay and desire to receive compensation through bids. If participants behave as suggested by the agent-based model, such an auction framework could function effectively in practice.

The simulations employed 10-minute netting intervals, chosen to align with the timing of the CHAPS LSM. Longer netting intervals could yield greater potential liquidity savings but may reduce auction success rates as more participants may be involved. Although the study did not vary the interval length—the focus being on demonstrating the feasibility of the auction approach—system operators could adjust the timing of netting cycles to suit their liquidity and operational objectives in real-life scenarios.

Pushing payments from failed auctions into subsequent intervals represents another important design choice. The analytical framework developed here allows for an explicit assessment of the trade-off between additional liquidity savings and increased settlement delays. By expressing both in comparable monetary terms, using an interest rate that reflects the opportunity cost of liquidity, operators can make transparent and well-informed design decisions.

While the analysis assumes participants manage intraday liquidity optimally, it does not model their decision to submit payments to the queue. Two benchmark cases were considered: one in which all payments are submitted, and another in which only payments between counterparties with expected reciprocal flows are included. The latter "filtered" case identifies payments between system participants that are included in payment cycles within each 10-minute period and uses these payments as a proxy for informed submission behaviour. Although participants would not have perfect foresight of these cycles, it is reasonable to assume they can anticipate which participants are likely to make reciprocal payments. Thus, while the filtered scenario may overstate liquidity savings, it serves as a useful upper bound.

Incorporating an auction-based LSM could provide payment system operators with new tools to enhance efficiency and resilience. The results presented here are, however, preliminary, and further testing and calibration with more data would be required before drawing system-specific conclusions or considering implementation.

Acknowledgements

The authors would like to thank the Central Reserve Bank of Peru for BCRP-RTGS data, the South African Reserve Bank for SAMOS and SADC-RTGS data, and the Swiss National Bank for SIC data. Their provision of anonymised transaction data was crucial for evaluating the auction-based LSM. We also gratefully acknowledge the valuable contributions of Marta Stangierska, Josiah Friesen, Guilherme Ribeiro da Silva and Aleksandar Lazic, who served as research assistants and played an essential role in the development and analysis of this work. The work was largely completed while Rodney Garratt was a technical adviser at the BIS. We wish to thank seminar participants at the Bank of Finland's Simulator Seminar, the Federal Reserve Bank of Cleveland, the BIS Next Generation Financial Markets Seminar Series and the University of California, Santa Barbara for helpful comments.

References

M. Bech. Intraday liquidity management: A tale of games banks play. *FRBNY Economic Policy Review*, September 2008. doi: 10.2139/ssrn.1141366. URL https://doi.org/10.2139/ssrn.1141366.

M. Bech and R. J. Garratt. The intraday liquidity management game. *Journal of Economic Theory*, 109(2):198–219, 2003. ISSN 0022-0531. doi: https://doi.org/10.1016/S0022-0531(03)00016-4. URL https://www.sciencedirect.com/science/article/pii/S0022053103000164. Festschrift for Karl Shell.

M. Bech and R. J. Garratt. Illiquidity in the interbank payment system following wide-scale disruptions. *Journal of Money, Credit and Banking*, 44(5):903–929, 2012. doi: https://doi.org/10.1111/j.1538-4616.2012.00515.x. URL https://onlinelibrary.wiley.com/doi/full/10.1111/j.1538-4616.2012.00515.x.

M. Bech and B. Hobijn. Technology diffusion within central banking: The case of real-time gross settlement. *International Journal of Central Banking*, September 2006. doi: 10.2139/ssrn.932596. URL https://doi.org/10.2139/ssrn.932596.

M. Bech, Y. Shimizu, and P. Wong. The quest for speed in payments. *BIS Quarterly Review*, March 2017. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2931564.

M. L. Bech, A. Martin, and J. J. McAndrews. Settlement liquidity and monetary policy implementation - lessons from the financial crisis. *Economic Policy Review*, 18(Mar): 3–20, None 2012. doi: None. URL https://ideas.repec.org/a/fip/fednep/y2012imarp3-20nv.18no.1.html.

E. Benos and G. Harper. Recycling is good for the liquidity environment: Why ending qe shouldn't stop banks from being able to make chaps payments. Bank Underground, May 2016. URL https://bankunderground.co.uk/2016/05/18/recycling-is-good-for-the-liquidity-environment-why-ending-qe-shouldnt-stop-banks-from-being-able-to-make-chaps-payments/.

A. Carstens. Inflation is back, challenging central banks. Technical report, Bank for International Settlements, June 2022. URL https://www.bis.org/speeches/sp220626.pdf.

K. Chatterjee and W. Samuelson. Bargaining under incomplete information. *Operations Research*, 31(5):835–851, 1983. ISSN 0030364X, 15265463. URL http://www.jstor.org/stable/170889.

CPMI. Developments in retail fast payments and implications for rtgs systems. Technical report, Bank for International Settlements, December 2021. URL https://www.bis.org/cpmi/publ/d201.pdf.

N. Davey and D. Grays. How has the liquidity savings mechanism reduced banks' intraday liquidity costs in chaps. *Quarterly Bulletin*, Q2:180–189, 2014. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2464096.

G. Galati. Settlement risk in foreign exchange markets and cls bank. *BIS Quarterly Review*, December 2002. URL https://www.bis.org/publ/qtrpdf/r_qt0212f.pdf.

R. J. Garratt. An application of shapley value cost allocation to liquidity savings mechanisms. *Journal of Money, Credit and Banking*, 54(6):1875–1888, 2021. doi: https://doi.org/10.1111/jmcb.12853. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jmcb.12853.

D. B. Johnson. Efficient algorithms for shortest paths in sparse networks. *J. ACM*, 24(1):1–13, Jan. 1977. ISSN 0004-5411. doi: 10.1145/321992.321993. URL https://doi.org/10.1145/321992.321993.

B. Norman. Liquidity saving in real-time gross settlement systems: An overview. Financial Stability Paper 7, Bank of England, May 2010. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1786686.

B. Norman, R. Shaw, and G. Speight. The history of interbank settlement arrangements: Exploring central banks' role in the payment system. Working Paper 412, Bank of England, June 2011. URL https://www.ecb.europa.eu/home/pdf/research/Working_Paper_412.pdf.

D. Seaward. Saving liquidity in a liquidity-abundant world: Why don't banks use less liquidity when making high-value payments? Bank Underground, February 2016. URL https://bankunderground.co.uk/2016/02/08/saving-liquidity-in-a-liquidity-abundant-world-why-dont-banks-use-less-liquidity-when-making-high-value-payments/.

Appendices

A Johnson's algorithm for cycle identification and data filtering

This section outlines the application of Johnson's algorithm to identify cycles in payments data and describes how this information is used to filter the dataset.

A.1 Cycle detection in time-partitioned payment networks

To capture the dynamic nature of payment flows, the dataset was divided into discrete time intervals (eg, 10-minute windows). Within each interval, a directed graph was constructed in which:

- **vertices** represent participants; and
- **edges** represent payments from one participant to another.

Using Johnson's algorithm, all *simple cycles* within each graph were identified, ie closed paths where no vertex is revisited. This process captures both simple bidirectional loops and more complex multi-participant cycles. The outcome is a set of participants involved in at least one cycle for each time window.

A.2 Data filtering based on cycle membership

After detecting cycles, the data set was filtered to retain only payments between participants identified as part of a cycle within each interval, specifically:

- If a participant is involved in at least one cycle, all payments *to and from* other cycle members are retained.
- Payments involving any non-cycle participants are excluded.
- This filtering is applied independently within each time interval.

A.3 Example scenario

Consider the payment flows shown in Graph 7. Participants A, B, C and D are identified as cyclical, whereas participant E is not part of any cycle:

- Payments between A, B, C, and D are retained.
- The payment from B to E is discarded, since E is not a cyclical participant.

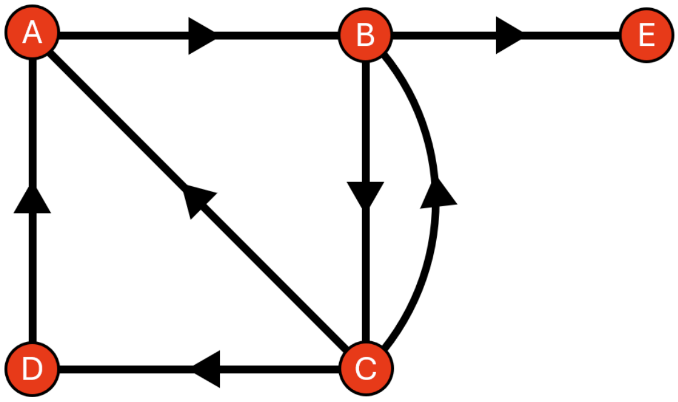

Graph 7: Directed graph of participants and payments. A–B–C–D–A form a cycle; E is not involved in any cycle.

B Quality check for agent-based model

Before application to payments data, the ABM was evaluated using a hypothetical scenario involving two participants, for which equilibrium bid functions can be computed. When queue participation is restricted to two participants, a simple auction can be simulated. Under certain distributional assumptions, the corresponding equilibrium bid functions can be derived.

With two system participants, one acts as the liquidity provider and the other as the liquidity receiver. The liquidity provider is treated as the seller (S), and the liquidity receiver as the buyer (B). The seller's and buyer's payoff functions are private information.

B.1 Individual steps involved in the auction

- **Step 1.** The system aggregates the queued payments of the two participants and computes their net positions.
- **Step 2.** The system sends a message to each participant indicating which payments will be netted and the amount of (net) liquidity each participant must provide. This amount equals zero for the participant in the net credit position and is positive for the counterparty in the net debit position.
- **Step 3.** The two participants submit their bids, denoted B_B and B_S .
- **Step 4.** If the bids overlap (ie $B_B \geq B_S$), then the allocation is implemented: the payments are settled on a net basis, and a side payment equal to $\frac{(B_B - B_S)}{2}$ is made from the buyer to the seller. Otherwise, the proposal is discarded, and the payments remain in the queue awaiting the next proposal.

B.2 Payoffs

$$\text{Buyer payoff} = \begin{cases} r_B - \frac{(B_B - B_S)}{2}, & \text{if } B_B \geq B_S \\ 0, & \text{otherwise} \end{cases}$$

$$\text{Seller payoff} = \begin{cases} r_S + \frac{(B_B - B_S)}{2}, & \text{if } B_B \geq B_S \\ 0, & \text{otherwise} \end{cases}$$

This setup can be mapped to the auction environment described by Chatterjee and Samuelson (1983) by defining bidder values as $v_B = r_B$ and $v_S = -r_S$. Within this framework, the equilibrium bid functions can be derived under additional assumptions. These theoretical bid functions are subsequently used to verify the solution of the ABM in this specific two-participant case.

Consider two payments: a \$1 million payment from participant 1 to participant 2 and a \$2 million payment from participant 2 to participant 1. The payments are assumed to be non-urgent, so that the benefit parameters b_i for each participant i are equal to zero. This configuration implies that participant 1 is the "buyer" and participant 2 is the "seller". It is further assumed that participants' liquidity costs or benefits do not depend on their settlement balances and that liquidity cost parameters c_B and c_S are drawn from the uniform distribution on $[0, 1]$. Consequently, bidders' values v_B and v_S are uniformly distributed on $[0, 1]$.

Following Chatterjee and Samuelson (1983), there exists a unique Bayesian Nash equilibrium in linear strategies, where:

$$B_B(v_B) = \frac{2}{3}v_B + \frac{1}{12} \quad \text{and} \quad B_S(v_S) = \frac{2}{3}v_S + \frac{1}{4}.$$

The two bidders reach an agreement whenever

$$B_B(v_B) \geq B_S(v_S),$$

which occurs if and only if

$$v_B \geq v_S + \frac{1}{4}.$$

Efficiency requires agreement whenever $r_B \geq r_S$, but no incentive-compatible mechanism fully achieves that outcome. In this setup, when $r_B \geq r_S$, agreement occurs 56.25% of the time. This relatively modest performance primarily reflects the uncertainty inherent in the uniform distribution. Tighter distributions of the liquidity cost parameters would yield better expected performance.

B.3 Results

The simulation generated 10,000 scenarios, resulting in 4,971 observations (cases in which the total surplus from settling the netted payments was positive). The results are presented in Table 14 and 15. The estimated bid functions closely approximate the theoretical bid functions.

Table 14: Buyer estimated bid function

Variables	**Coef**	**Std err**	**[95% conf Interval]**
Constant	0.06888	0.0298	[0.06304, 0.07472]
Slope	0.66904	0.00563	[0.65800, 0.68007]

Table 15: Seller estimated bid function

Variables	**Coef**	**Std err**	**[95% conf interval]**
Constant	0.25803	0.00311	[0.25193, 0.26412]
Slope	0.67663	0.00577	[0.66531, 0.68795]

Previous volumes in this series

All volumes are available on our website www.bis.org.