BANK OF CANADA
BANQUE DU CANADA

Staff Discussion Paper/Document d'analyse du personnel — 2020-9

Last updated: September 10, 2020

Liquidity Usage and Payment Delay Estimates of the New Canadian High Value Payments System

by Francisco Rivadeneyra and Nellie Zhang

Funds Management and Banking Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

riva@bankofcanada.ca, nzhang@bankofcanada.ca

ISSN 2369-9639

Acknowledgements

For their comments we thank Shaun Byck, James Chapman, Jorge Cruz-Lopez, Rod Garratt, Anneke Kosse, Wade McMahon, David Millard and Kevin Wai. The opinions here are of the authors and do not necessarily reflect those of the Bank of Canada. All errors remain our own.

Abstract

This paper presents simulation results for Canada's new large-value payments system: Lynx. We simulate the settlement process of Lynx using a large sample of payments observed in the current system (LVTS), taking the initial level of liquidity as given. We calculate the resulting liquidity usage, the payment delay and the shares of payments settled on a gross or net basis. The behaviour of participants (timing of payment submission) is assumed to remain the same as in LVTS. With an initial liquidity comparable to the collateral amount currently pledged in LVTS ($14.6 billion), Lynx FIFO Bypass would result in 28 minutes of average weighted delay and $17.3 billion of liquidity usage (the sum of intraday maximum net debit positions). Given this configuration, on average, $1.9 billion would be needed to clear non-urgent payments delayed until the end of the day, equivalent to 4.1 percent of payment value and 0.06 percent of volume. Doubling the amount of initial liquidity (to $29.3 billion) would result in 12 minutes of weighted delay. This basic configuration of Lynx requires a higher level of liquidity than LVTS and a plain-vanilla RTGS with pooled liquidity.

Bank topics: Payment clearing and settlement systems; Financial system regulation and policies; Financial services

JEL codes: C, C5, E42, E58

Non-technical summary

Canada is in the process of modernizing its core payments systems. As part of this process, the Large Value Transfer System (LVTS) will be replaced with a new Real-Time Gross Settlement (RTGS) system called Lynx consisting of two payment streams with separate liquidity saving mechanisms (LSMs) and separate collateral pools. LVTS is a hybrid system, with less urgent payments—the vast majority of payments—typically settling on a deferred net settlement (DNS) cover one basis. Naturally, like other DNS systems, LVTS requires less collateral to settle a given set of payments than a full RTGS system. However, given that it is desirable for Lynx to operate on as efficient a basis as possible, one of the most important questions for policy-makers is: How should LSMs in Lynx be designed?

This paper presents simulation results to aid in the design of the LSMs in Lynx. It is important to be able to measure the liquidity demands of the system and the corresponding amount of time it would take to settle the value and volume of transactions typically observed today in LVTS. To do so, we developed a simulation environment of the Lynx system using the description of its vendor.

We evaluated a variety of configurations of Lynx under several payment demand scenarios. With an initial level of collateral comparable to that pledged in LVTS today, Lynx with a FIFO Bypass configuration would result in 28 minutes of average weighted delay and $17.3 billion of liquidity usage (the sum of intraday maximum net debit positions). As expected, given that Lynx will be a full RTGS system rather than a partially DNS one, this basic configuration of Lynx would require a higher level of collateral than LVTS or a plain-vanilla RTGS with one liquidity pool instead of two.[1] This suggests that the scope that exists within the basic Lynx design to make use of various LSMs needs to be used.

[1]A plain-vanilla RTGS specifically refers to an RTGS system that is not equipped with any LSM.

1 Introduction

Canada is modernizing its core payments infrastructure. As part of this process, the Large Value Transfer System (LVTS) will be replaced by a new Real Time Gross Settlement (RTGS) system called Lynx.[2] LVTS is a hybrid cover one deferred net settlement (DNS)-RTGS system, deemed RTGS-equivalent thanks to a residual guarantee from the Bank of Canada.[3] Lynx, on the other hand, will be a pure RTGS with Liquidity Savings Mechanisms (LSMs), without the need of the residual guarantee. It will therefore require that participants pledge more collateral than they do in the current system. However, given its legislative mandate to ensure payments systems are efficient as well as safe, the Bank of Canada is interested in ways in which the system could be configured to be as efficient as possible in terms of liquidity usage.

This paper presents the first quantitative assessment of Lynx performed at the Bank of Canada. These results also help provide a recommendation for one important de- sign consideration of the LSM, that is the choice between a first in, first out (FIFO) and a FIFO Bypass configuration. The FIFO Bypass configuration yields substantially better results than the FIFO settlement sequence: for any given level of liquidity, more payments settle at an earlier time in the day. Our results provide empirical confirmation that, as anticipated, Lynx will have additional liquidity demands when compared with LVTS. If Lynx were to be operated with the collateral allocated currently to LVTS as initial liquidity, Lynx FIFO Bypass would result in a certain amount of delay. Using a random sample of 114 days of LVTS payments data, Lynx FIFO Bypass results in 28 minutes of average weighted delay, with $17.9 billion of used liquidity as measured by the sum of the intraday maximum net debit positions of all participants. For comparison, the level of liquidity allocated in LVTS in the same sample is $17.6 billion. Weighted delay is composed of the delay from intraday queued payments and the delay from payments forced to settle at the end of the day. Under the FIFO Bypass configuration, the payments delayed to the end of the day are 4.1 percent of value and 0.06 percent of volume. With double the initial liquidity, weighted delay would be reduced to 12 minutes. In this case, only 0.02 percent of the volume of payments would be delayed to the very end of the day. This discrepancy in value and volume indicates that the majority of the weighted

[2] The modernization program is led by Payments Canada. See Payments Canada (2016); and “High-Value Payments System” on the Payments Canada website.

[3] See Arjani and McVanel (2006) for a detailed description of LVTS. RTGS-equivalent means that clearing of payments is done on a gross basis in real time with legal finality. Net obligations are settled at the end of the day, as in deferred net settlement systems. The residual guarantee of the Bank of Canada is necessary to ensure real-time finality.

delay comes from a few large payments that are bypassed by the algorithm and therefore delayed until the end of the day.[4]

To benchmark the results, we compute liquidity-delay efficiency frontiers—the combinations of liquidity and delay that settle a set of payments in a given order—for Lynx and a simple RTGS (Rivadeneyra and Zhang 2020a). Theoretically, a simple RTGS with a single pool of liquidity could achieve a more efficient liquidity-delay frontier than Lynx. These theoretical minimums would be hard to attain in practice because of the uncertainty of the timing of payment flows. The exercises, however, reveal that the LSM of Lynx provides limited offsetting opportunities for payments to settle on a net basis.[5]

These results should be read as the first in a series of estimates of different configurations for the Lynx configuration process. While the simulation results of the basic Lynx configuration show higher liquidity requirements than LVTS, the exercises also suggest avenues to use available options to improve its design. One example, suggested by the benchmark results of the single-pool RTGS, would be to combine the payment submissions into one of the Lynx mechanisms with a single pool of liquidity. We perform these exercises in our follow-up paper (Rivadeneyra and Zhang 2020b).

Another caveat to the current results is that these exercises do not account for behavioural changes, as we have assumed that all payments are submitted at the same time they were submitted in LVTS. As of now, it is hard to ascertain what the effects of the incentives will be on the apportioning of liquidity and the timing of submissions. Our analysis, however, indicates that Lynx will likely require participants to actively manage their liquidity if they are to make most the efficient use of their collateral.

While the move to Lynx could increase liquidity and operational demands relative to LVTS, there are two important benefits from the change in risk model. Canada is currently the only member of the Committee on Payments and Market Infrastructure (CPMI) where the central bank underwrites some of the settlement risk of the wholesale payments system. Although the event where this risk could be realized is very remote, removing the residual guarantee would bring Canada to comply with the international standards set out in the Principles of Financial Market Infrastructures (PFMIs). Second, the move to Lynx could increase the resilience of the financial system by reducing the risk that LVTS could seize during crises. Since the vast majority of interbank payments in LVTS rely on bilateral credit lines that are extended voluntarily, these could be rapidly

[4]See Embree and Taylor (2015) for earlier simulation-based evidence of the additional required collateral from transitioning to a basic RTGS system. See also Byck and Heijmans (2020) for a simulation study on the use of alternative LSMs on LVTS data.

[5]Lynx will not net payments; instead, it will simultaneously offset payments. From a legal perspective, the gross amount of each payment is settled at the same time.

tightened during times of a crisis (Chapman et al. 2019). Since Lynx will not rely on this type of relationship, it would be more resilient to the sudden increase in the perceived risk of the participants.

The next section gives an overview of the proposed design of Lynx. Section 3 explains the simulation exercises. Section 4 shows the results, and Section 5 discusses the implications of the choice between FIFO and FIFO Bypass. Section 6 concludes with a discussion of the implications of our results. The Appendices explain the details of the samples and simulation methodology.

2 Proposed design of Lynx

This is a high-level overview of the settlement process of Lynx under FIFO Bypass rules. It is based on information provided by the vendor of the new system and from tri-party discussions between the Bank of Canada, Payments Canada and the vendor.[6] The proposed design of Lynx has two separate settlement mechanisms: one for urgent payments (the Urgent Payment Mechanism or UPM) and one for non-urgent payments (the Liquidity Optimization Mechanism or LOM). The UPM is just an RTGS with a simple FIFO queue for each participant.[7] The more complicated mechanism, where payment offsetting opportunities can be exploited, is the LOM. This mechanism is intended to settle non-urgent payments.

After a participant in the system submits a payment instruction, the settlement process in the LOM is the following:

1. **Gross liquidity and queue test.** A payment instruction will immediately settle if the payment is smaller than or equal to the available liquidity in the LOM *and* no other payments of the sending participant are in the queue. If both conditions are true, the liquidity gets consumed and the settlement balance is updated. On the contrary, if the payment value is greater than the available liquidity, or if the queue is non-empty, the payment will be queued.

2. **Event trigger called "Impact Intervention."** Payments remain in the queue until there is an event (more liquidity apportioned, received payment, cancellation, reordering). After such an event occurs, the payment in the queue is retried on a

[6] Note that our understanding and simulation implementation of the system have not been vetted by the vendor.

[7] See Committee on Payments and Market Infrastructures (CPMI) (2005) for a general description of different settlement mechanisms and queues.

FIFO Bypass order on a gross basis. If the payment value is smaller than or equal to the liquidity at that moment, the payment settles. The change in position of one participant has a cascading effect on other participants through the payments in their respective queues.

3. **Offsetting algorithm called "Gridlock Buster."** At the same time as the event mentioned above, the timer for a gridlock resolution algorithm is started. The remaining payments in the queue (which could include the original payment) will be tried with offsetting after the timer window is completed. Offsetting opportunities are used at this point. The Gridlock Buster is a complex process, proprietary of the vendor of Lynx, with many steps intended to find offsetting positions and resolve gridlock.

4. **Update of queue and liquidity positions.** After the gridlock algorithm is run, the payments that are chosen in this cycle are settled concurrently and the liquidity positions are updated. Note that not all payments that were in the queue when the gridlock ran will be included in the offsetting set because the gridlock algorithm might not have found a set of offsetting positions that included all payments in the queue. The payments that were not included in the set remain in the queue, and the settlement process continues again with step 1.

Note that the above settlement process does not allow a payment to be specifically directed so that it is settled exclusively by the offsetting algorithm. This means that a participant cannot flag a particular payment to be settled only on a net basis. This implies that a participant faces uncertainty as to how a payment in the LOM will settle: it might settle on a gross or a net basis. This depends on the liquidity provided by all participants in the system and the queue status at any point in time. Neither of these factors are known to participants.

The design states that payments are attempted to settle first on a gross basis if liquidity is available and the participant queue is empty. Furthermore, some payments could be tried twice on a gross basis by the Impact Intervention before they are attempted on a net basis by the Gridlock Buster. This implies that offsetting opportunities could be lost, even if the priority of a payment and the timing of its submission were intended to exploit offsetting opportunities.

3 Simulation methodology

To perform the simulation exercises, we developed a Lynx Simulator intended to replicate as closely as possible the flow of a payment instruction through the settlement process. To verify the accuracy, we compared the differences between our Lynx Simulator and a testing environment of Lynx that the vendor provided. The differences in the amount and timing of settlement, using more than 215,000 payment instructions over six simulation days, were negligible for practical purposes.

To complete a simulation exercise, we require making assumptions about: i) which payment instructions will be processed; ii) the timing of payment submissions; iii) the choice of UPM or LOM mechanism for each of these payments; iv) the level of initial liquidity that will be used; and v) the intraday management of liquidity.

3.1 Payment data

To examine the efficiency of Lynx over a multitude of scenarios, we use two large samples of LVTS payment instructions. For each scenario we vary the assumed level of initial liquidity to compute a variety of statistics of the settlement outcomes. The two samples of LVTS payments submission data contain 114 business days. The first is a random sample of the entire history of LVTS data. Days were selected randomly, so any given day in this sample should present a typical transaction day in LVTS. The second sample is the high-value sample: the 114 days with the highest payment values in LVTS. The size of the sample is equal to 5 percent of the available data. We use this size of sample so that we can report statistically meaningful results.[8] Regarding the timing of payment submissions, we make the following assumption.

Assumption 1. *The timing of payment submissions to Lynx is the same as in LVTS.*

3.2 Settlement mechanisms

Regarding the choice of settlement mechanism for each payment, we make the following two assumptions.

Assumption 2. *All T1 payments and payments to and from the Bank of Canada are submitted to the UPM.*

In the exercises, we assume that all the Tranche 1 (T1) payments and payments to and from Bank of Canada are directed to the UPM, based on the observation that these

[8]Refer to the Appendix A for more details of the samples.

are payments that need more immediate settlement or for which there is no tolerance for delay past a certain time, such as payments directed to financial market infrastructures. In the UPM, transactions are evaluated against risk controls upon their submission and will settle immediately if the payment-sending participant has sufficient liquidity. Otherwise, payments will be sent to this participant's UPM queue to wait for sufficient liquidity. Note that every participant has a queue in the UPM; these are individual queues, so no offsetting occurs between payments in different queues.

In the simulation of UPM, we do not use the queue, because that would require making assumptions of the participant's behaviour of making liquidity choices between mechanisms. The UPM will not offer central offsetting benefits anyway.[9] In other words, if a large-value urgent payment cannot pass risk controls, it might have to stay in the queue for an extended period: in reality, this is unlikely to happen to any urgent payment. Hence, in the simulation, participants are required to pledge more collateral whenever there is insufficient liquidity. Therefore, there is no payment delay in the simulated UPM.

Assumption 3. *All T2 payments are sent to the LOM.*

All the LVTS Tranche 2 (T2) payments are assumed to be non-urgent and are therefore sent to the LOM. Since they do not require immediate settlement, they can take advantage of the liquidity-savings features of the LOM. The LOM offers central queueing and a gridlock resolution algorithm that provide opportunities for payments from all participants to offset against each other. However, as explained before, in Lynx a payment will settle immediately upon submission if the sending participant has sufficient liquidity and there is no other transaction in its queue. In other words, if the queue is not empty, even if the participant has sufficient liquidity to settle this payment, it will be sent to the LOM queue.

3.3 Liquidity management

In the exercises, the UPM and the LOM settlement mechanisms are simulated separately, with separate liquidity pools. Intraday liquidity recycling is contained within each mechanism. In reality, after Lynx goes live, participants will be able to move liquidity across the settlement mechanisms; however, we avoided simulating this behaviour in our exercises. That means that in the simulation, a participant cannot use the liquidity received in the

[9] In Lynx, participants can choose not to use the queue in the UPM by pledging sufficient collateral at all times.

UPM to send payments in the LOM.[10]

We make three assumptions regarding the initial liquidity, intraday liquidity management, and end-of-day liquidity.

Assumption 4 (Initial Liquidity)**.** *For each day, the initial liquidity of a participant in the UPM is equal to its daily maximum intraday net debit position in the UPM. The initial liquidity in the LOM is a multiple (defined below) of the participant's maximum additional settlement obligation (MaxASO) for that day in the LVTS.*

In the UPM, we assumed that participants pledge exactly the amount of liquidity such that all payments are settled without any delay. To calculate this amount, we compute the maximum intraday net debit position for every day in the simulation. Therefore, by construction, there is no excess liquidity and no delay in the UPM.

For the LOM, participants start the day by pledging collateral equal to their LVTS T2 maximum additional settlement obligation (MaxASO).[11] We use MaxASO because it provides an estimate of the liquidity participants provided to LVTS based on the credit lines that were expected to allow settling of the typical value of incoming payments. Evidently LVTS is a very different system, but this level is an amount of collateral observed to have been allocated. We will vary this initial level in a systematic way to evaluate the tradeoff between liquidity and delay in Lynx. More of this below.

Assumption 5 (Intraday Liquidity)**.** *On an intraday basis, participants cannot pledge additional liquidity or transfer it between the UPM and LSM.*

In the LOM, once collateral is apportioned at the beginning of the day, it will not be changed intraday. Therefore payments that are delayed until the end of the day will require additional liquidity to settle at that point. We calculate and report this level of liquidity separately. We recognize that the previous two assumptions are strong constraints on behaviour. However, we think this is a useful benchmark because, as of this writing, we have not yet developed proper models of the participant responses to the new system.

Assumption 6 (End-of-Day Liquidity)**.** *If any payment has been delayed until the end of the payment cycle, the participant pledges additional liquidity equal to its multilateral net debit position.*

[10]When compared with the simple RTGS, we assume a single pool of liquidity. This provides an advantage to the RTGS regardless of not having an LSM.

[11]In LVTS, MaxASO is computed by multiplying the largest bilateral credit line a participant extends in the T2 with the system-wide percentage.

This assumption ensures that all payments are settled by the end of the payment cycle. To a certain extent, the handling of the end-of-day liquidity resembles the pre-settlement period of LVTS between 18:00 and 18:30, in which participants flatten their positions by lending to each other. The multilateral net position is a lower bound of additional liquidity requirements but avoids making assumptions on how the flattening would occur.

3.4 Liquidity and delay tradeoff

To examine the tradeoff between liquidity usage and payment delay in the LOM, we compute liquidity-delay efficiency frontiers, which are the pairs of liquidity usage and delay that settle a given set of payments in a given order (Rivadeneyra and Zhang 2020a). To compute these frontiers, we vary the level of initial liquidity, run the simulations for each level and recompute the outcomes of liquidity usage, payment delay and additional end-of-day liquidity, if necessary. To systematically vary the level of initial liquidity, we apply a factor α to the individual liquidity levels. This factor ranges between 1 and 2. Specifically, for participant i on day t, the level of initial liquidity, $C_{i,t}$, is computed as

$$C_{i,t} = \alpha \times \text{MaxASO}_{i,t}. \tag{1}$$

We vary α in increments of 0.2 so we compute six simulations in total. By varying MaxASO in this way, we preserve the proportions of collateral allocated to LVTS that are related to the value of payments a particular participant typically sends.

We define liquidity usage of the system as the sum of maximum net debit positions of all participants. Notice that in general, initial liquidity and liquidity usage will be different in the LOM depending on the intraday dynamics of incoming and outgoing payments.[12] For some participants on certain days, their level of initial liquidity might be larger than their liquidity usage. This indicates a buffer of liquidity in the LOM for those participants. A buffer at the system level does not imply that every participant is unconstrained on a given day. Therefore, even with buffers for some participants, we can observe delays and rejections. When initial liquidity is very low, most participants will be constrained and the central queue and the gridlock resolution algorithm will be used more frequently. In this case, the consumption of liquidity by some participants is mostly funded by the incoming payments.

[12]In the UPM, the initial liquidity and liquidity usage are identical because we chose initial to be the maximum net debit position.

3.5 Calculation of minimum initial liquidity: Iterative method

Recall that MaxASO is the amount of collateral that LVTS participants currently pledge in the system for T2 payments. Given the differences between LVTS T2 and an RTGS, it is unlikely that this level of initial liquidity without any intraday adjustment would settle all payments without delay. Therefore, to ensure that all payments in the data settle *before* they are delayed to the end of the day, we developed an iterative simulation method to calculate the level of initial liquidity that would ensure that all payments are settled before the end of the day. This method avoids the need to pledge additional liquidity at the end of the day to settle delayed payments.

We implement an iterative algorithm to calculate such level of liquidity. Starting from the level of liquidity equal to double the MaxASO for every participant, we calculate the liquidity shortfall for each participant that has any payment delayed by the end of the day. The algorithm adds this shortfall to the initial liquidity and repeats this process until there are no payments delayed by the end of the day. The algorithm also subtracts any excess liquidity from the initial level if a participant's intraday maximum net debit position is smaller than the beginning-of-day pledge, and it repeats this process as before. Therefore, in the end, when the algorithm finishes, the result is the minimum amount of liquidity required for every participant to settle all the payments on a given day.

To compare the iterative method with the standard simulation results, we create liquidity and delay curves by proportionally varying the minimum amount of liquidity that is generated by the iterative algorithm. We vary this liquidity by expanding it by one to two times for each participant. Note that by gradually increasing this minimum, every participant will have a buffer of liquidity and will not be a minimum anymore.

3.6 Lynx Straight-to-Queue variation

An important design feature of Lynx is that participants cannot select to settle a payment via Impact Intervention or the Gridlock Buster. To evaluate the implications of this design feature, we performed a second set of simulations, varying the setup of the payment submissions. In this straight-to-queue variation, all payments enter directly into the LOM queue without conditions. Even if funds are available and the central queue is empty, a payment will join in the queue and wait for the Impact Intervention or Gridlock Buster for offsetting opportunities. We call this "Lynx Straight-to-Queue."

Note that the Lynx Straight-to-Queue is not a straight-to-offsetting system. This means that although payments are sent to the queue, they might still settle on a gross basis from the Impact Intervention. A system allowing payments to settle only on a net basis would have required many additional modifications and assumptions.

3.7 RTGS benchmark

Lastly, to create a benchmark of the results, we simulate a pure RTGS system with and without a simple LSM. As in the Lynx simulation exercises, we create liquidity and delay efficiency frontiers by varying the initial level of liquidity and computing the associated delay. We compare the liquidity usage and delay of the Lynx simulations against the benchmark results.[13] One important difference with the Lynx simulations is that this RTGS benchmark is a single liquidity pool.

3.8 Additional assumptions and caveats

We made certain assumptions about some of the design options of Lynx, which could affect the results of the simulation. The assumptions are:

- The Gridlock Buster algorithm is set to run at five minute intervals.
- The LOM settlement sequence is set to FIFO Bypass.
- Lynx has an additional module called the Conditional Release Mechanism (CRM). This is a pre-settlement module prior to the settlement mechanisms such as the UPM and the LOM. It is designed mainly to allow payments to be released later, once conditions chosen by the participants are met. However, it can also be used for holding up the submission of payments that fail to meet a bilateral net send limit check. The CRM was not used in this study because it would have required a large amount of behavioural assumptions.

In summary, we simulate six systems. The configurations are: i) Lynx with some multiple of MaxASO as initial liquidity; ii) Lynx with some multiple of MaxASO as initial liquidity and the straight-to-queue submission to the LOM; iii) Lynx with minimum initial liquidity; iv) Lynx with minimum initial liquidity and the straight-to-queue submission to the LOM; v) RTGS; and vi) RTGS with LSM. This is summarized in Table 1.

[13] In the RTGS, all payments are tried on a gross basis with available liquidity. Liquidity is varied by reducing the level of liquidity that would have settled all payments without delay. Payments that fail initially upon submission are settled at the end of the day on a multilateral net basis. This simulation computes the minimum liquidity necessary for a given set of payments allowing for the maximum delay. The LSM added to the RTGS is similar to the Jumbo algorithm of LVTS without the size constraint. More details are available in Appendix B.

Table 1: Exercises performed. We simulated six systems for each sample of payments data (random and high-value). The differences between these exercises are determined by the level of initial liquidity used and variations of queues. For initial liquidity we use MaxASO or the iterative method to compute the minimum required. For the queues, in Lynx we used a straight-to-queue variation (which we call Lynx Straight-to-Queue) and in the RTGS we added the LVTS Jumbo algorithm as an LSM.

		Initial liquidity	Variations
1	Lynx	MaxASO	
2	Lynx Straight-to-Queue	MaxASO	Straight-to-queue
3	Lynx	Iterative method (min initial)	
4	Lynx Straight-to-Queue	Iterative method (min initial)	Straight-to-queue
5	RTGS	Minimum liquidity	
6	RTGS with LSM	Minimum liquidity	Jumbo algorithm from LVTS

4 Results

First, we present the liquidity-delay efficiency frontiers for each system and for each sample. Second, we present the tables of statistics that report initial liquidity, liquidity usage, various statistics of delay and the shares of payments settled on a gross and net basis. Third, to understand how liquidity usage is distributed, we show graphs of the liquidity apportioned by each participant in different simulations.

4.1 Liquidity and delay

The main simulation results are presented in figures 1, 2 and 3, which show the liquidity-delay frontiers based on the average of the maximum net debit position in every day in the sample, and the average of the value-weighted time of payments delayed in every day in the sample. We report the results for the random and high-value samples in green and red colors respectively. Each point in the curve is a simulation using a different level of initial liquidity. Tables 2 to 5 provide more details on the total number of payments delayed, the percentage of payment value and volume delayed till the end of the day, and the percentage of payment value and volume that is settled on a gross basis or a net basis (by the Gridlock Buster).

As expected, when initial liquidity is reduced, the amount of delay increases, resulting in the known tradeoff between liquidity and delay. The main simulation exercise of Lynx (Figure 1) shows the range of liquidity usage between $17 and $20 billion, with an associated weighted delay of 28 and 12 minutes for the random sample. For the high-value sample, the range is $25 and $28 billion, with an associated weighted delay of 24 and 10 minutes. Figure 1 also compares these curves with the simple RTGS with pooled liquidity. RTGS shows a better tradeoff because liquidity recycling is obviously enhanced in a single pool. The curves also show that sending payments directly to the queue (Lynx Straight-to-Queue) does not reduce the liquidity usage.[14] Finally, and very importantly, note that the average liquidity allocated to LVTS in this sample is $14.6 billion, indicating that if Lynx is used in the way we have assumed here, it will require more liquidity than LVTS.

The next important set of results is the one from the iterative calculation of the minimum liquidity required to settle all payments before the end of the day (Figure 2). These show that liquidity usage in the random would range between $20 and $24 billion, with an associated weighted delay between 8 minutes and 1 minute for the random sample. For the high-value sample, the range is $29 and $33 billion, with an associated weighted delay of 6 minutes and 1 minute. For ease of comparison, Figure 3 shows only the Lynx curves simulated with the MaxASO initial liquidity and with the iterative calculation.

To better understand the results, it is helpful to go over a few details in Table 2. Take the first row: this level of initial liquidity is equal to the amount of liquidity that participants apportion currently to T2 in LVTS. For the random sample, the initial liquidity in the LOM is $14.6 billion. Keep in mind that this does not include the liquidity for the UPM, which for this sample is $10.8 billion. With this amount of liquidity, the average weighted delay is 28 minutes. In this sample, on average 1,643 payments of LOM payments are queued. Recall that even if a payment is delayed by entering the queue, it could still be settled on a gross basis due to Impact Intervention. In fact, the last column shows that only 16.4 percent of LOM payments in value are settled on a net basis.[15] The payments that are settled on a net basis tend to be large payments, because the results show that the shares of value settled on a net basis is significantly larger than its volume. Lastly, notice that 4.1 percent of LOM payments in value do not settle within the day, but only 0.06 percent by value because of the FIFO Bypass settlement sequence. To settle these payments we calculate the amount of liquidity that would be necessary on a multilateral offsetting basis, which is on average $1.9 billion. We interpret this as the value of overnight loans that would have been traded to settle all delayed payments.

Now take the last row of the random sample in Table 2, which corresponds to the Lynx simulation, with level of initial liquidity equal to double the MaxASO.[16] For the random

[14]The explanation is that sending payments to the queue does not ensure that they will be settled with offsetting positions; in fact, they rarely are, because payments will be tried in the Impact Intervention first.

[15]To be more precise, this share should be somehow adjusted by the number of payments that settle at the end of the day.

[16]Note that the value of double the MaxASO is in close vicinity of the T2 net debit cap in LVTS (T2NDC) in the data.

sample, initial liquidity in the LOM is $29.3 billion. With this amount of liquidity, the amount of weighted delay is 12 minutes. In this simulation, only 708 payments enter the queues for the opportunity of settling by the gridlock resolution algorithm. In this case, only 0.1 percent of volume or 1.2 percent in value is settled on a net basis. This result is intuitive: when the liquidity level is high most payments settle on a gross basis upon submission and therefore are less likely to become queued. This has a cascading effect: since the central queue is empty more often, subsequent payments are tried on a gross basis more often before they enter the queue. In spite of having a large amount of initial liquidity, 1.2 percent of payments are still delayed until the end of the day, which requires $0.8 billion of liquidity to settle in the end of the cycle. Note that the difference between the initial liquidity and the intraday maximum net debit position (maxNDP) can be interpreted as liquidity buffers. It is evident from the last row of the table, which is the case of doubling MaxASO as initial liquidity in the high-value sample, that the buffer is much smaller than the same simulation case in the random sample.[17]

To understand how the different distribution of initial liquidity among participants has important effects on the outcomes of liquidity usage and delay, we analyze the results of the simulations using the iterative method in Figure 2 and Table 4. The iterative algorithm finds the minimum level of initial liquidity for each participant that ensures that no payments sent by that participant would be delayed by the end of the day. The iterative algorithm starts at the liquidity level that is equal to double the MaxASO of each participant, and tracks each participant's intraday net debit position at any given time. This results in a different distribution of liquidity usage for most participants.

As expected, weighted delay is greatly reduced because large-value payments account for the lion's share of payments that are settled at the end of the day in the first Lynx exercise. More precisely, the first row of Table 4 shows that the weighted delay is 8 minutes. The minimum initial liquidity found by the algorithm is $11.4 billion, which is much lower than the amount of double the MaxASO ($29.3 billion). In terms of liquidity usage, this exercise results in $20.6 billion, which is very close to the amount of used liquidity of $20 billion (maxNDP of $19.9 billion plus the $0.8 billion of the end of the day), in the previous Lynx simulation exercise that is not based on the iterative method.

Note that the simulations based on the iterative algorithm show an increase in the share of payments settled on a gross basis, although the initial liquidity is smaller (see the last four columns in Tables 2 and 4). This implies that a change in the initial liquidity available can lead to a different intraday liquidity distribution, which in turn results in

[17]Note that the difference of the initial liquidity and the intraday maxNDP is a rough approximation of the buffer. A true buffer should be calculated for each participant and for each day prior to averaging.

very different simulation outcomes.

4.2 Distribution of liquidity among participants

The simulation results show that a slightly lower level of initial liquidity (generated by the iterative algorithm) can yield a more efficient outcome of less settlement delay. This is because the iterative algorithm adjusts the collateral requirements at each participant's level to make sure all payments settle intraday; this liquidity allocation across the participants turned out to be the right proportions. To understand this result, we analyzed the distribution, across the days in the sample, of the intraday maximum net debit positions of each participant.

The box plots in Figure 9 show the distribution of maxNDP of both exercises for the random sample. The graphs show the mean, interquartile range and outliers for each participant. Figure 10 illustrates the same for the high-value sample. Comparing the top and bottom plots, we see that most of the differences in intraday liquidity usage come from the large banks, whereas the distributions are very similar for all small participants. In the high-value sample, the pattern is somewhat more pronounced, but on the other hand, the medians on most cases are quite close. This suggests that small changes in the liquidity usage can lead to significant improvements in settlement efficiency.

5 FIFO versus FIFO Bypass

An important system configuration is the choice between FIFO and FIFO Bypass for the LOM queue. To help in the policy recommendations of this choice, we also computed the four Lynx exercises with the FIFO queue configuration in the LOM, leaving all the other aspects of the simulation unchanged. As with the Bypass exercises, the Lynx FIFO simulation results are presented in a series of figures and tables. Figures 6, 7 and 8 show the liquidity-delay frontiers based on the average of the maximum net debit position in each day and the average of the value-weighted time of payments delayed in each day. Tables 6 to 9 provide additional details on the total number of payments delayed, the percentage of payment value and volume that is delayed till the end of day, and the percentage of payment value and volume that is settled on a gross basis or a net basis (by the Gridlock Buster).

Comparing the results of the FIFO and the FIFO Bypass simulations shows that Bypass always performs better for both samples and for any system variation and level of initial liquidity. This is not surprising given the setup of our exercises. Likewise, there is some empirical evidence of the efficiency improvements of a Bypass configuration. The logic behind this result is that by allowing participants to bypass large payments, the system reduces delay by permitting smaller payments to settle. Evidently there is a tradeoff in this configuration: by bypassing large payments and settling smaller ones, a participant cannot accumulate the liquidity to settle the large bypassed payments. This could increase the weighted delay if large payments are delayed for too long. This implies that a particular distribution of payment size and timing of submission to the system could have yielded the opposite result, for example, if large payments tend to be submitted early in the day.

To examine this further, Figure 11 shows that the size of payments in LVTS has become smaller over the past 16 years. The plot shows the percentiles of value plotted on a log scale. The reduction in payment size is evident in the 50^{th}, 75^{th}, 90^{th} and 95^{th} percentiles. The sharpest decline occurred during the 2008–10 financial crisis period. In general, regardless of the settlement sequence in a payments system, smaller transaction size promotes liquidity recycling and speeds up the settlement. For example, in the case of FIFO, the likelihood of big payments blocking the queue and preventing subsequent transactions from settling is decreased if the typical transaction is small in value. Given the distribution of payment sizes in LVTS, we conjecture that a FIFO Bypass rule will work better in Lynx than a simple FIFO sequencing, because a majority of payments in LVTS can settle before the few large payments keeping liquidity recycling throughout the day.

One last aspect to consider is the potential liquidity cost if large payments are generally time sensitive or even critical. If this is the case, then, to avoid penalties, participants will ensure they have enough liquidity to settle these payments with minimum delay. Therefore, the configuration between FIFO and FIFO Bypass would not matter much for delay but would increase the liquidity usage of the system. However, the relationship between size and priority has not been examined so we leave that for future simulations.

6 Concluding remarks

This paper presented a quantitative assessment of Lynx, the new large-value payments system in Canada. We performed a variety of exercises to examine the efficiency of various configurations. The key findings of the simulation exercises are the following.

For a typical day, Lynx liquidity usage could range between $17 and $20 billion, with an associated weighted delay of 28 and 12 minutes. For the high-value sample, the range is \$25 and \$28 billion, with an associated weighted delay of 24 and 10 minutes. Most of this delay comes from a few large delayed payments. The results indicate that Lynx, if implemented as in this paper and with the same behaviour from participants, could require more liquidity than the liquidity allocated to LVTS today.

The separation of liquidity into two pools (the UPM and the LOM) results in a significant loss of efficiency in Lynx. Even if the collateral requirement is increased to double the amount that LVTS participants currently pledge, settlement delay would be notably more than a simple RTGS operating on a single pool of liquidity. These results point towards ways to improve the configuration of Lynx, for example encouraging the participants to use the LOM mechanism to settle all payments, including the urgent. We perform these simulations in our follow up paper (Rivadeneyra and Zhang 2020b).

Although not available at this moment, a version of Lynx allowing a straight-to-LSM option that directs queued payments to settle in the Gridlock Buster only (preventing them from settling through Impact Intervention) could improve the delay and liquidity tradeoff of the system.

The results presented in this paper suggest that further analysis is needed to determine the best setup of Lynx within the available options. In future work we will examine several policy questions, for example: What would be the implications for liquidity efficiency, delay and liquidity management if all payments were sent to only one of the mechanisms as a single pool? And what are the implications of setting minimum liquidity requirements and establishing throughput guidelines?

References

Arjani, N. and D. McVanel. 2006. “A Primer on Canada's Large Value Transfer System.” Bank of Canada, March 1.

Byck, S. and R. Heijmans. 2020. “How Much Liquidity Would a Liquidity-Saving Mechanism Save If a Liquidity-Saving Mechanism Could Save Liquidity? A Simulation Approach for Canadas Large-Value Payment System.” Payments Canada Working Paper, February 18.

Chapman, J. T., M. Gofman and S. Jafri. 2019. “High-Frequency Analysis of Financial Stability.” Mimeo, University of Rochester.

Committee on Payments and Settlement Systems (CPSS). 2005. “New Developments in Large-Value Payment Systems.” Bank for International Settlements, Technical Report.

Embree, L. and V. Taylor. 2015. “Examining Full Collateral Coverage in Canada's Large Value Transfer System.” Bank of Canada Staff Working Paper No. 2015-29.

Payments Canada. 2016. “Developing a Vision for the Canadian Payment Ecosystem.” Technical report, Payments Canada White Paper.

Rivadeneyra, F. and N. Zhang. 2020a. “Efficiency Frontiers of Payments Systems.” Mimeo, Bank of Canada.

Rivadeneyra, F. and N. Zhang. 2020b. “Liquidity, Delay and Operational Demands in Lynx.” Mimeo, Bank of Canada.

Table 2: Liquidity usage, delay and netting of Lynx. This table presents the averages of the statistics across the days in each sample. The first column is the alpha multiples of MaxASO (from 1 to 2) used for the initial liquidity of the LOM mechanism used to settle all T2 payments plus the UPM liquidity. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary with alpha). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day. Total used liquidity is therefore the sum of columns 2 and 3. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns labelled Settled intraday are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq. Billion $	MaxNDP Billion $	EoD liq. Billion $	Delay hh:mm	Delayed payments (number)	Settled EoD		Settled intraday			
					Value	Volume	Value		Volume	
							Gross	Net	Gross	Net
Random sample										
14.6	17.3	1.9	00:28	1643	4.1	0.06	83.6	16.4	99.7	0.3
17.6	17.9	1.6	00:24	1372	3.4	0.05	87.1	12.9	99.8	0.2
20.5	18.4	1.4	00:20	1139	2.5	0.04	88.6	11.4	99.8	0.2
23.4	18.9	1.1	00:17	973	1.9	0.03	90.6	9.4	99.8	0.2
26.3	19.5	0.9	00:14	848	1.5	0.02	92.4	7.6	99.9	0.1
29.3	19.9	0.8	00:12	708	1.2	0.02	93.7	6.3	99.9	0.1
High-value sample										
15.2	25.2	2.7	00:24	2587	4.4	0.07	80.8	19.2	99.6	0.4
18.3	25.8	2.4	00:20	2187	3.5	0.05	83.9	16.1	99.7	0.3
21.3	26.4	2.1	00:17	1864	3.0	0.05	86.5	13.5	99.8	0.2
24.4	26.9	1.8	00:14	1549	2.4	0.03	88.9	11.1	99.8	0.2
27.4	27.4	1.6	00:12	1297	2.0	0.03	91.0	9.1	99.8	0.2
30.5	27.9	1.3	00:10	1147	1.7	0.02	92.8	7.2	99.9	0.1

Figure 1: Delay and liquidity of Lynx, Lynx Straight-to-Queue, RTGS and RTGS with LSM. This plot shows the liquidity and delay trade off of the systems for two samples. Green is the random sample, and red is the high-value sample. Samples are of 114 days. Delay time is the average of the value-weighted delay. The order of payment submissions is the same in all exercises. The dashed and dashed-dot lines are Lynx and Lynx Straight-to-Queue. In these simulations, participants' initial liquidity is a multiple of the MaxASO. The top point of each curve is equal to 1xMaxASO, and the bottom point is 2xMaxASO. In these two systems, some payments are delayed all day but settled calculating a round of multilateral netting at the end of the day. In Lynx Straight-to-Queue, payments are submitted immediately to the queue. The solid lines are the RTGS system, representing minimum amounts of liquidity and resulting delay without any LSM. The dotted lines are the RTGS with LSM, which adds the Jumbo algorithm of LVTS. Delayed payments in a first pass can settle before the end of the day via this netting algorithm. Additional details are reported in tables 2 and 3.

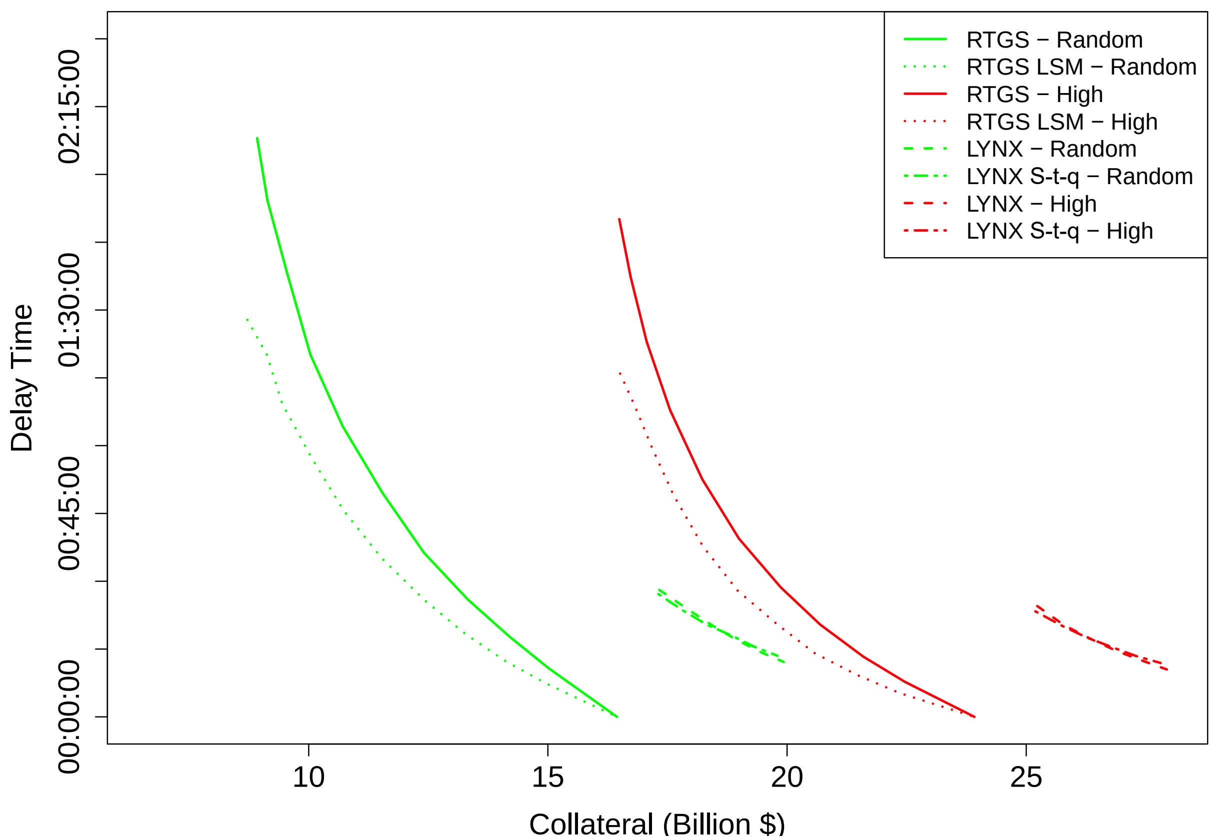

Figure 2: Delay and liquidity of Lynx and Lynx Straight-to-Queue (with the iterative method), RTGS and RTGS with LSM. This plot shows the liquidity and delay tradeoff of the systems for two samples. In green is the random sample and in red the high-value sample. Samples are of 114 days. Delay time is the average of the value-weighted delay. The order of payment submissions is the same in all exercises. The dashed and dashed-dot lines are Lynx and Lynx Straight-to-Queue respectively. In these simulations, participants' initial liquidity is the minimum required to settle all payments intraday. In Lynx Straight-to-Queue, payments are submitted immediately to the queue. The solid lines are the RTGS system, representing minimum amounts of liquidity and resulting delay without any LSM. The dotted lines are the RTGS with LSM, which adds the Jumbo algorithm of LVTS. Delayed payments in a first pass can settle before the end of the day via this netting algorithm. Additional details are reported in tables 4 and 5.

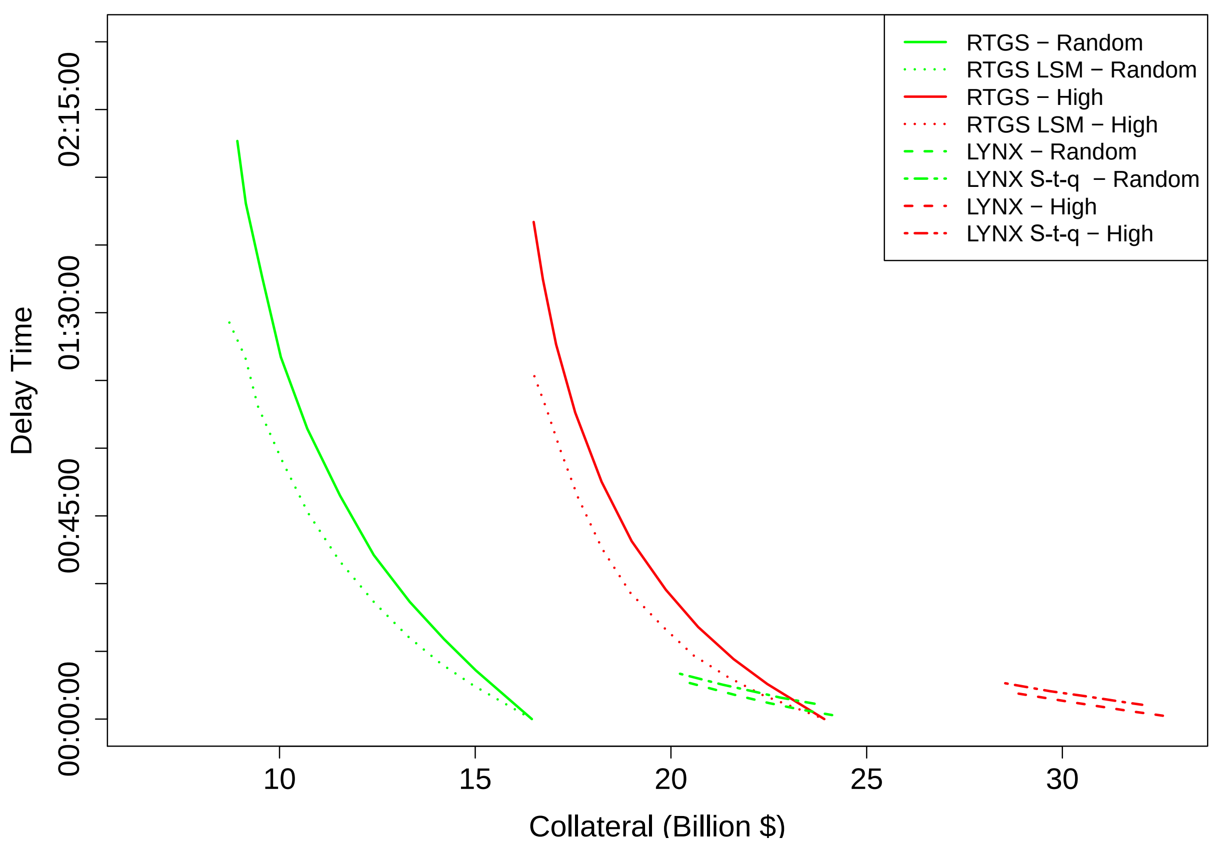

Figure 3: Delay and liquidity of Lynx and Lynx Straight-to-Queue with the MaxASO and iterative method simulations. This plot shows the curves in Figures 1 and 2 for ease of comparison. The bottom four curves are the ones using the iterative method. The upper four curves use MaxASO as initial liquidity.

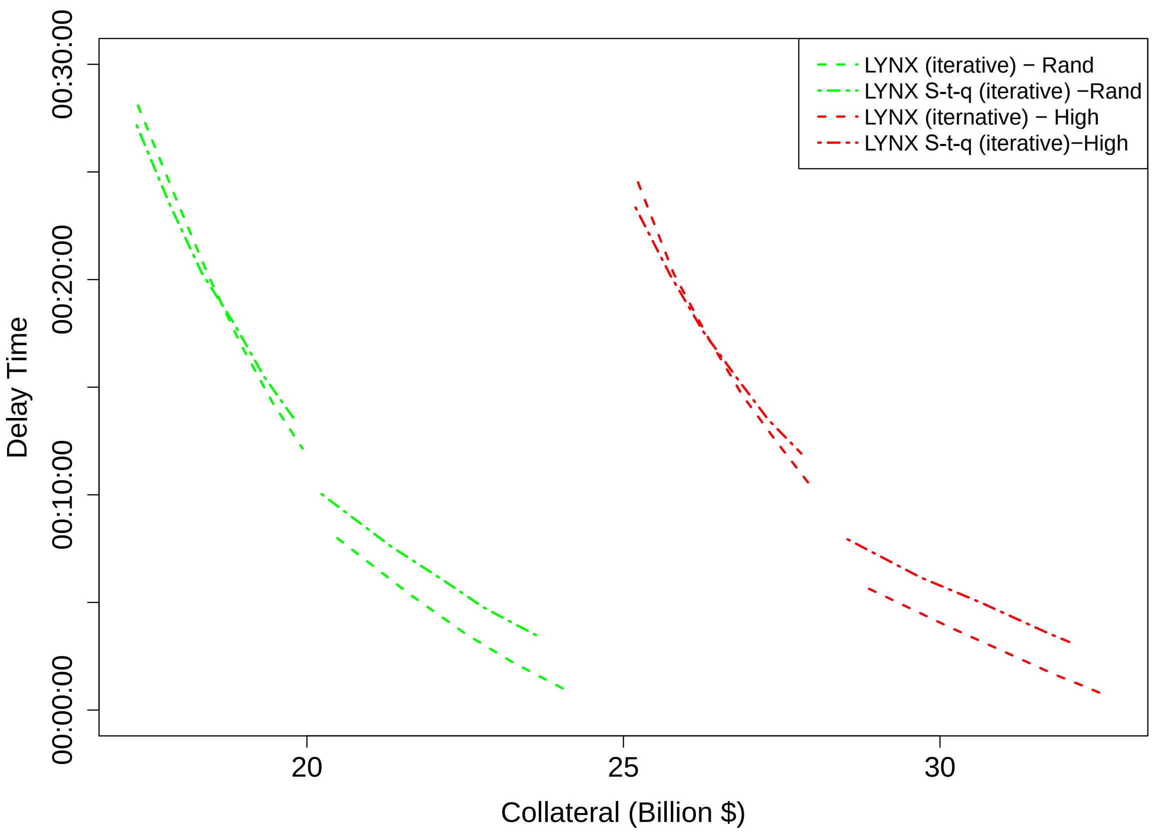

Figure 4: Lynx liquidity usage by participants in the random sample. This is the distribution of the LOM maxNDP in Lynx. The top plot shows the exercise when initial liquidity is double the MaxASO. The bottom plot is the exercise when initial liquidity is calculated iteratively.

(a) Initial liquidity 2xMaxASO

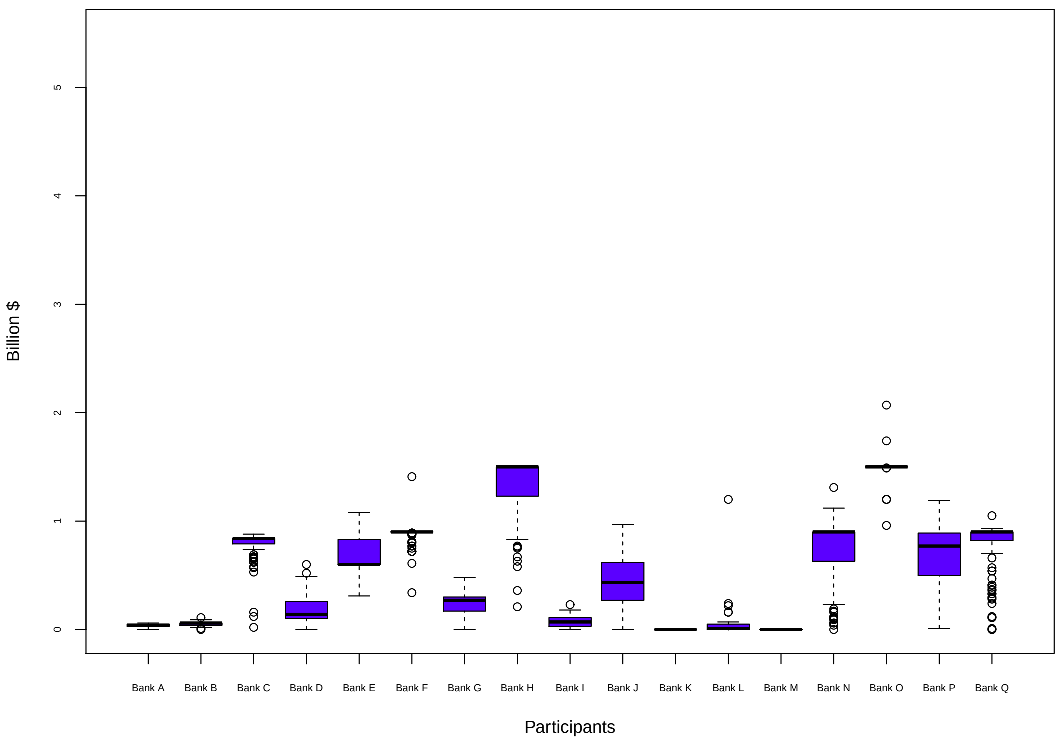

(b) Iterative method

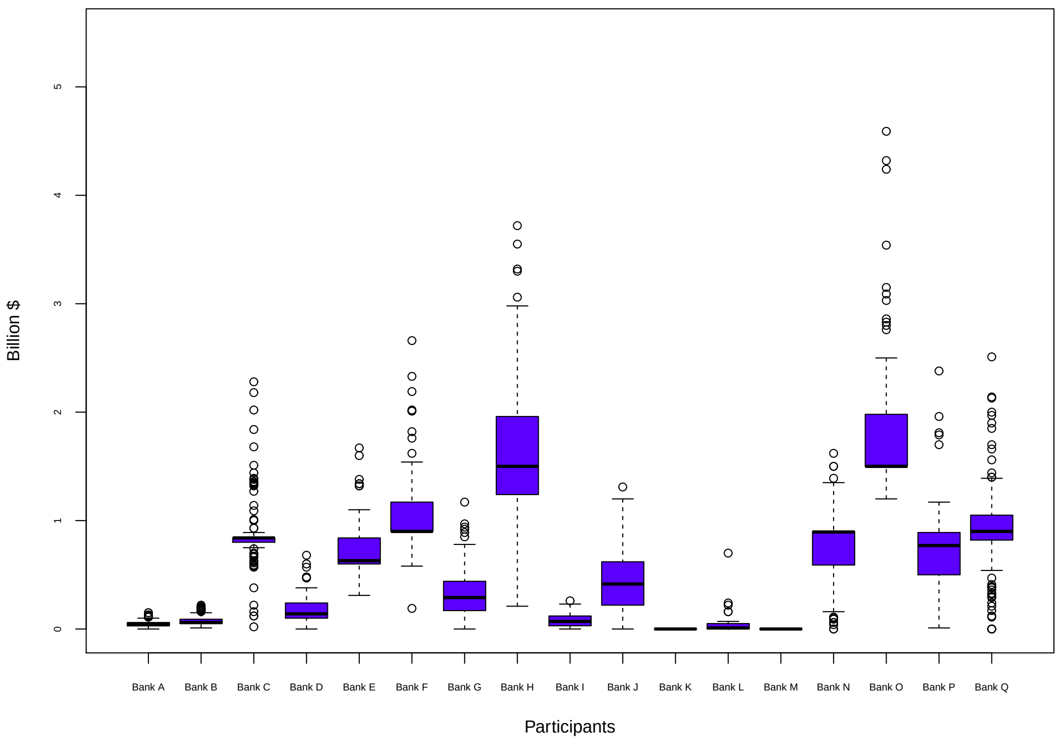

Figure 5: Lynx liquidity usage by participants in the high-value sample. This is the distribution of the LOM maxNDP in Lynx. The top plot shows the exercise when initial liquidity is double the MaxASO. The bottom plot is the exercise when initial liquidity is calculated iteratively.

(a) Initial liquidity 2xMaxASO

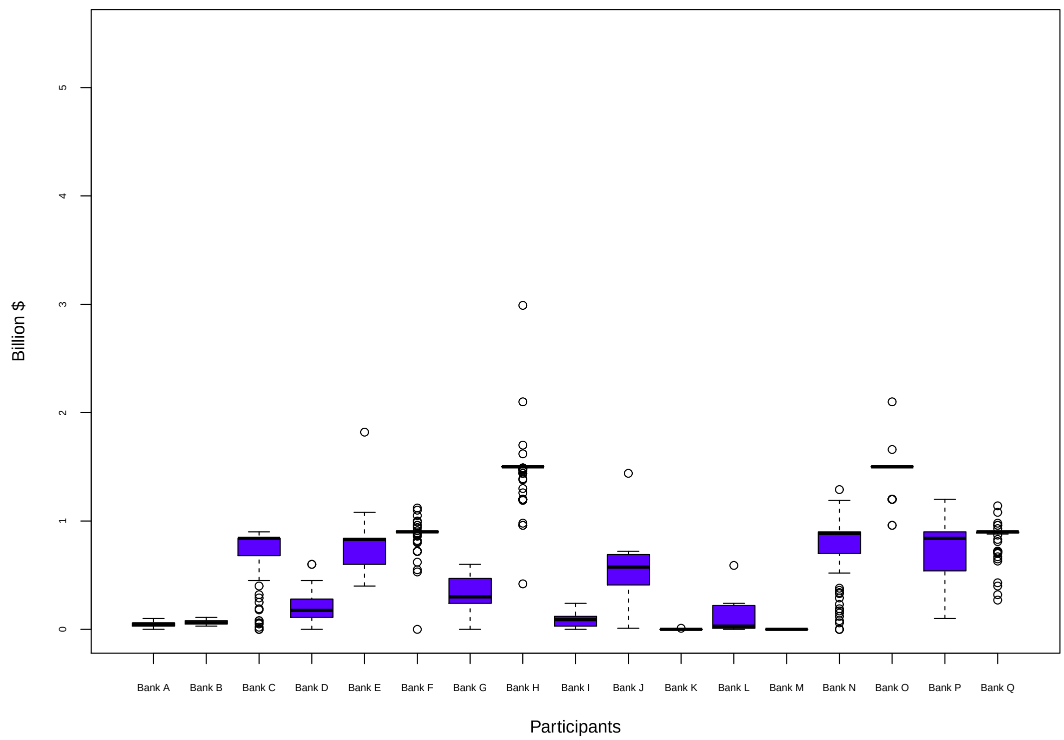

(b) Iterative method

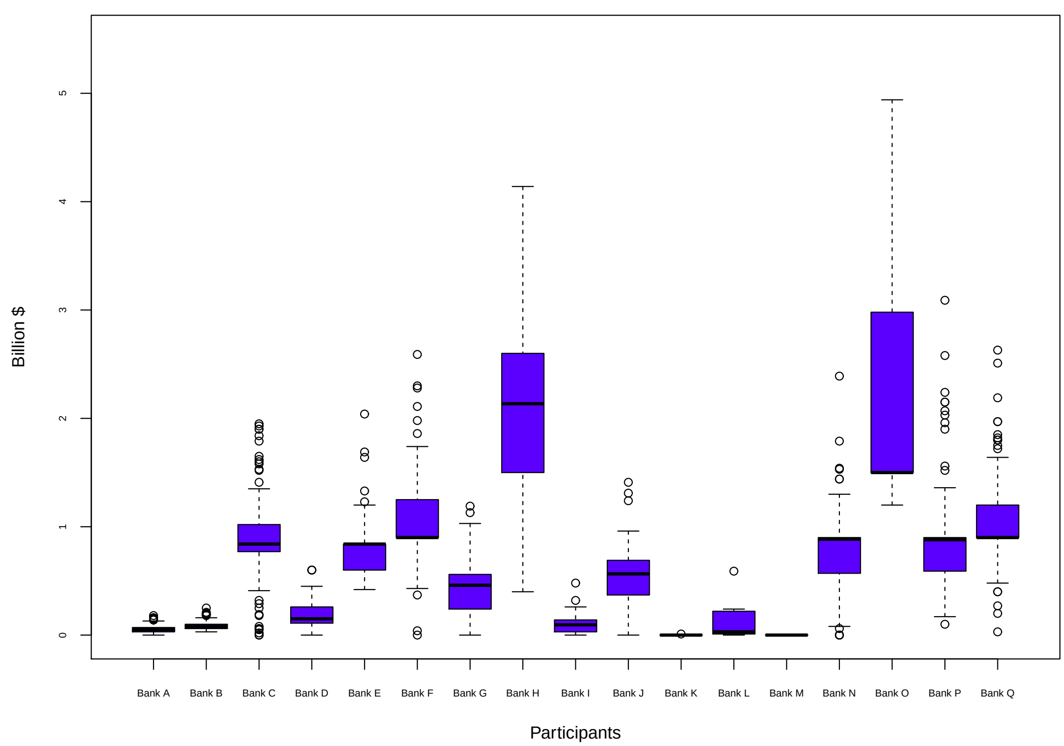

Table 3: Liquidity usage, delay and netting of Lynx Straight-to-Queue. This table presents the averages of the statistics across the days in each sample. The first column is the alpha multiples of MaxASO (from 1 to 2) used for the initial liquidity of the LOM mechanism used to settle all T2 payments plus the UPM liquidity. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary with alpha). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day. Total used liquidity is therefore the sum of columns 2 and 3. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns labelled Settled intraday are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq.	MaxNDP	EoD liq.	Delay	Delayed	Settled EoD		Settled intraday			
Billion $	Billion $	Billion $	hh:mm	payments	Value	Volume	Value		Volume	
				(number)			Gross	Net	Gross	Net
Random sample										
14.6	17.3	1.9	00:27	26801	3.5	0.06	16.5	83.5	8.2	91.8
17.6	17.8	1.6	00:23	26801	2.7	0.04	14.7	85.3	7.7	92.3
20.5	18.3	1.4	00:20	26801	2.2	0.04	12.9	87.1	7.3	92.7
23.4	18.9	1.1	00:18	26801	1.9	0.03	11.8	88.2	7.0	93.0
26.3	19.3	0.9	00:15	26801	1.5	0.03	10.4	89.6	6.8	93.2
29.3	19.8	0.8	00:13	26801	1.2	0.02	9.5	90.5	6.6	93.4
High-value sample										
15.2	25.2	2.7	00:23	34805	4.0	0.08	16.8	83.2	8.9	91.1
18.3	25.7	2.4	00:20	34805	3.2	0.06	15.2	84.8	8.4	91.6
21.3	26.3	2.1	00:17	34805	2.5	0.04	13.7	86.3	7.9	92.1
24.4	26.8	1.8	00:15	34805	2.2	0.03	12.6	87.4	7.6	92.4
27.4	27.3	1.6	00:13	34805	1.8	0.03	11.5	88.5	7.3	92.7
30.5	27.8	1.3	00:11	34805	1.5	0.02	10.6	89.4	7.1	92.9

Table 4: Liquidity usage, delay and netting of Lynx using the iterative solution. This table presents the averages of the statistics across the days in each sample. The first column is the multiples (from 1 to 2) of the minimum initial liquidity for T2 payments plus the UPM liquidity. The T2 minimum initial liquidity was calculated starting from a liquidity equivalent to double the MaxASO and iteratively varied for each participant until all payments are settled intraday. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary across rows). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day, which by construction is zero. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns, labelled Settled intraday, are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq.	MaxNDP	EoD liq.	Delay	Delayed	Settled EoD		Settled intraday			
Billion $	Billion $	Billion $	hh:mm	payments	Value	Volume	Value		Volume	
				(number)			gross	net	gross	net
Random sample										
11.4	20.5	0.0	00:08	503	0.0	0.0	96.1	3.9	99.94	0.06
13.7	21.6	0.0	00:05	342	0.0	0.0	97.6	2.4	99.97	0.03
15.9	22.5	0.0	00:04	250	0.0	0.0	98.7	1.3	99.98	0.02
18.2	23.2	0.0	00:02	180	0.0	0.0	99.1	0.9	99.99	0.01
20.5	23.8	0.0	00:01	131	0.0	0.0	99.5	0.5	99.99	0.01
22.8	24.1	0.0	00:01	86	0.0	0.0	99.7	0.3	100	0
High-value sample										
12.6	28.9	0.0	00:06	697	0.0	0.0	96.7	3.3	99.95	0.05
15.1	30.2	0.0	00:04	462	0.0	0.0	98.1	1.9	99.97	0.03
17.6	31.2	0.0	00:03	320	0.0	0.0	98.8	1.2	99.99	0.01
20.1	31.8	0.0	00:02	224	0.0	0.0	99.2	0.8	99.99	0.01
22.6	32.3	0.0	00:01	169	0.0	0.0	99.5	0.5	99.99	0.01
25.1	32.7	0.0	00:01	102	0.0	0.0	99.6	0.4	100	0

Table 5: Liquidity usage, delay and netting of Lynx Straight-to-Queue using the iterative solution. This table presents the averages of the statistics across the days in each sample. The first column is the multiples (from 1 to 2) of the minimum initial liquidity for T2 payments plus the UPM liquidity. The T2 minimum initial liquidity was calculated starting from a liquidity equivalent to double the MaxASO and iteratively varied for each participant until all payments are settled intraday. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary across rows). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day, which by construction is zero. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns, labelled Settled intraday, are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq.	MaxNDP	EoD liq.	Delay	Delayed	Settled EoD		Settled intraday			
Billion $	Billion $	Billion $	hh:mm	payments	Value	Volume	Value		Volume	
				(number)			gross	net	gross	net
Random sample										
11.3	20.2	0.0	00:10	26808	0.0	0.0	4.1	95.9	0.8	99.2
13.6	21.3	0.0	00:08	26808	0.0	0.0	2.9	97.1	0.5	99.5
15.8	22.1	0.0	00:06	26808	0.0	0.0	2.0	98.0	0.3	99.7
18.1	22.8	0.0	00:05	26808	0.0	0.0	1.2	98.8	0.2	99.8
20.3	23.3	0.0	00:04	26808	0.0	0.0	0.8	99.2	0.2	99.8
22.6	23.7	0.0	00:03	26808	0.0	0.0	0.5	99.5	0.1	99.9
High-value sample										
12.4	28.5	0.0	00:08	34814	0.0	0.0	3.4	96.6	0.8	99.2
14.8	29.7	0.0	00:06	34814	0.0	0.0	2.3	97.7	0.5	99.5
17.3	30.6	0.0	00:05	34814	0.0	0.0	1.6	98.4	0.3	99.7
19.8	31.3	0.0	00:04	34814	0.0	0.0	1.1	98.9	0.2	99.8
22.3	31.7	0.0	00:03	34814	0.0	0.0	0.6	99.4	0.1	99.9
24.7	32.4	0.0	00:03	34814	0.0	0.0	0.4	99.6	0.1	99.9

Table 6: Liquidity usage, delay and netting of Lynx FIFO. This table presents the averages of the statistics across the days in each sample. The first column are the multiples of MaxASO (from 1 to 2) used for the initial liquidity of the LOM mechanism used to settle all T2 payments plus the UPM liquidity. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary along the curve). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day. Total used liquidity is therefore the sum of columns 2 and 3. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns labelled Settled intraday are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq. Billion $	MaxNDP Billion $	EoD liq. Billion $	Delay hh:mm	Delayed payments (number)	Settled EoD Value	Settled EoD Volume	Settled intraday Value gross	Settled intraday Value net	Settled intraday Volume gross	Settled intraday Volume net
Random sample										
14.6	15.5	1.9	00:39	11,552	6.5	3.3	74.7	25.3	86.7	13.3
17.6	16.4	1.6	00:31	10,233	4.7	2.4	79.8	20.2	89.5	10.5
20.5	17.1	1.4	00:25	9,089	3.6	1.8	84.0	16.0	92.1	8.0
23.4	17.9	1.1	00:20	8,081	2.6	1.2	86.7	13.4	93.6	6.4
26.3	18.6	0.9	00:17	7,116	2.0	0.8	89.4	10.6	95.0	5.0
29.3	19.2	0.8	00:13	6,227	1.5	0.6	91.4	8.6	96.1	3.9
High-value sample										
15.2	22.5	2.7	00:36	16,259	7.3	3.2	70.6	29.4	83.4	16.6
18.3	23.4	2.4	00:28	14,429	5.3	2.2	75.9	24.2	86.9	13.1
21.3	24.3	2.1	00:23	12,767	4.2	1.6	81.0	19.0	90.3	9.7
24.4	25.1	1.8	00:18	11,285	3.1	1.1	84.7	15.3	92.6	7.4
27.4	25.9	1.6	00:15	10,012	2.4	0.9	87.8	12.2	94.3	5.7
30.5	26.7	1.3	00:12	8,862	1.9	0.7	90.3	9.7	95.6	4.4

Table 7: Liquidity usage, delay and netting of Lynx FIFO Straight-to-Queue. This table presents the averages of the statistics across the days in each sample. The first column are the alpha multiples of MaxASO (from 1 to 2) used for the initial liquidity of the LOM mechanism used to settle all T2 payments plus the UPM liquidity. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary with alpha). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day. Total used liquidity is therefore the sum of columns 2 and 3. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns labelled Settled intraday are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq.	MaxNDP	EoD liq.	Delay	Delayed	Settled EoD		Settled intraday			
Billion $	Billion $	Billion $	hh:mm	payments	Value	Volume	Value		Volume	
				(number)			gross	net	gross	net
Random sample										
14.6	15.4	1.9	00:41	26,802	6.5	3.3	8.1	91.9	7.8	92.2
17.6	16.3	1.6	00:33	26,802	4.7	2.4	6.5	93.5	7.0	93.0
20.5	17.0	1.4	00:27	26,802	3.6	1.8	5.6	94.4	6.6	93.4
23.4	17.8	1.1	00:23	26,802	2.6	1.2	4.7	95.3	6.1	94.0
26.3	18.5	0.9	00:19	26,801	2.0	0.8	4.2	95.8	5.8	94.2
29.3	19.1	0.8	00:16	26,801	1.5	0.6	3.8	96.2	5.7	94.3
High-value sample										
15.2	22.5	2.7	00:37	34,807	7.3	3.2	9.1	90.9	7.5	92.5
18.3	23.3	2.4	00:30	34,806	5.3	2.2	7.3	92.7	6.7	93.3
21.3	24.2	2.1	00:24	34,806	4.2	1.6	6.4	93.6	6.3	93.7
24.4	25.0	1.8	00:20	34,806	3.1	1.1	5.7	94.4	6.1	93.9
27.4	25.7	1.6	00:17	34,806	2.4	0.9	5.1	94.9	5.9	94.1
30.5	26.5	1.3	00:14	34,806	1.9	0.7	4.7	95.3	5.8	94.2

Table 8: Liquidity usage, delay and netting of Lynx FIFO using the iterative solution. This table presents the averages of the statistics across the days in each sample. The first column are the multiples (from 1 to 2) of the minimum initial liquidity for T2 payments plus the UPM liquidity. The T2 minimum initial liquidity was calculated starting from a liquidity equivalent to two times MaxASO and iteratively varied for each participant until all payments are settled intraday. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary across rows). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day which by construction is zero. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns labelled Settled intraday are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq.	MaxNDP	EoD liq.	Delay	Delayed	Settled EoD		Settled intraday			
Billion $	Billion $	Billion $	hh:mm	payments	Value	Volume	Value		Volume	
				(number)			gross	net	gross	net
Random sample										
11.5	20.6	0.0	00:08	4,492	0.0	0.0	95.2	4.8	98.0	2.1
13.8	21.8	0.0	00:05	3,290	0.0	0.0	97.4	2.7	99.0	1.0
16.1	22.6	0.0	00:03	2,456	0.0	0.0	98.3	1.7	99.3	0.7
18.4	23.3	0.0	00:02	1,814	0.0	0.0	99.0	1.0	99.7	0.3
20.7	23.8	0.0	00:01	1,343	0.0	0.0	99.5	0.5	99.9	0.1
23.0	24.2	0.0	00:00	973	0.0	0.0	99.7	0.3	99.9	0.1
High-value sample										
12.7	29.0	0.0	00:06	5,563	0.0	0.0	96.0	4.1	98.4	1.6
15.2	30.3	0.0	00:03	3,995	0.0	0.0	97.8	2.2	99.2	0.8
17.8	31.2	0.0	00:02	2,929	0.0	0.0	98.7	1.3	99.5	0.5
20.3	31.8	0.0	00:01	2,240	0.0	0.0	99.3	0.8	99.8	0.2
22.8	32.4	0.0	00:01	1,668	0.0	0.0	99.4	0.6	99.8	0.2
25.4	32.7	0.0	00:00	1,184	0.0	0.0	99.6	0.4	99.9	0.1

Table 9: Liquidity usage, delay and netting of Lynx FIFO Straight-to-Queue using the iterative solution. This table presents the averages of the statistics across the days in each sample. The first column are the multiples (from 1 to 2) of the minimum initial liquidity for T2 payments plus the UPM liquidity. The T2 minimum initial liquidity was calculated starting from a liquidity equivalent to two times MaxASO and iteratively varied for each participant until all payments are settled intraday. All T1 and BoC payments are settled without delay in the UPM mechanism. The second column is the sum of the max net debit position in the UPM and max net debit position in the LOM. The average max net debit position in the UPM is $10.8 and $17.7 billion in the random and high-value samples respectively (which do not vary across rows). EoD liq. is the liquidity required at the end of the day to settle all payments that remained in the queue at the end of the day which by construction is zero. The columns labelled Settled EoD are percentages of payments not settled intraday as a share of the payments sent to the LOM. The last four columns labelled Settled intraday are the percentages of the payments sent to the LOM settled on a gross or net basis before 18:00 hrs.

Initial liq. Billion $	MaxNDP Billion $	EoD liq. Billion $	Delay hh:mm	Delayed payments (number)	Settled EoD Value	Settled EoD Volume	Settled intraday Value gross	Settled intraday Value net	Settled intraday Volume gross	Settled intraday Volume net
Random sample										
11.5	20.4	0.0	00:10	26,815	0.0	0.0	0.0	100.0	0.0	100.0
13.8	21.5	0.0	00:07	26,816	0.0	0.0	0.0	100.0	0.0	100.0
16.1	22.3	0.0	00:06	26,816	0.0	0.0	0.0	100.0	0.0	100.0
18.4	22.9	0.0	00:04	26,816	0.0	0.0	0.0	100.0	0.0	100.0
20.7	23.3	0.0	00:03	26,816	0.0	0.0	0.0	100.0	0.0	100.0
23.0	23.7	0.0	00:03	26,816	0.0	0.0	0.0	100.0	0.0	100.0
High-value sample										
12.6	28.8	0.0	00:08	34,824	0.0	0.0	0.0	100.0	0.0	100.0
15.1	30.0	0.0	00:06	34,824	0.0	0.0	0.0	100.0	0.0	100.0
17.7	30.8	0.0	00:05	34,824	0.0	0.0	0.0	100.0	0.0	100.0
20.2	31.4	0.0	00:04	34,824	0.0	0.0	0.0	100.0	0.0	100.0
22.7	31.9	0.0	00:03	34,824	0.0	0.0	0.0	100.0	0.0	100.0
25.2	32.2	0.0	00:03	34,824	0.0	0.0	0.0	100.0	0.0	100.0

Figure 6: Delay and liquidity of Lynx FIFO, Lynx FIFO Straight-to-Queue, RTGS and RTGS with LSM. This plot shows the liquidity and delay tradeoff of the Lynx systems with FIFO configuration. In green is the random day sample and in red the highest value days. Samples are of 114 days. Delay time is the average of the value-weighted delay. The order of payment submissions is the same in all exercises. The dashed and dashed-dot lines are Lynx and Lynx Straight-to-Queue. In these simulations, participants initial liquidity is a multiple of the MaxASO. The top point of each curve is equal to 1xMaxASO and the bottom point is 2xMaxASO. In these two systems, some payments are delayed all day but settled calculating a round of multilateral netting at the end of the day. In Lynx Straight-to-Queue payments are submitted immediately to the queue. The solid lines are the RTGS system which represent minimum amounts of liquidity and resulting delay without any LSM. The dotted lines are the RTGS with LSM, which adds the Jumbo algorithm of LVTS. Delayed payments in a first pass can settle before the end of the day via this netting algorithm. Additional details are reported in Tables 6 and 7.

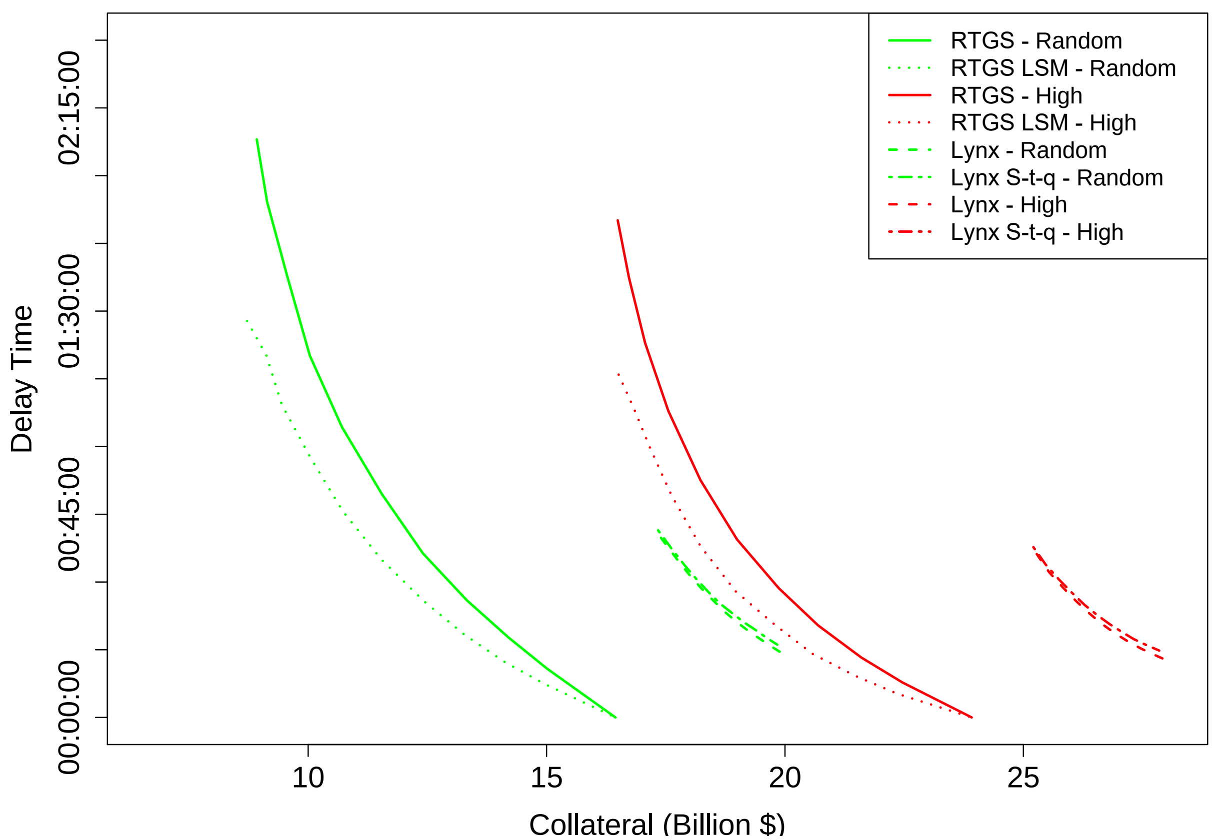

Figure 7: Delay and liquidity of Lynx FIFO and Lynx FIFO Straight-to-Queue (with the iterative method), RTGS and RTGS with LSM. This plot shows the liquidity and delay tradeoff of the Lynx systems with FIFO configuration. In green is the random day sample and in red the highest value days. Samples are of 114 days. Delay time is the average of the value-weighted delay. The order of payment submissions is the same in all exercises. The dashed and dashed-dot lines are Lynx and Lynx Straight-to-Queue. In these simulations, participants initial liquidity is the minimum required to settle all payments intraday. In Lynx Straight-to-Queue payments are submitted immediately to the queue. The solid lines are the RTGS system which represent minimum amounts of liquidity and resulting delay without any LSM. The dotted lines are the RTGS with LSM, which adds the Jumbo algorithm of LVTS. Delayed payments in a first pass can settle before the end of the day via this netting algorithm. Additional details are reported in Tables 8 and 9.

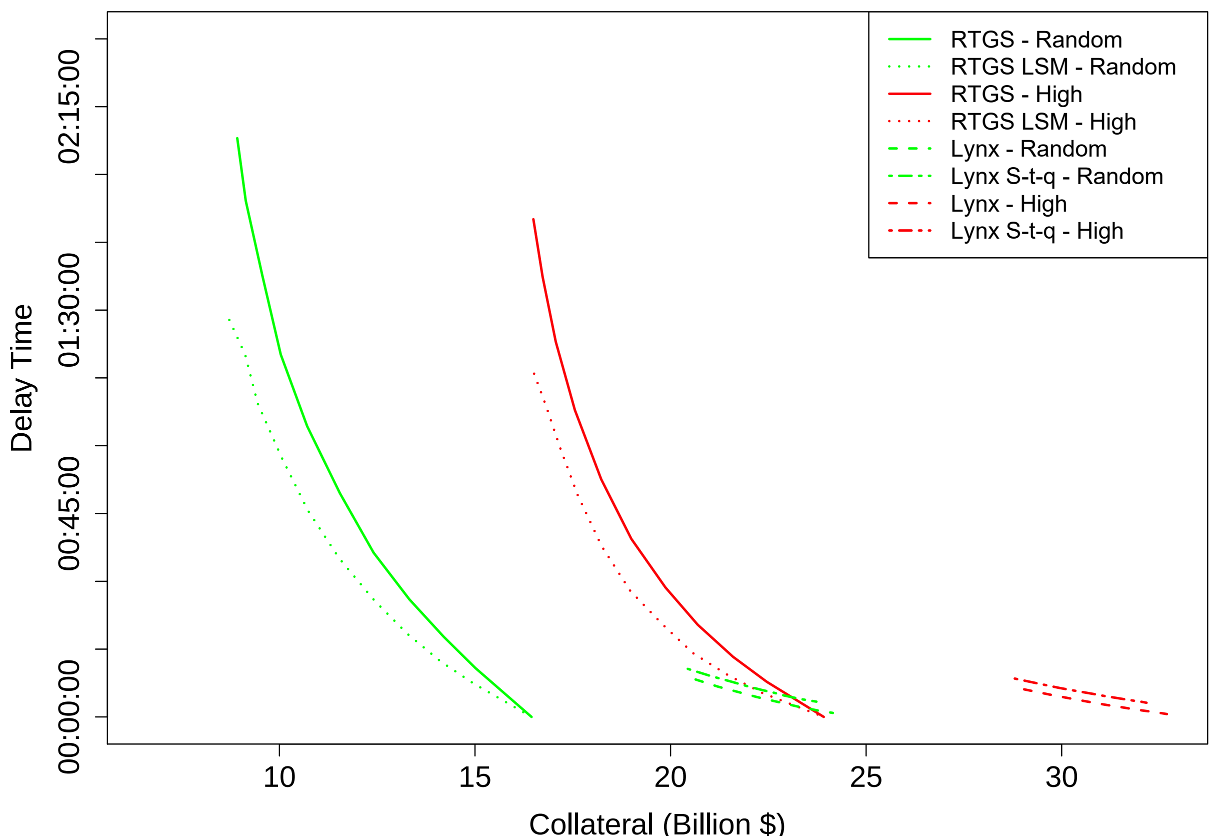

Figure 8: Delay and liquidity of Lynx FIFO and Lynx FIFO Straight-to-Queue with the MaxASO and iterative method simulations. This plot shows the curves in Figures 6 and 7 for ease of comparison. The bottom four curves are the ones using the iterative method. The upper four curves use MaxASO as initial liquidity.

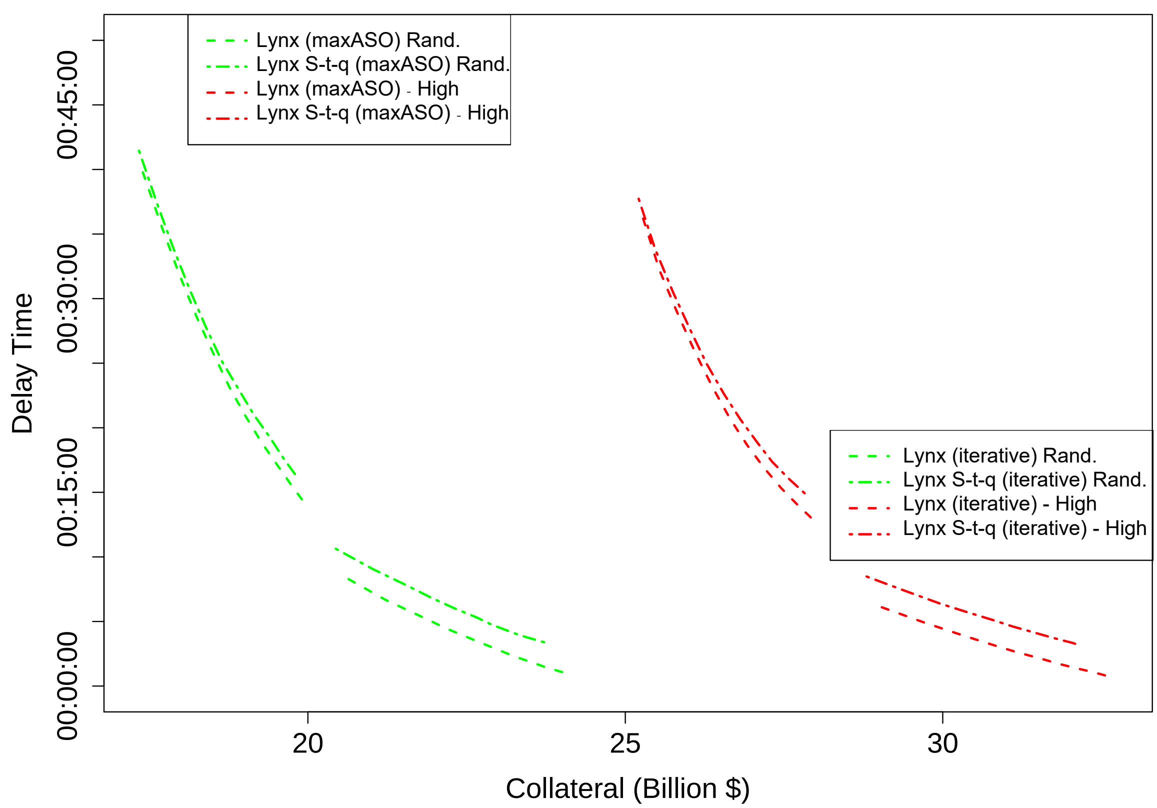

Figure 9: Lynx liquidity usage by participants in the random sample. This is the distribution of the LOM maxNDP in Lynx. Top shows the exercise when initial liquidity is two times MaxASO. Bottom is the exercise when initial liquidity is calculated iteratively.

(a) Initial liquidity 2xMaxASO

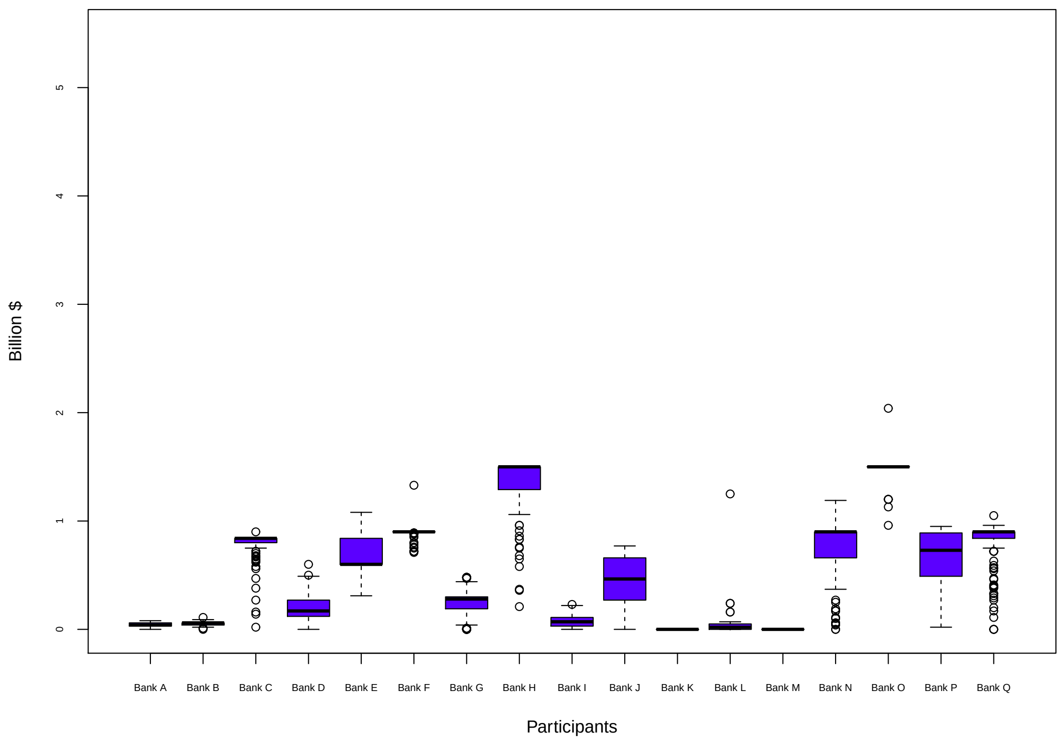

(b) Iterative method

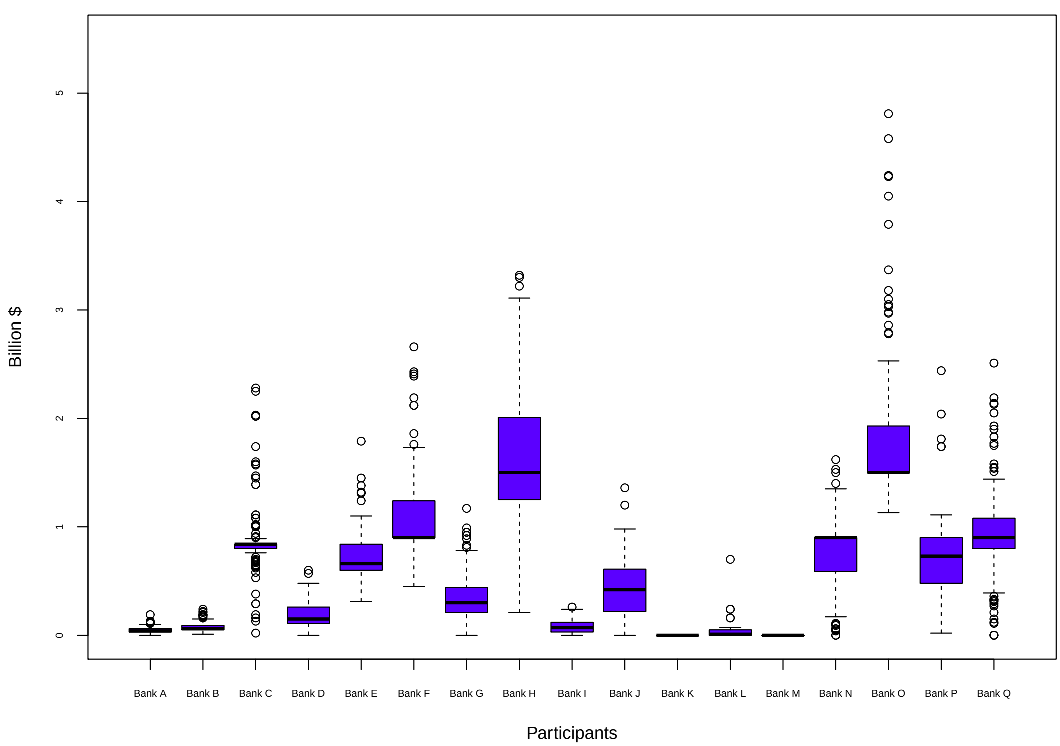

Figure 10: Lynx liquidity usage by participants in the high-value sample. This is the distribution of the LOM maxNDP in Lynx. Top shows the exercise when initial liquidity is two times MaxASO. Bottom is the exercise when initial liquidity is calculated iteratively.

(a) Initial liquidity 2xMaxASO

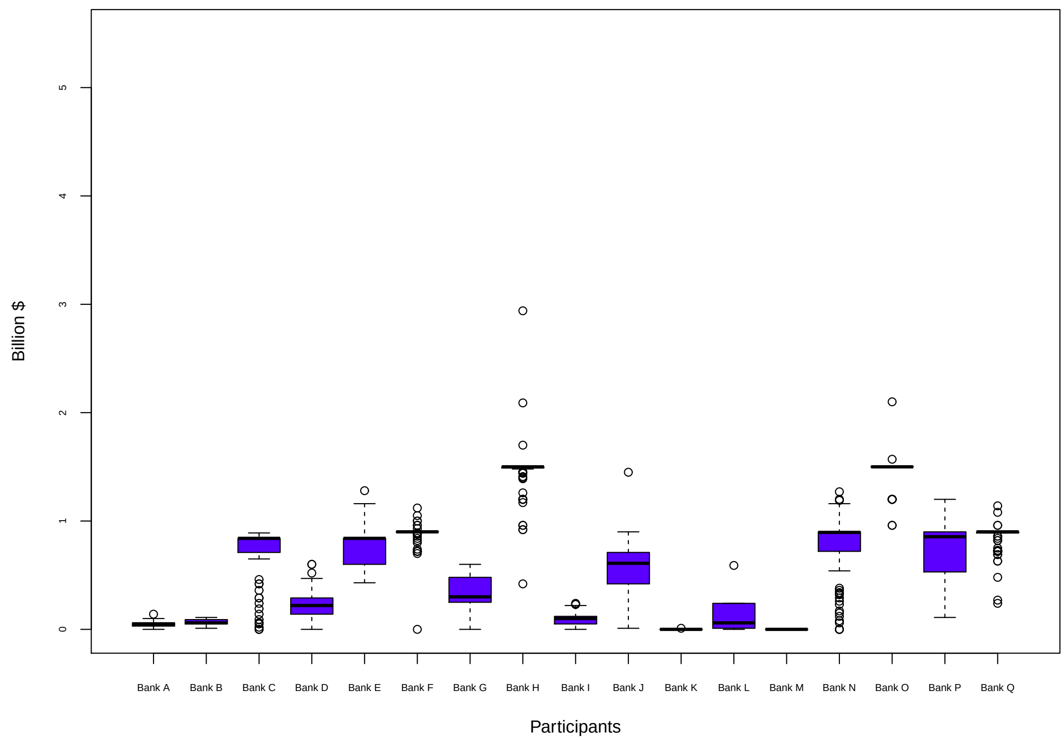

(b) Iterative method

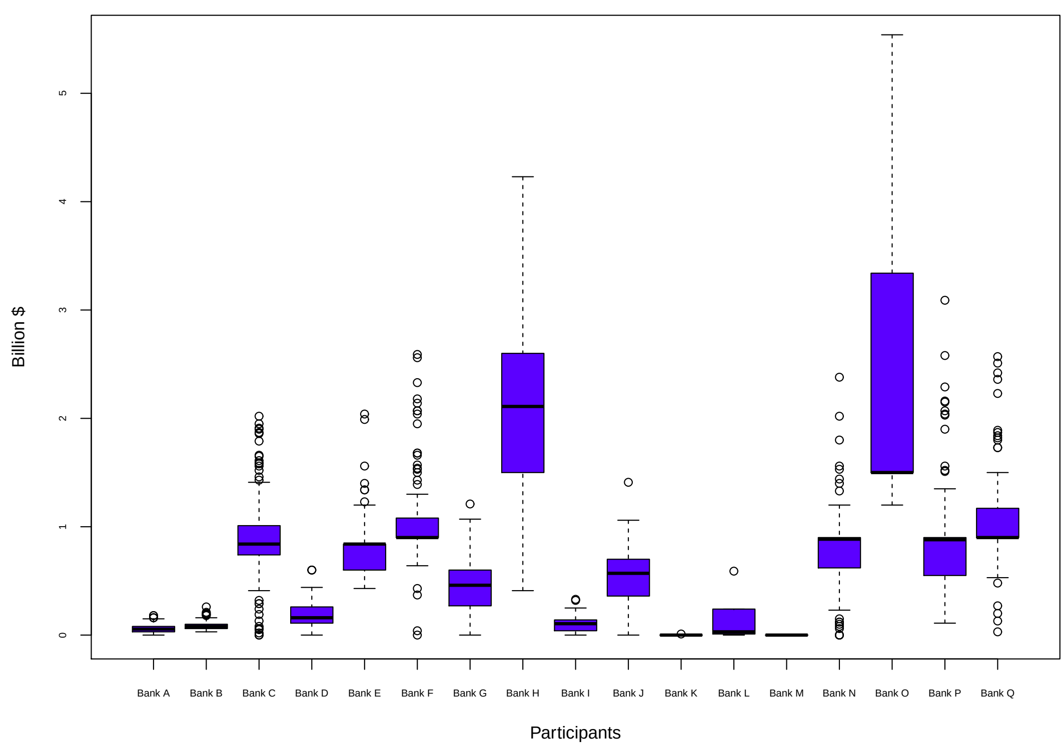

Figure 11: Percentiles of Payment Size in LVTS from June 2003 to March 2019. Percentiles of payment size in log scale. The series show a downward trend in payment size for the 50^{th}, 75^{th}, 90^{th} and 95^{th} percentiles. A sharp decline occurred during the 2008–2010 financial crisis. The fact that the mean values are much higher than the 90^{th} percentiles shows that the majority of LVTS payments are small in value. Note that the median payment has almost fallen one order of magnitude in size.

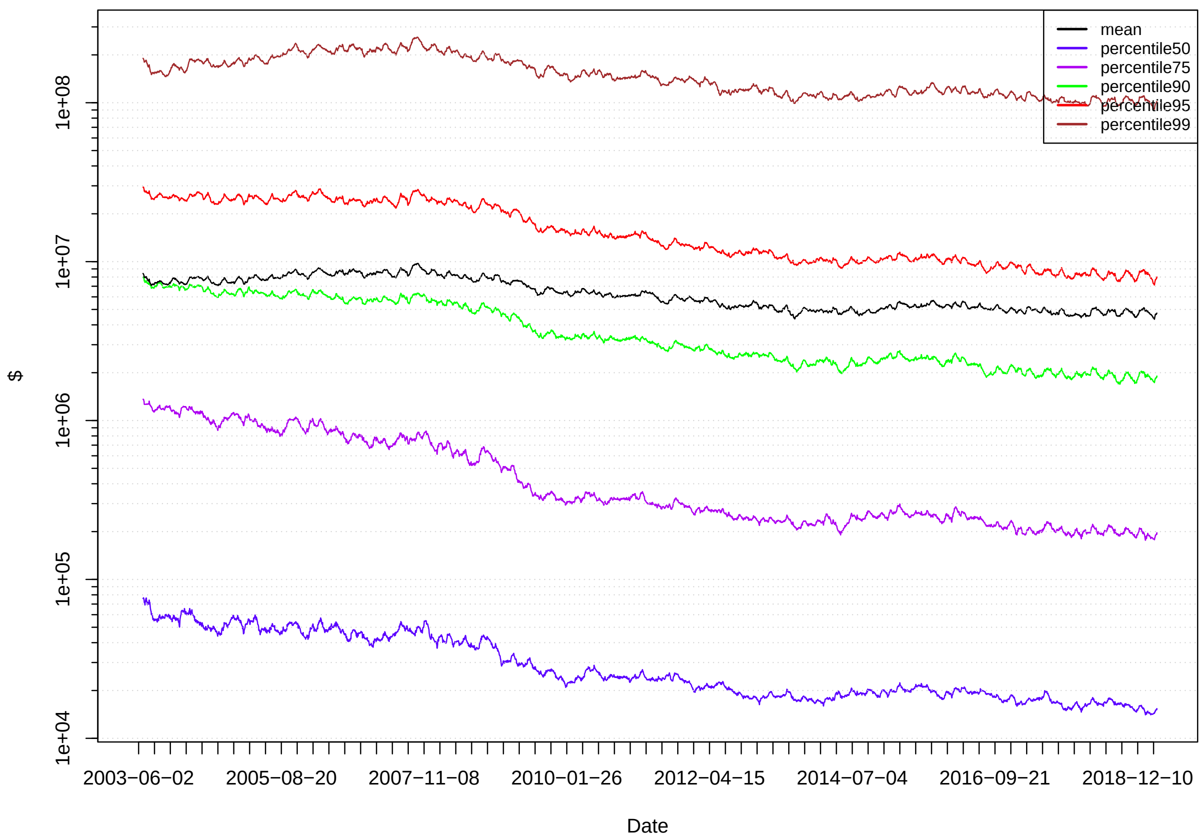

Appendix A Samples

The statistics are computed for the following samples of historical LVTS data. There are 17 participants in LVTS; on average, there are close to 35,000 payments every day. We choose to use 114 days of data, which is a 5 percent quantile of the entire sample, in each of the samples below. This allows us to have comparable standard errors of the calculations. The samples are the following.

- Normal: 114 random days in the entire sample. This sample should be the baseline for all statistics.
- High payment value: 114 days with volume in the top 5 percent quantile. These are days when LVTS processed larger than normal values. Typically, these days tend to stress systems, as they require more than normal liquidity provision.

Appendix B Methodology to calculate the liquidity-delay efficiency frontiers

Here we describe a summary of the methodology to calculate the liquidity-delay efficiency frontier. For more details see Rivadeneyra and Zhang (2020a). The procedure follows the next steps:

1. T1 and payments to and from the Bank of Canada should be settled without delay. To ensure this, we calculate the maximum net debit position for these payments.
2. For T2 payments, we perform a first-pass simulation, to calculate the maximum net debit position for these payments. Note that this is a different calculation than the previous step.
3. By summing the two T1 and T2 max net debit positions, we obtain the level of liquidity needed for a pure RTGS without any delay. This is the lower point in the RTGS curves in Figure 1.
4. To create the frontiers, we proportionally reduce the liquidity available for T2 payments. The liquidity for T1 and Bank of Canada payments is unchanged in these exercises. T2 payments settle in real time if liquidity is available. Otherwise, they will settle at the end of the day.

5. At the end of the day, the multilateral net debit positions are calculated. If needed, more liquidity is added to the T2 pool.

6. We included all payments from 12:00 a.m. to 6:30 p.m. The pre-settlement period, between 6:00 and 6:30 p.m. is included on a multilateral net basis.

Here are a few important observations:

- Although the collateral pools are calculated separately, the liquidity is shared between the two streams of payments.

- The assumption of end-of-day settlement is extreme, of course, but provides us with a consistent method to create the frontiers. With the extreme points, we can create mixes of these systems.