

Rivadeneira, Francisco; Zhang, Nellie

Working Paper

Payment coordination and liquidity efficiency in the new Canadian wholesale payments system

Bank of Canada Staff Discussion Paper, No. 2022-3

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Rivadeneira, Francisco; Zhang, Nellie (2022) : Payment coordination and liquidity efficiency in the new Canadian wholesale payments system, Bank of Canada Staff Discussion Paper, No. 2022-3, Bank of Canada, Ottawa, <https://doi.org/10.34989/sdp-2022-3>

This Version is available at:

<https://hdl.handle.net/10419/251033>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Payment Coordination and Liquidity Efficiency in the New Canadian Wholesale Payments System

by Francisco Rivadeneyra and Nellie Zhang

Banking and Payments Department
Bank of Canada

frivadeneyra@bankofcanada.ca, nzhang@bankofcanada.ca



Bank of Canada staff discussion papers are completed staff research studies on a wide variety of subjects relevant to central bank policy, produced independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

For their comments, we thank Shaun Byck, James Chapman, Rod Garratt, Nancy Goundry, Scott Kinnear, Anneke Kosse, Wade McMahon, David Millard and Kevin Wai. All errors remain our own.

Abstract

A new wholesale payments system will launch in Canada in 2021. This real-time gross settlement system called Lynx will have two types of settlement mechanisms, one allowing offsetting and the other not. This paper studies the decision problem of the Bank of Canada: which of the two settlement mechanisms should it use to send its payments. Using extensive simulation, we show that, mainly due to the benefits of liquidity pooling, Lynx would achieve its highest liquidity efficiency—even better than that of the current Large Value Transfer System (LVTS)—if all payments (urgent and non-urgent) from all participants were sent to the mechanism allowing offsetting. The minimum amount of liquidity required to settle all payments by critical deadlines is approximately \$10 billion, around half the amount of collateral that LVTS participants allocate (pre-COVID-19). Since time-critical payments sent to the offsetting mechanism could experience a delay, the high level of liquidity efficiency is accompanied by an increase in the number of participants' operational interventions (to pledge more collateral or to alter payment priorities) to ensure that those time-critical payments are never delayed. When coordination does not occur, liquidity efficiency can be far lower than in the LVTS. The results highlight that the Bank of Canada helping with coordination is more important than the specific choice of mechanism.

Topics: Payment, clearing and settlement systems

JEL codes: C, C5, E42, E58

1 Introduction

In 2021, Canada will launch a new wholesale payments system called Lynx.¹ The new system will be substantially different from the current Large Value Transfer System (LVTS). Lynx will be a real-time gross settlement (RTGS) system with two separate mechanisms: the liquidity-saving mechanism (LSM) and the urgent payment mechanism (UPM). In contrast, the LVTS is fundamentally a deferred net settlement (DNS) system that achieves real-time finality thanks to the residual guarantee offered by the Bank of Canada (see [Arjani and McVanel 2006](#)).

The transition from the LVTS to an RTGS-based system such as Lynx is a major one as it will change the incentives of financial institutions participating in the system, for example, their decisions on liquidity allocation and the timing of payment submission. As the institution in charge of ensuring the safety and efficiency of the system, the Bank of Canada has a unique interest in understanding how the new system can affect settlement outcomes like liquidity usage, payment delay and operational risks.

As an overseer of Lynx, the Bank of Canada is also a key direct participant in the system as the settlement agent for the Canadian federal government and other critical financial market infrastructures, such as the securities settlement system (CDSX) and the retail payments system (ACSS). Prior to the COVID-19 pandemic—for example, in 2019—the payment value that the Bank of Canada sent to participants was on average 11.58% of the system daily total, and most of these transactions were critical or time-sensitive.²

In this paper, we study the decision problem of the Bank of Canada: how should the Bank make use of the two settlement mechanisms in Lynx (the UPM and LSM) for its own payments? For example, should the Bank send a certain percentage of its payments through each settlement mechanism, or should it choose the mechanism by type of payment, or should there be some other decision rules? An overall guideline is that the Bank of Canada should base its decision on a choice that will yield the safest and most efficient outcome for the system.

This is a complex decision because, in any RTGS system, participants play an intraday liquidity-management game in which they trade off the cost of liquidity with the potential

¹Since 2014, Canada has been modernizing its core payments infrastructure, including wholesale and retail systems. The Bank of Canada (the Bank) has been working closely with Payments Canada and other stakeholders in the design and implementation of the new systems. For more details see [Chapman et al. \(2015\)](#).

²During the pandemic, due to the quantitative easing program, the share of payment value that the Bank of Canada sends to other LVTS participants has skyrocketed to 65%, underscoring the importance of the nation’s central bank in the large-value payments system.

cost of delay of their payments (Bech and Garratt 2003). The Bank of Canada, however, as the nation’s central bank, is not a strategic player in such games because it incurs no cost in creating liquidity and consequently has no incentives to delay any of its payments. Therefore, the Bank chooses settlement mechanisms as well as the timing of its payments simply to maximize the benefits of the system.

Since each settlement mechanism has its own separate pool of liquidity, the most important channel by which the choice of the Bank can affect the decisions of participants is in which mechanism a participant is expected to receive a payment from the Bank. Separate pools mean that the liquidity from an incoming payment can be used for an outgoing payment in the mechanism it was received. Otherwise, a participant would need to perform an operation to move liquidity between the two mechanisms. We call this an operational intervention because it can be time consuming and probably require human involvement. Another type of operational intervention is an increase of liquidity (in either mechanism) drawing from the pool of collateral that the participant has with the Bank.

We design three types of simulation scenarios to evaluate the effects of the Bank of Canada’s payment choices on the overall efficiency of the system’s liquidity, the extent of delay of payments and the number of required operational interventions. These scenarios are created based on combinations of the Bank of Canada’s choice and various possible actions that Lynx participants might take in response. Simulations of these scenarios are conducted using real LVTS data, and an average liquidity-delay efficiency frontier is constructed for each simulation case.³

Our main result is that Lynx would achieve its highest level of liquidity efficiency if every participant sent all payments (urgent and non-urgent) through the LSM. The average amount of intraday liquidity usage would range between \$9.1 billion and \$13.5 billion with an average weighted delay of payments varying between 9 and 24 minutes.⁴ The number of intraday operational interventions required to ensure urgent payments settle in time ranges between 5 and 34.

Similar results are observed for the simulation scenario where all payments including those sent by the Bank of Canada are exchanged in the UPM. The average amount of minimum intraday liquidity required would range between \$7.5 billion and \$12.4 billion, with an average weighted payment delay varying between 12 and 37 minutes. The number of intraday operational interventions would range between 7 and 55, and payment delay is

³See Rivadeneyra and Zhang 2019 for details of the methodology for constructing these frontiers.

⁴Liquidity usage is calculated as the sum of intraday maximum net debit positions incurred by each individual participant. It is a different measure than the amount of initial liquidity that participants choose to allocate at the beginning of the day.

marginally higher than in the simulation scenario where all transactions are conducted in the LSM. These results highlight one of our main propositions; that is, liquidity pooling can affect the efficiency of a payments system far more than the features of a settlement mechanism can, e.g., central-queuing functionality and payment-offsetting algorithms.

The results of other simulation exercises show a significantly higher demand for intraday liquidity, longer payment delays and more operational interventions. The scenario where the Bank of Canada chooses the UPM and other participants all respond by sending their urgent payments in the UPM shows a range of liquidity usage of between \$19 billion and \$23 billion and an average weighted payment delay of between 15 and 48 minutes. Altogether, the results of the imperfect-coordination simulation scenarios show the risk of the lack of coordination and highlight the importance of the Bank’s payments as a source of intraday liquidity for Lynx participants.

Overall, our results suggest that the Bank of Canada should announce and use the LSM for sending all its payments. In addition to this policy recommendation, our paper also contributes to the literature on payments system simulations by adding a behavioural component in our simulation design. Specifically, we model Lynx participants’ intraday payment and liquidity management by simulating their ongoing liquidity monitoring and operational intervention efforts to ensure all time-critical payments settle in time. We report in our results the number of intervention actions taken and the magnitude of such efforts represented by the dollar amounts of extra intraday collateral pledges.

Several previous studies have examined the trade-offs between liquidity efficiency and payment delay in the environment of payments systems. [Rivadeneira and Zhang 2020](#) performed the first simulation exercise on Lynx, using a large sample of historical LVTS payments data. In the paper, it is assumed that all Tranche 1 (T1) transactions take place in the UPM and all Tranche 2 (T2) payments are sent to the LSM. This assumption is motivated by the superficial similarity between T1 and the UPM, and between T2 and the LSM.⁵ In that study, the authors find that Lynx would require more intraday liquidity than the amount currently being allocated to the LVTS, which is confirmed in this paper and further explained by the fact that in the entire system of Lynx, including the LSM, net debit positions have to be fully collateralized.

A similar approach is taken by [Byck and Heijmans 2021](#). The authors use LVTS data in a simulated RTGS environment and calculate the liquidity-delay frontiers for multiple configurations of liquidity-saving mechanisms, for example, different queue-release

⁵In the LVTS, T1 is essentially an RTGS system where participants pledge collateral dollar-to-dollar for their net debit positions, whereas T2 is the true hybrid component that in essence is a DNS system featuring a unique risk-sharing model and a guarantee of real-time finality.

rules and different algorithms for bilateral and multilateral payment offsetting. The main findings also touch upon the importance of liquidity pooling.⁶

The rest of the paper proceeds as follows. Section 2 provides a brief overview of the new Canadian large-value payments system Lynx. Section 3 describes the simulation methodology used in this study, including the sample LVTS data, various experiment scenarios and the behaviour simulating method. Simulation results are presented in Section 4, followed by a discussion in Section 5. Section 6 concludes.

2 Overview of Lynx

Lynx is the new large-value payments system in Canada replacing the LVTS. Unlike the LVTS, which is an RTGS/DNS hybrid system, Lynx is a typical RTGS system equipped with liquidity-optimization algorithms. Payments can settle in one of two separate mechanisms: the UPM and the LSM.⁷ In principle, payments that require immediate settlement should be processed in the UPM, whereas non-urgent transactions should be submitted to the LSM. However, as demonstrated in our simulation exercises, the use of payment priorities and changes in allocated liquidity can be an alternative to ensure immediate settlement of a payment.

The UPM offers a simple queuing mechanism that releases queued payments on a strict first in, first out (FIFO) basis. A payment is added to the UPM queue if the sending participant does not have sufficient funds to settle this payment at the time of submission, and it is later released once funds are available. These UPM FIFO queues can be considered as centrally provided individual queues since there is neither multilateral nor bilateral offsetting between any queued payments.

By contrast, LSM has a central queuing system, which provides both multilateral and bilateral offsetting among all the queued payments. In the LSM, a payment settles immediately upon submission only if the sending participant has sufficient funds *and* there are no other payments of the same priority in the central queue; as long as the queue is not empty, any new payment arriving at the system is added and waits for offsetting opportunities.

In Lynx, the LSM also includes a “gridlock buster” algorithm triggered by certain changes of state in the system to find a batch of queued payments that can offset among

⁶For more on the previous simulation work regarding the LVTS, please see [Arjani 2006](#), which also explores the trade-off between settlement delay and intraday liquidity. More generally, [Martin and McAndrews 2008](#) and [Jurgilas and Martin 2013](#) examine the reasons why an LSM can increase efficiency in an RTGS.

⁷The Appendix discusses further details of the system.

themselves and settle concurrently. There is a time lag between the trigger and the gridlock buster algorithm runtime, and this interval is a system parameter that can be set between 5 and 60 minutes.⁸

Queued payments in the LSM can also settle through a process called impact intervention. Rather than offsetting payments, impact intervention checks each individual queued payment against the sending participant’s intraday liquidity. The process is triggered by any state of change in the LSM, such as an inter-account transfer of liquidity or an event of a participant receiving funds (either from incoming payments or from pledging more collateral).

The LSM offers two settlement sequences: strict FIFO or FIFO-bypass. Previous studies such as Rivadeneyra and Zhang (2020) show evidence that FIFO-bypass is a more efficient option in that it bypasses larger-value payments that fail the risk controls and allows smaller-value payments that meet the requirements to pass and settle first; faster liquidity recycling accelerates the entire settlement process in general.⁹

Within each settlement mechanism, every payment instruction in Lynx can contain a priority level from 1 to 99. At the time of this project, the proposed priority levels in Lynx are 1, 3, 5 and 99, with 1 being the highest priority and 99 being the lowest; the default priority level is set to 99, which means that all payments are assigned this base level if no higher priority is chosen.¹⁰

Priority 1 is reserved in Lynx as a “fast lane” for priority changes of urgent payments to settle at a critical moment. The idea is that there should never be any payments of priority 1 waiting in the central queue; therefore all payments that are increased to priority 1 should be able to settle at the moment of the priority change. Hence, the default highest priority in Lynx is proposed to be 3. In an urgent situation, Lynx participants are expected to adjust their intraday liquidity positions and increase the priority of urgent payments to 1, so that they can pass the risk controls and settle immediately.

In Lynx, payments of high priorities are always settled prior to those of low priorities. Within each priority, payment instructions in Lynx can be further prioritized by internal sequence numbers; all payments in Lynx sent by all participants at all times are ordered first by priority, second by internal sequence number, and subsequently by date and time.

⁸Studies conducted by Payments Canada show that a 5-minute interval for gridlock buster execution seems more efficient and optimal than other configurations.

⁹See Rivadeneyra and Zhang (2020) for additional details of the gridlock buster algorithm, impact intervention and FIFO-bypass settlement sequences.

¹⁰This priority system is different and separate from the current SWIFT message priority mechanism where payments are assigned as either urgent or normal.

3 Simulation methodology

We make use of our Lynx simulation environment, which replicates to the best of our knowledge the functioning of Lynx algorithms. Our simulator provides a good working model for Lynx since more than 98% of the transactions in the sample we tested are settled within one minute of the settlement time for the same payment in the vendor testing program (i.e., the actual system that will be in production). Therefore, we have a high degree of confidence in the simulation results.

The exercises are designed to examine the liquidity coordination effects from the Bank of Canada’s choice of settlement mechanism and the responses of participants. In terms of settlement sequence in each mechanism, the UPM provides an individual FIFO queue with neither bypass nor offsetting opportunities for each participant, whereas in the LSM there are two choices: strict FIFO or FIFO with bypass. And in this paper, we use the FIFO-bypass for the LSM configuration.

3.1 Payment data

To carry out the simulation in this study, we use sample LVTS data that span from February 7, 2019 to July 31, 2019, containing a total of 122 business days. During this period, around 39,173 transactions worth \$186 billion settled on average each day. At the time of the study, the sample period was chosen because it was the only time frame during which the High-Availability Banking System (HABS) data was available for analysis.¹¹

In the LVTS, there are no priority settings for payments, and therefore we do not have observable data on participants’ choices of payment priorities. In order to simulate the Lynx functionalities with priorities, we assign priority levels to the existing LVTS data based on our best knowledge of the nature of each payment. Using the HABS data, we are able to identify the different types of transactions, such as the payments for ACSS settlement, CDSX settlement, afternoon auctions of Receiver General (RG) account balances and bank note withdrawal settlement. Each transaction type is associated with an operational deadline by which payments for this particular transaction must settle.

In general, the critical payments sent to the Bank of Canada need to be completed before 3:30 p.m. because the Bank’s books close at that time; anything received after 3:30, with the exception of the PM RG auction, CDSX settlement and bank notes settlement, would be considered late and returned to the sending participant. In the simulation, all

¹¹HABS is the Bank of Canada’s back office system that participants use to pledge and allocate collateral and manage their intraday liquidity positions in the LVTS.

Table 1: Assigned priorities for LVTS transactions

Transaction type	Deadline	Priority assigned
ACSS	12:00:00	3
BoC outgoing	18:00:00	1
CDCC	07:45:00	3
CDIC	15:30:00	5
CDS entitlement	15:30:00	5
CDSX settle	17:00:00	3
RG PM auction	17:00:00	3
NES settle	16:00:00	3
CLS	CLS pay-in schedules	3
CLS same-day	14:00:00	5
CMHC	15:30:00	5
Foreign	15:30:00	5
Government	15:30:00	5
Other	15:30:00	99

the payments sent by the Bank of Canada are set to be priority 1 (P1) by convention.¹² All other payments by participants are assigned priorities 3 (P3), 5 (P5) or 99 (P99) at the time of submission, depending on their time sensitivity. Table 1 lists the assigned priorities for the LVTS transactions. We assign the highest standard priority (P3) to payments for the settlement of other financial market infrastructures as they are considered time critical. These payments are made for the settlement of the Automated Clearing and Settlement System (ACSS), Canadian Derivatives Clearing Corporation (CDCC), Canadian securities depository service (CDSX), CLS Bank payments, Note Exchange System (NES) and the Receiver General PM auction. Other payments dealing with federal government transactions—such as the Canada Mortgage and Housing Corporation (CMHC)—and foreign clients are marked P5. The rest of the payments, including the vast majority that are currently settling through T2, are defaulted to P99.

Table 2 reports the summary statistics of the LVTS payments for each priority we assign. In our proposed categorization, more than 99% of the T2 payments in the LVTS are assigned P99, and about 80% of T1 payments are marked as urgent transactions of P1 and P3. On average, less than 2% of LVTS payments in volume are labelled P1, P3 and

¹²As noted earlier, for other participants, P1 should always be reserved for emergency use, i.e., when a payment in a lower priority needs to be settled immediately. By altering its priority to P1, the sending participant essentially brings this urgent payment from the back of the queue to the front. In order for such a set-up to work, the P1 queue should always be kept empty. The Bank of Canada does not need to obey this rule because it can, in principle, generate unlimited liquidity, and its payments are never queued.

Table 2: Summary statistics of LVTS transactions with assigned priorities (%)

Priority	T1 (%)	T2 (%)	Transaction volume		Transaction value	
			Daily average (billion \$)	(%)	Daily average (billion \$)	(%)
1	64.68	0.00	254	0.65	22.79	12.24
3	16.04	0.13	114	0.29	9.07	4.88
5	11.48	0.31	165	0.42	13.47	7.24
99	7.80	99.55	38639	98.64	140.78	75.64

P5, whereas 24% of the transaction value falls into these three relatively urgent categories. This confirms the commonly held intuition that time-critical payments tend to be very large in value.

The distribution of the number and size of time-critical payments is especially relevant to the efficiency of queuing algorithms with bypass functionality, as these algorithms, like the LSM in Lynx, would tend to delay large payments more than small ones. Some of these large payments are time-sensitive; therefore, the algorithm might create additional need for operational interventions by the participant to avoid delaying critical payments beyond their cut-off time. On the other hand, the potential benefit of a queuing algorithm with a bypass option is that a larger total value of payments might be settled earlier than otherwise because finding offsetting opportunities is easier with a large number of smaller-value payments. This pattern does emerge when we compare the UPM and the LSM as the former does not have a bypass functionality.

In the simulation exercises, the timing of payment submissions in Lynx is assumed to be the same as the observed submission time in the LVTS. In the LVTS, there is no delay between submission and settlement with the exception of jumbo payments, which are very few.

3.2 Scenarios and behavioural assumptions

The Bank of Canada has two choices in regards to making payments in Lynx: (i) sending all of its payments in the UPM or (ii) sending all its payments through the LSM. It is indifferent between these two options as it does not need the liquidity from incoming transactions in order to send its own payments. Since the Bank is not a strategic player, we design three types of simulation scenarios based on the effects of the Bank’s choice and the participants’ numerous behaviours in response.

The first type are one-pool scenarios in which both the Bank of Canada and the Lynx participants coordinate all payments in either the UPM or the LSM. The second type

are tit-for-tat strategies where participants respond to the Bank of Canada’s choice by sending their urgent payments to the same mechanism and randomize the choice of mechanisms for the non-urgent. The third type are called non-coordinated because in these scenarios participants choose not to coordinate payments with the Bank of Canada. It is an unlikely situation; however, the results help paint a complete picture of the participants’ behavioural responses and demonstrate how crucial a role liquidity coordination plays in a payments system. We define urgent payments as those to the Bank of Canada as well as others marked P3 and P5.

These scenarios are designed to capture three incentives of participants: (i) minimizing liquidity costs, (ii) avoiding the cost of delay of time-sensitive payments and (iii) minimizing the number of operational interventions to ensure critical payments settle in time. Since by design the payments of the Bank of Canada are never delayed past their submission time regardless of the choice of mechanism, we assume other participants care about the choice of the Bank because it affects their ability to recycle liquidity in either mechanism or their need to intervene in the settlement operations to avoid delay of any time-sensitive payments.

In total, we propose the following seven scenarios:

1. **One-pool:** all payments are processed in one mechanism without exception.
 - All LSM: Participants and the Bank send all payments in the LSM.
 - All UPM: Participants and the Bank send all payments in the UPM.
2. **Tit-for-tat for urgent payments:** participants respond to the Bank’s choice of mechanism by sending all their urgent payments in the same mechanism.
 - Urgent in UPM: The Bank of Canada sends all payments in the UPM. All other participants send their urgent payments in the UPM, while directing the rest to the LSM. This scenario somewhat resembles the dichotomy observed in the LVTS today, where the Bank uses almost exclusively T1 and the participants also restrict their transactions with the Bank and other urgent payments to T1.
 - Urgent in LSM: The Bank of Canada sends all payments in the LSM. Other participants send their urgent payments in the LSM; the rest of the payments are sent randomly through either UPM or LSM. This scenario does not exactly mirror the scenario immediately above, because it is not sensible to assume that participants would send all of their non-urgent payments through UPM.

Table 3: Simulation exercises The first two columns are *one-pool* scenarios where the Bank and other Lynx direct participants send all their payments either in the UPM or LSM. The next two columns are the *tit-for-tat* scenarios where the Lynx participants respond to the Bank of Canada’s choice of settlement mechanism by sending their urgent payments in the same mechanism. The column labelled *Cross* is the scenario where participants do not coordinate with the Bank, but rather send their urgent payments in the UPM while receiving incoming payments from the Bank in the LSM. The last column lists the *non-coordinated* scenarios where Lynx participants randomize their choice of settlement mechanisms regardless of the Bank of Canada’s decision.

		One-pool		Tit-for-tat		Cross	Non-coordinated
BoC	Others	All LSM	All UPM	Urgent: UPM Rest: LSM	Urgent: LSM Rest: Random	Urgent: UPM Rest: LSM	All random
	All LSM	✓	n.a	n.a.	✓	✓	✓
All UPM	n.a	✓	✓ (LVTS-like)	n.a.	n.a.	n.a.	✓

3. **Non-coordinated:** Lynx participants do not coordinate with the choice of the Bank.

- Cross: The Bank of Canada sends all payments in the LSM, but other participants do the opposite, sending their urgent payments in the UPM, and the rest in the LSM. This scenario is known as *cross* because all the urgent payments in this case are sent before receipt of the free liquidity in the form of incoming payments from the Bank of Canada.
- BoC using LSM: The Bank sends all payments in the LSM, and other participants randomize their payments through either the UPM or the LSM.
- BoC using UPM: The Bank sends all payments in the UPM, and other participants randomize their payments through either the UPM or the LSM.

Table 3 summarizes the set of exercises we perform. Simulation results for the *one-pool* scenarios are reported in Table 4. Results for the *tit-for-tat* scenarios are described in Table 5. Results for the *non-coordinated* and *cross* scenarios are detailed in tables 6 and 7, respectively.

3.3 Initial liquidity and liquidity-delay frontiers

The level of initial liquidity in each mechanism is an important assumption for the simulation exercises. In the LSM, we assume that, at the beginning of the day, participants allocate the amount of liquidity that is equal to a multiple of their maximum additional settlement obligation (MaxASO) for that day as recorded in the LVTS data. The choice

of MaxASO is a natural one in this scenario since it is the amount of collateral that LVTS participants currently pledge in T2.

In the UPM, it is assumed that participants pledge no collateral at the beginning of the day. It may seem extreme; however, it does not impose a strong constraint overall on liquidity consumption.¹³ This is because in the simulation, as explained below, to avoid payments being delayed past their critical time, participants make intraday pledges of collateral whenever there is a shortage of funds. At the end of the day, these pledges achieve the same effect as if Lynx participants had pledged at the beginning of the day the exact amount of liquidity needed to settle all payments without any delay beyond their associated deadlines.

With this set-up, we ensure that the level of initial liquidity is the same for all the simulation scenarios even if, in some scenarios, only one of the two settlement mechanisms was used; thus, the simulation results are readily comparable across all scenarios.

Finally, in order to examine the trade-off between liquidity usage and settlement delay, for each defined scenario we repeat the simulation nine times at various levels of initial liquidity ranging from $0.5\times$ to $2.4\times$ of MaxASO. Specifically, we carry out 63 simulations in total, and each simulation yields a statistic of an average liquidity consumption, paired with the corresponding settlement delay; nine pairs of such statistics are then used to trace out an efficiency-delay frontier curve for each simulation scenario.

3.4 Simulation of behaviour

Incorporating participant behaviour has been a challenge in simulation exercises of payments systems. In this study, the focus of the simulation design is to make sure no payment is delayed past its critical deadline.¹⁴ To achieve this, we assume that Lynx participants take one or both of the following actions, if necessary, when a payment is still in the queue as the clock strikes the deadline: (i) escalating the priority level of this payment to 1 (by construction, there are no other payments in this priority level), and (ii) increasing the intraday collateral pledge by an amount that will pass the risk controls immediately upon the change in liquidity position. We evaluate the system efficiency

¹³The basic idea is that Lynx participants might use UPM only to a limited extent and instead direct most payments to the LSM for the offsetting opportunities; thus they may pledge low levels of collateral at the beginning of the day in UPM. Key questions are what might be the appropriate starting level of liquidity in UPM and how low is too low. For the lack of a better assumption, we resort to the assumption of zero.

¹⁴The critical times are obtained from the current LVTS operation timelines. For an example, the ACSS settlement pay-in time is 11:30 a.m.; hence, all transactions associated with ACSS settlement each have a deadline of 11:30. Critical times will be maintained once Lynx is in operation.

across different simulation scenarios by measuring how often Lynx participants have to engage in the intervention activities and how much extra collateral they need to pledge intraday.

Although the actual Lynx system will provide participants with flexibility to make inter-account transfers of liquidity between different settlement mechanisms, for simplicity, we do not model such behaviour in this paper.¹⁵ In the current exercises the UPM and the LSM are simulated individually with separate liquidity pools; thus intraday liquidity recycling is contained within each mechanism. In other words, in the simulations of this study a Lynx participant cannot use the liquidity received in the UPM to send payments in the LSM.

4 Results

We present our results in the categories of simulation scenarios. For each scenario, we report statistics of liquidity (initial allocation versus additional intraday pledges), the measure of settlement delay (value-weighted average delay), the amount of operational interventions (the number and the percentage of total payments that settle with the help of intervention), and the scales of liquidity efficiency (the liquidity efficiency ratio and the percentage of payments settled under five minutes from submission).

Due to the behaviour component of deadline control, every transaction settles by the end of day in the simulation; in other words, there are no rejected payments. Hence, the differences among the results across various simulation scenarios mainly reside in the measurement of settlement delay and the minimum quantity of intraday liquidity needed to settle the same set of transactions. We measure settlement delay by the average value-weighted delay, which is calculated as the summation of each payment value multiplied by the difference between its settlement time and submission time, and then divided by the total payment value of all delayed payments. This measure penalizes the high-value payments more than the low-value payments for the same amount of delay, the idea being that in general the consequence of delaying a high-value payment is supposedly more severe than that of delaying a low-value payment. Note that the delay measures are the averages over the delayed payments only, not over all transactions.

The minimum amount of liquidity needed to settle all the transactions during the day is the sum of the intraday maximum net debit positions (MaxNDP) of all Lynx participants. When there is only one settlement mechanism involved, it is simple to use

¹⁵This is an area to potentially explore in our future work.

the total MaxNDP of that mechanism as the measure. In the simulation scenarios where both the UPM and the LSM are utilized, the MaxNDP is computed independently for each settlement mechanism and then the sum of the two MaxNDPs is considered as the minimum amount of liquidity required in the entire system.¹⁶

We also report the number of operational interventions that are necessary to ensure the settlement of all payments on time. The number of interventions is reflected in statistics such as the amount of additional collateral pledged intraday and the number of times that participants have to allocate extra collateral and/or increase the priority ranking of the payments. These actions are counted separately, but they are closely related. Our results show that, in the vast majority of the cases, Lynx participants need to do both to settle an urgent payment. Occasionally, a payment is delayed only by it being in the lower rank and blocked by the higher-priority payments, not because of the lack of sufficient liquidity. The amount of extra collateral pledged is measured both in billion dollars and as a percentage of daily total value of all payments sent.

The liquidity efficiency ratio (LER) is defined as, on any given day, the ratio of the MaxNDP and the total transaction value of the system. It indicates the minimum liquidity requirement for sending every dollar of payments on average. A low LER points to a high efficiency of liquidity usage.

As defined in Section 3.2, each simulation type contains 2–3 different scenarios. The results of all scenarios are presented in this section, with the exception of the *Cross* scenario, which is discussed in Section 5. As shown in tables 4, 5 and 6, Lynx achieves higher efficiency if all payments are settled in one mechanism instead of two, regardless of which mechanism is used. In both scenarios of *All UPM* and *All LSM*, at every level of initial liquidity compared with all cases of other scenarios, we observe smaller average delays in settlement, lower liquidity requirements and fewer intraday operational interventions. This is not surprising; the results suggest that the benefits of liquidity pooling, which promotes liquidity recycling, is crucial to the efficiency of payments settlement—so significantly in fact that they seem to surpass any possible advantages that LSMs can offer.

If Lynx participants pledge an amount of collateral equal to MaxASO at the beginning of the day, i.e., the amount that they are currently allocating to T2 in the LVTS, and all the participants including the Bank of Canada send payments solely through the LSM,

¹⁶Note that this is quite different from moving along the timeline and choosing the maximum of the sum of the MaxNDP at any point of time, which is very likely a quantity smaller than the sum of max; the only time the two quantities are equal is when the two MaxNDPs peak at exactly the same time in each settlement mechanism.

Table 4: Liquidity usage, delay and intraday operational interventions for the one-pool scenarios This table presents the averages of various statistics over the 122 days in the data sample. The first column lists the multiples of MaxASO (ranging from 0.5 to 2.4) used as the initial amount of collateral that Lynx participants allocate at the beginning of the day to the settlement mechanisms in use. The second column is the minimum amount of liquidity required in the simulation to settle all the transactions by critical timelines; it is calculated as the sum of the MaxNDP of every participant in each settlement mechanism. The extra intraday collateral pledges are reported both in billion dollars and as a percentage of the total transaction value settled on each day. Likewise, the priority increases are recorded both in the total number of increases each day and as a percentage of the total transaction volume settled on that day. The liquidity efficiency ratio (LER) is defined as the minimum liquidity required for each dollar of payment settled. The last column lists the percentage of transactions in volume that are settled within five minutes of the submission time. The numbers in red indicate the results for the simulation case where Lynx participants pledge the collateral in the exact amount of MaxASO at the beginning of the day. The results in blue text are of the simulation case where the initial collateral pledge is equal to double the amount of MaxASO.

Initial liq. (billion \$)	Min liq. (billion \$)	Delay (h:mm:ss)	Extra collateral pledge (billion \$)	(%)	Priority increases (volume)	(%)	LER	Settled ≤ 5 min
One-Pool: All LSM								
2.61	9.12	0:23:40	6.61	0.04	34	0.09	0.06	97.45
5.23	10.09	0:18:39	5.25	0.03	19	0.05	0.06	98.22
6.28	10.59	0:16:46	4.83	0.03	15	0.04	0.07	98.48
7.33	11.07	0:15:10	4.44	0.03	11	0.03	0.07	98.69
8.37	11.60	0:13:37	4.10	0.02	10	0.03	0.07	98.86
9.41	12.04	0:12:12	3.73	0.02	8	0.02	0.08	99.01
10.46	12.62	0:10:52	3.52	0.02	7	0.02	0.08	99.13
11.51	13.05	0:09:47	3.18	0.02	6	0.02	0.08	99.24
12.56	13.54	0:08:51	2.90	0.02	5	0.01	0.09	99.34
One-Pool: All UPM								
2.61	7.48	0:37:20	4.94	0.03	55	0.14	0.05	56.34
5.23	8.61	0:28:10	3.69	0.02	28	0.07	0.05	64.87
6.28	9.10	0:24:51	3.26	0.02	23	0.06	0.06	68.11
7.33	9.63	0:22:04	2.90	0.02	19	0.05	0.06	70.90
8.37	10.18	0:19:34	2.56	0.02	16	0.04	0.06	73.64
9.41	10.73	0:17:22	2.27	0.01	12	0.03	0.07	75.95
10.46	11.26	0:15:28	1.99	0.01	10	0.03	0.07	78.08
11.51	11.83	0:13:42	1.78	0.01	8	0.02	0.08	80.21
12.56	12.39	0:12:07	1.60	0.01	7	0.02	0.08	82.17

then the settlement delay on average is around 18 minutes, and an extra \$5.25 billion of liquidity is needed intraday to ensure all payments settle by the designated deadlines. The minimum liquidity requirement in this *One-pool: All LSM* scenario is \$10.09 billion.

The scenario of *One-pool: All UPM* represents a pure RTGS system without central queuing. Compared with the simulation of *All LSM*, the results show a noticeably longer average settlement delay and larger numbers of incidents of intraday operational intervention. At all levels of initial liquidity, the minimum liquidity requirement is only slightly smaller, which might simply be an outcome of the trade-off between settlement cost and delay. Hence, one may argue that settling all payments in LSM is more efficient than utilizing UPM without any benefit of central queuing.

One noteworthy aspect of the results in comparing the two *one-pool* scenarios is that, at a given level of settlement delay (within the boundaries of critical deadlines), the amount of minimum liquidity to settle all payments without delay beyond critical deadlines does not vary much, no matter which single settlement mechanism is used and how much beginning-of-day liquidity the system starts with. If we endure an average settlement delay of around 18 minutes, perhaps arguably acceptable, then the minimum amount of liquidity required by Lynx will be around \$10 billion, that is, approximately two times of MaxASO.

Due to segmentation of liquidity pools between the LSM and the UPM, the results of two-pool scenarios show significant increases in all measures of system inefficiency. Specifically, we observe that if both settlement mechanisms are used and the Bank of Canada chooses to send all its payments through the LSM, the average settlement delay exceeds 2.5 hours in all cases at various levels of initial liquidity available; the minimum amount of liquidity required to settle every payment in time is found to be between \$20 and \$30 billion, and each day, Lynx participants would need to step in during the settlement process and increase the priority ranking of payments on average between 900 and 1100 times. The LER found in the two-pool scenarios are more than double the LER in the single-pool scenarios.

Among all the simulations under the *two-collateral-pool* type, however, there is a clear distinction in the results between the scenarios where the Bank of Canada sends its payments through the LSM and those where it chooses to use the UPM. It is interesting to note that Lynx seems to operate more efficiently when the Bank uses the UPM for sending payments, no matter what behavioural responses are carried out by other Lynx direct participants. Under the *tit-for-tat* type of scenario, the minimum amounts of liquidity to settle all payments in time are quite similar; however, the average delay for the cases where Lynx participants receive incoming payments from the Bank of Canada

Table 5: Liquidity usage, delay and intraday operational interventions for the tit-for-tat scenarios This table presents the averages of various statistics over the 122 days in the data sample. The first column lists the multiples of MaxASO (ranging from 0.5 to 2.4) used as the initial amount of collateral that Lynx participants allocate at the beginning of the day to the settlement mechanisms in use. The second column is the minimum amount of liquidity required in the simulation to settle all the transactions by critical timelines; it is calculated as the sum of the MaxNDP of every participant in each settlement mechanism. The extra intraday collateral pledges are reported both in billion dollars and as a percentage of the total transaction value settled on each day. Likewise, the priority increases are recorded both in the total number of increases each day and as a percentage of the total transaction volume settled on that day. The liquidity efficiency ratio (LER) is defined as the minimum liquidity required for each dollar of payment settled. The last column lists the percentage of transactions in volume that are settled within five minutes of the submission time. The numbers in red indicate the results for the simulation case where Lynx participants pledge the collateral in the exact amount of MaxASO at the beginning of day. The results in blue are of the simulation case where the initial collateral pledge is equal to double the amount of MaxASO.

Initial liq. (billion \$)	Min liq. (billion \$)	Delay (h:mm:ss)	Extra collateral pledge (billion \$)	(%)	Priority increases (volume)	(%)	LER	Settled ≤ 5 min
Tit-for-tat: Urgent in UPM								
2.61	19.28	0:48:32	16.72	0.10	129	0.34	0.12	95.90
5.23	20.17	0:33:13	15.27	0.09	124	0.32	0.13	97.60
6.28	20.59	0:29:46	14.79	0.09	123	0.32	0.13	98.01
7.33	21.00	0:26:23	14.33	0.09	122	0.32	0.13	98.26
8.37	21.49	0:23:38	13.98	0.09	120	0.31	0.14	98.47
9.41	21.93	0:21:19	13.59	0.08	119	0.31	0.14	98.68
10.46	22.38	0:19:26	13.29	0.08	119	0.31	0.14	98.80
11.51	22.83	0:17:24	13.00	0.08	117	0.30	0.14	98.94
12.56	23.21	0:15:50	12.70	0.08	117	0.30	0.15	99.04
Tit-for-tat: Urgent in LSM								
2.61	20.41	2:46:46	17.98	0.11	974	2.57	0.13	49.46
5.23	20.92	2:43:06	16.24	0.10	956	2.53	0.13	49.78
6.28	21.09	2:41:53	15.59	0.10	951	2.51	0.13	49.88
7.33	21.31	2:40:41	15.04	0.09	946	2.50	0.13	49.95
8.37	21.56	2:39:38	14.55	0.09	945	2.50	0.14	50.00
9.41	21.86	2:38:38	14.12	0.09	942	2.49	0.14	50.05
10.46	22.11	2:37:44	13.70	0.08	940	2.49	0.14	50.10
11.51	22.37	2:37:00	13.33	0.08	939	2.48	0.14	50.14
12.56	22.59	2:36:18	12.97	0.08	938	2.48	0.14	50.17

Table 6: Liquidity usage, delay and intraday operational interventions for the non-coordinated scenarios This table presents the averages of various statistics over the 122 days in the data sample. The first column lists the multiples of MaxASO (ranging from 0.5 to 2.4) used as the initial amount of collateral that Lynx participants allocate at the beginning of the day to the settlement mechanisms in use. The second column is the minimum amount of liquidity required in the simulation to settle all the transactions by critical timelines; it is calculated as the sum of MaxNDP of every participant in each settlement mechanism. The extra intraday collateral pledges are reported both in billion dollars and as a percentage of the total transaction value settled on each day. Likewise, the priority increases are recorded both in the total number of increases each day and as a percentage of the total transaction volume settled on that day. The liquidity efficiency ratio (LER) is defined as the minimum liquidity required for each dollar of payment settled. The last column lists the percentage of transactions in volume that are settled within five minutes of the submission time. The numbers in red indicate the results for the simulation case where Lynx participants pledge the collateral in the exact amount of MaxASO at the beginning of the day. The results in blue are of the simulation case where the initial collateral pledge is equal to double the amount of MaxASO.

Initial liq. (billion \$)	Min liq. (billion \$)	Delay (h:mm:ss)	Extra collateral pledge (billion \$)	(%)	Priority increases (volume)	(%)	LER	Settled ≤ 5 min
Non-coordinated: BoC using LSM								
2.61	25.91	2:45:42	23.64	0.15	1073	2.83	0.16	49.45
5.23	26.90	2:42:32	22.67	0.14	1067	2.82	0.17	49.70
6.28	27.21	2:41:33	22.33	0.14	1066	2.81	0.17	49.77
7.33	27.54	2:40:42	22.04	0.14	1065	2.81	0.17	49.81
8.37	27.83	2:39:57	21.76	0.14	1064	2.81	0.17	49.84
9.41	28.15	2:39:19	21.57	0.13	1063	2.81	0.18	49.88
10.46	28.39	2:38:44	21.36	0.13	1063	2.80	0.18	49.89
11.51	28.61	2:38:15	21.19	0.13	1062	2.80	0.18	49.91
12.56	28.85	2:37:51	21.06	0.13	1062	2.80	0.18	49.94
Non-coordinated: BoC using UPM								
2.61	25.26	1:29:23	22.70	0.14	326	0.84	0.16	78.63
5.23	25.07	1:00:06	20.12	0.12	311	0.80	0.16	80.26
6.28	25.02	0:53:39	19.18	0.12	308	0.79	0.16	80.55
7.33	24.88	0:48:29	18.18	0.11	305	0.79	0.16	80.81
8.37	24.84	0:44:06	17.34	0.11	302	0.78	0.16	80.98
9.41	24.86	0:40:59	16.57	0.10	300	0.77	0.16	81.13
10.46	24.90	0:38:00	15.88	0.10	297	0.77	0.16	81.27
11.51	24.94	0:35:51	15.19	0.09	295	0.76	0.16	81.36
12.56	25.04	0:33:53	14.60	0.09	294	0.76	0.16	81.44

through the UPM is well under 50 minutes, whereas when it sends payments through the LSM, the delays are all over two hours. A similar contrast can be found in the volume of intraday operational interventions; the number of increases in priority ranking to settle time-critical payments are around 950 on average per day in the *urgent in UPM* scenario versus well under 150 in the *urgent in LSM* scenario. The *all random* scenario shows the same distinction. This result suggests that, when both settlement mechanisms are used, the transactions sent by the Bank of Canada are a tremendously important source of liquidity for other Lynx direct participants, more so in the UPM than in the LSM.

In addition, there is also a noticeable contrast in the simulation results between the *tit-for-tat* and *non-coordinated* scenarios overall, which suggests that coordination in liquidity flows is of great importance to the efficiency of a payments system.


Figure 1 illustrates these simulation results more effectively in the space of liquidity usage versus settlement delay. Each frontier represents the most efficient outcome possible of each simulation; all the points lying above the frontier are considered less efficient for that particular scenario, because at any given level of settlement delay, a higher point above the curve marks a larger liquidity usage, and vice versa. Combinations of the minimum liquidity requirement and average settlement delay on each frontier curve construct the best solutions to this trade-off problem. That is, the closer the efficiency frontier is located to the origin of the chart, the more efficient a settlement outcome it represents. The figure shows that the two simulation scenarios of the *one-pool* type are found at the bottom left, closest to the origin, and are the most efficient of all scenarios.

The frontiers clearly show the trade-off between liquidity demand and settlement delay, within and across scenarios. There is a clear shift in collateral requirement as we move from one scenario to another. The *one-pool* type of scenarios consumes the least liquidity, roughly ranging from \$8 billion to \$13 billion. The *tit-for-tat* type of scenarios lies in the middle, showing that if the Bank of Canada and other Lynx participants coordinate urgent payments in either settlement mechanism, UPM or LSM, regardless of the difference in settlement delay, the minimum amount of liquidity needed to settle all payments in a timely manner is between \$19 billion and \$24 billion. In the *non-coordinated* type, which lacks any intended payment coordination, the liquidity demand is the highest among the three scenarios, all above \$25 billion in all cases.¹⁷

The two high-delay scenarios, the blue and purple lines at the top in Figure 1, are those where Bank of Canada chooses the LSM to send its payments. It suggests that

¹⁷The caveat is that we did not simulate the functionality where Lynx participants are able to move collateral between the UPM and the LSM. If such intraday re-allocation of liquidity is allowed in the simulation, then the differences would have been smaller.

Figure 1: Liquidity usage versus settlement delay (six simulation scenarios)



The *Cross* scenario is not included in this graph because there is no variation between liquidity usage and settlement delay. Please see Section 5 and Table 7 for details.

as long as there is non-negligible transaction volume in the UPM, the Bank of Canada sending payments in the LSM can result in a large amount of settlement delay in the overall system; the delay is most likely driven by the high-value payments in the UPM being held up in the strict FIFO queues, while at the same time blocking other, smaller-value payments from settling. The evidence shows that even if the Bank of Canada and all other Lynx direct participants coordinate urgent payments in the LSM (the blue line in Figure 1), the lack of incoming liquidity from the payments sent by the Bank has a detrimental effect on the settlement in the UPM. This point is further highlighted by the contrast between the two *tit-for-tat* scenarios (green and blue lines in the graph); coordinating important transactions—e.g., payments of P3 and P5—in one settlement mechanism helps to a certain extent, especially in terms of liquidity usage, but where the coordination takes place also matters.

Table 2 shows that on average P3 and P5 payments together account for 12% of the daily total transaction value. This does not seem a great deal; however, our simulation exercises indicate that it is a non-negligible amount and that where and how it settles in

Lynx makes a significant difference in system efficiency.

We also examine the results at the participants' level, and Figure 2 shows that in the *All LSM* scenario some Lynx participants might require relatively more intraday interventions than others. These are the total number of interventions across the time of day during the 122 days of the sample period, not an average.


Intraday interventions include two types of actions, increasing priority rankings and allocating extra collateral in the cases of shortage of liquidity; at different times, participants may need to act on only one of them or on both, to make sure all their payments settle in time. This plot demonstrates, for each Lynx participant, the number of seconds in a day when either type of intervention is performed; in other words, if multiple actions are needed for one or more than one transaction at exactly the same second, then these operations are counted as one. It is a conservative measure compared with the total number of interventions operated on every transaction needing help.

The figure illustrates that some Lynx participants are more operationally active than others at particular times. At a given time, the number of actions taken by one participant is stacked on top of the count of every other. This shows that in the early morning between 3:00 a.m. and 6:00 a.m., Bank C is quite busy making sure its CLS payments settle on time, whereas participant B dominates in the intervention operations during the hour of 2:00 a.m. Many more participants are found to be actively involved in the settlement process at noon, 2:00 p.m. and 5:00 p.m., relative to other times of day.

The amount of intraday interventions is further examined within each priority ranking. Detailed results are presented in tables 8, 9 and 10 in the Appendix. The numbers provide the same overview of the differences in settlement efficiency among all the simulation scenarios as discussed previously; the two scenarios in the one-pool type are significantly more efficient than any in the double-mechanism types, indicating far fewer intraday interventions in every priority category. When both settlement mechanisms are used, the scenarios where the Bank of Canada chooses the LSM to send its payments are remarkably worse than others, showing an enormous amount of intraday operational interventions needed to settle payments on time.

In the two scenarios where only one settlement mechanism is used, the number of interventions decreases in the descending order of the priority ranking. This is expected, because in Lynx payments of higher priority are always settled before the lower-priority payments; if liquidity is recycled and reused efficiently in the system, then it seems that the liquidity gained from the settlement of higher-priority transactions should be sufficient to fund the payments ranked lower. The fact that both scenarios in the one-pool type demonstrate this decreasing trend in the need of intraday interventions as settlement

Figure 2: Intraday operational interventions by participants (all LSM scenarios)



moves from high to low priorities is a sign of efficiency in liquidity usage.

The three tables of intraday operational interventions in each priority category provide unique insights into how settlement delay may occur in the proposed Lynx system and why liquidity cost is notably higher in some simulation scenarios than in others. For example, in the two-mechanism scenario, a large number of P99 payments needed extra help to settle because P99 payments are the least time-critical transactions among all and therefore assigned 6:00 p.m. as the deadline for settlement. The fact that a sizeable amount of P99 payments are delayed until the end of day and forced to settle at the same second indicates that the offsetting opportunities are somewhat lost in this case. They are also much more limited than in the preferred scenario where the settlement of these payments is scattered throughout the day. In the preferred scenario, liquidity from settling early submitted payments can potentially fund the payments that come into the system at a later time. In this case, most of the delayed P99 payments will have to settle individually, which undoubtedly increases the amount of overall liquidity required to settle every transaction before the close of the system.¹⁸

¹⁸Regardless of the reduction in batch settlement, in theory some end-of-day payments are still able to offset each other despite occurring at the same second. This is because impact intervention will still be

5 Discussion

We also carry out a simulation with an extreme and fairly unlikely scenario where the Bank of Canada and other Lynx direct participants send payments in opposite mechanisms. Specifically, in this *cross* scenario, in spite of being aware that the Bank of Canada sends all its payments in the LSM, other participants direct all their urgent payments (P3 and P5) in the UPM, the rest in the LSM.

This is the worst possible scenario, because in this case, Lynx participants cannot utilize the free liquidity (i.e., the incoming payments from the Bank of Canada) for making their urgent payments; the individual FIFO queues in the UPM are congested most of the day with payments waiting for settlement, and on the other hand, liquidity is found saturated in the LSM, a situation to definitely avoid in real operations. This exceedingly inefficient outcome is illustrated in Table 7.

The weighted average settlement delay in this scenario is as high as over five hours, and the minimum demand on intraday liquidity is as high as nearly \$32 billion. The LER is 0.2, comparatively higher than the ratios in all other scenarios, and the percentage of payments that settle within five minutes of submission is virtually zero, precisely 0.12%. Also notably, the simulation results in this scenario do not vary in response to different levels of initial collateral allocated to the LSM. This suggests that payment delays as well as the high demand for intraday liquidity all occur in the UPM, whereas half of the MaxASO is already more than sufficient in the LSM for the amount of transactions conducted there.¹⁹ The extreme inefficiency uncovered in this scenario partially results from the assumption in the simulation that Lynx participants allocate no collateral at the beginning of the day to the UPM. However, it would not change the fact that the liquidity provided by the incoming payments from the Bank of Canada turns out to be useless in the LSM.

Overall, the results in these simulation exercises suggest much higher efficiency in settlement if all participants predominantly use the LSM. It is also important to notice that, first, most of the efficiency gains in terms of liquidity demand and settlement delay do not originate from a particular choice of mechanism, but from payment coordination, liquidity pooling and recycling within one single operational domain. Second, such efficiency gains assume that Lynx participants have sufficient operational capabilities to ensure timely

launched on the payments that settle earlier than others within the same second; however, this internal order of settlement might not be the optimal order for maximizing the offsetting opportunities.

¹⁹This is further confirmed by examining the simulation results at the participant level, which is not presented in this paper. In this scenario, for every Lynx participant, the maximum net debit position in the LSM is zero.

Table 7: Liquidity usage, delay and intraday operational interventions for the cross scenario

This table presents the averages of various statistics over the 122 days in the data sample. The first column lists the multiples of MaxASO (ranging from 0.5 to 2.4) used as the initial amount of collateral that Lynx participants allocate at the beginning of the day to the settlement mechanisms in use. The second column is the minimum amount of liquidity required in the simulation to settle all the transactions by critical timelines; it is calculated as the sum of the MaxNDP of every participant in each settlement mechanism. The extra intraday collateral pledges are reported both in billion dollars and as a percentage of the total transaction value settled on each day. Likewise, the priority increases are recorded both in the total number of increases each day and as a percentage of the total transaction volume settled on that day. The liquidity efficiency ratio (LER) is defined as the minimum liquidity required for each dollar of payment settled. The last column lists the percentage of transactions in volume that are settled within five minutes of the submission time. The numbers in red indicate the results for the simulation case where Lynx participants pledge the collateral in the exact amount of MaxASO at the beginning of the day. The results in blue are of the simulation case where the initial collateral pledge is equal to double the amount of MaxASO.

Cross scenario: BoC using LSM, participants using UPM for urgent payments								
Initial liq. (billion \$)	Min liq. (billion \$)	Delay (h:mm:ss)	Extra collateral pledge (billion \$)	(%)	Priority increases (volume)	(%)	LER	Settled ≤ 5 min
2.61	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
5.23	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
6.28	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
7.33	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
8.37	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
9.41	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
10.46	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
11.51	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12
12.56	31.87	05:11:31	31.87	0.2	1411.72	3.77	0.2	0.12

settlement of urgent payments, and also require that the Bank of Canada and Payments Canada monitor the system to make sure there is no excessive usage of the UPM while the Bank of Canada sends payments in the LSM. In addition, enforcing certain throughput rules might be particularly useful because, compared with participants in the LVTS, Lynx participants will have different incentives for submitting payments earlier rather than later: Lynx is a relatively more costly system than the LVTS.

There are two important caveats to the results. The estimate of liquidity usage in one-pool scenarios should be considered as a minimum requirement. On the other hand, the two-mechanism scenarios do not consider potential improvements in efficiency from liquidity transfers between mechanisms. Future work could test the robustness of these simulation results. In particular, the potential impact of a decline in payment coordination among participants, possibly induced by the incentive to delay in payment submission, can be examined further. Another interesting subject for future study is the distributional impact of the FIFO-bypass option in the queue-release sequence on payments and/or participants.

6 Conclusion

In this paper, we conduct a variety of simulation exercises to study which of the two settlement mechanisms of Lynx the Bank of Canada should use to send its payments. First, our results show that both the Bank of Canada and all other Lynx participants should use only one settlement mechanism, instead of two in tandem, to take the most advantage of liquidity pooling.

Second, of the two mechanisms, the LSM is relatively more efficient than the UPM due to its central-queuing functionality and liquidity-saving algorithm. In particular, if all Lynx payments (urgent and non-urgent) are submitted and settled in the LSM, Lynx would achieve its highest efficiency characterized by the smallest settlement delay, lowest minimum liquidity requirement and fewest intraday operational interventions needed to settle transactions by their critical timelines.

When all transactions take place in the LSM, the Lynx system requires a minimum of around \$10 billion in liquidity with an associated average value-weighted delay of 19 minutes. This amount of liquidity is approximately twice the amount of the collateral that participants currently pledge in the LVTS. As to the system-wide delay of 19 minutes, one can argue that it is reasonable and acceptable, given that it is a worse-case measure weighted by the transaction value.

Third, the free liquidity injected through the payments sent by the Bank of Canada proves to be far more important in a system that does not offer central queuing and payment offsetting such as the UPM, than in a mechanism equipped with liquidity-saving features such as the LSM. In this study, we discover that, if both Lynx settlement mechanisms are used at the same time, the Bank of Canada and other participants should at least coordinate urgent payments in the UPM; the settlement outcome is evidently more efficient than the simulation cases where such coordination takes place in the LSM.

Fourth, the evidence further suggests that, if both the UPM and LSM are open for operations, and if the Bank of Canada has decided to send all its payments through the LSM, then it might be prudent to proactively monitor the participants' relative usage of the UPM. If a considerable amount of transactions are indeed settling through the UPM, Lynx could wind up far less efficient than otherwise; in this situation, it would actually be better for the Bank of Canada to send payments through the UPM, instead of the LSM.

Appendix

Additional features of the Lynx system

The settlement process discussed in the main body of the paper is the general settlement process in Lynx. In addition to the general process, Lynx offers a pre-settlement process called the conditional release mechanism (CRM) where payments settle only when some pre-defined conditions are met. The CRM is mainly designed for business validations, such as checking whether the system is open. Participants can use this additional functionality to store and release future-dated payments. Every payment instruction, or simply a payment, arriving at Lynx must first satisfy all the conditions in the CRM prior to being released to the general settlement process. If a payment fails to pass any one of the conditions, it is diarized and stays in the CRM warehouse until all the conditions are met. Since the CRM is unlikely to be utilized at the initial stages after the new system goes live, we exclude it from the simulation.

Recall that the general settlement process discussed in the body of the paper has two algorithms, the UPM and the LSM. In addition to these two mechanisms, there is a third one called the real-time settlement mechanism. This mechanism is as a storage of intraday liquidity where participants can reserve pools of funds and allocate liquidity to the UPM and LSM.

Intraday interventions within each priority ranking

Table 8: Intraday payment management for the one-pool scenarios This table presents how many times operational intervention actions are taken each day for all payments of different priorities. The numbers reported are the averages over 122 days in the data sample. For each priority group, the first column shows the numbers of priority increases; the second column shows the numbers of extra collateral pledged intraday to ensure all transactions settle by critical timelines; and the third column lists the numbers of seconds during the day when there is any action of operational intervention. In this data sample, there are on average 115 payments of priority 3 (P3), 166 payments of priority 5 (P5) and 38,639 payments of priority 99 (P99) every day.

Initial liq. (billion \$)	Priority 3			Priority 5			Priority 99		
	Pri. inc.	Col. inc.	Secs.	Pri. inc.	Col. inc.	Secs.	Pri. inc.	Col. inc.	Secs.
One-pool: all LSM									
2.61	28.76	28.76	6.57	5.70	5.70	3.61	0.10	0.10	0.09
5.23	14.13	14.13	4.11	5.03	5.03	3.22	0.03	0.03	0.03
6.28	10.83	10.83	3.43	4.78	4.78	3.08	0.02	0.02	0.02
7.33	6.78	6.78	2.89	4.42	4.42	2.84	0.02	0.02	0.02
8.37	5.93	5.93	2.56	4.21	4.21	2.79	0.03	0.03	0.03
9.41	4.74	4.74	2.20	3.89	3.89	2.58	0.02	0.02	0.02
10.46	3.68	3.68	1.89	3.58	3.58	2.33	0.01	0.01	0.01
11.51	2.96	2.96	1.76	3.30	3.30	2.16	0.01	0.01	0.01
12.56	2.43	2.43	1.49	3.10	3.10	2.04	0.02	0.02	0.02
One-pool: all UPM									
2.61	38.12	37.92	5.89	7.52	7.39	3.48	10.11	10.11	0.27
5.23	20.40	20.24	3.48	6.16	6.03	2.95	2.06	2.06	0.11
6.28	15.80	15.66	2.81	5.54	5.41	2.77	2.25	2.25	0.10
7.33	12.18	12.05	2.44	5.34	5.21	2.60	1.91	1.91	0.07
8.37	9.75	9.65	2.07	4.72	4.57	2.32	1.91	1.91	0.07
9.41	7.59	7.52	1.70	4.22	4.11	2.15	1.02	1.02	0.05
10.46	5.74	5.68	1.42	3.84	3.75	1.97	0.91	0.91	0.03
11.51	4.48	4.39	1.29	3.59	3.50	1.78	0.91	0.91	0.03
12.56	3.52	3.44	1.07	3.48	3.40	1.70	0.70	0.70	0.03

Table 9: Intraday payment management for the tit-for-tat scenarios This table presents how many times operational intervention actions are taken each day for all payments of different priorities. The numbers reported are the averages over 122 days in the data sample. For each priority group, the first column shows the numbers of priority increases; the second column shows the numbers of extra collateral pledged intraday to ensure all transactions settle by critical timelines; and the third column lists the numbers of seconds during the day when there is any action of operational intervention. In this data sample, there are on average 115 payments of priority 3 (P3), 166 payments of priority 5 (P5) and 38,639 payments of priority 99 (P99) every day.

Initial liq. (billion \$)	Priority 3			Priority 5			Priority 99		
	Pri. inc.	Col. inc.	Secs	Pri. inc.	Col. inc.	Secs	Pri. inc.	Col. inc.	Secs
Tit-for-tat: urgent in LSM									
2.61	30.02	30.02	7.55	6.66	6.66	4.19	937.86	937.62	13.83
5.23	15.25	15.25	5.03	5.84	5.84	3.66	935.08	934.84	13.48
6.28	11.93	11.93	4.35	5.39	5.39	3.47	934.31	934.07	13.36
7.33	7.74	7.74	3.66	5.15	5.15	3.30	934.00	933.76	13.25
8.37	6.78	6.78	3.14	4.79	4.79	3.04	933.61	933.37	13.19
9.41	5.45	5.45	2.78	4.52	4.52	2.83	932.54	932.30	13.13
10.46	4.33	4.33	2.43	4.26	4.26	2.64	932.07	931.83	13.06
11.51	3.52	3.52	2.25	3.89	3.89	2.43	931.68	931.44	13.00
12.56	3.02	3.02	2.02	3.70	3.70	2.33	931.39	931.16	12.95
Tit-for-tat: urgent in UPM									
2.61	76.44	76.3	15.8	33.97	33.78	8.41	19.44	19.44	6.27
5.23	76.44	76.3	15.8	33.97	33.78	8.41	13.69	13.69	4.86
6.28	76.44	76.3	15.8	33.97	33.78	8.41	12.79	12.79	4.39
7.33	76.44	76.3	15.8	33.97	33.78	8.41	11.93	11.93	3.91
8.37	76.44	76.3	15.8	33.97	33.78	8.41	10.44	10.44	3.53
9.41	76.44	76.3	15.8	33.97	33.78	8.41	9.31	9.31	3.23
10.46	76.44	76.3	15.8	33.97	33.78	8.41	8.80	8.80	2.98
11.51	76.44	76.3	15.8	33.97	33.78	8.41	7.13	7.13	2.67
12.56	76.44	76.3	15.8	33.97	33.78	8.41	7.23	7.23	2.38

Table 10: Intraday payment management for the non-coordinated scenarios This table presents how many times operational intervention actions are taken each day for all payments of different priorities. The numbers reported are the averages over 122 days in the data sample. For each priority group, the first column shows the numbers of priority increases; the second column shows the numbers of extra collateral pledged intraday to ensure all transactions settle by critical timelines; and the third column lists the numbers of seconds during the day when there is any action of operational intervention. In this data sample, there are on average 115 payments of priority 3 (P3), 166 payments of priority 5 (P5) and 38,639 payments of priority 99 (P99) every day.

Initial liq. (billion \$)	Priority 3			Priority 5			Priority 99		
	Pri. inc.	Col. inc.	Secs	Pri. inc.	Col. inc.	Secs	Pri. inc.	Col. inc.	Secs
Non-coordinated: BoC using LSM									
2.61	63.66	63.66	21.96	85.39	85.39	16.35	924.86	924.62	13.55
5.23	59.51	59.51	21.52	85.07	85.07	16.19	922.98	922.74	13.25
6.28	58.56	58.56	21.34	84.98	84.98	16.17	922.48	922.25	13.17
7.33	58.31	58.31	21.30	84.76	84.76	16.07	922.41	922.17	13.06
8.37	58.01	58.01	21.22	84.58	84.58	16.04	922.16	921.92	13.00
9.41	57.87	57.87	21.13	84.50	84.50	16.01	921.43	921.20	12.92
10.46	57.66	57.66	21.07	84.34	84.34	15.98	921.08	920.84	12.87
11.51	57.51	57.51	21.01	84.24	84.24	15.96	920.74	920.50	12.79
12.56	57.32	57.32	20.98	84.15	84.15	15.93	920.55	920.31	12.72
Non-coordinated: BoC using UPM									
2.61	45.49	45.44	13.62	14.48	14.46	8.29	266.98	266.92	12.75
5.23	38.11	38.07	10.92	10.61	10.60	6.55	262.69	262.62	11.85
6.28	36.86	36.81	10.37	9.70	9.68	6.05	261.71	261.65	11.46
7.33	36.25	36.20	10.02	8.98	8.97	5.64	259.86	259.80	10.95
8.37	35.48	35.43	9.57	8.51	8.49	5.41	258.29	258.22	10.66
9.41	34.98	34.93	9.31	7.85	7.84	4.97	257.17	257.11	10.07
10.46	34.44	34.39	8.91	7.59	7.57	4.75	255.52	255.46	9.75
11.51	34.09	34.04	8.65	7.07	7.06	4.52	254.81	254.75	9.53
12.56	33.77	33.72	8.44	6.80	6.78	4.29	253.90	253.84	9.33

References

- Arjani, N. 2006. “Examining the Trade-Off Between Settlement Delay and Intraday Liquidity in Canada’s LVTS: A Simulation Approach.” Bank of Canada Staff Working Paper No. 2006-20.
- Arjani, N. and D. McVanel. 2006. “A Primer on Canada’s Large Value Transfer System.” Bank of Canada, March 1.
- Bech, M. L. and R. Garratt. 2003. “The Intraday Liquidity Management Game.” *Journal of Economic Theory* 109 (2): 198–219.
- Byck, S. and R. Heijmans. 2021. “How Much Liquidity Would a Liquidity-Saving Mechanism Save if a Liquidity-Saving Mechanism Could Save Liquidity? A Simulation Approach for Canada’s Large-Value Payment System.” *Journal of Financial Market Infrastructures* 9 (3).
- Chapman, J., J. Chiu, S. Jafri and H. Pérez Saiz. 2015. “Public Policy Objectives and the Next Generation of CPA Systems: An Analytical Framework.” Bank of Canada Staff Discussion Paper No. 2015-06.
- Castro, P., A. Desai, H. Du, R. Garratt and F. Rivadeneyra. 2021. “Estimating Policy Functions in Payments Systems Using Reinforcement Learning.” Bank of Canada Staff Working Paper No. 2021-7.
- Jurgilas, M. and A. Martin. 2013. “Liquidity-Saving Mechanisms in Collateral-Based RTGS Payment Systems.” *Annals of Finance* 9 (1): 29–60.
- Martin, A. and J. McAndrews. 2008. “Liquidity-Saving Mechanisms.” *Journal of Monetary Economics* 55 (3): 554–567.
- Rivadeneyra, F. and N. Zhang. 2019. “Efficiency Frontiers of Payments Systems.” Bank of Canada mimeo.
- Rivadeneyra, F. and N. Zhang. 2020. “Liquidity Usage and Payment Delay Estimates of the New Canadian High Value Payments System.” Bank of Canada Staff Discussion Paper No. 2020-9.