Staff Working Paper/Document de travail du personnel—2023-26
Last updated: May 4, 2023

How Banks Create Gridlock to Save Liquidity in Canada's Large Value Payment System

by Rodney Garratt,[1] Zhentong Lu[2] and Phoebe Tian[2]

[1]Bank for International Settlements
rodney.garratt@bis.org

[2]Financial Stability Department
Bank of Canada
zlu@bankofcanada.ca, xtian@bankofcanada.ca

DOI: https://doi.org/10.34989/swp-2023-26 | **ISSN** 1701-9397

Acknowledgements

We thank Jason Allen, Radoslav Raykov, participants of the BoC-PayCan quarterly research workshop, BoC BBL seminar, the Bank of Finland 20th Simulator Seminar and the Economics of Payments XI conference for their helpful comments. The views expressed in this paper are the authors' and do not reflect those of any institutions.

Abstract

Using detailed data from Canada's new high-value payment system (HVPS), we show how participants of the system save liquidity by exploiting the new gridlock resolution arrangement. These observed behaviors are consistent with the equilibrium of a "gridlock game" that captures the key incentives that participants face in the system. The findings have important implications for the design of HVPSs and shed light on financial institutions' liquidity preference.

Topics: Financial institutions; Payment clearing and settlement systems
JEL codes: E42, E58, G21

Résumé

Au moyen de données détaillées provenant du nouveau système de paiement de grande valeur du Canada, nous montrons comment les participants au système économisent des liquidités grâce au mécanisme de résolution des blocages. Les comportements observés cadrent avec l'équilibre d'un « jeu de blocages » qui reproduit les principales mesures incitatives auxquelles les participants sont exposés dans le système. Nos résultats ont des conséquences importantes pour la conception des systèmes de paiement de grande valeur et mettent en lumière les préférences des institutions financières en matière de liquidités.

Sujets : Institutions financières; Systèmes de compensation et de règlement des paiements
Codes JEL : E42, E58, G21

1 Introduction

A high-value payment system (HVPS) is an integral part of any nation's financial system and plays a key role in the implementation of monetary policy.[1] Participants—typically financial institutions (FIs)—must actively manage their payment flows and account balances in real time to fulfill their payment obligations. In doing so, they face a trade-off between the cost of carrying abundant liquidity to avoid delays in payment processing and the cost of customer dissatisfaction if they run short of liquidity needed to make payments in a timely manner. Resolving this trade-off is complicated by the fact that participants' liquidity choices in a HVPS are interdependent, as one participant's outgoing liquidity is another bank's incoming liquidity. Thus, understanding the incentives of participants and the nature of their strategic interactions is vital in designing an efficient and safe HVPS.

In this paper, we examine the launch of a new HVPS in Canada called Lynx that substantially alters financial institutions' incentives to provide liquidity.[2] Whereas the previous system adopted liquidity pooling and risk sharing mechanisms to reduce banks' liquidity needs, the new system requires banks to provide liquidity up front for all payments, with an exception for banks that have insufficient liquidity available in the designated payment stream to make the payment. In this case, payments are queued and settled on a net basis, a process called gridlock resolution. Participants would like to save liquidity by queuing payments (which will be resolved by the gridlock resolution mechanism). However, now they cannot queue payments directly and can only do this indirectly by keeping their liquidity low, so this is a "friction" for them to access the queuing/gridlock resolution process. The "voluntary" queuing would give them direct control. The new system has two payment streams. Both are what payments professionals call real-time gross settlement streams

---

[1] HVPSs are used for settling large-value, critical, time-sensitive payments.

[2] Between February 1999 and August 2021, large-value payments in Canada were processed via LVTS (Large Value Transfer Systems). The upgrade to Lynx is an integral part of Canada's Payments Modernization initiative; see `https://modernization.payments.ca/` for a high-level overview of Payments Modernization. The modernization can be viewed as part of the global trend of adopting real-time gross settlement (RTGS) systems in different jurisdictions; see Bech et al. (2008) for an overview of this trend.

(RTGS), because payments are made on a gross basis and are final and irrevocable once processed. However, one stream includes a gridlock resolution mechanism (we denote this stream by RTGS^G thereafter) that has added functionality, and thereby dominates the pure RTGS stream. FIs quickly figured out a clever way to use both streams to their advantage.

In the Lynx system, gridlock resolution is activated only if there is insufficient liquidity in the payment stream to settle payments on a gross basis. Participants cannot voluntarily put payments into the gridlock queue. Hence, the only way for FIs to obtain the liquidity savings associated with netting in the gridlock resolution mechanism queue is to starve that stream of liquidity. By submitting more payments than liquidity to the RTGS stream with a gridlock resolution mechanism, FIs are able to trigger gridlock resolution and settle payments on a net basis.

We document the ability of FIs to recognize and coordinate on this mutually beneficial strategic action. Our analysis focuses on the first month in which the new Lynx system went live.[3] Lynx was launched on August 30, 2021. On the first day of Lynx, most FIs quickly learned the features of the two sub-systems, and chose to process most of their payments (more than 95% in terms of both volume and value) in the RTGS^G. Also, participants allocated large amounts of liquidity into RTGS^G so that the gridlock resolution was rarely engaged and the sub-system effectively became a pure RTGS. Recognizing an opportunity to save liquidity, on September 16, all the major participants jointly reduced their liquidity allocations to the RTGS^G. As expected, this joint action created more gridlocks and queued payments, and activation of the gridlock resolution algorithm led to settling queued payments on a net basis, thus achieving the desired liquidity savings. Engaging the gridlock mechanism lead to delay in settlement of some payments. However, our calculations show that this shift in behavior brought down the system-wide liquidity level by about 76% and caused only about 30 minutes' delay to the system.

---

[3] After the first month, FIs' liquidity management behavior reached to a new "steady state," so we focus on the first month, in which the most interesting adaptation process happened. Interested readers can consult our companion paper, Desai et al. (2022), which compares the steady-state payment patterns using similar sample periods from LVTS and Lynx.

These changes in participants' liquidity management behavior reveal FIs' key incentives and how their interactions shape the overall performance of the new system in terms of liquidity efficiency and throughput. We show that the key incentives can be captured by a parsimonious game theory model, which we call *gridlock game*. In this game, players, i.e., participants of the system, decide whether to allocate sufficient liquidity to settle payment(s) immediately or leave zero liquidity to get payments queued and let the gridlock resolution algorithm of the system settle them. The equilibrium outcome predicted by the gridlock game rationalizes participants' liquidity management behavior after September 16, which explains why the new liquidity management behavior persists. These findings are largely validated by our conversations with the cash managers from the major participants of Lynx.

Apart from furthering our understanding of the economic incentives of participants, results from this paper have policy implications for important issues with regard to the design of HVPSs in general. For example, our paper suggests that built-in voluntary queue functionality is a valued feature in RTGS systems, since participants clearly desire direct control over payment processing and liquidity management.

The rest of the paper is organized as follows. The literature review is provided in Section 2. Section 3 provides background information on the transition from the legacy LVTS to the new Lynx and discusses the key differences between the two systems. Section 4 describes the gridlock game and its equilibrium outcome. Section 5 documents the major changes in participants' liquidity management behavior and system-level implications. Section 6 concludes.

2 Related Literature

There have been numerous theoretical discussions of the liquidity-delay trade-off in the payment system literature; see Bech and Garratt (2003), Mills and Nesmith (2008) and Bech and Garratt (2012). These papers examine intraday payment timing in a world where priced intraday credit is available for liquidity shortfalls and strategic behavior relates to the timing of payments. There is also some empirical evidence of strategic behavior in payment timing. McAndrews and Potter (2002) study the behavior of Fedwire participants around the terrorist attacks on the US on September 11, 2001 ("9/11"). Similarly, Bech and Garratt (2012) show that coordination of early payment processing broke down in Fedwire both during 9/11 and in the immediate aftermath of the Lehman Brothers failure in 2008.

Here the emphasis is on the strategic allocation of liquidity. Most closely related to our work is Rivadeneyra and Zhang (2022), which provides simulation results for Lynx's liquidity efficiency under different assumptions of payment coordination on settlement mechanisms (i.e., payment streams). They conduct simulation experiments using payment data from the legacy system LVTS under different scenarios and behavioural assumptions, where banks may coordinate on one stream or are non-coordinated. They find that Lynx would achieve the highest liquidity efficiency if all participants use the RTGS^G stream, mainly due to the benefits of liquidity pooling and recycling. Our results confirm their main finding, and further show banks not only coordinate on payment streams, but also on liquidity levels and the extent to which they utilize the gridlock resolution in that stream.

We contribute to the growing literature on the benefits of liquidity-saving mechanisms (LSMs) in RTGS systems. Roberds (1999) examines the incentive effects of three different types of settlement rules on banks' portfolio decisions: gross settlement, net settlement, and gross settlement combined with queuing. He shows that the gross settlement combined with queuing can limit banks' "risk-shifting behavior" that appears in a gross settlement when settlement is delayed. Martin and McAndrews (2008) provide a theoretical analysis of liquidity LSMs. Participants of high-value payment systems trade off delay cost against liquidity cost, and they show that the design of an LSM has important implications. In some cases the addition of an LSM increases welfare, whereas in other cases it does not. Martin and McAndrews (2010) extend the previous model by introducing a noisy signal of bank's liquidity cost to study different specific designs of LSM. Jurgilas and Martin (2013) shows that an LSM allows banks to economise on collateral while also providing incentives to submit payments earlier; thus, introduction of the LSM always improves welfare. Atalay et al. (2010) quantify the efficient allocation in that environment. Galbiati and Soramaki (2010) model a stylized two-stream payment system where banks choose how much liquidity to post and which payments to route into each of two "streams": the RTGS stream, and an LSM stream. The paper clarifies and highlights the ways that LSMs are beneficial. Empirical studies on LSMs are rare, but Alexandrova-Kabadjova et al. (2022) provide a global comparison of particular LSM design features and describe the impact on the incentives and behavior of RTGS system participants around the world. Our paper contributes to the discussion of LSMs by demonstrating banks' desire to not only utilize an LSM, but take actions to effectively create one in a liquidity-abundant environment.

3 Background: Transition from LVTS to Lynx

3.1 Legacy System: LVTS

The LVTS was Canada's real-time electronic payments system for processing inter-bank large-value payments; it operated from February 1999 to August 2021.[4] It was the only "systemically important" payments system in Canada, operated by Payments Canada and overseen by the Bank of Canada (BoC). In 2021, the LVTS had 17 direct participants, including the BoC.

The LVTS consisted of two sub-systems, Tranche 1 or Tranche 2. Participants chose either one when sending a payment. Tranche 1 and Tranche 2 differed mainly in their distinct collateral requirements and risk control measures, reflecting their different loss-sharing arrangements in the event of default.

In Tranche 1, a participant could send a payment as long as its net debit position,

---

[4] The description of LVTS in this subsection is very brief; interested readers can consult Arjani and McVanel (2006) for a detailed account of the system.

calculated as the difference between all of the Tranche 1 payments it sent and those it received, was no greater than the collateral the participant had pledged to the BoC to back up its Tranche 1 payments. If the participant defaulted on its LVTS settlement obligations in Tranche 1, the collateral it pledged was used to cover any net negative position in its Tranche 1 account. For this reason, Tranche 1 payments were known as "defaulter pays."

In Tranche 2, at the beginning of each day, each participant granted bilateral credit limits (BCLs) to every other participant in the system; this represented the largest bilateral net exposure it was willing to accept with respect to the other participants. In addition, each participant was subject to a multilateral net debit cap, calculated as the sum of all of the BCLs extended to it and then multiplied by a specified system-wide percentage (SWP) set by the BoC. The multilateral net debit cap represented the maximum multilateral net debit position the participant could incur against all other participants during the trading day. Each participant pledged, to the BoC, collateral that was equal to the largest BCL it had extended to any other participants, multiplied by the SWP. If a participant defaulted on its final settlement obligation, the collateral pool was used to cover the defaulter's remaining amounts owing.[5] For this reason, Tranche 2 payments were referred as "survivor pays."

Payments were processed by the LVTS with finality in real time, while settlement of the system occured on a multilateral net basis at the end of each day. Immediate intraday finality was achieved by use of collateral to secure participants' intraday net debit (negative) positions and also by a residual guarantee provided by the Bank of Canada. In this respect, LVTS was a pseudo-RTGS system (with the collateral cost saving feature offered by Tranche 2), though it did not completely eliminate intraday credit risk among participants.

3.2 Payments Modernization: Lynx System

As an integral part of Canada's Payments Modernization initiative, Lynx has become Canada's new HVPS for processing large-value, time-critical payments. Lynx is an RTGS system and

---

[5] In the event of a participant default, the surviving participants' losses were determined based on the BCL that had been granted to the defaulter (Arjani and McVanel (2006)).

fully compliant with the Principles for Financial Market Infrastructures issued by the Committee on Payments and Market Infrastructures at the Bank for International Settlements.[6] In particular, its participants fully cover their credit risk exposures, which means the system no longer relies on either the "survivors-pay" collateral pool or the residual guarantee from the BoC in case of default.

As is typical in RTGS systems, the reduction of credit risk in Lynx comes at the cost of a substantial increase in intraday liquidity requirements. To mitigate this concern, Lynx offers two distinct payment streams with separate initial liquidity allocations: one is RTGS^G, a RTGS with gridlock resolution and the other is a pure RTGS.[7]

The RTGS^G stream is designed to provide liquidity efficiency through liquidity recycling and payment offsetting in the event of liquidity shortfalls. If a participant does not have sufficient liquidity to settle a payment immediately, the payment is automatically queued until sufficient liquidity becomes available, either through incoming payments or through transfers of additional funds into the stream. Liquidity recycling is enhanced by the use of a settlement sequence that bypasses a strict first-in, first-out (FIFO) ordering of queued payments. In addition, the RTGS^G stream employs a payment offsetting algorithm, called "Gridlock Buster," that runs periodically and attempts to identify queued payments that can be offset simultaneously. The pure RTGS does not have a gridlock resolution mechanism. Each stream is funded by separate intraday liquidity accounts, and intraday liquidity can be transferred between the streams; however, payments must be settled in the settlement mechanism to which they were submitted.[8]

---

[6] See `https://www.bis.org/cpmi/info_pfmi.htm` for details.

[7] In the Lynx documentation, the two streams are called "LSM (liquidity saving mechanism)" and "UPM (urgent payment mechanism)," respectively. We do not adopt this terminology, since referring to a gridlock resolution mechanism as an LSM, while true at a high level, can be misleading. While gridlock mechanisms do save liquidity when activated, they do not offer the voluntary queuing option that is typically associated with LSMs. Rather, gridlock has to be induced to save liquidity, which is the point of this paper.

[8] The description here is largely borrowed from Lynx documents published on the website of Payments Canada; see `https://www.payments.ca/high-value-payment-system-lynx`. Besides the core design features in terms of liquidity cost and settlement risk, the Lynx system adopts the global ISO 20022 messaging standard and builds on modern software and hardware technology, which provides enhanced cybersecurity and resiliency, as well as other new functionalities like automated dashboards and application programming interfaces (APIs). These improvements of Lynx over LVTS received very positive feedback from participants

3.3 Gridlock Resolution in Lynx

RTGS systems emerged in the 1980s to speed up wholesale payments, and are now standard around the world (see Bech and Hobijn (2007)). These systems, however, are prone to delay of payments due to their high liquidity requirements. A common response to the liquidity concern of a RTGS system is incorporating an LSM, i.e., a queuing arrangement, that allows a participant to submit payment orders to a queue waiting to be released when some conditions are met, e.g., the receipt of an offsetting payment. Such LSMs have been adopted by many RTGS systems, e.g., TARGET2 (for Euro payments), Japan's BOJ-Net and CHAPS (for the United Kingdom).

The key advantage of adopting an LSM in a RTGS system is that it can can increase liquidity efficiency without introducing credit risks to the system. Although an LSM might increase payment delays because of queuing, it can reduce participants' need for "internal queues"; thus, overall, the LSM may not increase the system throughput time, and can even shorten it in some cases.

The gridlock resolution arrangement in Lynx is a form of LSM and aims to improve liquidity efficiency of the system. However, this arrangement differs from the "canonical LSM" discussed above. The canonical LSM allows participants to voluntarily queue payments (even if they have sufficient liquidity), while the payments sent through the RTGS^G are queued only when the liquidity is insufficient and settled instantly otherwise. In other words, in canonical LSM, participants directly decide whether to send payments to the queue, while in Lynx they can only indirectly control queuing payments by adjusting their settlement balances.

3.4 Initial Liquidity Allocation in Lynx

By chance, Lynx was launched during the COVID-19 pandemic. At the time of launch the Bank of Canada was engaged in quantitative easing (QE). QE involves the central bank

---

(based on our conversations with them) and thus helped the transition.

buying government of Canada bonds from primary dealers (who are also Lynx participants) with settlement balances. At the beginning of each business day, the Bank of Canada makes an automatic payment of yesterday's settlement balance plus accrued interests to each participant. Crucially, the settlement balance payments always take place in RTGS^G. These settlement balances are so large that Lnyx participants have more than enough liquidity to meet all of their outgoing payments during the day. Effectively, without any active liquidity management, the large sum of settlement balances due to QE means that RTGS^G behaves like a pure RTGS, and there is no need for a gridlock resolution mechanism.

Thus, if participants want to obtain liquidity savings within Lynx, then they must intentionally lower their liquidity level in RTGS^G. This takes effort, because they need to actively transfer a large portion of their settlement balance from the RTGS^G stream to the pure RTGS stream in order to make their liquidity level in the RTGS^G stream low enough to get payments queued.

4 Gridlock Game

In this section, we provide a theoretical model that illustrates the strategic incentives of Lynx participants to conserve liquidity by creating gridlock. We consider a one-shot game in which banks have symmetric payments but differ in terms of their liquidity management costs; some banks find it prohibitively costly to manage liquidity allocations across payment streams intraday. This modelling assumption simplifies our analysis and reflects correspondence with system participants. In reality, banks decide repeatedly throughout the day how much liquidity to allocate across the two payment streams and into which stream to submit existing payment requests. However, the incentives that drive observed behavior can be illustrated in a one-shot model, in which banks make a single liquidity allocation decision and play simplified (binary) strategies when determining their payment allocations; they either place all payments in the RTGS^G stream or place all of them in the pure RTGS stream.[9] The essential feature of the model is that it captures two costs to strategically induce gridlock, the liquidity management cost and the delay cost associated with forcing payments to be queued and settled by the gridlock resolution mechanism.

4.1 Model

There are $i = 1, ..., n$ banks. Each bank starts off the day with requests to pay each other \$1. There are two payment streams, RTGS^G and RTGS. Each bank i decides how much initial liquidity I_i^J to allocate to each stream $J \in \{RTGS^G, RTGS\}$, at cost f per dollar, and which stream $J \in \{RTGS^G, RTGS\}$ to submit their payments. Payments submitted to stream RTGS^G or RTGS are processed immediately if there is enough available liquidity. If there is not enough liquidity in the RTGS^G stream, payments enter a netting queue that will allow payments to be processed on a net basis if the net amount of liquidity is available. The gridlock process takes some time to execute, so we assume banks bear a small per-dollar delay cost δ when payments are netted through this process. If there is not enough liquidity to cover net amounts in the RTGS^G stream or gross amounts in the RTGS stream at the end of the day, then banks incur a per-dollar cost d, which reflects the costs of delay and any liquidity management costs r needed to settle these payments.

The default liquidity allocation is the RTGS^G stream. If banks wish to move liquidity to the pure RTGS stream at the beginning of the day or remove liquidity from the system to save the liquidity cost f they must incur a liquidity management cost $c_i > 0$. There is no such explicit cost in reality. As mentioned above, this assumption is a short-cut to reflect the fact that some banks consider it to be costly to manage their liquidity allocations across streams throughout the day. Since we only consider one period, we capture this idea by a single transfer cost.[10]

---

[9] The assumption that players choose a single strategy to play against all of the other players is common in the literature on social coordination games (see Goyal and Vega-Redondo (2005), Goyal (2011)).

[10] In our conversations with the participants, we learned that another incentive to move potentially exces-

Let $s_i \in \{RTGS^G, RTGS\}$ denote that payment stream chosen by bank i for its payments to other banks. Then, a complete strategy for bank i with cost k is a triple $S_i = (I_i^{RTGS^G}(c_k), I_i^{RTGS}(c_k), s_i(c_k)) \in \boldsymbol{S_i} = R_+^2 \times \{RTGS^G, RTGS\}$ and a strategy profile is $S = (S_1 ...., S_n)$. Let $D_i^J(S) \geq 0$ denote the value of payments for bank i that remain unsettled in the RTGS^G and RTGS payment streams and let $N_i(S)$ denote the value of bank i's payments that enter the gridlock mechanism. Note that $D_i^J(S)$ and $N_i(S)$ depend on the the strategy profile of all players. Finally, if there is a net liquidity shortfall across streams at the end of the day, then banks have to borrow to complete these payments at rate $\ell > f$. Let $L_i(S) \geq 0$ denote any such shortfall for play i.

Bank i's payoff function is

$$P_i(S) = -f(I_i^{RTGS^G} + I_i^{RTGS}) - c_i \mathbf{1}_{I_i^{RTGS^G}=0} - N_i(S)\delta - (D_i^{RTGS^G}(S) + D_i^{RTGS}(S))d - L_i(S)\ell. \tag{1}$$

We refer to the collective choice of initial liquidity and the payment allocation decisions as the *Gridlock Game*. We can now state the following definition:

Definition 1 *A Nash equilibrium of the Gridlock Game is a strategy profile S^* such that for each bank i the strategy S_i^* maximizes $P_i(S_i, S_{-i}^*)$ over all $S_i \in \boldsymbol{S_i}$.*

4.2 Example 1

Suppose there are only two banks, 1 and 2. A bank will never allocate liquidity to the pure RTGS stream and not assign its payment to that stream (since then it incurs a liquidity cost with no potential benefit), nor will it assign its payment to the pure RTGS stream and not provide liquidity to that stream (since $\ell > f$), nor will it assign liquidity and payments

---

sive liquidity out of the RTGS^G stream is to keep intraday exposure under control. This is because nowadays many banks use an automated payment submission process (to the RTGS^G stream in Lynx), so a bank can incur a rather large exposure if it has a very high liquidity level in the system and for some reason the incoming payment flow slows down (e.g., operational disruptions to some participants). This incentive implies a potential benefit from moving liquidity out of the RTGS^G stream and thus a smaller (or even negative) c_i. But it is straightforward to verify that this potentially interesting modification does not affect the key implications of our theoretical model.

Table 1: Initial liquidity contribution & payoff in a Gridlock Game

		Bank 2	
		(1,0,RTGS^G)	(0,0,RTGS^G)
Bank 1	(1,0,RTGS^G)	$-f, -f$	$-f, -c_2$
	(0,0,RTGS^G)	$-c_1, -f$	$-\delta - c_1, -\delta - c_2$

to the pure RTGS stream (since this produces the same outcome as assigning liquidity and payments to the RTGS^G stream without incurring the cost $c_i > 0$). Hence, each bank has only two undominated strategies: $(1, 0, RTGS^G)$ and $(0, 0, RTGS)$. The reduced game in normal form is shown in Table 1.

The gridlock resolution process involves some small delay in payment processing, so banks must trade off liquidity savings against the cost of delay. In cases where the delay cost from entering gridlock resolution and the banks' liquidity management costs are low, relative to the cost of liquidity provision, we should expect to see banks attempting to induce gridlock. More precisely, provided that the delay cost $\delta + c_i < f$, this game has a unique Nash equilibrium $S = ((0, 0, RTGS^G), (0, 0, RTGS^G))$.

This example is highly stylized, but it should be apparent that even in more general cases, by limiting the amount of liquidity available to the RTGS^G stream, banks can cause their payments to enter the gridlock mechanism, which leads to netting and liquidity savings. We now provide a second example that shows, first, that the result in example 1 extends to the case of $n > 2$, and second, that incentives to induce gridlock remain even when some players find it prohibitively costly to actively manage their liquidity intraday.

4.3 Example 2

Suppose that there are $n > 2$ banks and each bank receives requests from its customers to pay \$1 to customers of each of the other $n - 1$ banks. As discussed in the model description, banks play simplified strategies in which they either provide $\$(n - 1)$ in the RTGS^G stream or the pure RTGS stream, or they provide \$0. We also assume for the purposes of this example that the liquidity choice is determined mechanically, given the decision to allocate liquidity to a particular payment stream. That is, we assume that if a bank decides to provide liquidity to a stream in which they have submitted payments, then the amount they provide is equal to the gross value of payments they submit.[11] As in example 1, there are only two undominated strategies, $(0, 0, RTGS^G)$ and $(n-1, 0, RTGS^G)$, in the game for any choice of the parameters satisfying $\ell > f$ and $c_i > 0$ for all $i = 1, ..., n$.

We make the following two observations.

Observation 1. Suppose $\frac{c_i}{n-1} < f - \delta$ *for* $i = 1, ..., n$. *Then it is a Nash equilibrium for all banks to submit their payments to the payment stream with gridlock resolution and provide no liquidity.*

The claim made in this observation is easily verified. Assume that players $j = 2, ..., n$ play $S_j^* = (0, 0, RTGS^G)$. Then player 1 gets a payoff of $-c - (n-1)\delta$ if she plays $S_1 = (0, 0, RTGS^G)$ and a payoff of $-f(n-1)$ if she plays $S_1 = (n-1, 0, RTGS^G)$. Since all the players are identical, evaluation of this case is enough to establish that it is a best response for any player to play the proposed strategy if the others do.

It should be clear that banks will continue to play a liquidity-saving strategy, even if they think others cannot. In fact, they have an even greater incentive to do so, because they can rely on incoming liquidity from players that do not attempt to save liquidity.

Observation 2. Suppose $c_i < f(n-1) - \delta(n-m)$ *for* $i = 1, ..., m$ *and* $c_i > f(n-1) - \delta(n-1-m)$ *for* $i = m+1, ..., n$. *There is a Nash equilibrium in which players* 1 *through* m *submit all their payments and no liquidity to the payment stream with gridlock resolution.*

To verify this claim, consider first the case where players $j = 2, ..., m$ play $S_j^* = (0, 0, RTGS^G)$

---

[11] This is a high estimate of what banks would typically provide; however, in this simple model it ensures that banks can fund payments in their selected stream, without delay, regardless of what other banks do. To capture the fact that banks typically count on incoming liquidity to fund some of their payments we can evaluate the predictions of the model, assuming the cost of liquidity provision is low.

and players $k = m + 1, ..., n$ play the strategy $S_k = (n - 1, 0, RTGS^G)$. Then player 1 gets a payoff of $-c - (n - m)\delta$ if she plays $S_1 = (0, 0, RTGS^G)$ and a payoff of $-f(n - 1)$ if she plays $S_1 = (n - 1, 0, RTGS^G)$. Hence, it is a best response for player 1 to play the proposed strategy given the strategies played by the others. Next consider the case where players $j = 1, ..., m$ play the strategy $S_j = (0, 0, RTGS^G)$ and that players $j = m + 2, ..., n$ play $S_j^* = (n - 1, 0, RTGS^G)$. Then player $m + 1$ gets a payoff of $-c - (n - 1 - m)\delta$ if she plays $S_1 = (0, 0, RTGS^G)$ and a payoff of $-f(n - 1)$ if she plays $S_1 = (n - 1, 0, RTGS^G)$. Hence, it is a best response for player $m + 1$ to play the proposed strategy given the strategies played by the others. Evaluation of these two cases is enough to establish that it is a best response for any player to play the proposed strategy if the others do.

In the next section, we will analyze the detailed data from Lynx and provide empirical evidence supporting the theoretical predictions.

5 Liquidity Management in the First Month of Lynx

5.1 First day in Lynx

To prepare for the transition from the legacy LVTS to Lynx, Payments Canada undertook training and testing plans starting in November 2019. Participants were provided with training sessions to see the Lynx application and perform hands-on exercises. Lynx officially began operations on Monday, August 30, 2021, with 16 direct participants at the time of launch, including BoC.[12] On the first day, there were minor technical issues but payments were successfully settled in the system.

To examine the adaptation process in detail, we divide the first day into three time windows: early morning (12:30 AM to 8:00 AM), late morning (8:00 AM to 12:00 PM) and afternoon (12:00 PM to 6:00 PM). Within each time window, we compare how payment value and volume change relative to the average levels over the last month of LVTS, and we

---

[12] A new participant, Citibank N.A., later joined Lynx on September 18, 2021

show results in Table 2.

As described in Section 3.2, Lynx offers two distinct payment streams with separate intraday liquidity accounts: RTGS^G stream (the stream with gridlock resolution) and the pure RTGS stream. Evaluating each of the two Lynx streams in isolation, RTGS^G stream completely dominates the pure RTGS stream; when Lynx was launched, virtually all payments (99.8% of volume and 96.4% of value) were sent through the RTGS^G stream; afterwards, payment volume and value became even higher, averaging 99.98% and 99.1% respectively.

Table 2: First day in Lynx

	Lynx		LVTS	
	RTGS	RTGS^G	Tranche 1	Tranche 2
Payment value				
Before 8AM	0.01	23.09	10.71	11.47
8AM-12PM	3.06	73.79	53.24	24.06
12PM to 6PM	3.09	67.20	43.18	15.05
Payment volume				
Before 8AM	24	12118	44.65	14029
8AM-12PM	29	17626	198.30	16220
12PM to 6PM	23	16703	114.75	11314
Liquidity allocation	75.904	123.281	8.096	32.919

This table shows aggregate payment values (in billion CAD) and volumes and liquidity allocation in the first day of Lynx (August 30, 2021) in comparison with average August numbers in LVTS (August 1-August 27, 2021). We exclude Bank of Canada's payment of overnight settlement balances in our calculation of payment value and volume. Liquidity allocation in Lynx is measured by available balances in each stream at 8AM, and it is measured by collateral apportioned in each stream in LVTS.

5.2 Evidence of Gridlock Creation

Figure 1: Intraday liquidity management

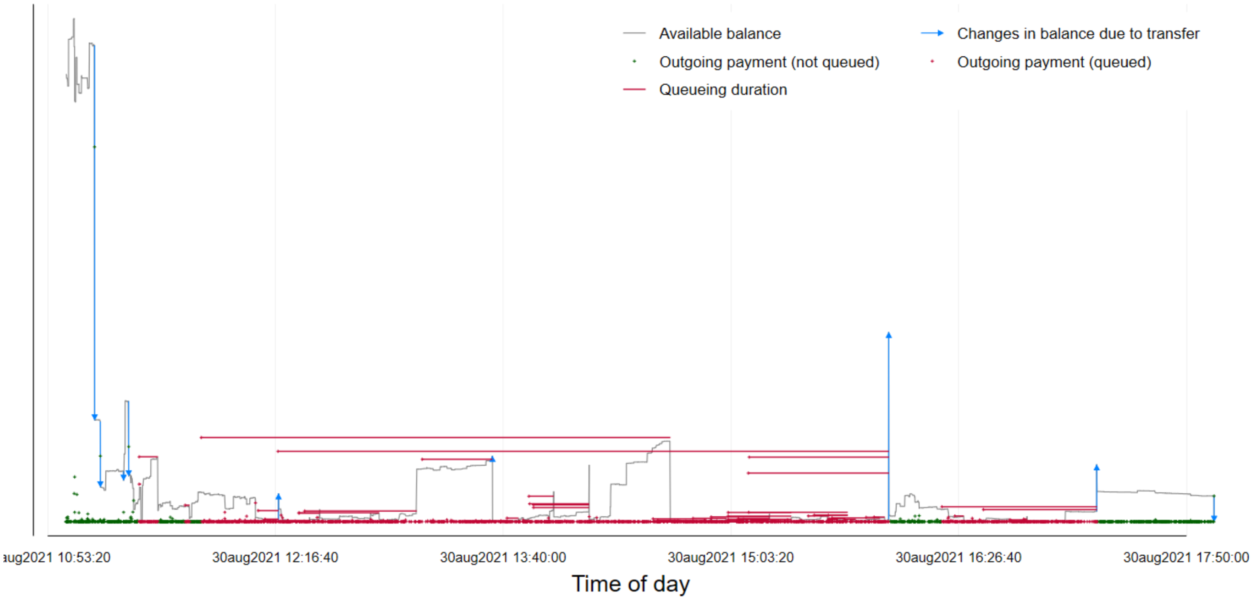

This figure shows one Lynx participant's liquidity management of their RTGS^G stream account on the first day of Lynx. The grey line shows available balance in their RTGS^G stream account at each second. The blue line with arrow indicates changes in balance due to the participant's internal transfer behaviors. The green dot represents an outgoing payment made by the participants that is not queued. The red dot represents submission of an outgoing payment that is queued, and the following horizontal line shows the queuing duration. For confidentiality reasons we cannot show the units on the y-axis.

In the RTGS^G stream, where the majority of payments took place, there is evidence of banks experimenting with the queuing and gridlock resolution function. Figure 1 illustrates one particular bank's liquidity management and gridlock creation behavior in the first day. In the figure, we plot outgoing payments along with the real-time balance in one bank's RTGS^G stream account on the first day. Payments that are not queued are represented by green dots; payments that are queued are shown by red dots, and each is followed by a horizontal line indicating the duration of queuing.[13] Changes in the account balance are caused either by (1) incoming and outgoing payments, or (2) transfers to and from other accounts, notably

---

[13] A red dot without a line means the payment gets queued and resolved almost immediately (duration is too short to see), typically as a result of FIFO-bypass.

the pure RTGS stream. We highlight balance changes due to transfers in blue arrows: a downward arrow means a withdrawal of funds from the RTGS^G account and an upward arrow represents an injection of funds into the RTGS^G account. In this way we can distinguish "passive" changes due to payments and "active" changes due to liquidity management.

At the beginning of the day, this participant has a rather high account balance, which makes all of its early morning payments go directly through without queuing. However, the participant makes a major fund withdrawal from the RTGS^G stream at around 11AM, and the lowered balance level leads to its first queued payments shortly after. Notice that there are numerous small red dots near the bottom of the graph, which represent the small-value payments that are automatically queued following the large-value queued payments that were submitted earlier. Small-value payments are resolved and released from the queue very quickly, due to the FIFO-bypass. The large-value queued payments, however, remain in the queue until incoming payments bring sufficient liquidity. It is also important to notice that this participant transfers funds into the RTGS^G stream account a couple of times to resolve a batch of queued payments at around 4PM and 5PM. These behaviors imply that the participant is actively exploring how to make use of the queuing by managing its liquidity level in the system. Such behavior is very typical: Figure 7 to Figure 11 in the Appendix show more examples throughout September.

The liquidity management behavior varies with participants. To examine this heterogeneity, we focus on the fund transfers during the core business hours of 8AM to 5PM, which indicate how actively a bank manages its account balance during the day. These transfers are typically small-value and for the purpose of micro-manage queuing, as opposed to the large-value transfers in the early morning session, which are typically used for setting baseline liquidity levels.[14] Figure 13 in the Appendix shows the distribution of the number of the liquidity management transfers on each day in the first three months after Lynx's launch. Although participants differ in their managing style, almost all participants engage in such

---

[14] See Figure 15 in the Appendix for a detailed decomposition of flows into and out of the RTGS^G stream over a business day.

micro-managing behaviors of their queues through internal transfers and liquidity controls. Moreover, such engagement intensifies after a coordinated shift that happens in the middle of September, which we shall document in the next subsection.

5.3 Shift on September 16

During the first two weeks after the launch of Lynx, Payments Canada organized daily drop-in sessions to provide support to participants in case they experienced issues related to the functionality and liquidity management in the new payment system. Since no major issues materialized, the sessions served mostly as a forum for communication and information sharing among participants. This convenient information exchange platform paved the way for the coordinated actions we witness later.

In these information sessions, Payment Canada shared system-level statistics with the participants, drawing attention to the high liquidity levels in the system that resulted in the under-utilization of queuing and gridlock resolution. This is mostly due to the special design and initial liquidity allocation of Lynx described in Section 3.4.

As illustrated by the gridlock game in Section 4, at least a subset of participants have incentive to take advantage of the queuing and gridlock resolution functionality by lowering the liquidity level in the RTGS^G stream. Such incentives were communicated during the sessions in early September and turned into action on September 16, 2021. In particular, starting from this day, most participants change their liquidity management behavior from passively sitting on high levels of settlement balances in the RTGS^G stream to actively transferring funds in and out of the RTGS^G stream to engage the queuing and gridlock resolution functionality, which resulted in a large drop in the overall liquidity usage in the system.

To quantitatively assess the shift of liquidity management behaviors on September 16, we measure each participant's daily liquidity level decision using its average balance in the RTGS^G stream during the core business hours (i.e., from 8:00AM to 5:00PM). This measurement is motivated by two empirical patterns about intraday liquidity management behavior. First, although banks can and do make internal transfers between their accounts throughout the whole business day, their transfers in the early morning session (before 8AM) are overwhelmingly larger (than the rest of the day) in the terms of value. Second, despite the constant flow of payments in and out of banks' RTGS^G accounts, the balance level from 8AM to 5PM remains relatively stable, as suggested in Figure 14 in the Appendix.

Together, these patterns suggest that a typical participant would complete major internal transfers in the early morning to set a base level of liquidity for the RTGS^G account, and carry out payment processing functions with occasional internal transfers to adjust liquidity levels if needed. Thus, we view the average balance during the core business hours in the RTGS^G account as a reasonable proxy for the daily liquidity level that a participant *intentionally* chooses for that business day.

Figure 2 shows the time series of total daily liquidity level aggregated across all participants from August 31 to the end of November in 2021. Overall, there is a significant decrease in liquidity levels after the coordinated efforts since September 16, 2021. The aggregate liquidity level dropped by 76%, and the post-change levels are largely sustained. Figure 3 shows the time series of daily liquidity level for six randomly chosen participants. The changes in liquidity levels are largely negative and significant for most banks, with the largest decreases being around 95%. However, there are four banks that did not reduce liquidity levels, with three of them having very low liquidity levels to begin with. As shown in Observation 2 of Section 4.3, the observed varying liquidity management behaviors of the banks can be rationalized by their heterogeneous operational costs and are consistent with the equilibrium of the gridlock game.

5.4 Comparing Pre- and Post-September 16

The shift in liquidity management behavior on September 16 led to a dramatic reduction in banks' account balances, and the new liquidity levels seem to be persistent without reverting

Figure 2: Aggregate liquidity levels of the RTGS^G stream account

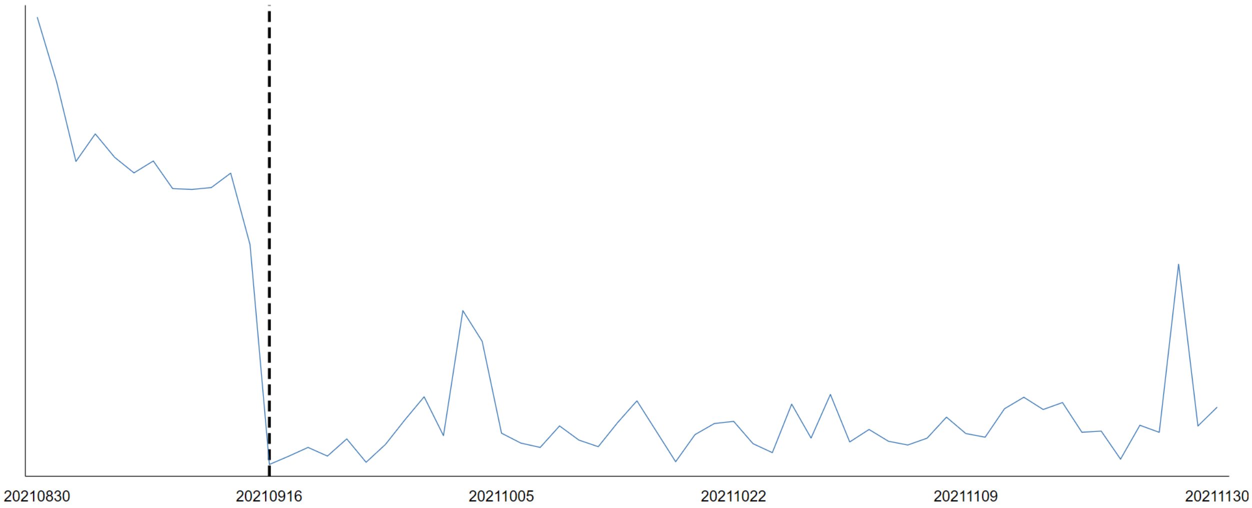

This figure shows aggregate liquidity levels across all Lynx participants' RTGS^G stream accounts over from August 31 to November 30, 2021. A participant's liquidity level in a day is measured by the average available balance in its RTGS^G account from 9AM to 5PM. The black dash vertical line indicates the day of change (September 16).

back to the previous states. In this section, we present the pre- and post-September 19 statistics on queued payments, gridlock resolution, and payment timing, as well as intraday liquidity usage and efficiency. We use a shorter sample period for the comparison, from 30 Aug 2021 to 24 Sept 2021. This is because a new participant joined Lynx at the end of September, which might cause non-trivial changes to the system. Table 3 presents the mean comparison results in this sample period.

5.4.1 Queued Value and Volume

At the system level, the average value of queued payments in each business day increased from 23.7 to 53.6 billion CAD after September 16, which is about a 127% increase. In terms of volume, the average number of queued payments increased by 117% from 4368 to 8739. We also check the fraction of queued payments in all payments settled in the day. Before September 16, queued payments make up 9.3% of all payments in terms of volume, and 14.3% in terms of value. After September 16, the proportion of queued payments increases

Figure 3: Liquidity levels of six randomly chosen participant's RTGS^G stream account

This figure shows liquidity levels in six randomly chosen Lynx participant's RTGS^G account from August 31 to November 30, 2021. A participant's liquidity level in a day is measured by the average available balance in its RTGS^G stream account from 9AM to 5PM. The numbers shown are standardized and anonymized to protect the participants' confidentiality. The black dash vertical line indicates the day of change, September 16.

Table 3: Mean comparison between pre- and post-September 16

	Mean (pre)	Mean (post)	Diff.	t-statistics
Aggregate RTGS^G stream account balance (bn)	88.214	12.768	-75.446	-13.307
Queued Payments				
Total value (bn)	23.672	53.625	29.953	5.787
Total volume (thousand)	4.358	8.739	4.381	4.493
Fraction in overall value	0.143	0.321	0.178	6.097
Fraction in overall volume	0.093	0.204	0.111	4.681
Queuing Duration (minute)[1]				
Total time spent in queues[2]	6458.981	15810.648	9351.667	4.541
Average queuing duration[3]	1.929	1.794	-0.135	-0.207
Value-weighted average queuing duration[4]	26.157	30.054	3.898	0.906
Gridlock Buster[5]				
Number of successful runs[6]	4.167	9.000	4.833	1.604
Volume of payments settled	8.500	28.714	20.214	2.289
Value of payments settled (bn)	1.258	4.191	2.933	2.304
Payment Settlement Time				
Time to settle 25% payment value	0.494	0.511	0.017	1.334
Time to settle 60% payment value	0.689	0.736	0.047	3.181
Time to settle 80% payment value	0.830	0.848	0.018	1.633
Value-weighted average settlement time[7]	0.639	0.657	0.018	2.401
Intraday Liquidity Usage and Efficiency				
Intraday liquidity used (bn)[8]	24.928	21.893	-3.035	-1.259
Total value of payments settled (bn)	182.182	185.823	3.642	0.349
Liquidity efficiency ratio[9]	7.476	8.627	1.151	2.202

Sample period: 30 Aug 2021 to 24 Sept 2021. All values are stated in CAD.

[1] The queuing duration of a queued payment is calculated by the time difference between the settled time and registered time.

[2] Total time spent in queues is the sum of queuing duration across all queued payments.

[3] Average queuing duration is total time spent in queues divided by the total number of queued payments in a day.

[4] Value-weighted average queuing duration is calculated by Equation (2).

[5] Gridlock Buster is a proprietary payment offsetting algorithm defined by Payment Canada.

[6] A successful Gridlock Buster run is recorded when two or more payments are identified by the payment offsetting algorithm and settled simultaneously in one batch.

[7] Value-weighted average settlement time is defined in Equation (3).

[8] Intraday liquidity used in the system is defined in Equation (6).

[9] Liquidity efficiency ratio of the system is defined in Equation (7).

to 20% in terms of volume, and 32.1% in terms of volume. These measures all suggest that queuing is roughly doubled after the joint reduction of liquidity levels in the RTGS^G stream.

Individual level changes of queued payments are rather heterogeneous, which is not surprising given the large disparity in participants' liquidity levels. A natural question is, how do the adjustments in liquidity levels translate to changes in queued payments? In particular, do the banks who withdraw more liquidity from the RTGS^G stream end up with more queued payments? To answer this question, we calculate the percentage change in liquidity level and queued payment value using log difference of pre- and post-September 16 means, i.e, $\log\left(\bar{X}^{post}/\bar{X}^{pre}\right)$. Figure 4 shows the scatter plot and fitted regression line. From the scatter plot, it is clear that banks with larger reduction in liquidity tend to experience larger increase in queued payments. A linear regression yields the negative and significant slope estimate -1.137 with t-statistics -5.14. Roughly speaking, a 1 % reduction in liquidity level is associated with a 1.13 % increase in the value of payments queued.

Figure 4: Percentage changes in liquidity level and queued value

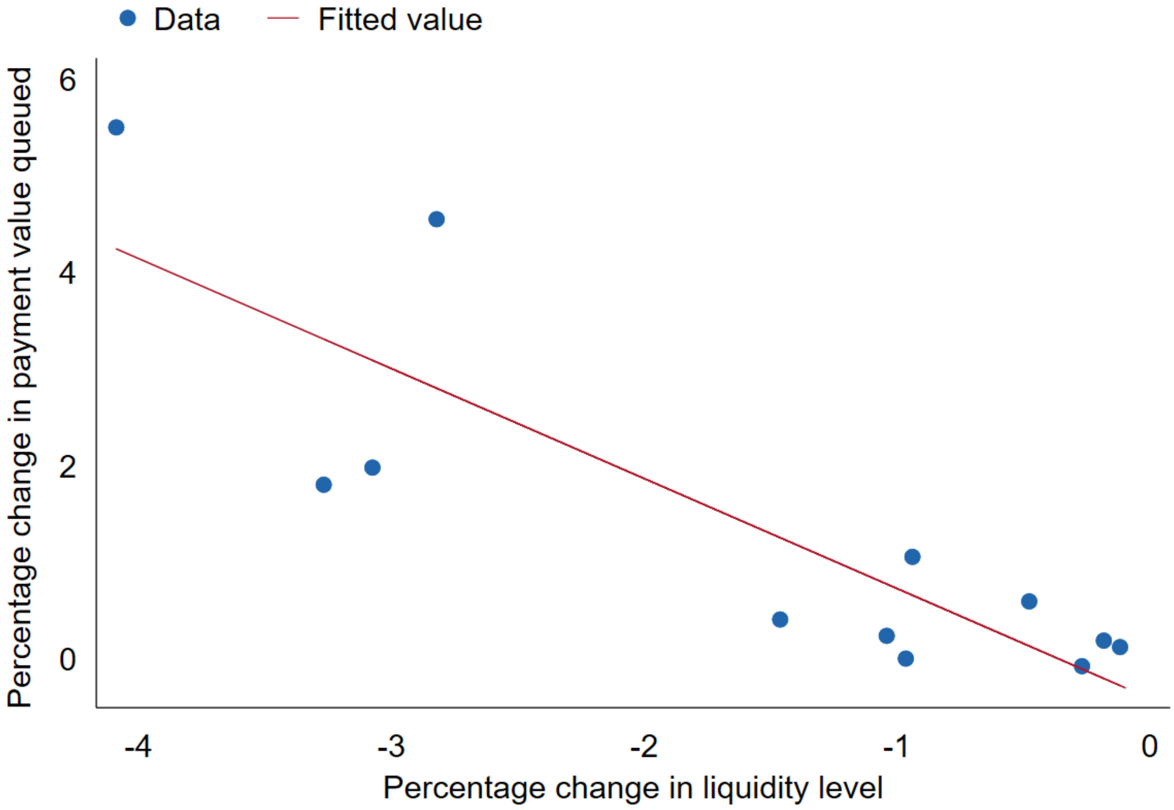

This figure shows the relationship between percentage changes in Lynx participants' percentage changes in liquidity levels and value queued after the shift on September 16, 2021. Liquidity level is measured by the average balance in the RTGS^G stream account from 8AM to 5PM, and value queued is the total amount of outgoing payments that are queued. Percentage change of liquidity level and value queued is calculated through log difference, i.e, $\log\left(X^{post}/X^{pre}\right)$. Each scatter point represents one participant, with some participants missing from the plot because of zero queuing in the pre-September 16 period. Fitted value is found by a linear regression of percentage change in payment value queued on percentage change in liquidity reduction. The slope estimate is -1.137 with t-statistics -5.14 and $R^2 = 0.726$.

5.4.2 Queuing Duration and Delay

With such drastic increase in queued payments, how much more congested is the system? To see this, we examine the time a payment spent in queues, i.e., queuing duration. We first add up all queuing durations to see the total amount of time spent in queues by all queued payments; this total more than doubled from 6,458 minutes to 15,810 minutes. This is not surprising given the doubling of the value and volume of queued payments. However, the average queuing duration (total queuing time divided by the number of queued payments) does not change much (the change is not statistically significant) after September 16. In other words, on average a queued payment does not take longer to settle after the behavioral change on September 16.

We also calculate the value-weighted average queuing duration (i.e., average queuing duration per dollar) on each day as

$$\tilde{D} = \frac{\sum_{k=1}^{K_q} P_k D_k}{\sum_{k=1}^{K_q} P_k} \tag{2}$$

where k is an index for queued payments, K_q is the total number of queued payments on that day, P_k is the dollar amount of queued payment k, and D_k is the time payment k spent in the queue (i.e., the difference between the settled time and registered time). Prior to September 16, the value-weighted average queuing duration is 26.16 minutes, which, as expected, is much higher than the simple average queuing duration because smaller payments typically get settled much faster than larger payment (especially with Lynx's FIFO bypass feature). The post-September 16 value-weighted average queuing duration is 30.05 minutes, slightly higher than pre-September 16, but the difference is not statistically significant.

Overall, the doubled number of queued payments lead to doubled total queuing duration, but conditional on being queued, the average time spent in queues per payment or per dollar does not significantly increase. This is because more queued payments lead to more netting opportunities, which counteracts the effect that more large-value queued payments tend to result in longer queuing duration.

5.4.3 Gridlock Resolution

As mentioned in Section 3.3, Lynx's RTGS^G stream uses two main approaches to resolve gridlocks: FIFO-bypass on the individual queue, and Gridlock Buster on the global queue. Gridlock Buster is a proprietary payment offsetting algorithm that runs periodically and attempts to identify queued payments that can be offset simultaneously. We record one successful Gridlock Buster run when at least two payments are settled simultaneously in one batch. Prior to September 16, the average number of successful Gridlock Buster runs is 4.167 per day, settling 8.5 payments on average. These payments make up about 0.2% of all queued payments in terms of volume, and around 5.3% in terms of value. This pattern is expected because queued payments settled by the Gridlock Buster are usually large so that they could not be resolved by FIFO-bypass. After the liquidity reduction on September 16, both the value and volume of payments settled through Gridlock Buster increase almost threefold, suggesting that more payment gridlocks are created so the algorithm can identify more payment offsetting opportunities.

5.4.4 System Throughput

Given the increased queuing after September 16, one may expect a longer delay in overall settlement time. Thus we would like to examine if participants can still achieve the three important Lynx throughput targets: completing 25%, 60%, and 80% of payment values before 10AM, 1PM, and 4:30PM, respectively. The main purpose of these throughput targets is to improve payment coordination and thus liquidity efficiency of the system.

Figure 5 compares the system-level throughput curve of pre- and post-September 16. A point on the throughput curve shows the proportion of total payment value completed at a time point of the day. The three orange dots represents the three recommended throughput targets at 25%, 60%, and 80%. The blue curve shows the median throughput of the pre-September 16 period, and the red curve shows the median in the post-September 16 period. The red curve is on the right of the blue curve, meaning that it takes a longer time to complete the same proportion of payments value after September 16. The difference around noon is about 20 minutes.[15]

Figure 5: System level throughput changes

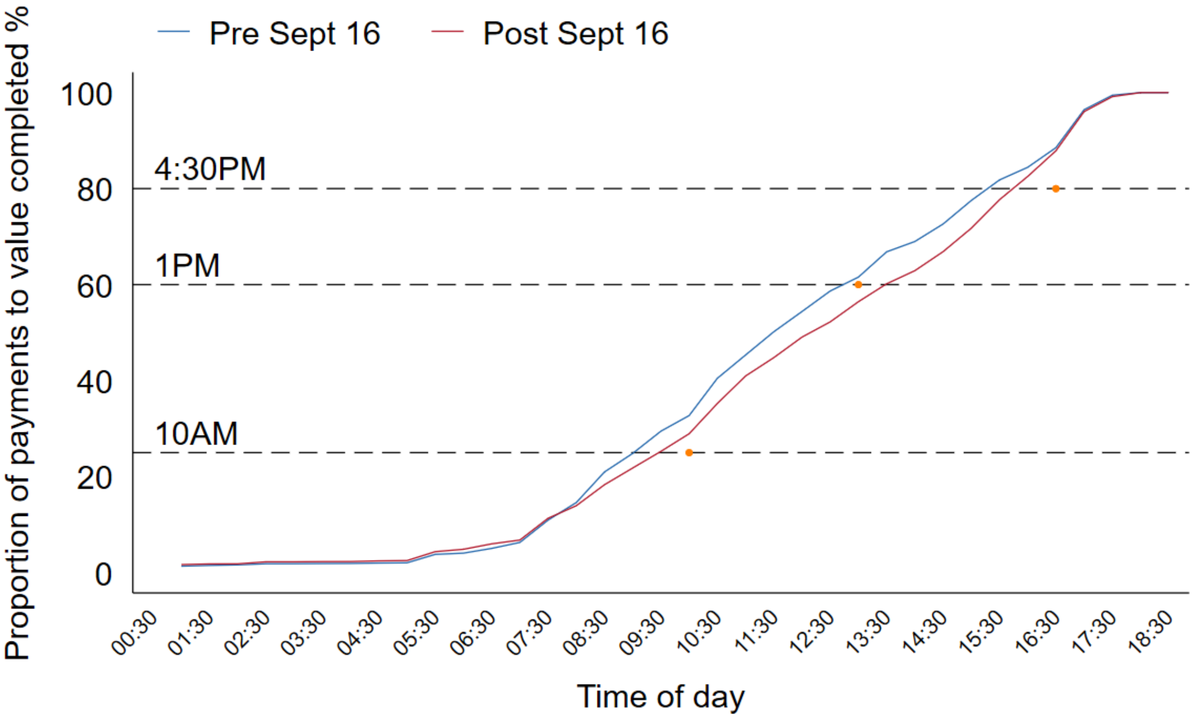

The graph shows payment value throughput before and after September 16. Value throughput at time t is defined as the proportion of payments value completed during a day. The blue line represents the median daily observation from September 1 to September 15, and the red line represents the median daily observation from September 16 to September 29. Orange dots represent the key throughput targets at 10AM, 1PM, and 4:30PM. Payments from Bank of Canada are excluded in our calculation.

In Table 3, we standardize the payment settlement time into an interval between 0 and 1, with 0.5 meaning halfway into the system's operation window in that day.[16] We find that the average settlement time at 60% (fraction of total value settled) is statistically significant, and the difference is around 50 minutes. We also calculate the value-weighted average settlement time for a whole day as

$$\tilde{T} = \frac{\sum_{k=1}^{K} P_k T_k}{\sum_{k=1}^{K} P_k} \tag{3}$$

---

[15] See Figure 12 in the Appendix for a histogram of the submission times of queued payments.

[16] Specifically, we transform a given time t by subtracting the start time 00:30 and dividing it by (18:30-00:30).

where the denominator is the total value sent through the system, and the numerator is the total settlement time weighted by payment value. Across participants, the mean of value-weighted average settlement time is around 12:00 prior to September 16, and it is 12:20 after the day, implying a roughly 20-minute delay.

5.4.5 Intraday Liquidity Usage and Efficiency

Intuitively, the creation and resolution of gridlocks should increase the liquidity efficiency of the system. To quantify the potential efficiency gain, we first calculate the intraday liquidity used by bank i in a day. Given a bank's incoming and outgoing payments, the amount of intraday liquidity used by that bank is equal to the amount of liquidity that the participant needs to have in place, in order to meet its payment obligations for the day. One common and ex-post measure of a bank i's payment obligations is the maximum cumulative net debit position that this bank attains during a day, i.e.,

$$N_i \equiv \max_t \{n_i(t), 0\}, \tag{4}$$

where $n_i(t)$ is the net debit position of bank i at the t-th second of the day

$$n_i(t) = \sum_{s=1}^{t} \sum_{i \neq j} \left[p^{i,j}(s) - p^{j,i}(s) \right], \tag{5}$$

and $p^{i,j}(s)$ is the amount that bank i sends to bank j in s-second of the day.

Given the amount of intraday liquidity used by bank i, N_i, we can then calculate the amount of intraday liquidity used in the system by summing across all participating banks:

$$\text{Total intraday liquidity used} = \sum_i N_i. \tag{6}$$

Table 3 shows the total intraday liquidity used in the system, which is around 25 billion CAD before September 16, and 21.9 billion CAD afterwards, though the difference is not

Figure 6: Time series of liquidity efficiency ratio

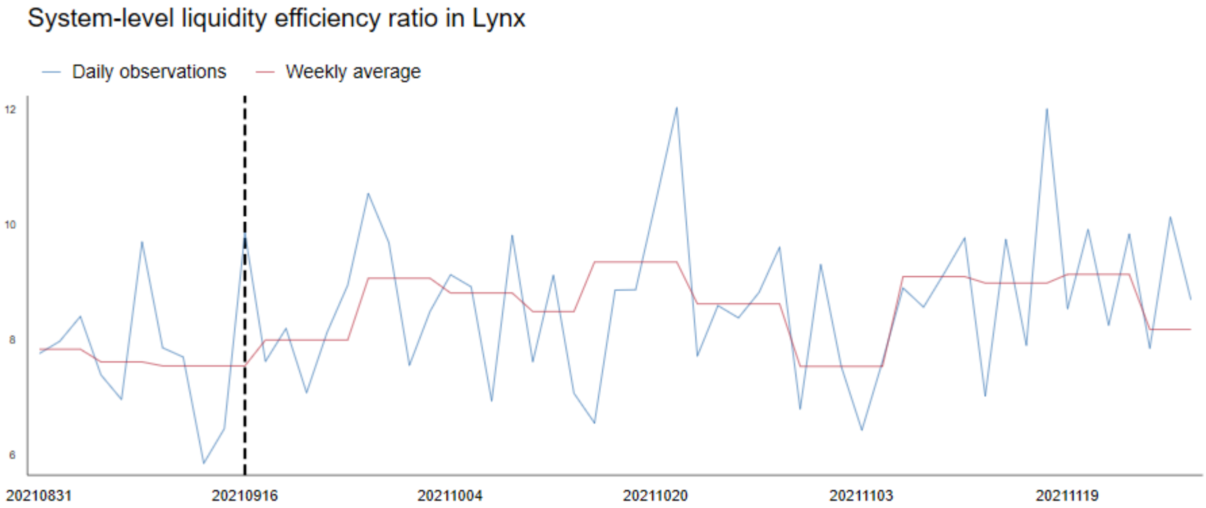

.

This figure shows the daily observations of liquidity efficiency ratio in Lynx, as well as the weekly average. The liquidity efficiency ratio in a day is calculated in Equation (7).

statistically significant. Next we use the ratio of total intraday liquidity used to total payment values to measure liquidity efficiency, as in Benos et al. (2018):

$$\text{Liquidity efficiency ratio} = \frac{\text{Total payment value settled}}{\text{Total intraday liquidity used}} \tag{7}$$

Figure 6 shows the time series of liquidity efficiency ratio in Lynx over a larger sample. The blue line shows the daily observations of liquidity efficiency ratio, and the red line shows the weekly average. The weekly average liquidity efficiency ratio rises almost immediately after the change on September 16, and continues to increase afterwards. From Table 3, the post-September 16 liquidity efficiency ratio is on average 15% higher than the pre-September 16 level and is statistically significant. This is consistent with our expectation that gridlock creation and resolution help settle more payments using fewer intraday liquidity, and thus the system becomes more liquidity-efficient after the joint effort of liquidity saving on September 16.

6 Conclusion

This paper documents banks' liquidity management behavior during the first month of Lynx, Canada's new HVPS. As participants transitioned from the legacy LVTS to this new RTGS system with gridlock resolution features, participants quickly realize that, without a voluntary queuing functionality in Lynx, they have to keep liquidity levels low and create gridlock in order to get payments queued and achieve liquidity savings. This revelation led to a simultaneous change in their liquidity management behaviours, which lowered system liquidity requirements by more than 70% on average by doubling the queuing of payments. Average payment settlement time was delayed by less than half hour while liquidity efficiency was drastically improved. We explain this finding by providing a game-theoretic model where such joint action is an outcome of a sustainable equilibrium.

The joint effort to reduce liquidity happened during a time when the central bank was injecting large amounts of excess liquidity in the system. Still, banks' incentive to save liquidity is strong, and we expect even more liquidity management effort going forward when the central bank reduces the liquidity injection. The implied trade-off between delay in payment timing and liquidity efficiency provides policy guidance for payment system oversight and implementation of monetary policy.

Our analysis also provides insights for the design of LSMs in HVPSs. Our finding that participants engage in liquidity management to micro-manage their queued payments shows their desire for a built-in voluntary queuing functionality, which allows any payment to be queued regardless of participants' liquidity levels. In the current system, a participant must ensure a low liquidity level in their account to queue payments. This calls for manual transferring of funds and thus incurs a cost of liquidity management. A voluntary queuing functionality allows participants to specify payments that they want to be queued, along with the condition(s) of releasing from the queue. The specification can be managed in a batch, or automated by applying predefined criteria. Such flexibility can potentially encourage participants to integrate their internal management of payment flow with external queues in the HVPS, which can improve the overall efficiency of the system.

References

Alexandrova-Kabadjova, Biliana, Anton Badev, Saulo Benchimol Bastoso, Evangelos Benos, Freddy Cepeda-López, James Chapman, Martin Diehl, Ioana Duca-Radu, Rodney Garratt, Ronald Heijmans, Anneke Kosse, Antoine Martin, Thomas Nellen, Thomas Nilsson, Jan Paulick, Andrei Pustelnikov, Francisco Rivadeneyra, Mario Rubem do Coutto Bastoso, and Sara Testi, "Intraday liquidity around the world," *Mimeo*, 2022.

Arjani, Neville and Darcey McVanel, "A primer on Canada's large value transfer system," Technical Report, Bank of Canada 2006.

Atalay, Enghin, Antoine Martin, and James McAndrews, "Quantifying the benefits of a liquidity-saving mechanism," *FRB of New York Staff Report*, 2010, (447).

Bech, Morten and Rodney Garratt, "The intraday liquidity management game," *Journal of Economic Theory*, 2003, *109* (2), 198–219.

_ **and** _ , "Illiquidity in the interbank payment system following wide-scale disruptions," *Journal of Money, Credit and Banking*, 2012, *44* (5), 903–929.

Bech, Morten L and Bart Hobijn, "Technology diffusion within central banking: The case of real-time gross settlement," *International Journal of Central Banking (September)*, 2007.

_ **, Christine Preisig, and Kimmo Soramaki**, "Global trends in large-value payments," *Economic Policy Review*, 2008, *14* (2).

Benos, Evangelos, Rodney J Garratt, and Peter Zimmerman, "The role of counterparty risk in CHAPS following the collapse of Lehman Brothers," *37th issue (December 2014) of the International Journal of Central Banking*, 2018, pp. 147–181.

Desai, Ajit, Zhentong Lu, Hiru Rodrigo, Jacob Sharples, Phoebe Tian, and Nellie

Zhang, "A quantitative assessment of the transition from LVTS to Lynx," Technical Report, Bank of Canada 2022.

Galbiati, Marco and Kimmo Soramaki, "Liquidity-saving mechanisms and bank behaviour," 2010.

Goyal, Sanjeev, "Chapter 15 - Learning in Networks," in Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, eds., *Handbook of Social Economics*, Vol. 1, North-Holland, 2011, pp. 679–727.

_ **and Fernando Vega-Redondo**, "Network formation and social coordination," *Games and Economic Behavior*, 2005, *50* (2), 178–207.

Jurgilas, Marius and Antoine Martin, "Liquidity-saving mechanisms in collateral-based RTGS payment systems," *Annals of Finance*, 2013, *9* (1), 29–60.

Martin, Antoine and James McAndrews, "Liquidity-saving mechanisms," *Journal of Monetary Economics*, 2008, *55* (3), 554–567.

_ **and** _ , "A study of competing designs for a liquidity-saving mechanism," *Journal of Banking & Finance*, 2010, *34* (8), 1818–1826.

McAndrews, James and Simon Potter, "Liquidity effects of the events of September 11, 2001," *Economic Policy Review*, 2002, *8* (2).

Mills, David C. and Travis D. Nesmith, "Risk and concentration in payment and securities settlement systems," *Journal of Monetary Economics*, 2008, *55* (3), 542–553.

Rivadeneyra, Francisco and Nellie Zhang, "Payment coordination and liquidity efficiency in the new Canadian wholesale payments system," 2022.

Roberds, William, *The incentive effects of settlement systems: a comparison of gross settlement, net settlement, and gross settlement with queuing*, Institute for Monetary and Economic Studies, Bank of Japan, 1999.

Appendix: Additional Figures

Figure 7: Real-time balance and outgoing payments in one bank's RTGS^G stream account

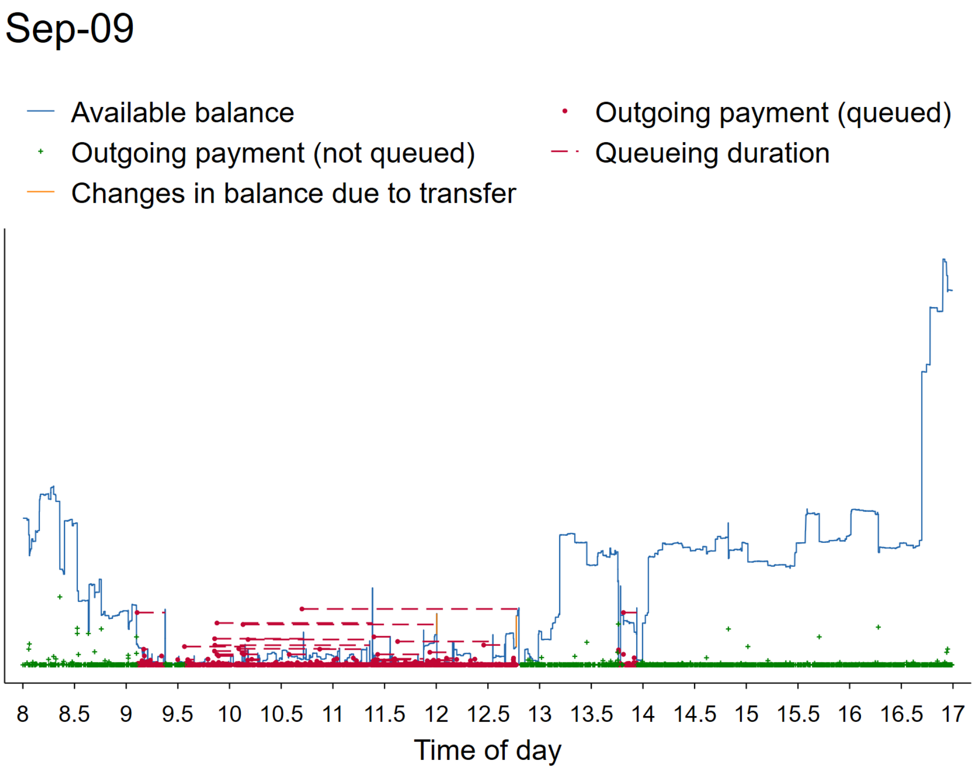

See note below Figure 1

Figure 8: Real-time balance and outgoing payments in one bank's RTGS^G stream account

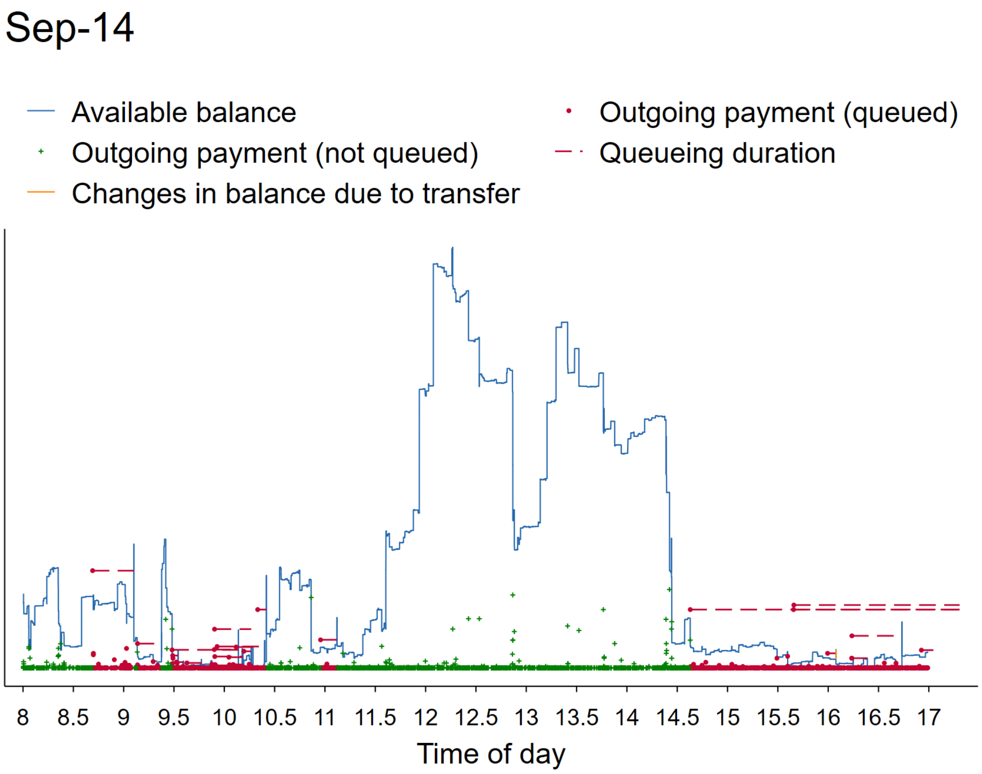

See note below Figure 1

Figure 9: Real-time balance and outgoing payments in one bank's RTGS^G stream account

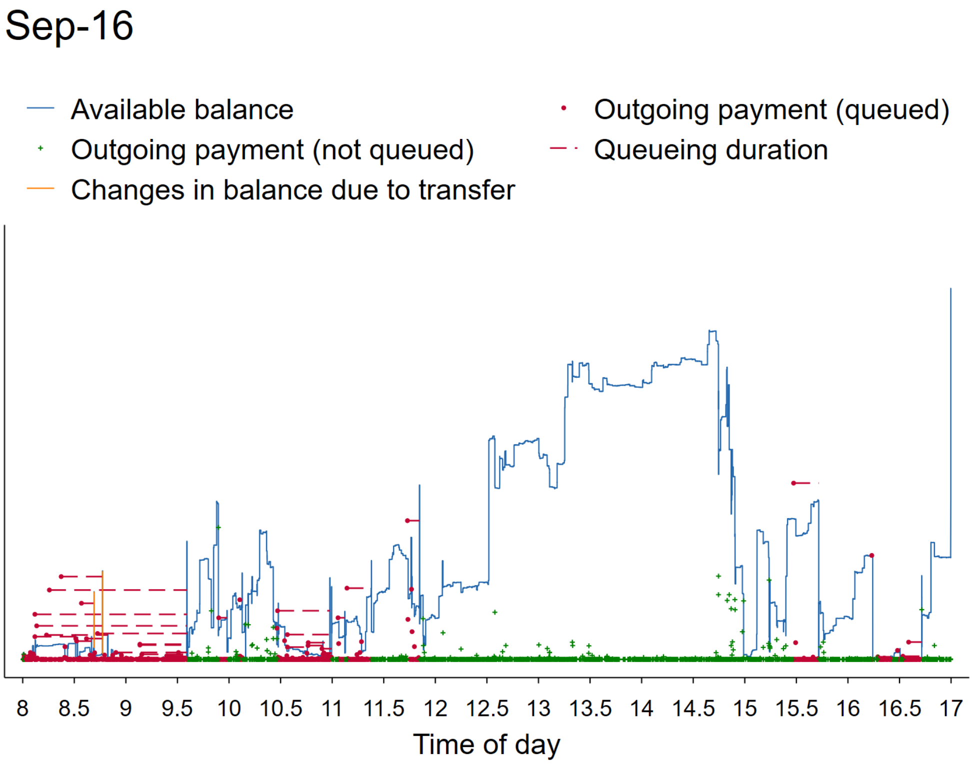

See note below Figure 1

Figure 10: Real-time balance and outgoing payments in one bank's RTGS^G stream account

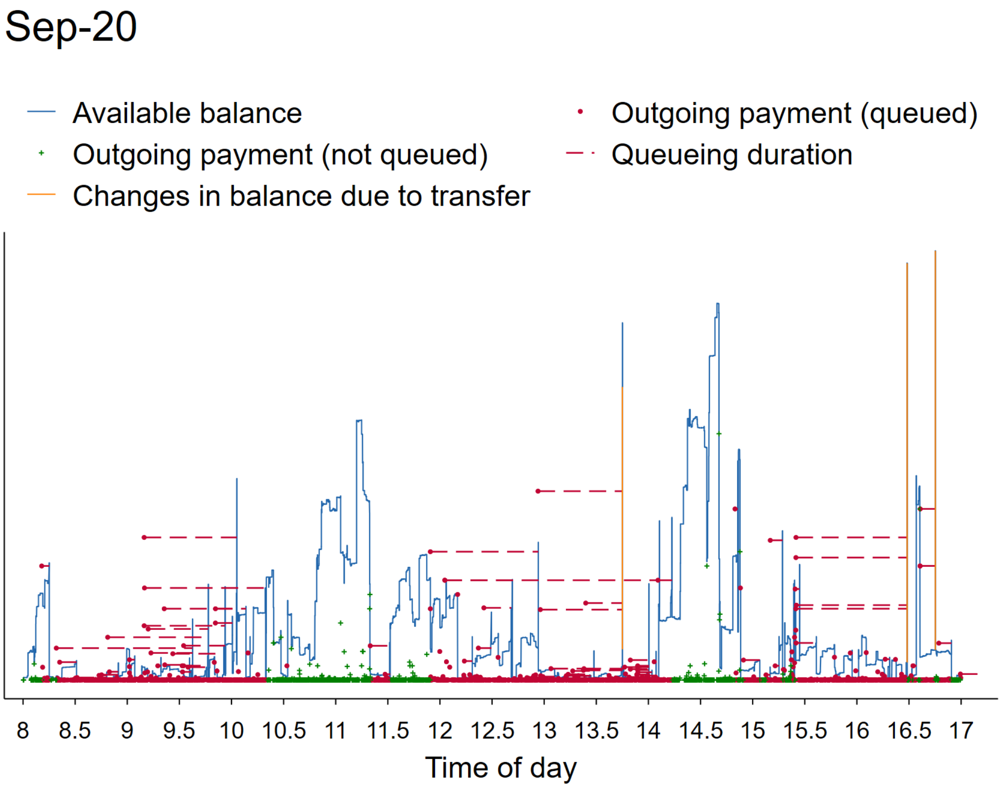

See note below Figure 1

Figure 11: Real-time balance and outgoing payments in one bank's RTGS^G stream account

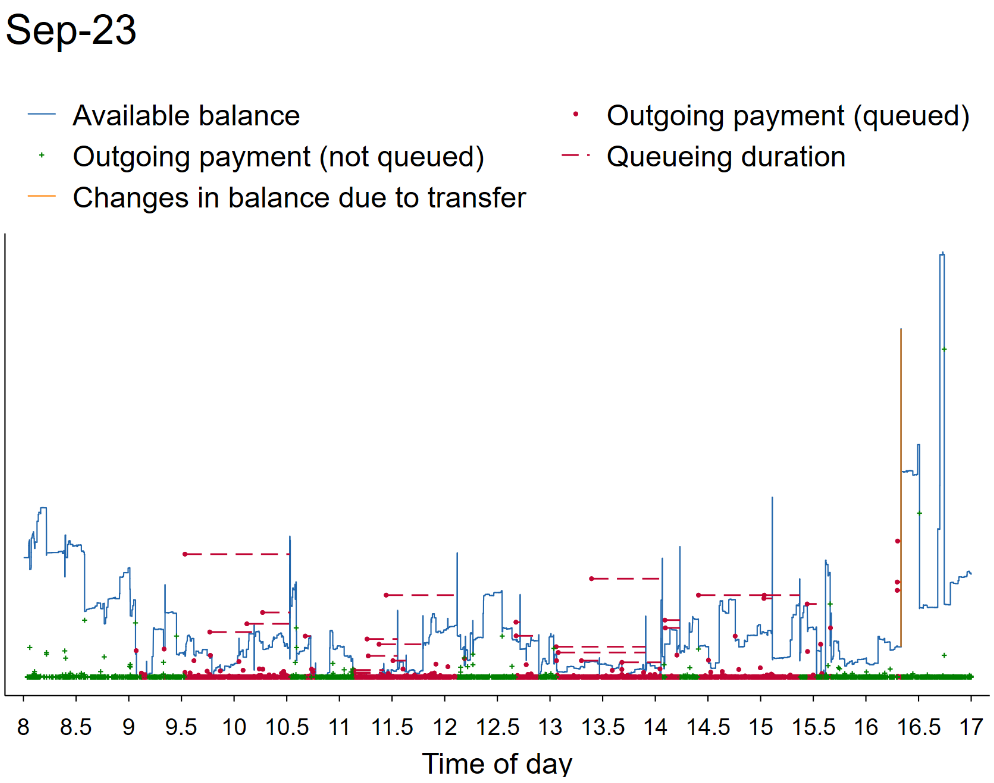

See note below Figure 1

Figure 12: Histogram of submission time of queued payments

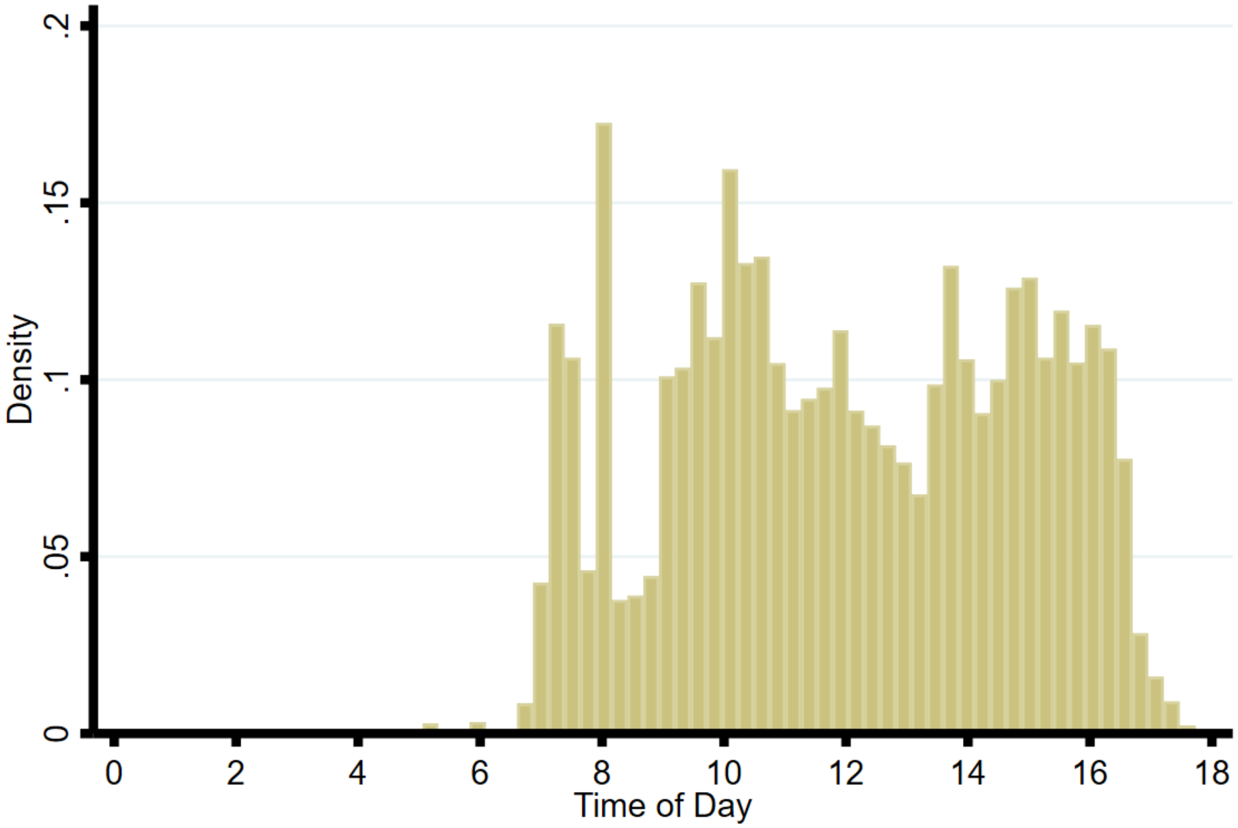

This graph shows the distribution of submission time of queued payments from August 31 to November 30, 2021.

Figure 13: Number of intraday internal transfers in Lynx

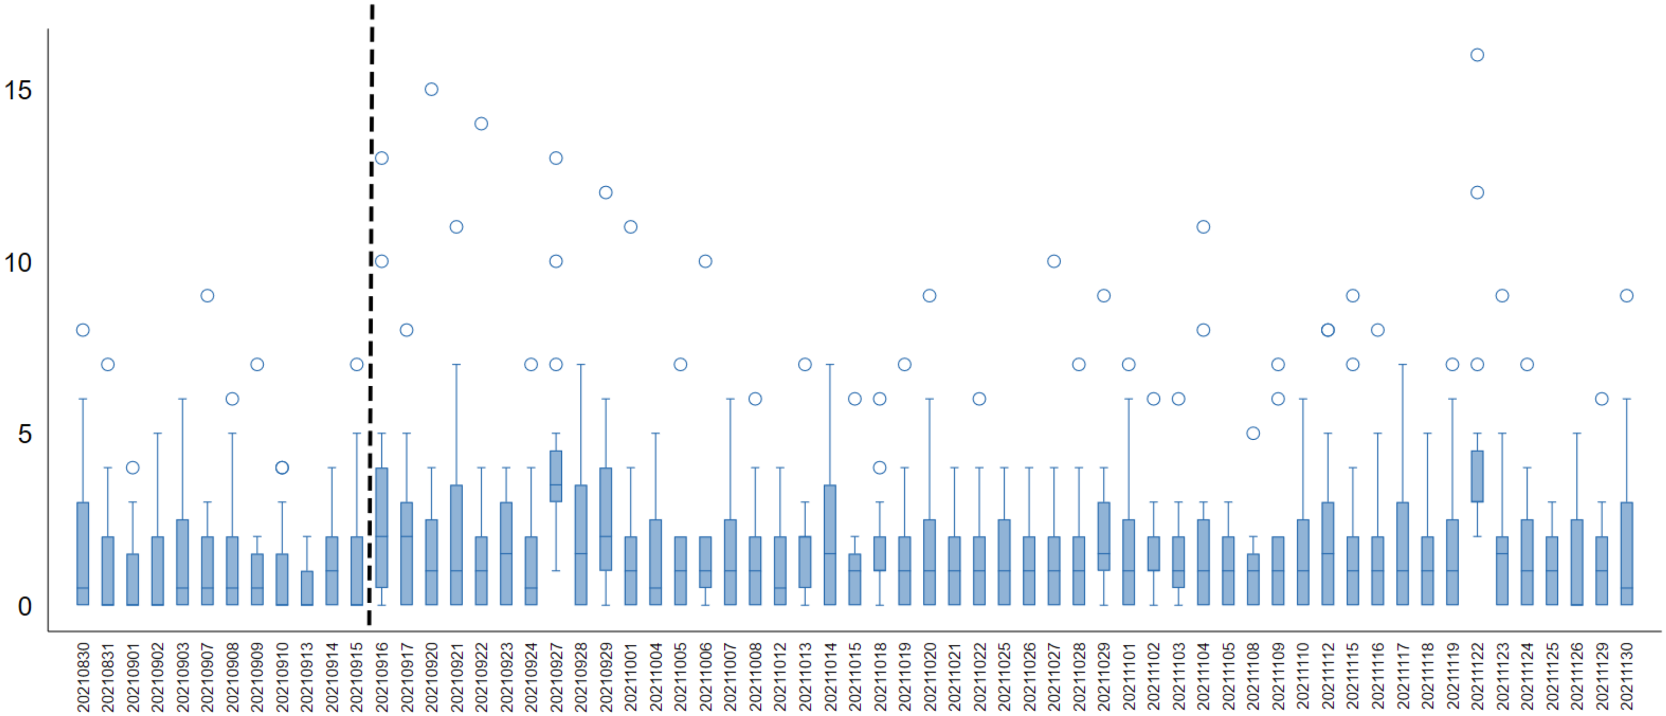

This is the box plot of the number of intraday internal transfers made by all participants (excluding BoC) from August 31 to November 30, 2021. Within each day, the box covers the 25th percentile to the 75th percentile, with the middle line of the box indicating the median. The vertical line covers the value from the 75th percentile plus 1.5 times the interquartile range (the difference between the 75th and the 25th percentile, i.e., the length of the box), to the 25th percentile minus 1.5 times the interquartile range. Scatter dots are outliers. An intraday internal transfer is defined as an transfer of fund within one participant's accounts from 8AM to 5PM.

Figure 14: Average account balance in the RTGS^G stream over a day

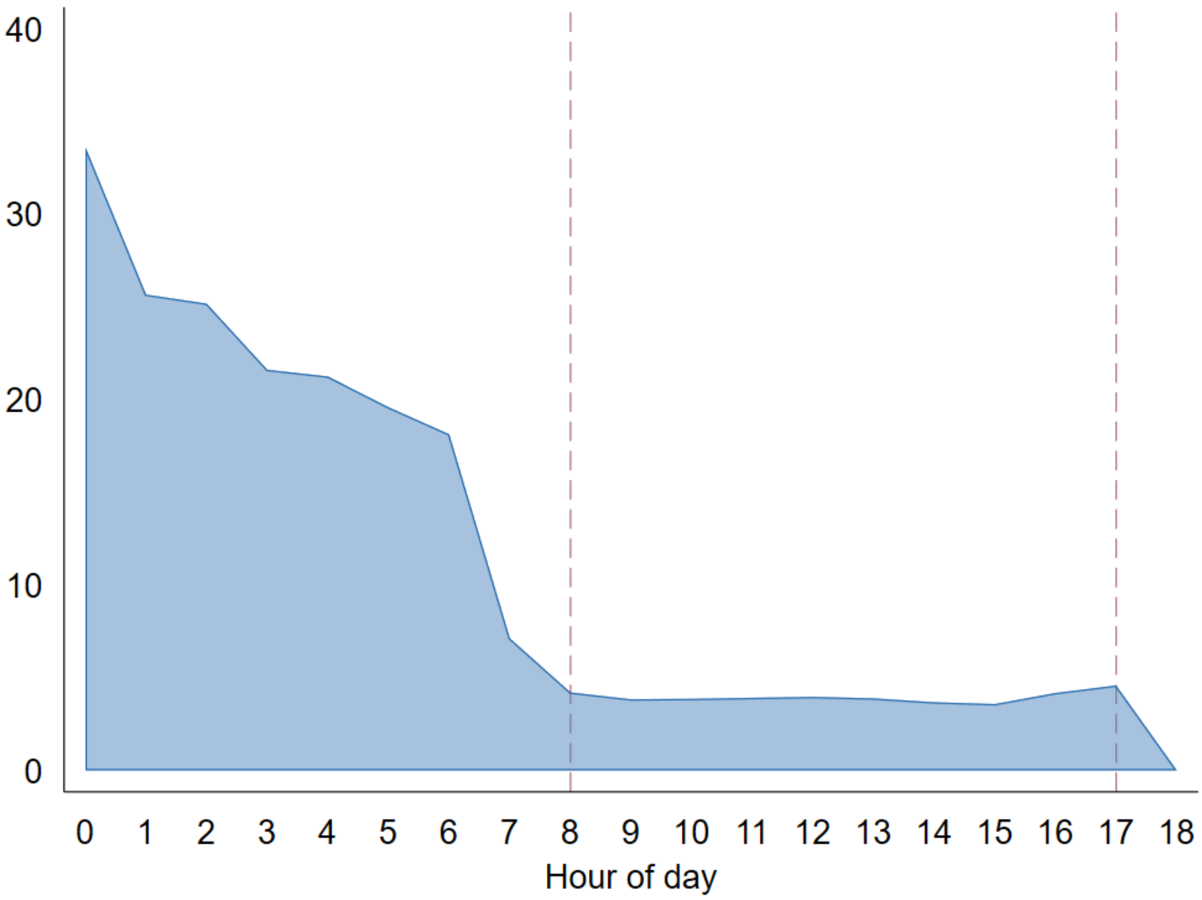

This graph shows the average hourly balance in RTGS^G accounts over a business day. Values are stated in CAD bn. The figure is obtained by averaging over all participants.

Figure 15: Decomposition of flows in and out of the RTGS^G stream during a typical day

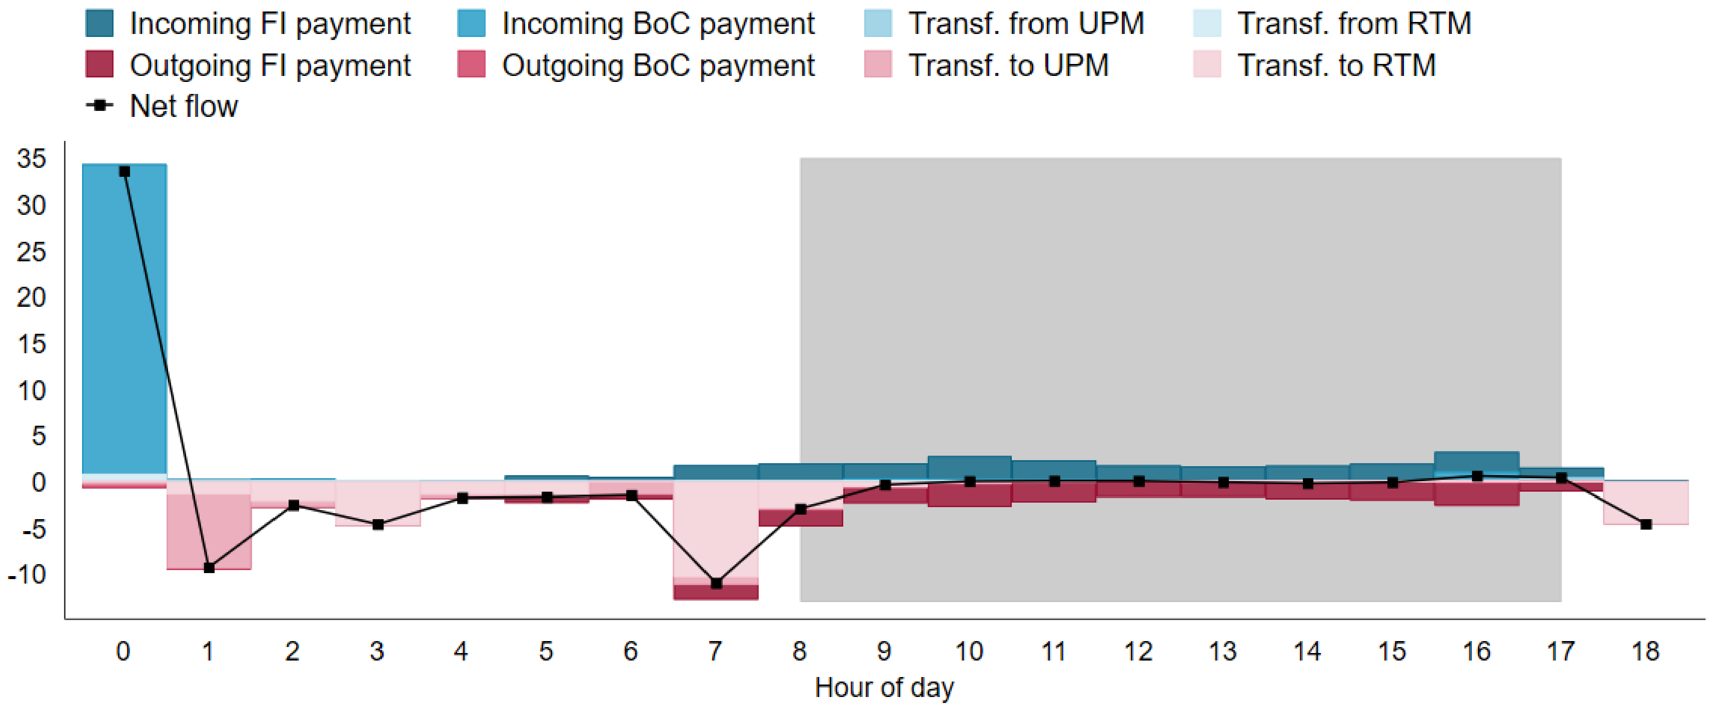

This graph shows the sources of inflows into the RTGS^G stream and destinies of outflows from the RTGS^G stream, by each hour during a day. “UPM” is synonymous with the pure RTGGSstream in this paper, and “RTM” is another account that banks typically use to pay and receive intraday loan from the central bank. The graph is obtained by averaging across participants. Values are stated in billion CAD.

Figure 16: Number of queued payment and gridlock buster runs

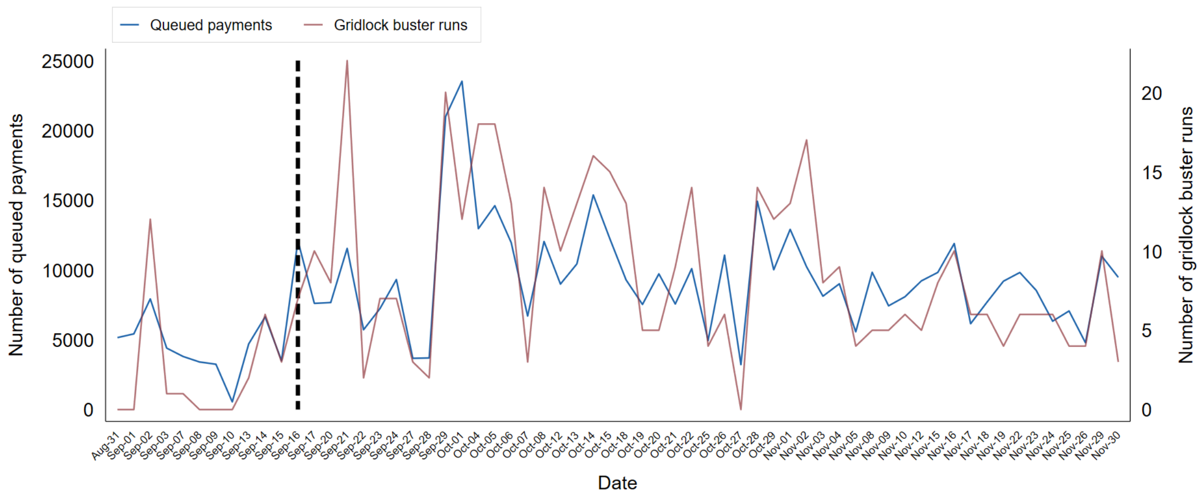

This graph shows the number of queued payments and gridlock buster runs from August 31 to November 30, 2021, with the black dashed line indicating the day of change in behavior, September 16.