Bank of Canada Banque du Canada

Working Paper 2006-20 / Document de travail 2006-20

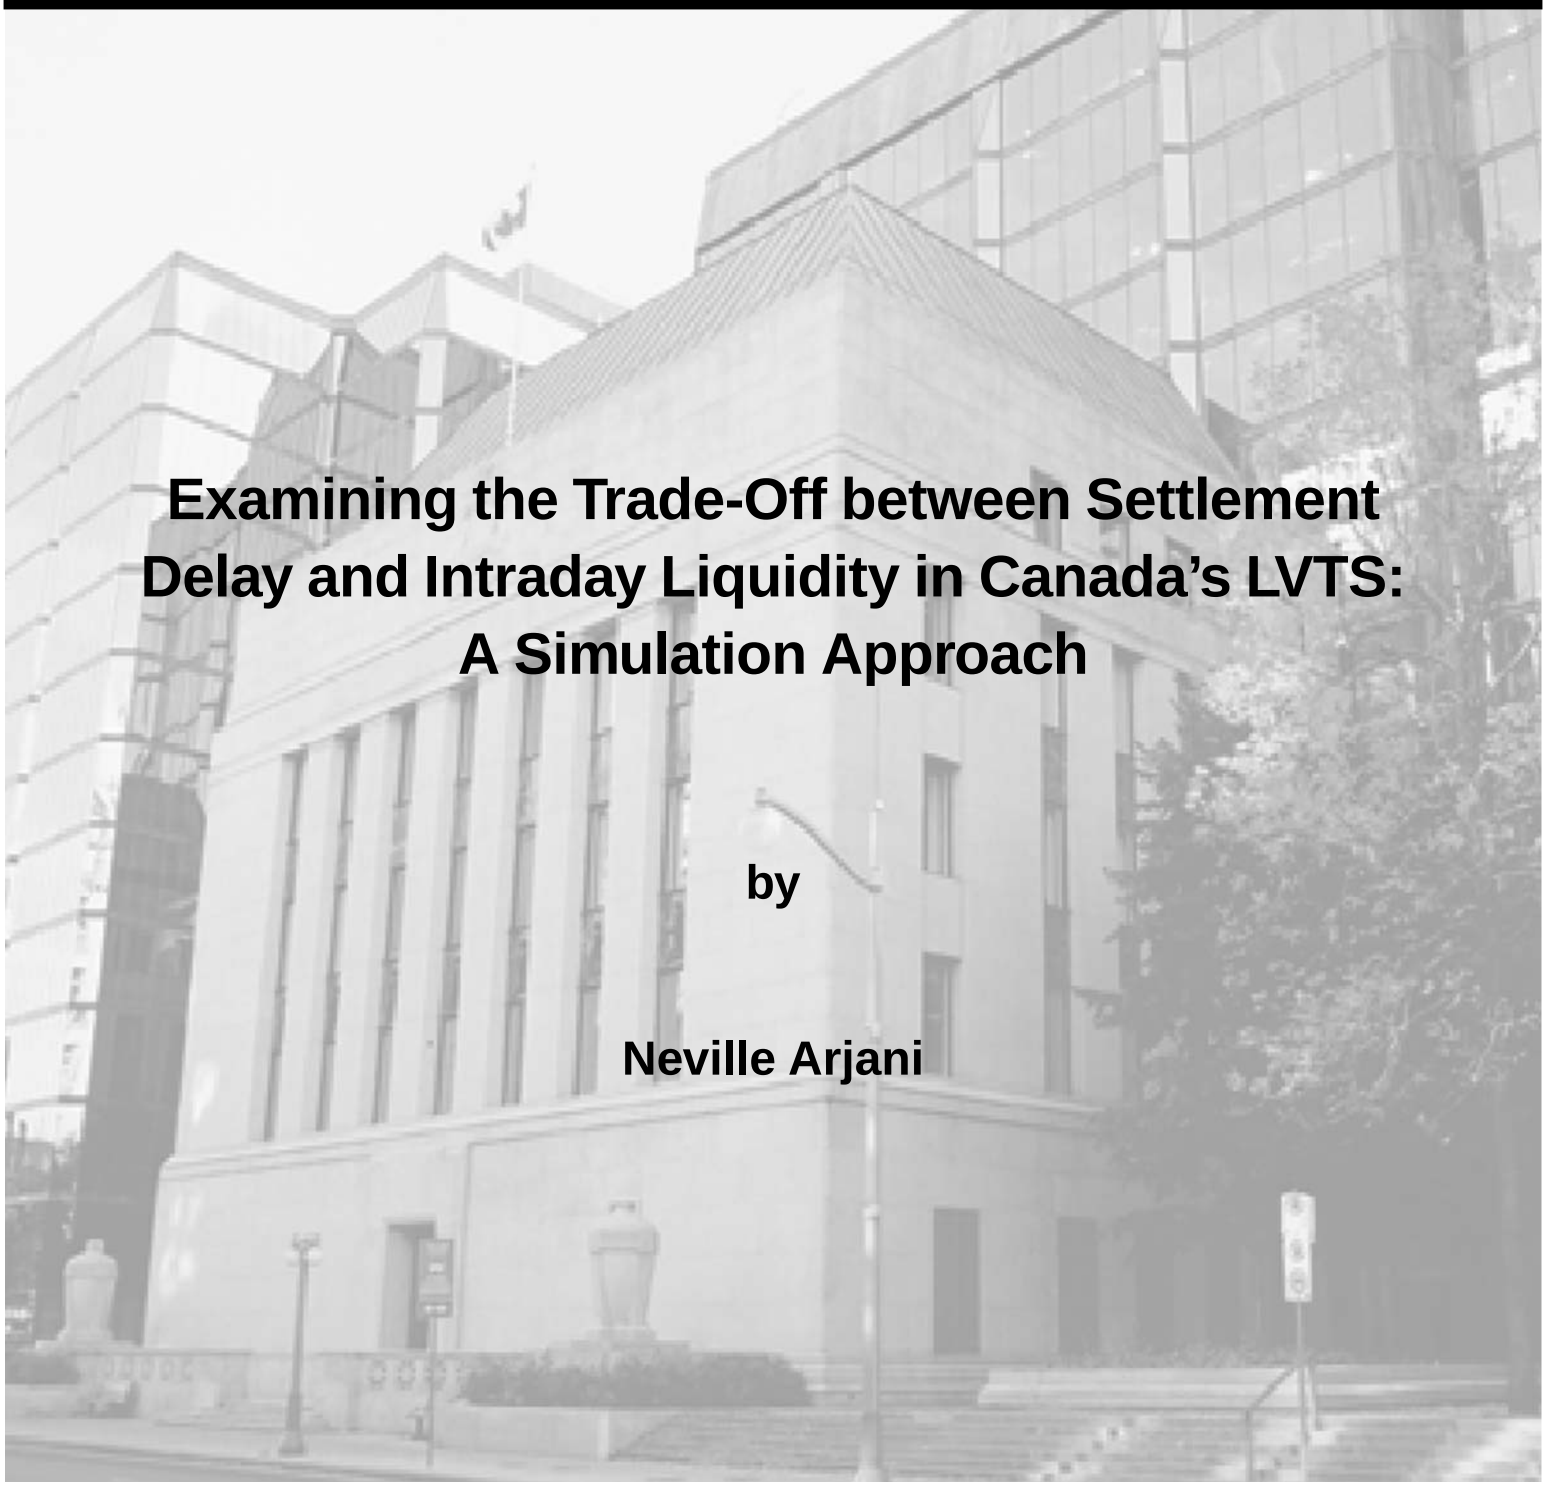

Examining the Trade-Off between Settlement Delay and Intraday Liquidity in Canada's LVTS: A Simulation Approach

by

Neville Arjani

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2006-20

June 2006

Examining the Trade-Off between Settlement Delay and Intraday Liquidity in Canada's LVTS: A Simulation Approach

by

Neville Arjani

Monetary and Financial Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
narjani@bankofcanada.ca

Contents

Acknowledgements

The author would like to thank Sean O'Connor, Walter Engert, Louise Hyland, Kim McPhail, Dinah Maclean, Lorraine Charbonneau, and seminar participants at the Bank of Canada and the Bank of Finland for their helpful comments and suggestions on this and earlier versions of the paper. Special thanks to the Bank of Finland, and especially Harry Leinonen, for developing the BoF-PSS2 payment system simulator and allowing other central banks to use it. The author is also grateful to the Canadian Payments Association for supplying the data used in this project.

Abstract

The author explores a fundamental trade-off that occurs between settlement delay and intraday liquidity in the daily operation of large-value payment systems (LVPS), with specific application to Canada's Large Value Transfer System (LVTS). To reduce settlement delay, participants generally must maintain greater intraday liquidity in the system. Intraday liquidity and settlement delay can be costly for LVPS participants, and improvements in the trade-off are desirable. The replacement of standard queuing arrangements with a complex queue-release algorithm represents one such improvement. These algorithms are expected to lower intraday liquidity needs and speed up payments processing in an LVPS. Simulation analysis is used to empirically test this proposition for the case of Canada's LVTS. The analysis is conducted using a payment system simulator developed by the Bank of Finland, called the BoF-PSS2. The author shows that increased use of the LVTS central queue (which contains a complex queue-release algorithm) reduces settlement delay associated with each level of intraday liquidity considered, relative to a standard queuing arrangement. Some important issues emerge from these results.

JEL classification: E47, G21
Bank classification: Payment, clearing, and settlement systems

Résumé

L'auteur examine un rapport d'arbitrage fondamental qui existe entre retard de règlement et liquidités intrajournalières dans les opérations quotidiennes traitées par les systèmes de transfert de gros paiements, notamment dans le cadre du Système de transfert de paiements de grande valeur (STPGV). Pour réduire le retard de règlement, les participants doivent généralement maintenir un degré élevé de liquidités intrajournalières dans le système. La tenue de ces liquidités et le retard de règlement pouvant être coûteux pour les participants aux systèmes de paiement, il est souhaitable que l'arbitrage soit amélioré. Le remplacement des modes habituels de gestion des opérations en attente par un régime doté d'un algorithme complexe de retrait de la file d'attente constitue l'une des améliorations possibles. Ce genre d'algorithme devrait permettre une baisse des besoins en liquidités intrajournalières et un traitement accéléré des opérations dans les systèmes de paiement. L'auteur recourt à une analyse par simulation pour tester empiriquement l'application de ce concept au STPGV du Canada. L'analyse est réalisée au moyen d'un simulateur de système de paiement, le BoF-PSS2, développé par la Banque de Finlande. L'auteur montre que le recours accru à une file d'attente centralisée dans le STPGV (qui contient un algorithme complexe de retrait de la file) a pour effet de diminuer le retard de règlement associé à chaque niveau de liquidités intrajournalières considéré, par rapport à un mode standard de gestion de la file d'attente. Les résultats obtenus mettent en lumière des questions importantes.

Classification JEL : E47, G21
Classification de la Banque : Systèmes de paiement, de compensation et de règlement

1. Introduction

A well-functioning large-value payment system (LVPS) is an integral component of any advanced financial system. In a market economy such as Canada's, virtually all economic transactions ultimately involve a transfer of funds between a buyer and a seller. An LVPS provides the electronic infrastructure necessary to facilitate such an exchange of funds between financial institutions in order to discharge large-value payment obligations on behalf of their own business and that of their customers. There are different designs of LVPS currently operating around the world, with each achieving a different balance between the minimization of systemic risk, the speed of payment settlement, and the liquidity and operational costs of settlement.

This paper examines a fundamental trade-off that occurs between settlement delay and intraday liquidity in the daily operation of an LVPS, with particular application to Canada's Large Value Transfer System (LVTS).[1] Settlement delay refers to a potential time lag occurring between a participant's intended submission of a payment to the system and when it is processed by the LVPS with finality.[2] Intraday liquidity refers to a participant's ability to meet its outgoing payment obligations immediately when intended. Generally speaking, to achieve shorter settlement delay, participants must maintain greater intraday liquidity in the system. When sufficient intraday liquidity is not maintained, payments will be queued and will be released only when the participant's liquidity position improves. Settlement delay, then, reflects the amount of time that a payment is queued before being processed by the system.

Intraday credit is an important source of liquidity. To control credit risk, grantors of intraday credit (typically, central banks) usually require eligible collateral, which is likely to entail a cost for participants. At the same time, settlement delay may also be expensive

[1] The LVTS is owned and operated by the Canadian Payments Association (CPA). For a more thorough description of the LVTS, including an overview of the Bank of Canada's multiple roles within the system, see Dingle (1998), and Arjani and McVanel (2006, forthcoming).

[2] Use of the term 'intended' is made so that this definition of settlement delay could apply to LVPS designs with and without a central queue. Under the latter design, a participant may intend to submit a payment to the LVPS at a certain time but, due to lack of intraday liquidity and the absence of a central queue, must hold the payment internally until it can be successfully processed by the system.

for participants. The cost of settlement delay may be borne both internally by the participant that delays sending the payment, and externally by the receiving participant. Participants generally must trade off the cost of settlement delay and the cost of intraday liquidity in conducting their daily payment operations. It follows that a reduction in the amount of intraday credit provision to participants will entail both a benefit and a cost: participants' liquidity (i.e., collateral) cost can be reduced, but possibly only at the expense of a higher settlement delay cost.

A simple graphical framework of the general risk-efficiency trade-off in payment systems, inspired by Berger, Hancock, and Marquardt (1996), is useful when thinking about the nature of the trade-off between settlement delay and intraday liquidity in an LVPS. Given the cost to participants of both settlement delay and intraday liquidity, improvements in the trade-off are desirable. An improvement is characterized by reduced settlement delay associated with each level of intraday liquidity, for the same value of payment activity. Innovations in LVPS design may make this possible. The replacement of standard queuing arrangements with a complex queue-release algorithm represents such an innovation. The potential benefit of such algorithms includes both lower liquidity needs for the release of queued payments and faster processing of these payments by the LVPS.

A simulation approach is used to empirically test the proposition that a complex queue-release algorithm can lower liquidity costs and speed payments processing relative to a standard queuing arrangement. Using actual intraday transaction and credit limit data, simulation analysis is employed to assess the nature of the current trade-off between settlement delay and intraday liquidity in the Canadian LVTS. Further, improvements in this trade-off can be sought by simulating an alternative LVTS environment in which current restrictions on use of the LVTS central queue are relaxed. The LVTS queue employs a complex queue-release algorithm that seeks to partially offset batches of queued payments on a multilateral basis throughout the day. However, under current system rules, participants' excessive use of the central queue is not encouraged.[3] Instead, standard internal queuing arrangements are typically employed by participants.

The analysis reveals that a trade-off does indeed exist between settlement delay and intraday liquidity in Canada's LVTS. Moreover, the results indicate that increased use of the central queue will reduce settlement delay in the LVTS for each level of intraday liquidity considered, according to three different settlement delay measures. Some important issues emerge from these results.

The remainder of this paper is as follows. Section 2 discusses the nature of the trade-off between settlement delay and intraday liquidity in greater detail. The graphical framework is presented in section 3, and potential improvements in the trade-off are also discussed in that section. Section 4 contains relevant background information on the LVTS. Section 5 provides a description of the data and an overview of the simulation methodology. Section 6 presents results from the simulations and a related discussion. Section 7 offers concluding remarks and some caveats to the analysis.

2. Settlement Delay and Intraday Liquidity in an LVPS

Participants in an LVPS typically maintain a daily schedule of payments, which they must send through the system on behalf of their own business and that of their customers. Included in this schedule is the time that each payment is due to be sent. For example, certain payments are considered 'time-sensitive' and thus have to be sent by a specific time during the day. The remaining majority of payments is considered 'non-time-sensitive' and simply must be sent by the end of the day. In practice, however, participants generally do not wait until the end of the day to submit all of their non-time-sensitive payments, for reasons that will be outlined below.

[3] See LVTS Rule No. 7 on the Canadian Payments Association website, available at <http://www.cdnpay.ca>. There are several hypothesized reasons for this. Perhaps the foremost reason pertains to the issue of whether queue transparency may cause participants to take on credit risk by crediting clients' accounts with expected incoming funds, prior to these payments actually being received. This was a major concern of central banks at the time the LVTS was being developed (see BIS 1997, and the discussion in section 6.2 of this paper).

In real-time gross settlement (RTGS) and RTGS-equivalent LVPS (such as Canada's LVTS), participants must maintain intraday funds in the system to send a payment to another bank. Hence, the concept of intraday liquidity in an LVPS specifically refers to a participant's ability to access sufficient intraday funds to meet its outgoing payment obligations in a timely manner. There are two main sources of intraday funds available to an LVPS participant: (i) funds acquired from other participants due to either regular transaction activity or through an overnight interbank loan arrangement, and (ii) funds acquired through an intraday credit extension. Incoming funds from regular transaction activity are the cheapest source of liquidity for participants, and it is expected that participant banks will try to use these funds as much as possible to finance their own payment activity.[4] For various reasons (e.g., the differing nature of individual participants' business), however, it may not always be possible for participants to coordinate their daily payment activity so that incoming payments largely finance their outgoing payment needs.

The inability of participants to perfectly coordinate their incoming and outgoing payment activity creates a role for the provision of intraday credit. Martin (2005) emphasizes the importance of intraday credit as a source of intraday funding for participants. The author argues that the coordination of incoming payments to meet outgoing obligations is often difficult (especially for time-sensitive payments), and therefore a well-designed LVPS should allow participants to acquire funds when necessary through intraday credit. Where intraday credit is available to participants on a free and unlimited basis, participants can borrow funds any time that a payment is due, thus eliminating potential settlement delay in the LVPS. However, although settlement delay would cease to exist in this case, lenders of intraday credit (typically central banks) could face large risk exposures vis-à-vis borrowers, which is not desirable from a public policy perspective. Consequently, intraday credit in RTGS and equivalent systems is not free and unlimited, but rather is often subject to net debit caps, (eligible) collateral requirements that typically entail an opportunity cost, and in certain cases an explicit interest charge (e.g., the U.S. Fedwire

---

[4] See McAndrews and Rajan (2000) and McAndrews and Potter (2002) for discussion and identification of this type of coordination behaviour among participants in the U.S. Fedwire system.

system). Maintaining intraday liquidity in the system can therefore be costly for participants.

Where a participant does not have sufficient funds available to meet a payment obligation upon intended submission, processing of the payment by the LVPS will be delayed. Settlement delay can be defined as a time lag occurring between a participant's intended submission of a payment to the LVPS, and when the payment is processed by the LVPS with finality; i.e., when intraday funds are exchanged between participants on an unconditional and irrevocable basis in order to discharge the payment obligation.[5] Payments that cannot be processed because of a participant's lack of intraday liquidity may be held in that participant's internal queue. Alternatively, these payments could be submitted to the LVPS and held in the system's central queue if one is available. Under standard queuing procedures, internally and centrally queued payments are released and processed by the LVPS on an *individual* basis when a sending participant's intraday liquidity improves to the extent that these payments can be passed.[6] The settlement delay associated with an individual payment essentially reflects the amount of time that the payment must wait in the queue before being processed by the LVPS.

Figure 1 provides a graphical characterization of settlement delay within the context of the life cycle of a large-value payment.[7]

[5] A key feature of RTGS and equivalent LVPS is that these systems offer immediate intraday finality. Payments in these systems are considered final upon being processed.
[6] This liquidity improvement could occur as a result of the participant receiving a payment, or gaining access to more intraday credit.
[7] This paper recognizes that achieving payment finality need not encompass the transfer of the settlement asset. Therefore, the notion of settlement delay applies equally to RTGS and RTGS-equivalent LVPS, where this transfer occurs on a multilateral net basis at the end of the day in the latter.

Figure 1: The Life Cycle of a Large-Value Payment

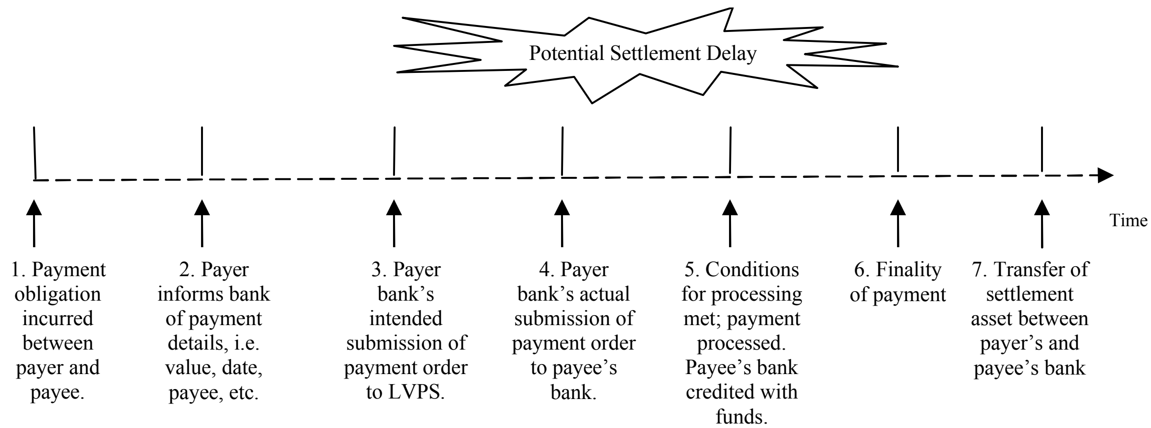

Just as there is a cost associated with maintaining intraday liquidity in the system, given the high speed and high value of daily payments processed by an LVPS, settlement delay may also entail a significant cost for participants. Further, the nature of this delay cost is likely to depend on whether a payment is time-sensitive or not. Time-sensitive payments may include those related to the final funds settlement of other important national and international clearing and settlement systems, large government receipts and disbursements, and also payments related to the daily implementation of monetary policy. A participant that is unable to meet a time-sensitive payment obligation when due may therefore face large internally borne costs because of the delay, such as reputation damage with its peers and, possibly, a loss of its clients' business. Explicit penalty charges may even be imposed by the system operator, since the delay of these payments could cause a disruption elsewhere in the financial system.

For the remaining majority of (non-time-sensitive) payments, there is no formal intraday deadline to submit these payments. It is not expected that a participant will incur an (immediate) reputation loss or penalty charge, nor a loss of its clients' business, if processing of these payments is delayed until the end of the day.[8] However, there may be

[8] Prolonged delay of non-time-sensitive payments is unlikely to cause reputation loss immediately, but such a loss could occur if repeated over time. In a relatively concentrated payments system like Canada's LVTS, participants maintain frequent communication with each other throughout the day, and are able to develop fairly accurate forecasts of certain incoming payment flows based on historical payment patterns with other participants. Thus, a participant that often delays its non-time-sensitive payments in favour of lower liquidity costs is unlikely to go unnoticed among its peers in the system.

other external costs imposed on the system in this case. Despite being non-time-sensitive, intended receiving banks may be expecting these payments by a certain time of day, and such a delay will result in a shortfall in their intraday funds position. If these participants are planning on using these funds to send their own payments, then they may have to incur additional liquidity costs in order to replace these funds on short notice. Where they cannot find other funds in time to meet their obligations, additional settlement delay is created in the system. Settlement delay created by one participant in an LVPS could quickly spread to others in the system. Moreover, a comparable disruption to the liquidity position of a receiving bank's client may also occur (where a delayed payment is ultimately intended for this customer), resulting in potentially broader consequences for economic activity.

Prolonged delay of these payments may also intensify the potential losses associated with other risks in the system, such as operational risk. An operational event (such as a computer outage that prevents one or more participants from sending payments) will likely have a larger impact in a case where a number of payments remain unprocessed at the time that the incident occurs.[9] At the same time, a large backlog of payments being submitted all at once to the LVPS late in the day could increase the potential likelihood that an operational event occurs in the first place. Lastly, where the potential for settlement delay could discourage use of an LVPS in favour of systems that are not as well risk proofed, the existence of settlement delay may translate to higher systemic risk in the broader financial system.

It follows that, to eliminate the potential costs associated with settlement delay, participants will likely have to borrow a large amount of intraday credit and thus incur high liquidity costs. Conversely, participants need not incur any intraday liquidity cost, but will then have to bear (possibly along with other participants in the system) the costs of the accompanying settlement delay. It is unlikely that participants will not maintain sufficient liquidity to meet their time-sensitive payment obligations, since the cost of

[9] Conversely, an operational disruption could also lead to settlement delay in an LVPS, since it may result in a participant's inability to send payments through the system. For this reason, contingency measures are usually available in an LVPS for the release of time-sensitive payments in the event of a disruption.

delaying these payments is very high. Consequently, the discussion of a trade-off between settlement delay and intraday liquidity may not apply to time-sensitive payments in practice. However, for non-time-sensitive payments, the trade-off is likely to exist. Since settlement delay may entail costs and repercussions for the system as a whole, any innovation in LVPS design that can increase settlement speed for a given level of intraday liquidity is desirable.

3. A Simple Graphical Framework

3.1 Description of the framework

The expected relationship between settlement delay and intraday liquidity in an LVPS is illustrated in Figure 2. Figure 2 is inspired by the concept of an 'efficient frontier' presented by Berger, Hancock, and Marquardt (1996).[10] This framework will help in interpreting the empirical results later in the paper.

Figure 2: The LVPS Delay-Liquidity Efficient Frontier

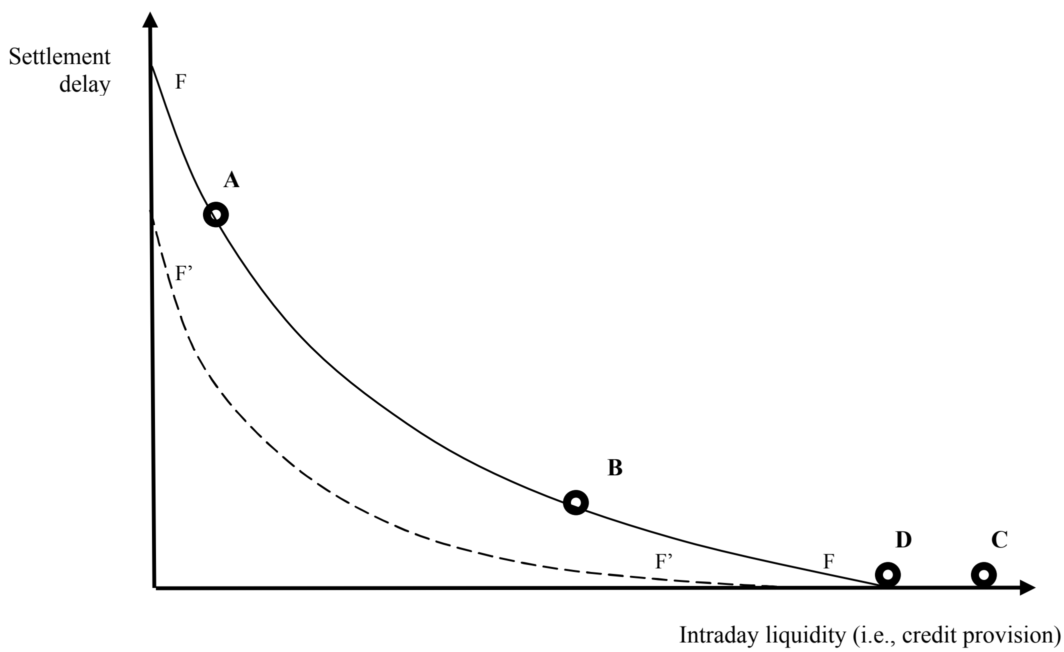

[10] In describing this framework, the terms 'intraday liquidity' and 'intraday credit' are used synonymously.

The framework is presented in delay-liquidity space. All points in the space represent possible settlement delay-intraday liquidity combinations necessary to produce a given level of payment activity. The vertical axis measures the magnitude of overall settlement delay in the LVPS, while the horizontal axis measures the provision of intraday credit. It is useful to think of the magnitude of settlement delay in an LVPS as reflecting both the number of payments entering the queue upon intended submission, and also each payment's duration in the queue until being processed. The trade-off is captured by the curve denoted FF, and this curve is generated based on the existing technology for processing payments (i.e., the existing LVPS design). Specifically, the curve shows how settlement delay and intraday credit provision can be traded off against each other for a given level of payment activity under current LVPS arrangements. The slope of FF captures the reduction in settlement delay that can be achieved by participants following a unit increase in the provision of intraday credit.

The decreasing convex shape of the trade-off curve reflects the assumption of diminishing marginal returns to liquidity. An increase in intraday credit provision is anticipated to have a lesser impact in terms of reduced settlement delay when moving further along the frontier from left to right. This assumption is attributed to the positively skewed nature of the distribution of individual payment values in an LVPS.[11] At a very low level of liquidity (point A), a small increase in intraday credit provision will lead to a higher reduction in settlement delay, since many smaller payments that would otherwise have been delayed can now be immediately processed upon intended submission. As intraday credit provision is continuously increased, it is expected that more payments will be processed upon intended submission and the delayed finality of these payments will be averted. However, even at higher levels of intraday credit provision (such as point B), it is expected that a few very large payments will still be delayed. Only a substantial injection of intraday credit would allow these payments to be processed immediately.

[11] For instance, in Canada's LVTS, the average payment value is around Can$7.5 million, while the median value is around Can$50,000. Moreover, the value of some payments in the LVTS is well over Can$100 million.

All combinations along the curve, and also above and to the right of the curve, represent feasible combinations of settlement delay and intraday liquidity for a given level of payment activity under the existing LVPS design. The trade-off curve is the most technologically efficient of these feasible combinations and, therefore, an LVPS is considered to be technically efficient if it is processing payments anywhere along the curve. This notion of efficiency captures the idea that, when operating along the curve, reductions in settlement delay can only be achieved by an increase in intraday credit provision, and vice versa, for a given level of payment activity. Processing the same level of payment activity at a point above, or to the right of, the trade-off curve represents inefficiency. For instance, producing at a point like C in Figure 2 means that intraday credit provision could be reduced and participants' liquidity costs lowered without causing any increase in settlement delay. In fact, intraday credit provision could be lowered from point C all the way to point D before any further reductions would lead to increased settlement delay in the LVPS. Point D represents the familiar upper bound of liquidity as described in Leinonen and Soramaki (1999, 2003). Points below the efficient frontier are currently unattainable given the existing LVPS technology, and can only be achieved through some form of innovation.

3.2 Innovation: A complex queue-release algorithm

Points below the trade-off curve are not attainable given the existing LVPS technology. An improvement that allows lower settlement delay for any given level of intraday liquidity, or vice versa, is required to attain such an outcome. The impact of this improvement appears in Figure 2 as a shift of the trade-off curve FF to its new position closer towards the origin at F'F'. Along the new curve, the same amount of payment activity can be produced with lower settlement delay for each level of intraday liquidity, and therefore at a lower overall cost to participants.

Such an improvement can be achieved through a technological innovation in LVPS design. Reductions in settlement delay can be achieved through either faster processing of queued payments or fewer payments entering the queue upon submission, where the latter may occur as a result of the former. Faster processing of queued payments means that intended receivers will obtain incoming funds more quickly, reducing the likelihood that their own subsequent outgoing payments will become queued upon submission. It is argued that the replacement of standard queuing arrangements with the introduction of central queuing that features a complex queue-release algorithm represents such an innovation. The benefit of these types of algorithms, in terms of both reduced settlement delay and intraday liquidity needs in an LVPS, are frequently highlighted throughout the payments literature. For example, see McAndrews and Trundle (2001); BIS (2005); Leinonen and Soramaki (1999); Bech and Soramaki (2001); Güntzer, Jungnickel, and Leclerc (1998); and Koponen and Soramaki (1998).

These algorithms are designed to simultaneously search for and offset batches of queued payments, thus serving as an effective coordination device for participants' incoming and outgoing payments. Recall that, under standard queuing procedures, payments are released from the queue *individually* when a participant's intraday liquidity is sufficient for them to be processed. In contrast, under central queuing with a complex queue-release algorithm, the simultaneous processing and release of a batch of queued payments is attempted at regular intraday intervals. In this latter case, LVPS participants no longer must wait to obtain sufficient intraday funds for their queued payments to be released individually, but rather they need only to hold the amount of intraday funds necessary to settle any net debit position resulting from the payment offset. The anticipated benefits to LVPS participants from this innovation include lower intraday liquidity needs and related costs for the release of queued payments, faster processing times for these queued payments, and a reduction in average intraday queue length, when compared with a standard queuing arrangement.

The addition of a complex queue-release algorithm will not necessarily represent a new development in all LVPS, since these algorithms have been used in some systems in the past as a gridlock resolution mechanism. However, over the past decade, increases in computing power have led to the improved design and more frequent use of these algorithms within an LVPS central queue. The complexity of these algorithms has also risen considerably; the choice of full or partial optimization is available and offsetting may take place on a bilateral and/or multilateral basis (BIS 2005).

To sum up, it is expected that the addition of a central queue with a complex queue-release algorithm will lead to an improvement in the trade-off between settlement delay and intraday liquidity in an LVPS, and will allow participants to complete the same level of payment activity at a lower overall cost, relative to a standard queuing arrangement.

4. Empirical Study: Estimating the Trade-Off in Canada's LVTS

This empirical exercise considers the trade-off between settlement delay and intraday liquidity in Canada's LVTS. Some questions that may arise are: What does the trade-off curve look like for the LVTS? Does it have the same shape as outlined above? Are there possible LVTS design changes, relating to queuing arrangements or otherwise, that could potentially improve this trade-off where the same level of payment activity can be processed with either reduced settlement delay or lower intraday liquidity needs, or both? The remainder of this paper is devoted to answering these questions using simulation analysis. Simulation analysis is a recent development in payment systems research. Simulation models are a valuable tool, since they often can be calibrated to replicate a specific LVPS environment. These models can then be used to assess the impact of changes in the structural arrangements and decision parameters of an LVPS without causing any costly disruption to the operation of the actual system.

4.1 Background on the LVTS[12]

The LVTS is an RTGS-equivalent system, where individual payment messages are processed on a gross basis in real time and settlement of the system occurs on a multilateral net basis at the end of the day. The LVTS's risk controls and collateral arrangements, coupled with a settlement guarantee provided by the Bank of Canada,

[12] Only LVTS background information relevant to the analysis is provided here. For more information on the LVTS, see Dingle (1998) and Arjani and McVanel (2006, forthcoming).

provide certainty of settlement for the system.[13] Certainty of settlement facilitates intraday finality for all individual payments sent through the LVTS. Recipients of LVTS payments can make use of these funds immediately upon receipt without any possibility that a payment will become unwound. The LVTS consists of two payment streams – Tranche 1 (T1) and Tranche 2 (T2) – and participants may use either stream when sending payments through the system. Each stream has its own real-time risk controls and collateral arrangements. The focus of this analysis is on the T2 payment stream since, due to its more economical collateral requirements relative to T1, it is the dominant stream for LVTS activity.[14]

Intraday liquidity in T2 is facilitated by T2 payments previously received and also by drawing on a T2 intraday line of credit. This intraday line of credit is subject to both a (indirect) collateral requirement and a net debit cap. Specifically, LVTS participants grant bilateral credit limits (BCLs) to each other, where the value of a BCL represents the maximum bilateral T2 net debit position that a grantee (credit line recipient) may incur vis-à-vis the grantor (credit line provider) at any time during the payment cycle. A participant's T2 intraday credit limit, known as its T2 net debit cap (T2NDC), is calculated as the sum of all BCLs granted to it by others in the system multiplied by a system-wide parameter (SWP), which is currently equal to 0.24.[15] The T2NDC represents the maximum multilateral T2 net debit position that a participant can incur during the LVTS payment cycle. The T2NDC of hypothetical bank n (where n = 1,..,N) is calculated as follows:

$$T2NDC^{n} = \sum_{j \neq n}^{N-1} BCL_{jn} \bullet SWP .$$

[13] In the extremely remote event of multiple participant defaults in the LVTS, and if collateral value pledged by participants to the Bank of Canada is not sufficient to cover the final net debit positions of all defaulters, the Bank stands ready to exercise its settlement guarantee by realizing on available collateral and absorbing any residual loss.

[14] Approximately 87 per cent of daily LVTS value and 98 per cent of daily LVTS volume are sent through the T2 payment stream, on average. T1 consists of mostly time-sensitive payments between LVTS participants and the Bank of Canada.

[15] The SWP is an exogenous parameter established by the CPA. When the LVTS began operations in February 1999, the SWP was equal to 0.30. Since then, it has been gradually reduced and has been equal to 0.24 since March 2000. The choice of SWP value (SWP < 1) reflects the effect of multilateral netting (Engert 1993). For information on the SWP, see LVTS Rule No. 2 on the Canadian Payments Association website, available at <http://www.cdnpay.ca>.

It follows that two real-time risk controls are applied to payments submitted to the T2 payment stream. A payment will be processed only if it does not result in the sending participant exceeding either its BCL vis-à-vis the receiver or its T2NDC.

A survivors-pay collateral pool is used in T2 to facilitate LVTS settlement in the event of participant default. Eligible collateral consists mainly of government securities, and also high-quality corporate debt. Participants are required to pledge T2 collateral equal to the value of the largest BCL that they grant to any other participant, multiplied by the SWP. The value of this T2 collateral obligation is referred to as a participant's maximum additional settlement obligation, or MaxASO. Essentially, a participant's MaxASO represents its maximum financial loss allocation as a result of another participant's default in the LVTS. Hypothetical bank *n*'s MaxASO is calculated as follows:

$$MaxASO^{n} = \max(BCL_{n,j \neq n}) \bullet SWP .$$

The LVTS employs a central queue. Submitted payments to the LVTS failing the real-time risk controls are stored in this queue.[16] The queue is equipped with an offsetting algorithm that runs at frequent intervals (every 15 minutes) throughout the payment cycle. This complex queue-release algorithm, called the Jumbo algorithm, searches for and offsets full or partial batches of queued payments on a multilateral and/or bilateral basis.[17] Payments successfully released by this mechanism are processed by the LVTS as normal. However, current LVTS rules state that excessive use of the central queue is not encouraged.[18] Instead, participants utilize internal queues to store payments that are unable to pass the real-time risk controls upon intended submission. Internally queued payments are typically resubmitted against the LVTS's risk controls (within a participant's internal LVTS workstation) *individually* on a bypass FIFO basis each time

[16] Payments are stored on a first-in first-out (FIFO) basis within each tranche type. Currently, only 'Jumbo' payments (>$100 million) failing the real-time risk controls become centrally queued in the LVTS.
[17] For queued T2 payments, the Jumbo algorithm applies partial offsetting on both a bilateral and multilateral basis over two stages. See Arjani and McVanel (2006, forthcoming) for more information on this algorithm.
[18] LVTS Rule No. 7 states that participants are able to track their bilateral and multilateral positions in real time through their internal LVTS workstations, and are expected not to submit payments that will fail the risk controls.

that its intraday liquidity position is increased.[19] If this process reveals that an internally queued payment can pass the risk controls, it is automatically released to the LVTS for processing.

4.2 Settlement delay and intraday liquidity in T2: Trade-off and improvement

Deciding on how to hypothetically impose a reduction in participants' intraday liquidity represents a key aspect of the analysis. For the LVTS T2 payment stream, one way to accomplish this is to constrain the intraday credit available to participants by lowering the value of the SWP.[20] As in the earlier discussion, a reduction of the SWP will entail both a benefit and a cost for LVTS participants, holding BCL values constant. The benefit is that a reduction in the value of the SWP will lower participants' T2 collateral requirement and related liquidity cost. However, assuming that no migration of payments from T2 to T1 occurs, reducing the SWP will likely also increase the level of settlement delay in the T2 payment stream. This is because participants' T2NDCs will decline, lowering T2 intraday liquidity in the system, and causing more payments to become queued upon their intended submission. Under current queuing arrangements, delayed payments will accumulate in participants' internal queues until the sending participants' T2 liquidity is sufficient for these payments to be processed by the LVTS.

The trade-off curve between settlement delay and intraday liquidity in the LVTS is expected to have a decreasing convex shape, as outlined in the earlier graphical framework. As the SWP is reduced further, overall settlement delay in the system is expected to rise at an increasing rate. Participants will become constrained by their T2NDC more quickly and frequently throughout the day when trying to send payments. In the extreme case, an SWP equal to zero will result in a state of payments deadlock where settlement delay reaches a maximum. No participant will have access to T2 intraday credit and therefore will not be able to incur a T2 net debit position.

[19] Under bypass-FIFO, a participant's first (earliest) queued payment will be retried against the risk controls. If it does not pass, this payment will be bypassed and the participant's second queued payment will be retried, and so on.
[20] Alternatively, such reductions in intraday credit availability can also be achieved through reductions in the value of BCLs that participants grant to each other, while maintaining the current SWP value of 0.24.

Consequently, no payments will be sent and all will remain unsettled in participants' internal queues until the end of the day.

It has been argued that an improvement in the trade-off between settlement delay and intraday liquidity can be achieved with the introduction of a complex queue-release algorithm in the central queue. The LVTS already contains a central queue with a partial offsetting algorithm, but use of this queue is currently discouraged. It is anticipated that, by allowing increased use of the LVTS central queue (and this algorithm), overall settlement delay could be reduced for each hypothetical level of T2 intraday credit provision. Under this alternative scenario, participants would no longer need to manage an internal payments queue, and instead would submit all payments to the LVTS at the time they are intended regardless of whether these payments could be immediately processed by the system. Release of these queued payments could then be attempted on a multilateral net basis rather than individually.[21] This proposed change in queuing regime is expected to increase the efficiency of the system since, even where the amount of T2 intraday credit available to participants (and related cost) is lowered, the processing time for queued payments can be faster, and average queue length could decrease, compared with current internal queuing arrangements.

In the next sections, a simulation approach will be utilized to shed light on the following questions:

- Under current internal queuing arrangements, what does the trade-off between settlement delay and intraday liquidity in the LVTS look like? Is it consistent with the assumptions of the graphical framework presented above?
- Could increased use of the LVTS central queue improve this trade-off? In other words, can the level of settlement delay associated with each amount of intraday credit be reduced for a given level of payment activity?

[21] The key benefit of central queuing compared with internal queuing is that multilateral offsetting of payments is possible only in the former case.

5. Data Description and Simulation Methodology

5.1 Description of data

Three months of LVTS T2 transaction and credit limit data have been extracted over the period July–September 2004. Transaction data include the date and time that each transaction was submitted to the LVTS, as well as the value of each payment and the counterparties involved in the transaction. It is assumed that the time stamp attached to each payment represents the intended submission time of the payment. Transactions data include only those payments processed by the LVTS, and do not include rejected or unsettled payments. Data on credit limits include the value of the T2NDC available to each participant, as well as the date and time that the value of the T2NDC is effective. These data represent 64 business days and approximately 1.05 million transactions, and are believed to be representative of normal LVTS activity. Table 1 provides a summary of the transaction data.[22]

Table 1: Summary of LVTS T2 Transaction Data

	Jul 2004	Aug 2004	Sep 2004
Total Value of T2 Payments (Can$ billion) (% of LVTS Total)	2,283.0 (87.8)	2,203.5 (87.9)	2,446.5 (86.3)
Total Volume of T2 Payments (% of LVTS Total)	349,948 (98.0)	344,357 (98.0)	356,676 (98.1)
Daily Average Value (Can$ billion)	108.7	100.2[23]	116.5
Daily Average Volume	16,664	15,653	16,985
Average Payment Value (Can$ million)	6.52	6.40	6.86
Median Payment Value (Can$)	42,436	40,377	45,719

[22] In addition, the Hirschman-Herfindahl index (HHI) suggests that payment activity over the sample period is somewhat concentrated. The HHI will vary between 0.50 (concentration among only two banks) and 1/N (equal distribution of payment activity among all participants), where N represents the number of banks in the sample. In this case, 1/N = 0.08. The average HHI value for the sample is 0.1944 and 0.1813 for T2 payments value and volume, respectively. A value in this range is consistent with payment activity being distributed evenly across approximately five to six banks. Indeed, the largest five Canadian banks account for between 85–90 per cent of daily LVTS value and volume.

[23] A lower average daily T2 payments value in August is expected given that the Canadian civic holiday occurs during this month. Total value reached only $6.9 billion on this holiday in 2004.

5.2 Simulation description and methodology

The simulation analysis is conducted using a payment and settlement simulator developed by the Bank of Finland (the BoF-PSS2). This software application is currently being used by over thirty central banks. It should be noted that the current version of the BoF-PSS2 does not contain BCL functionality, which is an important component of the LVTS.[24] As a result, the methodology in this paper includes the assumption that BCL values remain constant in light of proposed changes to LVTS rules on queue usage. Further, participants' payment-sending behaviour is also treated as exogenous and therefore the same transactions data are used throughout the analysis. Potential implications associated with these assumptions are addressed later in the paper.

Two batches of simulations will be run where each batch is intended to replicate a different LVPS design. In particular, batch one replicates the current internal queuing arrangement in the LVTS, while batch two replicates the alternative central queuing arrangement. Each batch consists of eight individual simulations (s = 1, 2,.., 8), where each simulation is distinguished by tighter constraints on participants' intraday liquidity. Changes in intraday liquidity are introduced by altering the value of each participant's T2NDC. Since it is assumed that BCLs remain constant, a reduction in each participant's T2NDC is achieved by hypothetically lowering the value of the SWP. Specifically, each individual participant *n*'s T2NDC in simulation *s* is calculated as follows:

$$T2NDC_s^n = SWP_s \bullet \sum_{j \neq n}^{N-1} BCL_{jn},$$

where $SWP_{1,\ldots,8}$ = 0.24, 0.21, 0.18, 0.15, 0.12, 0.09, 0.06, 0.03.[25]

[24] A revised version of the BoF-PSS2 is expected to be released in 2006 that will include both multilateral and bilateral credit limits functionality. Bank of Canada staff are involved in the development and testing of this new version.

[25] Transactions data include only processed payments under the current SWP value of 0.24. Thus, it is not possible to observe potential reductions in settlement delay from an SWP value greater than 0.24, due to a lack of readily available data on delayed or unsettled transactions for this SWP value.

In specifying the first batch of simulations, the objective is to mimic the participants' decision to either submit a payment to the LVTS for processing or hold the payment internally when sufficient intraday funds are unavailable. Settlement delay occurring in this batch represents payments being held internally by participants; i.e., the simulator's queue is replicating participants' internal queues. A bypass-FIFO queue-release algorithm is specified to imitate current internal queuing practices of LVTS participants. When this algorithm is applied, a participant's queued payments are resubmitted from the queue and retried against the risk controls on an individual bypass-FIFO basis whenever its intraday liquidity position improves. In the real LVTS, this occurs within the participant's internal workstation. Internally queued payments that can successfully pass the risk controls are assumed to be released from the participant's queue and submitted to the LVTS for processing. In interpreting the simulation results for this first batch, settled transactions are assumed to be those that participants were able to submit to the LVTS for processing, while unsettled transactions represent those remaining in participants' internal queues due to lack of intraday liquidity.

Specification of the second batch is intended to replicate a central queuing regime similar to that available in the LVTS. In these simulations, two queue-release algorithms are specified that closely match the LVTS's actual release mechanisms. The first of these algorithms is a FIFO (no bypass) queue-release algorithm, which resubmits a participant's centrally queued payments against the risk controls on an individual FIFO basis each time its intraday liquidity position improves. The second is a complex queue-release algorithm, which employs partial offsetting on a multilateral basis and is scheduled to run every twenty minutes, similar to the LVTS's Jumbo algorithm.[26] Settlement delay captured in this second batch of simulations is meant to represent payments being held in the system's central queue; i.e., the simulator's queue is

[26] At the time that the analysis was conducted, the frequency of the Jumbo algorithm was every 20 minutes. The frequency of this algorithm increased to every 15 minutes in December 2005. Since bilateral credit limit functionality is currently not incorporated in the simulation application, the partial offsetting algorithm used in the simulations does not exactly replicate the LVTS Jumbo algorithm for T2 payments. Despite this limitation, the results generated by the simulations are still expected to be useful and relevant. Further, in specifying this second batch of simulations, it is also assumed that the LVTS's queue expiry algorithm is no longer utilized and *all* payments failing the risk control check become centrally queued (not just 'Jumbo' payments).

replicating the LVTS central queue. In the simulation results for this batch, all payments in the sample are assumed to have been submitted to the LVTS at their intended time of submission, and unsettled transactions are those remaining in the central queue that cannot be processed due to a sender's lack of intraday liquidity.

Three alternative measures of settlement delay are calculated for each simulation within each batch. These measures are intended to capture the daily level of settlement delay associated with each amount of intraday credit provision under both the current and alternative queuing environments described above, for the same level of payment activity. They are described as follows:

1. Daily Proportion of Unsettled Transaction Value (PU):

$$PU_t^N = \left(\frac{Value\ of\ Unsettled\ Transactions_t^N}{Value\ of\ Submitted\ Transactions_t^N} \right)$$

This indicator is calculated on an aggregate level (i.e., across all participants) for each day t in the sample, where t = (1,…, 64). This measure represents the occurrence of the maximum settlement delay possible for a payment in this analysis. Unsettled transactions represent those that enter the queue upon intended submission and remain there until the end of the day.

2. Daily System-Wide Delay Indicator (DI):

$$DI_t^N = \left(\sum_{n=1}^{N} w^n r^n \right)$$

where $r^n = \left(\frac{\sum_{i=1}^{T} Q_i^n}{\sum_{i=1}^{T} V_i^n} \right)$ and $0 \leq w^n, r^n, DI^N \leq 1$

Adapted from Leinonen and Soramaki (1999) and commonly used in payment simulation analyses, this indicator is calculated on an aggregate level and is based on a weighted average of each individual (n) participant's daily delay indicator (r). This indicator (and the ratio r) can take on any value between 0 and 1, where a value of 0 is achieved when all payments are successfully processed by the LVPS upon intended submission and no settlement delay occurs. A value of 1 is calculated where all payments become queued upon intended submission and remain unsettled at the end of the day. Weights (w) are based on participants' average share of total transaction value over the 64-day sample period. Calculation of this measure requires dividing each LVTS business day into T=108 ten-minute intervals (i = 1,.., T). The numerator of r represents the sum of a participant's queued payment value (Q) over all T ten-minute intervals throughout the day. The denominator represents the sum of the cumulative value of a participant's submitted payments (V) over all T ten-minute intervals throughout the day. It follows that this indicator is influenced by both the value and delay duration of each payment in the queue calculated for each intraday interval.

3. Average Intraday (Interval) Queue Value (AQV):

$$AQV_t^N = \left(\frac{\sum_{i=1}^{T} Q_i^N}{T} \right)$$

This is an aggregate measure that calculates the average value of queued payments in an interval over day t. It is found by dividing the sum of total queued payment value (Q) over all T ten-minute intervals on each day by the number of intervals per day (T=108).

6. Simulation Results and Discussion

6.1 The delay-liquidity trade-off in the T2 payment stream

Simulation results for each of the three delay measures are presented in Figures 3 through 5. Two curves are presented in each graph, corresponding to each batch of simulations. The curve denoted 'Internal Queuing' portrays the simulation results estimated under current LVTS (internal) queuing arrangements. The curve denoted 'Central Queuing' depicts results estimated under the alternative LVTS (central) queuing environment.

Earlier hypotheses regarding the trade-off between settlement delay and intraday liquidity are confirmed by the simulation results. Under current LVTS queuing arrangements, a trade-off exists in the LVTS's T2 payment stream according to all three delay measures. Like the earlier graphical framework, the curve is convex; as intraday credit constraints are further tightened (by lowering the value of the SWP), participants' intraday liquidity becomes more scarce and settlement delay in the system rises at an increasing rate. The slope of this curve increases substantially at low amounts of intraday credit provision.

The introduction of a design innovation – allowing increased use of the LVTS central queue – results in an improvement to this trade-off and the curve shifts closer towards the origin, according to all three measures. Settlement delay associated with each level of intraday credit provision is reduced following the introduction of the partial offsetting algorithm. The relative benefit of partial offsetting (in terms of reduced delay) increases gradually as intraday liquidity is further constrained. At the SWP value of 0.06, the difference in settlement delay between the two queuing regimes is greatest. In this case, the average proportion of unsettled transactions value is reduced by 9 percentage points or about $10 billion (Figure 3), the system-wide delay indicator is reduced by 28 per cent (Figure 4), and the average intraday queue value is reduced by 29 per cent or about $1.6 billion (Figure 5), relative to the first batch of simulations.

Figure 3: Average Daily Proportion of Unsettled Transaction Value

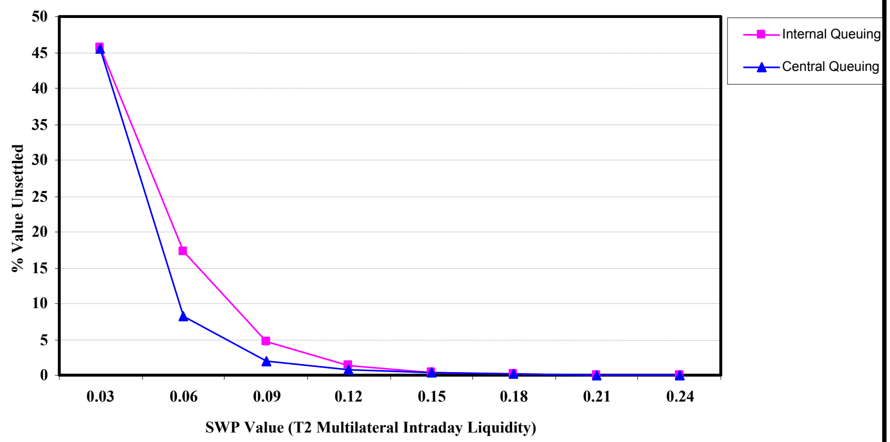

Figure 4: Average Daily System-Wide Payments Delay

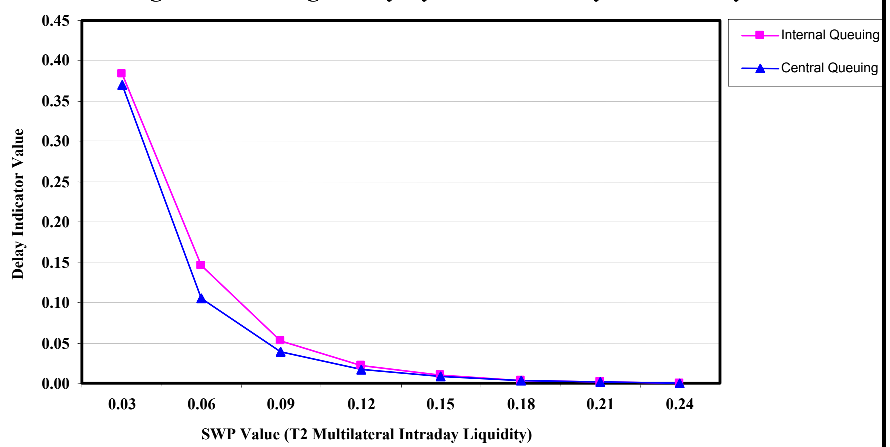

Figure 5: Average Intraday (Interval) Queue Value

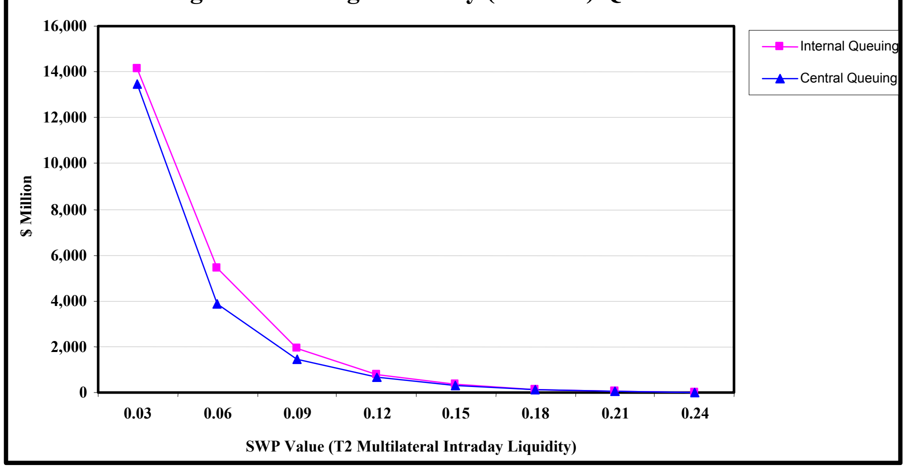

Gains from the alternative central queuing design begin to decline when the SWP is reduced beyond 0.06, as the system begins to approach a state of deadlock. When the SWP value is 0.03, settlement delay is only slightly reduced following the introduction of a partial offsetting algorithm, which could mean that participants' intraday liquidity levels are so low that only very small batches of queued payments can be processed each time this algorithm runs. At this level of SWP, close to half of all daily payment value remains unsettled on average under both queuing regimes (Figure 3).

The simulation results also reveal another finding that is closely related to the notion of technical efficiency described earlier. The above results suggest that, under current queuing arrangements, settlement delay in T2 increases when the SWP value is lowered from 0.24 to 0.21. However, it remains to be seen whether reductions in the SWP below 0.24 but still greater than 0.21 can be achieved without inducing any further settlement delay in the LVTS. In other words, can a lower amount of T2 intraday credit (and an associated reduction in T2 collateral requirements) be accommodated without increasing the level of settlement delay for payment activity during the three-month sample period, holding all other factors constant? If this were the case, it would be similar to operating at point C in the graphical framework. Indeed, the simulation results suggest that the current value of SWP (= 0.24) is needed to process payments in this sample, and cannot be reduced further without increasing the level of settlement delay. This is not necessarily a surprising result, since one might expect participants to conform to this value of SWP when sending payments through the system. A complete discussion of this analysis, including full details of the simulation methodology used, is provided in the appendix.

6.2 Discussion

Some interesting discussion points emerge from these results, suggesting areas for future research. First, the simulation results suggest that, under both existing LVTS queuing arrangements and also under the alternative central queuing arrangement, settlement delay in T2 will increase only marginally as the SWP is initially reduced from its current value of 0.24, holding all other factors constant. For example, a reduction in the SWP from 0.24 to 0.18 is estimated to increase the average proportion of unsettled daily transaction value by only 0.15 per cent under the current queuing regime and 0.14 per cent under a central queuing arrangement (Figure 3). Similar results are also observed according to the other two delay measures. Reducing the SWP entails a benefit for LVTS participants in the form of lower T2 collateral requirements and related liquidity cost, as has already been mentioned. Specifically, a reduction in the SWP to 0.18 reduces the aggregate value of T2 collateral required by about $750 million per day on average over the sample period, holding BCL values constant. On one particular day in the sample, the value of T2 collateral required is about $1 billion less when the SWP is equal to 0.18.

This raises the question as to whether a lower-cost combination of intraday credit provision and settlement delay currently exists for LVTS participants in the T2 payment stream.[27] Put differently, is it the case that the marginal settlement delay cost incurred by moving to an SWP value of 0.18 equals the marginal cost of additional intraday credit provision (and collateral) associated with the current value of 0.24? If the former cost is less than the latter, then lowering the SWP to 0.18 could lead to overall cost savings for participants. Of course, answering this question entails, among other things, the difficult task of quantifying the cost of the additional settlement delay associated with moving to a SWP value of 0.18.

Secondly, the analysis highlights the possible benefit of central queuing with a complex queue-release algorithm with respect to settlement delay and intraday credit provision. Nonetheless, participants face other types of risk and cost in the LVPS environment, and such a change in LVTS queuing arrangements could increase participants' other costs. For example, as outlined in BIS (1997), a possible implication of permitting unrestricted use of the central queue pertains to the issue of queue transparency, and specifically whether the reduction in settlement delay could be replaced by an increase in credit risk taken on by participants. A participant, upon observing an incoming payment in the central queue, may choose to provisionally credit its client's account with these expected funds before the payment actually arrives, thus exposing itself to credit risk until the

[27] Alternatively, the question could instead be posed as to whether current values of BCLs granted by participants to each other are cost-minimizing holding the current SWP value constant.

payment is successfully received. If these funds do not eventually arrive for some reason, the participant would seek to unwind this payment, which would be costly for both the participant and its client. This issue is pertinent to the LVTS because participants have the ability to track expected incoming and outgoing payments in the queue in real time through their internal participant workstations. Although details regarding client recipients of incoming queued payments are not included in these workstation reports, participants could informally access this information. However, it is not clear that LVTS participants would be willing to incur this credit risk in any case.[28]

7. Conclusions and Caveats

The objective of this paper has been to gain a better understanding of the trade-off between settlement delay and intraday liquidity in an LVPS, with a specific focus on the Canadian LVTS. Emphasis was placed on the provision of intraday credit as a source of participants' intraday liquidity in the system. Settlement delay and intraday liquidity entail costs for LVPS participants. A general graphical framework inspired by Berger, Hancock, and Marquardt (1996) is useful for understanding the nature of this trade-off in an LVPS. This framework illustrates how an innovation in LVPS queuing arrangements can improve this trade-off by lowering the combination of settlement delay and intraday liquidity necessary to produce a given level of payment activity. Such an innovation can improve the efficiency of the system, leading to overall cost savings for participants.

Simulation analysis represents a recent development in payment systems research and can be used to consider what this trade-off might look like in the case of Canada's LVTS. The empirical results indicate that a trade-off exists in the LVTS between settlement delay and intraday liquidity, and the decreasing convex shape of this trade-off is consistent with the earlier graphical framework. The simulation results also confirm findings from the payments literature regarding improvements in the trade-off.

[28] This credit-risk issue may also be avoided in the LVTS, since a client beneficiary of funds can always request a payment confirmation reference number (PCRN) from its participant bank. All payments processed by the LVTS are assigned a PCRN indicating that the payment has successfully passed all LVTS risk-control tests and is thus considered final and irrevocable. Upon obtaining the PCRN, the beneficiary does not have to worry about the funds being revoked at a later time.

Specifically, allowing increased use of the LVTS central queue (and the Jumbo algorithm) is expected to reduce settlement delay in the system for all levels of intraday liquidity considered, relative to current internal queuing arrangements.

Some important considerations emerge from these results. First, under both the current and proposed queuing regimes, a modest reduction in the SWP below its current value results in only a marginal increase in the level of settlement delay in the LVTS, while potentially providing substantial T2 collateral cost savings for system participants. Further research is necessary to quantify whether this collateral cost-saving benefit is worth the associated increase in settlement delay cost. It was also argued that, although increased use of the central queue is expected to reduce total settlement delay and liquidity costs for participants, this may result in a potential increase in credit risk taken on by participants. However, LVTS participants may not necessarily react to a change in LVTS queuing arrangements in this manner.

The simulation results outlined in this paper are believed to be useful and relevant. However, it must be emphasized that these results are preliminary, and the existence of certain caveats means that there is more work to be done. These caveats are discussed here with the intention of motivating further research. The first caveat relates to behavioural assumptions made throughout the analysis. Significant changes to LVTS queuing arrangements were proposed in the analysis. However, despite these changes, the current simulation methodology assumes that LVTS participants' payment sending and bilateral credit granting behaviour remains unchanged. One must question whether this is a realistic assumption. For example, following discussion in McAndrews and Trundle (2001), the availability of netting is likely to increase the incentive for participants to submit payments to the system earlier in the day, relative to these payments' current intended submission times, essentially increasing the scope for multilateral netting of payment messages. The benefit of netting is expected to increase with the number and value of payments in the queue at the time that it occurs. Anecdotal evidence suggests that LVTS participants typically receive information regarding outgoing payment requests well in advance of their intended submission time. Participants' collective submission of as many payments as early as possible to the system under a central queuing regime is anticipated to result in a greater turnover of intraday funds, a lesser need for costly intraday credit, and faster processing of these payments. This may result in a further downward shift of the trade-off curve closer to the origin, thus leading to further cost savings for participants.

At the same time, it is argued that participants, in granting BCLs to each other, strive to minimize the value of their T2 collateral requirement subject to achieving an established level of throughput efficiency; i.e., an acceptable level of settlement delay. It is likely that payment activity under current internal queuing arrangements may already reflect participants' acceptable levels of settlement delay. Thus, participants may not perceive the benefit of central queuing to be a further reduction in settlement delay, but instead may treat this as an opportunity to realize lower T2 collateral requirements (and costs) while maintaining the same level of settlement delay in the system. This suggests that, under the central queuing arrangement, participants may collectively choose to reduce the BCLs they grant to each other in order to achieve these cost savings. This reduction in BCLs is expected to continue to the extent that any decline in settlement delay resulting from increased use of the central queue is fully offset.[29]

To rigorously examine these propositions, more sophisticated tools are needed to study formal hypotheses regarding potential behavioural changes in response to alternative LVPS designs. The development of new theoretical and empirical models that capture the primary factors underlying participants' payment submission and bilateral credit granting decisions, as well as forthcoming developments in the BoF-PSS2 application, will help to fill this void.

A second caveat that follows closely a discussion found in Bedford, Millard, and Yang (2005) relates to the statistical robustness of the simulation findings. The simulation

---

[29] Initially, participants are not likely to know exactly how much BCLs must be reduced to achieve the same level of settlement delay under the alternative central queuing regime. Instead, this will be an iterative process that eventually converges to the equilibrium of a perfect offset. In the interim, it may be the case that participants 'overshoot' this target level of BCL reduction, temporarily resulting in a higher level of settlement delay in the system relative to the existing level.

analysis is intended to estimate the increase in settlement delay brought on by a reduction in LVTS participants' intraday liquidity over a three-month sample period. Point estimates of this impact for each amount of intraday liquidity are used to generate the trade-off curves presented in Figures 3 through 5. Previous internal research conducted by the Bank of Canada shows that annual LVTS payment activity is affected by specific calendar events and also monthly trends. Consequently, the estimated impact on settlement delay following reductions in intraday liquidity is expected to take on different values based on the specific dataset used in the analysis. Although using a three-month sample helps to capture the effect of certain monthly and quarterly calendar effects occurring during this period, there is a desire to reduce the risk of small-sample bias and to obtain more statistically robust results. For example, it has been observed that the same calendar event may yield a different effect on LVTS payment activity, depending on when it occurs throughout the year. Similarly, use of a single three-month sample may not capture the effect that semi-annual and/or annual calendar events may have on the simulation results. Nor will it capture the potential impact of monthly trends in LVTS T2 payment activity.

In order to achieve more statistically robust results, it is suggested that the same simulation methodology be repeated as many times as is feasible using real and/or artificially generated LVTS payment flow data over some fixed sample duration. Grouping the point estimates of the impact on settlement delay for each amount of intraday liquidity from all of the samples will facilitate generation of an empirical distribution of this potential impact. Figure 6 provides a hypothetical illustration of this result. It follows that the shape of the empirical distribution may be different for each amount of intraday liquidity. For example, the impact on settlement delay may be more volatile and will thus deviate from its mean value more often at lower amounts of intraday credit provision. The shape of the empirical distribution may also change over time.

Figure 6: Plotting Distribution of Settlement Delay Outcomes

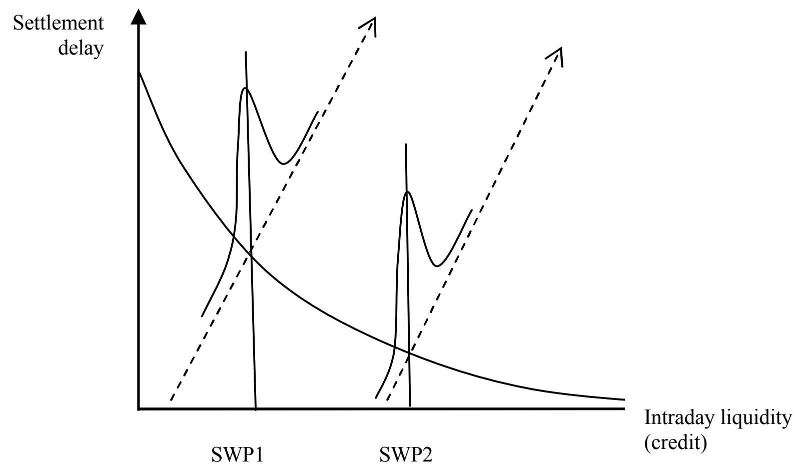

A third and final caveat pertains to the absence of BCL functionality in the current version of the BoF-PSS2. This absence creates the possibility that the estimated trade-off curves provided in Figures 3 through 5 represent a 'lower bound' of the impact on settlement delay resulting from reduced intraday liquidity. As the value of the SWP is reduced and payments become delayed upon failing the T2 multilateral risk-control test, intended receivers of these payments may consequently be prohibited from sending their own payments when due. All of this will result in added volatility in bilateral net positions, possibly to a point where some participants' bilateral net debit positions are greater than the BCLs granted to them. In the LVTS, this cannot occur due to a bilateral risk control test being applied to every payment, which guarantees that participants do not exceed their BCL vis-à-vis a receiving participant. Payments failing the bilateral risk-control test become queued until the sending participants' bilateral liquidity position improves. This added delay is not captured in the results generated by the current version of the simulator. This forces the assumption that all LVTS payments, when processed by the simulator, have passed not only the multilateral risk-control test, but also the bilateral risk-control test. It will be interesting to repeat the analysis again with Version 2.0 of BoF-PSS2, to compare how much greater is potential settlement delay in the system when bilateral risk controls are also taken into account.

References

Arjani, N. and D. McVanel. 2006. "A Primer on Canada's Large-Value Transfer System." Bank of Canada. Forthcoming.

Bank for International Settlements (BIS). 1997. "Real-Time Gross Settlement Systems." CPSS Publications No. 22.

———. 2005. "New Developments in Large-Value Payment Systems." CPSS Publications No. 67.

Bech, M. and K. Soramaki. 2001. "Gridlock Resolution in Interbank Payment Systems." Bank of Finland Discussion Paper. No. 9/2001.

Bedford, P., S. Millard, and J. Yang. 2005. "Analysing the Impact of Operational Incidents in Large-Value Payment Systems: A Simulation Approach." In *Liquidity, Risks and Speed in Payment and Settlement Systems – A Simulation Approach*, edited by H. Leinonen, 247–74. Bank of Finland Studies.

Berger, A.N., D. Hancock, and J. Marquardt. 1996. "A Framework for Analyzing Efficiency, Risks, Costs, and Innovations in the Payment System." *Journal of Money, Credit and Banking* 28(4): 696–732.

Dingle, J. 1998. "The LVTS – Canada's Large-Value Transfer System." *Bank of Canada Review* (Autumn): 39–55.

Engert, W. 1993. "Certainty of Settlement and Loss Allocation with a Minimum of Collateral." Bank of Canada Working Paper No. 93–14.

Güntzer, M.M., D. Jungnickel, and M. Leclerc. 1998. "Efficient Algorithms for the Clearing of Interbank Payments." *European Journal of Operational Research* 106: 212–19.

Koponen, R. and K. Soramaki. 1998. "Intraday Liquidity Needs in a Modern Interbank Payment System: A Simulation Approach." Bank of Finland Studies No. E:14/1998.

Leinonen, H. and K. Soramaki. 1999. "Optimizing Liquidity Usage and Settlement Speed in Payment Systems." Bank of Finland Discussion Paper No. 16/99.

———. 2003. "Simulating Interbank Payment and Securities Settlement Mechanisms with the BoF-PSS2 Simulator." Bank of Finland Discussion Paper No. 23/2003.

Martin, A. 2005. “Recent Evolution of Large-Value Payment Systems: Balancing Liquidity and Risk.” Federal Reserve Bank of Kansas City *Economic Review* (First Quarter): 33–57.

McAndrews, J.J. and S.M. Potter. 2002. “Liquidity Effects of the Events of September 11, 2001.” *Economic Policy Review*, Federal Reserve Bank of New York (November): 59–79.

McAndrews, J.J. and S. Rajan. 2000. “The Timing and Funding of Fedwire Funds Transfers.” *Economic Policy Review*, Federal Reserve Bank of New York (July): 17–32.

McAndrews, J.J. and J. Trundle. 2001. “New Payment System Designs: Causes and Consequences.” Bank of England *Financial Stability Review* (December): 127–36.

Appendix: Is the T2 Payment Stream Technically Efficient?

The objective of this supplemental analysis is to find the minimum SWP (call this SWP*) necessary to process all payments in the sample without delay, holding all other factors constant. It may be the case that SWP* < 0.24, which means that for existing levels of T2 intraday credit, and perhaps more importantly for participants, T2 collateral requirements could be lowered without inducing additional settlement delay during the three-month sample period (recall point C in Figure 2).

Simulation results produced by the BoF-PSS2 can provide insight into this issue. Treating participants' payment-sending behaviour as exogenous, a simulation is run using the same sample data but this time specifying unlimited intraday credit. Under this simulation scenario, all payments will pass the risk controls immediately upon submission and therefore no queuing algorithms need to be specified. The daily T2NDC that each participant actually needs in order for its payments to be passed without delay can be derived from these simulation results, and is equal to the largest multilateral net debit (negative) position incurred by each participant during the day. This value is defined as a participant's upper bound (UB) of T2 liquidity. The daily UB of T2 liquidity for each participant can then be used to calculate a value of SWP* that, when multiplied by the sum of the actual BCLs granted to each participant, will produce this UB value. It follows that the highest value of SWP* calculated for any participant on any day is considered the minimum SWP* value necessary to send all payments in the sample through the system without delay. This SWP* can then be compared with the current value of 0.24.

The results from this simulation analysis reveal that on 45 of the 64 days, SWP* reached 0.24 for at least one LVTS participant. This means that the current value of SWP was necessary for the immediate processing of T2 payment activity during this three-month sample period. Hence, further T2 collateral cost savings could not be realized without an increase in the level of settlement delay, holding payment activity constant. The results also indicate that the T2NDC constraint (when SWP=0.24) is binding more often for large LVTS participants (denoted ‘B5’ in Figure 7). Figure 7 shows that on 42 days in the sample, at least one of the major Canadian banks reached their T2NDC at some point in the day.

Figure 7 : Minimum SWP Required – B5 vs. S8

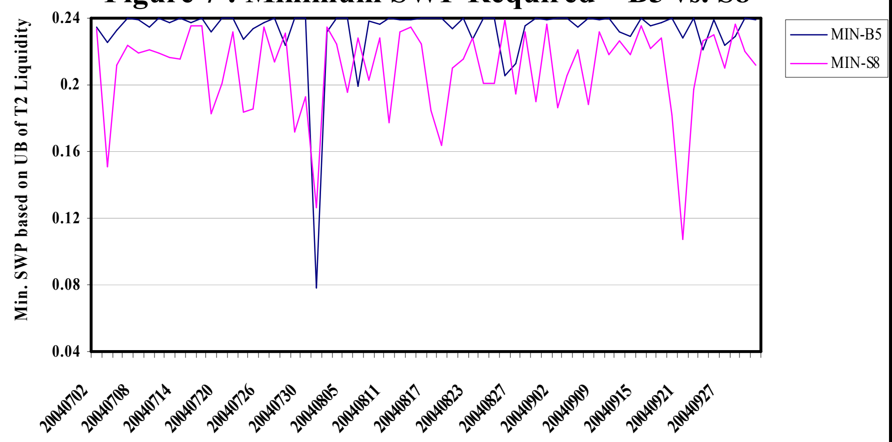

Focusing on the large LVTS participants, the simulation results show that, on these 42 days, four different institutions bumped up against their T2NDC at least once intraday. One of these participants reached its T2NDC at least once on 37 different days, while the three others reached this limit on 10, 2, and 1 day(s), respectively. The results also indicate that participants did not reach their T2NDC constraint at the same time each day. For example, regarding the first two large participants mentioned above, the LVTS day has been divided into four periods and the time that each of these participants reached its T2NDC has been located in the simulation results and tabulated. A summary of these findings is provided in Table 2.

Table 2: Percentage of Instances where T2NDC is Binding by Time of Day

Time of Day	Bank 1 (37 instances)	Bank 2 (10 instances)
00:30-06:00	0	0
06:00-12:00	19	0
12:00-17:00	73	40
17:00-18:30	8	60

It also deserves mention that, where a high number of instances occur within a certain period (e.g., 27 instances for Bank 1 during the interval between 12:00 and 17:00 hours), these occurrences typically do not take place at the same time within the interval, but rather were scattered throughout the period.

It is not necessarily surprising that SWP* reaches 0.24 on most days in the sample period. The gradual reduction of the SWP from 0.30 to 0.24 between February 1999 and March 2000 was influenced by participants' preferences, and this value has held steady at 0.24 since that time. Given participants' perceived contentment with this SWP value, one might expect participants to conform to it, meaning that they choose to structure their payment submission behaviour in a certain way so as to make full use of their available T2 intraday credit when sending payments through the system.

Some discussion is also warranted regarding results for the eight smaller LVTS participants (denoted 'S8' in Figure 7). On only four of the 45 days, SWP* reached 0.24 for one of these participants. Further, this occurred for a different participant in each of these four instances. There exist a variety of possible explanations for these results. It may be the case that larger LVTS participants, in sending a higher volume of payments earlier in the day, are 'subsidizing' smaller participants' intraday liquidity in the system, to the extent that smaller participants need to rely less on intraday credit as a source of funding for their outgoing payments. Indeed, SWP* was equal to zero (i.e., no T2 intraday credit was drawn upon) for at least one small participant on 18 of 45 days in the sample. In contrast, this did not occur on any day for large LVTS participants. A second possible explanation could be that, for various reasons, small LVTS participants may tend to bump up against their BCLs far more frequently relative to their T2NDC. Of course, further research is necessary before either of these explanations can be confirmed.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages, and are available on the Bank's website (see bottom of page). ***Les documents de travail sont généralement publiés dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue. On peut les consulter dans le site Web de la Banque du Canada, dont l'adresse est indiquée au bas de la page.***

2006

2006-19	Institutional Quality, Trade, and the Changing Distribution of World Income	B. Desroches and M. Francis
2006-18	Working Time over the 20th Century	A. Ueberfeldt
2006-17	Risk-Cost Frontier and Collateral Valuation in Securities Settlement Systems for Extreme Market Events	A. García and R. Gençay
2006-16	Benchmark Index of Risk Appetite	M. Misina
2006-15	LVTS, the Overnight Market, and Monetary Policy	N. Kamhi
2006-14	Forecasting Commodity Prices: GARCH, Jumps, and Mean Reversion	J.-T. Bernard, L. Khalaf, M. Kichian, and S. McMahon
2006-13	Guarding Against Large Policy Errors under Model Uncertainty	G. Cateau
2006-12	The Welfare Implications of Inflation versus Price-Level Targeting in a Two-Sector, Small Open Economy	E. Ortega and N. Rebei
2006-11	The Federal Reserve's Dual Mandate: A Time-Varying Monetary Policy Priority Index for the United States	R. Lalonde and N. Parent
2006-10	An Evaluation of Core Inflation Measures	J. Armour
2006-9	Monetary Policy in an Estimated DSGE Model with a Financial Accelerator	I. Christensen and A. Dib
2006-8	A Structural Error-Correction Model of Best Prices and Depths in the Foreign Exchange Limit Order Market	I. Lo and S.G. Sapp
2006-7	Ownership Concentration and Competition in Banking Markets	A. Lai and R. Solomon
2006-6	Regime Shifts in the Indicator Properties of Narrow Money in Canada	T. Chan, R. Djoudad, and J. Loi
2006-5	Are Currency Crises Low-State Equilibria? An Empirical, Three-Interest-Rate Model	C.M. Cornell and R.H. Solomon
2006-4	Forecasting Canadian Time Series with the New Keynesian Model	A. Dib, M. Gammoudi, and K. Moran

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
Telephone: 1 877 782-8248
(toll free in North America)
Email: publications@bankofcanada.ca
Website: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Téléphone : 1 877 782-8248 (sans frais en Amérique du Nord)
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca