PAYMENTS CANADA

A COMPUTATIONAL MODEL OF THE MARKET MICROSTRUCTURE OF BILATERAL CREDIT LIMITS IN PAYMENT SYSTEMS AND OTHER FINANCIAL MARKET INFRASTRUCTURES

2018-08-28

payments.ca

A Computational Model of the Market Microstructure of Bilateral Credit Limits in Payment Systems and other Financial Market Infrastructures

Oluwa**segun** Bewaji,*

*Oluwasegun Bewaji is currently with Payments Canada, Ottawa, Canada, e-mail: sbewaji@payments.ca.

Acknowledgements

The author gratefully acknowledges comments from colleagues at Payments Canada, Neville Arjani, Philippe-Olivier Blanchet, Anne Butler, Brendan Carley, Bill Cyr, Walter Engert, Viktoria Galociova, Michelle LeClair, Duangsuda Sopchokchai, Youcef Touizrar, and Ronald Heijmans at the Netherlands Central Bank. The author also gratefully acknowledges comments from James Chapman and other reviewers at the Bank of Canada, Professor, Howard M. Schwartz, PhD at Carleton University, as well as editorial assistance from Candice Baltare and Victoria McMullen

Disclaimer
The views presented in this paper are those of the author and do not necessarily reflect the views of Payments Canada

Abstract

This paper provides the first steps towards a theoretical framework through which optimisation decisions in payments and settlement systems can be assessed from a market microstructure perspective. In particular, the paper focuses on the application of agent-based computational economics and stochastic games in modelling the bilateral credit limit establishing behaviour of Participants in the Canadian Large Value Payments System. The data-driven stochastic game framework further illustrates how payments data, in conjunction with other financial market and credit data, can be used to assess emergent macroscopic outcomes in clearing and settlement systems from the underpinning interactions of autonomous decision making agents. The paper speaks to potential policy issues such as the effectiveness of the System-Wide Percentage, regulatory concerns about procyclicality and free-riding arising from the market microstructure behaviours, and design of the System.

Keywords: Agent-based computational economics, Market microstructure, High value payment systems, Bilateral credit limits, Intra-day liquidity management, stochastic games

JEL: C63, C73, D40, D47, D83, L10,

Contents

Executuve Summary

Market microstructure agent-based computational economics (ACE) models aim to tractably capture the structural dynamics of complex adaptive systems (CAS), which is often beyond the scope of conventional economic models. As such, ACE models are able to identify outcomes arising from feedback loops and trade-offs inherent to economic systems that traditional models are unable to handle as well. This paper applies the ACE approach to the understanding of participant decision making in financial market infrastructures with specific emphasis on the Canadian Large Value Transfer System (LVTS). More precisely, the paper disentangles the bilateral credit limit (BCL) decision making of LVTS participants into a series of incentive functions that combine to determine the pay-offs to participants in the LVTS. With these payoffs mapped to clearly defined incentive functions, the proposed ACE modelling framework can be empirically grounded. Moreover, being based on ACE principles, the theoretical framework is fully compliant with computability and complexity theories. That is, the decision making of modelled LVTS Participants is fully grounded in computationally feasible outcomes, and so there exists a set of algorithms that can solve the BCL decision-making problem in polynomial time and thus fully grounded in reality. Indeed, to this end, the proposed model illustrates how BCL establishment by LVTS Participants can be reduced to a multi-agent stochastic game for which there exist numerous reinforcement learning (RL), evolutionary computation (EC), artificial intelligence (AI), and machine learning (ML) based solutions.

From a regulatory and Principles of Financial Market Infrastructures (PFMI) standpoint, the proposed model highlights the complexity of the underlying market microstructure of the LVTS and FMIs more broadly. The empirically grounded theoretical model illustrates why the hitherto policy lever (i.e., the System Wide Percentage, SWP) may be less effective at affecting system liquidity and credit risk than previously believed. The model also draws attention to the extent to which the often overlooked collateral portfolio complexion of an LVTS Participant influences that Participant's intraday collateral/liquidity management. The value, to regulation and policy

making in the FMIs space and the LVTS, of a market microstructure evaluation is highlighted in the trade-off between credit risk and payment delay arising from System design. Specifically, under the risk model design of the LVTS, the ACE framework illustrates how it may be optimal for a Participant to continue to extend BCLs with respect to a distressed Participant, to ensure it can continue to receive timely payments from that Participant. This outcome is contrary to a perspective that assumes that BCLs will be withdrawn in times of stress. The ACE market microstructure model highlights this policy outcome since it is able to capture the System design implication of BCLs in the LVTS as collateral backed bets with the System against the default of a Participant's single largest counterparty credit exposure. A follow-on result of the ACE framework highlights that the design implementation of BCLs in the LVTS eliminates the incentive of Participants to engage in free-riding. Indeed, not only does the trade-off between credit risk and payment delay costs incentivise Participants to provide the System with liquidity, the intra-day liquidity management costs associated with the volatility in the spread between multilateral and bilateral risk controls, limit the extent to which free-riding is possible in the LVTS.

In summary, this analysis demonstrates the incentive-compatibility of the LVTS design with both efficient and stabilizing BCL-granting behaviour, as well as avoidance of free-riding in the LVTS. Moreover, with the focus of the analysis centred on BCL-granting behaviour and excludes Tranche 1 of the LVTS which is akin to a real-time gross settlement system (RTGS), an extension to the framework could assess the incentive-compatibility of RTGS design options given the upcoming implementation of the Lynx RTGS.

1 Introduction

Market microstructure centred agent-based computational economics (ACE) models are geared towards capturing the structural dynamics of complex adaptive systems (CAS) from not only the micro-behavioural, but also the institutional/rules perspective.[1] That is, such models assess evolving macroscopic outcomes from the microscopic standpoint of agents' behavioural incentives, their interactions, and other determinants of transaction costs, prices, quotes, volume, and trading behaviour that are inherent in institutions, rules and processes through which markets clear and settle on a daily and intra-day basis. The success of such models at capturing real world dynamics and the validation of their results is well documented in the financial markets literature where understanding market microstructures has gained the most attention.[2]

In this light, it is appropriate to consider high value payment systems (HVPSs) and other financial market infrastructures (FMIs) more broadly, as a CAS for a number of reasons. First, HVPSs are systems consisting of a sufficiently large number of interacting and adapting agents or financial intermediaries (both direct Participants and agency relationships). Each of these financial intermediaries, as corporations with internal business structures and processes (e.g. cash management, legal and compliance, global risk management, debt issuance pricing arising from creditor perceptions in capital markets such as through credit default swap (CDS) spreads), are themselves complex systems that continuously adapt to the multiplicity of environments they operate in. Thus the HVPSs are systems consisting of interacting and autonomous microscopic systems which can create outcomes at the macroscopic level that are influenced by internal in-

[1] Complex adaptive systems are complex macroscopic self-organising collections of relatively similar, partially connected and interacting micro-structures formed in order to adapt to changing environments and increase macro-structure survivability. Rather than models for predicting outcomes, CAS is a philosophical/theoretic framework for thinking about the world around us and thereby provides a variety of new options, giving the researcher more choice and freedom.

[2] See Bougheas and Kirman, 2014; Barde, 2016; Chen et al. 2011; Cona, 2008; Farmer and Joshi, 2002; Gallegati et al., 2011; Gilli and Winker, 2003; Kirman, 1991 & 2011; Kukacka and Barunik, 2016; LeBaron, 2005; Platt and Gebbie, 2016a &b

teractions between business units. Moreover, these interactions and adaptations, either at the microscopic level of the financial institution or at the macroscopic level of the HVPS, need not give rise to efficient equilibria.

Second, through these interactions between component agents, unplanned and potentially unpredictable system level regularities emerge which further feed back into the system and inform the interactions between the agents. In the Canadian context, from which this paper is focused as a case study, the Canadian HVPS, has, for example, been shown to exhibit properties associated with other systems that are considered to be self-similar.[3] Indeed, while Chapman, *et al*. (2016, *Mimeo*), do not cover the topic of self-similarity, the authors' results suggest that the Canada's Large Value Payment System (LVTS)[4] is self-similar with respect to the collateral pool usage as a result of multiple Participant defaults. That is, their results illustrated that exposures to the Bank of Canada in the LVTS under the residual guarantee became scale-independent and statistically exhibit a power law property at all levels of observation following the simulated default of five participants. This suggestion of self-similarity within high value payment systems including the LVTS should not come as a surprise given observations of self-similarity in other financial networks such as stock markets (see Campbell et al., 1997). Moreover, to the extent that the LVTS and other financial networks have been shown to exhibit scale-free properties (see Barabasi and Albert, 1999, Caldarelli et al, 2004, Champan and Zhang, 2010, and Chapman et al. 2011), it has further been proven by amongst others Song et al., (2005), Yook et al., (2005), Zhang et al. (2006), and Serrano et al., (2008) that properties of self-similarity are consistent with networks that exhibit scale-free degree distributions.[5]

As complex adaptive systems, one of the key challenges for researchers and policy makers in defining and assessing policy and design factors is the incompleteness and non-computability

[3]Self-similarity refers to the mathematical concept in which a system is exactly or approximately similar to one or more parts of itself.

[4]The Large Value Payment System (LVTS) is Canada's designated systemic HVPS utilised for amongst other things consumer wire transfers and interbank payments. A full review of the LVTS can be found in Arjani and McVanel (2006)

[5]Using a Bayesian implementation of the grid-search algorithm developed by Čopič, Jackson, and Kirman (2007), Champan and Zhang (2010) have shown that the LVTS consists of multiple community structures charaterised by a small Montreal-based cluster and a big Toronto-based group of the five big Canadian banks.

of outcomes within HVPSs. That is, the interaction of agents within HVPSs can produce novelty or surprises and Nash equilibria that challenge conventional wisdom in game theory where given fixed action sets innovation is not feasible. One such challenge is understanding LVTS Participants' decisions of setting Bilateral Credit Limits (BCLs). The ACE modelling framework permits look-through into the granular details underpinning the in-game pay-offs that define the evolution of outcomes within the LVTS in a computationally tractable way that is not immediately possible under the conventional game theoretic construct.[6] This allows for better identification of, or the potential for, market failure and appropriate policy responses; specifically in view of the international harmonization of guidelines as established under the 2012 and 2014 CCPMI-IOSOC Principles of Financial Market Infrastructures (PFMIs).In this light, the paper expands on the literature by merging both theoretical models (Bech and Garratt, 2006) and experimental simulations (Heemeijer and Heijmans, 2015, Addink et al., 2017) of financial institution interactions within financial market infrastructures such as HVPSs and how these interactions potentially give rise to market failures such as free-riding (see Diehl, 2013), payments timing (Heijmans and Heuver, 2014, Kaliontzoglou and Müller, 2016), and liquidity contractions and stress (Heijmans and Heuver, 2014, Alexandrova-Kabadjova et al., 2016). Moreover, just as Müller (2016) systematically assesses the different aspects of payments systems analysis and their significance from the perspective of central banks, this paper systematically breaks down the interactions between payment system participants and identifies possible areas of learning in the prudential regulation of FMIs, with emphasis on the LVTS. That is, the paper offers mechanisms through which FMIs' alignment with the PFMIs can be assessed in a more targeted fashion on the basis of incentives and other dynamics inherent in their market microstructures.

As a unique high value payment system structure, there has been little coverage in the literature on the market microstructure of the LVTS. This paper addresses this and contributes to the

[6] The theory of computation is a field of mathematical logic and computer science consisting of, computability theory, complexity theory, and Cantor-Gödel-Church-Turing quantum states, that addresses the questions of what can be truly computed mechanically, using for example mathematical algorithms, and their categorisation or classification according to the amount of time and space required to compute (see Breuer 2011, Velupillai and Zambelli 2012).

body of work on the LVTS by introducing an ACE model that captures the underpinning market microstructure of the system based on member incentives that drive BCL-setting decisions. In so doing, the paper captures granular details such as the pricing of credit risk, the costs of settlement delays, and the costs of collateral or liquidity (i.e., both the opportunity cost and the intra-day carry cost) that influence behaviour. The paper further describes how empirical data can be used to calibrate the theoretical model in a tractable manner. Indeed, just as León et al., (2016) argue that "studying isolated single-layer trading and registering networks yields a misleading perspective on the relations between and risks induced by partcipating financial institutions", assessing an FMI's ability to weather shocks in isolation of the underpinning market microstructure of the FMI may result in regulatory focus on the wrong policy levers or risk mitigation framework.

From a policy and oversight perspective the paper highlights the extent to which, changes to the system wide percentage (SWP), incentives for free-riding, and exogenous market rates impact BCL decisions. The framework proposed is also not limited in its application to the LVTS, but also extends to central counterparty (CCP) models that utilise BCLs in their loss sharing arrangements. For the CCP, these arrangements could include the transmission of losses from initial margins to variation margins and default fund contributions. Chande and St-Pierre (2016), propose the use of BCLs as part of the specification of loss-sharing arrangements to mitigate and smooth out margin increases at CCPs that would be considered as procyclical. This paper illustrates why the Chande and St-Pierre (2016) results may hold and why BCLs also help to prevent free-riding in loss sharing arrangements. Indeed, under appropriately designed risk modelling frameworks, the paper illustrates how it may be optimal for a participant to continue to extend BCLs with respect to a distressed participant, to ensure it can continue to receive timely payments from that participant. This outcome suggests that BCLs can in fact be countercyclical in providing CCPs with liquidity. The ACE market microstructure model also highlights the design implementation of BCLs in their ability to eliminate participants' incentives to engage in free-riding by linking this to their intra-day liquidity management costs.

Section 2 below provides a review of the LVTS as a level set for the dynamics presented in the later sections. Section 3 describes the underlying incentive functions behind the establishment of BCLs. In Section 4, these incentive functions are collated and the computational model introduced. Section 5 maps the various components of the theoretical framework to empirical data that provide the basis for model calibration. In Section 6, the next steps and approach in building out the ACE framework of modelling the setting of BCLs are described. Finally, the paper offers concluding remarks and potential policy implications along with future work in the development of the framework.

2 Overview of the LVTS

The LVTS is an electronic wire system that facilitates the transfer of Canadian-dollar payments between Payments Canada member financial institutions. The LVTS began operation in February of 1999 and is essential to the Canadian financial system, processing an average daily volume of approximately 22,000 payments equivalent to CA$171bn under a real-time net settlement model with final exchange of value at the end of day.

As of September 2017, sixteen financial institutions (FIs) and the Bank of Canada participate directly in the LVTS. These Participants provide LVTS payment agent services to other FIs, as well as domestic and foreign businesses and individuals, through contractual arrangements established between the Participant and its customers. The LVTS is characterized by two alternative payments streams ("tranches"), and Participants may use either stream to send a payment message through the LVTS. Each payment message is assessed against the applicable risk control tests and associated collateralisation scheme model, given the tranche it flows through, specified in the LVTS rules.[7]

[7] It should be noted that whilst a single LVTS settlement model underpins both streams, each stream is characterised by its own risk and collateralisation model.

At its core, the LVTS settlement model is defined by the concept of novation netting. This enables LVTS payments to settle in real-time given the applicable risk-control test.[8] Consequently, once a payment message passes the applicable LVTS risk control tests, the original bilateral net payment position between the sending Participant and the receiving Participant is extinguished, and replaced by a multilateral settlement obligation of the sending Participant vis-à-vis all other Participants in the System; importantly, this is also a conditional obligation of other Participants to make the owed Participant whole in the event of the sender's default.[9] At the end of each payment exchange period, it is these multilateral net positions that LVTS Participants settle, not the settlement of the individual payments processed by the LVTS.[10] Settlement in the LVTS is final once transfer of value ("central bank money") takes place over the books of the central bank through the settlement accounts LVTS Participants are required to maintain at the Bank of Canada.

2.1 LVTS Collateralisation and Intra-day Liquidity Model

Each of the two LVTS tranches has its own set of risk controls which combined guarantee mathematically that there is sufficient collateral value apportioned to the LVTS by Participants to ensure that settlement will take place in the event of the largest possible default ("Cover 1").[11] Counterparty credit risk within the LVTS arises from both the inability of the defaulting Participant to cover its end-of-day LVTS multilateral net position and the different collateralisation models. Under Tranche 1 (T1) payments are fully collateralised by the sending Participant; that is T1 payments will only pass the applicable risk control test if the value of the payment being sent does

[8] See CPA Bylaw No. 7 (sections 38 and 52).

[9] Once the individual payment message passes the applicable risk control, the underlying payment is deemed final and irrevocable.

[10] The multilateral net position is a single position, reflecting information on payments sent and received by a Participant in both tranches.

[11] In fact, available collateral can be and has been historically larger than the single largest default because not all maximum BCLs are extened to the same Participant. Moreover, exposure from any default is also reduced by the defaulter's own collateral.

not exceed the FI's collateral pledged plus its multilateral T1 position vis-à-vis the LVTS. In the event of a default, the Bank of Canada (the Bank) seizes the defaulting FI's collateral to cover that FI's end of cycle negative position. Therefore, T1 is a fully pre-funded *defaulter-pays* model. Conversely, Tranche 2 (T2) collateral requirements depend on the BCLs the FI has established that day; such that a Participant's T2 collateral pledge (known as a Participant's Maximum Additional Settlement Obligation or Max ASO) is set according to the largest BCL it extended during the day.

Before the start of each business day, Participants grant BCLs to each other, effectively setting the largest T2 bilateral net exposure they are willing to accept vis-à-vis another institution. The BCLs received by the Participant determine its Net Debit Cap (NDC) – the maximum net debit position that the Participant is able to incur during the payments exchange cycle. Consequently, the liquidity a Participant has available to send payments over the course of a typical exchange cycle is internally generated and a function of the BCLs it has received bilaterally, and multilaterally subject to the SWP and the intra-day flow of payments through T2.

Furthermore, unlike T1 where the sending Participant is required to collateralise the payments it sends, T2 payments are collateralised by the receiving Participant (who grants access to the system through the BCLs extended). In this regard, in the event of a default, should the defaulting Participant not have apportioned sufficient collateral to the LVTS to meet its time of default negative position (that is the defaulting FI has an Own Collateral Shortfall, OCS at default), the surviving Participants' ASO exposure to the outstanding balance at default is a function of the BCLs they have extended to the defaulting FI during the cycle.[12] The surviving Participants' exposure to the OCS is nevertheless capped at their respective Max ASOs. OCS exposures in excess of the

[12] If, following the LVTS pre-settlement period, a Participant enters LVTS settlement with a multilateral net debit position in the LVTS, it must apply for a discretionary advance from the Bank of Canada. If the Participant is unable to secure this discretionary collateral advance and therefore cannot settle their multilateral net debit position, they are in default for the purposes of the LVTS. In this case, the Bank of Canada seizes the defaulter's collateral and credits its settlement account with a non-discretionary advance of funds. If the amount of this advance is sufficient to cover the defaulting FI's negative multilateral net position, the Bank of Canada will immediately enable settlement. If, however, this advance is insufficient to cover the defaulter's multilateral net debit position, surviving Participants are obligated to meet an Additional Settlement Obligation (ASO) by advancing funds to the defaulting institution, secured by the collateral they posted to the Bank of Canada.

Max ASOs–which can only occur in the event of simultaneous Participant defaults– are covered by the Bank of Canada under the Residual Guarantee. Tranche 2 is therefore a *survivors-pay* scheme.

As with T1, T2 has controls built into the system to manage risk. For a T2 payment to pass the applicable risk control test, it must satisfy both the Bilateral Risk Control (BRC) and Multilateral Risk Control (MRC). More specifically, the payment must be (i) less than or equal to the difference between the BCL granted by the receiving Participant and the sending Participant's net bilateral T2 position vis-à-vis the receiving Participant,

$$P^T_{i,j} \leq \beta_{j,i} - \sum_{t=0}^{T-1} \left(P^t_{i,j} - P^t_{j,i}\right) \tag{1}$$

and (ii) less than or equal to the difference between the T2NDC of the sending Participant and its net multilateral T2 position vis-à-vis the system

$$P^T_{i,j} \leq \sum_{i \neq k \in N} \alpha\beta_{k,i} - \left[\sum_{i \neq k \in N} \sum_{t=0}^{T-1} \left(P^t_{i,k} - P^t_{k,i}\right)\right] \tag{2}$$

where N represents the set of all system Participants, $P^T_{i,j}$ is payment flow from Participant i to Participant j at time T, $\beta_{j,i}$ is the bilateral credit limit j extends to i,α is the system wide percentage. It is worth noting that the bilateral position, $\sum_{t=0}^{T-1} \left(P^t_{i,j} - P^t_{j,i}\right)$, at any given time t can be positive or negative. Where the bilateral position is positive, Participant i would have received more by way of payments from Participant j than it would have sent. Conversely, a negative bilateral position at time t would arise from Participant i having sent more value in payments to Participant j than it received.

Payments above $100m that breach the T2 risk controls are generally held in the LVTS central queue (the Jumbo Queue) until such time during the cycle that they can satisfy the controls, either as payment flows permit or receiving Participants are petitioned by the sender to increase the BCLs extended to it. Where BCLs are increased during the payment exchange cycle, such increases could potentially result in the sending Participant becoming the receiving Participant's largest BCL exposure and thus an increase in the sending Participant's Max ASO and T2 collateral requirement. Payments that breach the risk controls and do not trigger the Jumbo queue, will be rejected by the system.

The potential for payments not to satisfy all of the risk controls arises mathematically, at the limit, because there is a set of payments ($\mathbb{P} = \{P^t_{i,j} \, | BRC^t_{i,j} = P^t_{i,j} = MRC^t_{i,j} \}$) between any pair of Participants over the course of any given LVTS cycle large enough to just satisfy both the BRC and MRC. It follows therefore that breaches may occur because the receiving Participant extends a small BCL to the sending Participant relative to the confluence of bilateral payment flows and system-wide view of the optimal level of BCL to be extended to the sending Participant given payment flows. In other words, there is a set of bilateral payments ($\mathbb{P}^{\text{BRC}} = \{P^t_{i,j} \, | BRC^t_{i,j} < P^t_{i,j} \leq MRC^t_{i,j} \}$) between any pair of LVTS Participants over the course of a cycle for which bilateral risk controls will not be satisfied whilst multilateral risk controls are satisfied. Conversely, there is a set of bilateral payments ($\mathbb{P}^{\text{MRC}} = \{P^t_{i,j} \, | MRC^t_{i,j} < P^t_{i,j} \leq BRC^t_{i,j} \}$) for which the multilateral risk control is not satisfied but the bilateral risk control is satisfied. Such a breach might be the result of the receiving Participant extending BCLs in relative excess of what is optimal given system-wide expectations of the appropriate level of BCLs and payment flows. In this instance, the sending FI may need to obtain increased BCL extensions from multiple Participants.

As with other payment systems, given all intra-day payment flows result in Participants being either in positive or negative bilateral and multilateral positions, both individually (at Tranche level) and combined at the System level, collateralisation choices by Participants can be understood to be a zero-sum game. That is, as a closed system, the bilateral and multilateral positions within the LVTS sum to zero across all Participants and tranches.

Following the payment exchange across tranches, individual Participants may be in a multilateral net debit position, (i.e., they have sent more payments than they have received; they have a negative position), or a multilateral net credit position (i.e., they have received more payments than they have sent; they have a positive position).

End of cycle multilateral positions are settled over the books of the central bank such that negative positions in the LVTS at end of cycle imply securing overnight loans from the Bank of Canada at the Bank rate (top of a $50bps$ band). Conversely, positive end of cycle positions in the LVTS imply holding overnight deposits with the Bank and earning interest at the bottom end of the $50bps$ band. This creates an opportunity for position flattening prior to the end of the LVTS cycle (for example during Pre-Settlement at 6 pm) through maintaining expected multilateral positions close to zero or engaging with other Participants in the overnight market to lend or borrow at overnight interest rate somewhere within the $\pm 50bps$ band set by the central bank. These deals will tend to be struck at the overnight target rate which is the mid-point of the band.

LVTS Participants may also benefit from the use of Settlement Exchange Transactions (SETS) through which, intra-day exchanges of funds transactions are used to mitigate dislocation of settlement balances between the retail payments system the Automated Clearing and Settlement System (ACSS) and the LVTS. Using SET arrangements, Participants with projected positive end of cycle positions in the LVTS transfer funds to LVTS Participants with negative end of cycle positions in exchange for ACSS funds on the same business day.[13]

3 Incentive Functions Underlying BCL Decisions

Agents, the financial intermediaries within the LVTS, are specified as decision-making entities that possess an action set representing the universal set of all possible combinations of BCLs they can establish with respect to all other LVTS Participants. Each component of this universal set of actions, i.e., an individual vector of BCLs extended to all other agents, will be referred to as action profiles. The choice of the best action profile depends on idiosyncratic expectations

[13] It should be noted that due to the Bank of Canada utilising the LVTS as a platform for monetary policy implementation, positive and sometimes significantly large cash setting activity by the Bank each day to help alleviate small and transitory liquidity frictions and to support monetary policy implementation, results in all Participants (with the exception of the central bank) ending the LVTS cycle in a slightly positive position.

of the agents' current and future economic states as captured by the end of cycle rewards from participating in the LVTS, expectations of funding and liquidity risk, expectations of counterparty credit risk, collateral funding costs (both at start of day and intra-day) and the decisions of other Participants. It is further assumed that this decision about the choice of best action profile takes place at the beginning of the LVTS cycle.[14]

The associated trade-off between the costs and benefits of establishing BCLs are reflected in the agents' objective functions which they optimise by choosing the most appropriate action profile across all possible states. It is further assumed that, as computational agents, the ability of LVTS Participants to accurately predict end of cycle outcomes must be consistent with theory of computation. That is, the computational problem agents face when determining the appropriate level of BCLs to extend must necessarily fall under one or more of the many complexity classes and the associated algorithmic solutions given the bounds of computational time and space. In essence, agents are boundedly rational to the extent to which their decision making optimises their pay-offs according to update rules in a reasonable amount of time or computing resources. In the context of ACE, such update rules can take the form of arbitrary rules of thumb, or evolutionary computation and machine learning techniques such as, but not limited to, Reinforcement Learning, Belief Learning, Genetic Programming, Neural Networks, and Cellular Automata. The choice of update rule is problem–and research–objective dependent. It is important to note that unlike traditional economic models with rational agents who price perfectly (subject to functions and information frictions), under an ACE construct, systems are necessarily endogenously dynamic and evolving. Therefore, expected rewards upon which agents base their actions are typically not identical to the realised rewards; the latter being derived jointly from the actual collective actions of all Participants and actual observations of stochastic variables such as payment flows through the system. In the context of BCL decisions in the LVTS, these realised rewards then feed back into the BCL decision making process for the next round of BCL setting.

[14] The significance of this assumption will be expanded under the discussion on the intra-day cost of carry later in this section.

The following sequence of equations are the structural model of costs associated with the decision of whether to extend bilateral credit limits to other FIs within the LVTS; **Table 1** in Appendix A provides a listing of the variables used and their definitions.

3.1 Initial Cost of Liquidity

For any Participant, assuming each choice of BCL extension has the potential to be the largest BCL extended and would thus need to be collateralised, the initial or start of day cost is the implied opportunity cost to the collateral pledged by Participant j with respect to the reference Participant i. That is, the initial cost of liquidity for any day is the opportunity cost of the encumbered collateral for participation in the LVTS. This cost will, from a Participant's perspective, reflect applicable asset-class specific margin requirements or haircuts stipulated in the list of assets eligible as collateral under the Bank of Canada's Standing Liquidity Facility (SLF). Moreover, given that the collateral is encumbered and the extent to which assets used as collateral reflect a participant's portfolio management activity, the opportunity cost of the collateral will also reflect expected fluctuations in the fair market value of the collateral (i.e. a hold vs sell portfolio choice) as well as the expected spread between interest rate at which the cash management function of the Participant can borrow and lend collateral. For simplicity, the associated return from collateral portfolio management (η_j) is assumed to be exogenous and the initial cost of liquidity specified as

$$\lambda\left(\beta_{j,i}\right) = \alpha\beta_{j,i}\left(1+\eta_j\right) \tag{3}$$

The opportunity cost in this context is the implied gains foregone in other parts of Participant j's overall portfolio from the extending of BCLs to Participant i for the entirety of the LVTS cycle. The greater the BCL Participant j extends, the larger is the opportunity cost of granting BCLs.

Likewise, as the returns from collateral portfolio management and thus investment opportunities outside the payments system increases, so to will the implied opportunity cost of extending BCLs. Similarly, increases in the SWP, all else being equal, will result in a greater opportunity cost of liquidity provision.[15]

3.2 Intra-day Carry Cost of Collateral

The carry cost represents the intra-day cost of managing or adjusting BCLs. That is, the implied financing cost of increasing BCLs to another Participant intra-day or is the implied foregone intra-day investment gains arising from the over extension of BCLs to a given Participant.[16] The functional form for this intra-day carry cost of collateral is specified as:

$$c_{j,i}\left(\mu^{T}\right)=\sigma\left(\mu^{T}\right)\left[\prod_{x\in X}\left(1+w_{x}r_{x}\right)\right] \tag{4}$$

where

$$\mu_{j,i}^{T}=\left[\sum_{i\neq k\in N}\alpha\beta_{k,i}-\sum_{i\neq k\in N}\sum_{t=0}^{T-1}\left(P_{k,i}-P_{i,k}\right)\right]-\left[\beta_{j,i}-\sum_{t=0}^{T-1}\left(P_{j,i}-P_{i,j}\right)\right] \tag{5}$$

is the intra-day spread between the MRC and the BRC; the intra-day MRC-BRC spread (or risk controls spread), μ^{T}. Consequently, $\sigma\left(\mu^{T}\right)$ represents the average volatility in the MRC-BRC spread.

Since the carry cost is an implied cost which accounts for the cost of intra-day liquidity variations in the system, a Participant that efficiently manages its intra-day liquidity will in theory set BCLs

[15] It is worth noting that the initial cost of liquidity as modelled here does not account for any excess collateral pledged by System Participants over and beyond what they require for the clearing and settlement of payments. The empirical question as to why HVPS participants potentially pledge more collateral than they require has been covered by McPhail and Vakos (2003) but remains an open and intriguing question. However, to the extent than this collateral pledging behaviour is not tied to it is not tied to the intraday costs of liquidity management, such behaviour is not considered in this analysis. The approach taken here is nevertheless flexible enough to permit the inclusion of a model of excess collateral pledging as an additional incentive function.

[16] Given that BCLs can only be adjusted upwards during the course of a payment day, the foregone investment return is the opportunity cost of locked in excess collateral.

with reference to the multilateral risk controls. In other words, Participant j will look to minimise the intra-day MRC-BRC spread volatility, $\sigma\left(\mu^{T}\right) \in \mathbb{R}_{\geq 0} \forall \mu^{T}$, on a daily or cycle basis with respect to any reference Participant i. The intra-day volatility in the risk controls spread is thus a proxy measure of the extent to which individual Participants, subject to stochastic payment flows, over (or under) collateralise the system with respect to any given reference Participant relative to all other system Participants. For example, if it is assumed that bilateral payment flows follow a standard normal distribution (i.e.$(P_{j,i} - P_{i,j}) \sim N(0,1)$) then, a Participant j with a consistently large volatility in the intra-day risk controls spread against some reference Participant i would suggest that the individual BCL decision is poorly aligned with the BCL decisions of all other Participants with respect to the reference entity. In such a situation, the decision-making Participant would likely be faced with intra-day pressure to increase (or decrease in the following cycle) the BCLs it extends to the reference entity. Conversely, low volatility in the intra-day risk controls spread would suggest a tighter alignment of individual BCLs with the wider system level BCLs against the same reference entity. This in turn implies that calls to increase BCLs would likely occur at the multilateral level rather than at the bilateral, and the individual decision-making Participant faces less internal pressure, due to potential earnings foregone (or paid) on collateral funding sources, to reduce (or increase) BCLs proceeding cycles.

With volatility in the MRC-BRC spread being linked in the above mentioned manner to internal and external pressures to decrease or increase BCLs, either in the current or proceeding cycle, the intra-day carry cost of established BCLs positions can be measured and factored into the decision making of Participants. Thus, given the pool of funding sources (including but not limited to central bank overdraft facilities, interbank lending, trading or intermediation portfolio, etc) each Participant will have an associated asset mix and expected returns (w_x and $r_x$respectively).[17] The carry cost therefore reflects the market risk associated with the intra-day management of

[17]For the sake of simplicity, these asset mixes and expected returns are assumed to be exogenous. This assumption may be relaxed in further research once the initial model is fully calibrated.

encumbered collateral often overlooked in the literature.[18]

It should be noted that the intra-day carry cost of collateral is also subject to central bank, government, and other regulatory actions that result in an obscuring of the market pricing of assets and therefore can influence BCL and wider collateral decisions. Changes to the types of assets that are considered eligible collateral will influence the intra-day carry costs. For instance, the Bank of Canada's initiative to allow LVTS Participants to pledge their non-mortgage loan portfolio as LVTS collateral under certain conditions, which is otherwise composed of marketable securities that have collateral value outside of the LVTS, widened the eligible collateral pool and significantly reduced the influence of cost of the intra-day carry cost of collateral in the System. It is also noteworthy that, unlike in the case of the initial opportunity cost of collateral, all else being equal, an increase (or decrease) in the SWP has no impact on the intra-day cost of carry. This is because, if again it is assumed that bilateral payment flows are normally distributed, increasing (or decreasing) the SWP results in a scale increase (or decrease) in the lower bound of MRC-BRC spread and so too the mean spread whilst leaving the variation of the MRC-BRC spread around its mean are unchanged. This highlights the importance of market risk outside the payment system on the intra-day carry cost.

3.3 Counterparty Credit Risk Exposure

The survivor pay cost or counterparty credit risk exposure accruing to Participant j is the per dollar expected loss given the default of a reference Participant i

$$\psi_j(i) = \tau_i (1 - \phi_i) E \left[\sum_{i \neq k \in N} \sum_{t=0}^{T-1} (P_{k,i} - P_{i,k}) \right] b_{j,i} \tag{6}$$

[18] A further point of note is that by internalising through the intra-day carry cost of established BCL positions, LVTS Participants are faced with an endogenous charge for free-riding by undercutting of BCLs. Having to make intra-day adjustments to BCLs due to consistently under-extending BCLs will result in larger intra-day costs due to volatility in market returns.

where $b_{j,i} = \left[\frac{\beta_{j,i}}{\sum\limits_{i \neq k \in D \subseteq N} \beta_{k,i}} \right] \in B_j$.

The term $b_{j,i} \in (0, 1]$ is the survivor pay component of Participant j's exposure to Participant i's default. That is to say, given that Participant i defaults during the course of a payment cycle, the survivor pay component is Participant j's relative per dollar share of system-wide credit exposure to Participant i at default. While this relative share of exposure to default losses is not known to Participant j at the time it establishes BCLs with respect to Participant i, Participant j is able to approximately set an upper bound to which it is willing to be exposed to these losses. Consequently, $b_{j,i}$ represents a decision parameter which Participant j must set as part of establishing BCLs and B_j is the vector $B_j = \{b_{j,0}, b_{j,i}, \ldots, b_{j,N-1}\}$ which is a measure of its risk profile relative to the rest of System. [19]

The term $E\left[\sum\limits_{i \neq k \in N} \sum\limits_{t=0}^{T-1} (P_{k,i} - P_{i,k})\right]$ is the expected value of Participant i's stochastic multilateral net debit position at the time of default. It is important to note that the actual multilateral net debit position of agent i is unknown but assumed to be modelled or approximated by Participant j.

Finally, the terms τ_i and $(1 - \phi_i)$, respectively, are Participant i's probability of default and the proportion of its multilateral net position that is not recoverable from the collateral it pledged to the System. This reflects the fact that financial institutions do not necessarily default on their entire multilateral net debit position at the end of the cycle but only that portion over and above the value of the total collateral the defaulting Participant has apportioned to the System (i.e., their OCS).

From the counterparty credit risk incentive function, it is expected that as Participant j's exposure to agent i increases relative to all other LVTS Participants, the cost to agent j of a potential default by agent i also increases. Similarly, as the Participant i's default probability rises, due

[19] it is noteworthy that since $b_{j,i}$ is a share of loss exposure relative to the rest of the System, it represents Participant j's beliefs about the BCL setting behaviour of all other System Participants with respect to Participant i. $B_j = \{b_{j,0}, b_{j,i}, \ldots, b_{j,N-1}\}$ is therefore a proxy for how Participant j believes BCLs will be established across the System.

for example to a downgrade in credit ratings, Participant j's survivor pay credit risk exposure to i increases given extended BCLs. The survivor pay credit risk incentive function therefore drives Participant j to maintain as low an exposure to Participant i as possible in any payment cycle.[20]

3.4 Payment Delay Cost

The payment delay cost refers to the cost to the Participant of delays to inbound payments due to not extending sufficient BCLs to other Participants to ensure timely receipt of payments from those Participants. Such a cost could arise from the discontent from the recipient bank's clients or the recipient bank being unable to costlessly send payments back to the sender or other system Participants.[21]

$$d\left(\beta_{j,i}\right) = \left(\sum^{T-t'} Q_{i,j}^{T-t'} * \sum^{T-t'} P_{i,j}\right)\left(1 - exp^{-\sqrt{\frac{\alpha\beta_{j,i}^{1.5}+\left\langle P_{i,j}^{T}\right\rangle^{2}}{\alpha\beta_{j,i}^{2}+\left\langle P_{i,j}^{T}\right\rangle^{1.5}}}}\right) \tag{7}$$

where the payment delay cost multiplier is non-linear quadratic (hence raising to the powers of 1.5 and 2) and has the limit condition

$$\lim_{\beta_{j,i}\to 0}\left(1 - exp^{-\sqrt{\frac{\alpha\beta_{j,i}^{1.5}+\left\langle P_{i,j}^{T}\right\rangle^{2}}{\alpha\beta_{j,i}^{2}+\left\langle P_{i,j}^{T}\right\rangle^{1.5}}}}\right) \to 1$$

[20] The reader should be aware that the counterparty credit risk component can be extended by noting that, as currently defined, $\psi_j\left(i\right)$, is the contingent leg of a credit default swap (CDS) where, the premium leg is the intra-day net payments flows Participant j is able to exchange with agent i using recycled liquidity. Consequently, the decision to extend BCLs to Participant i can be seen as agent j selling a CDS to the System; thus it is betting against the default of Participant i. This is a critical point to note when assessing the underlying risk profile and collateral model in Tranche 2. That is, contrary to the popular misconceptions in the existing literature given the preponderance of RTGS systems, the existence of this embedded CDS implies that Participant j derives direct liquidity recycling benefits from the risk exposure to Participant i. As such the MNDP exposures Participant i is able to accumulate is a direct product of the BCL decisions, and thus the system-wide risk preferences, of all other Participants. The literature's over reliance on MNDPs to price risk in the LVTS, especially with respect to Tranche 2, is thus potentially erroneous.

[21] Note that it has been assumed that the only constraint is that payments which must or can be settled same day are settled. Moreover, even within that set of payments for which same day settlement is a requirement, the timing of the Participants' submission of those payments into the LVTS is assumed to be entirely at the Participants' discretion.

Figure 1: Inbound Payment Delay Cost

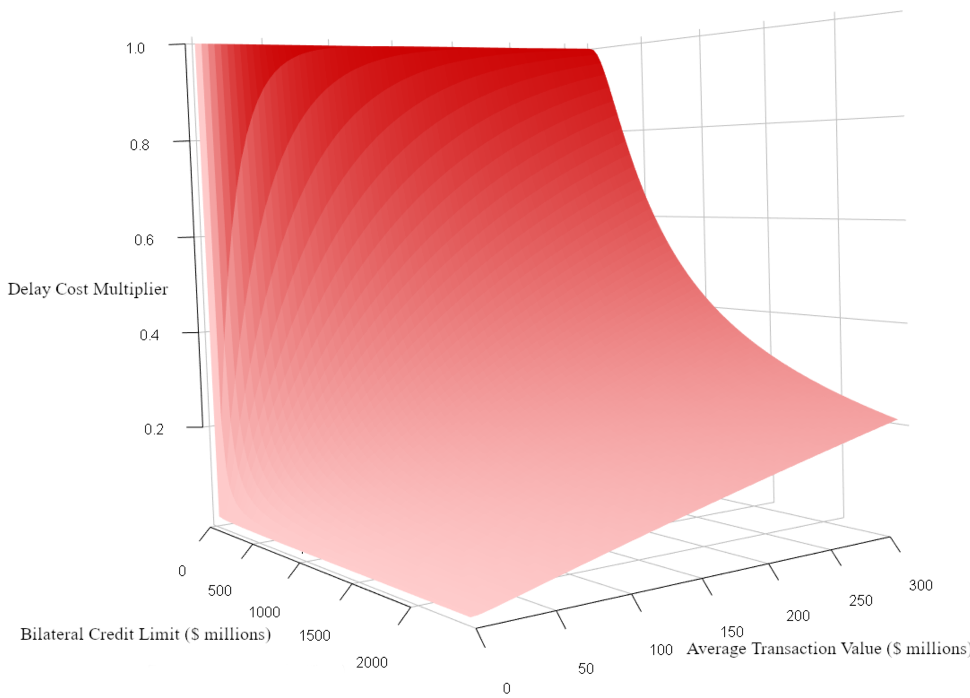

Figure 1 provides an illustration of the functional form of the cost to the receiving Participant for delays to inbound funds, given BCLs extended and the value of inbound payments.[22] The properties of delayed inbound funds availability cost are:

1. For a set of inbound payment values, as BCLs extended tend towards zero, the penalty cost of payment bottlenecks tends to 1; conversely, as BCLs tend to infinity, the penalty cost of payment bottlenecks asymptotically tends to 0.
2. At any level of BCL extension, as the inbound payments value tends to zero, the cost of a

[22] It should be noted that the exact functional form of the delay cost multiplier has been constructed arbitrarily with a view to be endogenous and dynamically set with payment flows between any pair of Participants while satisfying the limit condition of assuming a value between 0 and 1. A more precise functional form may be arrived at through model calibration. However, there is no loss of generality or applicability given this endogenous representation may be viewed as an extension of Galbiati and Soramäki (2008) and others who consider the delay cost as a constant and exogenously specified parameter.

potential payments bottleneck tends to zero. This follows from the declining probability of a bottleneck as payment values decline.

3. As transactions flow, relationships between Participants become stronger, either in terms of payment value or volume, the cost of a delay increases.

The third property of the delay cost function, highlights the potential for hidden tiering or clustering within the payments network that may give rise to reciprocal relationships. In other words, as either the volume ($\sum^{T-t'} Q_{i,j}^{T-t'}$) or value ($\sum^{T-t'} P_{i,j}$) of payments Participant j receives from Participant i increase, Participant j has an incentive to extend larger BCLs to Participant i in order to recycle the liquidity generated from those inbound payments to facilitate payments it has to make either back to Participant i or other system Participants. Likewise, Participant i, given the value or volume of transactions it receives from Participant j, has the incentive to extend BCLs to j so as to benefit from liquidity recycling arising from inbound payments. As a consequence, whilst there is no direct compulsion for Participants to extend similar BCLs to one another, there exists a built in possibility for reciprocal behaviour to emerge in the extension of BCLs.

3.5 Overnight Market Earnings

Since it is assumed that payments are stochastic, end of cycle multilateral net debit positons are likely to fluctuate over time and as such, Participants will typically be faced with alternating overnight market positions as either borrowers, lenders, or not engaging in overnight lending. End-of-cycle multilateral net debit positions can thus be thought of as representing the "states" Participants enter the overnight market in (i.e., as borrowers, lenders or non-participating). Moreover, given these states are quantifiable by virtue of being financial obligations as borrowers, lenders or non-participating, it is possible to quantify the rewards associated with these states. Representing the Participant's end of cycle multilateral net debit position states as, $s_j \in \mathbb{S} := \mathbb{R}$, the earnings from overnight lending to flatten positions can be specified as

$$\varpi\left(s_j\right) = \sum_{j \neq k \in N} \sum_{t=0}^{T} \left(P_{k,j}^{t} - P_{j,k}^{t}\right)\left(1 + r_f\right) \tag{8}$$

where $\varpi\left(s_j\right) > 0$ means the Participant is lending in the overnight market, $\varpi\left(s_j\right) < 0$ implies the Participant is a borrower, and is non-participating in overnight lending if $\varpi\left(s_j\right) = 0$. Moreover, since the bilateral risk control in Equation 1 stipulates that the maximum payment value that Participant j is able to receive from any other Participant k at any point during the LVTS cycle is the sum of the BCL it granted to Participant k and its bilateral net credit position via-a-vis Participant k , Equation 1 can be substituted into Equation 8 for $P_{k,j}^{t}$ such that,

$$\varpi\left(s_j\right) = \sum_{j \neq k \in N} \sum_{t=0}^{T} \left[\left(\beta_{j,k} + \sum_{t=0}^{T-1} \left(P_{k,j}^{t} - P_{j,k}^{t}\right)\right) - P_{j,k}^{t}\right]\left(1 + r_f\right) \tag{9}$$

Thus, Participant j's end of cycle multilateral net debit position and by extension, its ability to garner income in the overnight market, is a function of the BCLs it extends during the LVTS cycle. This would further imply that, regardless of the SWP and credit risk exposure, Participant j, by virtue of simple System's rules on payment process, is incentivised to establish as much in the way of BCLs as possible (other things being equal).

4 Computational Model of BCL Decisions

Breaking down the BCL extending decision of LVTS Participants into incentive functions as described above, the intra-day liquidity management problem faced by financial institutions in the LVTS can be reduced to a multi-agent coordination game. That is, subject to stochastic payment flows, the actions of other Participants, and market or policy factors highlighted, this coordination problem entails each Participant deciding at the start of each LVTS cycle, how much by way

PAYMENTS CANADA

of shared loss exposure to assume and BCLs to extend to others in the LVTS.[23]

Given the stochastic component of this coordination game is influenced by Participants' own decisions, the setting of BCLs can be thought of as a multi-agent Markov Decision Process (MDP) or more generally a stochastic game.[24] In the BCL stochastic game, actions, $a_j \in \mathbb{A} := \mathbb{R}\ \forall j \neq i \in N$, taken by each Participant j represent the vector of BCLs and loss sharing exposure set by that Participant with respect to all other Participants i. Similarly, the vector $g \in \mathbb{G}$ is the joint set[25] of individual actions played by all Participants. Furthermore, the process of an agent transitioning from state s to some future state $s' \in \mathbb{S}$ follows a transition model $\mathbb{T}(s, g, s')$ with probability $p(s' \mid a \in g,\ s)$. Combining the incentive functions, an LVTS Participant's single stage reward function for establishing BCLs in relation to other Participants is specified as

$$R_a^j(s, g, s') := \varpi(s) - \sum_{j \neq i \in N} \left[\lambda(\beta_{j,i}) + c\left(\mu^T\right) + \psi_j(i) + d(\beta_{j,i})\right] \tag{10}$$

Moreover, given the repeated nature of the BCL stochastic game with long-lived agents, each Participant j will, over an infinite horizon, choose a_j for any state s that maximises the discounted long-run reward of providing liquidity in the LVTS. By maximising discounted long-run rewards that are quantifiable in dollar terms over an infinite horizon, Participants are seen as making a distinction between short-term and long-term rewards and their associated policies.[26] Thus Participant j's expected reward, for action $a_j \in g$ in state s is given by the Q-function

$$Q_j^\pi(s, a_j \in g) = E\left\{\sum_{s=1}^{\infty} \gamma^s R^j_{a_j \in g} \middle| s = s, a = a_j, \pi\right\},\ \gamma \in (0, 1] \tag{11}$$

[23]Whilst it is operationally useful to distinguish between "Standing" vs "Intra-day" BCLs, at a theoretically efficient fixed point that may be considered the equilibrium point, if the BCL decision at the start of day was appropriate, there should be no intra-day changes in BCL extension. Where such intra-day BCL changes occur, these should be to address unforeseen extraordinary payments or events.

[24]See Bauerle and Rieder (2011) and Hu and Yue (2008) for more details on MDPs in finance. For a review of stochastic games and multi-agent systems see Condon (1992), Filar and Vrieze (1997), Mertens and Neyman (1981), Neyman and Sorin (2003), Schwartz (2014), Shapley (1953)

[25]The space of all possible joint actions, i.e. the universal set of all possible combinations of action profiles is denoted by $\mathbb{G} = \mathbb{A}_1 \times \mathbb{A}_2 \times \mathbb{A}_3 \times \ldots \times \mathbb{A}_N$

[26]Tsitsiklis and Van Roy (2002) have shown for temporal difference learning that as the rate at which future rewards are discounted approaches 1, the value function produced by infinite horizon discounted models converges to the differential value function generated by average reward models.

or

$$Q_j^{\pi}(s, a_j \in g) = \sum_{s' \in \mathbb{S}} \mathbb{T}\left(s, g, s'\right)\left[R_{a_j \in g}^{j}\left(s, g, s'\right) + \gamma V_j^{\pi}(s')\right], \ \gamma \in (0, 1| \tag{12}$$

Likewise Participant j's expected reward, in state s given the action $a_j \in g$ will be given by the value function

$$V_j^{\pi}(s) = E\left\{\sum_{s=1}^{\infty} \gamma^s R_{a_j \in g}^{j} \middle| s = s', \pi\right\}, \ \gamma \in (0, 1| \tag{13}$$

or

$$V_j^{\pi}(s) = \sum_{a_j \in g} \pi\left(s, a_j \in g\right) \sum_{s'} \mathbb{T}\left(s, g, s'\right)\left[R_{a_j \in g}^{j}\left(s, g, s'\right) + \gamma V_j^{\pi}(s')\right], \ \gamma \in (0, 1| \tag{14}$$

and

$$V_j^{\pi}(s) = \sum_{a_j \in g} \pi\left(s, a_j \in g\right) Q_j^{\pi}(s, a_j \in g) \tag{15}$$

The recursive Equations 11 and 13 are Bellman equations for which numerous reinforcement learning and evolutionary computing approaches can be applied to computationally determine the optimal policy, $\pi^* = \left\langle \pi_j^*, \pi_{-j}^* \right\rangle$. In the stochastic game setting, π represents the joint policy such that π_j is Participant j's policy response to the policy π_{-j} of all other Participants, i.e. $\pi_j \in BR_j\left(\pi_{-j}\right)$. Moreover, for all possible states, $\pi_j^* \in BR_j\left(\pi_{-j}\right)$ is Participant j's best response policy to other Participants' policy if and only if

$$\forall \pi_j \in \left[\mathbb{S} \times \mathbb{M}\left(\mathbb{A}_j\right)\right], s \in \mathbb{S} \ \ V_j^{\left\langle \pi_j^*, \pi_{-j} \right\rangle}(s) \geq V_j^{\left\langle \pi_j, \pi_{-j} \right\rangle}(s) \tag{16}$$

where $\mathbb{M}\left(\mathbb{A}_j\right)$, as in matrix games, is the matrix row vector of the set of all possible BCL establishing actions Participant j can apply. The Nash equilibrium is therefore the vector of policies across all Participants, such that all the contained policies are the best response policies upon which no Participant can improve by changing policies. This is represented more formerly as:

$$\forall j \in N, \ \ \pi_j \in BR_j\left(\pi_{-j}\right) \tag{17}$$

5 Data

In this section, empirical data is mapped to the various components of the model presented above to illustrate how the ACE framework is empirically grounded and relevant to the data. Starting with the initial cost of liquidity, for simplicity this can be assumed to be the spread between the 3-month Canadian Interbank Offer Rate and the 3-month Canadian treasury bill rate. In reality the initial cost of liquidity will be Participant specific and reflect margins derived from the composition of the collateral portfolio pledged to the Bank of Canada. This follows from standard convention of using short-term domestic sovereign debt to proxy the risk-free rate because of the market assumption that the probability of a highly rated sovereigns such as Canada, currently AAA rated at the time of writing, defaulting on its obligations to be close to zero. Furthermore, the typical large size and deep liquidity of the market for short-term treasuries contribute to the perception of safety.

The intraday cost of carry is mapped from a market risk perspective to daily financial and commodities market index returns in the industries considered key reflections of the Canadian economy and balance sheet exposures of LVTS Participants (i.e. agriculture and fisheries, forestry, metals, energy, TSX Composite and the Canadian All Bond Index). With regard to the MRC-BRC spread generated at a one minute frequency from LVTS transactions data, it is observed that the daily volatility in spreads (see **Figure 2**) tend to be relatively tight, falling between an upper and lower bounds. There are also observable trends in the data between certain pairings of Participants (the tuple of the BCL setting Participant and the reference Participant) suggesting that the dispersion in risk controls are pairwise dependent and can increase or decrease over time.

The magnitude of the daily volatility reflects the size of payments flows between pairs of institutions. Larger value flows between a pair of Participants will, all else being equal, result in larger variations around the mean MRC-BRC spread for that pair. However, where pairwise spreads do not continuously widen or tighten over time, daily volatility between these pairs has gener-

Figure 2: Example Intra-day Spreads between Multilateral and Bilateral Risk Controls

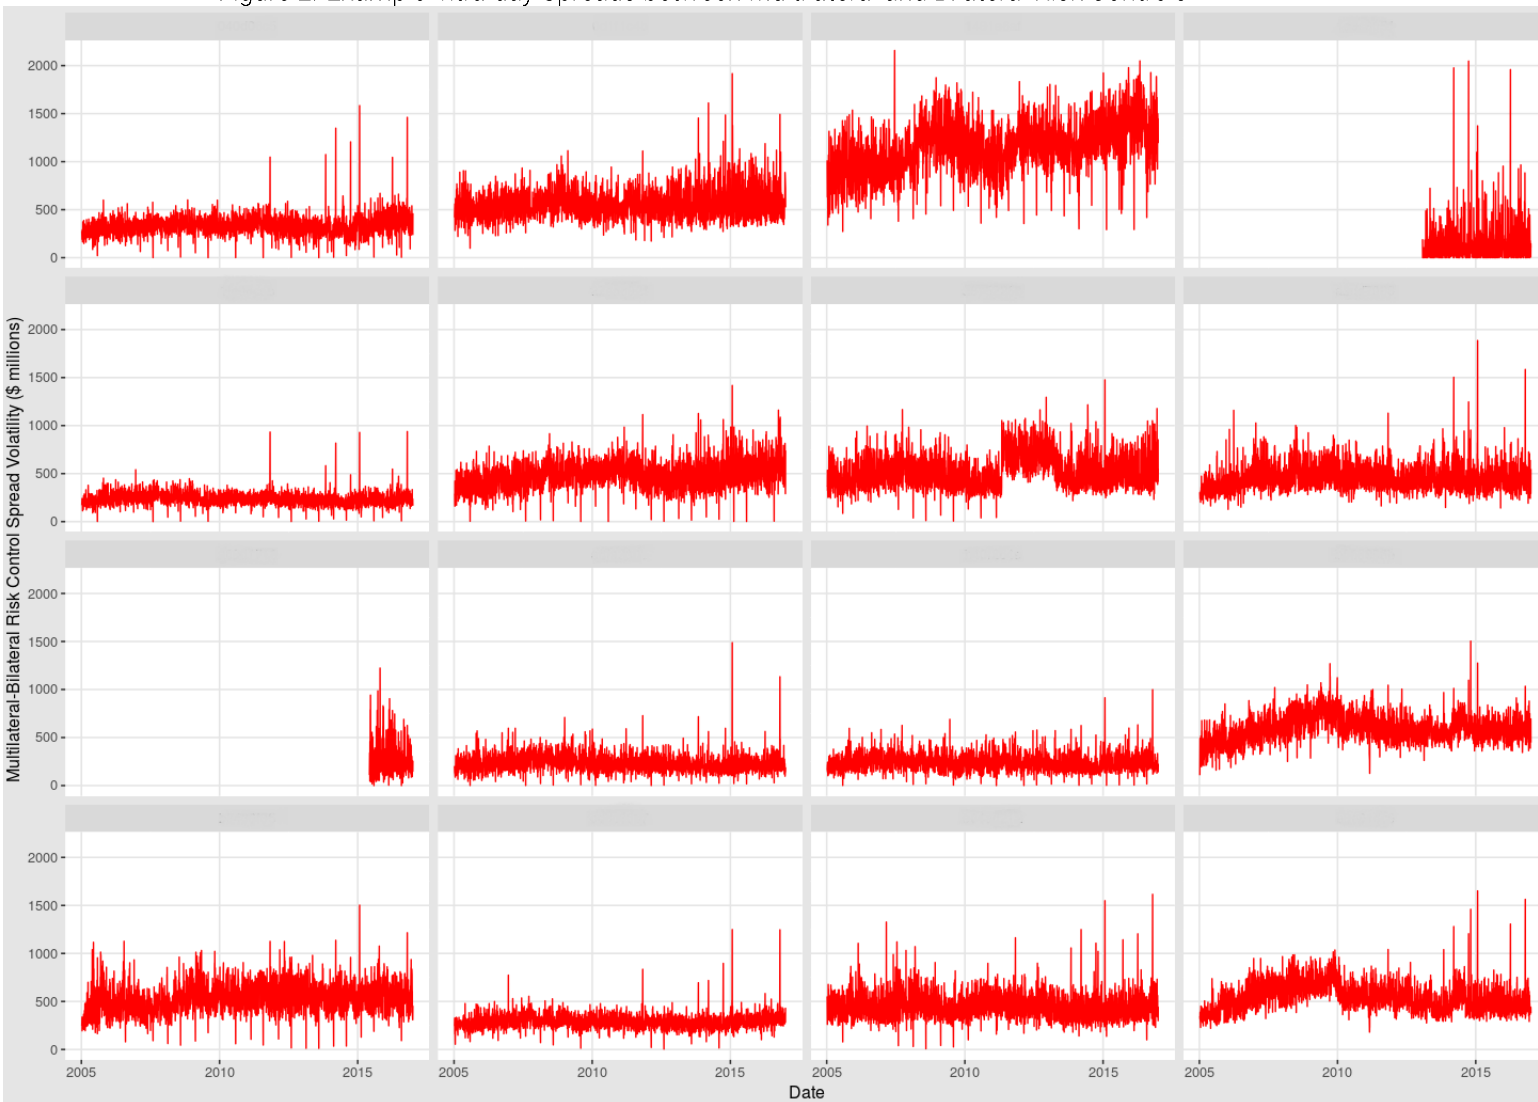

The figure plots the daily volatility in intra-day spread between the multilateral and bilateral risk controls (MRC-BRC spread) associated with payments sent by a single LVTS Participant (the reference Participant *i*) to all other LVTS Participants *j* between January 2005 and December 2016. Each series represents the volatility in pairwise MRC-BRC spread between the reference Participant and a receiving Participant.

ally fluctuated around its long-term mean. Observations within certain pairs of increasing daily volatility over time suggests a widening of the MRC-BRC spread and thus reduced alignment between the BCL setting Participant in the pair and the wider LVTS with respect to the reference Participant. Conversely, empirically observed declines in pairwise spread volatility suggest that for those pairs, there is improving coordination between the BCL setting entity and the wider LVTS.

Default probabilities for each of the LVTS Participants are based on the one-year default probability tables published by the Standard and Poor's (S&P) in its Annual Global Corporate Default Study And Rating Transitions. Of the sixteen Participants (excluding the Bank of Canada) in the LVTS over the period between January 2005 and December 2016, eight had ratings of between double-A minus and double-A. The lowest observed rating was triple-B minus. One Participant is a crown corporation and thus assumed to be triple-A rated in accordance with the provincial government under which it was established. Default probabilities associated with these ratings ranged from 0.00% to 0.28% with an average of 0.05%.

The extent to which payment flow relationships may influence delay cost modelling can be illustrated through the data. The network diagram in **Figure 3**, depicts the relationship between financial institutions in the LVTS with respect to payment flows. Under the force atlas layout, nodes at the center of the graph and their proximity represent those LVTS Participants with the strongest connectivity to one another and to other Participants in the network. The nodes farther away from the central core and other nodes are less connected to other Participants in the network.

In the network, the colour scale of the nodes maps the extent to which a particular Participant is, on average, a net receiver or a net sender of payments. At the extremes of the spectrum, red nodes are net senders whereas blue nodes are net receivers; as nodes get lighter in colour the degree to which they are net senders or net receivers diminishes. Likewise, the thickness and colour of the edges in and out of each of the nodes depicts the relative proportion of payments

Figure 3: LVTS Participant Payment Flow Relationships

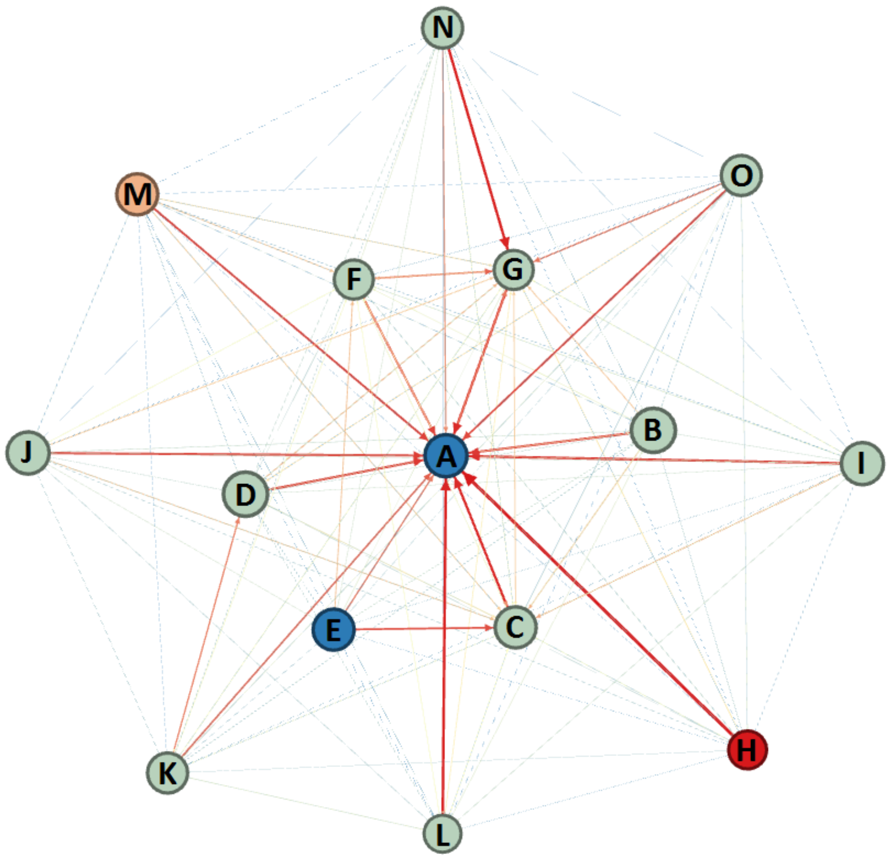

The depicted Force Atlas network layout is derived from sender to receiver payment flows by value between 2005 and 2016.

from the source node to the sink node. The thicker and closer to red the edge is, the larger the proportion of the value of payments from the sending financial institution to the receiving institution relative to the value of payments sent to all other Participants in the LVTS. The concentration of red edges with fairly similar thickness at the centre of the network indicates that the LVTS has underpinning core-periphery characteristics.[27] This implies that payment to and from the group of LVTS Participants at the core have greater weight and will incur greater delay costs than those

[27] Chapman and Zhang (2010) cover these core-periphery characteristics in more detail.

to and from Participants at the periphery.

End of cycle positions upon which individual Participant's state transition probabilities derive are taken as the combined final tranche 1 and tranche 2 multilateral net debit positions prior to 6 PM (**Figure 4**). Since these positions are elements of the set of real numbers, for computational ease the number of possible states may be limited to ten or more categories specified as percentiles (5%, 10%, 20%, 25%, 40%, 50%, 60%, 75%, 80%, 95%) of the distribution of end of cycle MNDPs. Overnight lending or borrowing upon the realisation of end of cycle positions within each of these states is assumed to be done at the Canadian Overnight Repo Rate (CORRA).

Using the last recorded MNDP prior to 6 PM pre-settlement ensures that Bank of Canada cash-setting to neutralise government activities in the System, monetary policy implementation, and other interbank flattening activities which may obfuscate final settlement positions are excluded from end-of-cycle position calculations as these represent monetary policy activity and not regular payment settlement.[28] An interesting observation from the end of day multilateral net debit positions data is that the various Participants' series appear stationary and strongly mean reverting with a mean of zero.[29] The exception to this is the Bank of Canada, which typically maintains a target cash setting position each day, and therefore a negative closing positions, of approximately $2.5bn and $5-6bn during the global financial crisis and recovery between 2008 and 2011.

[28] See Arjani and McVanel, 2006 and Kamhi, 2006 for an overview of cash-settlement and monetary policy actions in the LVTS.

[29] This outcome is not entirely surprising since the Bank of Canada's Standing Loan Facility (SLF) policy creates an incentive for LVTS participants to flatten out their end of day positions through interbank lending.

Figure 4: LVTS Pre-Settlement End of Day Multilateral Positions (millions)

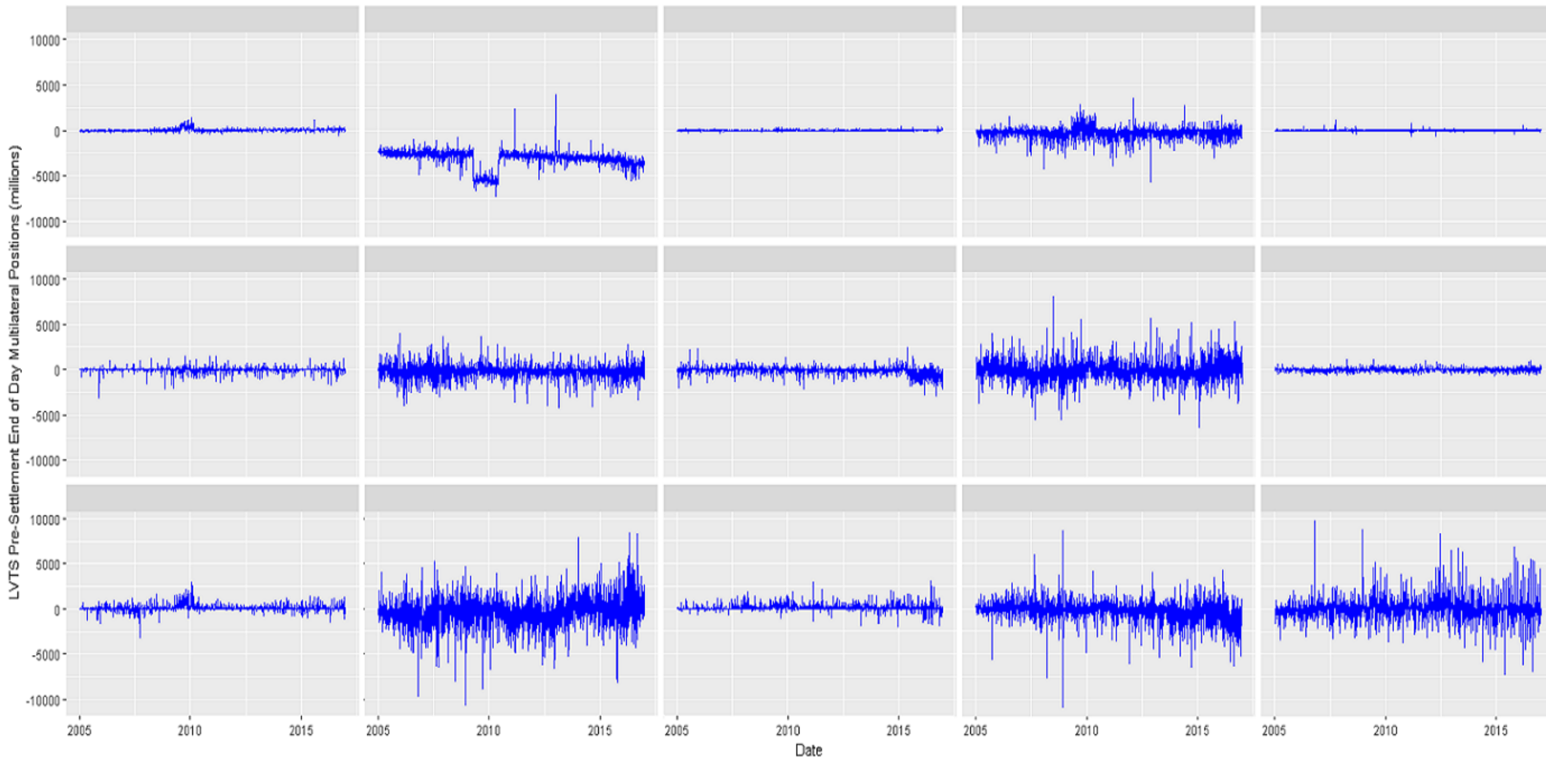

The figure plots the daily end of cycle multilateral net debit positions of all Participants in the LVTS 2005 and 2016. Each series represents a single Participant's end of cycle position.

6 Next Steps: Developing the Simulation Laboratory

Having specified the theoretical framework and mapped the various incentive functions to real world data, the next steps will be to develop an agent-based simulation laboratory which can be calibrated using the data. The development of this agent-based simulation laboratory necessarily entails the definition and instantiation of the agents. According the theoretical framework of BCL establishment as a stochastic game, the Participants in the model are autonomous decision-making agents that formulate a rule sets of actions subject to their current state, available actions, and perceptions of the actions taken by others. Moreover, these rule sets are adapted and amended as the Participants learn through experience gained from their repeated interaction with other LVTS Participants. This of course means the simulation laboratory must capture the modularity in the specification of the Participants. Agents most therefore be treated as self-contained, identifiable, discrete individuals with a set of characteristics or attributes, behaviours, and decision-making capabilities. In Section 3 these characteristics, attributes, and behaviours defined under the incentive functions. Their decision-making capabilities were specified in Section 4 wherein the set BCLs and obtain continuous feedback relating their BCL setting and information from their environment (the LVTS and broader markets) as they interact with other LVTS Participants.

In Figure 5 the interaction between LVTS participants and their interaction with the LVTS is depicted in multi-layer grid space and the simulation laboratory could be constructed in this fashion. Under this grid space layout, agents' positions on the grid space could represent their proximity to one another with respect to which LVTS Participants form closer bonds through the setting of BCLs given information received from the environment. Initially, Participants may be placed randomly on the grid, however, as they interact and learn from the information they receive from the environment, they would adapt their BCL setting behaviours, and thus update their grid positions,

Figure 5: LVTS Participants Interacting in a Grid Topology Agent Space

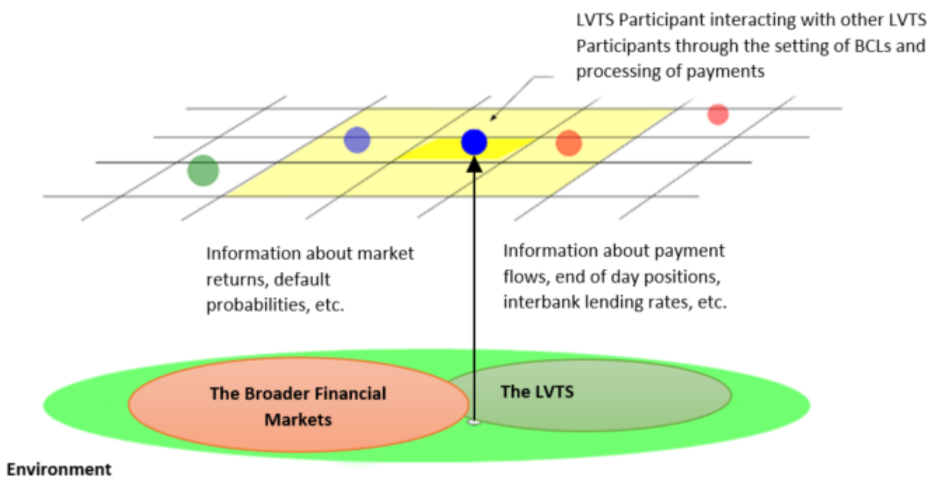

according the Bellman update rules presented in Section 4. It is worth noting that subject to calibration the emerging grid positions in steady state should converge on empirically observed BCL setting behaviours and associated payment flows. The simulation laboratory construction should also help policy makers better understand the dynamics influencing emergent BCL relationships, the impact of policy decisions on these relationships, and under which conditions these relationships may break down.

Presently, the development of this simulation laboratory is ongoing utilising the Java programming language, the JAS agent-based simulation Java library, and the Brown University Reinforcement Learning and Planning (BURLAP) Java library for implementing MDPs and general classes of stochastic games. The JAS library is used primarily to visualise, define, structure, and schedule simulation environments, events, and time.[30] BURLAP on the other hand is the engine behind the agents' ability to learn and their formulation of update rules in the BCL decision-making stochastic game.[31] Coding is still at an early stage and further research will be required to build

[30] The JAS library and associated documentation are available at http://jaslibrary.sourceforge.net/
[31] The BURLAP is available at http://burlap.cs.brown.edu/

out and calibrate the theoretical framework into the simulation laboratory. Moreover, once complete and results are forthcoming, these results will require empirical validation and the theoretical framework updated accordingly.[32] Finally, though specific to the LVTS, the framework and simulation laboratory could in addition be easily calibrated according to the incentive structures at play within other FMIs.

7 Concluding Remarks

This paper has described a theoretical and computationally tractable framework to assess the self-organising complexity that underpins financial market infrastructures and more specifically the Canadian Large Value Payment System from the perspective of bilateral credit limit setting decisions of the System's Participants. The paper breaks down the various incentives and trade-offs that, at the market microstructure level, drive regularities observed at the macro-level. These incentives were specified along four core categories (liquidity risk, market risk, credit risk, and settlement delay) that impact the survivor pay scheme component of the LVTS. These incentives were mapped to market and credit data as well as payments data; illustrating the relative ease with which market microstructure focused agent-based computational economics models can be empirically grounded whilst capturing the fundamental dynamics of settlement systems and other financial market infrastructures. Moreover, at the time of writing, this was the first paper to truly identify and model the microstructures of the LVTS and the economic decision making of Participants.

Viewed from the market microstructure perspective, a number of fundamental policy considerations have been highlighted. Firstly, and surprisingly, the System Wide Percentage which still is the primary tool regulators and system operators had to control liquidity provision in the LVTS,

[32]For the thorough literature on the empirical validation of agent-based computational economics models, see Iowa State University Professor Leigh Tesfastion's ACE resource repository at http://www2.econ.iastate.edu/tesfatsi/empvalid.htm

may in fact be limited in its effectiveness. This is due to the relative trade-off between the impact of the SWP on the initial cost of liquidity provision and its influence on the intra-day carry cost of collateral. Whilst the opportunity cost of liquidity provision is negatively impacted by increases in the SWP, the minimal impact this has on the intra-day carry cost of liquidity provision implies policy actions through the SWP may only lead to BCL behaviour changes in so far as risk-free rate and other financial market rates permit.

The intra-day carry cost of collateral was also shown to impose an embedded cost to Participants for attempting to free-ride by consistently setting BCLs low relative to other Participants in the LVTS given payment flows. Large intra-day swings in the volatility of the MRC-BRC spread that are a result of intra-day adjustments to BCLs will incur a cost and as such there is little incentive to free-ride in the LVTS through the undercutting of BCLs, given payment flows. Additionally, the theoretical observation that a Participant's ability to garner income in the market for overnight lending is influenced by the liquidity it is able to draw from inbound payments intra-day through its BCL-establishing decisions.

The market microstructure framework also illustrates that, whilst credit risk during times of economic stress may induce LVTS Participants to cut BCLs they extend to others, this is tempered by the cost of settlement delays that such cuts may give rise to. This again highlights the trade-off between minimising exposure to credit risk and settlement delay. Participants with a history of larger payment flows (value or volume) with one another may be more willing to increase their BCLs to one another even with heightened expectations of default amongst themselves.

Finally, it is worth noting that this paper does not speak to collateral setting across the entire LVTS. By just covering Tranche 2 of the LVTS it overlooks the collateral pledge in Tranche 1. This has been done to focus the paper on the market microstructure of Tranche 2 and how Participant incentives influence their decisions to extend bilateral credit limits. While the inclusion of Tranche 1 payments and collateral will impact the results, the described framework does not lose generality. Indeed, the framework can be readily extended to include the incentive structures underpinning Tranche 1 collateral pledging that is more akin to real time gross settlement systems about which there is an extensive body of literature. The inclusion of Tranche 1 brings with it an additional layer of complexity and could enrich the modelling. Further research will be required to build this out, especially as efforts continue to implement the Lynx real-time gross settlement system. In such a setting, an ACE market microstructure approach will help identify the impact of incentives in competing system designs and on system outcomes such as the emergence of cliques, free-riding, payment delays, and collateral pledging.

8 References

1. Abbink, K., Bosman, R., Heijmans, R., and van Winden, F. (2017). Disruptions in Large Value Payment Systems: An Experimental Approach. Internationl Journal of Central Banking, Vol. 13 No. 4

2. Alexandrova-Kabadjova, B, Heijmans, R., Garcia-Ochoa, L., and Serguieva, A. (2016), "Intra-day Liquidity Flows within the Financial Market Infrastructures in Mexico", pp. 191-208, in "Analyzing the Economics of Financial Market Infrastructures" ed: Martin Diehl (Deutsche Bundesbank, Germany), Biliana Alexandrova-Kabadjova (Banco de México, Mexico), Richard Heuver (De Nederlandsche Bank, The Netherlands) and Serafín Martínez-Jaramillo (Banco de México, Mexico)

3. Arjani, N. and McVanel, D. (2006), "A Primer on Canada's Large Value Transfer System". Bank of Canada Working Paper 2006

4. Barde, (2016) "Direct Calibration and Comparison of Agent-Based Herding Models of Financial Markets". Journal of Economic Dynamics and Control 73:329-353;

5. Bech, M. and Garratt, R. (2006). Illiquidity in the Interbank Payment System Following Wide-scale Disruptions. Federal Reserve Bank of New York Staff Reports, 239.

6. Bougheas, S. and Kirman, A (2014). "Complex Financial Networks and Systemic Risk: A Review," Discussion Papers 2014/04, University of Nottingham, Centre for Finance, Credit and Macroeconomics (CFCM).

7. Breuer, T. (2011), A Gödel-Turing Perspective on Quantum States Indistinguishable from Inside.

8. Caldarelli G., Battiston S., Garlaschelli D., and Catanzaro M. (2004) Emergence of Complexity in Financial Networks. In: Ben-Naim E., Frauenfelder H., Toroczkai Z. (eds) Complex Networks. Lecture Notes in Physics, vol 650. Springer, Berlin, Heidelberg

9. Campbell, J. Y., Lo, A. W-C., and MacKinlay, A. C. (1997) "Econometrics of Financial Markets", Princeton, NJ: Princeton University Press. ISBN 978-0691043012

10. Chande, N. and Labelle St-Pierre, N (2016) "Using Speed and Credit Limits to Address the Procyclicality of Initial Margin at Central Counterparties", Bank of Canada Staff Discussion Paper 2016-18. September 2016

11. Chapman, J. Embree, L., Roberts, T., and Zhang, N. (2011) "Payment Networks: A Review of Recent Research", Bank of Canada Review Winter 2010–2011

12. Chapman, J., Garratt, R., and Zhang, N. (2016) "The Bank of Canada's Exposure to Default Risk in Canada's Large Value Transfer System", Mimeo.

13. Chapman, J., and Zhang, N. (2010) "Estimating the Structure of the Payment Network in the LVTS: An Application of Estimating Communities in Network Data"; Bank of Canada Working Paper 2010/13.

14. Chen, S-H., Kampouridis, M. and Tsang, E. (2011), "Microstructure Dynamics and Agent-Based Financial Markets" In: Multi Agent Based Simulation XI, Volume 6532 of the series Lecture Notes in Computer Science pp 121-135.

15. Cona, J. (2008) "A Theory of Market Microstructure Abstractions of Agent Based Models with Insights into Limit Order Markets The Static Theory".

16. Condon, A. (1992). "The complexity of stochastic games". Information and Computation. 96: 203–224.

17. Diehl, M. (2013). "Measuring Free Riding in Large-Value Payment Systems: The Case of TARGET2", Journal of Financial Market Infrastructures, 1(3):31–53.

18. Farmer J. D., and Joshi, S. (2002) "The Price Dynamics of Common Trading Strategies". Journal of Economic Behaviour Organ 49:149-171.

19. Filar, J. and Vrieze, K. (1997). Competitive Markov Decision Processes. Springer-Verlag.

20. Galbiati and Soramäki (2008). "An Agent-Based Model of Payment Systems". Bank of England Working Paper No. 352

21. Gallegati, M., Giulioni, G., Kirman, A., and Palestrini, A., (2011). "What's That Got to do with the Price of Fish? Buyers Behavior on the Ancona Fish Market," Journal of Economic Behavior & Organization, Elsevier, vol. 80(1), 20-33.

22. Gilli M, Winker P (2003) "A Global Optimization Heuristic for Estimating Agent Based Models" Computational Statistics & Data Analysis 42:299-312.

23. Heijmans, R., and R. Heuver (2014). "Is this Bank Ill? The Diagnosis of Doctor TARGET2." Journal of Financial Market Infrastructures 2 (3): 31– 53.

24. Heemeijer, P., Heijmans, R. (2015). Central bank intervention in large value payment systems: An experimental approach. De Nederlandsche Bank Working Paper, 466.

25. Kaliontzoglou, A. and Müller, A. (2016) "Implementation Aspects of Indicators Related to Payments Timing", pp 169-190, in "Analyzing the Economics of Financial Market Infrastructures" ed: Martin Diehl (Deutsche Bundesbank, Germany), Biliana Alexandrova-Kabadjova (Banco de México, Mexico), Richard Heuver (De Nederlandsche Bank, The Netherlands) and Serafín Martínez-Jaramillo (Banco de México, Mexico)

26. Kamhi, N. (2006) "LVTS, the Overnight Market, and Monetary Policy". Bank of Canada Working Paper 2006-15

27. Kirman, A. (1991) "Epidemics of Opinion and Speculative Bubbles in Financial Markets". In: Taylor M (ed) Money and Financial Markets. Blackwell, Oxford, pp 354-368.

28. Kirman, A., (2011). "Learning in Agent-based Models," Eastern Economic Journal, Palgrave Macmillan; Eastern Economic Association, vol. 37(1), 20-27.

29. Kukacka J, and Barunik, J. (2016) "Estimation of Financial Agent-Based Models with Simulated Maximum Likelihood" SSRN:2783663.
30. LeBaron B (2005) "Agent-based Computational Finance" In: Judd, K. L., and Tesfatsion, L. (eds) The Handbook of Computational Economics, Vol. 2. Elsevier, Amsterdam, pp 1187-1233.
31. León, C., Perez, V. J. and Renneboog, L., (2014), "A Multi-Layer Network of the Sovereign Securities Market", Borradores de Economia No. 840, Banco De La Republica, Available at SSRN: https://ssrn.com/abstract=2532461 or http://dx.doi.org/10.2139/ssrn.2532461.
32. Mertens, J. F. & Neyman, A. (1981). "Stochastic Games". International Journal of Game Theory. 10 (2): 53–66.
33. Müller, A. (2016) "Payment System Analysis: The Central Bank Perspective", pp 209-224, in "Analyzing the Economics of Financial Market Infrastructures" ed: Martin Diehl (Deutsche Bundesbank, Germany), Biliana Alexandrova-Kabadjova (Banco de México, Mexico), Richard Heuver (De Nederlandsche Bank, The Netherlands) and Serafín Martínez-Jaramillo (Banco de México, Mexico)
34. Neyman, A. and Sorin, S. (2003). "Stochastic Games and Applications". Dordrecht: Kluwer Academic Press.
35. Platt, D. F. and Gebbie, T. J. (2016a) "Can Agent-Based Models Probe Market Microstructure?" .
36. Platt, D. F. and Gebbie, T. J. (2016b) "The Problem of Calibrating a Simple Agent-Based Model of High-Frequency Trading".
37. Reid, C., (2007) "The Canadian Overnight Market: Recent Evolution and Structural Changes", Bank of Canada Review, Spring 2007.
38. Schwartz, H. M. (2014). "Multi-Agent Machine Learning: A Reinforcement Approach". Wiley.
39. Shapley, L. S. (1953). "Stochastic Games". PNAS. 39 (10): 1095–1100.
40. Serrano, M. A., Krioukov, D., and Boguna, M. (2008) Self-similarity of complex networks and hidden metric spaces, Physical Review Letters 100, 078701, 10.1103/PhysRevLett.100.078701
41. Song, C., Havlin, S., and Makse, H. A. (2005) "Self-Similarity of Complex Networks", Nature 433, 392-395 (27 January 2005) doi:10.1038/nature03248
42. Velupillai, K.V. and Zambelli, S. (2012), "Computability and Algorithmic Complexity in Economics".
43. Yook, S-H., Radicchi, F., and Meyer-Ortmanns, H. (2005) Self-similar Scale-free Networks and Disassortativity,
44. Zhang, Z., Zhou, S., and Zou, T. (2006) Self-similarity, small-world, scale-free scaling, Disassortativity, and robustness in hierarchical lattices, The European Physical Journal B (EPJ B) 56, 259-271, DOI: 10.1140/epjb/e2007-00107-6

9 Appendix

Table 1: List of Symbols

Variable	Description	Type
α	system wide percentage	exogenous
$\beta_{j,i}$	bilateral credit limit Participant j extends to reference Participant i	endogenous decision variable
$b_{j,i}$	Participant j's relative share of the total system loss exposure to a default by Participant i that	endogenous
B_j	Participant j's risk profile in terms of its relative credit exposure to all other Participants in the system	endogenous
η_j	return from collateral portfolio management	exogenous
μ^T	intra-day liquidity shortfall or over-supply	endogenous
$P_{j,i}$	payment flow value from Participant j to Participant i	exogenous
x	collateral funding source asset/liability mix	exogenous
w_x	weight assigned to collateral funding asset/liability x	exogenous
r_x	rate of return on collateral funding asset/liability x	exogenous
τ_i	Participant$i$'s default probability	exogenous
ϕ_i	reflects the recovery rate on multilateral net debit positions Participant i accumulates over a cycle	endogenous
N	the set of all Participants in the system	exogenous
$Q_{i,j}$	the set of payments volume from Participant i to Participant j that are internally queued by i	endogenous
D	the subset of all N agents in the system that extend BCLs to agent i	endogenous
a_j	BCL granting action profile. This is the vector of BCLs Participant j grants to all other Participants in the system	endogenous decision variable
$\mathbb{A}$	the universal set of all possible BCL granting action profiles	exogenous
g	the joint set of BCL granting action profiles of all the agents at a given state or point in time	endogenous
$\mathbb{G}$	universal set of all possible joint BCL granting action profiles	exogenous
s	the current state of the agent	endogenous
$\mathbb{S}$	the universal set of all possible states	endogenous
$\mathbb{T}$	the state transition model	endogenous
$R_a^j(s, s')$	reward Participant j receives for taking an action a in state s that results in the system moving to state s'	endogenous
$V_{a\in g}^j(s)$	Expected value Participant j derives for taking an action a in state s where the system-wide joint action is g	endogenous
γ	discount factor	exogenous
$\mathbb{M}$	strategy matrix of a matrix game	exogenous