Latin American Journal of Central Banking 5 (2024) 100115

Contents lists available at ScienceDirect

Latin American Journal of Central Banking

journal homepage: www.elsevier.com/locate/latcb

A computational model of bilateral credit limits in payment systems and other financial market infrastructures

Oluwasegun Bewaji

Canadian Payments Association (d.b.a Payments Canada), Ottawa, Canada[1]

ARTICLE INFO

JEL classification:
C63
C73
D40
D47
L10

Keywords:
Agent-based computational economics
Market microstructure
High value payment systems
Financial market infrastructures
Bilateral credit limits
Intra-day liquidity management
Stochastic games
Multi-agent reinforcement learning

ABSTRACT

This paper provides the first steps towards a theoretical and structural modelling framework through which optimal decision making in financial market infrastructures such as payments clearing 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 loss sharing arrangements within financial market infrastructures such as the Canadian Large Value Payments System (LVTS). With specific focus on the LVTS, the paper presents a structural model where the payments system represents a market in which bilateral credit limits are the pricing mechanisms for intraday liquidity provisioning and the credit risk arising from the loss sharing arrangement. 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 policy levers such as the System-Wide Percentage, regulatory concerns around procyclicality and free-riding arising from the market microstructure behaviours, and design of the System.

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.[2] That is, such computational 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 intraday basis. Tesfatsion (2022), provides an overview of the axiomatic characterisation of complete agent-based modelling (c-ABM), of which ACE is a specialisation, and evolution of the origins of ACE. Tesfatsion (2022) further points to the philosophical underpinnings of agents and system level dynamics in ACE models as dating back to the (Ryle, 1949) exposition of the acts of learning, remembering, imagining, knowing, or

Linkedin: segunbewaji.
E-mail address: sbewaji@payments.ca.

[1] https://www.payments.ca/

[2] 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.

https://doi.org/10.1016/j.latcb.2023.100115
Received 28 April 2023; Received in revised form 4 October 2023; Accepted 15 November 2023
Available online 19 December 2023

willing as not merely being indicative of some hidden mental processes or complex sequences of intellectual operations, but rather as the means through which those mental processes or intellectual operations are themselves defined. This learning and adaptation need not be in real-time as decision making occurs but may occur in self-referential games or simulations (see Markose, 2005, 2016, 2021 and Harré et al., 2017. Such simulations include but are not limited to the use of cellular automata, reinforcement learning, genetic algorithms and the emergence of large-scale agent-based models. 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.[3]

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 are themselves complex systems that continuously adapt to the multiplicity of environments they operate in.[4]

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 interactions 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 to inform future 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.[5] Indeed, while Chapman et al. (2016), *Mimeo*), do not cover the topic of self-similarity, the authors' results suggest that Canada's Large Value Transfer System (LVTS)[6] 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; Chapman and Zhang, 2010; Chapman et al., 2011), it has further been proven by, among 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.[7]

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 of outcomes within HVPSs.[8] 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) through which counterparty credit risk exposures at default are valued and allocated within the LVTS' loss sharing arrangement. It is important to note at the outset that the LVTS is no longer in operation and has been superseded in September 2021 by the Lynx HVPS in providing wholesale payment services in Canada. However, the BCLs under the LVTS loss sharing remain applicable to other FMIs, including central counterparty (CCP) clearing. This is because BCLs represent an internal valuation of the trade-off between system liquidity, as captured by bilateral and multilateral payment flows, and the counterparty credit risk appetite of system participants in relation to one another. Using BCLs in the LVTS as a pricing mechanism, this paper presents a structural model through which this valuation of the trade-off between system liquidity and counterparty credit risk exposure in loss sharing arrangements can be assessed in a tractable fashion.

Indeed, for the CCP, loss sharing arrangements similar to the LVTS 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 up and why BCLs also help to prevent free-riding in loss sharing arrangements. Moreover, 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.

[3] 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).

[4] Examples of these internal business structures and processes include cash management, legal and compliance, global risk management and, debt issuance pricing arising from creditor perceptions in capital markets such as through credit default swap (CDS) spreads.

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

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

[7] Using a Bayesian implementation of the grid-search algorithm developed by Čopič et al. (2009), Chapman and Zhang (2010) have shown that the LVTS consists of multiple community structures characterised by a small Montreal-based cluster and a big Toronto-based group of the five big Canadian banks.

[8] 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).

This paper thus presents a computational model as a tool to assess liquidity management decisions in clearing and settlement systems taking into consideration the market microstructure, including counterparty credit risk arising from loss sharing arrangements. Mechanisms are developed to evaluate the implementation of internationally harmonised guidelines as established under the 2012 and 2014 CCPMI-IOSOC Principles of Financial Market Infrastructures (PFMIs). In this regard, for a financial authority to be able to assess the extent to which FMIs are in alignment with the PFMIs, a framework able to incorporate the dynamic of the interactions among participants is required. Specifically, the computational model presented here allows for better identification of, or the potential for, market failure and appropriate policy responses. 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; Abbink et al., 2017) of financial institution interactions within financial market infrastructures 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 combines approaches in the literature to offer mechanisms through which FMIs' alignment with the PFMIs can be assessed in a more targeted and structural fashion on the basis of incentives and other dynamics inherent in their microstructures.

The ACE modelling framework permits look-through into the microscopic details underpinning the in-game pay-offs and trade-offs that define the evolution of outcomes in a computationally tractable way that is not immediately possible under the conventional game theoretic construct. Indeed, the ACE approach permits the structural assessment of the decision parameters driving the interactions among heterogeneous agents and their impacts on systemic outcomes within the payments system. The focus of this paper on the incentive structures underpinning decision making of LVTS participants when establishing BCLs and thus pricing the trade-off between credit risk and their intraday liquidity management, is consistent with the Epstein (2001) utilisation of an ACE model to simulate various social scenarios and investigate the impact of social norms such as dyadic efficiency on decision making and thoughtless conformity to entrenched norms among strongly connected agents. Likewise, Nowak et al. (2000) used a two-agent evolutionary ABM to determine the role incentive structures of fairness vs. rationality play in dominating responders' decision making and the emergence of game outcomes. Tesfatsion (2001) extended this to show how ACE models can capture the emergence of interaction networks exhibiting dominant and weaker attractors within labour markets given the underlying market structure. Indeed this paper contributes to the FMI literature by providing a structural model for understanding a potential preference for maximising liquidity over minimising credit risk or vice versa may result in cooperative outcomes across the system or within strongly connected participants. 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 intraday carry cost) that influence behaviour.

Moreover, since the ACE modelling framework presented in this paper is entirely defined by the LVTS rules and FI incentives, the approach is consistent with Axelrod (1984, 1986) who explored the emergence, evolution and propagation of norms such as reciprocity within an environment. Likewise, the ACE model is presented in this paper develops a structural basis upon which the emergence of behavioural norms can be understood in relation to participants in FMIs. It makes no presuppositions about payment flows in the LVTS. Instead, it is flexible enough to ingest whatever the intraday payment flows may be. Based on rules and incentives, it assesses the evolution of BCL decisions across all Participants. This includes the potential emergence of reciprocity in BCL decisions.

Specifically, in accordance with the taxonomy for ACE models outlined in Tesfatsion (2017), this paper employs a model consisting of financial institutions as decision-making agents, and their liquidity provision to the FMI for the timely settlement of payments as the durable goods or financial assets whose quantity and valuation is determined. Furthermore, FMIs may be viewed as a market for the channelling of liquidity in the clearing and settlement of payments, securities, derivatives, or other financial transactions. Moreover, as noted by Diehl (2016), FMIs enable fast and smooth movements, i.e. re-allocation of funds and liquidity across regions, durations, risk levels and currencies. They channel relevant information about payment streams, investment trends, risk perceptions, liquidity status and systemic risk. Consequently, decision-making by agents in FMIs pertaining to the timing of payments, establishment of BCLs or net sender limits (NSLs) reflect the valuation they place on the liquidity they provide to enable clearing and settlement activities within the FMI.[9] For example, in times of economic stress where liquidity is more expensive, FMI participants may choose to minimise the collateral the pledge to the system and rely more heavily on payments they receive from others to settle their payments. Likewise, if they perceive a higher counterparty credit risk than they are willing to accept relative to the liquidity received from bilateral payment flows, they may choose to withhold BCLs causing payments to be delayed or rejected. With respect to NSLs, these may be set sufficiently low as to create gridlock where payments are only sent after receipt of payments from other participants.

The paper further describes how empirical data can be used to calibrate the theoretical model in a tractable manner. Indeed, just as Leon et al. (2014) argue that studying isolated single-layer trading and registering networks yields a misleading perspective on the relations between and risks induced by participating financial institutions'', assessing an FMI's ability to weather shocks in

[9] As of the June 16, 2015 definition by the Committee on Payments and Market Infrastructures (CPMI), a sender limit is a limit, set either by a participant or by the FMI, that is sometimes used as an additional condition for settlement, to restrict credit exposures. It has also become a feature in some FMIs providing continuous intraday finality to control the outflow of settlement funds by setting sender limits. The net sender limit is therefore a limit on the net bilateral flows out of a participant (the sender) to another participant (the receiver).

isolation of the underpinning market microstructure of the FMI may result in regulatory focus on the wrong policy levers or risk mitigation framework.

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. 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 was an electronic wire system that facilitated the transfer of Canadian-dollar payments between Payments Canada's member financial institutions. The LVTS in operation between February of 1999 and September of 2021 was 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 participated 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 characterised 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.[10]

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.[11] 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.[12] 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.[13] 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 intraday 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").[14] 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 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, collateral pledged under Tranche 2 (T2) of the LVTS are a function of the largest single BCLs the FI has extended to all other Participants that day. The Participant's T2 collateral pledge (also known as its Maximum Additional Settlement Obligation or Max ASO) is set at the start of each business day and specifies the largest T2 bilateral net exposure it is willing to accept vis-à-vis another institution. All BCLs, including the largest, can be adjusted upwards at any point throughout the day and the largest BCL counterparty exposure can also be changed as a result of such intraday adjustments. Increasing BCLs intraday where the largest counterparty exposure is unchanged will not result in additional collateral requirements however, the adjustment may impact the composition of counterparty credit risk exposures in the event of a default. An adjustment that increases the largest BCL exposure will involve raising additional collateral intraday from outside the LVTS. Again, it is important to note that unlike NSLs in RTGS systems, BCLs are a double-edged sword in the sense they permit the FI granting the BCL to receive payments up to the limit, but also expose that FI to counterparty credit risk in the event the FI to which the BCL was granted defaults. 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

[10] 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.

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

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

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

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

exchange cycle is internally generated and a function of the BCLs it has received bilaterally and multilaterally subject to the SWP and the intraday 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, if a default occurs, the defaulting Participant may not have provided enough collateral to the LVTS to cover its negative position at the time of default. The residual between the collateral pledged and outstanding obligations at the time of default is known as an Own Collateral Shortfall (OCS). This OCS is allocated to surviving Participants as part of the counterparty credit risk loss sharing based on their relative share of BCLs extended to the defaulting FI during the cycle.[15] The surviving Participants' exposure to the OCS is nevertheless capped at their respective Max ASOs. OCS exposures in excess of the 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 by executing risk control tests on each payment to assess and control the extent to which LVTS can become negative relative to other Participants and tie all multilateral net debit positions to the amount of collateral available to settle payments; thus enforcing the Net Debit Cap for the Tranche. For a T2 payment to pass the applicable multilateral Net Debit Cap risk control test, it must satisfy both the Bilateral Risk Control (BRC) and Multilateral Risk Control (MRC). These risk controls link the payment to the counterparty credit risk tolerance of Participants and thus liquidity headroom available to send the payment within T2 of the LVTS at a bilateral (BRC) and multilateral (MRC) level. More specifically, to ensure that there is enough liquidity within T2 to execute the payment and therefore not breach the multilateral T2NDC, the risk control tests verify that the payment is (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,

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

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

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

where N represents the set of all system Participants, $P_{i,j}^{T}$ 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, $\sum_{t=0}^{T-1} \left(P_{i,j}^{t} - P_{j,i}^{t} \right)$, 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 $100 m 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 ($\mathbb{P} = \left\{ P_{i,j}^{t} \middle| BRC_{i,j}^{t} = P_{i,j}^{t} = MRC_{i,j}^{t} \right\}$) 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 ($\mathbb{P}^{\mathrm{BRC}} = \left\{ P_{i,j}^{t} \middle| BRC_{i,j}^{t} < P_{i,j}^{t} \leq MRC_{i,j}^{t} \right\}$) 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 ($\mathbb{P}^{\mathrm{MRC}} = \left\{ P_{i,j}^{t} \middle| MRC_{i,j}^{t} < P_{i,j}^{t} \leq BRC_{i,j}^{t} \right\}$) 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.

It is important to reiterate that even though overall counterparty credit risk exposure is capped at the Max ASO, by linking the counterparty credit risk exposure in the loss sharing scheme to the BCLs extended, a surviving financial institution may still have

[15] 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.

the largest credit risk exposure to a defaulting institution if the BCLs it extended to the defaulting FI exceeds those extended by all other surviving FIs. That is, even if bank A's largest BCL exposure, say $25, is to bank B and it extended $15 to a defaulting bank E, if all other banks extended $10 or less in BCLs to bank E, then bank A would still hold the largest counterparty risk exposure to bank E of all the surviving banks. Consequently, a Participant's decision to extend BCLs is not only an assessment of expected bilateral payment flows and counterparty credit risk appetite towards the Participant to whom it granted the largest BCLs, but also the implied pricing of intraday liquidity and counterparty credit risk relative to all other Participants in the System. The nature of this loss sharing arrangement thus highlights the underlying complexity of the LVTS. That is, BCLs not only underpin the collateralisation of bilateral and multilateral payments flows under T2, but also form the basis upon which Participants price counterparty credit risk across all their bilateral and multilateral exposures in the System. The nature of the loss sharing arrangement, which is a function of the BCLs extended to each potential default Participant within the LVTS, also implies that BCLs are extended by each Participant subject to its counterparty credit risk exposure appetite towards each of the other Participants not merely its largest exposure. The objective of each Participant is therefore to set the optimal level of BCLs it extends to all other Participants at the start of the day to maximise payment flows while minimising counterparty credit risk exposures to each of its counterparties. This optimal allocation of BCLs is set to ensure that the Participant can receive and send payments from and to other Participants over the course of the day without having to adjust these BCLs at cost on an intraday basis.

As with other payment systems, given all intraday 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 50 bps 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 50 bps band (Reid, 2007). 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 ± 50 bps 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, intraday 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.[16]

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 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 intraday) 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.[17]

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

[16] 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.

[17] The significance of this assumption will be expanded under the discussion on the intraday cost of carry later in this section.

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.

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 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).[18] 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_x) and high quality liquid asset mix (w_x) are assumed to be exogenous and the initial cost of liquidity specified as

$$\lambda\left(\beta_{j,i}\right) = \alpha\beta_{j,i}\left[\prod_{x\in X}\left(1 + w_x\eta_x\right)\right] \tag{3}$$

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.[19]

3.2. Intraday carry cost of collateral

The carry cost represents the intraday cost of managing, reallocating or adjusting BCLs. That is, the implied financing cost of increasing BCLs to another Participant intraday or is the implied foregone intraday investment gains arising from the over extension of BCLs to a given Participant.[20] The functional form for this intraday 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 intraday spread or liquidity headroom between the MRC (Eq. (2)) and the BRC (Eq. (1)). Consequently, $\sigma\left(\mu^T\right)$ represents the volatility in the MRC-BRC spread over the course of a day or some predefined intraday window. Note that since the MRC and BRC are executed against each successive payment to ensure there is no breach of the multilateral T2NDC tests, the value of the liquidity headroom under both MRC and BRC are time varying and based on payment flows and BCLs extended at the time each new payment message is submitted to the LVTS. Thus, $\sigma\left(\mu^T\right)$ can be measured as the standard deviation in the spread between the BRC and MRC. This volatility should be fairly stable with time only changing for structural changes in relationships between each pair of Participants. Spikes in the volatility beyond the normal would indicate the existence of intraday liquidity shortages or over-supply. While the latter may represent an over collateralisation of the System, an intraday liquidity shortage may result in a payment failing either risk control and being rejected, delayed or requiring further liquidity injection by raising BCLs and therefore collateral intraday.

---

[18] Stipulations as to asset-class and maturity margin requirements are maintained by the Bank of Canada and the Bank retains the right to apply additional margins for securities without a reliable market price, especially in times of market instability. As of the time if writing this paper, these SLF margin requirements can be found at the URL: https://www.bankofcanada.ca/2019/07/assets-eligible-collateral-standing-liquidity-facility-290719/.

[19] 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.

[20] 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.

Since the carry cost is an implied cost which accounts for the cost of intraday liquidity injections in the system, a Participant that efficiently manages its intraday liquidity will in theory set BCLs with reference to the bilateral and multilateral risk controls. In other words, Participant j will look to minimise the intraday 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 intraday 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 one assumes that bilateral payment flows follow a standard normal distribution (i.e. $\left(P_{j,i} - P_{i,j}\right) \sim N(0,1)$) then, a Participant j with a consistently large volatility in the intraday 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 intraday pressure to increase (or decrease in the following cycle) the BCLs it extends to the reference entity. Conversely, low volatility in the intraday 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 intraday 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 high quality liquid asset mix accepted as collateral by the central bank and the intraday cost of raising the high quality liquid assets (w_x and $r_x$respectively).[21]For example, a financial institution may be less willing to have high quality liquid assets such as Government of Canada guaranteed – stripped coupons and residuals encumbered as collateral in the payment system, particularly in times of stress, if it deems the demand for those assets to be sufficiently high in other competing areas such as intraday derivatives margining purposes, balance sheet management against the liquidity coverage ratio and net stable funding ratio under Basel III or securities lending transactions. Thus, any collateral apportionment to the payment system will come at the cost of intraday activities in these other business areas in which the collateral could be used. Moreover, the encumbrance of assets pledge and allocated for intraday liquidity to support payments may not be consistently applied across all institutions meaning the relative cost of encumbrance will differ from one institution to the next. The intraday carry cost of collateral therefore reflects the market risk to the individual Participant associated with the intraday management of encumbered collateral often overlooked in the literature but this collateral optimisation problem is something that larger banks, at least, employ internal groups to solve on a daily basis. Essentially, even in circumstances where there exist a glut in collateral banks will still assess which assets are used to support what activity while staying within regulatory requirements.[22]

It should be noted that the intraday carry cost of collateral is also subject to central bank, government, other regulatory, and central counterparty actions that result in an obscuring or influencing of the market pricing of assets and therefore can drive BCL and wider collateral decisions. The coalescence or divergence between what central banks and central counterparties deem as eligible collateral will impact demand for these assets and thus the intraday reallocation costs. Changes to the types of assets that are considered eligible collateral will influence the intraday carry costs. For instance, the Bank of Canada's initiative to allow LVTS Participants to pledge their non-mortgage loan portfolio (NMLP) 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 intraday carry cost of collateral in the System by reducing pressure on demand for other higher quality asset classes in the HQLA eligibility pool.

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 intraday 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 deviation of the MRC-BRC spread around its mean unchanged.

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

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

[21] For the sake of simplicity, these high quality asset mixes and expected returns, though constrained to the pool of asset classes the central bank accepts as collateral, are assumed to be exogenous. This assumption may be relaxed in further research once the initial model is fully calibrated. Indeed, given that these asset mixes must be made up of asset classes permitted as collateral by the central bank, the impact of changing costs of raising these assets intraday especially in times of financial market stress will be of interest to understand.

[22] A further point of note is that by internalising through the intraday carry cost of established BCL positions, LVTS Participants are faced with an endogenous charge on any incentive to engage in free-riding by undercutting of BCLs. Attempting to undercut BCL extensions potentially leave the Participant subject to having to make intraday adjustments to its BCLs that could require intraday collateral injections into the LVTS and therefore raise an exposure to intraday market volatility.

where $b_{j,i} = \left[\frac{\beta_{j,i}}{\sum_{i \neq k \in D \subseteq N} \beta_{k,i}} \right] \in B_j$. 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 (i.e. Participant j's Max ASO). 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.[23]

The term $E\left[\sum_{i \neq k \in N} \sum_{t=0}^{T-1} \left(P_{k,i} - P_{i,k}\right)\right]$ 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 $\tau_i \in Y_j = \{\tau_0, \tau_i, \ldots, \tau_{N-1}\}$ and $(1 - \phi_i)$ where $\phi_i \in \Phi_j = \{\phi_0, \phi_i, \ldots, \phi_{N-1}\}$ for $i \neq j \in N$, 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). The vectors Y_j and Φ_j represent the information Participant j possesses with respect to its perceptions of the counterparty credit risk it faces from all other LVTS Participants at the time it makes its BCL extension decisions. It can therefore be seen that the vector of bilateral credit limits B_j represents Participant j's counterparty risk exposure preferences across the entire System.

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 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.[24] In other words, even though Participant j's collateral pledge is tied to its largest BCL, Participant j's still needs to manage its counterparty credit risk exposure to any Participant i since it may not wish to hold the largest relative exposure to Participant i in the event of default if Participant i was not its largest BCL. Participant j will therefore extend BCLs to all other Participant i that are not its largest BCL exposure so as to ensure bilateral payment flows between itself and those Participants while avoiding a large exposure to the default of those Participants relative to the all other BCLs granted to the defaulting Participant i.

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.[25]

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

[23] 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.

[24] The reader should be aware that the counterparty credit risk component can be extended by noting that, as currently defined, $\psi_j(i)$, is the contingent leg of a credit default swap (CDS) where, the premium leg is the intraday 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.

[25] 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.

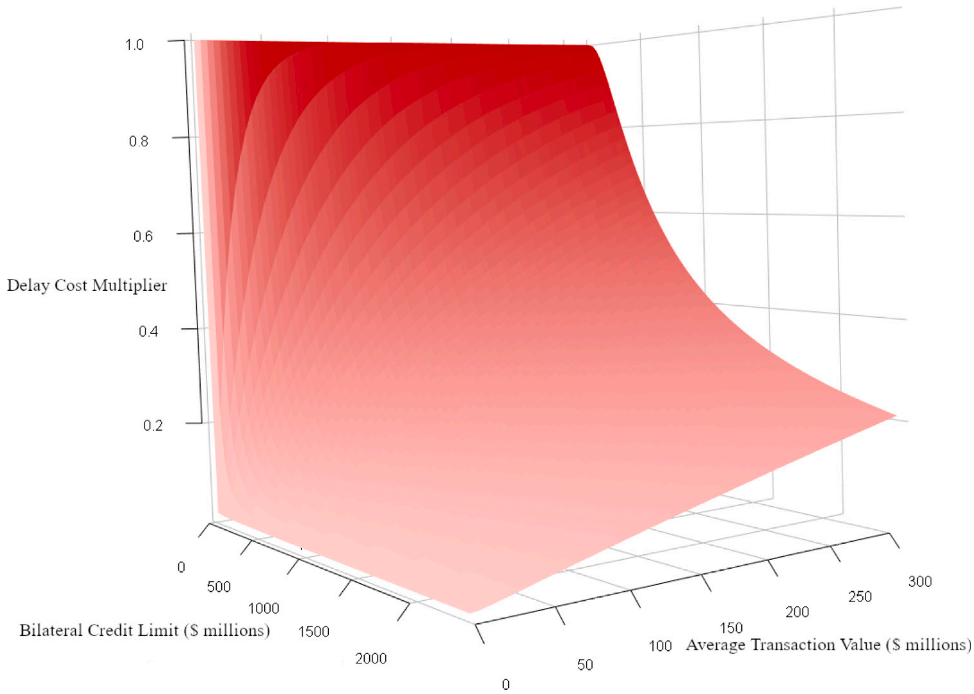

Fig. 1. Inbound payment delay cost.

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

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

Fig. 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.[26] 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 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, and though seemingly counter intuitive, the delay cost penalty acts as a mechanism counteracting Participant j's proclivity not to extend BLCs as a result of the counterparty credit risk exposures it faces under the survivor pay credit risk incentive function introduced in Section 3.3 above.

[26] 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.

To understand this, one should remember that under the BCL mechanism, a Participant extends limits to all other Participants and collateral is apportionment as a fraction of the largest BCL it extended. Therefore, unlike NSLs in real-time gross settlement (RTGS) systems that limit the potential for FIs to free-ride and become liquidity sinks, BCLs manage a trade-off between liquidity provisioning to the CCP and the counterparty credit risk exposure FIs are willing to assume in relation to all other FIs from loss sharing. Therefore, while NSLs can constrain system liquidity due to the potential for free-riding, BCLs expand liquidity subject to counterparty credit risk appetite. Under economic conditions with tightening liquidity such as the 2008–2009 Global Financial Crisis (GFC), NSLs can serve to amplify system liquidity constraints. By contrast, given that BCLs are liquidity generating subject to expectations of counterparty default, they may serve to keep system liquidity high even in times of tighter liquidity conditions. Indeed, while the Bank of Canada was proactive in providing support to the Canadian financial markets during the GFC, for only a brief period at the end of 2008 was there any take-up of the central bank's 1-month term loans to LVTS Participants. Moreover, this up-take peaked in early December 2008, at just over $4 billion, but gradually subsided to zero after the end of 2008; thus the LVTS proved resilient to the GFC (see Zorn et al., 2009 and Chande and St-Pierre, 2016). Bewaji (2019), further illustrated why BCLs could be economically beneficial in moments of economic stress by numerically pricing the credit default swap implied in the use of BCLs and the liquidity they generate. The research showed that at the height of the GFC, the implied credit default swap remained in-the-money for entities with corporate bond ratings of BBB+ and above. It was therefore beneficial for LVTS Participants to assume the credit risk as a trade-off for the liquidity.

Moreover, 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. Indeed, this is a result Chapman and Zhang (2010) empirically highlight as they find well coordinated Toronto and Montreal based cliques each consisting of at most 5 banks. Bech et al. (2008), also find BCLs to be aligned with the likely liquidity flows and speed at which payments are processed by the banks. In other words, there exists evidence of geographical coordination between LVTS Participants that is reflected in liquidity flows and the speed of those flows. It is also worth noting that while the LVTS is more of a real time net settlement system as opposed to a real time gross settlement system, similar intraday liquidity driven coordination is present in real time gross settlement systems that employ net sender limits despite such systems having no credit risk. Indeed, Becher et al. (2008) in assessing the Clearing House Automated Payment System (CHAPS) sought to empirically disentangle the effects of centralised throughput guidelines and the use of bilateral net sender limits in enforcing coordination on a particular time of the payment day that overcome the tendency for financial institutions to delay payments until the end of the day. Ledrut (2007) looking at the impact of the operational failure of participants within the Dutch interbank payment system also finds that payment system participants will control their exposures to an operational failure stricken bank using internal bilateral limits based on the degree to which there is reciprocity in the value of bilateral payments with the stricken bank.

3.5. Overnight market earnings

Since it is assumed that payments are stochastic, end of cycle multilateral net debit positions 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 Eq. (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, Eq. (1) can be substituted into Eq. (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 intraday 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 of shared loss exposure to assume and BCLs to extend to others in the LVTS.[27]

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.[28] 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[29] 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 $\mathbb{T}\left(s, g, s'\right)$ with probability $p\left(s' \mid a \in g,\ s\right)$.

Combining the incentive functions, an LVTS Participant's single stage reward function for establishing BCLs in relation to other Participants is specified as

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

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.[30] Thus Participant j's expected reward, for action $a_j \in g$ in state s is given by the Q-function

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

or

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

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

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

or

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

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 Eqs. (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:

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

[27] Whilst it is operationally useful to distinguish between "Standing" vs "intraday" 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 intraday changes in BCL extension. Where such intraday BCL changes occur, these should be to address unforeseen extraordinary payments or events.

[28] See Bäuerle 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).

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

[30] 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.

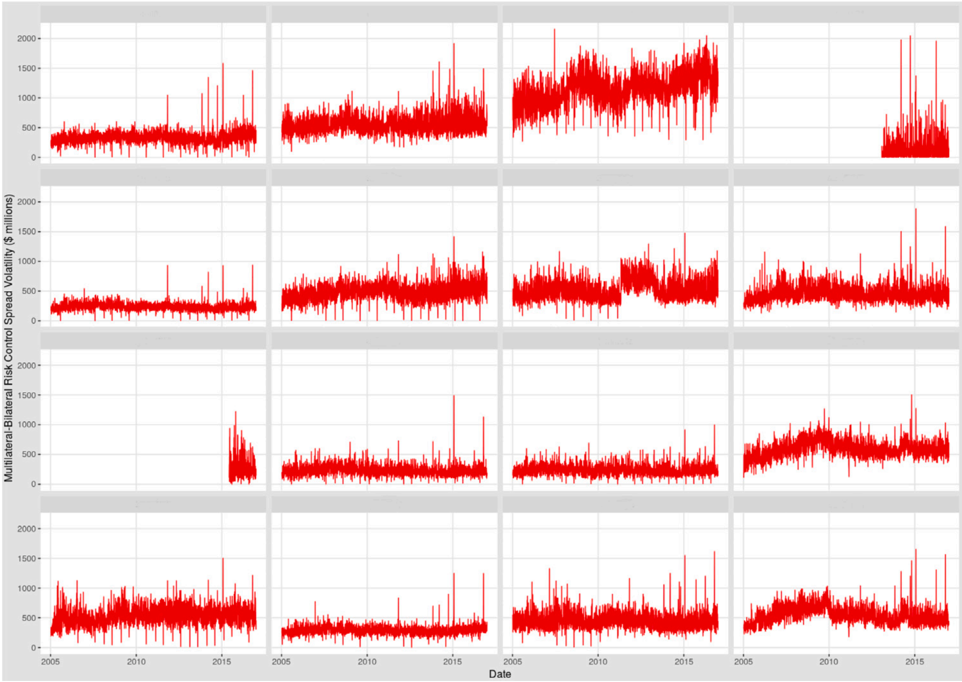

Fig. 2. Example intraday spreads between multilateral and bilateral risk controls.
The figure plots the daily volatility in intraday spread between the multilateral and bilateral risk controls (MRC-BRC spread) associated with payments sent by a single LVTS Participant (the reference Participant *i*) 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.

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 Fig. 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 generally 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%.

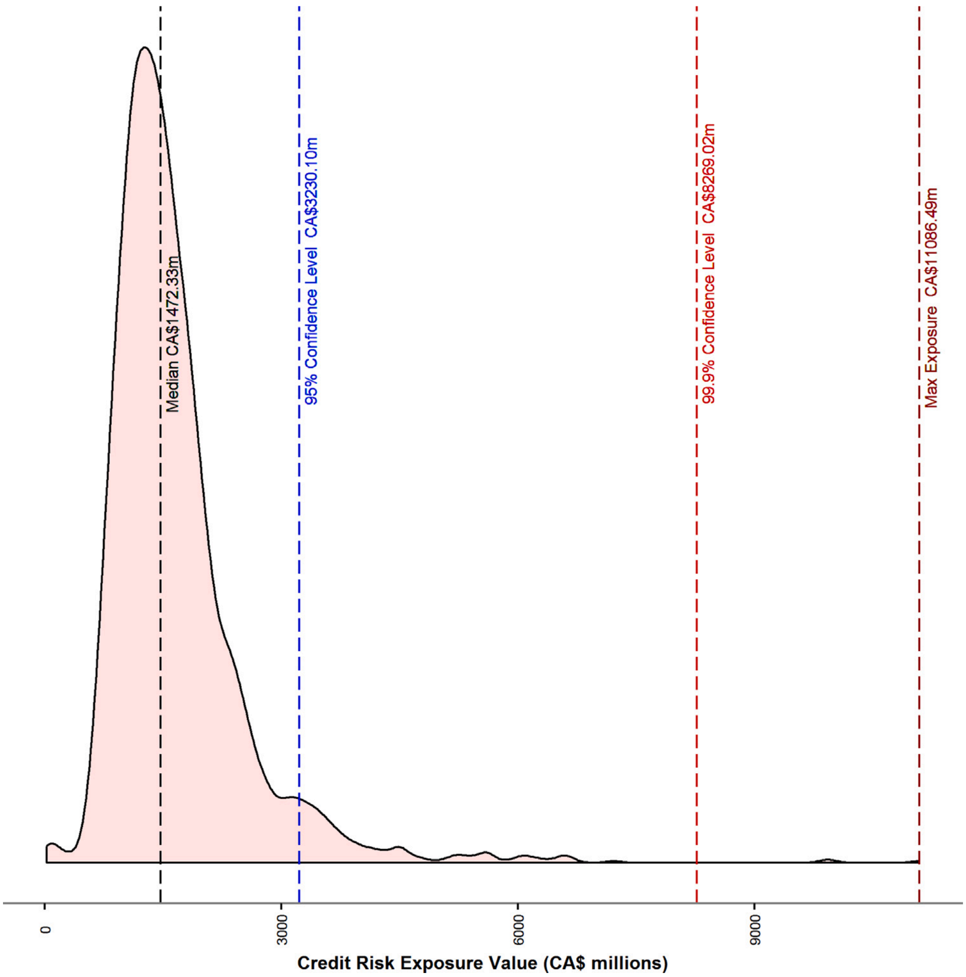

Fig. 3. LVTS simulated multiple default total Maximum Additional Settlement Obligations (Max ASO).
The depicted total Max ASOs present extreme levels of potential counter party risk exposures within the LVTS between 2005 and 2016. The exposures are generated from the simulated default of multiple LVTS Participants based on their multilateral net debit positions over rolling 5 min windows. These Max ASOs should be viewed as extreme levels of counter party credit risk exposure because they are calculated against the time payment messages were sent to the LVTS as opposed to when those payments actually settled. As such, these exposures also include payments that were rejected or settled through multilateral offsetting within the LVTS Jumbo Queue.

In addition to the default probabilities, the parametrisation of credit exposures can be mapped to empirical total net additional settlement obligations in the event of multiple defaults. As illustrated in Fig. 3, the distribution of surviving Participants' Max ASO at the extremes based on the time payments are sent as opposed to when they settled in the event of multiple simulated LVTS defaults between January 2005 and December 2016 ranged between $0 and $11bn with a median value of $1.47bn and 99.9% of multiple defaults resulting in ASOs under $8.27bn. While at the extremes, these Max ASO exposures appear high, they must be taken in contrasts to the shared collateral pool arising from the BCL implementation within the LVTS that on average represented between $7–8bn in coverage against these exposures between 2005 and 2016. Indeed the system collateral generated under the BCLs are seen through the simulated defaults to cover 99% of all multiple default scenarios. These multiple defaults are simulated over rolling 5 min windows by assuming that all Participants with multilateral net debit positions below zero default and surviving Participants' Max ASO exposures subsequently calculated for the 5 min window under consideration. The Max ASO exposures from the simulated multiple defaults do not control for which Participants default or how many default. Consequently, only a single default may have occurred within one 5-minute window while 7 may have occurred in the subsequent 5 min and those 7 defaulting Participants could include the single defaulting Participant from the previous 5-minute window.

It worth noting that by using the time the payments are sent as opposed to the actual settlement time, these multiple default Max ASO exposures represent an over estimation of counterparty credit risk exposures. By overlooking the settlement time of the payments, the simulated multiple default scenarios do not parse out payments that were rejected by the LVTS for failing the multilateral and bilateral risk controls or conditions required to be accepted into the Jumbo Queue. The use of the time payment messages are sent rather than their settlement time also does not account for those payments that would have been forced to

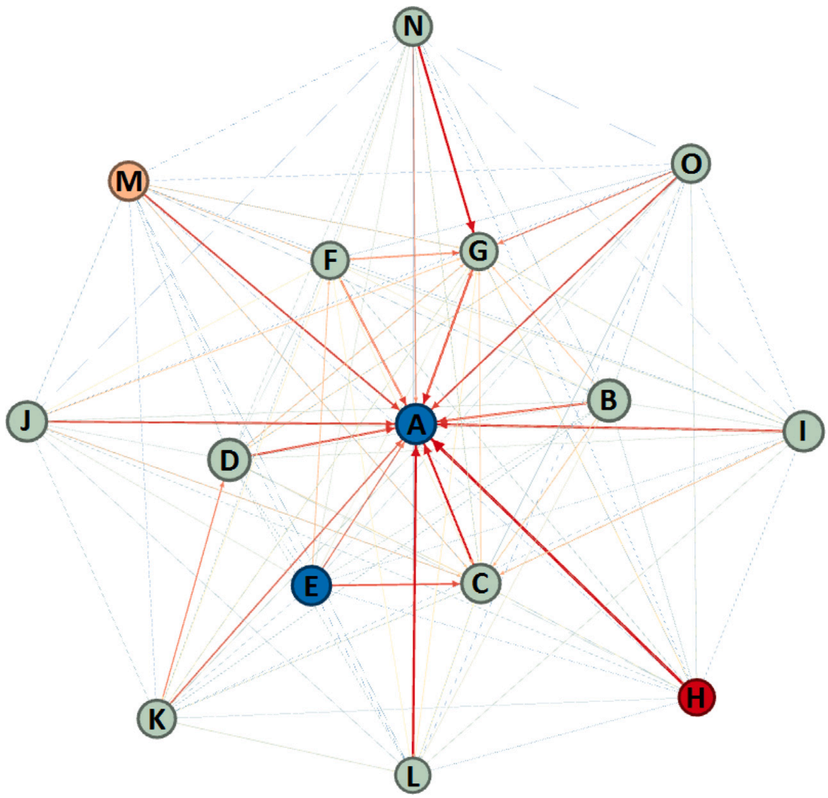

Fig. 4. LVTS participant payment flow relationships.
The depicted Force Atlas network layout is derived from sender to receiver payment flows by value between 2005 and 2016.

settle using multilateral offsetting under the LVTS' Jumbo Queue or released from the Jumbo Queue once the payment was able to satisfy both the multilateral and bilateral risk controls.[31] The simulated counterparty credit risk exposures depicted in Figure 3 can therefore be viewed as extreme outcomes as the actual exposures once the Jumbo Queue and multilateral and bilateral risk controls are accounted for would be significantly lower.

The extent to which payment flow relationships may influence delay cost modelling can be illustrated through the data. The network diagram in Fig. 4, depicts the relationship between financial institutions in the LVTS with respect to payment flows. Under the force atlas layout, nodes at the centre 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 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.[32] 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 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 (Fig. 5). 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

[31] The non-accounting for Jumbo Queue and rejected payments in the simulated defaults inevitably result in double counting of payments when generating the simulated Total Max ASOs. For the current purposes this double counting is acceptable as the objective is to identify a range of potential extremes rather than actual feasible exposures given risk controls.

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

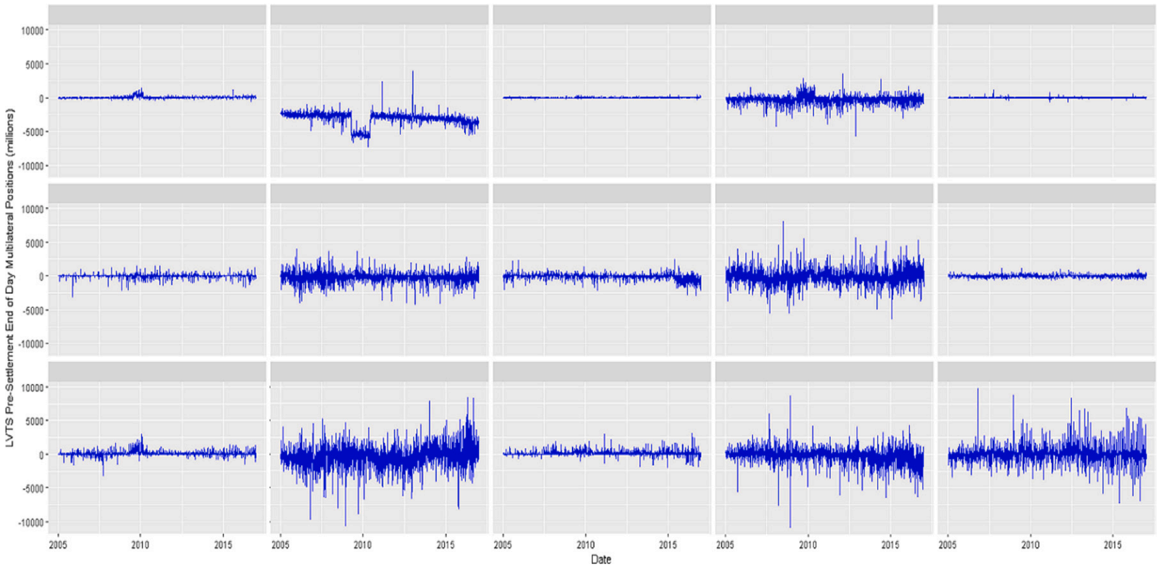

Fig. 5. LVTS pre-settlement end of day multilateral positions (millions).
The figure plots the daily end of cycle multilateral net debit positions of all Participants in the LVTS 2005 and 2016. Each series represents a single Participant's end of cycle position.

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.[33] 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.[34] 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.

6. 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, 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 intraday 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 intraday 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 intraday 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 intraday swings in the volatility

33 See (Arjani and McVanel, 2006; Kamhi, 2006) for an overview of cash-settlement and monetary policy actions in the LVTS.

34 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.

of the MRC-BRC spread that are a result of intraday 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 intraday 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 among 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 and the loss sharing arrangement within the LVTS. While the inclusion of Tranche 1 payments and collateral will impact the results, the described framework does not lose generality. Indeed, now that the LVTS has been replaced by the Lynx RTGS system, a similar incentives driven framework of liquidity management can be readily extended Lynx. Lynx brings with it additional layers of complexity that include the utilisation of liquidity savings mechanisms and net send limits. In fact, rather than BCLs being used as the decision or pricing parameter, in an RTGS such as Lynx, this could be replaced by NSLs. This would enable the comparison of decision making and systemic outcomes in Lynx and LVTS under certain scenarios that could help policymakers further identify areas of vulnerability and opportunities to enhance Lynx. Further research will be required to build this out, especially as FIs continue to familiarise themselves with Lynx and as they recover from the high liquidity environment of quantitative easing during the 2020 to 2023 COVID-19 pandemic and pandemic recovery. 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.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix

See Table 1.

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	intraday 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

(*continued on next page*)

Table 1 (*continued*).

Variable	Description	Type
$\mathbb{G}$	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

References

Abbink, K., Bosman, R., Heijmans, R., van Winden, F., 2017. Disruptions in large value payment systems: An experimental approach. Int. J. Central Bank. 13 (4).

Alexandrova-Kabadjova, B., Heijmans, R., Garcia-Ochoa, L., Serguieva, A., 2016. Intraday liquidity flows within the financial market infrastructures in Mexico. In: Martin Diehl (Deutsche Bundesbank, Germany), Biliana Alexandrova-Kabadjova (Banco de México, Mexico), Richard Heuver (De Nederlandsche Bank, The Netherlands), Serafín Martínez-Jaramillo (Banco de México, Mexico) (Eds.), Analyzing the Economics of Financial Market Infrastructures. pp. 191–208.

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

Axelrod, R., 1984. The Evolution of Cooperation. Basic Books Inc., New York, NY.

Axelrod, R., 1986. An evolutionary approach to norms. Am. Political Sci. Rev. 80, 1095–1111.

Barabasi, AL., Albert, R., 1999. Emergence of scaling in random networks. Science 1999 286 (4), 509–512. http://dx.doi.org/10.1126/science.286.5439.5.

Barde, 2016. Direct calibration and comparison of agent-based herding models of financial markets. J. Econom. Dynam. Control 73, 329–353.

Bäuerle, N., Rieder, U., 2011. Markov Decision Processes with applications to finance. Springer-Verlag, Heidelberg.

Bech, M., Chapman, J., Garratt, R., 2008. Which Bank is the Central Bank? an Application of Markov Theory to the Canadian Large Value Transfer System. Bank of Canada Working Paper 2008/42.

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

Becher, C., Galbiati, M., Tudela, M., 2008. The Timing and Funding of CHAPS Sterling Payments. Federal Reserve Bank of New York, pp. 113–133.

Bewaji, O., 2019. Procyclicality and risk-based access: valuing the embedded credit default swap of employing bilateral credit limits in financial market infrastructures. J. Financial Mark. Infrastruct. 7 (3), 45–72. http://dx.doi.org/10.21314/jfmi.2018.109.

Bougheas, S., 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).

Breuer, T., 2011. A Gödel-Turing perspective on quantum states indistinguishable from inside.

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

Campbell, J.Y., Lo, A.W.-C., MacKinlay, A.C., 1997. Econometrics of Financial Markets. Princeton University Press, Princeton, NJ, ISBN: 978-0691043012; Chande, N., 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.

Chande, N., St-Pierre, N. L., 2016. Using Speed and Credit Limits to Address the Procyclicality of Initial Margin at Central Counterparties. Bank of Canada Staff discussion paper, 1914-0568; 2016/18.

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

Chapman, J., Garratt, R., Zhang, N., 2016. The Bank of Canada's Exposure To Default Risk in Canada's Large Value Transfer System. Mimeo.

Chapman, J., 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.

Chen, S.-H., Kampouridis, M., Tsang, E., 2011. Microstructure dynamics and agent-based financial markets. In: Multi Agent Based Simulation XI. In: Lecture Notes in Computer Science, Vol. 6532, pp. 121–135.

Cona, J., 2008. A theory of market microstructure abstractions of agent based models with insights into limit order markets the static theory.

Condon, A., 1992. The complexity of stochastic games. Inform. and Comput. 96, 203–224.

Čopič, J., Jackson, M., Kirman, A, 2009. Identifying community structures from network data via maximum likelihood methods. The B.E. Journal of Theoretical Economics 2009 9 (1), 1–40. http://dx.doi.org/10.2202/1935-1704.1523.

Diehl, M., 2013. Measuring free riding in large-value payment systems: The case of TARGET2. J. Financial Mark. Infrastruct. 1 (3), 31–53.

Diehl, M., 2016. Financial market infrastructures: The backbone of financial systems. In: Analyzing the Economics of Financial Market Infrastructures. pp. 1–19. http://dx.doi.org/10.4018/978-1-4666-8745-5.ch001.

Epstein, J.M., 2001. Learning to be thoughtless: Social norms and individual competition. Comput. Econ. 18, 9–24.

Farmer, J.D., Joshi, S., 2002. The price dynamics of common trading strategies. J. Econ. Behav. Organ 49, 149–171.

Filar, J., Vrieze, K., 1997. Competitive Markov Decision Processes. Springer-Verlag.

Galbiati, M., Soramäki, K., 2008. An Agent-Based Model of Payment Systems. Bank of England Working Paper No. 352.

Gallegati, M., Giulioni, G., Kirman, A., Palestrini, A., 2011. What's that got to do with the price of fish? Buyers behavior on the Ancona Fish Market. J. Econ. Behav. Organ. 80 (1), 20–33, Elsevier.

Gilli, M., Winker, P., 2003. A global optimization heuristic for estimating agent based models. Comput. Statist. Data Anal. 42, 299–312.

Harré, M., Lizier, J.T., Boschetti, F., 2017. Self-referential basis of undecidable dynamics: from the liar paradox and the halting problem to the edge of chaos. Phys. Life Rev. 31 (20), http://dx.doi.org/10.1016/j.plrev.2018.12.003.

Heemeijer, P., Heijmans, R., 2015. Central Bank Intervention in Large Value Payment Systems: An Experimental Approach. De Nederlandsche Bank Working Paper, 466.

Heijmans, R., Heuver, R., 2014. Is this bank ill? The diagnosis of doctor TARGET2. J. Financial Mark. Infrastruct. 2 (3), 31–53.

Hu, Q., Yue, W., 2008. Markov Decision Processes with Their Applications - Advances in Mechanics and Mathematics. Springer-Verlag, New York.

Kaliontzoglou, A., Müller, A., 2016. Implementation aspects of indicators related to payments timing. In: Martin Diehl (Deutsche Bundesbank, Germany), Biliana Alexandrova-Kabadjova (Banco de México, Mexico), Richard Heuver (De Nederlandsche Bank, The Netherlands), Serafín Martínez-Jaramillo (Banco de México, Mexico) (Eds.), Analyzing the Economics of Financial Market Infrastructures. pp. 169–190.

Kamhi, N., 2006. LVTS, the Overnight Market, and Monetary Policy. Bank of Canada Working Paper 2006-15.
Kirman, A., 1991. Epidemics of opinion and speculative bubbles in financial markets. In: M., Taylor (Ed.), Money and Financial Markets. Blackwell, Oxford, pp. 354–368.
Kirman, A., 2011. Learning in agent-based models. East. Econ. J. 37 (1), 20–27, Palgrave Macmillan; Eastern Economic Association.
Kukacka, J., Barunik, J., 2016. Estimation of financial agent-based models with simulated maximum likelihood. SSRN:2783663.
LeBaron, B., 2005. Agent-based computational finance. In: Judd, K.L., Tesfatsion, L. (Eds.), The Handbook of Computational Economics, Vol. 2. Elsevier, Amsterdam, pp. 1187–1233.
Ledrut, E., 2007. Simulating Retaliation in Payment Systems: Can Banks Control their Exposure to a Failing Participant?. DNB Working Paper, No. 133/2007.
Leon, C., Perez, V.J., 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.
Markose, S., 2005. Computability and evolutionary complexity: Markets as complex adaptive systems (CAS). Econ. J. 115 (504), F159–F192. http://dx.doi.org/10.1111/j.1468-0297.2005.01000.x.
Markose, S., 2016. The gödelian foundations of self-reference, the liar and incompleteness: Arms racein complex strategic innovation. In: Pinto, Alberto A., Gamba, Elvio Accinelli, Yannacopoulos, Athanasios N., Hervés-Beloso, Carlos (Eds.), Trends in Mathematical Economics Dialogues Between Southern Europe and Latin America. Springer, ISBN: 978-3-319-32541-5, http://dx.doi.org/10.1007/978-3-319-32543-9.
Markose, S., 2021. Genomic intelligence as über bio-cybersecurity: The gödel sentence in immuno-cognitive systems. Entropy 2021 23 (4), 405. http://dx.doi.org/10.3390/e23040405.
McPhail, K., Vakos, A, 2003. Excess collateral in the lvts: how much is too much?. Bank of Canada Working paper 1192-54342003-36.
Mertens, J.F., Neyman, A., 1981. Stochastic games. Internat. J. Game Theory 10 (2), 53–66.
Müller, A., 2016. Payment system analysis: The central bank perspective. In: Martin Diehl (Deutsche Bundesbank, Germany), Biliana Alexandrova-Kabadjova (Banco de México, Mexico), Richard Heuver (De Nederlandsche Bank, The Netherlands), Serafín Martínez-Jaramillo (Banco de México, Mexico) (Eds.), Analyzing the Economics of Financial Market Infrastructures. pp. 209–224.
Neyman, A., Sorin, S., 2003. Stochastic Games and Applications. Kluwer Academic Press, Dordrecht.
Nowak, M.A., Page, K.M., Karl, S., 2000. Fairness versus reason in the ultimatum game. Science 289, 1773–1775.
Platt, D.F., Gebbie, T.J., 2016a. Can agent-based models probe market microstructure?
Platt, D.F., Gebbie, T.J., 2016b. The problem of calibrating a simple agent-based model of high-frequency trading.
Reid, C., 2007. The Canadian Overnight Market: Recent Evolution and Structural Changes. Bank of Canada Review.
Ryle, G., 1949. The concept of mind. In: Hutchinsons University Library. Hutchinsons House, London, UK, https://archive.org/details/conceptofmind032022mbp.
Schwartz, H.M., 2014. Multi-Agent Machine Learning: A Reinforcement Approach. Wiley.
Serrano, M.A., Krioukov, D., Boguna, M., 2008. Self-similarity of complex networks and hidden metric spaces. Phys. Rev. Lett. 100, 078701. http://dx.doi.org/10.1103/PhysRevLett.100.078701.
Shapley, L.S., 1953. Stochastic games. Proc. Natl. Acad. Sci. USA 39 (10), 1095–1100.
Song, C., Havlin, S., Makse, H.A., 2005. Self-similarity of complex networks. Nature 433, 392–395. http://dx.doi.org/10.1038/nature03248.
Tesfatsion, L., 2001. Structure, behavior, and market power in an evolutionary labor market with adaptive search. J. Econom. Dynam. Control 25 (3–4), 419–457.
Tesfatsion, L., 2017. Modeling economic systems as locally-constructive sequential games. J. Econ. Methodol. 24 (4), 384–409.
Tesfatsion, L., 2022. Agent-Based Computational Economics: Overview and Brief History. Economics Working Paper No. 21004, ISU Digital Repository, Iowa State University, Ames, IA.
Tsitsiklis, J.N., Van Roy, B, 2002. On average versus discounted reward temporal-difference learning. Machine Learning 2002 49, 179–191. http://dx.doi.org/10.1023/A:1017980312899.
Velupillai, K.V., Zambelli, S., 2012. Computability and algorithmic complexity in economics.
Yook, S.H., Radicchi, F., Meyer-Ortmanns, H., 2005. Self-similar scale-free networks and Disassortativity. Physical Review E 72 (4), 1–4. http://dx.doi.org/10.1103/physreve.72.045105.
Zhang, Z., Zhou, S., Zou, T., 2006. Self-similarity, small-world, scale-free scaling, Disassortativity, and robustness in hierarchical lattices. Eur. Phys. J. B 56, 259–271. http://dx.doi.org/10.1140/epjb/e2007-00107-6.
Zorn, L., Wilkins, C.A., Engert, W., 2009. Bank of Canada liquidity actions in response to the financial market turmoil. In: Bank of Canada Review. Bank of Canada, pp. 7–26.