BANK OF CANADA
BANQUE DU CANADA

Discussion Paper/Document d'analyse
2009-13

Network Analysis and Canada's Large Value Transfer System

by Lana Embree and Tom Roberts

www.bank-banque-canada.ca

Bank of Canada Discussion Paper 2009-13

December 2009

Network Analysis and Canada's Large Value Transfer System

by

Lana Embree and Tom Roberts

Funds Management and Banking Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
lembree@bankofcanada.ca
troberts@bankofcanada.ca

ISSN 1914-0568

Acknowledgements

We would like to thank Meyer Aaron, Nikil Chande, James Chapman, Jonathan Chiu, and Miguel Molico for useful comments and suggestions.

Abstract

Analysis of the characteristics and structure of a network of financial institutions can provide insight into the complex relationships and interdependencies that exist in a payment, clearing, and settlement system (PCSS), and allow an intuitive understanding of the PCSS's efficiency, stability, and resiliency. The authors review the literature related to the PCSS network and describe the daily and intraday network structure of payment activity in the Large Value Transfer System (LVTS), which is an integral component of Canada's financial system. The picture that emerges confirms that the LVTS is highly centralized among a few key participants, similar to other PCSSs, based on the network of relationships. This could heighten the systemic importance of these participants, and the susceptibility of the system to financial contagion. There have been small variations in the relative importance of individual banks, but the network-wide pattern of relationships has remained stable from 2004 to 2008, even through the credit crisis.

JEL classification: D85, G10
Bank classification: Payment, clearing, and settlement systems; Financial stability

Résumé

L'analyse des caractéristiques et de la structure d'un réseau d'institutions financières peut jeter un éclairage intéressant sur les relations et interdépendances complexes qui existent au sein d'un système de paiement, de compensation et de règlement (SPCR), et permettre de comprendre intuitivement l'efficience, la stabilité et la résilience d'un tel système. Les auteurs passent en revue la littérature se rapportant aux réseaux de SPCR et décrivent la structure en réseau des opérations de paiement journalières et intrajournalières du Système de transfert de paiements de grande valeur (STPGV), qui fait partie intégrante du système financier canadien. Le tableau qui se dégage de leur étude confirme que, d'après le réseau des relations, le STPGV est fortement centralisé autour de quelques participants clés, à l'exemple d'autres SPCR. Cette configuration pourrait accentuer l'importance systémique des participants clés et la vulnérabilité du système à la contagion financière. L'importance relative des banques individuelles a varié légèrement entre 2004 et 2008, mais le profil des relations dans l'ensemble du réseau est demeuré stable, même pendant la crise du crédit.

Classification JEL : D85, G10
Classification de la Banque : Systèmes de paiement, de compensation et de règlement; Stabilité financière

1 Introduction

In financial systems, the network of relationships between financial institutions (FIs) could have important policy implications in terms of counteracting vulnerability to financial shocks and contagion. Central banks and regulators may have disregarded the network aspects of the financial system, even as individual FIs adopted homogeneous business and risk-management practices that resulted in a lack of diversity in the industry, possibly contributing to the fragility that the 2007–09 credit crisis exposed (Haldane 2009). The systemic importance of key participants, which is a critical factor for shocks and contagion, also arguably depends on their interconnectedness in the financial network. This was the view held for AIG, for which the U.S. Federal Reserve engineered a public buyout in September 2008 (*Financial Times* 2008).[1] Payment, clearing, and settlement systems (PCSSs), in particular, are an integral component of the financial system; they represent the platform on which the last stages of a financial transfer are completed. The PCSS network structure thus offers a relevant and timely perspective on broader financial relationships.

This paper reviews the literature on networks, emphasizing how it might relate to PCSSs, particularly Canada's Large Value Transfer System (LVTS); we then empirically interpret and describe the LVTS as a network of FIs connected by payment flows. Network analysis gives us a way to study and understand the complex arrangement of relationships, and interdependencies, that exist in a PCSS. Examining the network characteristics of a payments system can give us an intuitive understanding of the trade-off between efficiency and stability, which is a prominent concern.

The literature review covers three streams: the first describes other PCSS networks, which we contrast with the LVTS. The second stream explores how network characteristics can influence PCSSs when they face shocks or different system parameters. For example, theoretical work specifically on PCSSs that relates to network relationships has examined payment coordination and payment queuing. The third stream of literature models network formation. The identification of algorithms that replicate observed structures might offer insight into the plausibility of factors that influence formation.

The empirical section uses an approach similar to that found in the recent literature of Soramäki et al. (2006), Inaoka et al. (2004), and Boss et al. (2003), which aims to describe financial systems through quantitative network analysis. We describe stylized facts of the LVTS payment flow network over the April 2004 to December 2008 period;

[1] AIG had sold a large amount of credit-default swaps on fixed-income securities tainted by the subprime-mortgage-lending crisis, and had an extensive global network of counterparties (Walsh and de la Merced 2008).

the picture that emerges confirms that the structure of relationships in the LVTS is highly centralized among a few key participants. This type of hub-and-spoke structure is similar to other PCSSs. Its small-world structure could increase susceptibility to financial contagion, while having implications for resiliency, efficiency, and adaptability to shocks.[2] The network-wide pattern of relationships has been stable, even if there have been small changes in the relative importance of individual banks. An econometric analysis does not give conclusive results for the influence of the credit crisis on network structure.

Section 2 provides a brief summary of the network concepts and measures used in this study. Section 3 reviews the literature. Section 4 describes our data and their limitations, and section 5 the empirical methodology used. Section 6 describes the results. Section 7 discusses our findings and suggests further empirical research. Section 8 offers some conclusions.

2 Summary of Network Concepts and Measures

A network can be described as a structure of nodes connected by links. In the context of a payments system such as the LVTS, nodes represent FIs, and links can represent payment flows, bilateral credit limits, loans, or other variables describing a relationship between two FIs over a given observation period.[3] Social network analysis involves different terminology than economists commonly use, so we will begin by defining the main terms and measures used in this study.[4] These network measures describe characteristics of PCSSs that might help us to understand, for example, the resiliency and efficiency of the network.

Figure 1 depicts a network between four nodes in both matrix and graphical form. In graphical form, an arrow represents a *directed link* going from one node to another.[5] *Undirected links* can also describe network relationships; however, we will focus on the

[2] The LVTS eases the implementation of Canadian monetary policy, and FIs use the LVTS to transfer funds and settle payments related to the delivery of financial assets. These transactions enable the allocation of capital, management of risk exposure, and conversion of savings into investments, which are essential for economic growth. In 2008, about 22,593 payments with a total value of Can$181.6 billion, on average, flowed through the LVTS each business day; see Arjani and McVanel (2006) for an overview of the LVTS. International organizations such as the Bank for International Settlements have recognized the LVTS as being liquidity efficient and comparable in safety to the more traditional real-time gross settlement (RTGS) systems (CPSS 2005, 36).

[3] The technological infrastructure underlying the LVTS is a static network itself, apart from system outages and the entry of State Street on 18 October 2004. The network of payment flows, however, is more dynamic, and displays variations according to the time of day and over the calendar year.

[4] Wasserman and Faust (1994) and Jackson (2006) also provide introductions to the measures used in social network analysis.

[5] In PCSSs, directed links are used to represent, for example, which node is the payment sender or credit grantor.

directed case, where links are not necessarily reciprocal. A matrix element $m_{ij} = 1$ represents a directed link. Taking the example of payment flows in a PCSS, $m_{ij} = 1$ represents the occurrence of a payment from sender i to receiver j. The case of $m_{ij} = 0$ represents an absence of any payment flow from i to j for that observation period. The appendix more formally describes how network statistics can be calculated from this matrix.

Figure 1: An Example Network Matrix and Graph

$$\begin{array}{ccc} & & \text{IN} \\ & & \begin{array}{cccc} A & B & C & D \end{array} \\ \text{OUT} & \begin{array}{c} A \\ B \\ C \\ D \end{array} & \begin{pmatrix} 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix} \end{array}$$

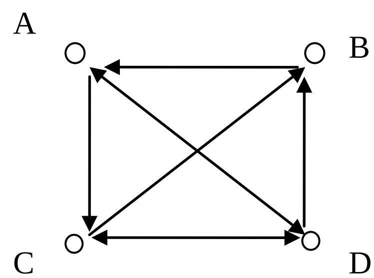

Figure 1 shows a network of unweighted links, where a link between two nodes either exists or does not; a weighted link would indicate the intensity of the link. For example, in a PCSS, a network of weighted links could represent the value of the payments that participants make to each other or a transition probability matrix for liquidity movement between banks.[6] Many network measures simplify relationships to unweighted links, as in Figure 1, which is a reasonable place to start.

Some of the basic network measures, such as connectivity, reciprocity, and degree, relate to a counting of links. *Connectivity* calculates the number of actual links that exist in the network divided by the total number of potential links, giving a measure of network completeness. In a PCSS, connectivity is a measure of activity and indicates whether participants are transacting, in general, with many or only a few of the other participants. In Figure 1, there are 8 observed links and 12 potential links; thus, connectivity is 2/3.

For directed graphs, *reciprocity* measures the percentage of observed links that also have a link travelling in the opposite direction; as the word implies, this is the percentage of links that are reciprocated. In Figure 1, A and D have a reciprocal relationship and C and D have a reciprocal relationship; thus, four links are reciprocated. There are eight

[6] In Bech, Chapman, and Garratt (2008), a transition probability matrix reflects the probability that a dollar will be sent from each LVTS participant to any other participant. They use the stationary equilibrium resulting from this matrix to identify the central participant in terms of who is likely to hold the most liquidity.

observed links in Figure 1, so the reciprocity in this example is 1/2. In the LVTS, reciprocity could provide information about payment coordination, which influences liquidity needs. For example, if there is a reciprocal relationship between two participants, one of the participants is using a payment received from the other as a source of liquidity to send a return payment, rather than drawing on its bilateral credit limit (BCL).[7]

Degree is a count of a node's direct relationships, also known as its neighbours. The *in-degree* of node *i* is the number of links that end at *i*, while the out-degree of *i* is the number of links that start with *i*. For example, in Figure 1, node D has an in-degree of 2 and an out-degree of 3. In a PCSS, degree indicates the number of counterparties that a participant interacts with during a given period.

The *degree distribution* is a frequency distribution that describes how many nodes have any given degree in the population of nodes. Figure 2 depicts the out-degree distribution of the example network in Figure 1. Empirical degree distributions often seem to resemble a *power-law* distribution, both in the natural world and the social sciences. For example, Barabási and Albert (1999) find power-law distributions for article citations in the scientific literature, and for the numbers of links to web pages. Power-law distributions are *scale free* – for a given ratio of two values in the distribution, the relative frequency of encountering the two values does not change. For example, with a power-law exponent of 2, a node of degree 6 is four times less frequent than a node of degree 3; a node of degree 10 is four times less frequent than a node of degree 5. Power-law networks have a few high-degree nodes, and many nodes of low degree, although high-degree nodes are more common than in a Gaussian distribution. Soramäki et al. (2006) show that, in some fields, scale-free networks are robust to the removal of random nodes but are vulnerable to becoming disconnected by the removal of highly connected nodes. However, they also note that the structure of a network and the mechanisms for contagion differ across networks and fields and therefore the vulnerability to targeted node removal might not necessarily apply to PCSSs. They note that, in a PCSS, understanding how liquidity flows through the system will help to assess the robustness of the network.

[7] In the context of a PCSS, liquidity is the ability to meet payment obligations, by having funds or credit available to complete a transaction.

Figure 2: Out-Degree Distribution of Figure 1

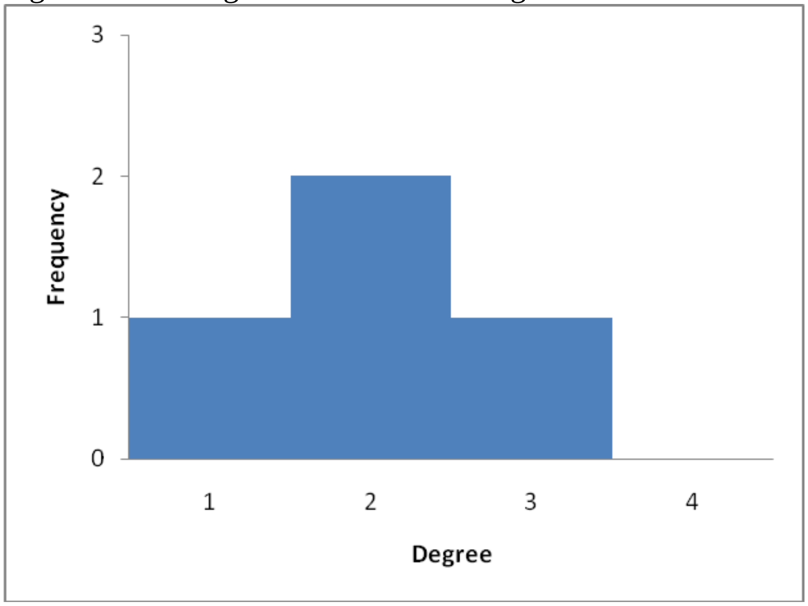

The *clustering coefficient* measures the proportion of a node's neighbours that are also neighbours to one another. In other words, it measures the probability that if node *i* is linked to nodes *j* and *k*, then *j* and *k* are also linked. In Figure 1, node C has two neighbours, B and D.[8] Two links can exist between B and D: a link from B to D and one travelling in the opposite direction. One of these possible links exists in the example graph; therefore, node C's clustering coefficient is 1/2. A clustering coefficient for the network averages these coefficients across individual nodes. The network clustering coefficient in Figure 1 is 1/2.[9]

Paths help to measure how close nodes are to one another at any given time. A path is a sequence of nodes and links beginning and ending with nodes, where any link or node is not included more than once. The *length* of a path is measured by its number of links. An example path in Figure 1 begins at node A, goes through node D and then ends at node B, for a length of two. In the LVTS, all participants can make payments to all others, so paths do not reflect the course that payments can travel; however, a path can reflect the course that liquidity or contagion could follow.[10]

---

[8] The clustering coefficient is often defined for out-links of nodes. Using that approach, nodes B and D are C's neighbours because a link travels from C to these nodes. Node A is not C's neighbour because the link travels from A to C; however, node C is A's neighbour. This occurs because Figure 1 is a directed graph, where relationships are not necessarily reciprocal.

[9] Node A's clustering coefficient is 1, while nodes C and D both have clustering coefficients of 0.5 and node B has a clustering coefficient of 0 because it has only one direct neighbour. Soramäki et al. (2006) note that, by definition, a node with degree 1 has a clustering coefficient of zero.

[10] Allen and Gale (2000) model financial contagion that follows interbank claims (loans).

The *distance* between a pair of nodes is the length of the shortest path connecting them. For example, in Figure 1, the distance from node C to A is two. In PCSSs, distance might provide an indication of how a shock affecting one node could travel through links to affect another. The *average path length* is the average distance for any combination of two nodes in the network. In Figure 1, the average path length is 1.33.[11] In the LVTS, an average path length of one indicates that all participants have sent a payment to all others. A longer path length indicates that activity is concentrated among fewer pairs of participants. When there is no path between a pair of nodes, that distance is undefined, so we calculate distance-based measures for the *giant strongly connected component* (GSCC). The GSCC is the largest subnetwork where all members have both ingoing and outgoing links to the rest of the GSCC as a whole. There is a directed path between any two nodes that are members of the GSCC.

Eccentricity and *diameter* provide further descriptions of the distance between nodes. The *eccentricity* of a node is the distance to the node that is the farthest away. In Figure 1, node C's eccentricity is two, since it is two links away from node A and all other nodes are closer. The *diameter* is the largest eccentricity in the network, and indicates the maximum distance between any two nodes in the network. In the example in Figure 1, the diameter is two because the distance between all pairs of nodes is two or less. Eccentricity and diameter can provide an indication of how easily or quickly an event affecting one node could potentially affect the other nodes in the network. For example, if one participant ceases to send payments, those with direct relationships might find themselves short of liquidity sooner than those who have only indirect relationships with that participant.

Centrality helps to indicate the importance of a given node to the network as a whole. We measure *degree*, *value*, and *betweenness* centrality.[12] For PCSS studies, centrality would ideally identify systemically important participants, which for example might be important for the smooth flow of liquidity through the system.[13] Also, a disruption or default might have a larger impact on the network if a central participant experienced the disruption or default.

Degree centrality equates centrality directly with the degree of a node and hence does not consider indirect relationships. Applying this measure, the most central node is the one with the most direct relationships. For this paper, we have defined *value* centrality to mean a node's share of the total gross value of payments sent and received, which also

---

[11] In the example in Figure 1, there are eight pairs of nodes that have a path length of 1 and four pairs of nodes with a path length of 2. Thus the average path length is 16/12 or 1.33.

[12] Other centrality measures are discussed in Wasserman and Faust (1994), Jackson (2006), and Borgatti (2005).

[13] von Peter (2007) uses centrality to identify international banking centres.

considers only direct relationships. A bank that was a counterparty to all transactions would have the maximum centrality for either degree or value centrality.

Betweenness centrality considers both direct and indirect relationships. The betweenness centrality of a node is the probability that the node is used as an intermediary on the shortest path between any two other nodes; it aims to measure how important a node is for flows between other nodes in the network. Different centrality measures could allow different patterns to become apparent, even though they are based on the same links. In the case of LVTS payment flows, centrality measures have strong positive correlations.

Centralization provides a summary statistic of the extent to which the network is centralized; it takes the centrality scores of individual nodes and sums up their deviations from the maximum centrality score. Centralization is normalized by the "theoretical maximum deviation" to give a value in the range of zero to one: zero represents a network where all nodes have equal centrality, and one represents a hub-and-spoke network where a single node has maximum centrality and all other nodes are equally peripheral. Centralization could be useful for comparing across networks or across time. Figure 3 shows a decentralized network on the left, where all nodes have links to all other nodes. Using any of the centrality measures, the nodes are all equally central, and the centralization value of this network is zero. In the hub-and-spoke network on the right, the hub node is central, and this network has a centralization value of one.

Figure 3: Two Example Networks

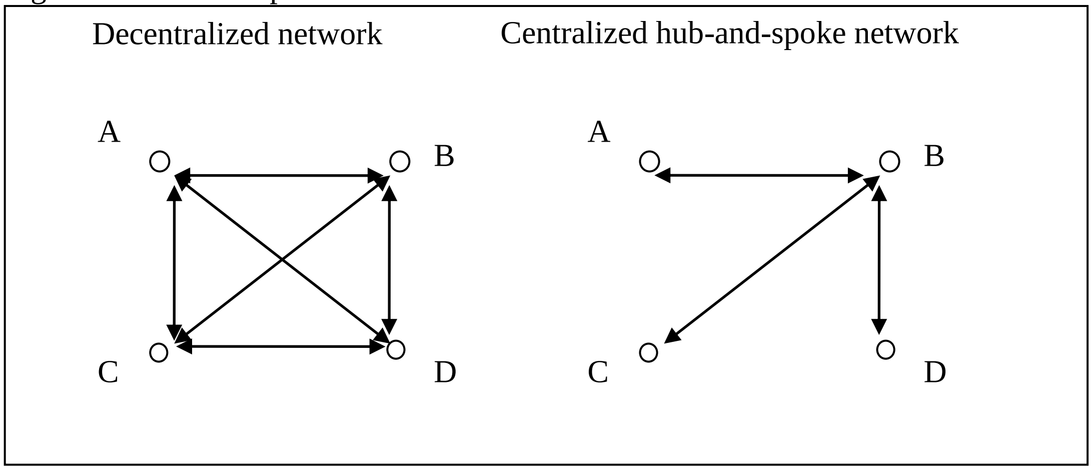

A caveat to these statistics is that they do not incorporate the varying importance of individual links, for example, by reflecting payment values or transition probabilities. The distribution of individual transaction payment values is highly skewed to the right, with the majority of payments having relatively low value; a later section will discuss the similar case of bilateral links. Meanwhile, the bulk of value that FIs transfer through the

LVTS occurs in the smaller number of high-value payments.[14] A possible extension of this analysis is to formulate new measures that suitably capture the importance of links relative to the size of an FI. Value centrality is one measure that simply reports on a bank's share of gross payment flows.

Neither do network statistics of payment activity represent financial exposures, which would be important for financial contagion. But a network analysis of payment flows could be valuable as a tool to understand broad patterns of activity, when detailed data of financial exposures are not available or manageable: it would reveal which participants play an important role, and the avenues for liquidity flows, which relate to financial contagion.

3 Related Literature

Network analysis first appeared in statistical physics and sociology; the existence of network-like structures is common in both the natural and human world. Examples include river systems, the Internet, and networks of human relations (Inaoka et al. 2004). Though the network literature is multidisciplinary, we focus on economics papers and other papers that could be relevant for understanding PCSSs.[15]

We cover three streams of literature: the first includes a small number of empirical papers that have used social network measures to study financial system networks. Empirical results give a reference point for theoretical work, and allow us to compare the network structure of the LVTS to other PCSSs. The second stream of literature comprises theoretical models of PCSSs and examines how network characteristics of the system influence outcomes. The third stream models network formation; these models are not necessarily found in the economics literature, but we discuss their potential applications to PCSSs. Models of network formation can contribute to our understanding of why certain network structures arise under different conditions, which could permit insight into participant behaviour.

3.1 Empirical network studies of PCSSs

Soramäki et al. (2006) describe the topology of interbank payment flows between participants in the Fedwire Funds Service, which is the large-value payment system (LVPS) at the core of the U.S. financial system; Inaoka et al. (2004) describe payment flows in BOJ-NET, which is Japan's LVPS; Boss et al. (2003) examine the network of

[14] In 2008, the average LVTS payment value was $8.04 million, while the median was $44.6 thousand. The extreme values for a single payment were close to $13.4 billion, and less than a dollar. Fedwire had an average of $3 million and a median of $30 thousand for the first quarter of 2004.

[15] Jackson (2006) summarizes and discusses the economic applications of network analysis. Haldane (2009) draws insights from the literature in other fields to examine the credit crisis.

interbank loans in Austria.[16] Figure 4 depicts these networks, and Table 1 summarizes some of their results, along with the corresponding measures for the LVTS. In contrast to the usual network algorithms in the literature, Bech, Chapman, and Garratt (2008) use Markov theory to identify central participants in the LVTS, and Chapman and Zhang (2009) identify communities within the LVTS, based on the probability of different bank interactions.

In Figure 4, points or circles represent banks, and lines represent payment activity between those banks. Fedwire has over 9.5 thousand participants, of which about 6.5 thousand might be active on a given day; therefore, in the top left, we can see only the basic shape of a dense core of the 1 January 2004 network, along with the rest of the network, which becomes increasingly less connected, away from that core. The core of Fedwire in the top right is the part of the payments system that accounts for 75 per cent of the value transferred. We can see about four banks that appear to be more strongly connected than others. Network measures might have further identified their relative systemic importance, possibly contributing to different policy prescriptions for those banks during a crisis.

[16] Additionally, there are many similar internal documents at other central banks.

Table 1: Comparison of International Network Structures

	LVTS payments	Fedwire payments	Austrian interbank loans	BOJ-NET payments
Time period	Daily. April 2004–December 2008.[a]	Daily. First quarter 2004.	Quarterly single-month periods between 2000 and 2003.	Monthly. June 2001.
Payment statistics, U.S.-dollar total for 2004[b]				
Value (trillion)[c]	25.4	478.9		188.8
Volume (million)[c]	4.4	125.1		5.2
Avg. value (million)[c]	5.8	3.8		36.5
Value as % of GDP[c]	2,558	4,099		4,099
Network statistics (variable time frames)				
Number of nodes[d]	14	5,086 ±123	883	354
Number of links		76,614 ±6,151		1,727
Connectivity (%)	69.2 ± 3.3	0.3 ±0.01		2.76
Reciprocity (%)	89.3 ± 2.5	21.5 ±0.03		
Average path length	1.31 ±0.03	2.62 ±0.02	2.59 ±0.02	
Average eccentricity	1.84 ±0.07	4.67 ±0.33		
Average diameter	2.01 ±0.07	6.6 ±0.5		
Clustering coefficient (%)	84.3 ± 1.5	53 ±1	12 ±1	
Power-law exponent		2.11 ±0.01 2.15 ±0.01[e]	2.01[f]	2.1

a. The LVTS measures use a threshold of 1 per cent of the lower Tranche 2 Net Debit Cap (T2NDC) for each bilateral pair of banks, averaged over the study period. A participant's T2NDC is the limit on their multilateral net debit position, and relates to the bilateral credit limits granted to it by other participants. See section 5 for further details.

b. The average Can$/US$ exchange rate for 2004 was 1.3012 (CPSS 2008).

c. CPSS 2008.

d. The number of nodes in BOJ-NET is the number of banks for which there are 21 or more transactions with at least one other bank over the 1-month period studied. The number of nodes in Fedwire on a given day is the number of nodes that are part of that day's GSCC. The number of nodes in Boss et al. (2003) is the number of banks. The number of nodes in the LVTS excludes the Bank of Canada.

e. The first power-law exponent listed in the Fedwire study corresponds to the out-degree distribution, and the second to the in-degree distribution.

f. The power-law exponent reported for this study is for undirected links. When considering directed links, the results are 1.73 for in-degree and 3.11 for out-degree.

Sources: Inaoka et al. (2004), Soramäki et al. (2006), Boss et al. (2003).

Network analyses of payments systems typically find that a small number of FIs are highly connected in their payment flows, and thus act as hubs for the system. Most FIs have much fewer links, and only a small proportion of potential links in the network actually exist. For Fedwire, the connectivity of the GSCC is 0.3 ± 0.01 per cent, while in BOJ-NET connectivity is 2.76 per cent. Despite low connectivity, the path length between any two FIs is often short due to the presence of hubs. The average path length for Fedwire is 2.62 ± 0.02, while for loans in the Austrian interbank system it is 2.59 ± 0.02.

Figure 4: International LVPS Networks

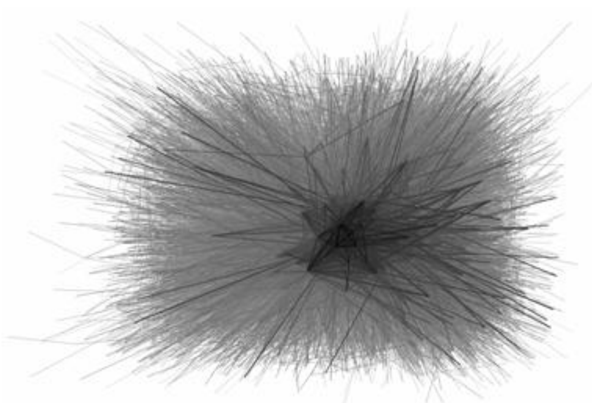

Fedwire interbank payment network.
1 January 2004 (Soramäki et al. 2006).

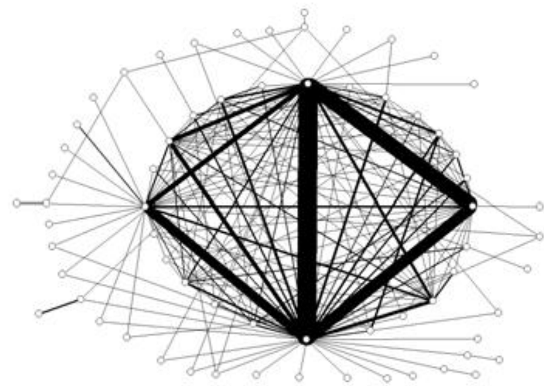

Fedwire core accounting for 75 per cent of value transferred. 1 January 2004 (Soramäki et al. 2006).

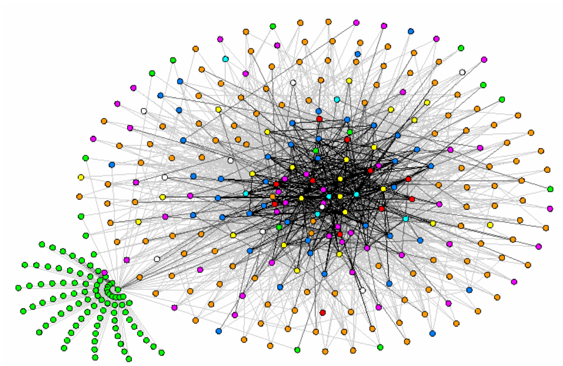

BOJ-NET (Inaoka et al. 2004).

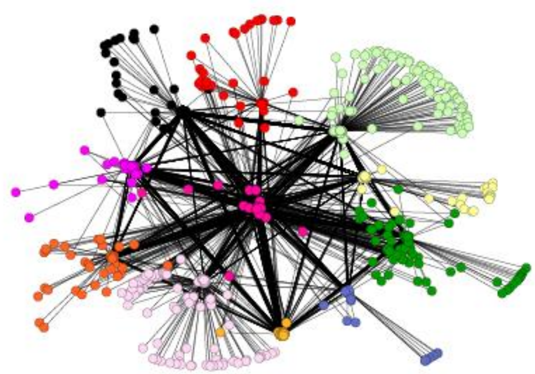

The banking network of Austria (Boss et al. 2003).

Figure 5 shows the 14 LVTS participants apart from the Bank of Canada, with average daily payment flows ranging from thin pink lines, for small values of Can$10 million or less, to thick black lines for larger values. The figure illustrates some general points about the LVTS network of average payment activity in 2008: (i) even a small payments system like the LVTS would be difficult to describe and understand without the right approach; (ii) the LVTS has high connectivity of around 69.2 ± 3.3 per cent on a daily basis; (iii) there is considerable heterogeneity in the strength of links, and a participant’s interconnectedness in the LVTS might not correspond exactly to its size; (iv) much like what we can gather from the picture of Fedwire’s core in Figure 4, this figure suggests that the systemic importance of a participant depends partly on its interconnectedness, with only a small number of banks that are heavily connected with thick lines.

Figure 5: LVTS Average Daily Gross Payment Flows, 2008

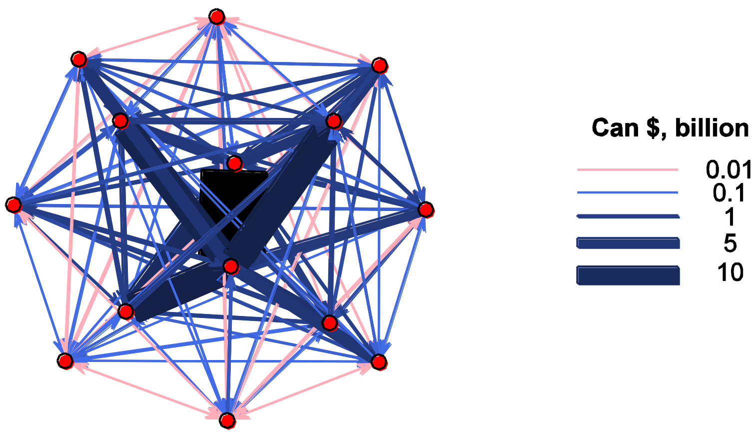

The average path length in the LVTS is 1.31 ± 0.03; despite Fedwire's much lower connectivity, its average path length is only twice that of the LVTS, because of the role of network hubs. In the end, we could describe all of these systems as *small worlds*, which exhibit high clustering and short distances between banks, even when the prevalence of direct relationships is low. These short distances can mean that it takes only a small number of steps for liquidity to circulate between any two banks, or for disruptions to spread.

In Fedwire, half of the nodes have four or fewer outgoing links, while the maximum degree for a node reaches 1,922 ± 121 (out-degree) and 2,097 ±115 (in-degree), with over 5,086 ± 123 participants active on a given day. The degree distribution for Fedwire follows a power-law distribution, for degrees greater than 10, which is similar for both in- and out-degrees. Maximum-likelihood estimates of the power-law exponents are 2.15 ± 0.01 for in-degree and 2.11 ± 0.01 for out-degree (Soramäki et al. 2006). Boss et al. (2003) find a power-law exponent of 2.01 for undirected links, while Inaoka et al. (2004) estimate an exponent of 2.1 and find a fractal structure similar to other network structures such as river basins or the Internet. The smaller size of the LVTS does not afford much opportunity for such an assessment, but, for a threshold where we examine only links that have a value of at least 40 per cent of the lower Tranche 2 Net Debit Cap (T2NDC) for a given pair of banks,[17] the big five banks have an average degree over three times greater than the other banks. These hubs would seem to have the greatest potential for propagating financial contagion, because of the network's reliance on their interconnecting role.

[17] In section 5 we more fully describe the thresholds used in this study.

Networks with power-law distributions in other disciplines have been found to be robust to random failures, but not to a failure of a hub (Albert, Jeong, and Barabási 2000; Crucitti et al. 2004). Whether this result also holds for payments systems depends on the mechanisms of contagion and adaptation by FIs (Soramäki et al. 2006). Nevertheless, the network structure of payments systems suggests that they have evolved primarily with efficiency in mind, rather than resiliency to extreme events (Inaoka et al. 2004). Some banks would presumably take on the role of hubs because of efficiency advantages, whether this was to result from economies of scale, information asymmetries, history, or other reasons. An ability to intermediate information and liquidity flows could be a further advantage for a hub bank, such that a network structure is self-reinforcing, much like in the preferential attachment models we will discuss later. Conversely, inefficiencies or disadvantages that impede the growth of smaller institutions in the network might have a similar self-reinforcing effect.

A complete network might create resilience through risk dispersion by spreading the impact of a shock safely across many nodes (Allen and Gale 2000), but, at the same time, it could increase vulnerability to shocks that exceed a threshold beyond which the system is no longer self-repairing (Haldane 2009). Increased financial network complexity, partly due to structured financial products, can amplify uncertainty because the precise source and location of underlying claims becomes unknown. If there is also homogeneity in business and risk-management practices, a relatively small initial financial disturbance could result in herd behaviour and a cascade of losses that transmits easily through a highly connected network (Haldane 2009).

Fedwire reciprocity is low on any given day, at around 22 ± 0.3 per cent, and is uncorrelated with the number of nodes and links, or the volume and value of payments. The clustering coefficient is 53 ± 1 per cent, which is high relative to the clustering of Austrian interbank loans at 12 ± 1 per cent. Clustering in the Austrian system might be relatively low because the cost of link formation for smaller Austrian branch banks, which mostly rely on a head institution, might not justify the benefits of diversification (Boss et al. 2003). Daily reciprocity and clustering in the LVTS are high, at around 89.3 ± 2.5 and 84.3 ± 1.5. The tendency for reciprocity to be high relative to connectivity reflects the bilateral coordination of payments and its importance for managing liquidity.

More involved algorithms have further identified network characteristics of the LVTS. Bech, Chapman, and Garratt (2008) apply Markov theory to identify the central participants in the LVTS. In that paper, the central participant is the one that we expect to hold the most liquidity, where the network structure represents a transition probability matrix of liquidity movement that determines liquidity holdings. Chapman and Zhang (2009) use transactions-level data to identify where LVTS participants transact more heavily with other participants in the same community than with those in different communities. Using two different measures of transaction intensity, the authors find that there could be communities within the LVTS based on either transaction size or the participants' geographical locations.[18]

Papers that characterize an existing network are an important start to further research on PCSSs. Network descriptions can provide input for modelling how networks form, which we discuss in section 3.3. In addition, there is evidence that network structure can influence outcomes, as shown in the models we describe in the next section. Descriptive studies of PCSSs can contribute to the discussion and intuitive understanding of observed outcomes.

3.2 Models of PCSSs

Existing papers that incorporate a network into PCSS modelling take the network as given. These models help researchers understand how the existing network influences economic outcomes; the models do not attempt to explain why a certain network exists. In this section, we describe several papers: Bech and Garratt (2006) examine disruptions, Allen and Gale (2000) study a liquidity shock and whether it leads to contagion, Gale and Kariv (2007) examine the network of trading relationships and what features can lead to market breakdowns, and Martin and McAndrews (2007) model payment queuing.[19] These papers demonstrate that network characteristics can influence economic outcomes, and that such networks can be complementary to economic modelling.

Bech and Garratt (2006) study the equilibrium effects of a disruption in the ability of a few participants to send payments, and explore the impact of the network on coordination. They use a game-theoretic model with a large bank and several small banks. Banks choose to send payments in the morning or the afternoon, depending on the cost of delay and liquidity. When each bank chooses whether to make payment in the morning or afternoon, they consider the action taken by the large bank, and the number of small banks making their payments in the morning, since this influences the cost of liquidity. Banks obtain liquidity in part from incoming payments; if there is a substantial shift in payments to the afternoon, a bank might also prefer to pay in the afternoon rather than face the costlier alternative of intraday credit to fund a morning payment.

The model begins in the efficient equilibrium where all participants pay in the morning; a disruption then temporarily forces a subset of banks to pay in the afternoon. An alternative equilibrium occurs if all participants pay in the afternoon; this equilibrium is

[18] Chapman and Zhang (2009) note that geography itself might not be the cause of the community, but rather that geography might serve as a proxy for other characteristics or business relationships that lead to the formation of the community.
[19] The Bech and Garratt (2006) and Martin and McAndrews (2007) models are both based on the game-theoretic model of payment coordination in Bech and Garratt (2003).

inefficient because of delay costs. The main question that Bech and Garratt (2006) address is whether the disruption will cause the banks to move to the inefficient afternoon equilibrium. The outcome depends on the decisions taken by each bank, and in particular the large bank, since the incentives each bank faces to move to one equilibrium or the other depend on whether the other participants are paying in the morning or the afternoon.

The network of the system plays an important role: when the large bank is bigger, or when small banks do not often transact among themselves, the large bank influences the outcome more. If the small banks transact only with the large bank, we observe a hub-and-spoke network, with the large bank in the centre; if small banks also transact among themselves, we observe a more complete network. It is not surprising that the influence of the large bank is greater in the hub-and-spoke network, because small banks are entirely reliant on the large bank for any incoming liquidity that they receive.

Allen and Gale (2000) examine the role of the network between banks in financial contagion. In this model, banks use an interbank market to redistribute liquidity among themselves. When a bank experiences a liquidity shock, the network topology will influence how the shock might spread to other banks. If the network is complete, all banks will experience some loss. Dispersing the loss across all banks can diminish the impact of the crisis on each bank and reduce the likelihood of contagion. Conversely, when the interbank market is incomplete, the immediate counterparty exposure could be concentrated on fewer banks, leading to a greater risk of contagion where additional banks succumb to the shock.

Gale and Kariv (2007) model a market for trade, where an assumed network determines which nodes can trade. In a complete network, all nodes can trade with any other node in the network, whereas if the network is incomplete some nodes cannot trade directly with each other. Using this approach, the LVTS would be a complete network, since all participants can trade with each other, although LVTS participants would be incompletely connected to domestic or international systems beyond the LVTS. The authors suggest that an incomplete network might occur if asymmetric information leads participants to transact only with participants that they “know and trust.” An incomplete network could also occur where an intermediary acts as central counterparty and guarantor. Gale and Kariv (2007) show that transaction costs could also lead participants to limit the number of counterparties, or lead small traders to rely on an intermediary broker. In other words, a move to efficiency could result in some nodes taking the role of intermediary in an incomplete network. The authors argue that market breakdowns can occur if the network structure changes, representing changes to which participants can trade and act as intermediaries.

While the issue of financial contagion comes to the forefront during most financial crises, the efficiency and reliability of operations is another concern for payments systems. The arrangements that determine when queued payments settle provide one example. Martin and McAndrews (2007) model two different queuing mechanisms and study how the network of payments influences which queue generates the highest participant welfare. In their model, participants decide whether to make payments in the morning, the afternoon, or to place them in a queue in the morning.

The network of payments, and how they offset, determines whether payments are released from the queue. Martin and McAndrews (2007) examine two network types. In the first network, all payments offset bilaterally. In this case, a payment in the queue is released if the offsetting payment is also in the queue or if the offsetting payment is made in the morning. In the second network, all payments form a circle. In this case, a payment in the queue can be released only if the sender has received a payment in the morning, or is the recipient of a queued payment that is released by a payment made in the morning. The network type influences the probability of settlement for a payment in the queue, and therefore influences the incentive to place payments in the queue and the resulting welfare that the queue provides. Other parameters, such as the probability of time-sensitive payments and of liquidity shocks, also influence queue usage and welfare.

Many models do not incorporate agent behaviour, which would adapt to changes in policy or other events. Because of the way it aggregates interactions, network analysis can allow for easier interpretation of whether economic shocks result in behavioural changes. One possible approach for payments systems that does consider changes to participant behaviour, and the influence of network effects, is agent-based modelling.[20] This approach will be of particular interest for the study of participant responses to policy changes or actions, or disruptions such as operational outages.

The models discussed in this section demonstrate that a network structure can influence economic outcomes in a PCSS. Because of the network's effect on the results of shocks and the impact of system parameters, such as the queuing mechanism, it is important to understand the network structure. But these models do not try to explain why a certain network exists. The next section describes models of network formation that do address this question.

3.3 Models of network formation

Theoretical models, in various fields, have offered explanations for how networks form. We describe several of those models in this section, focusing on how they can be used to

[20] Work has begun at the Bank of Canada to develop an agent-based model of the LVTS (Labelle St-Pierre 2009).

better understand PCSS networks. We discuss random formation, preferential attachment, and strategic formation models. Each of these classes of theoretical model can provide insight into how network characteristics might be determined. Additionally, some of the models could provide insight into how the network might change in response to a shock, where these changes reflect participant behaviour.

Models of network formation can help us to understand what network will form under certain conditions and shocks to the system. Models of random network formation assume that links randomly form between a fixed number of nodes. The Erdös-Rényi (1960) model is a basic prototype where each pair of nodes has an unconditional probability *p* of being linked. When the number of nodes is large, the degree distribution is approximated by the Poisson distribution (Jackson 2006).

Random models have been extended, or adapted, to account for the observed clustering and degree distribution found in real-life networks, which the purely random model does not generate. These extensions typically involve non-random elements; for example, Watts and Strogatz (1998) and Newman and Watts (1999) begin with clustering and randomly change or add links, which allow these models to account for observed clustering and short path lengths. Newman, Strogatz, and Watts (2001) model random graphs while imposing an arbitrary probability distribution for the node's degrees.

An example of how random models could relate to a PCSS is that the times at which LVTS participants receive payment requests from clients could be thought of as the outcome of a random process that would generate a network of payment flows. Because payments are not necessarily submitted to the LVTS when the request is received, the observed LVTS payment times will not perfectly reflect the random process. But Newman, Strogatz, and Watts (2001) suggest that random network models can act as benchmarks for observed networks. The benchmark can help identify characteristics in the observed network that are caused by participant incentives rather than randomness, and researchers can then begin to identify what those incentives might be.

Preferential attachment models also provide an explanation for how networks form. Barabási and Albert (1999) describe the basic design of preferential attachment models. They are typically characterized by growth: each period, a new node joins the network and forms links with existing nodes. In preferential attachment models, nodes are more likely to form links with high-degree nodes than with low-degree nodes. In a network of friendships, for example, a new person would often prefer to befriend a well-connected person, because having many friends-of-friends confers benefits in the form of social encounters, or more potential friends, in addition to whatever intrinsic value the well-connected person might provide (Brueckner 2004). In a financial network, indirect links could increase market access and trading opportunities through counterparty intermediation.

Various model designs can produce preferential attachment; in some models, the probability of linking with node *i* is proportional to *i*'s degree, as in Barabási and Albert (1999). Klemm and Eguíluz (2002a, b) create a preferential attachment model where nodes are deactivated based on their degree. When nodes are deactivated, they no longer gain links as new nodes join the network; deactivation could be difficult to interpret for PCSSs. The preferential attachment models result in high-degree nodes acquiring more links than low-degree nodes, which allows these models to produce power-law degree distributions.

Extensions or alterations to the basic preferential attachment models more closely replicate observed network structures. For example, Klemm and Eguíluz (2002b) model a preferential attachment model where there is a probability that each new link will be randomly changed to link to an inactive node. In a model by Jackson and Rogers (2007), nodes form links randomly but also form additional links with some immediate neighbours of the randomly encountered nodes. These extensions allow preferential attachment models to produce commonly observed real-world network characteristics, such as small distances, small diameters, and large clustering coefficients.

Preferential attachment models rely on growth in the number of nodes; however, PCSS networks do not generally grow. Preferential attachment models might be useful for PCSSs if they accurately explain the existing network, or if the existing network has features that behave like growth. For example, we might observe networks that resemble growth models if participants move sequentially. Also, preferential attachment networks tend towards a hub-and-spoke type of structure, where many low-degree nodes connect to a high-degree node, similar to what is found in the empirical studies described in section 3.1.

Strategic models also explain network formation. In these models, a payoff function describes the costs and benefits accruing to a node from a given network; thus, links will form only if they provide a strategic benefit. Bala and Goyal (2000) outline a general payoff function that can be adapted to reflect the incentives faced by nodes, such as indirect links creating value (Bala and Goyal 2000; Jackson and Wolinsky 1996), or different costs, depending on which nodes link (Jackson and Rogers 2005). In PCSSs, a participant might face different costs for an interbank loan based on whom they are borrowing from, because the lender has incomplete information about the participant's creditworthiness.[21]

[21] Furfine (1999) discusses how asymmetric information leads to relationship banking in Fedwire.

When payoffs are modelled explicitly, as in these strategic formation models, we can study the welfare that a network creates, and thus can compare networks based on total welfare; the network that provides the highest total welfare is the socially optimal network. Links create externalities if a node's payoff depends on the entire structure of the network. Jackson (2006) explains that externalities mean that the socially optimal network will not be stable if it does not align with individual incentives.

Strategic formation models need to define a transition process that dictates how nodes are permitted to change their links. Bala and Goyal (2000) describe a transition process where nodes can observe the network in the previous period and determine whether to add or sever links. This basic process can be altered; for example, Jackson and Watts (2002) allow just one link to change per period, while Bala and Goyal (2000) allow multiple links to change. Finally, some studies, such as Herings, Mauleon, and Vannetelbosch (2006) and Page, Wooders, and Kamat (2005), allow for forward-looking nodes. In models with forward-looking nodes, the transition process permits nodes to consider the effect of their choices on the subsequent choices of other nodes.

Strategic formation models also need to specify a stability definition. The stability definition determines when the network will stop changing. A stability definition might be a Nash equilibrium where all nodes play their best response given the actions taken by other nodes. Nash equilibrium is used in Bala and Goyal (2000), whereas pairwise stability and core stability are modelled in Jackson and Watts (2002). Bloch and Jackson (2006) discuss some of the stability definitions used in models of network formation.

Strategic models might be the most useful of the formation models for explaining PCSS networks. Many of the PCSS participants' actions are the result of strategic decisions. For example, FIs make strategic decisions when granting BCLs. Participants likely choose a BCL to maximize the benefit of receiving payments, and resulting higher liquidity, subject to the cost of credit exposure and collateral. Similarly, interbank loans are strategic decisions balancing interest earned against credit risk. In addition, the decision regarding when to submit payments could be modelled as a strategic decision weighing the costs and benefits of delaying payments. However, some incentives faced by PCSS participants could be difficult to accurately identify and model.

Some work has begun to model network characteristics as strategic decisions; for example, Chapman, Chiu, and Molico (2008) model the decision to become a direct participant in the payments system or to tier through a clearing agent (CA). In this model, the cost structure and information structure are important determinants of the network. Certain small participants do not have public credit histories. If these participants choose direct participation, their counterparties will not lend to them, since the counterparties cannot assess whether the small participant is a safe debtor. This is inefficient, because of the resulting liquidity costs. These small participants can instead choose to tier under a CA who has a public credit history. If they tier under a CA, the frequent interactions between the CA and the small participant mean that the CA can monitor the small participant's creditworthiness and optimally choose whether that participant can borrow. Small participants who are safe debtors will be able to borrow and save liquidity costs when they tier under a CA.

The strategic models are also useful for PCSSs because these models might allow us to better understand how the network will differ if the modelled incentives faced by the participants change, due to either a shock or a change in the system parameters.

4 Data

LVTS data enable us to identify the sender, recipient, payment submission time, amount, and tranche, from April 2004 to December 2008, for 1,196 daily observations.[22] The Canadian Payments Association (CPA), which operates the LVTS, provides the data to the Bank of Canada. The fifteen LVTS participants include the six major Canadian banks, four smaller Canadian FIs, four subsidiaries of foreign-based banks, and the Bank of Canada.

There are limitations to the data. The Canadian financial system extends beyond the LVTS such that direct participants act as correspondent banks for other FIs or organizations. Hence, the data describe only activity passing through a part of the core of the Canadian financial system.[23] There is no way to identify the original sender or final intended recipient, in the case of intermediated banking, and the motivation for payment flows is usually unknown.[24]

In 2008, daily LVTS transaction values averaged around Can$181.6 billion, with daily transaction volumes of around 22.6 thousand. The maximum daily value of Can$268.6 billion was on 4 September 2007, and the maximum volume of 39,303 was on 2 July 2008. The average transaction value was Can$8.04 million in 2008. Tranche 2 in 2008 accounted for 85.2 per cent of the LVTS value and 98.6 per cent of the volume, and had an average transaction value of Can$6.94 million. During the day, value peaks

[22] Participants in the LVTS can send a payment through one of two payment streams. In Tranche 1, they collateralize any credit usage on a dollar-for-dollar defaulter-pays basis, while Tranche 2 uses a survivors-pay collateral system. Tranche 1 accounts for about 10 to 15 per cent of payment value, and 1.5 per cent of payment volume. LVTS participants typically use Tranche 1 for time-sensitive payments, and payments to or from the Bank of Canada. Tranche 2 economizes on collateral requirements and is thus the main LVTS payment stream (Arjani and McVanel 2006).

[23] Neither would these data encompass activity netted out in the CDSX and Automated Clearing Settlement System (ACSS) payments systems.

[24] It is possible, though, to filter out certain types of transactions, by matching their expected characteristics to payments. Furfine (1999) and Hendry and Kamhi (2007) are thus able to characterize the markets for federal funds in Fedwire, or uncollateralized overnight loans in the LVTS, respectively.

anywhere between 16:00 and around 17:00, and tapers off by 18:00, with modest bumps closer to midday, as shown in Figure 6.

The number of transactions shows the strongest peaks at 06:00, 07:00, and 08:00 (Figure 6); the first peak, at 06:00, could result from the release of overnight payments at the start of the daily LVTS processing cycle for general payments exchange. On U.S. holidays in the study period, total daily value and volume in the LVTS decline by an average of 61 per cent and 34 per cent, respectively. Civic holidays correspond to an average decline in daily value of 95 per cent from the mean and a decline in transactions volume of 82 per cent. The August civic holiday is observed in all Canadian provinces except Quebec, but banks continue operations to deal with CLS Bank and Quebec-based institutions. KPSS tests suggest that levels or logs of value and volume are trend stationary; transactions volumes have a certain amount of monthly periodicity, from the pattern of autocorrelations in the linearly detrended component.

The big five banks account for about 75 per cent of the transacted value for both tranches in 2008. A Herfindahl-Hirschman index measure of gross payment flow shares among the participants gives daily values of between 1,495 and 1,876 with a mean of 1,695, which indicates that some banks play a much more prominent role than others do.[25] The results of other measures of centralization specific to network analysis are reported later; in contrast, payment shares and the Herfindahl-Hirschman index do not consider network structure.

[25] The Herfindahl-Hirschman index is calculated by $hh = share_1^2 + share_2^2 + \cdots + share_n^2$ with shares in percentages, such that 10,000 indicates a pure monopoly.

Figure 6: Average Total Value and Transactions in the LVTS, 10 Min. Intervals

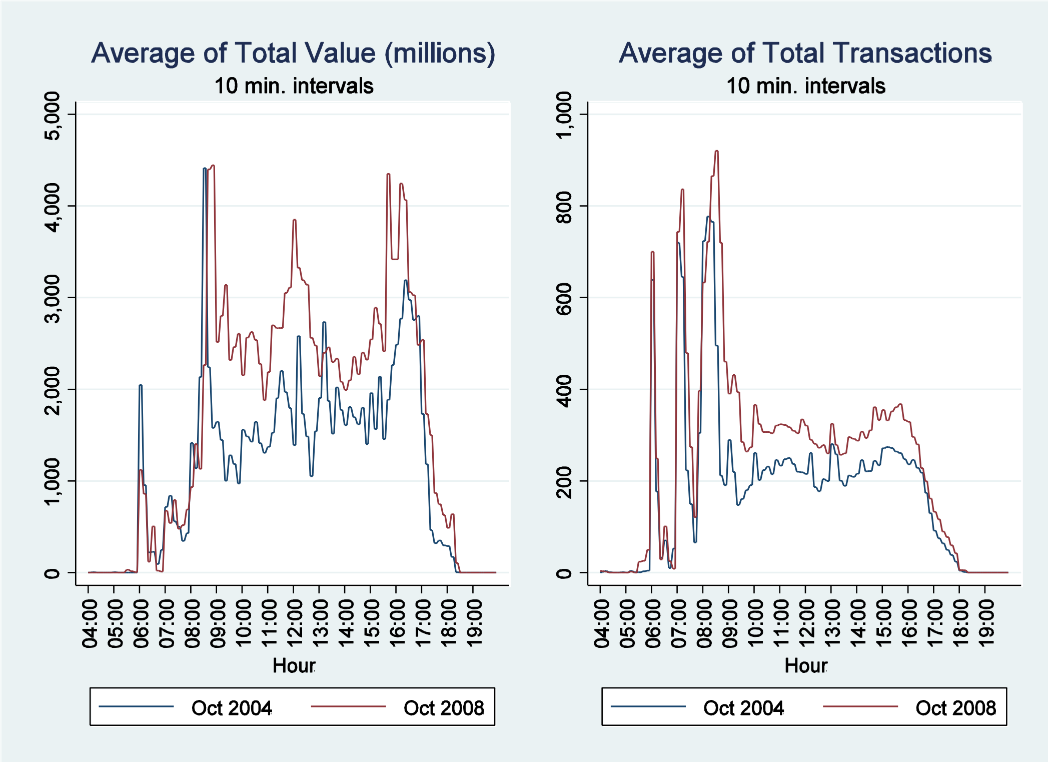

5 Methodology

The LVTS is a relatively small payments system, such that a network analysis of daily payment activity has a different character than for other systems; the LVTS network is almost completely connected, and, based on centrality measures, it would be difficult to distinguish between participants, without a threshold. We therefore examine both a daily frequency and intraday sampling frequencies of 5, 10, 15, 20, 30, and 60 minutes, and use thresholds for the daily case.[26]

The analysis focuses on Tranche 2 transactions, which represent the majority of value and volume in the LVTS, and excludes the Bank of Canada. In Tranche 1, a participant must fully collateralize its debit limit, resulting in different risk characteristics than in Tranche 2, which instead uses a survivors-pay mechanism that would fully collateralize the largest single default. The Bank of Canada relies mainly on Tranche 1 and serves a much different role relative to the private banks.[27]

[26] In this paper, we use the statistical computing software R and the social network analysis tools (Butts 2009) that are available for that software.
[27] The inclusion of Tranche 1 and the Bank of Canada results in connectivity, reciprocity, and clustering that are somewhat lower, reflecting the less-frequent involvement of the Bank of Canada in the payments system. Distance-based measures are not greatly affected, while degree and closeness centralization are higher, which is not surprising if the network has become more sparse.

The existence of a link depends on the length of the data sampling frequency and any minimum threshold of payment value that qualifies as a link. Obviously, there is some discretion according to the choice of sampling frequency and threshold. An intraday frequency reveals more detail about FI interrelationships, compared to a daily frequency. Nevertheless, it is also possible that a daily frequency, in some cases, provides a more-relevant and less-arbitrary metric. A day is a clearly divided unit of time and a basis for the payment cycle.

An arbitrary threshold is an incomplete substitute for measures that take a more complete account of the weights of different links, but it still helps to filter out links that are unimportant for the network. All bilateral payment values during a sample interval greater than a threshold become links in the network. An intraday frequency also reflects the variability in relationship strength, because, with a higher frequency, links disappear more quickly, unless they recur.

Any thresholds we impose should account for the asymmetry in the amounts that would be significant for different banks. One way to do this uses a participant's T2NDC, since this is a rough gauge of the estimation by other banks of the T2NDC recipient's credit-bearing capacity. The T2NDC jumps discontinuously, so the threshold uses the mean over the study period. Figure 7 illustrates the distribution of daily non-zero bilateral payment values as a proportion of the smaller T2NDC of the two participants in each bilateral relationship.[28] The largest bilateral relationships evidently exist among the big five banks, who are able to use collateral more efficiently, since gross payment values reach a higher multiple of their T2NDCs.[29] Links between other banks and the big five are the next most important, followed by the much smaller links among the non-big five banks. This reflects the type of hub-and-spoke network that we typically see in payments systems.

[28] We use a 30 per cent system-wide percentage (SWP) throughout the data period, for consistency in the threshold calculations. BCLs do not change significantly to compensate for the SWP increase from 24 per cent to 30 per cent which occurred on 1 May 2008.

[29] It is possible that large-bank thresholds should account for whether they process payments and net out exposures more quickly, which we do not consider in this paper.

Figure 7: Bilateral Links as a Ratio of the Pair's Lower T2NDC

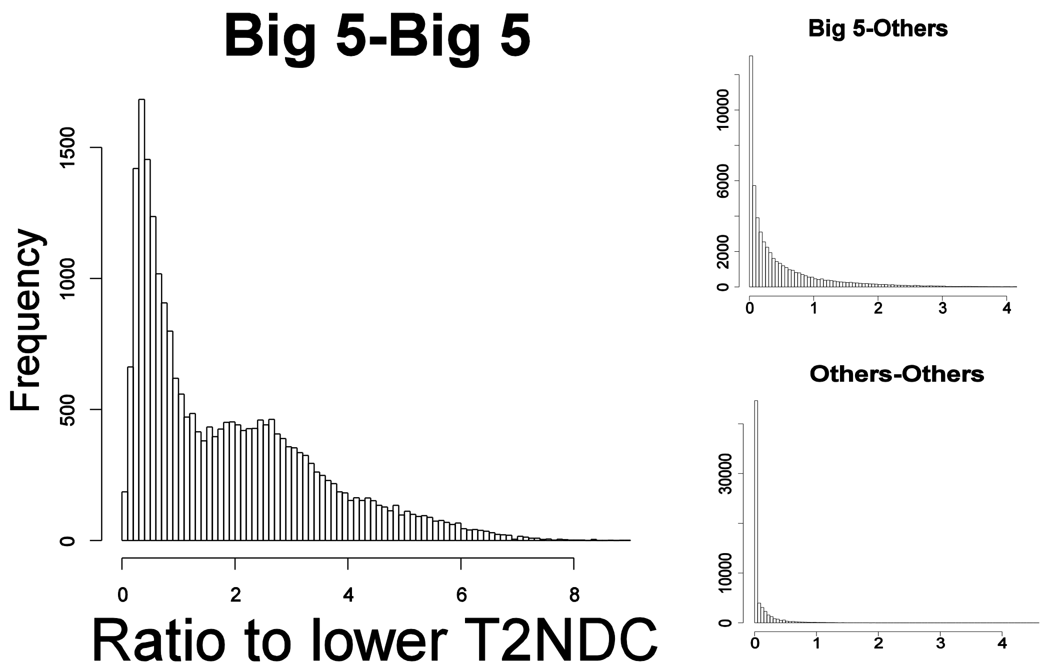

Immediate interpretation of network statistics can be difficult; for example, the immediacy of network connections alone cannot determine the risk of contagion. Some nodes and links might easily disappear without disruption, especially if the underlying transactions can take place with an alternate counterparty. Other financial relationships, even if less frequent, could involve more risk exposure, and could thus be a source of vulnerability. We could expect, though, that such exposures would more likely exist between FIs that interact regularly, and network distance might affect the speed of contagion; intermediate nodes would first have to be affected before a shock could reach subsequent nodes. A consistently high centrality measure for a certain node might also be an indication of excessive reliance on that node, from a stability or liquidity recycling perspective.

6 Results

Network statistics at the frequency of a day reflect that the LVTS is a small, dense core of the Canadian financial system. Statistics for the financial crisis, from August 2007 onwards, do not vary greatly from statistics for the entire sample period (Figure 9), which points to the stability of the LVTS network and the reliability of the payments system.

A 1 per cent threshold for the minimum T2NDC of each pair of banks translates to a daily payment value threshold of close to Can$10 million for an average pair, and reduces mean connectivity from 89.9 per cent to 69.2 (± 3.2 standard deviation) per cent, but reduces the mean total value of daily payments by only 0.03 per cent. A 40 per cent threshold reduces mean connectivity to 28.9 (± 3.5 standard deviation) per cent, and reduces mean total daily value by 6.9 per cent. Figure 7 shows that such a threshold strongly reduces links among the non-big five banks, but has only a small effect on links among the big five, because there are fewer links between the big five in the left side of the distribution. Table 2 summarizes the main measures for the 1 per cent and 40 per cent thresholds.

Figure 8 shows the effect of a range of thresholds on the means of a few main network measures, which mostly translates to vertical shifts in the lines shown in Figure 9. The threshold of 1 per cent shows the most noticeable incremental change, with connectivity falling by almost 25 per cent and degree centralization jumping by more than threefold. Because the thresholds filter out links mostly among small participants, network centrality measures reflect a hub-and-spoke characteristic that we might expect in a network with big and small participants.

It would appear that payments in the LVTS tend to be bilaterally coordinated, because reciprocity stays relatively flat across thresholds and averages 89.3 ± 2.5 per cent for the study period (Table 2). Subsequent results use a 1 per cent threshold, unless otherwise noted. The mean clustering coefficient is similarly high, at 84.3 ± 1.5 per cent, compared to other models of PCSSs described in section 3.2. The mean average path length is 1.31 ± 0.03, and the mean average eccentricity is 1.84 ± 0.07. Most of the network also stays connected, even when the threshold increases, as shown by the GSCC. Section 6.1 describes the results for intraday sampling frequencies.

Figure 8: Effect of Thresholds on Averages

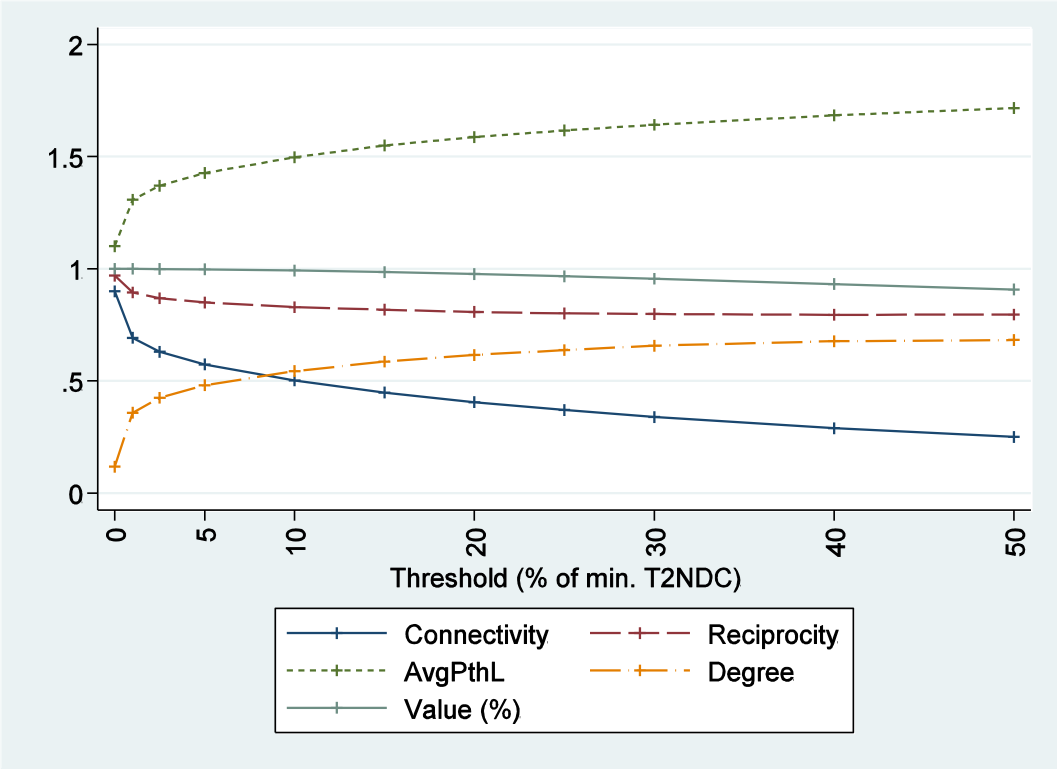

Table 2: Summary of Tranche 2 Network Measures, April 2004–December 2008[a]

Variable	Mean	Std. dev.	Min	Max
(Threshold = 1% of lower T2NDC, no holidays)				
Connectivity	0.692	0.032	0.577	0.791
Reciprocity	0.893	0.025	0.803	0.977
AvgPthL	1.308	0.032	1.209	1.423
Diameter	2.005	0.072	2	3
GSCC	13.885	0.319	13	14
Eccentricity	1.841	0.069	1.615	2.143
Clustering	0.843	0.015	0.781	0.883
Betweenness	0.062	0.017	0.023	0.134
Degree	0.358	0.037	0.244	0.481
(Threshold = 40% of lower T2NDC, no holidays)				
Connectivity	0.284	0.035	0.038	0.401
Reciprocity	0.793	0.057	0.571	0.967
AvgPthL	1.688	0.066	1.489	2
Diameter	2.718	0.505	2	4
GSCC	12.524	1.033	4	14
Eccentricity	2.184	0.266	1.800	3.385
Clustering	0.573	0.094	0	0.854
Betweenness	0.426	0.093	0.018	0.779
Degree	0.676	0.066	0.135	0.872
(No Threshold, Tranche 2)				
Total value	140.99m	23315.42	88600.148	230.49m
Total trans	19650.33	3420.00	13344	38787
Value cent.	0.300	0.017	0.231	0.343
Herfindahl	0.186	0.008	0.163	0.209

a. The Bank of Canada is not included in the network.

With no threshold, visible changes in the statistics include a small downward adjustment of connectivity, and an upward adjustment of the average path length, following the entry of State Street into the LVTS network on 18 October 2004. The effect seems to mostly disappear with a threshold, however, which is surprising in that the new participant does not change the pre-existing pattern. In Figure 9, the shift occurs only with a brief transition, after which there is a slight trend upwards for connectivity, up until late 2008, probably as financial activity generally increased.

Outliers occur on civic holidays, not shown in Figure 9, where the average path length spikes up to 2.22 and connectivity and reciprocity drop to 16.5 and 33.3 per cent on average, respectively. Diameter for any given day is nearly always two, and the daily GSCC includes all participants, with the exception of civic and U.S. holidays. The spike on 18 February 2008 corresponds to President's Day in the United States. Notwithstanding these outliers, daily network measures test as covariance stationary, although they exhibit significant partial autocorrelations, to varying degrees. In other words, there is a persistence to variations that occur in the network structure in the short term, even while there is never a great departure from the long-run average.

Figure 9: Network Statistics (1% Threshold of T2NDC), April 2004–December 2008

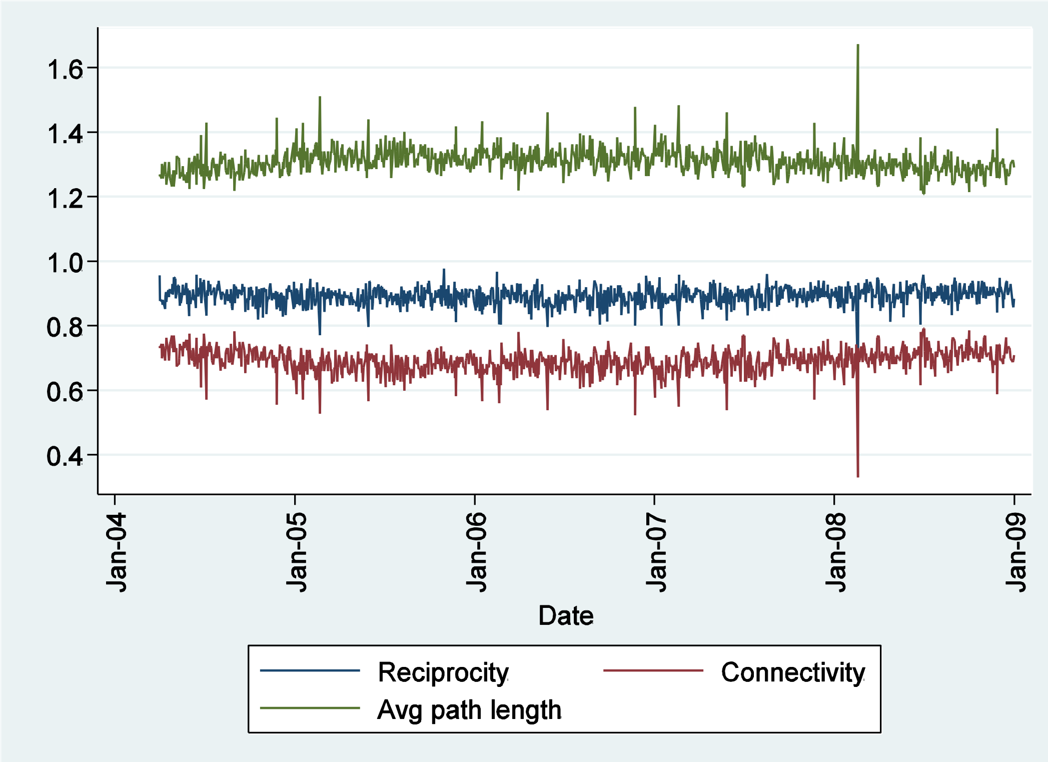

Correlations between network statistics tell us how certain network properties relate to one another, possibly leading to ideas about the nature of interbank relationships. In many cases, however, a common construction from the same underlying data could influence a correlation. For example, additional links increase connectivity; if the LVTS is already highly connected, reciprocity would also likely increase, while path lengths would decrease, since additional links add to the density of connections. If the network is sparse, we should expect path lengths to increase and reciprocity to fall as new links connect peripheral nodes that do not yet connect to the GSCC.[30] With a daily sampling interval, the former case of high connectivity materializes.

Table 3 summarizes correlations at a daily frequency with a 1 per cent threshold for the 2004–08 sample period.[31] Connectivity shows correlations of 0.302 with total value and 0.476 with the total volume of transactions. Value centralization is uncorrelated with network measures at a 1 per cent threshold, including non-weighted centralization measures, but correlates positively with total value and volume. With a 40 per cent threshold, connectivity shows correlations of 0.139 with reciprocity, -0.352 with average

[30] These facts remind us of the importance of context in interpreting correlations, which might not be constant with respect to the values of the variables.
[31] This table excludes outlying civic holidays and U.S. holidays, which does affect certain correlations, but might provide a better indication of the normal relationships that exist.

path length (which still confirms the expected signs of a highly connected network where additional links do not expand the network but, instead, fill in missing links), 0.645 with total value, and 0.584 with total transactions. Even at a 40 per cent threshold, connectivity negatively correlates with centrality. Reciprocity also correlates positively with total value and volume.

Table 3: Correlations, Excluding Civic and U.S. Holidays, April 2004–December 2008

	Conn.	Recip.	AvgPthL	Betw.	Value[a]	Volume[a]
1% T2NDC Threshold						
Connectivity	1					
Reciprocity	0.37	1				
AvgPthL	0.9999	-0.3703	1			
Betweenness	0.8052	-0.361	0.8045	1		
Total value[a]	0.3021	0.2108	-0.3023	0.4381	1	
Total volume[a]	0.4758	0.2631	-0.4758	-0.5253	0.7864	1
Value cent.[a]	0.0137	-0.0179	0.0142	0.0515	0.2565	0.1737
40% T2NDC Threshold						
Connectivity	1					
Reciprocity	0.1389	1				
AvgPthL	0.3523	-0.1715	1			
Betweenness	-0.184	0.0174	0.3903	1		
Total value[a]	0.6448	0.2368	-0.2547	0.0689	1	
Total volume[a]	0.5843	0.1512	-0.1928	-0.034	0.7864	1
Value cent.[a]	0.1299	0.1841	0.0492	0.3193	0.2565	0.1737

[a] No threshold is imposed.

These correlations confirm our intuition that a higher volume of transactions corresponds to more links in the network, and consequently shorter path lengths. The association of higher total value with more value centralization could relate to higher variability in the activity of the major banks, which account for the majority of value. Higher-value liquidity movement might tend to take place bilaterally, rather than circularly, through a larger number of connected participants; this would involve easier coordination, and is a tentative explanation for the positive correlation of reciprocity and value. The intraday situation provides another case to consider, and reinforces the daily result that indicators of increased payment activity typically increase with one another, when compensating for the time-of-day effect.

Fedwire connectivity is strongly positively correlated with value and volume, while reciprocity is not significantly correlated with value, volume, or connectivity (Soramäki et al. 2006), which points to a case where the bilateral coordination of payments does not vary with these other characteristics, in contrast to the LVTS. In Fedwire, reciprocity also correlates negatively with network size, possibly because non-reciprocal links might often expand the network.

6.1 Intraday sampling frequencies

Along with the stability implied by daily network measures, Figures 10 and 11 show the few trends and shifts in intraday activity for one-month moving averages.[32] Connectivity during the 07:00 interval declined noticeably from 2004 to 2008, while connectivity and reciprocity during the 08:00 interval increased, reflecting a shift in morning payment activity. During other times of the day, there was a slight increase, as financial activity generally increased. In Figure 11, average path length has declined slightly, but especially between 07:00 and 09:00. Total value rose during most parts of the day, but declined between 16:00 and 18:00 in 2008, after increasing from 2005 to 2007. This development is arguably good from a social perspective, since it means that there is less delay and vulnerability to an end-of-day interruption of payment settlement due to an operational event. We might have instead expected that participants would push a greater share of their payments towards the end of the day in 2008, because of uncertainty and concerns about liquidity during the financial crisis.

Figure 10: Intraday Connectivity and Reciprocity, April 2004–December 2008

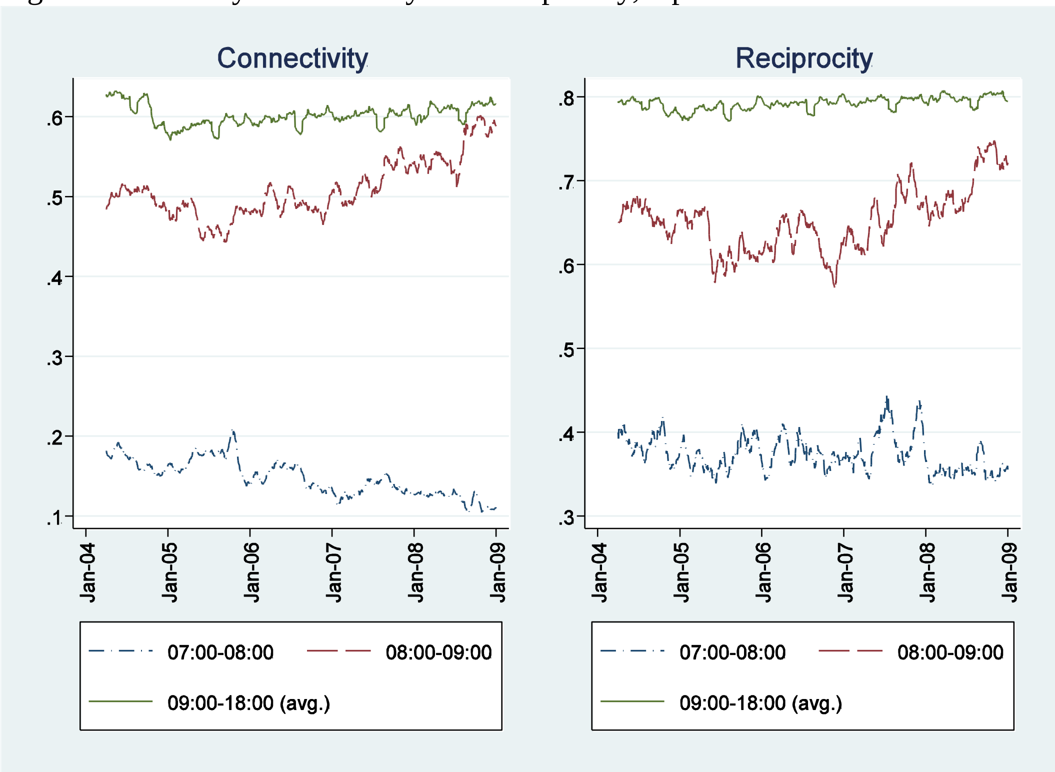

[32] Without averaging, there would be too much variability for trends to be visible in the measures.

Figure 11: Other Intraday Network Stats, April 2004–December 2008

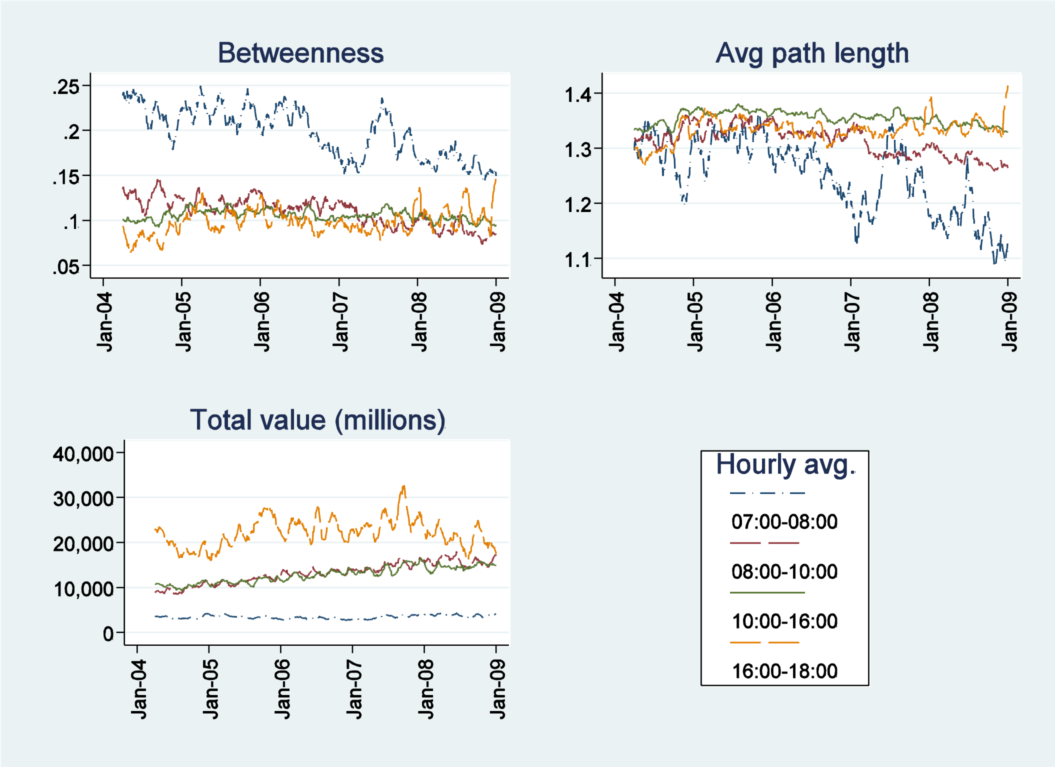

For most network measures, different frequencies depict a comparable intraday evolution of network structure, as Figure 12 shows with the example of January 2007. In the cases of diameter, average path length, and betweenness centralization, the hourly frequency captures a different pattern, with higher values near the start and end of the general LVTS payment exchange, which is not evident from the other frequencies.[33]

[33] The general payment exchange of the daily cycle occurs between 06:00 and 18:00.

Figure 12: Intraday Network Structure, January 2007

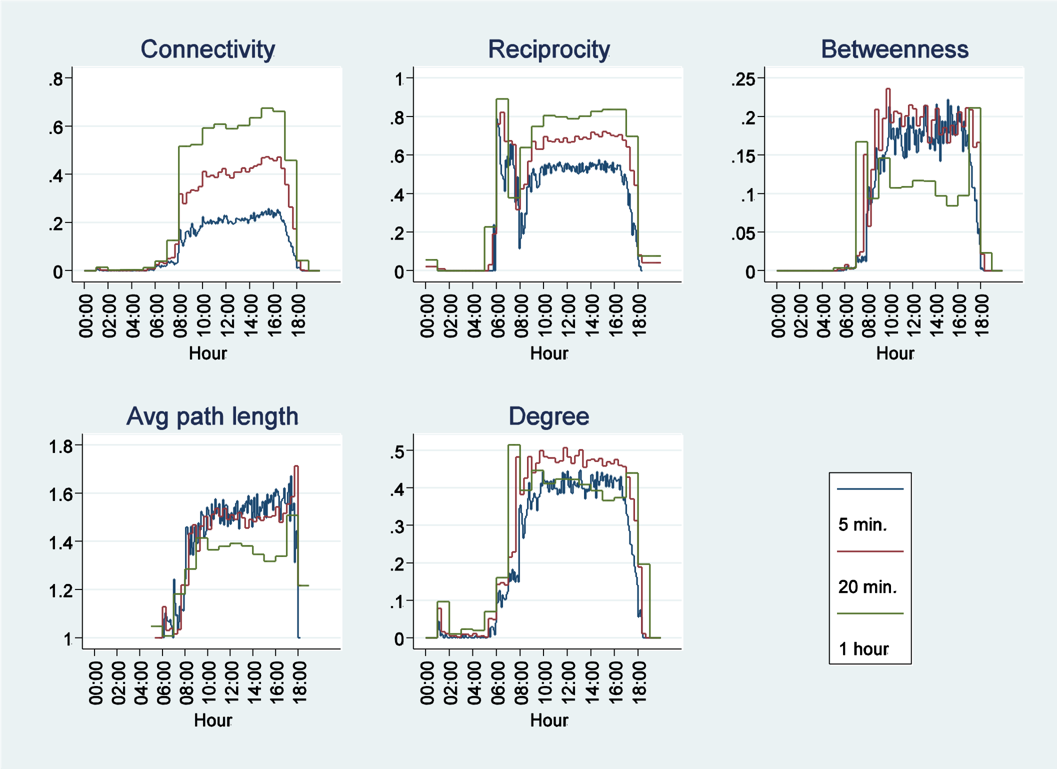

Figure 12 also shows that connectivity strictly increases as a sampling interval becomes longer, while reciprocity instead shows the coordination payments, because it remains relatively high regardless of frequency, especially between 06:00 and 08:00. Earlier in the day, there could be a higher liquidity cost to not coordinating payments; higher reciprocity at the start of day is consistent with this.

6.2 Centralization

Identifying which banks are important in a financial system, and why, is a central theme for financial overseers. Centrality, in network analysis, attempts to capture a bank's importance for network interconnectedness. Although, in principle, this might not necessarily correspond to size, the correlation between a bank's betweenness centrality and its share of total gross LVTS payments value is 0.869, while for degree centrality the correlation is 0.756.

There are no dramatic changes in overall network centralization for the middle part of the day from 10:00 to 16:00 across the sample period of April 2004–December 2008, as can be seen in Figure 11 for betweenness. Centralization between 07:00 and 10:00 has declined, while in 2008 there was a rise between 16:00 and 18:00. For a given day, betweenness and degree centralization are higher in the morning period from 07:00 to

10:00 and near the end of normal operations during the 17:00 interval, shown in Figure 12. This suggests that the prominent role of key participants is somewhat more pronounced at certain times of the day, namely near the start of general payment exchange and in the late afternoon.

Figure 13: Betweenness Centrality, 1% T2NDC Threshold, April 2004–December 2008

Figure 14: Betweenness Centrality, 40% T2NDC Threshold, April 2004–December 2008

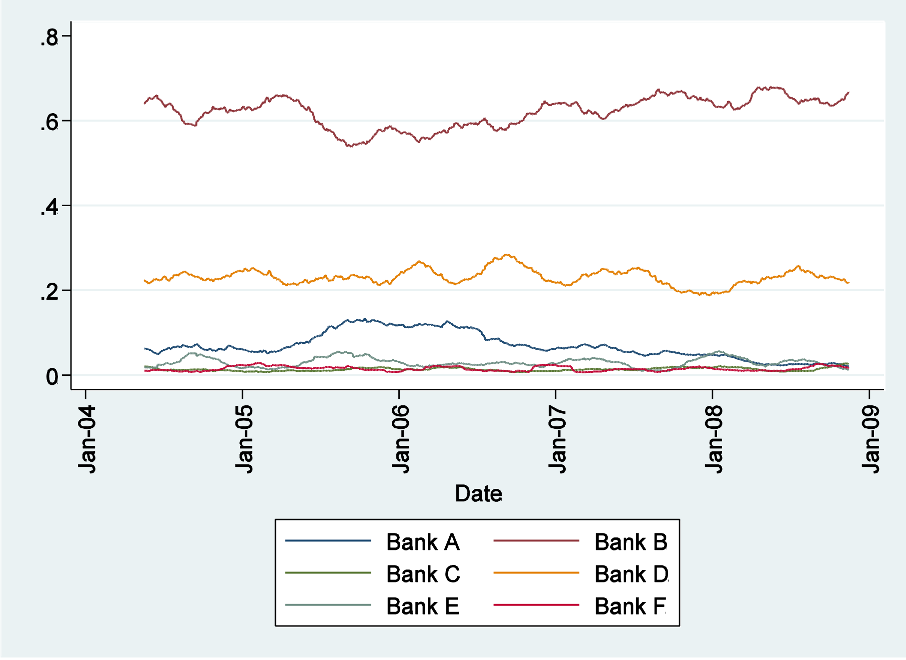

Shifts in the centrality of individual participants depend on the threshold; Figure 13, which filters out only the least significant links, shows a convergence among the five most central of the big six banks. On the other hand, as Figure 14 shows, a higher (40 per cent) T2NDC threshold emphasizes the network centrality of Banks B and D, while the other four banks converge together towards a much lower centrality.[34]

The removal of a big six bank from the network can have an effect on average connectivity ranging from a reduction of about 15 per cent to over 45 per cent, at a 40 per cent threshold. This assumes that there are no knock-on effects and that participants do not adapt; for example, indirect participants would ordinarily find another direct participant to tier through.[35] Reciprocity varies much less across the bank removal scenarios, again reflecting the bilateral coordination of payments. The effect of the removal of a big six bank on the number of banks that are a part of the GSCC ranges from a reduction of about 15 per cent to over 30 per cent.

Simulation studies of the LVTS have shown that the failure of a single participant is unlikely to create losses within the payments system that threaten the solvency of surviving participants (Ball and Engert 2007). From a broader financial network perspective, the failure of a major bank would undoubtedly have a heavy impact, although we lack a reference point to say whether the impact could be of systemic importance. Figure 15 shows that, in comparison with a failure of the most connected bank, a failure of at least two of the other big six banks, or of six or seven of the smallest participants, would have to take place to have a comparable effect on network connectivity, for the study period. This could overstate the importance of the smaller participants, since a 40 per cent threshold filters lower-value links, but otherwise does not take account of the value of different links.

[34] Other methods could also determine centrality, such as that of Bech, Chapman, and Garratt (2008).
[35] In this exercise, the payments to and from the removed participant are set to zero in order to observe a base case for how the network characteristics might change.

Figure 15: Effect of Removing Banks on Connectivity (40% Threshold)

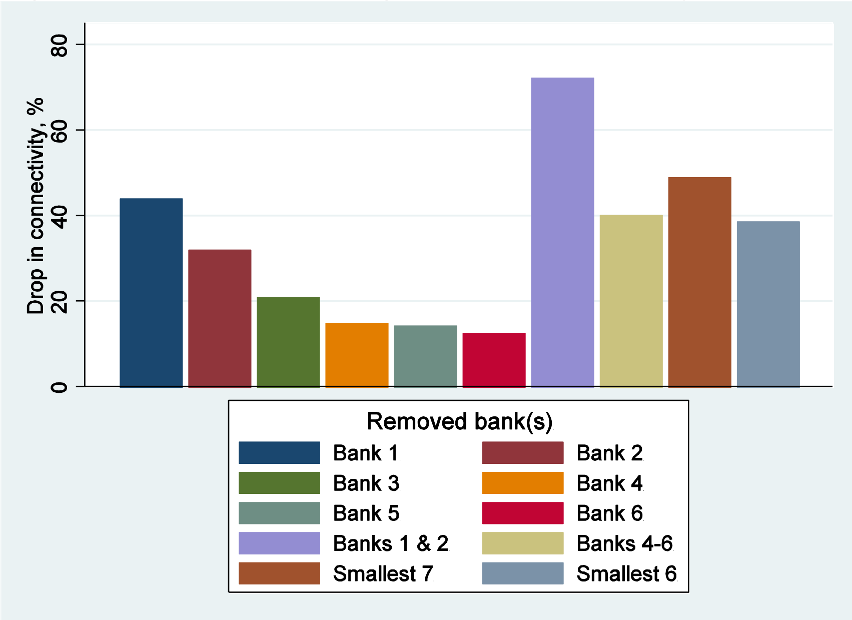

6.3 Regression analysis

We know about the importance of network structure for stability and resiliency, so a further goal is to empirically learn more about structure determinants. Our simple approach using regressions of daily network connectivity and betweenness centralization demonstrates the relevance of calendar effects, and LVTS value and volume; however, nothing conclusive can be said about any economic variables.

One idea that relates to the strategic models we saw earlier is that, if there were higher general costs to developing and maintaining a financial relationship, we would expect fewer links. Credit risk, reflected in market spreads, is such a cost, and could reduce a bank's willingness to conduct business with some FIs. If some banks had an advantage with regard to the asymmetric information involved in credit risk, the sparser network could tend more strongly towards a hub-and-spoke structure. The easing of both the Bank of Canada overnight rate and the federal funds rate also coincided with increased credit risk concern, which this policy action was meant to counteract.

In a model regressing connectivity on calendar effects, the overnight rate, the federal funds rate, logs of LVTS value and volume, the spread between one-month bankers' acceptances and treasury bills, and a time trend, the results do not convincingly show a strong influence from any financial variables, once the time trend is included.[36] Thresholds above 15–20 per cent, which reveal more economically important links, give the most promising results; for example, the spread between one-month bankers' acceptances and treasury bills is sometimes statistically significant, but not consistently in line with our hypothesis. Models that regress individual bank centralities on credit default swap spreads, the bank's share of payments, and a time trend are similarly inconclusive.

Factors that affect network structure could be hard to measure – for example, the historical and efficiency advantages of individual banks. Further problems include the aggregation of different kinds of payment activity, the coarseness of network measures, and delays between economic stimuli and transaction settlement times, which would impede our ability to find statistical relationships.

Figure 16: Economic Variables, April 2004–December 2008

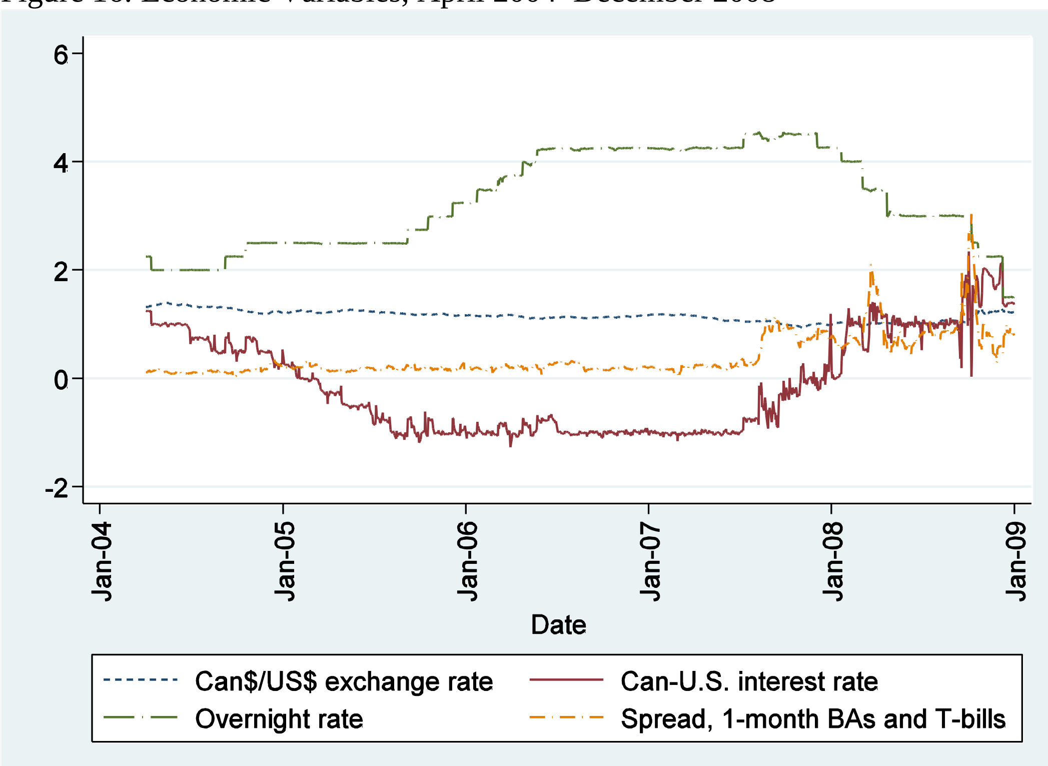

[36] Figure 16 shows general trends, with the increase in the overnight rate from 2.50 per cent in September 2005 to 4.25 per cent in May 2006, the gradual appreciation of the Canadian dollar from $1.396 Can$/US$ in May 2004 to $0.922 in November 2007, and the increase in the federal funds rate from 1 per cent in April 2004 to 5.25 per cent in July 2005, reflected in the decline of the Canada/U.S. interest rate differential. Table 4 summarizes the main statistics of the economic regressors.

Table 4: Summary of Economic Variables, April 2004–December 2008

Variable	Mean	Std. dev.	Min	Max
Overnight rate	3.295	0.881	1.491	4.552
Spread, overnight rate, and fed rate	0.123	0.903	-1.27	2.342
Fed rate	3.419	1.578	0.09	5.41
Fed rate, 1st diff.	0.001	0.123	-0.95	1.05
Overnight rate, 1st diff.	0.001	0.043	-0.746	0.262
Spread, 1-month T-bills and BAs	0.391	0.39	0	3.03
Spread, 1 month, 1st diff.	0.001	0.073	-0.72	0.8
Can$/US$ FX rate	1.149	0.103	0.917	1.397

7 Discussion

At a basic level, networks can allow an intuitive understanding and provide a graphical depiction of how the financial system trades off stability against efficiency. Theory has diverging views over whether complete networks of interbank claims withstand disruptions better than incomplete or centralized hub-and-spoke structures (Allen and Gale 2000; Haldane 2009). Ultimately, the influence of network structure could depend on the type of shock: a complete network might easily absorb an isolated disruption by rerouting its activity, but could succumb to a larger disruption through a chain of feedback effects. We have seen from intraday frequencies or daily activity above a given threshold that payment flows in the LVTS have a network structure highly centralized in a small number of key participants, but the network still retains substantial complexity. The Canadian financial system could therefore be vulnerable to either a specific bank failure or feedback-driven contagion, to the extent that the LVTS reflects the general financial network.

In other words, stability can have different meanings: a financial system could be resilient to small disruptions even when it is vulnerable to low-probability, high-impact events, or vice versa. If a hub-and-spoke financial system has a strong bank in the centre, most disruptions will have no chance to propagate through the system. If the bank in the centre were to succumb to one of these disruptions, however, it would likely bring down the rest of the system.

Financial system regulators and overseers will want to know what PCSS networks could mean for policy. While payments systems have remained stable and robust throughout the credit crisis, owing to their sound risk controls, network centrality measures might be useful for assessing participant systemic importance in a more general sense, for the financial system. The network effect of removing a bank or a combination of banks further illustrates this potential aspect of a disruption. Even though the consequences of a bank failure are unknown in advance, interconnectedness complicates the problem, such that a relative ranking of bank centralities would help prioritize central bank policy responses in the case of a future financial crisis. The viability of such an analysis would depend on the quality and breadth of individual bank financial data available to a central bank.

The social networks literature explores different technical conditions and how they might result in a range of network equilibria, but it does less to address the case of how agents might adjust to persistent shocks. Policy-makers for PCSSs, though, are likely more concerned about events that threaten to derail a financial system, in the form of bank failures and catastrophic operational events. Arguably, the social networks literature uses a stylization of networks that might not always be pertinent to a PCSS. The generality of some of these papers and lack of specific context could be contributing factors. The PCSS network-related literature, on the other hand, does identify pertinent trade-offs relevant to its context.

The literature emphasizes that network structure is a key determinant of the payoffs ensuing to an agent (Jackson and Watts 2002). This conclusion should extend to networks in PCSSs: a bank's role in the network could affect its access to credit and information, its ability to financially intermediate, and its risk exposure. A small FI that links to a large well-connected bank might gain something from the indirect links, such as easier financial market access. For example, a small FI in Canada can use correspondent banking with a direct LVTS participant as a relatively cost-effective means of transacting through CLS or with other direct LVTS participants and their respective correspondent bankees. Possibly, some of the large bank's benefit from connectedness could residually benefit the smaller bank. The tiered model of Chapman, Chiu, and Molico (2008) illustrates another payoff of indirect relationships, in that small agents without a public credit history benefit from settling through a clearing agent.

The regression models are inconclusive regarding the idea that higher credit-risk concern may have reduced links in the network. Neither can fluctuations in an individual bank's centrality be attributed to changes in its credit default swap spread. A direct connection might, in any event, be difficult to detect with respect to LVTS data, because of data aggregation, lack of information about payment rationale, or the time that could elapse between a payment request or a financial stimulus and the settlement of a consequent payment.

Network analysis still remains abstract. For network measures such as centralization in the LVTS, their validity partly comes down to what seems "right," leaving us with a task of interpretation that is not much simpler than before, although neither should we overinterpret such stylized measures. The implications of network measures also depend on the links that are studied; for example, the interpretation of a network of payments could be different from one of loans. For the case of paths and distances, it seems reasonable to argue that an existing network path opens the door to financial contagion, but it is less certain that a given network representation could help us to quantify this contagion risk.

The idea that reciprocity relates to payment coordination is supported by its positive correlation with the ratio of payments that net bilaterally, relative to total gross payments. The gradual daily buildup in the connectedness of the network, meanwhile, as shown in Figure 12, only hints at the concentration of payment value that takes place towards the end of the day. Regardless, connectivity conveys a different kind of information – namely, the extent to which participants engage with a range of counterparties in the system.

Network analysis uses arbitrary criteria for defining a link; weighted links provide an alternative, but could come at the cost of simplicity. For example, it would be difficult to define network distance based on weighted links.[37] Each sampling frequency is also static and does not relate to periods before or after. The choice of sampling frequency is subjective and the meaning of a relationship might not relate entirely to the frequency of interaction.

More fundamentally, the rationale for payment activity is not identified in LVTS payment flows data, and, as a consequence, network statistics can appear to share this abstraction. Knowing the actual network of financial exposures, if possible, would provide us with a better understanding of the financial system, but these data are often not readily available. Finally, LVTS data alone cannot tell us about the network of the broader Canadian financial system. Nevertheless, Jackson (2006) concludes that "the wide variety of settings where network structure is an important determinant of behaviour makes it clear that this is one of the most wide open and important areas for further study."

7.1 Future empirical research

Network measures could provide a means to determine whether operational events and other major disruptions have a persistent impact on participant interactions. Alternatively, the network for specific transaction types that are of particular importance might give a clearer view of key relationships. Furfine (1999) and Hendry and Kamhi (2007), for example, are able to identify and characterize the markets for uncollateralized overnight loans, in Fedwire and the LVTS, respectively. Their methodologies find pairs of

[37] Wasserman and Faust (1994) briefly discuss weighted measures, which are less common in the literature than measures where links are unweighted.

transactions that could plausibly represent loans and next-day repayments, based on whether the implied interest rate is reasonable. Hendry and Kamhi (2007) also identify transactions that likely represent sales/purchases of U.S. dollars. A network analysis could be applied to such filtered data, in order to focus on specific markets.

A focus on the network of specific transaction types of particular importance might also give a better understanding of its liquidity. Liquidity is often proxied by trading volumes and bid/ask spreads, variables that might not provide reliable predictions for changes to liquidity due to a market disruption. In contrast, a network approach can help to emphasize the potential availability of counterparties for a trade. A network approach might also be a useful input in evaluating the need for central counterparty services for a given security type, in a securities clearing and settlement system such as CDSX.

LVTS data cannot tell us about financial linkages involving indirect participants; arguably, this is an obstacle to our understanding of the financial system and a challenge to determining the systemic importance of banks. The Bank of England has made use of the 2003 Clearing House Automated Payment System (CHAPS) Traffic Survey, which consists of a sample of CHAPS payments for five days in February 2003. For each payment in the sample, the data contain the value and purpose of the payment, the day and the time at which it was sent, and codes identifying the bank sending the payment and the bank receiving the payment (whether or not they are direct members of CHAPS). Codes also identify the direct members that execute the payment. A comparable survey in Canada could be beneficial for understanding the Canadian financial system.

8 Conclusions

By examining participant interactions and identifying key participants, network analysis can provide insight into a payments system's efficiency, stability, and resiliency. It allows us to summarize the complex web of interrelationships between participants, consisting, for example, of payments, loans, and credit lines. In this paper, we introduce the PCSS-related network literature and describe the network structure of LVTS payment flows, in an effort to comment on the aforementioned issues.

The literature highlights that network characteristics play an important role in the outcomes observed in PCSSs, such as payment coordination, or the system parameters that achieve social optimality, which helps motivate a need to understand PCSS networks. The theoretical literature also indicates how network formation can be modelled. Some of these models – in particular, strategic formation models – might be useful for understanding how a network changes in response to a shock.

Describing a PCSS network is a first empirical step for understanding, and potentially modelling, the network. The network structure of the Canadian LVTS represents a relatively small, densely connected core. Thresholds and an intraday analysis reinforce the fact that only a few LVTS participants act as highly connected hubs, a pattern similar to the LVPSs of other countries. The resiliency of network hubs could be a decisive factor in the promotion of financial stability, since these hubs intermediate many pathways that system liquidity, and possibly contagion, could follow. Additionally, the removal of a key participant can have an effect on network connectivity comparable to the removal of several other participants, which would seem to support the possibility of systemic importance. The LVTS network remained stable throughout the credit crisis, though, and regression modelling does not show any effect from financial variables that relate to credit-risk concern.

Progress in network analysis will both contribute to, and benefit from, more clarity on what questions it can feasibly answer or inform with data. Limitations remain, partly due to the arbitrary nature of many network measures. Suggested areas for future empirical research include an examination of networks for important transaction types such as overnight loans, and research on the broader Canadian financial system, should survey data become available.

References

Albert, R., H. Jeong, and A.-L. Barabási. 2000. “Error and Attack Tolerance of Complex Networks.” *Nature* 406: 378–82.

Allen, F. and D. Gale. 2000. “Financial Contagion.” *Journal of Political Economy* 108 (1): 1–33.

Arjani, N. and D. McVanel. 2006. “A Primer on Canada’s Large Value Transfer System.” Bank of Canada. Available at <http://www.bankofcanada.ca/en/financial/lvts_neville.pdf>.

Bala, V. and S. Goyal. 2000. “A Noncooperative Model of Network Formation.” *Econometrica* 68 (5): 1181–1229.

Ball, D. and W. Engert. 2007. “Unanticipated Defaults and Losses in Canada’s Large-Value Payments System, Revisited.” Bank of Canada Discussion Paper No. 2007-5.

Barabási, A. and R. Albert. 1999. “Emergence of Scaling in Random Networks.” *Science* 286: 509–12.

Bech, M., J. Chapman, and R. Garratt. 2008. “Which Bank is the ‘Central’ Bank? An Application of Markov Theory to the Canadian Large Value Transfer System.” Bank of Canada Working Paper No. 2008-42.

Bech, M. and R. Garratt. 2003. “The Intraday Liquidity Management Game.” *Journal of Economic Theory* 109 (2): 198–219.

———. 2006. “Illiquidity in the Interbank Payment System Following Wide-Scale Disruptions.” Federal Reserve Bank of New York Staff Reports No. 239.

Bloch, F. and M. Jackson. 2006. “Definitions of Equilibrium in Network Formation Games.” *International Journal of Game Theory* 34 (3): 305–18.

Borgatti, S. P. 2005. “Centrality and Network Flow.” *Social Networks* 27 (1): 55–71.

Boss, M., H. Elsinger, M. Summer, and S. Thurner. 2003. “The Network Topology of the Interbank Market.” Santa Fe Institute Working Paper No. 03-10-054.

Brueckner, J. K. 2004. “Friendship Networks.” Paper No. 184 presented at the 2004 North American Winter Meetings held by the Econometric Society, San Diego, 3–5 January 2004.

Butts, C. T. 2009. “sna: Tools for Social Network Analysis.” Contributed Packages. The Comprehensive R Archive Network.

Chapman, J., J. Chiu, and M. Molico. 2008. “A Model of Tiered Settlement Networks.” Bank of Canada Working Paper No. 2008-12.

Chapman, J. T. E. and N. Y. N. Zhang. 2009. “Estimating the Structure of the Payment Network in the LVTS: An Application of Estimating Communities in Network Data.” Paper presented at the Canadian Economics Association 43rd Annual Conference, University of Toronto, 28–31 May 2009.

Committee on Payment and Settlement Systems (CPSS). 2005. “New Developments in Large Value Payment Systems.” CPSS Publication No 67. Basel: Bank for International Settlements.

———. 2008. “Statistics on Payment and Settlement Systems in Selected Countries.” CPSS Publication No. 82. Basel: Bank for International Settlements.

Crucitti, P., V. Latora, M. Marchiori, and A. Rapisarda. 2004. “Error and Attack Tolerance of Complex Networks.” *Physica A* 340 (1–3): 388–94.

Erdös, P. and A. Rényi. 1960. “On the Evolution of Random Graphs.” *Publications of the Mathematical Institute of the Hungarian Academy of Sciences* 5: 17–61.

Financial Times. “Rescuing AIG.” Lex column. 18 September 2008.

Furfine, C. H. 1999. “The Microstructure of the Federal Funds Market.” *Financial Markets, Institutions and Instruments* 8 (5): 24–44.

Gale, D. M. and S. Kariv. 2007. “Financial Networks.” *American Economic Review* 97 (2): 99–103.

Haldane, A. G. 2009. “Rethinking the Financial Network.” Speech delivered at the Financial Student Association, Amsterdam, the Netherlands, 28 April.

Hendry, S. and N. Kamhi. 2007. “Uncollateralized Overnight Loans Settled in LVTS.” Bank of Canada Working Paper No. 2007-11.

Herings, J. J., A. Mauleon, and V. Vannetelbosch. 2006. “Farsightedly Stable Networks.” Maastricht Research School of Economics of Technology and Organizations (METEOR) Discussion Paper No. 2006-46.

Inaoka, H., T. Ninomiya, K. Taniguchi, T. Shimizu, and H. Takayasu. 2004. “Fractal Network Derived from Banking Transaction – An Analysis of Network Structures Formed by Financial Institutions.” Bank of Japan Working Paper Series No. 04-E-04.

Jackson, M. O. 2006. “The Economics of Social Networks.” In *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, volume 1, edited by R. Blundell, W. Newey, and T. Persson, 1–56. Papers and commentary presented at the

9th World Congress of the Econometric Society, London, England, 19–24 August 2005. Cambridge: Cambridge University Press.

Jackson, M. O. and B. W. Rogers. 2005. "The Economics of Small Worlds." *Journal of the European Economic Association* 3 (2–3): 617–27.

———. 2007. "Meeting Strangers and Friends of Friends: How Random Are Social Networks?" *American Economic Review* 97 (3): 890–915.

Jackson, M. O. and A. Watts. 2002. "The Evolution of Social and Economic Networks." *Journal of Economic Theory* 106: 265–95.

Jackson, M. O. and A. Wolinsky. 1996. "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* 71: 44–74.

Klemm, K. and V. M. Eguíluz. 2002a. "Growing Scale-Free Networks with Small-World Behavior." *Physical Review E* 65 (5): 1–4.

———. 2002b. "Highly Clustered Scale-Free Networks." *Physical Review E* 65 (3): 1–5.

Labelle St-Pierre, N. 2009. "When Do Operational Events Become a Systemic Concern: An Agent-Based Model of the LVTS." Paper presented at the Agent-Based Modeling for Banking and Finance workshop in Torino, Italy, 9–11 February 2009.

Martin, A. and J. McAndrews. 2007. "Liquidity Saving Mechanisms." Federal Reserve Bank of New York Staff Reports No. 282.

Newman, M. E. J., S. H. Strogatz, and D. J. Watts. 2001. "Random Graphs with Arbitrary Degree Distributions and their Applications." *Physical Review E* 64 (2): 1–17.

Newman, M. E. J. and D. J. Watts. 1999. "Scaling and Percolation in the Small-World Network Model." *Physical Review E* 60 (6): 7332–42.

Page, F. H. Jr., M. Wooders, and S. Kamat. 2005. "Networks and Farsighted Stability." *Journal of Economic Theory* 120 (2): 257–69.

Soramäki, K., M. L. Bech, J. Arnold, R. J. Glass, and W. E. Beyeler. 2006. "The Topology of Interbank Payment Flows." Federal Reserve Bank of New York Staff Reports No. 243.

von Peter, G. 2007. "International Banking Centres: A Network Perspective." *BIS Quarterly Review* (December): 33–45.

Walsh, M. W. and M. J. de la Merced. "Regulators Scramble to Head Off Collapse of AIG." *International Herald Tribune*, 17 September 2008.

Wasserman, S. and K. Faust. 1994. *Social Network Analysis: Methods and Applications.* Cambridge: Cambridge University Press.

Watts, D. J. and S. Strogatz. 1998. "Collective Dynamics of 'Small-World' Networks." *Nature* 393: 440–42.

Appendix: Definitions of Network Statistics

This appendix briefly outlines formulas that derive network statistics. For more information, see Soramäki et al. (2006), Jackson (2006), or Butts (2009).

A.1 Connection-Based Measures

Connectivity is $c = \frac{m}{N(N-1)}$ and reciprocity is $r = \frac{m_r}{m}$,

where

- N is the number of nodes
- m is the number of directed links
- m_r is the number of reciprocated links

The clustering coefficient is $\Gamma_i = \frac{m_{nn,i}}{k_i(k_i-1)}$,

and the network's average clustering coefficient is $\sum_i \Gamma_i$,

where

- k_i is the number of neighbours of node i
- $m_{nn,i}$ is the number of directed links among neighbours of node i

A.2 Distance-Based Measures

The distance d_{ij} is the length of the shortest path between i and j and can be computed from the power matrices M^p of the adjacency matrix M that represents the existence of links between nodes, with $p \in \mathrm{I}, 1 \leq p \leq N-1$. The distance d_{ij} is the value of p where the matrix element $(M^p)_{ij}$ first becomes non-zero.

The average path length of a node is $l_i = \frac{1}{N-1}\sum_{j,j\neq i} d_{ij}$. The average path length of a network is $L = \sum_i l_i$.

Eccentricity is $\varepsilon_i = \max_j d_{ij}$.

Diameter is $D = \max_i \varepsilon_i$.

A.3 Centrality and Centralization

Note that centrality measures are often rescaled so that the centralities of all nodes in the network sum to one.

Degree centrality is $C_D(n_i) = \frac{d(n_i)}{N-1}$,

where

- $d(n_i)$ is the degree (either in-, out-, or the average of both) of node n_i

Closeness centrality is $C_C(n_i) = \dfrac{(N-1)}{\sum_i \sum_{j, j \neq i} d_{ij}}$,

where

- d_{ij} is the distance between nodes *i* and *j*

Betweenness centrality is $C_B(n_v) = \dfrac{\sum_{i, i \neq v; } \sum_{j, j \neq i, j \neq v} \dfrac{g_{ivj}}{g_{ij}}}{(N-1)(N-2)}$,

where

- g_{ij} is the number of shortest paths from *i* to *j*
- g_{ivj} is the number of shortest paths from *i* to *j*, through *v*

The centralization measure is $\varsigma_{(C_Z)} = \dfrac{\sum_i \left| \max_v C_Z(n_v) - C_Z(n_i) \right|}{TMAXDEV_{C_Z}}$,

where

- $C_Z(n_i)$ is the chosen centrality measure, for node *i*
- $TMAXDEV_{C_Z}$ is the theoretical maximum deviation of the numerator, conditional on *N*, used to normalize the centralization measure