BANK OF CANADA
BANQUE DU CANADA

Discussion Paper/Document d'analyse
2011-7

The Impact of Operational Events on the Network Structure of the LVTS

by Tom Roberts

www.bank-banque-canada.ca

Bank of Canada Discussion Paper 2011-7

August 2011

The Impact of Operational Events on the Network Structure of the LVTS

by

Tom Roberts

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

ISSN 1914-0568

Acknowledgements

I would like to thank James Chapman, Miguel Molico and Nicholas Labelle St-Pierre for useful comments and suggestions, and Geoffrey Hill for research assistance. A special thank you goes to Christian Belisle for his knowledge of operational events, and related participant behaviour, in the LVTS. Additional thanks goes to Mike Hoganson of the Canadian Payments Association.

Abstract

The author uses a quantitative network analysis approach to assess how participants in the Large Value Transfer System (LVTS) respond to partial outages at other banks. Despite the limited number of operational events, benchmarks can be established. For example, the effect of a partial outage at a big six bank that has a net payment balance of $1 billion is estimated to correspond to a decline in connectivity of about 6 or 7 per cent, due to other participants not sending payments to the problem participant. This suggests that participants tend to perceive a partial outage at a counterparty as inconvenient enough to warrant a delaying of payments, at least to the problem participant. There is no strong evidence of systemic effects, whereby participants also delay payments to non-problem counterparties, in an effort to maintain a reasonable liquidity position. Notable events mostly occurred in 2004 or 2005, and the incidence of operational events did not increase over the sample period. The improvement of contingency measures, such as the use of the LVTS Direct Network, possibly contributed in this regard.

JEL classification: C49, G14, G21
Bank classification: Payment, clearing, and settlement systems

Résumé

L'auteur utilise une approche quantitative d'analyse des réseaux afin d'évaluer comment les institutions qui font partie du Système de transfert de paiements de grande valeur (STPGV) réagissent en cas de panne partielle chez d'autres participants. Malgré le nombre limité d'incidents opérationnels à s'être produits, il est possible d'établir des balises. À titre d'exemple, on estime que l'effet d'une panne partielle dans une des six grandes banques canadiennes qui aurait un solde net de paiement de un milliard de dollars correspond à une perte de connectivité d'environ 6 ou 7 %, attribuable à l'interruption des paiements envoyés à cette institution par les autres. Un tel résultat indique que ce type d'incident est généralement assez préoccupant aux yeux des participants pour justifier un report des paiements, du moins à la contrepartie en cause. Rien ne laisse valablement supposer la présence d'effets systémiques, qui amèneraient les participants à différer leurs paiements même à ceux qui fonctionnent normalement, dans un effort pour maintenir l'équilibre de leur situation de trésorerie. Les incidents opérationnels d'importance sont survenus pour la plupart en 2004 ou en 2005, et leur incidence n'a pas augmenté au cours de la période étudiée. L'amélioration des mesures de contingence, comme le recours au réseau direct du STPGV, n'est peut-être pas étrangère à cette évolution.

Classification JEL : C49, G14, G21
Classification de la Banque : Systèmes de paiement, de compensation et de règlement

1 Introduction

The Bank of Canada has previously studied the systemic effects of a participant outage within the Large Value Transfer System (LVTS), through simulation exercises. The LVTS is of particular interest because it is part of the core of Canada's financial infrastructure, and eases the implementation of monetary policy.[1] A limitation of simulation exercises, however, has been that they often do not account for adaptive participant behaviour, or instead make ad hoc assumptions in this regard. This study uses a quantitative network analysis approach to assess how LVTS participant behaviour, as reflected by LVTS network structure, has changed when encountering a disruption in the system. Network analysis has the advantage of giving a summarized, system-wide perspective, rather than a disorganized view of a multitude of individual actions.

Operational events, and specifically partial outages, could have the immediate effect of compromising payment activity, especially with respect to the problem participant.[2] A partial outage denotes an event where the participant can receive payments but is either unable to send payments or experiences a slowdown in its ability to do so. This situation creates the risk of a liquidity trap, and is the most common type of participant outage, since payments are cleared and settled in the LVTS, but they are sent from a participant's facilities − unless a participant is offline or suspended, the LVTS will continue to process payments sent to the participant.[3] There might also be a system-wide effect of reduced payments activity, if participants become concerned about their overall liquidity position.

This paper examines partial outage effects on payment activity, focusing on a network perspective. It is a natural extension of Embree and Roberts (2009), which described and analyzed the network structure of LVTS payment activity. Network analysis is the study of how a set of agents link to one another, in terms of connectivity, reciprocity, distance and centrality, described in section 3. In a payment, clearing and settlement system, this can mean payment flows among participants, the approach this study uses (thus, the idea of connectivity in network analysis is distinct from that of connectivity in terms of IT systems).[4] Partial outages provide an opportunity to test how payment disruptions can affect this structure, due to the unexpected and exogenous nature of an outage.

Initial Fisher sign tests on connectivity support the hypothesis that a precautionary reaction with respect to problem participants could be taking place.

[1]See Arjani and McVanel (2006) for an overview of the LVTS.

[2]An operational event can be defined as an event where human error, fraud, management failure, deficiencies in information systems or internal controls, or disruption from external events such as natural disasters, terrorism or pandemics, results in unexpected losses. Business reputation and the resource cost of resolving an event could be other concerns. The LVTS operational events in this study, however, have mostly resulted in an interruption or slowdown to operations.

[3]Conversely, an inability to "receive" payments, due to a connection problem between a participant's systems and the LVTS, would mean that only the participant's systems would not record the received payments.

[4]Please refer to Embree and Roberts (2009) for an introduction to the main measures.

Precautionary behaviour means that participants stop payments to another participant experiencing a partial outage, very soon after the Canadian Payments Association (CPA) sends a notification. For connectivity, which is a network-wide measure, the effect is noticeable only when the problem participant is a big six bank. Directly counting the number of participants that continue to send payments to a problem participant (the "in-degree") reveals, though, that a decline occurs for both large and small problem participants. The number of counterparties almost always falls below a problem participant's mean in-degree. However, in absolute terms, the effect is the greatest with respect to the big six banks.

The econometric models then take network connectivity and reciprocity as the dependent variables and use calendar events, payment activity, lagged terms and characteristics of the operational event as explanatory variables. Although calendar and time-of-day effects are the most important, network connectivity and reciprocity nevertheless decline, depending on the type of event (either a payment slowdown or interruption) and the identity of the problem participant. A positive net payment balance at the problem participant during the partial outage, indicating the absorption of liquidity, also corresponds to reduced network connectivity. The effect of a partial outage at a big six bank that has a net balance of $1 billion is estimated to correspond to a decline in connectivity of about 6 or 7 per cent. This estimate is above and beyond the reduction in payment activity owing to the problem participant not sending payments.

Given that the problem participant is only 1 of 14 participants (not counting the Bank of Canada), this can be seen as a moderate but meaningful tendency for participants to exercise precaution toward a problem participant. The implication is that participants tend to perceive a partial outage at a counterparty as inconvenient enough to warrant a delaying of payments, at least to the problem counterparty. This could be interpreted as an indication of active management by banks of their payment activity or intraday liquidity.

Further possibilities exist for behavioural reactions to an operational event. Participants could take a precautionary stance toward the rest of the network, as well, to avoid a disadvantageous liquidity position. This would lead to a contraction in the subnetwork that excludes the problem participant. The results show that this scenario might have occurred, to an extent, in a couple of the small number of partial outages at major banks. But, more typically, participants seem to continue to send payments as usual to other non-problem participants. To cause systemic disruption, an operational event would likely have to be considerably more serious than the type of simple outage experienced in the LVTS so far.

Operational risk analysis in the past has helped the Bank of Canada and others to develop contingency measures. These allow participants to work around a failure in a system; an example is the LVTS Direct Network, which has been available for use during outages since December 2005. Through the Direct Network, for instance, the small number of payments that are of particularly high value can be sent to alleviate liquidity pressure. A precautionary delay of payments, however, depends on notification by the problem participant, and the CPA, of the disruption. Admittedly, a delay in notification can often occur, but available contingency measures seem adequate, based on the events in this data set.[5]

The reliability of the LVTS means there is a scarcity of serious events to study. Lack of data is a common problem in operational risk research, and does not make it easy to draw conclusions about high-impact events, which are the most infrequent. Past events might also not provide good indicators of how other participants will respond in the future, due to changes in technology, sources of risk, risk management, contingency measures or simply because a particular low-frequency, high-impact event might not have a precedent. A quantitative approach to operational risk thus cannot provide a definitive analysis, but can potentially contribute to discussions and the formulation of scenarios related to operational risk.

Section 2 summarizes the related literature, and section 3 briefly describes the LVTS network structure, and the measures used in this paper. Section 4 describes the data and section 5 the methodology. Section 6 reports the results and section 7 offers some conclusions. The appendix provides a few network measure formulas.

2 Related Literature

As one part of their study, Soramäki et al. (2006) examine the impact on the Fedwire network of the 11 September 2001 terrorist attacks in the United States. They find that massive damage to property and communications systems made it more difficult, or impossible, for some Fedwire participants to send payments, and connectivity dropped by up to 13 per cent, about four times the standard deviation. The initial payment disruption had a subsequent effect on other participants that were then short on incoming payments. Cheung (2002) describes the impact of 11 September on payment flows in the LVTS, using half-hourly multilateral balances for two weeks in January and February of 2002, contrasted with 11 September 2001. Payment volume dropped by more than a standard deviation between 10 a.m. and noon on that day.

Theoretical models on coordination in payments systems search for explanations for broad patterns in payment activity. Bech and Garratt (2006) have an n-player game theory model where banks choose to send payments in the morning or afternoon. They find that a disruption that forces some banks to pay in the afternoon could lead to an afternoon equilibrium, where all banks delay payments, because of reduced liquidity from fewer incoming payments

[5] If the network had otherwise seemed unable to handle the events affecting a key participant, for example, the Bank of Canada might have needed to be more proactive in coordinating contingency measures that reinforce the available system resources for that participant. The economic rationale for further central bank support would be that private costs and benefits in preventing delays may not be aligned with social costs and benefits for the network. Operational risk creates costs for the participant itself, but also for counterparties, which will not be fully accounted for in private decisions to mitigate operational risk.

in the morning.[6] From a social perspective, it is better to avoid delay and gridlock.

Merrouche and Schanz (2009) use a related 2-player, 3-period (morning, afternoon, evening) game-theory model, to predict that banks will delay a payment to a counterparty experiencing an outage only in the morning, because they are still uncertain of any urgent payment instructions they might receive in the afternoon. Based on their sample of eight 2007 outages in the Clearing House Automated Payment System (CHAPS) at the five main banks, they conclude that the reaction to outages is stronger in the morning than in the afternoon, and that payment activity among non-problem banks is unaffected by an outage.

3 LVTS Network Structure

As Figure 1 illustrates, a network is a structure of nodes connected by links. In the context of a payments system such as the LVTS, these nodes normally represent individual financial institutions, and links can represent payment flows, bilateral credit limits or other variables describing a relationship between two financial institutions over a given observation period. A link can be weighted and directed, for example, according to the volume or value of payments that occur between the sender and recipient that it connects.

Average gross daily payment activity for subgroups in the LVTS over the course of 2008 are depicted in Figure 1. It emphasizes the network dominance of the three largest banks, who made an average of over $58.7 billion in gross payments among themselves per day. The next largest links were between the three largest banks and the remaining big six banks, which account for $50.6 billion. Payments between the largest three banks and foreign subsidiaries accounted for another $29.7 billion, while other payment activity is comparatively low in value. The average daily LVTS payment value in 2008, including payments involving the Bank of Canada, was $181.6 billion.

In the network without aggregation into subgroups, most bilateral relationships involve less than $1 billion in payment activity per day; most relationships above that amount involve at least one of the three largest banks. Some bilateral gross flows between smaller participants and the three largest banks have a magnitude comparable to those among the other big six.

Typical network measures simplify the payment flows to unweighted links. High-frequency intraday data make this simplification more palatable, since the variability of transaction activity by participant is still captured. Payment activity is extremely variable (i.e., there are few high-value payments and many low-value payments); the transformation of these data into network measures can make it more amenable to modelling.

Connectivity describes the proportion of actual links to potential links,

[6] An empirical issue is that it is difficult to disentangle the strategic decision of delaying payments from the possibility that many instructions could simply be received later in the day (Mills and Nesmith 2006).

while *reciprocity* is the proportion of reciprocated links to actual links. This paper focuses on connectivity and reciprocity, due to their simplicity and interpretability. The *in-* and *out-degrees* of a node count its number of ingoing and outgoing links, respectively. The distance between two nodes is the length of the shortest path between those two nodes, measured by the number of nodes that must be travelled to, to reach the other node. When there is no path between two nodes, that distance is undefined, so distance-based measures are calculated for the *giant strongly connected component*. This is the largest subnetwork where all members have both ingoing and outgoing links to the rest of the subnetwork as a whole.

Apart from being measures of payment activity, network measures could have further interpretations. Lower connectivity could imply lower overall liquidity in the financial system, in the sense that fewer counterparties are engaged in actively transacting at a particular moment. We might also interpret connectivity as a characterization of the type of interdependence that exists. For instance, high connectivity could imply that financial or operational stress at a problem participant would be more evenly distributed across participants, instead of the case where one counterparty takes the full weight of exposure. Reciprocity, on the other hand, suggests the extent that payments tend toward bilateral exchanges. Such coordination could help participants manage their intraday liquidity.

One way to measure *centrality* is with the betweenness algorithm. The *betweenness* centrality of a node is the probability that the node is used as an intermediary on a shortest path between any two other nodes. *Centralization* is a summary statistic; 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 0 (low) to 1 (high).

Network statistics with a daily sampling frequency describe the LVTS as densely connected. The high connectivity, reciprocity and short average path lengths reflect how the LVTS already encompasses a main core of the Canadian financial system. This contrasts with the United States, where many peripheral players in financial markets are members of Fedwire.

Intraday statistics better illustrate the stratification that exists even within this small group of LVTS participants. Different intraday sampling frequencies typically depict a similar evolution of network structure, shown in Figure 2. Connectivity increases quickly in the morning after 8 a.m., before slowing to a gradual increase, peaking at closer to 5 p.m. Reciprocity peaks at around 6 a.m. to 7 a.m. With betweenness centralization, an hourly frequency shows greater centralization at around 7 a.m. and 5 p.m.

Calendar and time-of-day effects are the most important variables for explaining network structure. Embree and Roberts (2009) did not find that increased credit-risk concerns during the financial crisis had a statistically significant effect on connectivity.

Simple network measures are the most readily interpretable in the context of this study. Centrality and distance-based measures depend on the size of the giant strongly connected component, whereas connectivity and reciprocity have a meaning that is independent of other features of the network.

4 LVTS and Operational Events Data

LVTS transaction-by-transaction payments data, for the 942 business days from April 2004 to December 2007, are available to the Bank of Canada through the CPA. The LVTS data include the sender, receiver, settlement time and tranche, for each payment.[7] From April 2004 to December 2007, about 170 operational events related to the LVTS occurred, most of which were minor events such as late payments and incorrect manual computer entries. Participant outages have been relatively infrequent. Records contain a brief 1-line description of the event, the start time, the duration, the name of the affected participant and the severity,[8] among other details.

Of these events, 28 originated from the Bank of Canada, 6 originated from either CDSX, the CPA or SWIFT, and 6 had an "other" origin, such as telecommunications. Participants apart from the Bank of Canada accounted for the rest of the events. The cause of an event was attributed to the following sources: 64 for process and procedures, 36 for human error, 8 for external sources such as hydro and telecommunications, and 65 for internal systems issues such as systems connectivity, hardware and software. These classifications are inevitably subjective, and any given event could have more than one contributing cause.

Results included an inability to send payments (21 cases), a delay in a particular type of payment (auction, CLS, bank note, participant or settlement payment, 56 cases), the withholding of payments (19), insufficient collateral (6), constraint by a credit limit (2), other (46), or no significant impact (20), with the possibility of multiple effects. Events involving an overall decline in payments from the problem participant are likeliest to affect general LVTS behaviour, rather than the delay of a particular payment.

Other qualifications remain for the Bank of Canada's LVTS operational events data, apart from those mentioned in the introduction. Records do not always give a complete range of times for different stages of the event timeline.[9] Additionally, we are primarily interested in knowing when other participants become aware that there is a problem, through a CPA notification, since this designates the *disclosed phase* when they might change their behaviour in re-

[7] The payments data do not cover the time periods of two notable events: the terrorist attacks of 11 September 2001, and the electricity blackouts in August 2003. Detailed payments data are generally not available earlier than April 2004. This paper extends the data set only to 2007 due to the manual demands of verifying the operational events data for consistency with respect to the payments data and CPA notification emails.

[8] The severity of an operational event is difficult to quantify, particularly when none of these events resulted in appreciable financial losses. In any case, the large majority of events, with no noticeable impact, are rated by the Bank as negligible.

[9] These stages include when the event occurred, when the origin participant became aware, when the origin participant notified the CPA, when the CPA notified others, when the problem was resolved and when the CPA notified other participants of this resolution.

sponse to the new information. Thus we are not always able to confirm exact start times and durations of the disclosed phase for some events.

A visual depiction of a partial outage illustrates the elements involved. Figure 3 shows the timeline for the main stages of an operational event, superimposed in red lines on the connectivity of the LVTS, the connectivity of the subnetwork that excludes the problem participant, the mean connectivity of that subnetwork, and the region between the mean positive and negative deviations from the mean connectivity shaded in grey.[10] The red bars at the bottom indicate the in-degree, the number of links directed toward the problem participant, on the right y-axis. Figure 4 graphs the corresponding situation for the total value transacted.

The figures reveal that if network-wide reactions to an operational event are present, they are subtle; visually, a reaction in LVTS or subnetwork connectivity is impossible to distinguish from random fluctuations in LVTS activity. Bank A's partial outage in 2004 leaves a noticeable effect on network structure relative to other events; the connectivity of the LVTS and the subnetwork without Bank A stays close to the lower end of the average negative deviation from the mean. The total value transacted in the network jumps shortly after CPA notification of the problem's resolution.

A change in LVTS behaviour directly with respect to the problem participant is more visible. Bank A in-degree and in-value are abnormally low during the 2004 event, but spike briefly after CPA notification of the problem's resolution; most partial outages at other major banks show more temporary hesitations in problem participant in-degree, but the value of payments to these problem participants does reach a low level. There is still variability in payment value, however, which a network approach obscures. This serves as a reminder that network analysis can be a blunt tool, possibly better suited for detecting systemic effects or major developments in how participants interact over a longer time period.

4.1 Filtering of operational events

This paper focuses on events involving an interruption in payments from the problem participant, such as partial outages, and other events that prompted a notification from the CPA, such as when the problem participant experiences a slowdown in its ability to make payments. Other events might involve an interruption in certain payments; however, for an event to cause a deliberate and general reaction among LVTS participants, there would plausibly have to have been some form of alert from the CPA.[11] The analysis further excludes the following types of events:

- Events the duration of which occur mainly outside regular business hours (before 8 a.m. and after 6 p.m.), when it would be impossible to gauge

[10]The distribution is not necessarily symmetrical around the mean.

[11]Impact categorizations give largely insignificant results in the models, and usually do not correspond to any aggregate reduction in payments made by the problem participant, except for those that involve an outage or an inability to send.

participant behaviour because activity is normally very low during these periods.[12] Although CLS pay-ins in the early morning are important, they would more suitably be covered by a different type of analysis.

- Events of short duration (typically less than 30 minutes), if we cannot be reasonably certain that the CPA sent out a notification.

- A small number of other events where information in the LVTS participant report seems to be too inconsistent with actual payment activity.[13]

4.2 CPA email notifications

Bank of Canada economists have previously identified the stages that take place for any given event; initially, there is the undisclosed phase when an event has occurred and the rest of the LVTS is not yet aware of the problem.[14] When the CPA notifies the LVTS participants of the event, the disclosed phase begins. The Canadian Payments Association sends out a brief email that identifies the problem participant, and whether that participant is experiencing an inability to send/receive payments, or experiencing a slowdown in payments. Occasionally, the email gives an estimate of the time to recovery.[15] After disclosure, other participants still might not suspend payments to the problem participant, if they do not expect the problem to be serious or protracted enough to warrant precautionary behaviour.

An absence of a CPA email record does not guarantee that an email was not sent. If an email could have been sent, because of the nature of the problem and its duration, or if the time stamp on the email is simply missing, then it would be preferable not to discard this event. Informed judgment, based on payment activity and typical time delays for stages of the event timeline, helps to determine a reasonable time for these notifications. In practice, though,

[12]Too many interacted variables would be required to appropriately handle such events in a model, given the total number of usable events. LVTS network connectivity is relatively high during daytime business hours, but it is significantly more sparse outside of that period. Regardless, there would not be any informative behavioural changes, if payment activity is already minimal.

[13]Examples include the following: (1) the event description states that the problem participant was "unable to send payments," even though there was no interruption or CPA email; the LVTS participant report might describe a different type of problem; (2) the event description states that the problem participant was "unable to send payments," even when a CPA email states that the problem participant was unable to *receive* payments, and payment activity continued.

[14]It seems likely that other participants will not immediately infer that there is a problem, although it could be possible; for example, if there had been a pre-agreed coordination of payments, or if communication occurs without CPA intermediation.

[15]The participant is expected to notify the CPA within 15 minutes of detecting a partial outage (Canadian Payments Association 2011), and subsequently the CPA aims to notify other LVTS participants on a best-efforts basis. As shown in Table 4, actual delays can vary relative to these recommendations. Delays can occur because the participant might not notice the problem immediately, or because it does not report it immediately.

for most significant events that might have had a general effect on payment activity, an email time stamp can be found.[16]

This leaves us with 54 partial outages that Table 1 summarizes by participant subgroup and event type. Forty-one of these events involve an actual interruption in payments from the problem participant, while the remaining 13 either involve a slowdown in payment transmission or otherwise result in a CPA email stating that the participant was unable to send. For simplicity, this paper will often refer to all of these events as "partial outages." The three largest banks represent seven of these events. Table 1 should help to emphasize the limited number of events, especially if further restrictions are imposed; 19 events last less than an hour, and only three events simultaneously involve one of the largest three banks, an actual interruption in payments and a duration of greater than an hour. Table 2 similarly shows that the Bank of Canada rated only five events as having medium severity, while all others were either of low or negligible severity. Figure 5 and Table 3 show the distribution of start times and event calendar quarters. There were 17 of these events in 2005, 9 in 2006, and 17 in the first half of 2007 but only 1 in the second half of that year.

Table 4 summarizes the duration of the different stages in the event timeline, based on available records. Participants, on average, take 22 minutes to notify the CPA after detection of a problem, compared to a CPA guideline of 15 minutes. The CPA has no set time requirement for notifying LVTS participants after an alert from a problem participant, and operates on a best-efforts basis. The CPA took an average of 15 minutes after the receipt of an alert to notify LVTS participants both for a disruption and resolution. There is variation in how long notification takes, or how long the disclosed and undisclosed phases of an event can last. The average event duration is close to two hours. The average durations of the undisclosed and disclosed periods were 70 minutes and 68 minutes, respectively, based on available reported times.

5 Methodology

Disturbances in network measures such as connectivity and reciprocity, as well as the total value and volume of transactions before, during and after events, could potentially indicate altered participant behaviour with regards to an operational event. The distribution of these network statistics for the relevant time of day and calendar period provides a basis of comparison.

The analysis starts with a Fisher sign test, which compares the number of times that connectivity is either above or below its mean during specific partial outages. A 1-month moving average of connectivity for the relevant 5-minute time interval is used as the mean. This descriptive test provides a non-parametric way to evaluate individual events for precautionary behaviour toward a problem participant. It gives an initial sense of whether network

[16] Bowman (1983) emphasizes the importance of ascertaining the timing of an event accurately; an examination of LVTS payment activity suggests that a misalignment of start and end times would affect the analysis considerably.

characteristics changed, for specific cases.

The econometric models then seek to explain how network structure, and thus participant behaviour, relates in a general way to partial outages. Connectivity and reciprocity are the dependent variables, while a dummy variable represents the ongoing disclosed phase of an operational event.[17] The models control for holiday effects, time dependence (by demeaning network variables), the type of origin participant, type of problem (either a payment interruption or a payment slowdown) and Bank of Canada-rated event severity. Payment activity in the LVTS is strongly characterized by a daily pattern of activity, so, to control for this time dependence, the network variables are demeaned by a 1-month moving average of each particular 5-minute interval. A 1-month moving average is used instead of differencing with a 24-hour lag, because of the volatility in the series. For the connectivity model estimation, connectivity was also corrected for the effect of a halt in payments *from* the problem participant, such that the model estimates the response from other participants.

The analysis uses a 5-minute sampling frequency, meaning that a network structure snapshot is taken of the payment activity that accumulates over each 5-minute period. Embree and Roberts (2009) determine that sampling frequencies between five minutes and an hour portray a similar network structure evolution over the course of a day. Due to the short duration of most operational events, a 5-minute frequency is likely the most suitable for this study.

This paper focuses on Tranche 2 payments, since that is the standard means for LVTS participants to make payments efficiently. Tranche 1 payments are less numerous, and typically involve the Bank of Canada, so they would not measurably affect the network analysis, but can be of significant amounts; they are therefore considered on an event-by-event basis.

5.1 Fisher sign tests

Rather than jumping straight to an econometric model that perhaps naïvely presumes to draw general conclusions about the impact of certain operational event attributes, the Fisher sign test is a starting point for comparing network connectivity during a particular operational event to normal LVTS conditions.

The Fisher sign test compares the number of times a variable is either above or below its mean, and assumes that, under normal conditions, there is an independent and equal probability for either the high or low state in each period. Given this assumption, the probability of drawing a particular number of observations above the mean corresponds to the binomial distribution. With respect to network connectivity, the assumption of a binomial distribution is reasonable in that close to half of the observations for a given time of day are above the mean; the result from the outage period can be compared to the rest of the day, as a benchmark. Although the Fisher test does make simplifications, it gives an interpretable case-focused alternative to a model, and allows us to

[17] Thus the actual period of time when the event was taking place is not represented in the models, but rather the time of other participants' awareness of a problem.

recognize that any given operational event could have an effect that does not correspond to the general aspects of other disruptions.

The following equation calculates the P-value for the Fisher sign test by summing the probabilities of obtaining n^+ or fewer observations greater than the mean, for a binomial distribution with N observations:

$$P - val = \frac{\sum_{i=0}^{n^+} \binom{N}{i}}{2^N},$$

where n^+ is the actual number of observations greater than the mean, for the event (Weisstein 1999). The mean is taken to be the 1-month moving average of connectivity for that specific 5-minute interval of the day, since this would be a more relevant comparison than a mean for the entire sample.

5.2 Model specification

The models search for an overall effect on the network structure, including the direct consequence that other participants might suspend payments to the problem participant, which would materialize as a drop in connectivity.

Importantly, we would expect that participant reactions would depend on whether the problem participant has already absorbed a substantial level of liquidity. Although the other participants would not be aware of this balance, the rest of the LVTS will have an aggregate balance complementary to that of the problem participant. The model thus includes an interacted variable representing the net accumulated Tranche 1 and Tranche 2 balance for the current problem participant, during the disclosed phase of any operational events. For any other non-event time, this variable is set to zero, since there is no problem participant. The model can be summarized as follows:

$$c_t = \alpha + CIV_t\beta_{civ} + US_t\beta_{US} + IP\_BIG6_t\beta_{B6} + IP\_OTHER_t\beta_{OTHER}$$

$$+SP_t\beta_{SP} + ETR\_lte30_t\beta_{ETR} + BAL\_OE_t\beta_{BAL} + \varepsilon_t,$$

where c_t is the connectivity of the network, demeaned by the 1-month moving average of the relevant time interval to remove most of the time dependence in the series; *CIV* and *US* are dummy vectors for civic and U.S. holidays, respectively; *IP_BIG6* is a dummy for an interruption of payments from a big six bank; *IP_OTHER* is likewise for an interruption of payments from a non-big six bank; *ETR_lte30* is a dummy for an estimated time to recovery of less than or equal to 30 minutes in the CPA email; *SP* represents a slowdown of payments; *BAL_OE* is the net payment balance of the problem participant (if there is an outage for that time interval); and ε_t is an error term.

The inclusion of two or three lags of the dependent variable, and autoregressive and moving-average terms of one lag, results in residuals that have minimal autocorrelation and that are not rejected by Bartlett's periodogram-based test for white noise. Estimation proceeds with this approach, even if it could result in conservative estimates; persistence in the data could occur for a number of reasons that would vary across time. A low-activity day will likely have many adjoining 5-minute intervals with fewer payments than average, but the lagged terms will also capture part of the effect that could otherwise be legitimately attributed to a partial outage. Furthermore, the estimation of coefficients for the lagged terms will be driven by periods of normal activity, because partial outages affect only a small part of the data.

We might also expect a change when a problem continues to be unresolved after a certain period of time, or soon after normal activity is restored. Including the appropriate variables does not offer much improvement, though. Bank economists have previously observed a brief spike in the volume of payments to the problem participant after resumption of normal activity, but this particular effect does not noticeably translate to a similar overcompensation in the network structure after an event. Neither does an attempt to include intervals of elapsed time during the disclosed phase improve the model.

The basic model that examines indirect effects of an operational event is as follows:

$$c_t^{-i} = \alpha + CIV_t\beta_{civ} + US_t\beta_{US} + IP_{it}\beta_{IP_i} + SP_{it}\beta_{SP_i} + \varepsilon_t,$$

where c_t^{-i} is the connectivity of the subnetwork that excludes the problem participant i. IP_i and SP_i are specific to problem participant i in this case, while ε_t is an error term.

6 Results

Table 8 summarizes the results of the Fisher test for nine partial outages, along with the evolution in the net balance position of the problem participant for both Tranches 1 and 2, at the start of the problem, the start of the disclosed phase and the end of the problem.[18]

The test result for connectivity is insignificant for the small-participant events, while it is significant for most large-participant events when looking at the entire LVTS network, at a 5 per cent level or higher. In contrast, during the rest of the event day, close to half of the observations are above the mean (45 per cent for large banks and 42 per cent for the others). Problem participant in-degree (not shown), as a measure of the direct response toward the problem participant, is always lower than the mean, for the largest six banks, and usually also for other banks.

The effect on the subnetwork of non-problem participants is only somewhat significant for a couple of events at larger banks at, or close to, a 10 per cent level. One such case occurred despite the bank's previously mentioned negative net position of over \$1 billion throughout its 2007 partial outage. For other events, the problem participant often does accumulate liquidity; in one case (not shown), the bank absorbed around \$1.43 billion in payments before the CPA was able to disclose the problem to other participants. The bank had

[18] It is possible that the problem participant's net balance would decline if they were to use alternative measures (e.g., manual processing) to make payments.

reached a balance of $1.52 billion by the time it resolved its problem soon after. Compared to the value of LVTS payments transacted by a big participant, this level of liquidity absorption would not be of systemic concern.

These findings lend support to the model results that a disruption at a participant leads to a precautionary reduction from others, while the effect is noticeable on a network-wide scale only for the largest banks. Events that involved a notable positive accumulation of liquidity seem to have mostly occurred in 2004 or 2005.

6.1 Model analysis

Tables 5 and 6, respectively, show results of model estimation for two key statistics: connectivity and reciprocity. We should interpret the individual coefficients with the caveat that these estimates correspond to a limited number of operational events, most of which have a negligible effect on the system.

The connectivity model in Table 5 results in an expected negative sign for the net-balance variable, significant at a 1 per cent level. The coefficient implies that for a $1 billion net position at a problem participant, the network's connectivity tends to decline by 0.0054; combined with the coefficient of -0.0075 for a payment interruption from a big six bank (also significant at a 1 per cent level), and an average connectivity of 0.188 during the regular business day (8 a.m. to 6 p.m.), this would represent a decline in network connectivity (in levels) of about 6 or 7 per cent, above and beyond the reduction in payment activity owing to the problem participant not sending payments. This is about half of a standard deviation, for a 1-hour average of connectivity (the disclosed phase is, on average, 68 minutes), across most given periods of a day (e.g., the standard deviation of connectivity at 12 p.m., across all days).

This effect should be taken in the context that the problem participant is only 1 of 14 participants (not counting the Bank of Canada), and that the dependent variables were corrected for the effect of a halt in payments from the problem participant, such that the model estimates the response from other participants. Considering these factors, the 6 or 7 per cent reduction could be seen as a moderate but meaningful tendency of participants to exercise precaution toward a problem participant. It suggests that banks actively manage their payment activity throughout the day, and do not send payments with complete obliviousness. Presumably, one motivation for this could be to maintain a reasonable liquidity position. For connectivity, which is a network-wide measure, this precaution is mainly noticeable with respect to big six banks; the dummy coefficient for other banks was not significant.

Arguably, the lagged terms will capture a part of an operational event's effect; persistence in the data could occur for many reasons, which will likely vary over time. It seems likely that an operational event would be a relatively important explanation for some of this persistence, while the event is occurring, so it could be reasonable to also estimate the model without lagged terms. Excluding these terms, in a comparable ordinary least squares model, the estimated effect of a partial outage (significant at a 1 per cent level) would be more than twice as large, at 13 per cent. This is mainly due to an increase in coefficient for a payment interruption from a big six bank, to -0.02.[19] The general pattern could easily be overridden in particular cases. For example, a major bank was in a negative net position of between $2.5 and $1 billion during a 2007 partial outage, but participants appear to have exercised precaution, based on payment activity and the corresponding effect on connectivity.

Most participants have not experienced slowdown events, according to our data, and the slow-payments coefficient is not significant. An estimated time to recovery of less than 30 minutes has a small positive effect (significant at the 10 per cent level), whereas if the problem participant is not one of the largest six banks, the network-wide effect is not perceptible. An F-test shows that the event variables are significant as a group.

Table 6 shows a decline in reciprocity for a problem at a big six bank of about -0.056, significant at the 1 per cent level, which would represent about 12 per cent of the network's reciprocity, during the regular business day, and about 1.1 standard deviations. This is not surprising, if others occasionally continue to send to a problem participant that has ceased making payments. Thus, while the connectivity result suggests precautionary behaviour, the reciprocity model shows that this precaution does not entail a complete halting of payments.[20] A drop in reciprocity among non-problem participants could imply further reduced coordination in the payments system. This would not be desirable, if net exposures persist longer. But it does not seem that reciprocity falls among non-problem participants, based on the modelling of reciprocity in the subnetwork that excludes the problem participant.

Table 7 summarizes the results of the subnetwork models that look for an impact on connectivity, in levels, between non-problem participants. Since these models are specific to a particular participant's operational events, they are even further limited by the small number of disruptions. Based on this limited evidence, the effect on the subnetwork that excludes the problem participant is not statistically significant, especially when the model includes lagged terms (lagged dependent variables, and autoregressive and moving-average terms of one lag). There is also no apparent indirect effect from operational events on reciprocity among the subnetwork of non-problem participants. Without the lagged terms, though, the coefficients for two large banks suggest a decline in connectivity in the non-problem subnetwork of about 4 and 10 per cent, respectively, indicating that indirect network-wide effects on payment activity are also possible, even for a partial outage of one or two hours. Such an effect would, in theory, reflect a deterioration in payments coordination because of participants' reluctance to risk a decline in their liquidity positions.

---

[19]Only an R^2 of 0.033 can be attributed to the model's variables, without lagged terms, but this should be expected given that there were 54 events of limited duration in a data set spanning 942 days.

[20]The proportional effect between connectivity and reciprocity is not directly comparable. Larger banks likely account for a greater share of total reciprocity, due to their centrality in the network. Also, connectivity was adjusted to compensate for the effect of a halt in payments from the problem participant. For reciprocity, this adjustment would make less sense, because its purpose is to illustrate bilateral coordination.

Table 7 also shows the effect on the in-degree to the problem participant. About 46 per cent fewer participants tend to send payments to the problem participant in an average period, compared to the relevant average in-degree, when it is one of the largest six banks. The effect on other banks is similarly large, in percentage terms, with a 48 per cent reduction, on average, although the in-degree for these other banks is usually less than a third that of the largest six banks. These results confirm that bilateral precautionary behaviour occurs with respect to a problem participant, even when the problem participant is not one of the largest banks. The implication is that participants tend to perceive a partial outage at a counterparty as potentially inconvenient enough to warrant a delaying of payments, at least to the problem counterparty. In the case of the big six problem participants, though, the network-wide effect is perceptible, as shown in Table 5.

Models relating Tranche 1 in- or out-payments to the problem participant do not reveal a marked tendency to use Tranche 1 during an operational event. However, in specific cases this can occur. With a couple of major bank partial outages, problem participant use of Tranche 1 made a significant difference for the participant's outstanding balance.

6.2 Discussion

Partial outages at the largest five or six banks consistently had the most impact on participant behaviour, and thus network structure. A precautionary reaction toward the problem participant would be desirable, to the extent that it prevents an absorption of liquidity by this participant, which is not able to send payments itself. A large positive balance position for the problem participant at time of disclosure does also seem to be a meaningful factor for reducing connectivity. Yet Table 8 gives an example of where a major bank's large negative balance position is not enough to dissuade other participants' precautionary reactions, during the decline in connectivity of the 2007 event.

The results of this study do not suggest that options for the handling of operational events are inadequate, or that the Bank of Canada needs to increase operational assistance to LVTS participants. Previous improvements in contingency measures possibly helped in this regard. Given the apparent precautionary behaviour with respect to big-bank outages, unless the bank had already reached a critical accumulation of funds by the start of the disclosed phase, it would appear unlikely that the situation would reach that point after a CPA notification. When an event does occur, contingency measures can reduce the likely impact of the disruption and the corresponding reactions from other participants. Manual entry of payments is one such option when system problems prevent automatic processing.

Examples include the manual processing of payments during three of a smaller bank's events in 2006 and 2007. In addition, at least one major bank took advantage of the LVTS Direct Network to send a CDSX settlement payment of close to $2 billion during a 2007 outage. If participants prioritize the small fraction of payments that are of high value, these interim measures should be effective.

If there was concern for the possibility of a liquidity trap, the system could notify participants more quickly of a disruption through automatic and high-frequency tests of system functionality. This test could verify that all participants are able to send and receive payments. A standardized electronic form that the problem participant uses to promptly notify the CPA about an operational event, which then triggers an automatic notification for other LVTS participants, might also reduce the length of the undisclosed phase. It could be reasonable, though, for a participant to delay its notification to the CPA, if this enables it to assess a problem and verify that it interferes with the functionality of the LVTS. If that is the case, the relevant CPA guideline could be revised.

Continuing on the theme of analyzing the network's response to external shocks, future research could investigate the network's response to Bank of Canada policy actions or changes, including repo transactions, term lending or the expanded eligibility of collateral. These measures were intended to reduce the cost of intraday liquidity, which, following Bech and Garratt (2006), could affect morning versus afternoon payments coordination equilibrium. Disruptions originating at a payment, clearing and settlement system or the Bank of Canada represent another category of operational events.

6.3 Limitations

The LVTS functions very well, and there are few serious and persistent disruptions to the system; the lack of suitable data is a common problem for the study of operational risk. This could limit the robustness of the results, especially given idiosyncratic events that are heterogeneous in the problem participant, the time of day and the nature of the problem. Because there might be only a couple of events, if any, for a given participant, one should proceed with caution before generalizing the results.

Participants might not perceive the need to alter their payment behaviour during an operational event affecting another participant, since many problems can be resolved in a short amount of time. This would appear to be the case for many of the events in our data set, especially with respect to small problem-participants.

It is also unclear how the experience of routine events informs us of the repercussions of more serious outages. We should remain cautious in extrapolating from past experience to make conjectures about low-probability, high-impact events in Canada's payments systems.

A complete understanding of the behaviour in the LVTS requires more than network measures. A small number of payments made through alternate means could be sufficient to significantly change the net accumulated balance of the problem participant, which would not be reflected in a network analysis. Nevertheless, a network approach could give a better appreciation of the possible systemic implications of disruptions, which admittedly did not arise in the events examined in this study.

7 Conclusions

The findings suggest that other participants tend to withhold payments to the problem participant, but using a network approach, the effect is mainly perceptible for the big six banks. The main econometric model illustrates that participants tend to withhold more payments when the problem participant has already absorbed a great deal of liquidity: for a $1 billion net position at a big six bank experiencing a partial outage, connectivity is estimated to decline by about 6 or 7 per cent. Given that there are 14 participants in the LVTS, this can be viewed as a moderate and meaningful effect on LVTS activity. Arguably, it suggests that banks actively manage their liquidity throughout the day, and do not send payments with complete obliviousness. The implication is that participants perceive a partial outage at a counterparty as potentially inconvenient enough to warrant a delaying of payments, at least to the problem counterparty.

The effects of a partial outage, on only a couple of occasions, might have influenced how non-problem participants interact amongst themselves, based on the subnetwork model and Fisher sign tests, but it is difficult to say that there has been much change in behaviour among non-problem participants. Overall, a partial outage in the LVTS would not seem likely to cause a systemic disruption. The incidence of operational events did not increase over the sample period, and cases that involved an accumulation of liquidity seem to have mostly occurred in 2004 or 2005. The improvement of contingency measures, such as the use of the LVTS Direct Network, might have helped in this regard.

Network analysis is a blunt tool for dealing with participant-specific reactions, but can help to assess the broader impact of a disruption on the structure of participant interactions. Another limitation for studies on operational risk is the lack of data, and the difficulty in predicting the implications of a low-probability, high-impact event based on past experience.

If the occurrence of a sudden liquidity trap is a genuine concern, then real-time automatic monitoring of system functionality could be a solution. However, it is not clear that this would be warranted, given that a small number of manually processed payments can alleviate a liquidity trap.

References

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

Bech, M. and R. Garratt. 2006. "Illiquidity in the Interbank Payment System following Wide-Scale Disruptions." Federal Reserve Bank of New York Staff Reports No. 239.

Bowman, R. G. 1983. "Understanding and Conducting Event Studies" *Journal of Business Finance & Accounting* 10 (4): 561−84.

Canadian Payments Association. 2011. "LVTS Rule 12: Emergency Conditions." Toronto.

Cheung, L. 2002. "Understanding Intraday Payment Flows in the Large Value Transfer System." Bank of Canada *Financial System Review* (December): 49−52.

Embree, L. and T. Roberts. 2009. "Network Analysis and Canada's Large Value Transfer System." Bank of Canada Discussion Paper No. 2009-13.

Jackson, M. O. 2006. "The Economics of Social Networks." In *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society*, 1−56. Edited by R. Blundell, W. Newey and T. Persson. Cambridge: Cambridge University Press.

Merrouche, O. and J. Schanz. 2009. "Banks' Intraday Liquidity Management during Operational Outages: Theory and Evidence from the UK Payment System." Bank of England Working Paper No. 370.

Mills Jr., D. C. and T. D. Nesmith. 2006. "Risk and Concentration in Payment and Securities Settlement Systems." Federal Reserve Board FEDS Working Paper No. 62.

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.

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

Weisstein, E. W. 1999. "Fisher Sign Test." From MathWorld–A Wolfram Web Resource. Available at http://mathworld.wolfram.com/FisherSignTest.html.

Table 1: Event Type by Participant

	Payment Impact				
Origin	No interruption		Interruption		
	Duration				
	<hour	>hour	<hour	>hour	Total
Big six	2	2	10	7	21
Others	2	7	5	19	33
Individual Banks					
Median	0	0	0.5	1	2.5
Min.	0	0	0	0	0
Max.	2	4	5	6	10
Total	4	9	15	26	54

Table 2: Actual Event Severity by Participant

	Actual Severity						
	Negligible		Low		Medium		
	Duration						
	<hour	>hour	<hour	>hour	<hour	>hour	Total
Big six	8	2	4	4		3	21
Others	5	11	2	13		2	33
Individual Banks							
Median	0	0	0	1	0	0	2.5
Min.	0	0	0	0	0	0	0
Max.	5	4	2	4	0	1	10
Total	13	13	6	17	0	5	54

Table 3: Distribution of Events by Year and Quarter

	Quarter				
Year	1	2	3	4	Total
2004	0	4	4	2	10
2005	4	4	3	6	17
2006	2	2	1	4	9
2007	6	11	1	0	18
Total	12	21	9	12	54

Table 4: Summary of Event Timeline Phases (minutes)

Variable	**Mean**	**Std. dev.**	**Min.**	**Max.**
Time for participant to notice problem	31	76	-5	376
Time for participant to notify CPA	22	50	0	268
Time for CPA to notify others	15	15	0	66
Time for CPA to confirm resolution	10	17	-81	44
Undisclosed period	70	139	-10	673
Disclosed period	68	57	7	298
Actual duration of event	128	163	12	750

Table 5: Estimation Results: Connectivity, Demeaned by 1-Month Moving Average

Variable	**Coefficient**	**(Std. err.)**	**P $> \|t\|$**
Equation 1: Connectivity			
Civic holiday	-0.0404	(0.0023)	0.0000
U.S. holiday	-0.0104	(0.0015)	0.0000
Payment interruption, big six	-0.0075	(0.0021)	0.0010
Payment interruption, other	0.0016	(0.0017)	0.3380
Payment slowdown	0.0005	(0.0031)	0.8650
ETR less than 30 minutes	0.0084	(0.0038)	0.0260
Prblm-Bank Balance ($bn)	-0.0054	(0.0000)	0.0080
Depvar, 1st lag	0.1562	(0.0017)	0.0000
Depvar, 2nd lag	0.0173	(0.0016)	0.0000
Intercept	0.0001	(0.0002)	0.7780
Equation 2: ARMA			
L.ar	0.9770	(0.0006)	0.0000
L.ma	-0.9174	(0.0013)	0.0000
Equation 3: sigma			
Intercept	0.0276	(0.0000)	0.0000
N	226318		
Log-likelihood	491545.9		
$\chi^2_{(11)}$	4330508.8		

Table 6: Estimation Results: Reciprocity, Demeaned by 1-Month Moving Average

Variable	**Coefficient**	**(Std. err.)**	**P > |t|**
Equation 1: Reciprocity			
Civic holiday	-0.3057	(0.0077)	0.0000
U.S. holiday	-0.0647	(0.0034)	0.0000
Payment interruption, big six	-0.0557	(0.0146)	0.0000
Payment interruption, other	0.0165	(0.0117)	0.1560
Payment slowdown	0.0012	(0.0230)	0.9570
ETR less than 30 minutes	0.0119	(0.0334)	0.7210
Prblm-Bank Balance	-0.0053	(0.0130)	0.6850
Depvar, 1st lag	0.0839	(0.0019)	0.0000
Depvar, 2nd lag	0.0181	(0.0019)	0.0000
Depvar, 3rd lag	0.0152	(0.0019)	0.0000
Intercept	0.0089	(0.0006)	0.0000
Equation 2: ARMA			
L.ar	0.9792	(0.0026)	0.0000
L.ma	-0.9644	(0.0034)	0.0000
Equation 3: sigma			
Intercept	0.1491	(0.0002)	0.0000
N		134968	
Log-likelihood		65330.7	
$\chi^2_{(12)}$		363609.2	

Table 7: Partial Outage Effects on Subnetworks Excluding Problem Participant (For big six)

	Connectivity		In-degree	
	Mean	Coefficient	Coefficient	% Effect
LVTS	0.1946			
(Bank 1)	0.1720	-0.002	-2.07**	-47
(Bank 2)	0.1709	0.004	-1.54**	-36
(Bank 3)	0.1961	0.004	-1.22**	-48
(Bank 4)	0.1425	-0.008	-2.19**	-36
(Bank 5)	0.1678	-0.002	-2.87**	-65
(Bank 6)	0.1760	n/a	n/a	

Note: ** = 1% significance level

Table 8: Fisher Sign Tests, and Liquidity Positions by Operational Event

Participant	Event year	Disclosure time[a]	C > C_mean[b]		$C^{-i} > C^{-i}$_mean[c]		Net balance, mil. (T1 & T2)		
			n^+/N^d	P-value	n^+/N^d	P-value	Start	Disclosure	End
(Big six)	2004	8h07	2/24	0.000	5/24	0.003	-219	397	840
(Big six)	2004	10h25[e]	2/12	0.019	5/12	0.387	185	583	683
Other	2004	15h45	10/22	0.416	13/22	0.857	-18	-16	-16
(Big six)	2005	10h09	0/42	0.000	17/42	0.140	121	933	1,375
(Big six)	2005	15h50[e]	1/11	0.006	5/11	0.500	-804	423	427
Other	2006	11h45[e]	5/13	0.291	7/13	0.709	485	546	704
(Big six)	2007	8h21	3/28	0.000	13/28	0.425	359	372	751
(Big six)	2007	16h18	4/16	0.038	5/16	0.105	-2,513	-1,006	-1,531
Other	2007	9h06	10/17	0.834	10/17	0.834	-440	-422	-421

[a] Start time of the disclosed phase, when other participants receive a CPA email notification.
[b] The 1-month moving average of connectivity for that time of day.
[c] The 1-month moving average of connectivity, excluding the problem participant, for that time of day.
[d] The fraction of 5-minute periods in the disclosed phase where connectivity exceeds the 1-month moving average for that time of day.
[e] Estimated (and rounded) time of CPA email notification.

Figure 1: Average Daily Value in 2008, Can$ billions, for LVTS subgroups

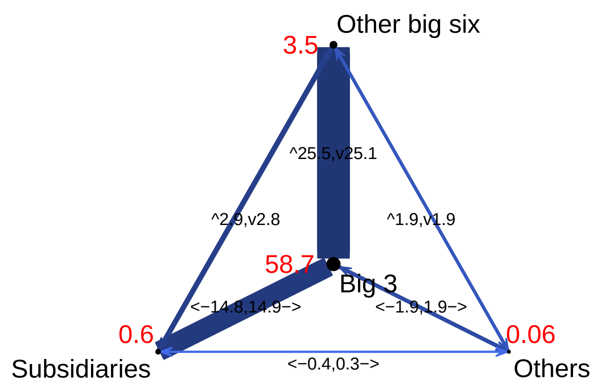

Figure 2: Network Statistics for the LVTS, January 2007

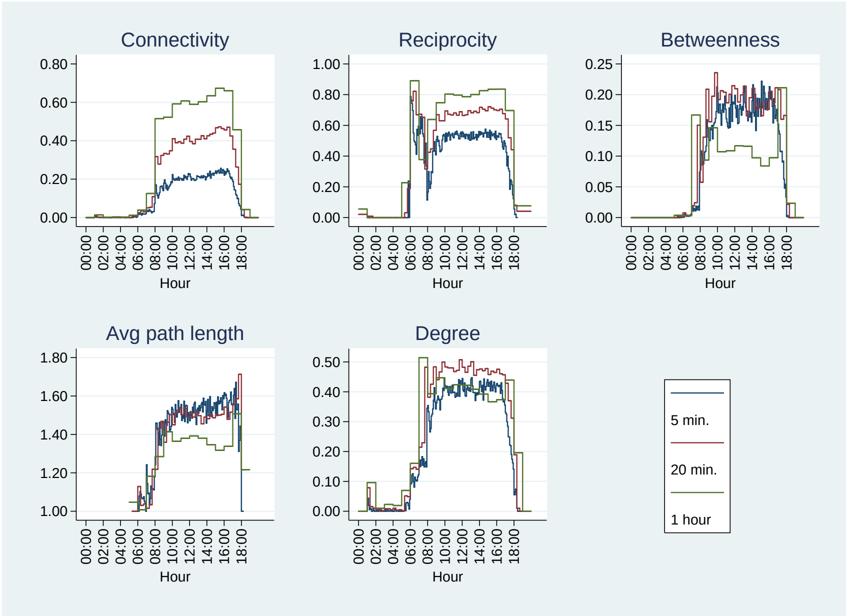

Figure 3: Large Bank Partial Outage, 2004

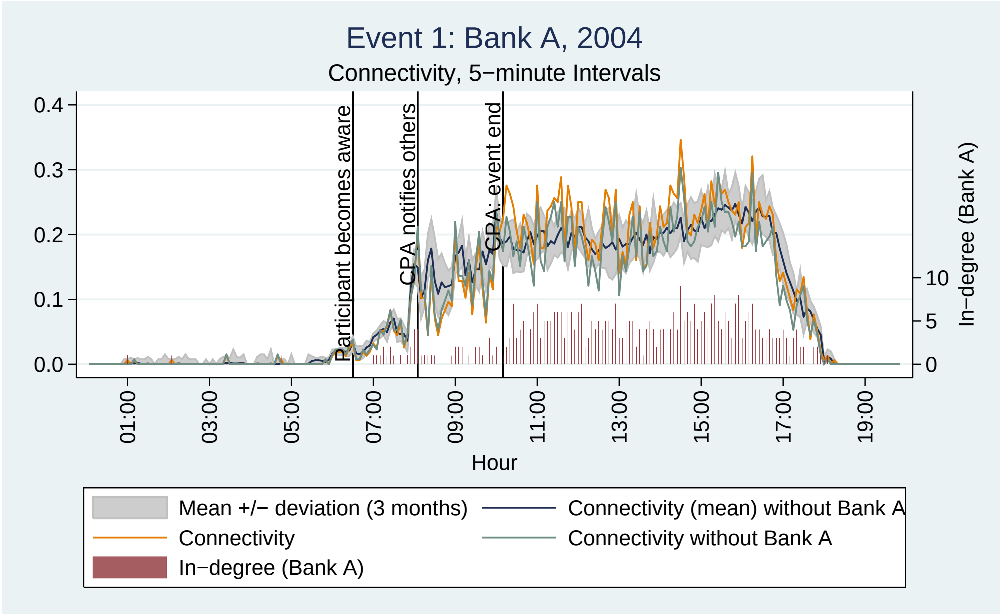

Vertical black lines depict when the participant becomes aware of the problem, when the CPA notifies others, and when the CPA notifies others of the event end. Other important stages of the event, such as when the problem occurs and when the participant notifies the CPA, are not depicted for space reasons, and to emphasize the disclosure period to other participants. Lines show the connectivity of the LVTS, the connectivity of the subnetwork that excludes the problem participant, and the mean connectivity of that subnetwork; the region between the mean positive and negative deviations from the mean connectivity is shaded in grey. The red bars indicate the in-degree, the number of links directed toward the problem participant, on the right y-axis.

Figure 4: Large Bank Partial Outage, 2004

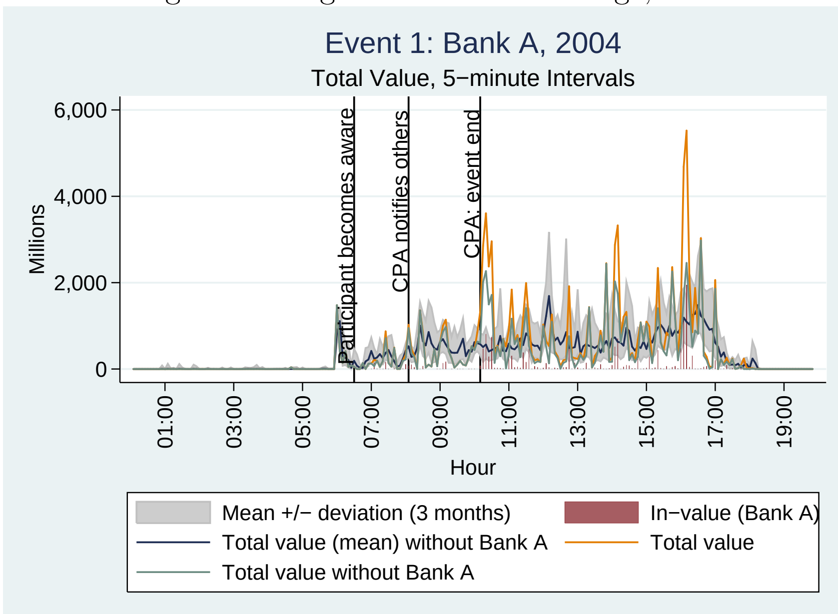

Figure 5: Distribution of Event Start Times

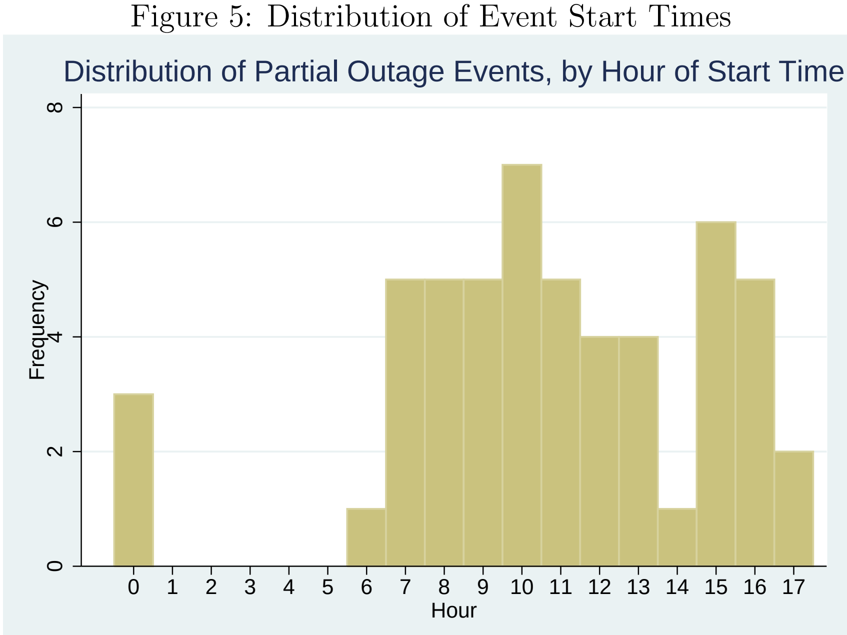

Appendix A: Definitions of Measures

This appendix defines the measures used in the main text. For more information, refer to Soramäki et al. (2006), or Embree and Roberts (2009). Wasserman and Faust (1994), and Jackson (2006), also offer overviews of social network analysis.

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

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 presence of links between nodes, with $p \in \Im, 1 \leq p \leq N-1$. 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 $l_i = \frac{1}{n-1} \sum_{j \neq i} d_{ij}$. The average path length of a network is $L = \frac{1}{n} \sum_i l_i$.

A.3 Centrality and Centralization

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

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

Betweenness centrality is $C_B(n_v) = \frac{\sum_{i \neq v} \sum_{j \neq i, j \neq v} \frac{g_{ivj}}{g_{ij}}}{(N-1)(N-2)}$ 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 $\zeta_{C_Z} = \frac{\sum_i |\max_v C_Z(n_v) - C_Z(n_i)|}{\mathrm{TMAXDEV_{C_Z}}}$ where

- $C_Z(n_i)$ is the chosen centrality measure, for node i
- $TMAXDEV_{C_Z}$ is the theoretical maximum deviation of the numerator, based on the number of nodes in the network, used to normalize the centralization measure