BANK OF CANADA
BANQUE DU CANADA
CELEBRATING 75 YEARS
CÉLÉBRONS 75 ANS

Working Paper/Document de travail
2010-13

Estimating the Structure of the Payment Network in the LVTS: An Application of Estimating Communities in Network Data

by James T. E. Chapman and Yinan Zhang

www.bank-banque-canada.ca

Bank of Canada Working Paper 2010-13

June 2010

Estimating the Structure of the Payment Network in the LVTS: An Application of Estimating Communities in Network Data

by

James T. E. Chapman and Yinan Zhang

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

ISSN 1701-9397

Acknowledgements

We would like to thank our colleagues at the Bank of Canada for the helpful comments and suggestions. We are also grateful to the seminar participants at the Laboratory for Aggregate Economics and Finance (LAEF) 2008 conference of Payments and Networks and at the Canadian Economics Association (CEA) 43rd Annual Conference 2009 for valuable discussions.

Abstract

In the Canadian large value payment system an important goal is to understand how liquidity is transferred through the system and hence how efficient the system is in settling payments. Understanding the structure of the underlying network of relationships between participants in the payment system is a crucial step in achieving the goal.

The set of nodes in any given network can be partitioned into a number of groups (or "communities"). Usually, the partition is not directly observable and must be inferred from the observed data of interaction flows between all nodes. In this paper we use the statistical model of Čopič, Jackson, and Kirman (2007) to estimate the most likely partition in the network of business relationships in the LVTS. Specifically, we estimate from the LVTS transactions data different "communities" formed by the direct participants in the system.

Using various measures of transaction intensity, we uncover communities of participants that are based on both transaction amount and their physical locations. More importantly these communities were not easily discernible in previous studies of LVTS data since previous studies did not take into account the network (or transitive) aspects of the data.

JEL classification: C11, D85, G20
Bank classification: Payment, clearing, and settlement systems; Financial stability

Résumé

Il est important de bien saisir comment les liquidités transitent dans le Système de transfert de paiements de grande valeur (STPGV) canadien et, ainsi, d'évaluer l'efficience avec laquelle les paiements sont réglés. Pour ce faire, il est essentiel de comprendre la structure du réseau de relations sous-jacent qui lie entre eux les participants au système.

Les nœuds que comporte un réseau, quel qu'il soit, se partagent en un certain nombre de groupes (ou « communautés »). Le plus souvent, cette répartition n'est pas directement observable et doit être déduite des données observées sur les interactions entre tous les nœuds. Dans notre étude, nous utilisons le modèle statistique de Čopič, Jackson et Kirman (2007) pour estimer la répartition la plus probable des nœuds du réseau de relations opérationnelles à l'intérieur du STPGV. Plus précisément, nous cernons, à partir des données sur les opérations, les différentes communautés que forment les participants directs au système.

En faisant appel à diverses mesures de l'intensité des opérations, nous mettons au jour des communautés de participants fondées tant sur la valeur de leurs opérations que sur leur emplacement géographique. Il s'agit là d'une avancée importante, car les études

antérieures sur les données du STPGV permettaient difficilement de discerner les communautés du fait qu’elles ne tenaient pas compte des aspects transitifs des données.

Classification JEL : C11, D85, G20
Classification de la Banque : Systèmes de paiement, de compensation et de règlement; Stabilité financière

1 Introduction

The linkages between participants in a payment system are of interest to payment system designers and regulators for two main reasons. First, it is through these links that financial contagion could potentially spread.[1] Second, the flow of liquidity between participants is affected by the system of linkages between participants[2]. In this paper we use a dataset of transactional level data between direct participants in the Canadian Large Value Transfer System (LVTS) to estimate a simple model of the underlying structure of business links between LVTS participants.

The key contribution of this paper is the identification of different partitions of the set of LVTS direct participants who transact with each other, within the same group (i.e. the same component of a partition or the same "community"), more intensively than with any participant outside the community. Using two measures of transaction intensity that depend on average and maximum transaction linkages, we uncover communities of participants that are based on both transaction amount and their physical locations. More importantly, part of the community structure found in this paper is not easily discernible in previous studies of LVTS data since previous studies do not take into account the network (or transitive) aspects of the data. This implies that two banks which are similar in other observable ways, but belong to different communities, may have very different impacts on the LVTS due to their relative importance in their communities. This study adds to the literature on network topology in payment systems pioneered by Soramaki, Bech, Arnold, Glass, and Beyeler (2007).

We apply the model of Čopič, Jackson, and Kirman (2007) to LVTS data using a Bayesian implementation.[3] In the model, the probabilities of system direct participants forming a link within a group and across different groups are assumed to follow binomial distributions. Posterior densities of these probabilities are constructed based on both observed transactions between each pair of participants and

[1] See for example the work of Allen and Gale (2000). Bech and Garratt (2006) also show that the degree of interconnectedness between banks can affect the resiliency of a payment system.

[2] Recent work by O'Conner, Chapman, and Millar (2008) also makes this point.

[3] An alternative methodology to discover network structure developed by Bergstrom (2007) uses a minimum descriptive approach to describe the structure of the network.

a potential maximum capacity of bilateral transactions (that could have been observed). A Markov chain Monte Carlo (MCMC) sampler is then used to compare all different partitions of the network probabilistically.

Taking a network view allows us to account for features that observed bilateral payment flows cannot take into account. For example consider four banks (A, B, C, D). It could be that A and D are the largest banks in the system according to the transaction volume. But by looking at all the flows in system, it may be that A, B, and C send the majority of their flows to one of the other two banks in this triad. In general, we would conclude that banks A and D are the most important participants. But from a systemic perspective it might be that D is relatively less important to the system than the coalition of A, B, and C, even though the latter two are smaller in terms of the transaction amount. This is because the flow of funds between the banks in the coalition is being taken into account. It should be noted that the results, as well as the interpretation of the results, depend on the measure of transaction intensity used.

We try two measures of transaction intensity. In the first case the amount of transactions is equal to the average daily total value of payments sent from one participant to another, and the maximum potential bilateral linkage for a participant is then estimated as the sum of the maximum amount of daily gross payment inflow and the maximum bilateral credit limit (BCL) granted.[4] Under the second measure, the amount of interactions among participants is defined to be the number of days on which one bank's daily transaction value sent to the other bank exceeds the average payment value the sending bank sends to all other system members on the same day; and it follows that the maximum amount of linkages between any two participants in this context is the total number of LVTS-operating days over the sample period.

Using these measures, we uncover two of the most likely partitions of the LVTS network, and in both five big Canadian banks are classified in the same community. One of the partitions also reveals a second community of some small and large participants that conduct most of their businesses in the same province (Quebec).

[4]In a recent work, Bech, Chapman, and Garratt (2010) use BCLs to define the network structure of LVTS to capture participants' liquidity redistribution behaviours.

In the next section (section 2) we describe the statistical model and the Bayesian MCMC sampler that are used to estimate the communities in the LVTS. We then describe the data set of LVTS payment instructions to which we apply the model in section 3. The results are presented in section 4 and we conclude in section 5.

2 Empirical Model

2.1 The Model

The empirical model we use to identify community structures from interaction flow data of a network draws heavily on the work of Čopič, Jackson, and Kirman (2007).

One key assumption of this model is that the probability of nodes of the same community interacting with each other is different from, and strictly higher than, the probability of interactions among nodes of different communities. The simplest model specification defines that these probabilities are the same for all nodes.

The model is described as follows.

Define a set of nodes $N = \{1, 2, ..., n\}$ which will represent banks in the LVTS. Over the sample period, we observe links between all pairs of nodes; these links form a network on the set N that can be represented by a matrix $o \in \mathbb{Z}_+^{n\times n}$. The total amount of links any two nodes can have is limited by a maximum capacity between the two nodes over the sample period. Let matrix $c \in \mathbb{Z}_+^{n\times n}$ denote the set of interaction capacities on N.

The set of nodes N can be partitioned into a number of groups (or communities), and let π be one partition of N. For any partition π and node $i \in N$, $k_\pi(i)$ denotes the element of π containing i. Any two nodes within the same group are more likely to form links with each other than with any other node that is not in that group. We parametrize the probabilities of forming a link inside a group and outside (across) groups as p_{in} and p_{out} respectively. As mentioned above, the key assumption is that $1 \geq p_{in} > p_{out} \geq 0$.

Therefore, given the maximum number of links c_{ij} that node i and node j can possibly have, the probability mass function of an observed number of links o_{ij} between the two nodes, if i and j are in the same group, is:

$$p(o_{ij}|p_{in}, p_{out}, c_{ij}) = \binom{c_{ij}}{o_{ij}} p_{in}^{o_{ij}} (1 - p_{in})^{c_{ij} - o_{ij}}, \tag{1}$$

and if i and j are not in the same group is:

$$p(o_{ij}|p_{in}, p_{out}, c_{ij}) = \binom{c_{ij}}{o_{ij}} p_{out}^{o_{ij}} (1 - p_{out})^{c_{ij} - o_{ij}}, \tag{2}$$

Both of them follow binomial distributions.

Let's assume that the probability distribution of the number of links formed between one pair of nodes is independent of the distribution between any other pair in the network. Thus, given a community structure π, we can easily calculate the probability of having the observed network data, matrix o, by multiplying the two probability functions across all pairs of nodes.

$$L_{o,c}(\pi|p_{in}, p_{out}) = B \prod_{i \in N} \left[\left(\prod_{j \in k_\pi(i)} (p_{in})^{o_{ij}} (1 - p_{in})^{c_{ij} - o_{ij}} \right) \left(\prod_{\substack{j \in N \\ j \notin k_\pi(i)}} (p_{out})^{o_{ij}} (1 - p_{out})^{c_{ij} - o_{ij}} \right) \right] \tag{3}$$

where B is a constant of binomial coefficients.

Taking logs of both sides, we get the log-likelihood function as follows. We ignore the constant $\log(B)$) when comparing the log-likelihood scores between two partitions.

$$\begin{aligned} \ell_{o,c}(\pi|p_{in}, p_{out}) &= log\Big(L_{o,c}(\pi|p_{in}, p_{out})\Big) \\ &= q_1 S^{In(\pi)}(o) + q_2 S^{In(\pi)}(c) + q_3 S^{Out(\pi)}(o) + q_4 S^{Out(\pi)}(c) \end{aligned} \tag{4}$$

where $In(\pi) = \{ij|i \in N, j \in k_\pi(i)\}$ denotes the set of all pairs of nodes that are in the same group under the community structure π; $Out(\pi) = \{ij|i \in N, j \notin k_\pi(i)\}$ is the set of all pairs of nodes that are not in the same group under π. $S^{In(\pi)}(o) = \sum\limits_{ij \in In(\pi)} o_{ij}$ and $S^{In(\pi)}(c) = \sum\limits_{ij \in In(\pi)} c_{ij}$ are the sum of links, between all pairs of nodes that are in the same group, in the observed network o and in the capacity network

c respectively. $S^{Out(\pi)}(o) = \sum_{ij \in Out(\pi)} o_{ij}$ and $S^{Out(\pi)}(c) = \sum_{ij \in Out(\pi)} c_{ij}$ are defined in the same fashion as immediately above. We let $q_1 = log\left(\frac{p_{in}}{1-p_{in}}\right)$; $q_2 = log(1 - p_{in})$; $q_3 = log\left(\frac{p_{out}}{1-p_{out}}\right)$ and $q_4 = log(1 - p_{out})$.[5]

2.2 Bayesian Implementation

The set of all possible partitions of a given network is usually a large unordered set. For example, a network of 14 nodes can be partitioned in more than 190 million different ways. An appropriate Bayesian MCMC sampler provides us with convenience to compare all these partitions probabilistically.

We use the Gibbs sampler complemented by the Metropolis-Hasting algorithm to construct our posterior simulators.[6] In general, the Gibbs sampler divides the set of all unknown features of the model θ_A into groups, where A denotes the model; and then it generates a sequence of samples for every group by repeatedly making drawings from well-defined conditional distributions. In this study, θ_A consists of three parameters: p_{in}, p_{out} and the network structure π.

The MCMC simulation theory has established that if the output chain converges to a fixed point, then the simulated sequence $\{\theta_A^{(m)}\}_{m=1}^{\infty}$, where m is the *m-step* iteration, should be a representative of that invariant joint posterior distribution of all the parameters in the model.

2.2.1 Priors and Posteriors

We assume uniform probability across all possible partitions of the network, which reflects our limited a priori knowledge about the network structure,

$$p(\pi|A) = \text{constant},$$

this is a proper uninformative prior since the set of all possible partitions is a, very large, finite number. The likelihood function of a given partition π, conditional on

[5] An interested reader can find a complete characterization of this maximum likelihood model, e.g. its theorems and properties, please see Čopič, Jackson, and Kirman (2007).

[6] An interested reader can find a text book on this method, e.g. *Contemporary Bayesian Econometrics and Statistics*, by John Geweke.

the values of p_{in} and p_{out}, is $L_{o,c}(\pi|p_{in}, p_{out})$ in Equation 3. Hence, the log-posterior of π is obtained as follows.

$$\begin{aligned} \log\left(p(\pi|p_{in}, p_{out}, y^o, A)\right) &\propto \ell_{o,c}(\pi|p_{in}, p_{out}, y^o, A) + \log(p(\pi|A)) \\ &\propto \ell_{o,c}(\pi|p_{in}, p_{out}, y^o, A) \qquad\qquad \text{(i.e. Equation (4))} \end{aligned}$$

where y^o denotes the set of the observed data collected on the network, which in this study includes the actual linkages among all the nodes, matrix o and the interaction capacities matrix c.

Our priors for p_{in} and p_{out} are selected to be the Beta distribution, which is the conjugate prior of the binomial distribution,

$$\begin{aligned} p(p_{in}|A) &= \frac{1}{B(\alpha,\beta)} p_{in}^{\alpha-1}(1-p_{in})^{\beta-1} \\ p(p_{out}|A) &= \frac{1}{B(\alpha,\beta)} p_{out}^{\alpha-1}(1-p_{out})^{\beta-1}, \end{aligned}$$

where $B(\alpha, \beta)$ is the Beta function on the two shape parameters α and β. We chose prior Beta distributions since these are flexible distributions (e.g. it can be multimodal) and are conjugate; therefore we think this choice is without loss of generality.

Conditional on the data y^o and a given partition π, the likelihood functions of p_{in} and p_{out} can be derived , respectively, from Equation (1) and (2). These conditional likelihood functions contain complete information of p_{in} and p_{out} that resides in the data y^o.

$$\begin{aligned} L(p_{in}|p_{out}, \pi, y^o, A) &\propto \prod_{i\in N} \prod_{j\in k_\pi(i)} (p_{in})^{o_{ij}} (1-p_{in})^{c_{ij}-o_{ij}} \\ L(p_{out}|p_{in}, \pi, y^o, A) &\propto \prod_{i\in N} \prod_{\substack{j\in N \\ j\notin k_\pi(i)}} (p_{out})^{o_{ij}} (1-p_{out})^{c_{ij}-o_{ij}} \end{aligned}$$

Combining the priors and the likelihood functions, we get the posteriors of p_{in}

and p_{out} as follows.

$$p(p_{in}|p_{out},\pi,y^o,A) \propto (p_{in})^{\alpha+\left(\sum_{i\in N}\sum_{j\in k_\pi(i)} o_{ij}\right)-1}(1-p_{in})^{\beta+\left(\sum_{i\in N}\sum_{j\in k_\pi(i)} (c_{ij}-o_{ij})\right)-1} \tag{5}$$

$$p(p_{out}|p_{in},\pi,y^o,A) \propto (p_{out})^{\alpha+\left(\sum_{i\in N}\sum_{\substack{j\in N\\ j\notin k_\pi(i)}} o_{ij}\right)-1}(1-p_{out})^{\beta+\left(\sum_{i\in N}\sum_{\substack{j\in N\\ j\notin k_\pi(i)}} (c_{ij}-o_{ij})\right)-1} \tag{6}$$

If we assume $\alpha = 1$ and $\beta = 1$ in our prior beliefs about p_{in} and p_{out}, then the posteriors of the two parameters are Beta distributions with the shape parameters being:

$$\left[\left(\sum_{i\in N}\sum_{j\in k_\pi(i)} o_{ij}\right)+1, \left(\sum_{i\in N}\sum_{j\in k_\pi(i)} (c_{ij}-o_{ij})\right)+1\right] \qquad \text{(for } p_{in}\text{)}$$

$$\left[\left(\sum_{i\in N}\sum_{\substack{j\in N\\ j\notin k_\pi(i)}} o_{ij}\right)+1, \left(\sum_{i\in N}\sum_{\substack{j\in N\\ j\notin k_\pi(i)}} (c_{ij}-o_{ij})\right)+1\right] \qquad \text{(for } p_{out}\text{)}$$

Taking logs of both sides of Equation (5) and (6), we get the log-posteriors of p_{in} and p_{out}.

$$\begin{aligned}\log\Big(p(p_{in}|p_{out},\pi,y^o,A)\Big) \propto & \left[\alpha + \left(\sum_{i\in N}\sum_{j\in k_\pi(i)} o_{ij}\right)-1\right]\log(p_{in})\\ &+\left[\beta + \left(\sum_{i\in N}\sum_{j\in k_\pi(i)} (c_{ij}-o_{ij})\right)-1\right]\log(1-p_{in})\end{aligned}$$

$$\begin{aligned}\log\Big(p(p_{out}|p_{in},\pi,y^o,A)\Big) \propto & \left[\alpha + \left(\sum_{i\in N}\sum_{\substack{j\in N\\ j\notin k_\pi(i)}} o_{ij}\right)-1\right]\log(p_{out})\\ &+\left[\beta + \left(\sum_{i\in N}\sum_{\substack{j\in N\\ j\notin k_\pi(i)}} (c_{ij}-o_{ij})\right)-1\right]\log(1-p_{out})\end{aligned}$$

2.2.2 Metropolis within Gibbs

The use of a combination of the Gibbs sampler and Metropolis-Hasting algorithms is an effective and powerful solution to sampling issues so that it is straightforward to sample parameters in group 1, for instance, from distribution conditional on the parameters in group 2, but the reverse conditional distribution is intractable. In this study, both p_{in} and p_{out} are assumed to have conjugate priors, i.e. simple Beta distributions, and therefore, Markov chains of these two variables are obtained using standard Gibbs sampling procedure.

On the other hand, it is much less obvious of how to generate samples from the distribution of partitions conditional on p_{in} and p_{out}. Hence, we use the Metropolis-Hasting algorithm to generate new candidates for the partition and to form a Markov chain accordingly.[7] In this paper, the probability of move, also known as the Metropolis-Hasting Ratio (M-H Ratio) is

$$\min\left[\frac{p(\pi^{(*)})q(\pi^{(*)},\pi^{(m)})}{p(\pi^{(m)})q(\pi^{(m)},\pi^{(*)})},1\right]$$

where $\pi^{(m)}$ is the partition at the *m-step* iteration; $\pi^{(*)}$ is a new partition proposed for the next-step iteration; $q(\pi^{(m)},\pi^{(*)})$ is the *transition probability density function*, also known as the *candidate-generating function*, which characterizes the partition moving from $\pi^{(m)}$ to $\pi^{(*)}$; the target density function, $p(\pi^{(\cdot)})$ of any given partition π is characterized by the log-likelihood function shown above in the model specification: $\ell_{o,c}(\pi)$.

Alternatively, in logarithm, our M-H Ratio can be written as:

$$\min\left[\exp\left\{\ell_{o,c}(\pi^{(*)})+\log\left(q(\pi^{(*)},\pi^{(m)})\right)-\ell_{o,c}(\pi^{(m)})-\log\left(q(\pi^{(m)},\pi^{(*)})\right)\right\},1\right]$$

We formulate our transition probability density function $q(\pi^{(m)},\pi^{(*)})$ on the basis of the idea that, given an existing partition of the network, a new partition can be generated by one node random-walking out of a source component to a different destination component; and more importantly, the probability for the node to random walk the reverse path is nonzero. The function is constructed as follows.

Let π be a partition of the set of nodes in the given network, and π' represent a different partition of the network distinct from π (resulted from moving one node

[7]An interested reader can lean more details about the Metropolis-Hasting algorithm from Chib and Greenberg (1995).

at a time). C denotes the set of components in a given partition and C_{ij} is the j^{th} element of component i. $\#C$ represents the cardinality of the components in a given partition, and $\#C_i$ is the cardinality of the elements in component i.

Then we identify three probabilities underlying every transition of a partition. $P_1 = P(C_i|\pi) = \frac{1}{\#C}$ denotes the probability of choosing a source component i; $P_2 = P(C_{ij}|C_i, \pi) = \frac{1}{\#C_i}$ is the probability of choosing C_{ij} as a candidate to move; and we let $P_3 = P(C_k|C_i, C_{ij}, \pi) = \begin{cases} \frac{1}{\#C} & \text{if } \#C_i > 1 \\ \frac{1}{\#C-1} & \text{if } \#C_i = 1 \end{cases}$ be the probability of choosing a destination component for placing C_{ij}.

The reason for $P_3 = \frac{1}{\#C}$ when $\#C_i > 1$ is that the element C_{ij} selected to be moved has a possibility of forming its own singleton community. Hence, the transition probability of the partition of the network changing from π to π' is calculated as:

$$\begin{aligned} P(\pi'|\pi) &= \text{the probability of moving } C_{ij} \text{ from } C_i \text{ to } C_k \\ &= P_1 \times P_2 \times P_3 \\ &= P(C_i|\pi) \times P(C_{ij}|C_i, \pi) \times P(C_k|C_i, C_{ij}, \pi) \\ &= \begin{cases} \frac{1}{(\#C)^2 \times \#C_i} & \text{if } \#C_i > 1 \\ \frac{1}{(\#C)^2 - \#C} & \text{if } \#C_i = 1 \end{cases} \end{aligned}$$

In summary, the Metropolis-Hasting algorithm used in this paper can be summarized as follows. We conduct the simulation for for $m = 1, 2, \ldots, M$ iterations. For each iteration, we generate a partition candidate $\pi^{(*)}$ from the candidate-generating function $q(\pi^{(m)}, \cdot)$ and u from the uniform distribution $\mathcal{U}(0, 1)$. If $u \leq$ M-H Ratio, then $\pi^{(*)}$ is chosen to be the partition for the next-step iteration; otherwise, we let $\pi^{(m+1)} = \pi^{(m)}$. In the end, we get a chain of values $\{\pi^{(1)}, \pi^{(2)}, \ldots, \pi^{(M)}\}$.

2.2.3 MCMC Simulation

The Markov chain Monte Carlo simulation begins with some arbitrary initial values for the two parameters in the model: (1) $\pi^{(0)}$ is arbitrarily set to be any partition, for instance, a network of 14 singleton communities; and (2) $p_{out}^{(0)}$ (or $p_{in}^{(0)}$; it does not matter which one is chosen) is randomly drawn from its posterior beta distribution

$beta(\cdot|\pi^{(0)}, y^o, A)$; We then successively make drawings from the following conditional distributions:

$$
\begin{aligned}
& p_{in}^{(0)} \sim beta(\cdot|\pi^{(0)}, y^o, A, p_{out}^{(0)} : p_{in}^{(0)} > p_{out}^{(0)}) \\
& \pi^{(1)} = \text{chosen by the M-H algorithm}|p_{in}^{(0)}, p_{out}^{(0)} \\
& \dots \\
& p_{out}^{(m)} \sim beta(\cdot|\pi^{(m)}, y^o, A, p_{in}^{(m-1)} : p_{in}^{(m-1)} > p_{out}^{(m)}) \\
& p_{in}^{(m)} \sim beta(\cdot|\pi^{(m)}, y^o, A, p_{out}^{(m)} : p_{in}^{(m)} > p_{out}^{(m)}) \\
& \pi^{(m+1)} = \text{chosen by the M-H algorithm}|p_{in}^{(m)}, p_{out}^{(m)} \\
& \dots
\end{aligned}
$$

The output of the simulation is a Markov chain of the partitions and the two probabilities: p_{in} and p_{out}, generated after an initial burn-in period. The estimate of the probability of each different partition of the network is simply the proportion of the observed times that partition is sampled.

3 Data

The Large Value Transfer System (LVTS) is the main Canadian large value payment system. There are fifteen direct system members, including the Bank of Canada. The fourteen private financial institutions (FIs) are listed in Table 1. The majority of these FIs can also be categorized into two groups: Montreal-based versus Toronto-based, depending on where the majority of their operations are located.

We use the actual data from LVTS operations that includes all individual transactions and the bilateral credit limits (BCL). This payment system consists of fourteen direct member institutions plus the Bank of Canada. The LVTS provides two channels to submit a payment: an RTGS-equivalent process known as Tranche one and a second method which is a hybrid of RTGS and DNS systems known as Tranche two (T2).[8] We will focus on Tranche two since the majority of LVTS transaction value flow through this payment stream.

[8]RTGS and DNS stand for *Real Time Gross Settlement* and *Deferred Net Settlement*, respectively.

Table 1: Direct Participants in LVTS

Royal Bank of Canada
Bank of Montreal
Canadian Imperial Bank of Commerce
The Toronto-Dominion Bank
The Bank of Nova Scotia
National Bank of Canada
Alberta Treasury Branches
Bank of America National Association
BNP Paribas (Canada)
La Caisse Centrale Desjardins du Qubec
Credit Union Central of Canada
HSBC Bank Canada
Laurentian Bank of Canada
State Street Bank and Trust Company

Our sample consists of all T2 transactions as well as BCL's in the LVTS from June 1st, 2005 to December 31st, 2007. The start of the sample was chosen based on an internal analysis that suggests that the system had returned to a stationary equilibrium around May 2005, following the entry of the newest LVTS direct participant, State Street Bank and Trust Company, in October 2004. Due to the deteriorating financial crisis in 2008 and 2009, more recent data is not used for showing the main results that are aimed to depict the community structure of the LVTS under normal financial circumstances.[9] This sample period spans 650 business days and it provides us with the largest possible sample size which helps us take into account low-frequency features of the data such as year-end and quarter-end effects.

The econometric model described in Section 2 requires a measure of interaction flows between every two nodes in the network and the potential maximal amount of linkages between the two nodes (i.e. the capacity of that edge). We denote the matrix of actual transaction flows between any two LVTS direct participants as $\mathbf{o}$ and the matrix of trading capacities as $\mathbf{c}$.

There is no definite way of defining these two matrices, and any logical measure of transaction intensity can be considered. We construct our o and c matrices in two

[9] In addition, a major parameter change in the LVTS took effect on May 1, 2008, i.e. an increase of the System Wide Percentage from 24% to 30%. It is possible that this change may affect certain aspects of participant's payment behaviours.

different ways.

Under the Liquidity Measure, the **o** matrix is a matrix of daily gross transaction value that one bank sends to the other averaged over the sample period. The **c** matrix, the maximum capacity of transactions between any two participants is estimated by the sum of the maximum BCL and the maximum daily total transaction value (chosen independently of the maximum BCL) that one participant receives from the other during the sample period. Under this measure, the entries of **c** matrix represent the upper bound of bilateral liquidity available to each LVTS participant, and $\frac{\mathbf{o_{ij}}}{\mathbf{c_{ij}}}$ can be interpreted as Bank i's average liquidity-consumption ratio. The idea is that the higher the ratio, the more intensively the payment-sending participant transacts with the payment-receiving participant.[10]

For the Averages Measure, the transaction flow in **o** matrix between every pair of participants is estimated by the number of days in the sample on which one participant's daily transaction value sent to the other participant exceeds the average flow on the same day from that sending bank to all other system members. The maximum linkages between any two participants is well-defined under this measure, so every entry in **c** matrix is the total number of LVTS-operating days over the sample period. The transaction intensity is thus reflected in the frequency at which each participant's bilateral daily gross payment outflow is larger in value than its multilateral average over the sample period.

4 Results

We apply the maximum likelihood model to the historical LVTS operations data. Using the two measures of transaction intensity specified in Section 3, we uncover two different community structures of business relationships among LVTS direct participants. Under the both measures, the "big five" Canadian banks are invariably found in the same group; the averages measure also reveals a second community composed

[10]To be 100% correct, multilateral liquidity constraints should also be considered in constructing the measure. However, the focus on bilateral liquidity flows should be sufficient in this case because the empirical model proposed by Čopič, Jackson, and Kirman (2007) is founded on bilateral interactions between pairs of nodes in a system.

of participants that are based in the same physical location.

We conduct three simulations for each measure of transaction intensity, using different staring values of the partition of the LVTS. The three starting values of the partition we used are: (i) a network of 14 singleton communities; (ii) every LVTS participant being in one and the same group; and (iii) a community structure featured by 7 singletons, a small Montreal-based cluster and a big Toronto-based group of the five big Canadian banks. The third partition is in fact one of the main results that we obtained from implementing the grid-search algorithm developed by Čopič, Jackson, and Kirman (2007).

Under each measure of transaction intensity, the results from all simulations are consolidated into one output chain for every parameter. Each simulation runs for 5×10^5 iterations, and we execute a burn-in by discarding the initial 10% of the output chain in each simulation.

4.1 Liquidity Measure

Using the liquidity measure of transaction intensity in the LVTS, we discover a cluster of five Toronto-based Canadian banks and that every other participant forms its own singleton community. As shown in *Figure 1*, T labels the participants whose operations are based in *Toronto*, M represents *Montreal*-based banks, and O means *Other* geographical locations. Large circles represent large participants and different communities are shaded with different colors to show the clustering. The positioning of the circles in the graph is randomly chosen and irrelevant in describing the results.

This partition is consistent with the anecdotal practice of grouping LVTS participants based on their transaction sizes in the system, i.e. large banks versus small participants. Indeed, during the sample period of this study, the payment value sent by the five big banks accounts for 82.24% of the system throughput.

All 1.35×10^6 post-burn-in observations point to this community structure without exception, and the posterior average values for the two parameters are $\bar{p}_{in} = 0.334223$ and $\bar{p}_{out} = 0.121034$. The difference between the two means is tested at four significance levels of 5%, 1%, 0.5% and 0.1%, and the null hypothesis of the difference

Figure 1: Community Structure of the LVTS under Liquidity Measure

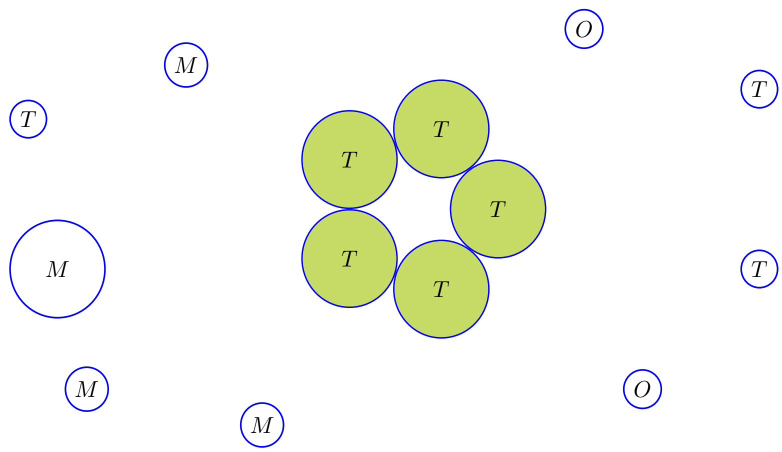

being zero is rejected in all cases. [11]

The 95% highest probability density (HPD) interval for the MCMC draws of p_{in} and p_{out} are $[0.3342210, 0.3342244]$ and $[0.1210333, 0.1210353]$, respectively.

4.2 Averages Measure

Under the Averages measure, the strength of directed links between any two participants is estimated by comparing, for each day in the sample, the payment sender's bilateral transaction outflow with its multilateral average. The idea is that if two participants both send payments to each other consistently more than what they transact in the system on average (in terms of the dollar value), then they should be considered having a closer-than-average business relationship and thus belong to the same community.

This altered perspective results in different partitions of the network. Our main result shows a partition consisting of two clusters of LVTS direct participants. In addition to the large community composed of the five big Canadian banks, the same as uncovered by the liquidity measure of transaction intensity, there is a second

[11] The fact that the community structure does not change for all iterations of the MCMC run is not surprising given the amount of data relative to the community structure; this is expanded upon in the discussion below.

cluster of two Montreal-based participants, one large in size, the other small. The remaining seven LVTS participants still form their own singleton communities, and the whole community structure is portrayed in *Figure 2*.

Figure 2: Community Structure of the LVTS under Averages Measure

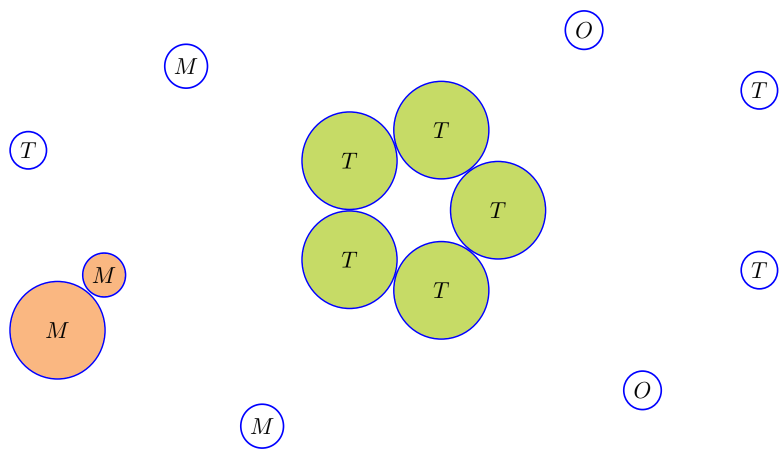

This partition of the LVTS network accounts for 66.67% of all 1.35×10^6 MCMC draws in the combined simulations, This community structure tells us that, during the sample period, the chances that larger value of bilateral payments flow between the participants within the same group is higher than that occurring between any pair of the participants across communities. For example, the network data shows that it is more common that the large Montreal-based bank sent more payments in value on a day to one of the small participants based in the same province than its average outflow on that day, and vice versa. The same is true with the five big Toronto-based banks in the larger community.

The finding can be potentially useful. For instance, in an event of either of the participants in the Montreal cluster experiencing an operational disturbance, the discovered structure, based on the transaction intensity among participants, suggests that the other participant would be more likely and/or severely affected than any other system member on average.

This community structure is meaningful also in the sense that there seems to be a geographic component in the structure of LVTS payment flows. LVTS is the large-value system in Canada used to settle funds positions in other payments (and securities) clearing and settlement systems, and these funds positions are initiated and driven by real business funds transfers between numerous economic agents. It is not hard to see that the same physical operating location can result in a higher degree of business association between LVTS direct participants.

The rest 33.33% of the observations in the simulation point to a slightly different partition of the network. That is, the Montreal-based group and the five big Toronto banks cluster together and form one big community, with other participants remaining singletons.[12]

The posterior average estimated values for the two parameters are $\bar{p}_{in} = 0.678916$ and $\bar{p}_{out} = 0.177031$. Tested at the significance level of 0.1%, the two values are shown statistically different from each other in all cases. The 95% HPD intervals for p_{in} and p_{out} are in this case [0.5272393, 0.7591251] and [0.1651885, 0.1842609] respectively.

4.3 Discussions

We carried out two assessments to ensure the correctness of our findings. First, we implemented a classical statistical approach, i.e. the grid-search algorithm developed by Čopič, Jackson, and Kirman (2007), to the same sample of the LVTS data, and the results found are identical to what is shown above.[13] Second, we conducted a formal verification of the posterior simulators, known as the *Joint Distribution Test* (introduced by Geweke (2004)), to confirm that our simulation results do not contain any analytical and/or computing errors.[14]

In the Joint Distribution test, the test functions used are the mean and variance of p_{in}, p_{out} and the number of components averaged over all observed partitions, as well as the covariances across the three parameters. We use 2.5×10^5 iterations of each of the *marginal-conditional* and *successive-conditional* simulators. *Table 2*

[12]We applied Geweke's *Joint Distribution Test* to make sure of no analytical and/or computing errors in the MCMC simulations. Details of the test are provided in the next section (section 4.3). Therefore, we conjectured that this might be a multi-modal surface. The difficulty in the chain convergence can potentially be resolved in future work by introducing a hierarchal structure into the model.

[13]For details of the algorithm, please see Čopič, Jackson, and Kirman (2007).

[14]For details of the Joint Distribution test, please see Geweke (2004).

presents the test results showing that none of the nine test statistics failed the test at all four levels of significance.

Table 2: Summary of the Joint Distribution Test Results

	Tests Failing (yes or no) at p-value=			
Test Statistic	*0.05*	*0.01*	*0.005*	*0.001*
p_{in}	no	no	no	no
p_{out}	no	no	no	no
$mean(\#C)$	no	no	no	no
p_{in}^2	no	no	no	no
p_{out}^2	no	no	no	no
$[mean(\#C)]^2$	no	no	no	no
$p_{in} \times p_{out}$	no	no	no	no
$p_{in} \times mean(\#C)$	no	no	no	no
$p_{out} \times mean(\#C)$	no	no	no	no

One interesting result from the joint distribution test is the amount of information used to pin down the partition. The way to see this is that each day contains information on $14(14-1)$ links between banks this implies we have 118300 data points to pin down three parameters. Thinking of it this way it is clear that there is a large amount of data to pin down a given partition.

The empirical results found in this study highlight two main points.[15] [16]First, the measure of capacity and transaction intensity has a direct effect on what kind of community structures (as well as the interpretation of these structures) that can be uncovered by this maximum likelihood (ML) method. Second, for a highly heterogeneous network such as LVTS, a simple model with only two parameters is probably much too simple to capture the patterns of relationships, despite the LVTS being

[15]In addition to the data discussed in the paper, we also experimented with various other data samples: (i) five weekday sub-samples (i.e. Mondays, Tuesdays, Wednesdays, Thursdays and Fridays) of the same base period (June 1st, 2005 - December 31st, 2007), (ii) pre-financial-crisis phase (July 2006 - June 2007), and (iii) the financial-crisis period (September 2008 - July 2009). In all these case, the *Averages* measure revealed results identical to those presented earlier in the paper. Under the *Liquidity* measure, the results are unchanged for most of the new data samples, however, there are two exceptions. During the period of financial crisis, the result shows a cluster of three Toronto-based banks with all remaining participants being singletons. And in the case of Tuesdays, we found an even smaller Toronto cluster made of two banks only with the remaining banks as singletons; this result is probably related to the fact that most statutory holidays fall on Monday and therefore transaction intensity in LVTS may show differently.

[16]Future work can look into the different results that may arise if we exclude client-initiated payments from the data and focus on inter-bank transactions only.

a truly small network. This also in part explains why a second small community appears under the *Averages* measure but is not visible under the *Liquidity* measure. The *liquidity* measure constructs the interaction intensity among LVTS participants by completely drawing on transaction value and the intraday liquidity, whereas the *Averages* measure is liquidity-based but transformed. Hence, when transaction value is examined directly, it is very likely that the payment flows to and from the five big Toronto-based banks are so strongly dominant that other small communities in the network will be overwhelmed.

Since the Averages measure uncovers communities of banks based on their geographic proximity, we conducted the following two sensitivity studies. First, we restricted the sample to a subset of the banks to see whether or not a richer model is needed. Second, using Čopič, Jackson, and Kirman (2007)'s classical statistical approach, we examined the partitions which are "likely" but not the ML estimate of the partition.

We removed the larger Toronto community from the data and reran the MCMC simulator on the remaining subset of banks. While these results should be used with caution they do provide an idea of whether or not a richer model (i.e. more parameters for inter- and intra-community transactions) is needed; since we are now essentially allowing the two parameters p_{in} and p_{out} for the remaining subset of LVTS participants to vary independently of the data on the five large banks.

The results from the re-estimation show that the geographic structure is still in the data but another one bank based in the same province now joins the Montreal group and the remaining small participants form singleton communities. The estimated values of p_{in} and p_{out} in this case are 0.3629 and 0.0173 respectively; which are remarkably different from the estimation obtained on the full set of data. Both the intra-community and the inter-community transaction probabilities are much lower than what are shown in the full data set, which implies that the results from the full-sample estimation are mostly driven by the transaction flows to, from and within the large Toronto community. This result suggests that a richer model is probably needed to reveal a hierarchical structure of the LVTS community, i.e. by estimating extra parameters (additional p_{in}'s and p_{out}'s) at different levels of communities.

One advantage of using Čopič, Jackson, and Kirman (2007)'s grid-search estimation method is that it provides a complete ranking of all possible partitions of a network. Hence, we also examine the second and the third most likely community structures of the LVTS under each of the two measures of transaction intensity. Under the Liquidity measure, the second best partition found is the same as shown in *Figure 2*. That is, in addition to the group of five big Toronto banks, the partition also contains a small component comprised of two LVTS participants whose operations are based in Montreal. Under the averages measure, the second most likely partition is similar to the results from the sub-sample re-estimation above. More specifically, it shows a big cluster of five large Toronto participants, a small Montreal group that consists of three financial institutions and the rest of system members forming their singleton communities.

Both sensitivity analyses suggest that the results found in this study are in general useful and valuable; however, it is likely that a richer specification of the model would help find a complex community structure that more accurately describes the structure of the business relationships among the LVTS direct participants.

5 Concluding Remarks

In this paper, we implemented on the LVTS social network the techniques developed by Čopič, Jackson, and Kirman (2007), a new network-partitioning model based on likelihood estimation.

This likelihood approach provides a unique probabilistic perspective on identifying community structures from sophisticated network data, by emphasizing what a community structure is and how a particular partition can be formed. More specifically, it allows us to find transaction patterns among LVTS participants solely based on the outcomes of payment exchanges in the system.

The major finding is that using an appropriate measure of transaction intensity, we find that the most likely partition of the network of LVTS direct participants consists of two clusters. One is a group of five large Toronto-based banks, and the other is a small community of financial institutions whose operations are based in the city of Montreal.

Our results show that no matter how the transaction intensity among LVTS participants is defined, the cluster of the five big banks is always strongly dominant in the network. The fact that it is difficult to identify additional communities among the remaining small participants suggests that we should extend the model to introduce hierarchies to the network structure. For example, we can add an extra pair of parameters p_{in} and p_{out} for small participants. Both the LVTS network data and the results from this study show that it is very unlikely for the small LVTS participants to have the same probabilities of interacting with other system members within the same group and/or across the group as do large banks.

Other steps in our future work include changing the model specifications (not based on strict binomial distribution, which will broaden the scope of defining the o and c matrices), using daily frequency data, and comparing this likelihood method with that of Bergstrom (2007).

References

Allen, F., and D. Gale (2000): "Financial Contagion," *journal of political economy*, 108(1), 1–33.

Bech, M. L., J. T. Chapman, and R. J. Garratt (2010): "Which bank is the "central" bank?," *Journal of Monetary Economics*, 57(3), 352–263.

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

Bergstrom, C. (2007): "Eigenfactor: Measuring the Value and Prestige of Scholarly Journals," *C&RL News*, 68(5).

Chib, S., and E. Greenberg (1995): "Understanding the Metropolis-Hastings Algorithm," *Journal of the American Statistical Association*, 49(4), 327–335.

Čopič, J., M. O. Jackson, and A. Kirman (2007): "Identifying Community Structures from Network Data via Maximum Likelihood Methods," Mimeo Standford University.

Geweke, J. (2004): "Getting it Right: Joint Distribution Tests of Posterior Simulators," *Journal of the American Statistical Association*, 99(467).

O'Conner, S., J. T. Chapman, and K. Millar (2008): "Liquidity Efficiency and Distribution in the LVTS: Non-Neutrality of System Changes Under Network Asymmetry," Bank of Canada Discussion Paper 2008-11.

Soramaki, K., M. L. Bech, J. Arnold, R. J. Glass, and W. E. Beyeler (2007): "The Topology of Interbank Payment Flows," *Statistical Mechanics and its Applications*, 379(1), 317–333.