Physica A 424 (2015) 44–51

Contents lists available at ScienceDirect

Physica A

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

Mapping change in the overnight money market

Morten L. Bech [a], Carl T. Bergstrom [b,c], Martin Rosvall [e], Rodney J. Garratt [d,*]

[a] *Bank for International Settlements, Switzerland*
[b] *Department of Biology, University of Washington, United States*
[c] *External Faculty, Santa Fe Institute, United States*
[d] *Federal Reserve Bank of New York, United States*
[e] *Integrated Science Lab, Department of Physics, Umeå University, Sweden*

HIGHLIGHTS

- We use a network clustering approach known as the *map equation*.
- We identify changes in lending patterns in the US overnight money market.
- Dramatic changes in lending patterns occur after the Federal Reserve begins paying interest on reserve balances.
- Analysis of micro-scale rates of change suggests these changes were triggered by the collapse of Lehman Brothers.

ARTICLE INFO

Article history:
Received 7 May 2013
Received in revised form 27 June 2014
Available online 5 December 2014

Keywords:
Network
Money market
Topology
Mapping change
Alluvial diagram

ABSTRACT

We use an information-theoretic approach to describe changes in lending relationships between financial institutions around the time of the Lehman Brothers failure. Unlike previous work that conducts maximum likelihood estimation on undirected networks our analysis distinguishes between borrowers and lenders and looks for broader lending relationships (multi-bank lending cycles) that extend beyond the immediate counterparties. We detect significant changes in lending patterns following implementation of the Interest on Required and Excess Reserves policy by the Federal Reserve in October 2008. Analysis of micro-scale rates of change in the data suggests these changes were triggered by the collapse of Lehman Brothers a few weeks before.

1. Introduction

The overnight money market is an important part of the US financial system. Banks and certain other financial institutions use this market to reallocate liquidity among each other. As such, the market is the marginal source of funding for many financial institutions and the rates at which funds trade play a key role for monetary policy implementation, financing arrangements, and in the extension of credit in the economy.

In this paper, we demonstrate how a network clustering technique known as the *map equation* [1,2], that has been used successfully to shed light on the structure of scientific publishing, can be used to simplify and highlight important aspects of the overnight money market. Our analysis focuses on the period of unprecedented stress that hit the US financial system in the autumn of 2008. We show how the map equation can be used to better understand how the market responded to key events that occurred during this period, including the collapse of Lehman Brothers and the introduction of the Interest

* Corresponding author.
E-mail address: Rodney.Garratt@ny.frb.org (R.J. Garratt).

http://dx.doi.org/10.1016/j.physa.2014.11.034
0378-4371/

on Required and Excess Reserves (IOER) policy initiated by the Federal Reserve, and form testable hypotheses on behavioral aspects of market participants.

Analysis of the overnight money market is complicated by the fact that there are several hundred active participants in this market [3,4]. Hence, it is difficult to identify patterns or changes in behavior without advanced techniques. We model the flow of payments generated by loans between institutions as a large weighted and directed network. Each node of this network represents an individual institution. Each link represents a loan; the value of loans between node A and node B determines the weight on the link from A to B. Given a network partition, the map equation measures the per-step description length of movements of flow on a network. Minimizing the map equation over all possible network partitions gives an optimal clustering with respect to the dynamics on the network. In the original formulation of the map equation, the flow was induced by a random walker guided by the directed and weighted links of a network. Because here the weighted and directed links already represent flow, which need not be ergodic, we measure the description length of the raw data and do not encode the movements necessary to make the flow ergodic [5,6]. In this way, we are able capture the actual payment flows associated with lending activity and reveal clusters of banks for which there is long persistence time.

To distinguish meaningful structural change from mere noise in the data, we perform parametric bootstrap re-sampling of the networks and assess the significance of each cluster [7]. Our maps reveal significant, tractable changes in clustering over the most intense period of the recent financial crisis. In particular, there are dramatic changes in lending patterns following the decision by the Federal Reserve to start paying interest on depository institutions' required and excess reserve balances.[1] Changes in lending patterns identified by the map equation may be a culmination of changes that began at an earlier point in time. Analysis of micro-scale rates of change in the data suggests the collapse of Lehman Brothers was the driving force behind the clustering changes we observed.

2. Related literature

Refs. [8,9] identify the network structure of two types of financial networks using the maximum likelihood approach developed in Ref. [10]. These works determine the network structure that is most likely to generate observed flows assuming that flows within groups are more likely than flows across groups. Both studies apply the technique to small networks: Ref. [8] considers the sterling unsecured loan market in the United Kingdom during 2006–2008 (12–13 banks), while Ref. [9] considers payment flows through the Canadian Large Value Transfer System from (14 banks). Ref. [11] examines interbank lending for the much larger, German banking system using a related approach in which a core–periphery structure is assumed, i.e., specific borrowing and lending relationships are assumed that apply to core versus non-core banks, and the assignment of banks to the core is selected which minimizes errors.[2] Both Refs. [9,11] consider pre-crisis data. Ref. [9] finds that a relatively stable, core–periphery structure emerges endogenously in the Canadian payment system with five core banks. Ref. [11] finds a relatively stable core size from 1999 until the middle of 2006 at which point the core size fell sharply from around 45 banks to 35 banks. Ref. [8] identifies a large increase in the core size and overall connectivity in the sterling unsecured loan market following Lehman's default.

The main difference between our approach and the maximum likelihood approach is that clustering arrangements in the latter are based solely on each node's direct connections and all links are undirected—meaning that no distinction is made between the borrower and the lender. This approach is appropriate for identifying changes in direct lending relationships, but it may not detect broader dependencies that exist through intermediate banks, nor does it help identify how funding shocks might propagate through the network. Under the map equation approach, both the direction and weights of all network linkages are utilized and direct and indirect linkages matter. Bank A will not tend to be clustered together with bank B if the lending relationship is unidirectional unless there is a broader lending cycle which connects the two banks through other banks. We argue that this approach may be preferred for identifying areas of funding risk. Moreover, instead of encoding a random walker that occasionally teleports to a random node for an ergodic solution, as in the original formulation of the map equation, we use a version that only encodes actual flows between banks without any need for teleportation [6]. In this way, our results only depend on the data and not on any parameter choice.

3. Data

Unfortunately, no data sources cover the entire overnight money market, at least in the US.[3] This is in part due to the over-the-counter nature of the market and the different platforms used to clear and settle transactions. Here, we look at

[1] Coinciding with IOER, the Federal Reserve's balance sheet expanded dramatically as reserves were added to the system, initially through credit and liquidity facilities created to support the financial markets during the crisis, and then by the Quantitative Easing. Since the fall of 2008, depository institutions in the United States have accumulated over $2.5 trillion in excess reserves, compared to the pre-crisis level in single digit billions. For more on IOER see http://www.federalreserve.gov/monetarypolicy/reqresbalances.htm.

[2] For a larger network, like the one considered in Ref. [11] (around 2000 banks), it is not feasible to examine all possible clustering combinations. The added structure imposed by Ref. [11] allows for a greedy algorithm.

[3] See http://www.federalreserve.gov/newsevents/press/bcreg/20130625a.htm for details on new data collection requirements related to selected money market instruments.

a subset of the market. The subset consists of overnight loans that settled over the Federal Reserve's large value payment system (Fedwire). Overnight loans are not specifically identified in the Fedwire transaction journal and thus we extract an estimate of these loans from using the Furfine algorithm [12]. The basic idea behind the Furfine algorithm is as follows: a pair of payments between two accounts in Fedwire on consecutive business days is identified as a potential overnight money market loan if the amount transferred on the first day is reversed on the second day with a plausible amount of interest.

Past work has used the Furfine algorithm to extract Federal Funds transactions, which are a subset of overnight loans.[4] However, the Research Group of the Federal Reserve Bank of New York has recently concluded that the output of its algorithm based on the work of Furfine may not be a reliable method of identifying federal funds transactions.[5] This paper therefore refers to the transactions that are identified using the Research Group's algorithm as overnight or term loans made or intermediated by banks. Our reference to overnight or term loans made or intermediated by banks in this paper to describe the output of the Research Group's algorithm is not intended to be and should not be understood to be a substitute for or to refer to federal funds transactions.

Potential problems with the Furfine algorithm are detailed in Ref. [16]. A couple of issues are relevant here. First, we are not able to tell whether a transaction is an unsecured loan (e.g. federal funds or euro–dollar) or the cash leg of bilateral repurchase agreement. Second, we are not able in all cases to pin down the ultimate counterparties to a transaction. Overnight transactions flowing in and out of Fedwire accounts can either be driven by proprietary trading by the account holders or due to customer activity. For example, a transaction between say a US broker–dealer and a bank in Europe (which both do not have Fedwire access) is attributed to the two Fedwire participants that manage the respective US dollar funds on behalf of their clients.

The former issue is less of a problem here, as we are interested in the flow of funds between institutions rather than the rates at which they trade.[6] The second issue is potentially more problematic as links from a node may represent relationships between more than two financial institutions. For that reason, we focus in the subsequent discussion on accounts of Fedwire participants for which counterparty misidentification has been found to be less of an issue.[7]

The Furfine algorithm estimates that almost 900 accounts were actively settling loans during the second half of 2008, but many of these accounts were not active every day. Hence, using only a single trading day's data to construct lending networks may tend to produce networks that are too sparse to reflect interesting counterparty relationships in the periphery. Instead, we aggregate the data over multiple days. The number of days chosen reflects a trade-off between being long enough to identify lending relationships and being short enough not to conceal structural changes. Specifically, we choose the two-week maintenance periods that the Federal Reserve applied for computing reserve requirements for the majority of banks at the time.[8] Reserve requirements are the amount of funds that banks must hold in reserve against specified deposit liabilities.[9] If banks fail to meet their reserve requirements then charges will be applied by the Federal Reserve. Maintenance periods start on a Thursday and end on a Wednesday. Banks that lend to each other within the two weeks will be connected in our networks and weights on links will reflect the amount of activity.

For the purposes of illustrating changes in lending relationships we show every maintenance period starting July 3, 2008 and ending December 31, 2008. The collapse of Lehman Brothers occurred during the maintenance period from September 11 to September 24, 2008, labeled Sep 24 in the diagrams. The IOER policy began at the start of the maintenance period which ran from October 9 to October 22, 2008 and is labeled Oct 22 in the diagrams.

4. Method and results

Details of the clustering algorithm are available in Rosvall and Bergstrom [1]; see also Garratt et al. [17] for an application to international claims and liabilities between countries. A full description of how to construct the alluvial diagrams we use to show structural changes in lending are found in Rosvall and Bergstrom [7]. The alluvial diagrams in Fig. 1 show the progression of change in clustering using the 13 two-week maintenance periods that cover the second half of 2008. Each column in the alluvial diagram represents the clustered network of a period, labeled by the last day of the period. Each block in a column represents a cluster in the network and the height of each block represents the total value of loans that pass through institutions in the cluster. Confidentiality requirements prevent us from identifying individual members of the clusters. Consequently, each cluster is named by an anonymized 3 character code of the institution with the highest transaction volume in the cluster, e.g. "KA8", and followed by ", . . ." if there are more than one bank in the cluster.

Clusters are stacked in order from largest to smallest. Darker shading represents nodes that are assigned with statistical significance, while lighter shading represents non-significant assignments. Significantly identified clusters are separated by

4 Refs. [13,4,14,15], among others.

5 It should be noted that for its calculation of the effective federal funds rate, the Federal Reserve Bank of New York relies on different sources of data, not on the algorithm output.

6 For this reason, we are not concerned if the algorithm picks up term loans. Ref. [16] reports that conversations with Fed funds traders suggest that term loan activity is likely to be small.

7 For example, Ref. [16] finds that the Furfine algorithm may perform well for government-sponsored enterprises (GSEs).

8 Effective June 27, 2013, a common two-week maintenance period for all depository institutions was created.

9 See http://www.federalreserve.gov/monetarypolicy/reservereq.htm.

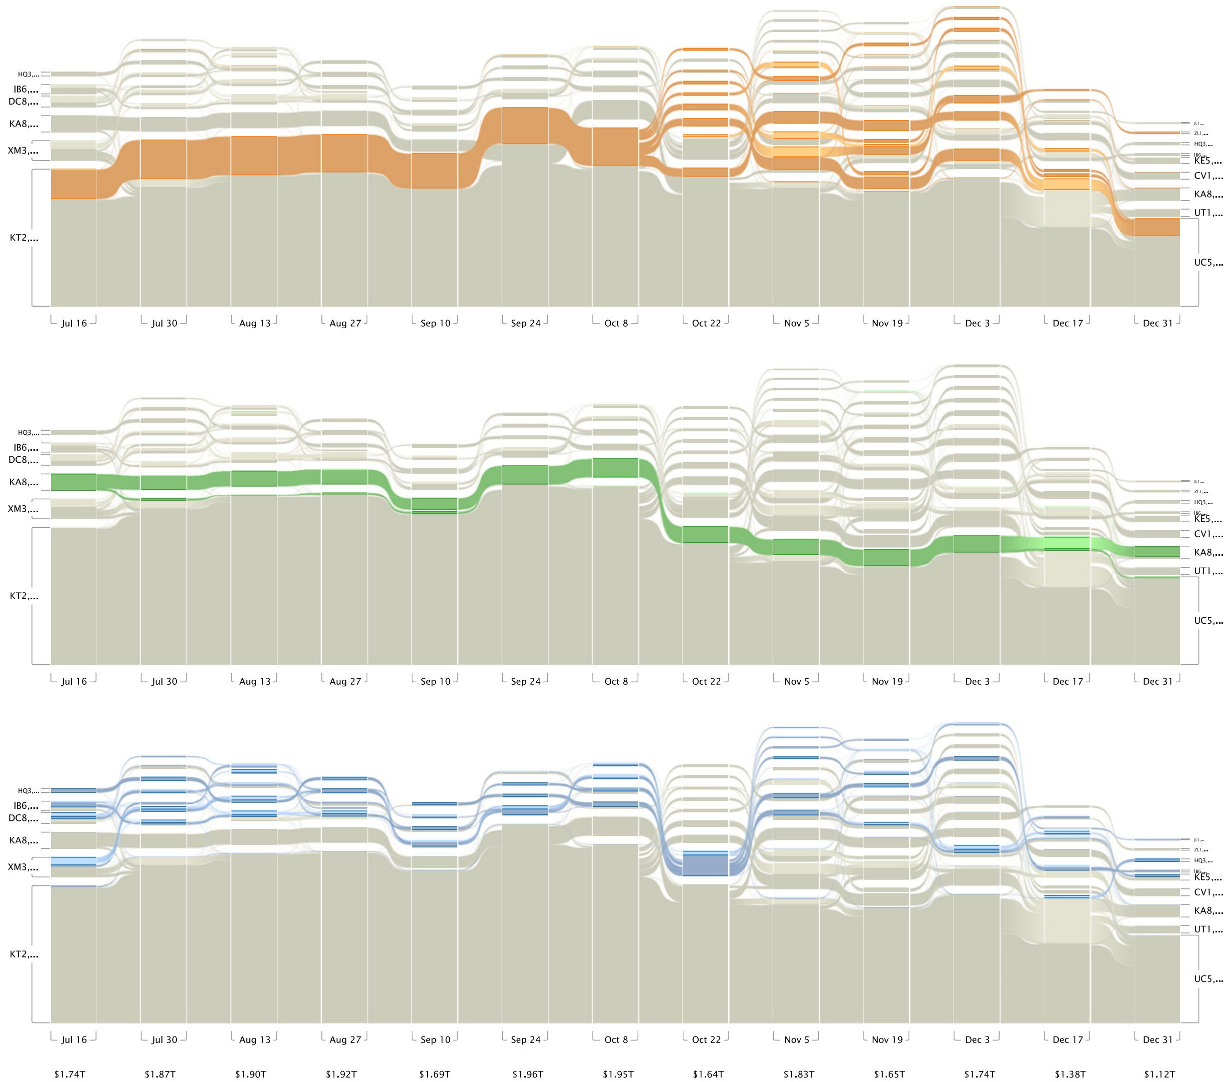

Fig. 1. Mapping change of payment flows driven by overnight money market activity from July 2008 to December 2008. From top to bottom, the alluvial diagrams highlight three structural changes: banks that segment in October (orange), a group of banks that intermittently separate from the largest cluster up until October 8th and are consistently included in it afterwards (green), and banks that combine to form a single cluster in early October (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

thick (white) vertical spaces. Changes in the clustering structure from one time period to the next are represented by the mergers and divergences that occur in the streamlines linking the blocks from one period to the next.

We perform two exercises to illustrate the changes reflected in the clustering analysis. First, we highlight groups of institutions that were together, either as part of a larger cluster or as their own cluster, at the start of the sample period and ask: how are these institutions clustered in the future? Second, we highlight banks that were clustered together on a certain date and ask: which clusters did they come from and where do they go?

Two illustrations of the first exercise are shown by orange and green streamlines in Fig. 1. Banks highlighted in orange in the top panel of Fig. 1 formed part of the largest cluster until the Oct 22 maintenance period.[10] However, in the maintenance periods subsequent to October 8th these banks separate from the main cluster, and combine into several smaller clusters. Fig. 2 provides a more detailed description of the changes in flow patterns that led to this change. Banks highlighted in orange during the Oct 5 maintenance period were included in the UC5, . . . cluster reflecting the fact that these banks had sufficiently strong cyclical lending relationships with the other banks in this giant cluster. However, during the Oct 22 maintenance period the flows between these institutions and the other members of the cluster became weaker and

[10] The orange-highlighted banks are grouped together within the largest cluster in all maintenance periods prior to Oct 22 for presentation purposes only. There is no special relationship between these banks within the larger cluster up to that point.

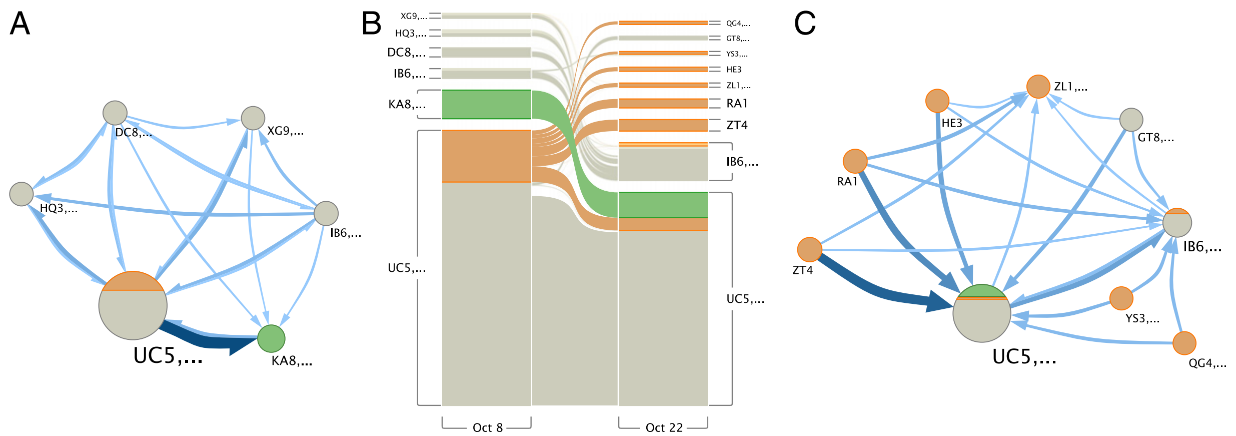

Fig. 2. Mapping change of payment flows driven by overnight money market activity in late September and October 2008. For every two-week period we partitioned close to 900 banks connected by more than 2000 transaction links into clusters of banks; each period is labeled by its last day. The maps in panels A and C show the most important clusters in two consecutive maintenance periods: September 25–October 8, and October 9–October 22. The size of the nodes represents the total value of loans that passes through banks within each cluster, and the size and color of the arrows represent the total payment flows between the clusters. To illustrate the structural change, the diagrams highlight in orange all banks that split off and in green all banks that merge with the main cluster during October. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

more unidirectional, as reflected in panel C. Consequently, these institutions were clustered into smaller subgroups over the Oct 22 maintenance period. Panel A shows a weak cyclical relationship between clusters IB6, . . ., DC8, . . ., HQ3, . . ., and XG9, . . ., which strengthened in the Oct 22 maintenance period leading to their merger in that maintenance period. Finally, Panel B shows the inclusion of cluster KA8, . . . into the large UC5, . . . cluster in the Oct 22 maintenance period. This results from a strengthening of the bilateral flows between the two groups of institutions over two maintenance periods.

Returning to Fig. 1, the banks highlighted in green and identified by the label KA8, . . . in the bottom panel of this figure were clustered together by the algorithm over the Jul 16 maintenance period. If we follow these banks over time we see that in five out of seven of the maintenance periods preceding the Oct 22 maintenance period, these institutions are clustered together in a single module reflecting the existence of strong lending cycles between these institutions. However, from the Oct 22 maintenance period onwards, these institutions are absorbed by the largest cluster. This means that the borrowing and lending behavior of these institutions is no longer highly concentrated amongst its members, but rather significant borrowing and lending activity is occurring within the broader collection of institutions included in the largest cluster.

The bottom panel of Fig. 1 highlights in blue a cluster of institutions that formed in the Oct 22 maintenance period and tracks them back in time to the beginning of the sample period and forward to the end. The interesting observation here is that this clustering seemed to be unique over the sample period. Namely, there was no other period in which anything approximating this relatively large and significant cluster formed. The banking groups highlighted in blue on the Oct 22 maintenance period came from several smaller clusters that coexisted in the previous maintenance periods. Interestingly, this change in behavior was very short-lived. In the following months, these institutions dispersed again into separate clusters.

5. Discussion

After an interlude of relative calm following the rescue of Bear Stearns in March of 2008, concerns about the profitability and asset quality of financial institutions started to mount again over the summer. These tensions came to a head with the bankruptcy of Lehman Brothers in the early-morning hours of Monday, September 15. During the tumultuous days that followed, any hopes that the Lehman bankruptcy was the end of the trouble quickly dissipated. If not for the multitude of actions undertaken by public authorities the Lehman debacle might easily have been the beginning of the end.[11]

A byproduct of the Federal Reserve's interventions was that the level of reserve balances exploded from $10 billion on average during August of 2008 to $850 billion by year end.[12] Unlike earlier in the financial crisis, the Federal Reserve was not able to sterilize the increase in reserves balances from these operations by selling US Treasuries due to the sheer size of injections required, and this had implications for the overnight money market. Ref. [19] states: "With the banking system awash in funds, the rate at which banks were willing to buy and sell these funds – the federal funds rate – dipped well

11 See Ref. [18] for a description of these events and an account of the steps taken by the federal reserve and other government agencies to ease investor concerns and support US banks and other companies.

12 Federal Reserve Statistical Release, H.4.1, Factors Affecting Reserve Balances, Historical Data, Table 9. http://www.federalreserve.gov/releases/h41/hist/h41hist9.pdf.

Table 1
Hypothetical patterns and clusters.

$t =$	Config	Cluster
0	a	1
1	a	1
2	b	1
3	g	1
4	h	1
5	i	2
6	j	2
7	j	2

below the intended policy target rate set by the Federal Open Market Committee . . .. This situation created a tension for the Federal Reserve: while the increases in liquidity would prove to help improve market functioning, these increases were also exerting downward pressure on the federal funds rate.'' (p. 1)

These events motivated the Federal Reserve to start paying interest on required and excess reserve balances held by banks. The new policy began on October 9, 2008. However, as discussed by Ref. [19], government-sponsored enterprises (GSEs), which are significant sellers of funds on a daily basis, are exempt.[13] Ref. [19] states: ''This heterogeneity across participants . . . created a segmented market with different rate dynamics.'' (p. 2) Moreover, ''. . . a combination of financial consolidation, credit losses, and changes to risk management practices has led at least some GSEs to limit their number of counterparties in the money market and to tighten credit lines.'' (p. 3)

Our maps reveal several aspects of these extraordinary events. While we must refrain from discussing individual institutions, we can identify some general trends. The collection of institutions highlighted in orange is dominated by a set of Federal Home Loan Banks and a number of small and medium sized banks. This cluster disintegrates shortly after the implementation of interest on reserves, likely due to the fact that the lending between the institutions was significantly reduced as the amount of reserves supplied by the Federal Reserve increased. The blue cluster which forms only during the Oct 22 maintenance period is comprised by another Federal Home Loan Bank and a number of banks that tend to be located in the same geographic region as the Home Loan Bank. We speculate that this cluster reflects the fact that the Home Loan bank may have started to intermediate funds between its members by borrowing funds from some and making overnight advances (i.e., collateralized loans) to others during this period. The fairly stable green cluster, which is subsumed into the large cluster after the implementation of interest on reserves, is dominated by a prominent GSE and one large money center bank. The break down of the cluster may reflect the reduction (or even elimination) of the lending relationship to the particular bank by the GSE.

Based on our descriptive analysis, it appears that the Lehman bankruptcy was less of a transformative event for the overnight money market than the implementation of interest of reserves. However, we do not draw this conclusion. While the alluvial diagrams are very nice for showing the general patterns of how lending takes place in each period, they do not necessarily reveal the onset of change in the system. Changes in clustering patterns reflect tipping points at which the cumulative effect of multiple small changes in flows constitute a significant change. For example, suppose that lending configurations a, b, c, d, e, f, g, h all generate a system with module structure of type 1, and configurations i, j, k, l, m, n, o, p all generate a system with module structure of type 2.

Suppose at the following times we are at the following places in Table 1. Looking at the cluster structure, the change appears to have occurred between periods 4 and 5. But measured in terms of the change in underlying configurations, the most rapid change occurs from period 2 to 3, where the system leapt from configuration b to configuration g in a single period. This sets up the possibility of a small change – h to i – tipping the cluster from structure 1 to structure 2. In this example – and perhaps in our data as well – changes in the cluster structure diagram can lag the initial changes in the underlying system that we are trying to comprehend.

In testing hypotheses about changes in lending patterns we should look more closely at the raw data. Next we describe a method that allows us to further differentiate between the two hypotheses that (1) Lehman's failure is associated with a shake-up in lending patterns and (2) paying interest on reserves is associated with the shakeup.

An $n \times n$ lending matrix is specified by a unique vector of length n^2. Thus each time period corresponds to a vector in n^2 space. To measure the amount of change in the system, we look at how much the angle between the vector at each time t changes going to time $t + 1$. This is a standard approach in network theory, known as cosine distance (see Ref. [20]). The results are shown in Fig. 3.

The orange trace compares the distance between each time period and a fixed time period, here Sep 10. This shows us that the change we see in the active periods of interest moves continually away from where we were before, rather than moving away and then returning to some approximation of the original state.

The black trace shows us the velocity of change from one period to the next over time. We see that the largest change occurs in the Oct 8 maintenance period, which covers the time period from September 25 to October 8, after collapse of

13 Ref. [16] find that the Furfine algorithm does a good job in terms of identifying GSEs counterparts to overnight transactions.

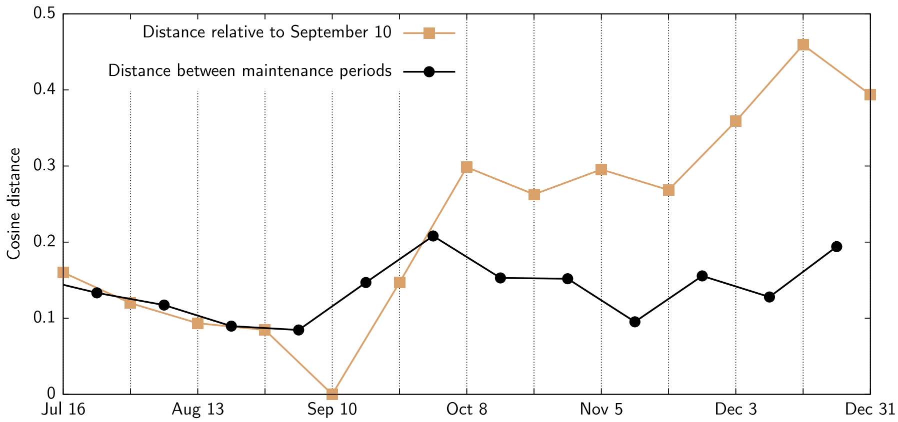

Fig. 3. Cosine distance of payment flows for the overnight money market.

Lehman and before the Federal Reserve starts paying interest on reserves. Is it possible that people were anticipating the change in Federal Reserve policy? Or was paying interest on reserves a surprise to everyone, so that the change from September 25 to October 8 must really be ascribed to other factors, such as the collapse of Lehman? The fact that the policy of paying interest on reserves was not announced until October 6, 2008 suggests it was that latter.[14] Thus, while changes in borrowing and lending patterns do not fully reveal themselves in our maps until after the implementation of IOER, examination of the micro-scale rates of change strongly suggests that the collapse of Lehman Brothers was the driving force.

6. Concluding remarks

Advanced network techniques can help stakeholders in the financial system to understand its structural features and to analyze the impact of transformative events. As illustrated here, the map equation appears to be a very useful tool for understanding funding flows. The lending flows in the overnight money market changed in a fundamental way between September 10 and October 22, 2008, and the alluvial diagrams reveal this clearly.

Prior to performing this analysis, we had suspected that the Lehman bankruptcy in the fall of 2008 would have been the most important shock to the underlying flow dynamics of the overnight money market. The clustering maps suggested that the IOER policy may have had a more pronounced impact. However, in evaluating such hypotheses it is important to recognize that changes in clustering of flow reflect not the instantaneous rate of change of a system, but rather the cumulative effects of such changes. The occurrence of important structural changes tells us that the system has reached a tipping point at which the description length of one clustering configuration falls below another. Thus the underlying shifts in network flows that precipitate these structural changes may have been initiated well before the tipping point was reached. As a consequence, considerable caution is required when deriving causal inferences from alluvial diagrams. While the major structural changes observed here follow the institution of interest of reserves, this does not in and of itself imply a causal link.

Our analysis of micro-scale rates of change in the raw data suggests that the observed changes in lending patterns resulted from the cumulative effects of lending changes that were precipitated by the collapse of Lehman and accumulated over time, reaching the tipping point only after the IOER policy began. This does not rule out the possibility that IOER was a contributing factor, and it may have helped perpetuate the observed changes. Institutions that dissolved lending relationship following the collapse of Lehman may have seen no reason to resume them in a world where they could earn interest on reserves without participating in the overnight money market. It is not possible to completely disentangle the effects of these two, overlapping events.

We advocate the use of two tools for analyzing the structure of financial networks. Each complements the other and neither is sufficient in isolation. The map equation reveals network structure and changes can be visualized via alluvial diagrams. However, these maps do not reveal details of micro-scale rates of change that precipitated change. The cosine distance analysis is useful in this regard, but itself is not informative: it reveals change, but not the content of that change.

Acknowledgments

This paper was written while MLB was at the Federal Reserve Bank of New York and RJG was there as part of the Occasional Visitor program. We thank Vijay Narasiman for excellent research assistance and an anonymous reviewer for

[14] The Financial Services Regulatory Relief Act of 2006 originally authorized the Federal Reserve to begin paying interest on reserves beginning October 1, 2011. The date was pushed up due to the crisis.

helpful comments and suggestions that improved the paper. MR was supported by the Swedish Research Council grant 2009-5344. CTB was supported by National Science Foundation grant SBE-0915005 and by US NIGMS MIDAS Center for Communicable Disease Dynamics 1U54GM088588 at Harvard University. RJG was supported by NSF grant SES-0851357. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York, the Federal Reserve System or the Bank for International Settlements.

References

[1] M. Rosvall, C. Bergstrom, Maps of random walks on complex networks reveal community structure, Proc. Natl. Acad. Sci. 105 (2008) 1118.
[2] M. Rosvall, D. Axelsson, C.T. Bergstrom, The map equation, European J. Phys. 178 (2010) 13–23.
[3] K. Soramäki, M. Bech, J. Arnold, R. Glass, W. Beyeler, The topology of interbank payment flows, Physica A 379 (2007) 317–333.
[4] M. Bech, E. Atalay, The topology of the federal funds market, Physica A 389 (2010) 5223–5246.
[5] M. Rosvall, C. Bergstrom, Multilevel compression of random walks on networks reveals hierarchical organization in large integrated systems, PLoS One 6 (2011) e18209.
[6] R. Lambiotte, M. Rosval, Ranking and clustering of networks with smart teleportation, Phys. Rev. E 85 (2012).
[7] M. Rosvall, C. Bergstrom, Mapping change in large networks, PLoS One 5 (2010) e8694.
[8] A. Wetherilt, P. Zimmerman, K. Soramäki, The sterling unsecured loan market during 2006–08: insights from network theory, Bank of England Working Papers No. 398, 2010.
[9] J. Chapman, Y. Zhang, Estimating the structure of the payment network in the LVTS: an application of estimating communities in network data, Bank of Canada Working Paper 2010-13, 2010.
[10] J. Copic, M. Jackson, A. Kirman, Identifying community structures from network data via maximum likelihood methods, B. E. J. Theor. Econ. 9 (2006).
[11] B. Craig, G. von Peter, Interbank tiering and money center banks, BIS Working Paper No. 322, 2010.
[12] C. Furfine, The microstructure of the federal funds market, Finan. Mark. Inst. Instrum. 8 (1999) 24–44.
[13] A. Ashcraft, D. Duffie, Systemic illiquidity in the federal funds market, Amer. Econ. Rev. 97 (2007) 221–225.
[14] L. Bartolini, S. Hilton, J. McAndrews, Settlement delays in the money market, J. Bank. Finance 34 (2010) 934–945.
[15] G. Afonso, A. Kovner, A. Schoar, Stressed, not frozen: the federal funds market in the financial crisis, J. Finance 66 (2011) 1109–1139.
[16] O. Armantier, A. Copeland, Assessing the quality of "Furfine-based" algorithms, Federal Reserve Bank of New York Staff Reports No. 575, 2012.
[17] R.J. Garratt, L. Mahadeva, K. Svirydzenka, Mapping systemic risk in the international banking network, Bank of England Working Paper No. 413, 2011.
[18] M. Bech, T. Rice, Profits and balance sheet developments at US commercial banks in 2008, Federal Reserve Bull. (2009).
[19] M. Bech, E. Klee, The mechanics of a graceful exit: interest on reserves and segmentation in the federal funds market, J. Monetary Econ. 58 (2011) 415–431.
[20] E. Yan, Y. Ding, Scholarly network similarities: how bibliographic coupling networks, citation networks, cocitation networks, topical networks, coauthorship networks, and coword networks relate to each other, J. Am. Soc. Inf. Sci. Technol. 63 (2012) 1313–1326.