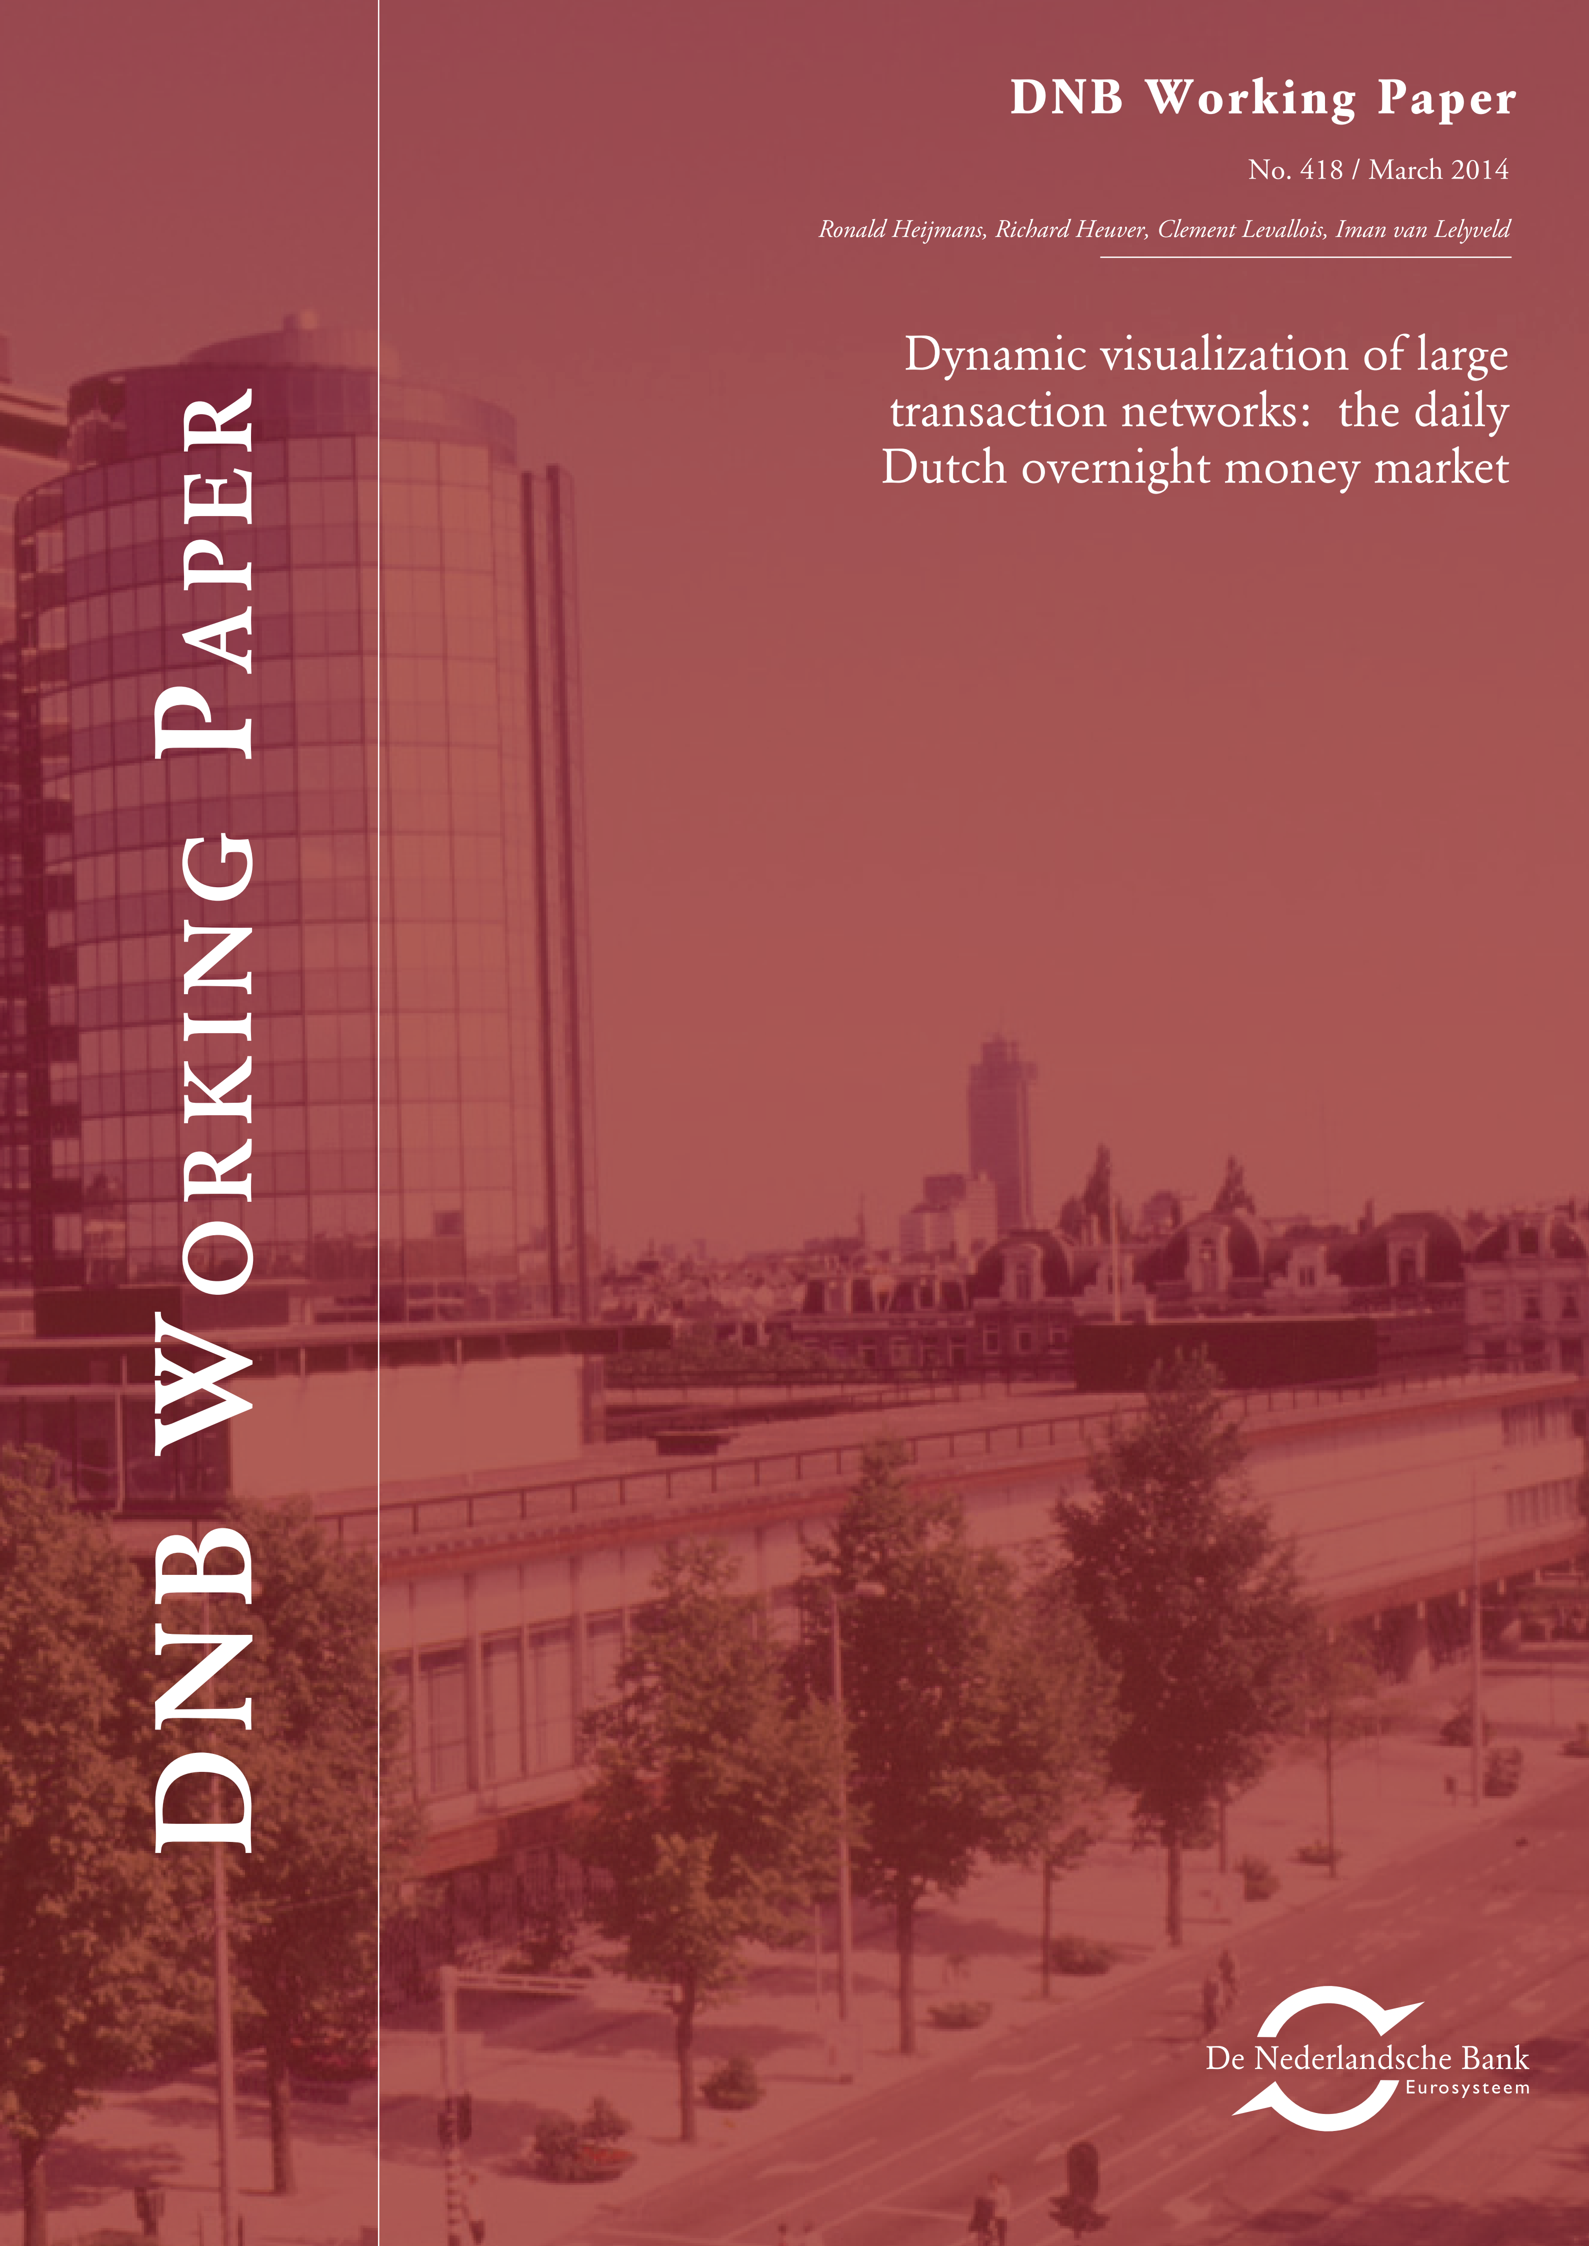

DNB Working Paper

No. 418 / March 2014

Ronald Heijmans, Richard Heuver, Clement Levallois, Iman van Lelyveld

Dynamic visualization of large transaction networks: the daily Dutch overnight money market

DNB Working Paper

De Nederlandsche Bank
Eurosysteem

Dynamic visualization of large transaction networks: the daily Dutch overnight money market

Ronald Heijmans, Richard Heuver, Clement Levallois, Iman van Lelyveld *

* Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank.

Working Paper No. 418

March 2014

De Nederlandsche Bank NV
P.O. Box 98
1000 AB AMSTERDAM
The Netherlands

Dynamic visualization of large transaction networks: the daily Dutch overnight money market

Ronald Heijmans[a], Richard Heuver[a], Clement Levallois[b], Iman van Lelyveld[a,c]

[a] *De Nederlandsche Bank, PO Box 98, 1000 AB Amsterdam, the Netherlands*
[b] *EMLYON Business School, 23 avenue Guy de Collonges, 69130 Ecully, France*
[c] *Bank for International Settlements, Centralbahnplatz 2, Basel, Switzerland*

Abstract

This paper shows how large data sets can be visualized in a dynamic way to support exploratory research, highlight econometric results or provide early warning information. The case studies included in this paper case are based on the payments and unsecured money market transaction data of the Dutch part of the Eurosystem's large value payment system, TARGET2. We show how animation facilitates analysis at three different levels. First, animation shows how the market macrostructure develops. Second, it enables individual banks that are of interest to be followed. Finally, it facilitates a comparison of the same market at different moments in time and of different markets (such as countries) at the same moment in time.

Keywords: interbank network, visualisation
JEL classification: D23, E42, E44, E58, G01, G2, G21.

☆The views expressed in this paper are those of the authors and do not necessarily represent the views of the institutions with which they are affiliated. Corresponding author: Richard Heuver richard.heuver@dnb.nl

Thursday 10th April, 2014

1. Introduction

Visual exploration reveals insights that may be overlooked when using traditional forms of numerical analysis. For example, visual exploration may be used to discover emergent properties in large network structures found in social sciences (Hausmann and Hidalgo, 2013), biology (Wong, 2012) and physics (Borgatti et al., 2009; Borgani et al., 2004). In the words of statistician John Tukey, the picture-examining eye is the best finder we have of the wholly unanticipated (Tukey, 1980).

In 1980, when Tukey wrote these words, visual exploration was generally limited to making drawing on graph paper. Dramatic progress has been made since, providing conceptual clarity and practical techniques to map networks and explore their properties interactively. However, two fundamental areas remain challenging: the size of data sets and effective ways to handle longitudinal (or dynamic) data sets.[1]

The general view is that effective visual exploration requires graphs to display no more than several hundred nodes and a few thousand edges.[2] Beyond these dimensions, most graphing solutions are no longer effective because the display becomes too cluttered to discern relevant patterns. An alternative approach is to aggregate data points until the network is reduced to a manageable size. This approach is valid when the aggregated units themselves are of interest. However, in many cases the interactions and motives at the micro-level are of prime interest and these interactions remain hidden at the aggregated visualization level.

Longitudinal data pose additional difficulties. There are many examples of how to design static graphs. However, there are few examples of how to design dynamic visualizations that provide useful information.[3] Technologically speak-

[1]The terms 'longitudinal' and 'dynamic' both refer to the time-evolving aspect of data sets.

[2]The term 'nodes' refers to the members in a network and are sometimes also called units, entities, persons, individuals, organizations or simply points. The term 'edges' refers to the links between nodes and are often also called relations, links, arcs, ties or interactions.

[3]See for instance the seminal work of Bertin (originally published in French in 1967 but available in English as Bertin (2010) and popularized by Edward Tufte (see `http:\www.edwardtufte.com`).

ing, it is more challenging to represent a dynamic network than a static network. This is due to the fact that dynamic network representation requires more complex data structures and software architectures, and it involves an increase in the computational load. This calls for creative solutions to represent dynamic data sets as static images. The time dimension can, for instance, be represented by a static graph whose horizontal axis depicts the time (e.g. Figure 4 of Pröpper et al. (2013), cf Schwabish (2014)) or, alternatively, by several snapshots of the network at different moments in time shown side by side (e.g. Bech et al., 2011). While these methods can be used effectively (Moody and Mucha, 2013), the scope for visual exploration is greatly reduced when set against an animated view.

Our paper aims to report on recent advances in the visualization of large dynamic networks and to highlight the analytical value of these types of representation. For illustrative purposes, we have used the network of financial transactions between banks recorded in the Dutch part of TARGET and TARGET2, the Eurosystem's real time gross settlement system for large value payments (1999-2013).

The case studies in this paper focus on 1) monitoring the interbank payment system and 2) analyzing the unsecured interbank money market. Monitoring the *interbank payment system* is important because this payment system plays a key role in the economy. Many financial transactions of agents (such as banks, companies and consumers) are settled directly or indirectly in this system. Large disruptions in this system may therefore have a serious impact on the real economy.[4] We show how visualization of the payment network can assist in monitoring daily and weekly payment connectivity patterns.

The reason for discussing the *interbank money market* as a case study is threefold. First, this market plays an important role for central banks in steering interest rates. Second, the unsecured interbank money market is an over-the-counter (OTC) market for which bilateral transaction level data are hard to come by. That

[4]Given the importance of interbank payment systems, they must meet high standards (CPSS, 2012).

is why there has been relatively limited research in this area. However, such data have recently become available (see Heijmans et al., 2010; Arciero et al., 2013), providing a time series of virtually all bilateral unsecured interbank money market transactions since these data were recorded. At the time of writing we were not aware of any other papers that show developments in this key market. Third, the European interbank market has seen considerable changes since the fall of Lehman Brothers in 2008 and the sovereign debt crisis that followed later on.[5] Such variability suggests that dynamic visualization can support the reconstruction of these developments and contribute to their explanation.

Our visualizations show in detail how distress affected the market players after the collapse of Lehman Brothers in the fall of 2008. Following the Lehman collapse, many authorities around the world were worried about the interbank market completely freezing up. However, due to the absence of extensive money market transaction data with sufficient coverage, authorities were unable to make a full assessment of the liquidity deficit. Later on, during the European sovereign debt crisis involving the GIIPS countries (Greece, Italy, Ireland, Portugal and Spain), similar uncertainties emerged. Banks located in the GIIPS countries faced increased sovereign risk premiums while cross-border liquidity flows to these countries declined (BIS, 2012). A visual representation of this market would certainly have been helpful at the time.

This paper is laid out as follows. First, Section 2 reviews the state of the art in visualizing large dynamic networks. Section 3 describes the key role of the interbank money market in the financial sector. Section 4 describes our case studies and Section 5 concludes the paper with recommendations for further work in this area.

[5] See ECB (2010b) and van Riet (2010) for changes in Eurosystem monetary policy during the crisis and Abbassi et al. (2013) for the effects on the interbank market.

2. A review of dynamic visualizations of complex networks

2.1. Longitudinal social networks

The development of software that can help visualize evolving networks is relatively recent (see Table 1). Starting in the early 2000s, several software programs introduced features to create step-wise animations (for instance, SoNIA and Visone).[6] The step-wise technique used by these programs is well suited to represent the dynamics of social networks, where the set of nodes is typically stable while the relations between nodes evolve at regular intervals. One example is a network comprising a fixed group of students whose friendship relations are recorded every month.[7]

The typical animation produced is a series of snapshots of the network at successive moments in time. This raises the question as to how to transit from one snapshot to the next (i.e. the preservation of the 'mental map'). If two successive snapshots are very dissimilar (as is the case in the above example where student friendships are fickle and are subject to significant changes from one month to the next), the sense of coherence and continuity in the network structure is lost to the viewer (Brandes and Wagner, 1997). Each software program has its own approach to dealing with this important issue by using different transition effects ('tweening' in animation terms). For instance, SoNIA and Visone's representation shows distinct jumps from one moment in time to the next.[8]

2.2. Large networks

Advances in computational methods in the last ten years have facilitated numerical analysis of very large networks. For example, researchers at the Università degli Studi di Milano examined 721 million active Facebook users and the 69 billion friendship links connecting them.[9] Standard desktop computer applications are now capable of processing networks comprising hundreds of millions

[6] Erten et al. (See 2003) and Baur and Schank (2008).

[7] See Snijders (2001) for the general approach and Knecht (2005) for an illustration

[8] SoNIA: `http://tinyurl.com/Sonia-viz`, Visone: `http://tinyurl.com/Visone-viz`.

[9] Ugander et al. (2011).

Table 1: Software for large scale or dynamic network visualisations

	Distinctive features	Dynamic	Number of nodes and edges	Initial release (active development)
Cytoscape[1]	Extensible with plugins	Experimental feature via a plugin	1,000+	2002 (yes)
Gephi[2]	Extensible with plugins	Yes	100,000+	2006 (yes)
GraphInsight[3]	Emphasis on 3D rendering	No	100,000+	2012 (yes)
NodeXL[4]	Based on Microsoft Excel, extensible with plugins	Experimental feature via a plugin	1,000+	2006 (yes)
SoNIA[5]	First software to handle dynamics	Yes	100+	2002 (no)
Visone CREEN edition[6]	Pajek Interface for dynamic networks	Yes	1,000+	2008 (no)
Financial Network Analyzer [7]	Online service	Yes	1000+	2010 (yes)

Note: All these packages are available with a license for academic research. See for an extensive overview Gort and Gort (2007). 1) http://www.cytoscape.org/, 2) http://www.gephi.org, 3) http://www.graphinsight.com/, 4) http://nodexl.codeplex.com/, 5) http://www.stanford.edu/group/sonia/ and 6) http://www.leydesdorff.net/visone/, 7) http://www.fna.fi/.

of nodes and edges.[10] However, *visual* network analysis remains restricted to much smaller networks given that heavy computational loads are involved in displaying a large number of objects with visual attributes (position, color, label and shape). Until recently, screen-based network visualizations (created with desktop applications such as NodeXL, Netdraw or GUESS) were limited to a maximum of several hundred nodes and as many edges. However, in 2012 the GraphInsight desktop application was enhanced to enable the visualization of graphs with 5 million nodes and 4 million edges including a 3D rendering engine (Nicolini and Dallachiesa, 2013). Due to the limited performance of web browsers, web-based visualizations have traditionally been even more limited in size. However, the performance of Java script libraries has improved, currently enabling static visualization of networks comprising several thousands of nodes and edges (Kashcha, 2013).

2.3. Transaction data and high-frequency data networks

Networks of agents with high-frequency connections do not easily fit into the framework of longitudinal social network visualizations. Although this is only a matter of frequency level (the connections are merely more frequent), there are important technical and conceptual consequences to consider in achieving effective visualizations.

One example is a network of individuals' mobile communications, essentially a network that is not characterized by a 'wave of observations'. Instead, the network is a continuously evolving set of agents with many short-lived connections (Blondel et al., 2012). Similarly, it may be better to view a network of financial agents engaged in high-frequency trading as a continuously evolving network rather than a collection of snapshots. As a consequence, the animation of such networks cannot be conceptualized as a series of successive graph layouts: to visualize a network with new connections forming every minute over a period of six months would imply the computation of 259,200 graph layouts, not to men-

[10]See `http://tinyurl.com/Pajek-viz`.

tion the difficulty of aligning each consecutive snapshot for the preservation of the mental map.

A solution to the problem of representing high-frequency networks is to give a primordial role to the algorithmic procedure that lays out the network and to make the algorithm run continuously (i.e. without a stopping condition). Once the algorithm is launched, and without the user intervening, the recalculation and any repositioning of nodes on the canvas will indefinitely, even when an equilibrium is reached (which is when the position of nodes no longer changes significantly).

Such a continuously running procedure supports the use of a sliding time window displaying the longitudinal data: the algorithm calculates the position of only the nodes in this 'visible' range of the graph. As the window moves, nodes and edges with out-of-range time stamps leave the view (as they are no longer in the visible range). At the same time, new nodes and edges start to appear (which were previously outside the visible range). The layout thus adjusts smoothly as the algorithm processes these incremental changes.

2.4. A solution for dynamic visualizations of large transactions networks

Continuously running layout algorithms provide a practical solution for dynamic visualizations large transaction networks. However, the vast majority of layouts are conceived as time-limited procedures. They either stop when the end of a sequence of steps is reached (for example, OpenOrd: Shawn et al., 2011) or when a measure of layout quality reaches a certain threshold (such as an energy minimum in a spring model, as in Kamada and Kawai (1989)). In some cases a simple modification allows these algorithms to run continuously, for example by removing the condition on quality improvement included in the Kamada and Kawai (1989) algorithm.

Given the availability of such algorithms, the challenge for software engineering is in the computational and memory resources required for the representation of an animated network, especially when the network scales in size. Gephi (see Bastian et al., 2009) is a software program that includes continuously running algorithms. Its software architecture can accommodate such algorithms in a large-scale dynamic transaction network environment.[11] Two continuously running layouts are available: Fruchterman and Reingold (1991) and ForceAtlas2 (Jacomy et al., 2012). The Fruchterman-Reingold and ForceAtlas2 are force-directed layout algorithms. Force-directed layout algorithms consider nodes as magnets repulsing each other and the edges as springs attracting the nodes they connect, by direct analogy with magnetic and mechanical forces. According to this algorithm, the sum of the force vectors determines the direction a node moves into. The step width, which is a constant, determines how far a node moves in a single step. When the energy of the system is minimized, the nodes stop moving and the system reaches its equilibrium state. Given that the step width is a constant, this algorithm has a drawback in that the system may never reach its equilibrium.

In the following sections, we demonstrate how dynamic visualizations of large volumes of transactional data are now possible, and we report on the lessons learned and the remaining obstacles in the exploratory analysis of complex dynamic networks.

3. Data visualization for the interbank money market

3.1. Monetary policy and the interbank money market

The Eurosystem's monetary policy operations have a strong impact on the interbank money market. By setting reserve requirements, central banks steer the overall tightness of liquidity and the main refinancing rate, i.e. the rates that

[11] A beta version of Gephi was launched in 2008. The dynamic capabilities were introduced in 2010, followed by further improvements in robustness in version 0.8.1 beta (2012) (`http://tinyurl.com/Gephi-viz`).

banks pay to each other on their loans.[12] These rates, in turn, drive the rates that banks charge for lending to the real economy, i.e. households and businesses. By the execution of its policy, the Eurosystem aims to fulfill its mandate of safeguarding price stability.[13]

The central bank can change several parameters in its monetary policy operating procedure to steer market rates, most importantly the target rate and the marginal lending and deposit facility.[14] The Eurosystem's target rate traditionally acts as a benchmark for the rates at which commercial banks lend to each other on the interbank market. The unsecured interbank market is an OTC market and hence the rates agreed on bilaterally by the banks active on this market are not publicly known.[15] The market averages of unsecured overnight lending rates of the major banks, however, are visible through the Euro OverNight Index Average (EONIA). EONIA is computed as the value weighted average interest rate of all overnight unsecured loans reported by the contributing euro area panel banks.[16] By monitoring borrowing and lending between banks on the interbank market (i.e. at which interest rate and in which volumes), the central bank can gain a

[12] In the interbank money market, banks with a liquidity deficit find banks with a surplus to fund their daily operations. By trading in the interbank market they can fulfill their reserve requirements, which is the average amount of liquidity a bank is required to have in its account with the central bank during a predefined period (the reserve maintenance period). The interbank market also provides insurance against inter-temporal liquidity shocks (Bhattacharya and Gale, 1987). Such liquidity shocks could for instance be due to the uncertainty in the timing of depositors' consumption (Allen and Gale, 2000).

[13] See ECB (2010a) for an overview of the Eurosystem's monerary policy in the first 10 years of the euro.

[14] See ECB (2011) for more details.

[15] An unsecured loan is a loan for which a bank does not receive collateral to secure the loan in case the borrowing counterparty defaults. In other words, the bank trusts that the counterparty will repay the loan principal plus the agreed interest when the loan is due. This in contrast to the secured loan market where the counterparty will provide collateral to secure the loan. Overnight means from today until the next business day.

[16] In December 2013 the panel of banks contributing to EONIA consisted of 34 banks.

better understanding of how its monetary policy affects the financial market.[17]

3.2. Available data on interbank money market

Detailed transaction data about the unsecured money market are not readily available for the majority of markets, including for the euro area. By applying an algorithm, however, researchers have identified the size and interest rates of the loans that banks settle in large value payment systems data.[18]

In his seminal contribution Furfine (1999) provides information on the micro structure of the market and looks at relationship between bank size and participation in the funds market. He discovers that even the largest banks are frequently net sellers of funds. He also provides first insights into the network of trading patterns. Abbassi et al. (2013) study the liquidity and relationships in the euro area overnight money market, using a data set constructed with a variation on the algorithm developed by Arciero et al. (2013). They find that for the European sovereign debt crisis, borrowers from countries with most severe sovereign problems were less likely to obtain interbank funding. Millard and Polenghi (2004) study the relationship between the overnight interbank unsecured loan market and the British large value payment system, CHAPS. They find that 22% of all

[17]The importance of the interbank market is widely noted in the literature. Cocco et al. (2009) show that relationships between banks are important for the ability to access interbank market liquidity. Due to the bilateral nature of this market, banks are able to establish and maintain such relationships. Apart from access to liquidity, relationships do matter for both smaller and larger banks for receiving better terms both when borrowing and when lending. The model presented by Carlin et al. (2007) shows that with repeated interaction, cooperation among banks is an equilibrium outcome that involves refraining from predation and that allows those with a larger imbalance in their liquidity position to borrow at more favourable prices than they would otherwise. Acharya and Merrouche (2013) study the liquidity demand of large (settlement) banks in the UK and its effect on the Sterling Money Markets before and during the sub-prime crisis of 2007-08. They find that liquidity holdings of the large settlement banks experienced on average a 30% increase in the period immediately following 9 August, 2007, the day when money markets froze, igniting the crisis.

[18]The algorithm is straightforward: two transactions, first, from bank A to bank B on day t and, second, a slightly higher value transaction from bank B back to bank A on day $t+n$. The difference between the two transactions has to be a plausible interest rate for the duration n. By setting the right boundary conditions of the loan value and the plausibility of the interest rates, unsecured loans can be identified.

payments in CHAPS are associated with money market deals. The majority of the activity relates to 4 (out of the 13) banks in the system at the time of the analysis. Over the course of the day the number of transactions related to loans is divided relatively evenly, but the corresponding value increases.[19]

Building on Furfine (1999), who was the first to develop such an algorithm for the American market, we apply the algorithm of Arciero et al. (2013), resulting in a data set of all the loans exchanged by commercial banks active on to the Dutch part of the euro area money market. This algorithm has been thoroughly validated, identifies loans up to one year and is suitable for the entire euro area.

A clear aspect missing in standard analysis of interbank market activity is the ability to quickly visualize developments including recent developments. A visualization is particularly useful when the payment system has many participants. The European system counts 1,100 account holders, mostly banks (CPSS, 2012). Graphs like in Heijmans et al. (2010) and Heijmans and Heuver (2011) help to monitor the market as a whole (the whole country) or at the level of individual banks.

The above studies help to understand the dynamics of the interbank market up to a certain point, given that the analysis is done at a relatively high level of aggregation. As a consequence, patterns and trends might be concealed at the level of individual banks.

Creating a dynamic representation of the interbank market at the most disaggregated level, individual loan transactions between any two banks, as we will present in Section 4, enables the ‘picture-examining-eye’ to identify transient, fine-grained aspects of the activity on the market. The viewer can get a sense of what regular market activity looks like by observing the patterns of connections between banks. Compared with a static visualization, a dynamic animation provides additional visual clues. A video provides a sense of the speed of the changes unfolding in the market. If a regime switch occurs, the complete transi-

[19] Other studies are Guggenheim et al. (2010) and Akram and Christophersen (2010) who study the Swiss (up to three months) and Norwegian (overnight) unsecured money market, respectively.

tion is visible as it happens, which is lost if only the start and end time slices are shown.

The knowledge gained from observing normal market conditions can be used to identify abnormal connectivity patterns involving individual banks or groups of banks. This exploratory phase of the analysis could lead to a confirmatory phase of analysis, where standard econometric procedures are used to test the significance and estimate the magnitude of the phenomenon that the visualization brought into focus. Conversely, the visualization could also be used to show effects that first have been shown to be statistically (and economically) significant. In the next section, we illustrate the feasibility and usefulness of dynamic network visualization using interbank network data.

4. Case study: Some examples of visualizations

This section provides four examples of animations that are useful in understanding and monitoring payments systems and interbank markets.

4.1. Daily payments

Video 1 shows the first example of dynamic network visualization that focuses on patterns in payments between banks. We selected all payment transactions in the Dutch part of the European TARGET2 interbank payment system for five consecutive business days in June 2013. TARGET2 is the largest euro interbank payment system of the Economic and Monetary Union. It is owned and operated by the Eurosystem.[20] In 2012 the system had 999 direct participants.[21] The daily average number of transactions settled was over 354 thousand with an average daily value of almost EUR 2,500 billion, resulting in an average transaction value of EUR 7.1 million. However, two thirds of transactions had a value below EUR 50,000 and 11% of the transactions had a value of more than EUR 1 million. In other words, the size distribution is very skewed. Banks and their

[20] The system usually operates between 7:00 and 18:00.

[21] Source ECB website.

payment transactions are represented by yellow nodes and cyan edges, respectively. Visualising several sequential days shows the waxing and waning of the daily networks. The ForceAtlas2 layout runs continuously. As a result, the positions of the nodes are recalculated on an ongoing basis ((Jacomy et al., 2012)). Bank labels and the direction of the edges have been omitted to ensure confidentiality. For in-house analysis, displaying bank names, transaction volumes and the direction of edges enables users to answer questions about the behaviour of individual banks: Who are the counterparties at any particular moment of the day? What is the direction of these transactions and what is their value?
The bottom of Video 1 shows a timeline with a sliding selection marker. The selection marker has been set to a three-hour window. This selection can be played automatically across the timeline, resulting in a constantly moving selection of transactions visible within this three-hour window. Both the window and speed of the selection can be changed. The timeline also shows a line graph of the number of banks active throughout the day.

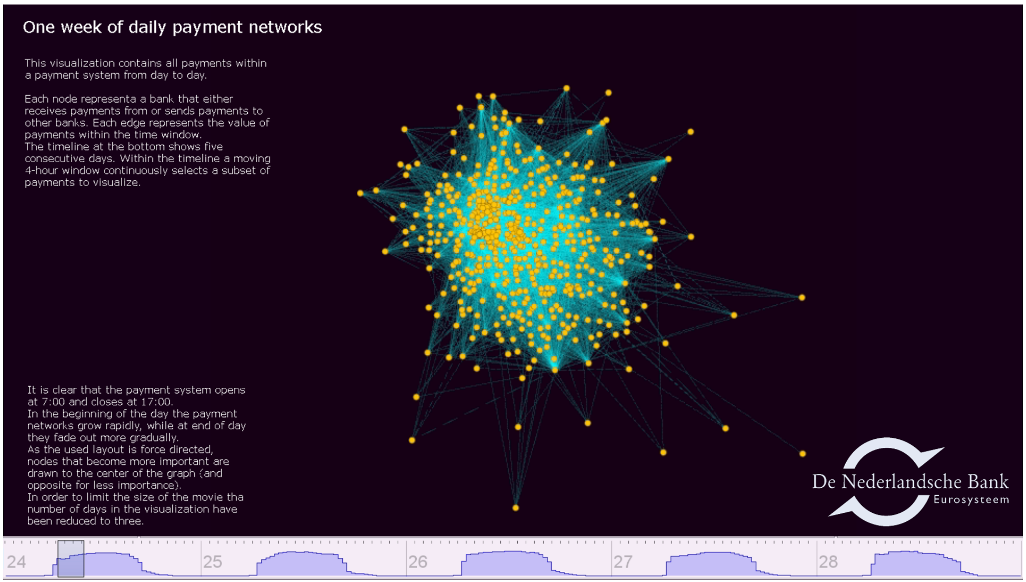

Video 1: All payment transactions within one week.

To start the video, please click on the figure or use the following link `http://www.dnb.nl/binaries/01_one_week.mov`.

Each day at 7:00, the system opens with an explosion of activity as many banks pre-schedule their payments. These scheduled payments are settled immediately if the paying bank's account balance is sufficient.[22] At the end of the business day, the network decreases significantly in size and gradually fades out after 16:30. After 17:00, banks are not allowed to send in client payments (i.e. payments executed on behalf of a bank's clients). Payment instructions sent in between 17:00 and 18:00 relate to payments made by banks on their own behalf, which are often related to money market transactions.

We observe that banks are heterogeneous in their payment behaviour. Some banks are very active sending and receiving relatively large payment amounts throughout the whole day. These banks remain located in the centre of the graph. Other banks are less active, remaining on the outside of the graph. Also, there are banks whose activity changes in the course of the day. Finally, we can observe that banks active during the last hour of the business day are generally also active throughout the entire business day.

The animation facilitates monitoring of the interbank payment market at three different levels. First, we can identify any changes in the market's macrostructure. Is the network stable (in terms of the number of active banks and their connectivity) or does it seem to be splitting up into subgroups? This could prompt further investigation. Second, at a micro level, we can follow an individual bank and replay the timeline for the day, the week or any other time interval deemed relevant, focusing on the transactions made by this bank and identifying its neighbors in the network. Moreover, we can easily switch back and forth between the visualization and the underlying data. Third, we can use the animation for comparative purposes, by evaluating e.g. how animations of the same payment system differs across days, or how animations of different types of payments differ (e.g.for example, client payments versus payments on a bank's own behalf) or how national markets differ at a particular moment in time (for example, the Netherlands compared with other Eurosystem countries).

[22] According to Pröpper et al. (2013) who also study the network topology of the Dutch payment system in graphical, albeit static terms.

Our second case study in Video 2 visualizes an ‘average day’ using the same week’s worth of data as used in our first case study. This visualization helps to investigate structural peaks in the course of the day, identify behavioural patterns of participants and analyze interdependencies between participants over a longer time period. The nodes in this graph have been both sized and colored according to time-changing characteristics. The size of the node is proportional to the size of payments within the time window. The color of the nodes reflects the actual payment balance: red for banks that have sent more liquidity than they received, and green for banks that have received more value than they sent to others.

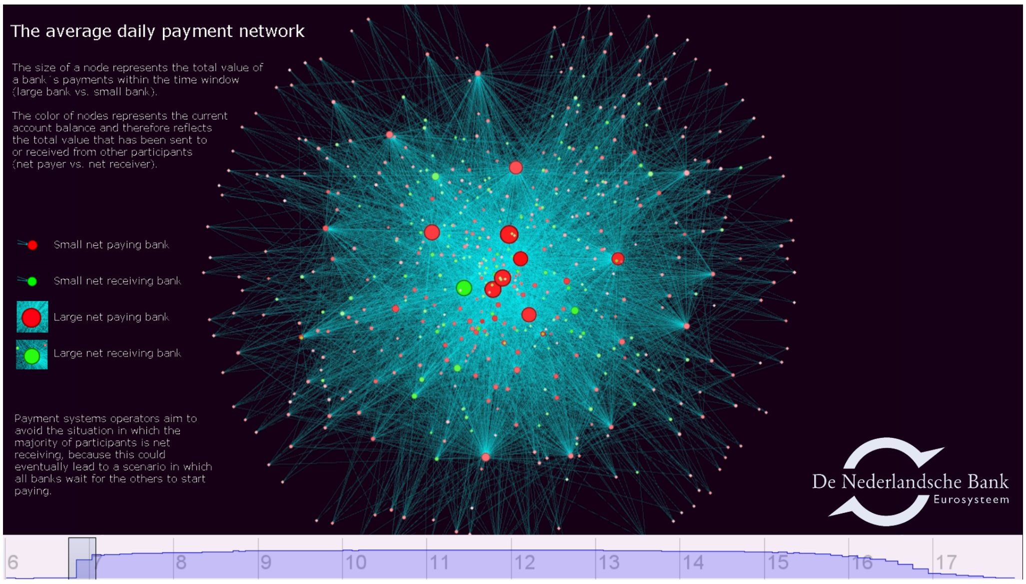

Video 2: Payment transactions within one week converted into one average day.

To start the video, please click on the figure or use the following link `http://www.dnb.nl/binaries/02_one_week_into_one_day.mov`.

Again, the ForceAtlas2 layout is applied continuously to the network shown in Video 2. Contrary to Video 1, however, we have applied the ‘Dissuade Hubs’ parameter.[23] This parameter decreases the attraction of nodes (banks) with large

[23] This parameter was introduced by the developers of ForceAtlas2 in the context of the study of networks formed by webpages linked by URL references (see Jacomy et al., 2012). In this context, an authority is a webpage receiving many links from other pages, while a hub is a webpage with many outbound links. It is interesting to note that this domain-specific feature has proven relevant in the financial domain.

out-degrees (those banks that send many payments to other banks). As a result, such banks are pushed to the periphery of the network (where banks making few transactions are also positioned). Banks with large in-degrees (banks receiving many payments from other banks) are not affected by this parameter; they remain at the centre of the network. Hence, a clear distinction is achieved between small and large players in the market, and between major payment-making banks and payment-receiving banks in the market, at a glance and at any point of time. This distinction is especially useful to spot banks that tend to hoard liquidity, i.e. banks that obtain funding by creating a time lag between payments received and payments made. Central banks pay attention to this type of behaviour as it could lead to a situation where participants in the system wait for other participants to initiate payments, effectively creating a gridlock.[24]

4.2. Dutch money market and central bank liquidity

Video 3 consists of the network of Dutch overnight money market transactions. In order to provide a fuller picture of the system's liquidity as a whole, all liquidity transactions to and from the central bank have also been included.[25] Money market trades and central bank transactions are presented by pink and green edges, respectively.[26] To highlight the relative roles of the central bank and the market, the central bank's position has been frozen on the left of the screen (large green node), while the centre of gravity for the unsecured interbank market has been frozen on the right. Bank nodes are coloured according to their activities, as follows:

- pink nodes: banks that solely participate in the money market, without using central bank facilities.

[24] See Bech (2006) for a game-theoretical model on the timing of payments in a payment system and Abbink et al. (2010) for an experimental game based on this model.

[25] These liquidity transactions can either involve the central bank providing liquidity to banks (if banks need liquidity, in the form of monetary loans or overnight liquidity) or banks sending liquidity to the central bank (if they have a liquidity surplus, in the form of overnight deposits or fixed term deposits).

[26] These are directed transactions (i.e. money is lent by one node to another node). For confidentiality reasons they are not shown in the video.

- green nodes: banks that solely use central bank facilities, without participating in the money market.
- white nodes: banks that neither participate in the money market nor use central bank facilities.

The size and type of payments determine whether a node is either attracted to the money market on the right, or to the central bank on the left, or whether it is located in the periphery in the absence of activity.

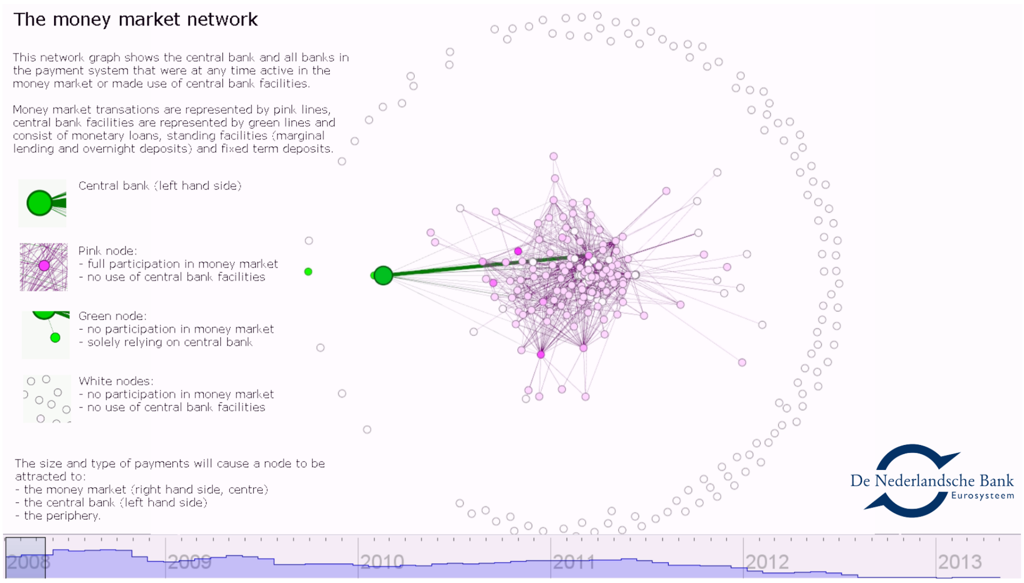

Video 3: Dutch overnight money market and central bank liquidity.

To start the video, please click on the figure or use the following link `http://www.dnb.nl/binaries/03_money_market_liquidity.mov`.

Video 3 spans the period 2008-2013, allowing us to analyse changes in the network in the recent crisis periods. For reference purposes, important events are included in the timeline at the bottom.[27]

[27] The following crisis events are marked yellow on the timeline: the Lehman failure (September 2008), sovereign debt problems of Greece (May 2010), Ireland (October 2010), Italy (August 2011) and Spain (May 2012). Other relevant changes that may have influenced the money markets are the three-year long term refinancing operations (LTROs, 22 December 2011 and 1 March 2012, respectively, marked green) and the reduction in the Eurosystem overnight deposit rate from 25 basis points to zero on 12 July 2012, marked cyan.

The video starts by presenting a vibrant well-connected money market network in 2008, as evidenced by the multitude of pink-colored edges. At the same time most banks also had a small connection with their central bank. The video shows that the money market network became more sparse following the collapse of Lehman. This decline in connectivity became increasingly apparent in 2009, but rebounded slightly in the following two years. The market dropped significantly at the start of 2012, with central bank connections becoming stronger. This decline strongly manifested itself after July 2012. At that point, the Eurosystem overnight deposit rate was reduced to zero, resulting in a final state with a very small money market network, very strong connectivity between banks and the central bank, and many banks no longer participating.

Arguably, the information on the activities of banks can also be presented in traditional charts or tables. However, the use of a dynamic network animation provides several additional layers of information. First, we can observe the trading partners of each bank over time. Lending and borrowing (directionality) can be represented by arrow-pointed edges (not shown for confidentiality reasons). Second, the thickness of the edges can reflect the amount traded or the interest rate, possibly against a benchmark. In this case, we show the edges relative to the EONIA reference rate. As a result we capture the structure of the interbank financial system over time based on the aggregation of individual transactions. Third, we can differentiate between different classes of actors according to relevant attributes and observe how the roles of these classes differ in the network. This video focuses on the activity of commercial banks in relation to the central bank. Freezing the position of the central bank on the left of the screen and allowing the commercial banks freedom to move, has enabled us to highlight the commercial banks that are tightly connected to the central bank and relatively loosely connected to other banks, i.e. banks that are placed outside the herd, located closer to the central bank. Other banks are strongly connected to the central bank through liquidity transactions (as shown by the thick green edges reflecting financial flows involving the central bank) and have dense connections with other banks, which places them within the herd of banks. These relations evolve over time: this animation highlights that the market shrank considerably in 2011 and accelerated in 2012. The banks that remained active in the market continued to be strongly connected to the central bank, effectively showing a market dominated by the links with the central bank. While the video offers the complete stream of events, snapshots can be taken for use in traditional presentations and for detailed analysis of the network's composition at a given time.

4.3. Maturity and cross-border flows in the money market

The last video, Video 4, provides an example of dynamic visualization that combines both geographic and loan level information. It features a daily network of money market transactions. The money market's cross-border transactions have been aggregated to their respective country nodes. These country nodes have been fixed to a large circle that depicts the geographical location relative to the Netherlands. The remaining nodes represent the banks that reside in the Dutch part of TARGET2; they are colored green.

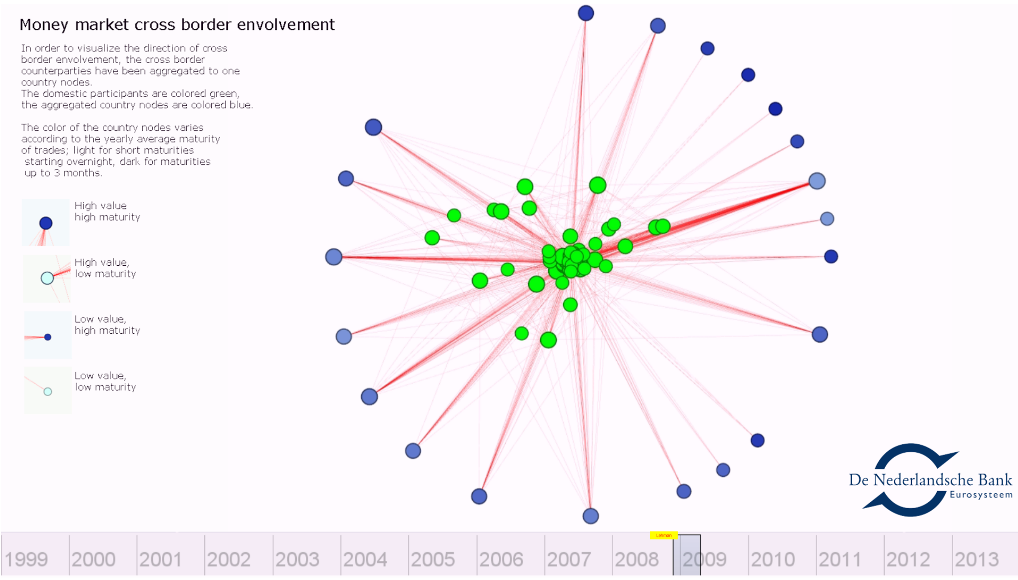

Video 4: Cross-border money market including evolving maturity.

To start the video, please click on the figure or use the following link `http://www.dnb.nl/binaries/04_money_market_cross_border_maturity.mov`.

As visual attributes we have used the color and size of nodes. Size is used to show the overall value of a node's loans. Color is used for each country node to visualize the average maturity (i.e. due date of a loan) of all money market contracts. Short-term average maturities (i.e. overnight) are colored light blue, while longer maturity loans (up to three months) are colored dark blue. Averages have been calculated based on twelve-month moving windows. Again, bank and country labels are not shown, but this feature can be enabled for detailed analysis and monitoring.

Here, the ForceAtlas2 layout is used. Given that the position of banks is recalculated on an ongoing basis, their location is determined by the actual volume of loans to and from each other as well as by the location of the cross border country nodes. For a bank with relatively strong relationships with participants in one country, the location of the node will tend to be towards that country's node.

Banks concentrate their cross border transactions on particular countries. Therefore, not all banks are located in the center of the circle. We also note that the geography of preferred trading partners evolves over time. Spanning 1999 to 2013, the dynamic network visualization reveals multiple shifts in the position of banks towards different country nodes. The visualization highlights two groups of banks. One group has strong ties with countries on the Atlantic shores, while the other has stronger ties with Germany. These two groups were unsettled at the onset of the financial crisis in 2008 when a marked shift in the direction of Germany took place. This type of information is relevant for central banks as it reveals the level of trust placed in different countries by Dutch money market participants. We also note that several banks have strong cross-border activities but make relatively few transactions with other banks in TARGET2-NL. Visually, these banks can be seen flocking to the periphery of the circle. In risk terms, such banks might be exposed to country-specific risk. This is another valuable piece of information for authorities because it is helpful in anticipating possible difficulties at these banks if banks from other euro area countries went into default. The ForceAtlas2 algorithm positions small banks that are heavily (and often solely) related to a certain counterparty very close to that counterparty.

Accordingly, small banks 'follow' their larger counterparties. As a consequence, one major bank increasing its business with a country can result in a group of bank nodes shifting from one country towards this particular country, dragging in its wake the many small banks otherwise attached to it. This provides an insight as to how cross-border transactions can cause ripple effects if there is a shock affecting a major bank.

5. Discussion and conclusions

This paper shows how big data can be visualized in a dynamic way using case studies based on interbank payment transactions. First, we show how the interbank payment system could be monitored by looking at typical payment patterns over several days or, by collapsing the sample, for an average day.
Second, we zoom in on an important subset of interbank payments: unsecured interbank money market transactions, using a long time series from 1999 to 2013.

This paper aims to demonstrate 1) the feasibility of dynamic visualization of large data sets and 2) the potential gains from dynamic visualizations in terms of both analysis and presentation. We show that a regular laptop computer can be used to represent large data sets without aggregating these data. In addition, it is possible to create visualisations of large dynamic networks in a relatively routine manner. By virtue of recent developments in software engineering, the need for programming expertise to create network visualisations has been reduced considerably. While scripting remains useful to transform tabular data into a network structure, visualizations can be created with point and click software. As a result the visual approach to network analysis can become part of the regular analysts' toolkit.

The advantages of visualizations of big data as provided in this paper are twofold. First, all individual transactions between banks (or in general interactions between nodes) can be visualized. This means no initial loss of information as a result of aggregation at the start of the analysis. In payment systems analyzes daily aggregates are often used whereas in bank supervision monthly or even quarterly data are generally used. The obvious disadvantage of transaction data is that the visualization may become too noisy to identify anything meaningful. We countered this disadvantage by using different time windows to smooth excessive volatility. Second, our approach enables multiple dimensions to be included in one visualization. As regards our payment system monitoring case study, we show all banks and their flows in one overall view that changes continuously over time, and we show the number of active banks in the chosen time interval.

As regards the interbank money market case study, we also provide information on average interest rates paid to a certain country. It is also possible to analyze the directed network, however, we do not show this in the paper for confidentiality reasons.

In other words, visualization allows us to examine the 'unexpected' given that raw data can be presented at the start of an analysis showing multiple dimensions of the data. This would not be possible using two-dimensional static graphs or econometrics.

When would this tool be useful? According to John Tukey, we can distinguish between exploratory and confirmatory modes of analysis. We are agnostic as to the relative value of these two modes and believe that both are necessary to obtain sound results. In exploratory mode, there is no precisely defined causal model, and at this stage we want to get a description of key parameters of the data set. The exploratory function partly overlaps with the function of monitoring. In both forms of analysis, we are looking for things that are out of the ordinary, facing the challenge of managing the unexpected. At this stage, leaving ample room for discovery is essential. Freezing one or more parameters at the start might make us miss important new insights that might not be adequately captured by our ex-ante model. Visualizing the data set using different spatial layouts that can be replayed at varying speeds enables us to gain an understanding of the broad structure of the, and how subsets of the data set fit into the overall picture. Trends, regime transitions, dispersion of attribute values, and outliers can all be made clearly visible. As regards network visualizations, the network properties of the data set are made visible. Insights gained at this stage encourage the development of hypotheses that can be tested in the confirmatory mode of the analysis. Confirmatory analysis relies on the traditional analytical tools of the scientific domain under consideration. For financial data, such tools may include time series analysis and a variety of econometric frameworks. At this stage too, visualizations prove useful. They assist in presenting results in an intelligible manner to non-specialists in the domain, thereby increasing the audience's engagement with the subject.

References

Abbassi, P., Gabrieli, S., and Georg, C. (2013). A Network View on Money Market Freezes. *Deutsche Bundesbank and Banque de France, mimeo*.

Abbink, K., Bosman, R., Heijmans, R., and van Winden, F. (2010). Disruptions in Large Value Payment Systems: An Experimental Approach. *De Nederlandsche Bank Working Paper*, 263.

Acharya, V. V. and Merrouche, O. (2013). Precautionary Hoarding of Liquidity and Inter-Bank Markets : Evidence from the Sub-prime Crisis. *Review of Finance*, 17(1):107–160.

Akram, F. A. and Christophersen, C. (2010). Interbank Overnight Rates - Gains from Systemic Importance. *Norges Bank Working Paper*, 11.

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

Arciero, L., Heijmans, R., Heuver, R., Massarenti, M., Picillo, C., and Vacirca, F. (2013). How to Measure the Unsecured Money Market? The Eurosystem's Implementation and Validation using TARGET2 Data. *De Nederlandsche Bank Working Paper 369*.

Bastian, M., S., H., and Jacomy, M. (2009). Gephi: An Open Source Software for Exploring and Manipulating Networks. In *American Journal of Sociology*, Proceedings of the Third International AAAI Conference on Weblogs and Social Media, pages 361–362.

Baur, M. and Schank, T. (2008). *Dynamic graph drawing in Visone*. Interner Bericht. Fakultät für Informatik, Universität Karlsruhe.

Bech, M. L., Bergstrom, C. T., Garratt, R. J., and Rosvall, M. (2011). Mapping Change in the Federal Funds Market. *Federal Reserve Bank of New York Working Paper*, (507).

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

Bertin, J. (2010). *Semiology of Graphics: Diagrams, Networks, Maps*. ESRI press.

Bhattacharya, S. and Gale, D. (1987). Preference Shocks, Liquidity and Central Bank Policy. In Singleton, W. and Barnett, J., editors, *New Approaches to Monetary Theory, Cambridge University Press*, pages 69–88. New York.

BIS (2012). Quarterly review.

Blondel, V., Esch, M., Chan, C., Clerot, F., Deville, P., Huens, E., Morlot, F., Smoreda, Z., and Ziemlicki, C. (2012). Data for Development: the D4D Challenge on Mobile Phone Data. *arXiv preprint arXiv:1210.0137v2*.

Borgani, S.and Murante, G., Springel, V., Diaferio, A., Dolag, K., Moscardini, L., Tormen, G., Tornatore, L., and Tozzi, P. (2004). X-ray properties of galaxy clusters and groups from a cosmological hydrodynamical simulation. *Monthly Notices of the Royal Astronomical Society*, 348(3):1078–1096.

Borgatti, S., Mehra, A., Brass, D., and Labianca, G. (2009). Network Analysis in the Social Sciences. *Science*, 323(892).

Brandes, U. and Wagner, D. (1997). A Bayesian Paradigma for Dynamic Layout. In Di Battista, G., editor, *Proceedings of the 5th Symposium on Graph Drawing*, Lecture Notes in Computer Science.

Carlin, B., Lobo, M., and Viswanathan, S. (2007). Episodic Liquidity Crisis: The Effect of Predatory and Cooperative Trading. *Journal of Finance*, 62:2235–2274.

Cocco, J., Gomes, F., and Martins, N. (2009). Lending Relationships in the Interbank Market. *Journal of Financial Intermediation*, 18:24–48.

CPSS (2012). Principles for Financial Market Infrastructures: Disclosure Framework and Assessment Methodology. *Bank for International Settlements*.

ECB (2010a). Monetary Policy Transmission in the Euro Area, a Decade after the Introduction of the Euro. *ECB Monthly Bulletin*, May:83–94.

ECB (2010b). The ECB's Response to the Financial Crisis. *ECB Monthly Bulletin*, May:59–74.

ECB (2011). The implementation of Monetary Policy Operations in the Euro Area. *General Documentation on Eurosystem Monetary Policy Instruments and Procedures*.

Erten, C., Harding, P., Kobourov, S., Wampler, K., and Yee, G. (2003). GraphAEL: Graph Animations with Evolving Layouts. In Liotta, G., editor, *Graph Drawing: 11th International Symposium, GD 2003 Perugia, Italy, September 21-24, 2003 Revised Papers*, pages 98–110. Springer Berlin Heidelberg.

Fruchterman, T. and Reingold, E. (1991). Graph Drawing by Force Directed Placement. *Software: Practice and Experience*, 21:1129–1164.

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

Gort, A. and Gort, J. (2007). Survey of Network Visualization Tools. *Defense R and D Canada Contract Report*, CR2007-280.

Guggenheim, B., Kraenzlin, S., and Schumacher, S. (2010). Solving the Puzzle: Evidence of the Unsecured Swiss Franc Money Market. *Swiss National Bank Working Paper*, 5.

Hausmann, R. and Hidalgo, C. (2013). *The Atlas of Economic Complexity: Mapping Paths to Prosperity*. Cambridge, Mass.: The MIT Press.

Heijmans, R. and Heuver, R. (2011). Is this Bank Ill? The Diagnosis of Doctor TARGET2. *De Nederlandsche Bank Working Paper*, 316.

Heijmans, R., Heuver, R., and Walraven, D. (2010). Monitoring the Unsecured Interbank Money Market Using TARGET2 Data. *De Nederlandsche Bank Working Paper*, 276.

Jacomy, M., Heymann, S., Venturini, T., and Bastian, M. (2012). ForceAtlas2, A Continuous Graph Layout Algorithm for Handy Network Vizualization.

Kamada, T. and Kawai, S. (1989). An Algorithm for Drawing General Undirected Graphs. *Information Processing letters*, 31:7–15.

Kashcha, A. (2013). "VivaGraphJs" Github Repository.

Knecht, A. (2005). Simulation from Knecht Friendship Data – Four Wave Study on Adolescent's Behavior and Relations using an Actor-Driven Statistical Model in SIENA.

Millard, S. and Polenghi, M. (2004). The Relationship between the Overnight Interbank Unsecured Loan Market and the CHAPS Sterling System. *Bank of England Quarterly bulletin*, Spring.

Moody, J. and Mucha, P. (2013). Portrait of Political Party Polarization. *Network Science*, 1:119 – 121.

Nicolini, C. and Dallachiesa, M. (2013). GraphInsight: An Interactive Visualization System for Graph Data Exploration.

Pröpper, M., van Lelyveld, I., and Heijmans, R. (2013). Network Dynamics of TOP Payments. *Journal of Financial Market Infrastructures*, 1(3):3–29.

Schwabish, J. (2014). An economist's guide to visualizing data. *Journal of Economic Perspective*, 1:209–234.

Shawn, M., Brown, W., Klavans, R., and Boyack, K. (2011). OpenOrd: An Open-source Toolbox for Large Graph Layout. In *Visualization and Data Analysis*. Proceedings SPIE7868.

Snijders, T. (2001). *The Statistical Evaluation of Social Network Dynamics*. Boston and London: Basil Blackwell.

Tukey, J. (1980). We Need Both Exploratory and Confirmatory. *The American Statistician*, 34:23–25.

Ugander, J., Karrer, B., Backstrom, L., and Marlow, C. (2011). The Anatomy of the Facebook Social Graph. *ArXiv e-prints*, 1111.4503.

van Riet, A. E. (2010). Euro Fiscal Policy and the Crisis. *ECB occasional paper series*, 109.

Wong, B. (2012). Point of View: Visualizing biological data. *Nature Methods*, 9(1121).

Previous DNB Working Papers in 2014

No. 406 **Raymond Chaudron and Jakob de Haan**, Identifying and dating systemic banking crises using incidence and size of bank failures
No. 407 **Ayako Saiki and Sunghyun Henry Kim**, Business cycle synchronization and vertical trade integration: A case study of the Eurozone and East Asia
No. 408 **Emmanuel de Veirman and Andrew Levin**, Cyclical changes in firm volatility
No. 409 **Carlos Arango, Yassine Bouhdaoui, David Bounie, Martina Eschelbach and Lola Hernández**, Cash management and payment choices: A simulation model with international comparisons
No. 410 **Dennis Veltrop and Jakob de Haan**, I just cannot get you out of my head: Regulatory capture of financial sector supervisors
No. 411 **Agnieszka Markiewicz and Andreas Pick**, Adaptive learning and survey data
No. 412 **Michael Ehrmann and David-Jan Jansen**, It hurts (stock prices) when your team is about to lose a soccer match
No. 413 **Richhild Moessner, Jakob de Haan and David-Jan Jansen**, The effect of the zero lower bound, forward guidance and unconventional monetary policy on interest rate sensitivity to economic news in Sweden
No. 414 **Dirk Broeders, An Chen and Birgit Koos**, Utility-equivalence of pension security mechanisms
No. 415 **Irma Hindrayanto, Siem Jan Koopman and Jasper de Winter**, Nowcasting and forecasting economic growth in the euro area using principal components
No. 416 **Richhild Moessner**, Effects of ECB balance sheet policy announcements on inflation expectations
No. 417 **Irma Hindrayanto, Jan Jacobs and Denise Osborn**, On trend-cycle-seasonal interactions