

Towards a Network Description of Interbank Payment Flows

Marc Pröpper^{*}, Iman van Lelyveld^{**} and Ronald Heijmans^{*}

PRELIMINARY, PLEASE DO NOT QUOTE WITHOUT CONSULTATION

March 2008

Abstract

We present the application of network theory in analysing payment networks with specific attention to systemic stability. As an example we use Dutch payment data which represent an average system in Europe. We first show that payments are well behaved through time and that there are no structural imbalances in the system. Since the Dutch market is highly concentrated, we zoom in on behaviour of the largest participants. Subsequently, we analyse the properties of prominent network measures against time. As payment links are temporary by nature the role of time is crucial. Finally, we show that the network is susceptible to directed attacks and that the recent ‘sub prime’ turmoil in credit markets has not materially affected the network structure.

JEL codes: G1, E5

Key words: network, topology, interbank, payment, systemic risk, financial stability

^{*} De Nederlandsche Bank, Supervision Policy Division, P.O. Box 98, 1000 AB Amsterdam, The Netherlands. Email: m.h.propper@dnb.nl and i.p.van.lelyveld@dnb.nl.

[♦] Nijmegen School of Management, Radboud University, P.O. Box 9108, 6500 HK, Nijmegen, The Netherlands.

^{*} De Nederlandsche Bank, Payments Policy Division. Email: r.h.heijmans@dnb.nl.

The contents of this working paper reflect the opinions of the individual authors and do not necessarily reflect the views of De Nederlandsche Bank. The authors would like to thank Hans Brits, Maarten Gelderman, Jan Groenewold, Walter Stortelder and Kimmo Soramäki for their useful comments, and Gwan Tjia, Richard Heuver and Odette Min Verlee for their statistical support.

INTRODUCTION

From the late 1990s the study of the topological structure of random networks has gained momentum. Empirical observations from large and rapidly evolving networks like the World Wide Web (Albert (1999)) and the Internet (Faloutsos, et al. (1999)) brought to light a surprising compactness ('small world phenomenon') and relatively many highly connected networks elements (Dorogovtsev and Mendes (2003)). These findings have shifted attention away from classical, static networks (Erdős and Rényi (1959)) towards growing networks.¹ An important property of the latter is their robustness against random failures. At the same time, however, they are vulnerable against directed attacks (Albert, et al. (1999)).

The ideas of network theory can be applied to the field of economics, for example to study the risk of widespread propagation of financial distress (*systemic risk*). There is a vast literature analysing the interactions of various financial markets such as equity and bond markets.² A small but growing literature examines the riskiness of interbank markets, where banks exchange relatively short term and largely unsecured funds.³ These papers, however, do not focus on the network topology of the markets. Inaoka, et al. (2004) and Soramäki, et al. (2007) have started to describe large value financial payment systems (in Japan and the US, respectively) from a network perspective. Another example is Bech and Garrat (2006) which analyses the effects of a wide-scale disruption on the functioning of the interbank payment system. Our paper builds on this literature, and adds to it by illustrating (1) the influence of the chosen time frame on the properties of the payment network and (2) the central role of highly connected banks in the functioning of the payment network. We use data from the Dutch large value payment system, an active, medium sized network.

Importantly, in contrast to for instance analyses of interbank exposures (e.g. Boss, et al. (2004) or van Lelyveld and Liedorp (2006)), payments networks are by definition short lived: as soon as the payment is settled, the visible, recorded connection between banks disappears. As we analyse and discuss in more detail in (Pröpper, et al. (2007)), this affects our understanding of what constitutes a network. Here we show that the *time frame* used to compute the network measures materially affects the outcomes. A proper understanding of the evolution of the network is important from a risk management perspective. Ultimately, the purpose is to get an understanding of the level of stability or, alternatively, vulnerability of the system to random or directed failures and to systemic risk.

¹ The former, equilibrium random networks, have Poisson degree distributions (the degree of a node is the number of its links). The latter, non-equilibrium random networks, may under the right conditions result in fat-tailed, scale-free degree distributions close to a power law. This is the case when they are governed by (a linear kind of) preferential attachment which means that new network elements are more likely to attach themselves to elements that are already highly connected (Barabási and Albert (1999)).

² See Pericoli and Sbracia (2003) for an overview.

³ See Allen and Gale (2000) for a theoretical characterisation of these markets and van Lelyveld and Liedorp (2006) and the references therein for an empirical analysis. The general finding is that interbank markets are from a systemic stability point of view relatively safe.

Network theory equips us with promising tools. The questions we will answer are tackled by first studying the (time development of) network structure of the payment system in terms of commonly used network properties like the size of the payment network, the connectivity between banks, distances in the network, the distributions of connections between banks and network correlations. This for example allows us to take a first, indirect peek at the risks that the system faces, by assessing the importance to the network of the most highly connected banks. In addition, we present results of the impact on the payment network of the sub prime crisis (in 2007) which caused worldwide turmoil on financial markets.

The set-up of our paper is straightforward. We start with a short description of the institutional detail of the Dutch large value payment system (TOP), the technical details of the data set and an international comparison of aggregate key figures. Next, we discuss the intraday behaviour of the system. Then we investigate whether there are structural imbalances (between individual –or groups of– participants) in the system. After this examination of the basic properties of our data we examine the build up of the network over time. First we analyse the development over time of commonly used network measures in the literature. Second, we analyse the vulnerability, or, alternatively formulated, systemic stability, of the system. We report the impact on the network structure and on the key system figures of removing the ten most highly connected participants in the data set. In addition, we analyse whether the recent ‘sub prime’ crisis in credit markets has affected the network properties of the payment system. We end with the conclusions.

THE DUTCH PAYMENT SYSTEM

Since 1999 the Dutch large value payment system (TOP) is part of the European system for euro-denominated payments, TARGET.⁴ TOP is restricted to a limited set of participants, mainly banks. Connections to participants in other TARGET countries takes place through TARGET.⁵

For the system to function properly it is essential that participants have sufficient funds so that payments can be made without delay. Intraday credit provided by DNB (secured by collateral) facilitates a smooth functioning of the payment system and prevents gridlock.⁶ In the Netherlands, commercial banks permanently hold (pledged) collateral at the central bank, generally at a relatively stable level during the year.⁷

⁴ Trans-European Automated Real-time Gross settlement Express Transfer system.

⁵ By the end of 2007, TARGET will be replaced by TARGET2, which is a single shared platform. All local systems, including TOP, will then migrate to TARGET2 in three phases. Technically, the latter is a centralised system, but legally it is a decentralised system in which each country designates its own component system.

⁶ See Ledrut (2006) for a discussion of the optimal provision of intraday liquidity.

⁷ In addition to maintaining a collateral pool, it is possible to place collateral through repo transactions. When the credit balance becomes insufficient, collateral is brought in and the balance is raised, usually in the morning. At the end of the day, the transaction is reversed.

Regular opening hours are from 7h until 18h and during these hours the payment system processes all transaction types. In addition, there is an evening settlement period from 19:30h to 22h, which is used for settling ancillary system batches and not for standard domestic transactions and cross border (TARGET) payments. Incidentally, the latter two types of transactions make up for more than 80% of the value transferred (Oord and Lin (2005)).

We analyse a data set consisting of one year of transaction data from the Dutch large value payment system, running from June 2005 to May 2006 (257 business days).⁸ Transactions carried out during evening settlement are excluded. No standard domestic, domestic correspondent bank and cross border transactions through TARGET are carried out during evening settlement. We use the settlement time rather than the moment a transaction is entered into the system in our analysis. Participants with more than one account are consolidated. Payments between two accounts of a single participant are therefore not included in the analysis. Also, due to the limitations of the dataset, cross border transactions are not analysed on a participant, but on a country level. In short, we analyse a network of participants, not of accounts, and some participants are countries rather than banks.

Table 1 shows daily averages on numbers of participants, transaction volumes, values transferred and (average) transaction values for the Top (NL), TARGET (EU), CHAPS (UK) and Fedwire (US) payment systems. The TOP figures are presented with and without evening settlement.⁹ They include incoming and outgoing cross border payments through TARGET. The numbers show that TARGET and Fedwire are both large payment systems of the same order of magnitude. The Dutch domestic system is clearly smaller; only the average transaction value is relatively high.

Table 1: Key figures on daily payment characteristics Top (NL), TARGET (EU), CHAPS (UK) and Fedwire (US).

	<i>Top</i> (without / with evening settlement)	<i>TARGET</i>	<i>CHAPS</i>	<i>Fedwire</i>
<i>Measurement Period</i>	6/2005-5/2006	2005	2005	2005
<i>Participants</i>	155 ¹⁰	10,197	Not available	6,819
<i>Of which direct participants</i>	100	1,126	15	not available
<i>Transactions (x1000)</i>	15.1 / 18.1	312	116	519
<i>Value (in billion €)</i>	151 / 173	1,987	297	1,634
<i>Transaction value (in million €)</i>	9.9 / 9.5 ¹¹	6.4	2.6	3.1

Source: Top (DNB), Target (ECB bluebook), CHAPS and Fedwire BIS (2007).

⁸ Processing of data has been done in Java by extending graph data structures from Goodrich and Tamassia (2006).

⁹ Numbers including evening settlement are relevant for section 'Net value transferred'; numbers without evening settlement are relevant for intraday payment behaviour discussed in section 'Intraday dynamics'.

¹⁰ The number of active participants in the measurement period amounts to 129 (or 131 with evening settlement)

¹¹ All payments within a second from and to the same participant are aggregated. When every payment is treated separately, the average value decreases to approximately €7.5 million. In case the incoming cross border payments are excluded the average payment value is €6.5 million.

TRADITIONAL CHARACTERISTICS

Now we first turn to an examination of ‘traditional’ characteristics of the payment system. The network measures discussed later will not render these traditional measures obsolete: they are complementary measures. We will first look into the intraday dynamics. Relatively stable dynamics are important when considering different time frames for network measures later in this study. Secondly we will analyse the net payment position of participants over different horizons. This should tell us whether there are persistent net payers or receivers and thereby indicate if directions in payment links between participants matter when studying network properties. We will also examine the role of the three biggest banks in the Netherlands. Market concentration is high and therefore the behaviour of the large banks is an important determinant of the overall market structure.

Intraday dynamics

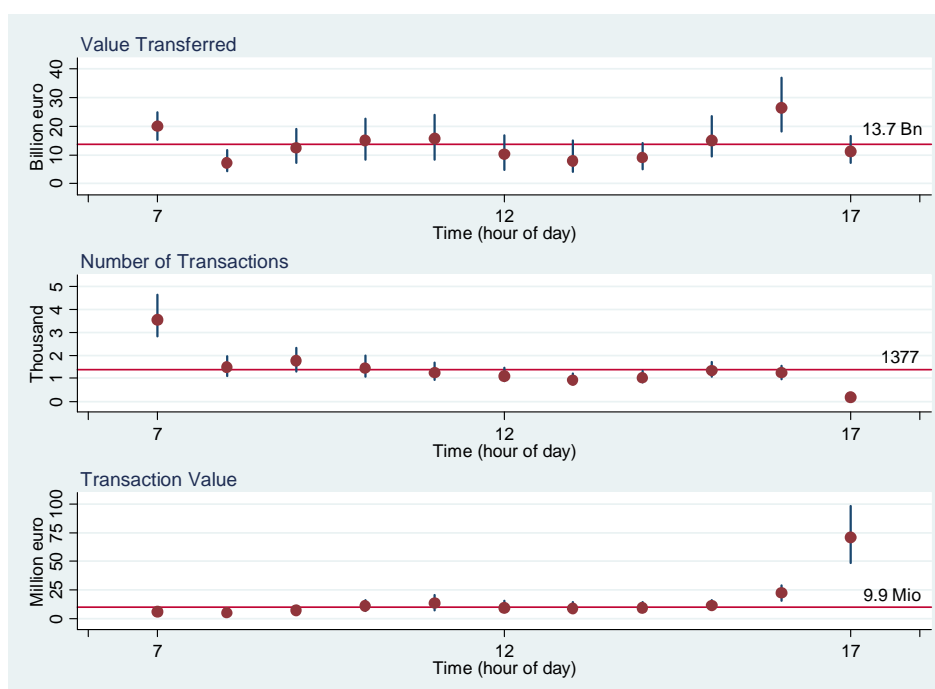
Figure 1 displays for each business hour the average value transferred, the number of transactions processed and the transaction value. The first pane shows Dutch banks are willing to pay early in the day: *value transferred* peaks (€20.0 billion) during the first business hour.¹² This is due to payments entered the day before. Numbers strongly increase to an all day high (€26.6 billion) between 16h and 17h. Some of this activity is the result of banks that need to level their balances as a result of the intraday credit used or to fulfil their cash reserve requirements. These payments are usually few in number but relatively large in value. In the last business hour, from 17h to 18h, only transactions between banks are processed (no retail orders), but most banks usually do not wait until the last hour to close their balance of the day, and finish before 17h. Therefore, value transferred slumps in the last hour. Closer inspection of the distributions around the means revealed they are fairly symmetrical. This is also the case for the distributions of the number of transactions and the transaction values.

The second pane, showing the *number of transactions*, illustrates that on average about 3,500 transactions are processed in the first hour, almost one every second. The 5% and 95% percentile values range from 2,800 to 4,600. The rest of the day transactions take place in smaller numbers (between 900 and 1,750 transactions) along a somewhat downward sloping trend against time that abruptly drops to very low numbers in the last hour (5% percentile \approx 200). The average number of transactions per hour and per day respectively amount to 1,377 and 15,148. The distribution range is significant and comparable through time (between 45% to 65% of the average value for the specific hour).

Finally, the third pane shows the *transaction value*. The average transaction value during the day amounts to €9.9 million but in the last two hours of business it increases strongly to respectively €22

¹² Against an average value transferred per hour (day) of €13.7(151) billion.

Figure 1: Average value transferred, number of transactions processed and transaction value during regular opening hours.



Note: the averages for a particular hour (over all of the 257 business days) are denoted by a dot; the bars run from the 5% low to the 95% high percentiles of the observations. The horizontal lines depict daily averages over regular opening hours.

million and €71 million. The last hour, however, hardly contributes to the overall average due to the small number of transactions.

Overall we observed payment characteristics do change during the day. Payment behaviour in the beginning and the end of the day differ most from the rest of the day. In the morning many payments seem to be driven by ‘good customer’ behaviour to pay early while payments during the last hour very likely reflect liquidity decisions mainly.

NET VALUE TRANSFER AND CIRCULAR FLOWS

Introduction

The payment system is a closed system: for the total system there is no net displacement in value. Structural net payment flows could still exist between (groups of) participants within the time frame of our data set. Conceptually, at least two independent types of structural flows can be distinguished. Individual participants can either have a non-zero net transfer to the rest of the system or circular flows may exist between participants. In the latter case, the net value transfer of each individual participant in the circular flow may still be zero.¹³ If participants become dependent on such structural flows then

¹³ Note that in a typical payments network net value transfer from any bank to the remaining banks is limited as they have to maintain an average reserve value on their account during the reserve period. In case of an

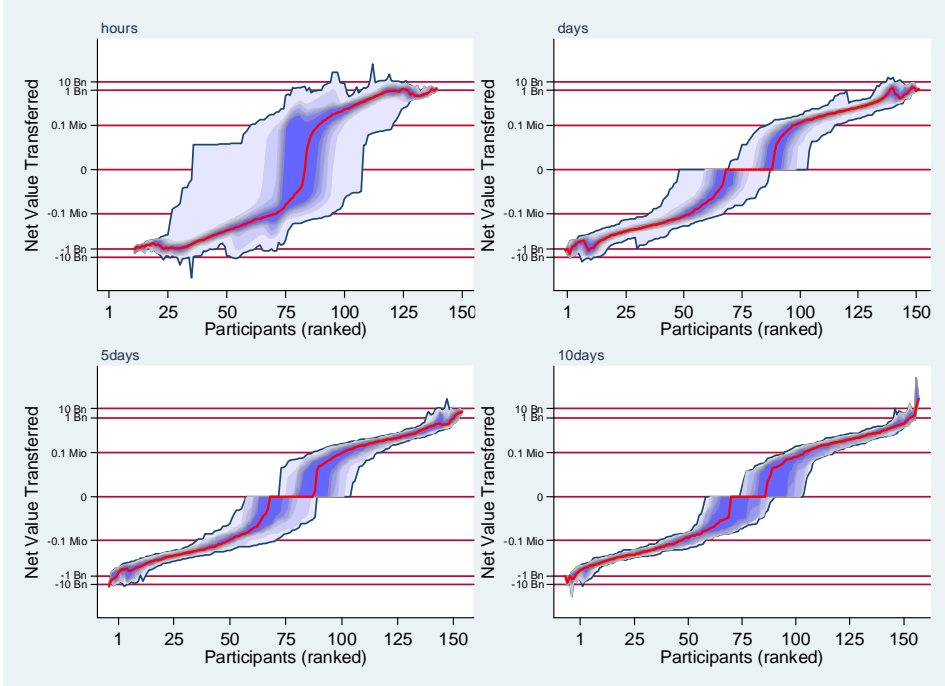
this is likely to increase sensitivity to disruptions in the system. We will show that individual participants generally manage their balance actively towards the rest of the system. However, when we focus on the three largest banks in the system and on cross-border TARGET payments at the country level, the payment system will prove to contain significant structural circular flows.

Net value transfer by individual participants

To capture possible structural behaviour between individual participants and the rest of the system, i.e. net value transfer, the participants are first ranked by their net value transferred during “time snapshots” of fixed length (respectively for one hour, one day, five days and ten days). For time snapshots of one day this results in 257 rankings of participants. We align these rankings at the middle participant of each ranking. Then, net value distributions *across* the rankings are determined.

It is important to emphasize this procedure includes all participant institutions, but none of the TARGET country accounts (for it would be hard to see balance management for these accounts). The latter issue will resurface when we analyze structural circular flows between countries below. The current analysis should not be labelled ‘domestic’, though, since accurate determination of the net value transfers of individual institutions requires inclusion of all payments, including cross-border transactions to and from other TARGET countries. In addition, some of the largest foreign banks have their own participant numbers in the system.

Figure 2: Net value rankings of individual participants; distributions for snapshots of 1 hour, 1 day, 5 and 10 days.



insufficient average at the end of the period the participant will obtain funds (a loan) from another participant or pledge more collateral.

Note: Displayed in each pane are median (in red), maximum and minimum values and in different colour shades the 5th, 10th, 15th, 20th, 25th, 75th, 80th, 85th, 90th and 95th percentiles.

Figure 2 shows the resulting distributions on the logarithmic y-axis against participants on the x-axis.^{14,15} The figure points at fairly consistent net payment behaviour over time, with many smaller participants and a few larger participants. For the one hour period, the distribution width is clearly wider, decreasing for longer periods. Remarkably, the maximum and minimum net value amounts remain more or less constant for all four time periods (maximum between €15 billion and €18 billion; minimum between €-13 billion and €-22 billion).¹⁶ For the entire period of 257 days (not shown) the maximum and minimum values are even smaller, respectively €13 billion and €-8 billion. Total gross value on the other hand increases from on average €9.8 billion for the one hour period to on average €173 billion for the one day period (see also Table 1) and €44,575 billion over all 257 days. The relative net displacement of value (net value/gross value) therefore strongly decreases as the time period increases. The net flow in a year's time is negligible in comparison with the gross value transferred and we will thus proceed abstracting from the direction of payments.

To illustrate further, measured over all 257 days the ten largest participants (gross value) account for 75% of total gross value transferred. Besides large banks, these include the TARGET countries Germany, UK, Belgium and France. The maximum ratio of net value over gross value amounts to only 0.36% in this group of ten individual participants. So, while there is a substantial flow of funds to and from these participants (banks and countries) the net value displacement can almost be ignored. Overall, this leads to the conclusion that the system does not contain a group of significant structural net receivers or payers (cf. Furfine (1999)). It suggests all significant participants actively manage their balance over time.

Cross-border TARGET payments (net circular flows)

Table 2 illustrates total net cross-border payments from the Dutch TOP payment system through TARGET to the other TARGET countries, from June 2005 to May 2006. In aggregate, there was a negligible net inflow of €12.6 billion (only 0.04% of the total gross cross-border payment flows of €29,831 billion). On the country level, however, large net outflows take place to DE and in particular the UK (together €513 billion). These are largely balanced by large net inflows of similar, but opposite size from IT, ES, BE, EU/ECB and FR (together €472 billion). The Netherlands prove to be part of an international structural circular net flow between TARGET countries. For similar considerations can also be made for the other TARGET countries. The UK for example (also shown in the table) has a

¹⁴ The order and total number of active participants on the x-axis will vary over snapshots.

¹⁵ Unless explicitly stated otherwise, logarithms in this paper will have base ten.

¹⁶ The maximum (minimum) net value need not be at the far right (left) end of the x-axis. The maximum (minimum) net values across all snapshots may be linked to a snapshot that is not the largest in terms of the number of participants.

large net outflow of €1,758 billion to DE, FR, EU/ECB, SE, ES and DK and a large net inflow of €1,776 billion from BE, IT, NL and LUX. The system as a whole is a closed system.

Rosati and Secola (2005) analyze *gross* cross-border large value payment flows through TARGET and speak of a ‘tiered market structure for liquidity’. The big countries (DE, FR and GB) are at the center of a core integrated market which also comprises of IT, BE and NL. The other, smaller countries form a sort of ‘periphery’. They also mention the basic idea that significant and stable payment patterns between countries could entail dependencies resulting in possible channels of contagion of liquidity tensions. Here we add to this analysis evidence for the existence of large circular *net* flows between TARGET countries (order of size of GDP).¹⁷

Table 2: Cross-border net value transfers from The Netherlands and UK through TARGET to other TARGET countries (6/2005-5/2006).

<i>Net value transfer From NL</i>		<i>Net value transfer From UK</i>	
<i>To country</i>	<i>Value (in bn €)</i>	<i>To country</i>	<i>Value (in bn €)</i>
GB	414	DE	827
DE	99	FR	619
IE	37	EU	104
SE	4	SE	92
PT	2	ES	62
DK	14	DK	53
GR	16	GR	29
AT	16	AT	16
FI	22	FI	5
LU	24	PT	12
FR	76	IE	19
EU	83	LU	69
BE	97	NL	414
ES	103	IT	427
IT	113	BE	866
Total	13	Total	0

Source: ECB

The three largest banks in the system (gross circular flows)

The three largest banks in the system play a dominant role and are involved in 52% of all transactions and 63% of total value transferred (see also section ‘Vulnerability of the network structure’). These figures emphasize their important role in balance management over time. Active balance management by each of the three banks separately is suggested, besides in actual numbers of net value transfer, by the presence of relatively high, negative autocorrelations (-0.44, -0.44 and -0.29 for banks A, B and C respectively) when the net value series of each bank is shifted over one day. Positive (negative) net

¹⁷ These net flows should be accompanied by opposite streams in other markets, such as the capital markets and markets for goods and services.

value positions are likely to be followed by negative (positive) net value positions the next day. These autocorrelations do not extend beyond one day.

Gross payment flows between the three largest banks in the system over the data period of one year reveal relatively large structural net flows exist between each pair of two banks. In particular, a large net flow of around €90 billion exist between one bank (A) and the other two banks (B and C). One circular flow in the system can therefore be identified to run from bank A to banks B and C, from there to the rest of the system and from the rest of the system back to bank A. As the payment system is a closed system, no circular flow runs solely between the three banks. These observations add to the idea that the large banks form the central core of the network and are natural counterparts for the other banks in the system.

Further work could be directed at identifying the most important circular flows in the system and at investigating the implications they have on the stability of the payment system. This could reveal possible vulnerabilities in the system. One conclusion we can already draw is that that the network is directed of nature. Since the current study principally wants to reveal the time development of network measures, we have, however, chosen to proceed by analysing the system in the next sections from the simplified perspective of an undirected network.

NETWORK MEASURES

Introduction

The previous sections have described the payment system from a traditional perspective in terms of transactions processed and values transferred. This has given insight in the behaviour of individual participants and of the system as a whole. Now the perspective will shift towards describing the payment system in terms of its *network* properties.

A network (or graph in mathematics) is a set of connections (links) between pairs of objects (nodes). Examples of real-life networks are numerous and include social networks, communications networks, transportation networks, biological networks, the World Wide Web, the Internet and financial networks. In a payments network, like the one studied here, the participants form the nodes and transactions establish links between the nodes. Within each time period considered, a link between two nodes is created by the first transaction between them. Subsequent transactions add weight to the link in terms of the number of transactions processed and the additional value transferred over this link.¹⁸ Every pair of nodes can be connected by two opposing links since the individual transactions contain a clear direction from payer to receiver, i.e. the network is a directed network.¹⁹ Finally, a path is an

¹⁸ In this paper link weights are not taken into account in determining network properties (to prevent subjectivity) and all links are thus considered equivalent.

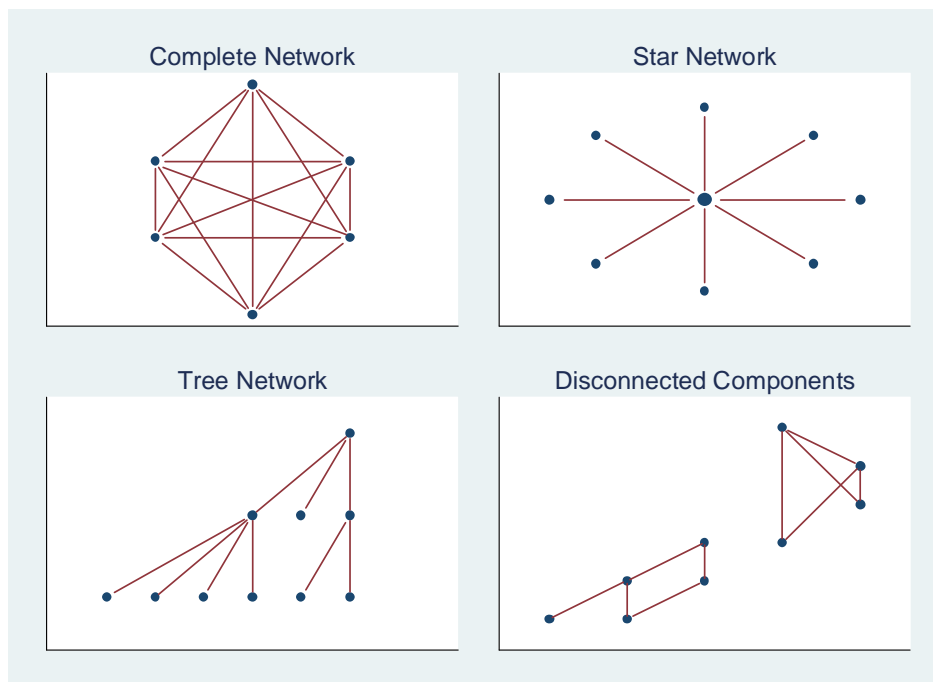
¹⁹ The payment network is also a ‘simple network’. It contains no self-loops (payer = receiver) nor parallel or multiple links from one sender to one other receiver. The links of the network form a set of node pairs (not just a

alternating sequence of nodes and links such that each link is incident to its predecessor and successor nodes. A path can be directed (along directed links) or undirected (along undirected links).

For illustrative purposes graphical representations of some basic (undirected) network types are shown in Figure 3. These include a complete network, a star network, a tree network and a network with two disconnected components. In a complete network all nodes are connected to all other nodes by a link. In a tree network all nodes are connected by exactly one path (no loops or cycles). A star network is a network in which the nodes connect to a central node called the hub. In a network component all nodes are connected by at least one path. A network is connected if it consists of a single component; if a network is not fully connected it consists of two or more components.

For an introduction into the theory of random networks and the treatment of real-life examples and extensive lists of reference material the interested reader is referred to Albert and Barabási (2002),

Figure 3: Basic types of (undirected) networks



Dorogovtsev and Mendes (2003) and Newman (2003). Statistical descriptions of financial networks are scarce in literature, however, which probably relates to the confidentiality of the transaction data. This may be especially true for payment systems. Exceptions are Soramäki, et al. (2007), Lublóy (2006), and Inaoka, et al. (2004).

collection, see Goodrich and Tamassia (2006). For some applications this sense of direction is not essential. In that case, connections between nodes are formed by a single, undirected link (undirected network). This may for instance be the case when the establishment of a contact by a transaction is important, but not the direction of the transaction.

Evolution of network properties

Payment systems are dynamic networks of which the number of nodes and links can vary greatly over time. The actual transfer of money only creates a temporary link. As we have argued in Pröpper, et al. (2007), the choice of timescale for the statistical description of network properties is important and, for the Dutch case, the network properties of the dominant network component are representative for the whole payments network after about a 10 minute time period.²⁰

The network measures we analyse are explained in the Appendix.²¹ They include network size in nodes and links, connectivity, reciprocity, path length, eccentricity, degree, degree correlation, degree distribution, nearest and second nearest neighbours, and clustering. The treatment of these properties aims at giving more insight in the topological structure of the network. Figure 4 displays the development over time of the various network measures *for the dominant network component*. For each of the measures, the x-axis represents the different time periods investigated (i.e. 1, 3, 5, 10, 30 minutes, 1, 3, 5 hours, 1, 3, 5, 15, 257 days) in minutes. The use of a logarithmic scale enables coverage of all time periods. It requires careful interpretation of the figures, though, to get a good understanding of the high rates of development for short time periods (\leq one hour) and lower rates for periods beyond one day. The discussion of the results is largely restricted to the relatively variable outcomes for the one hour time period and the relatively stable outcomes for the time period of one day. The former generally represent intraday network properties well, the latter network properties for periods from and beyond one day.

The figure shows major developments take place mostly in the first hour of network formation. From one hour to one day the network grows more gradually. The *size* of the network measured 88 ± 6 nodes on an hourly basis and 129 ± 5 nodes on a daily basis (top left). During the whole period of 257 days (only) 183 nodes have been active in total. These numbers characterize a small-size network, also with respect to other investigated banking networks.²² In Inaoka, et al. (2004) the total number of banks amounted to 354, in Boss, et al. (2004) to about 900 and in Bech and Soramäki (2005) more than 5000 banks made up the system.

On an hourly basis 326 ± 76 directed *links* were found between the nodes (top middle). On a daily basis there are almost four times as many links: 1182 ± 61 . The number of links increases at a higher rate than the number of nodes, but the number of *possible* links increases with the number of nodes *squared*. Over the whole period a mere 12% of the possible number of links between nodes ($183 * 182 = 33306$) actually did become a link, to a total of 4079 links for one or more transactions.

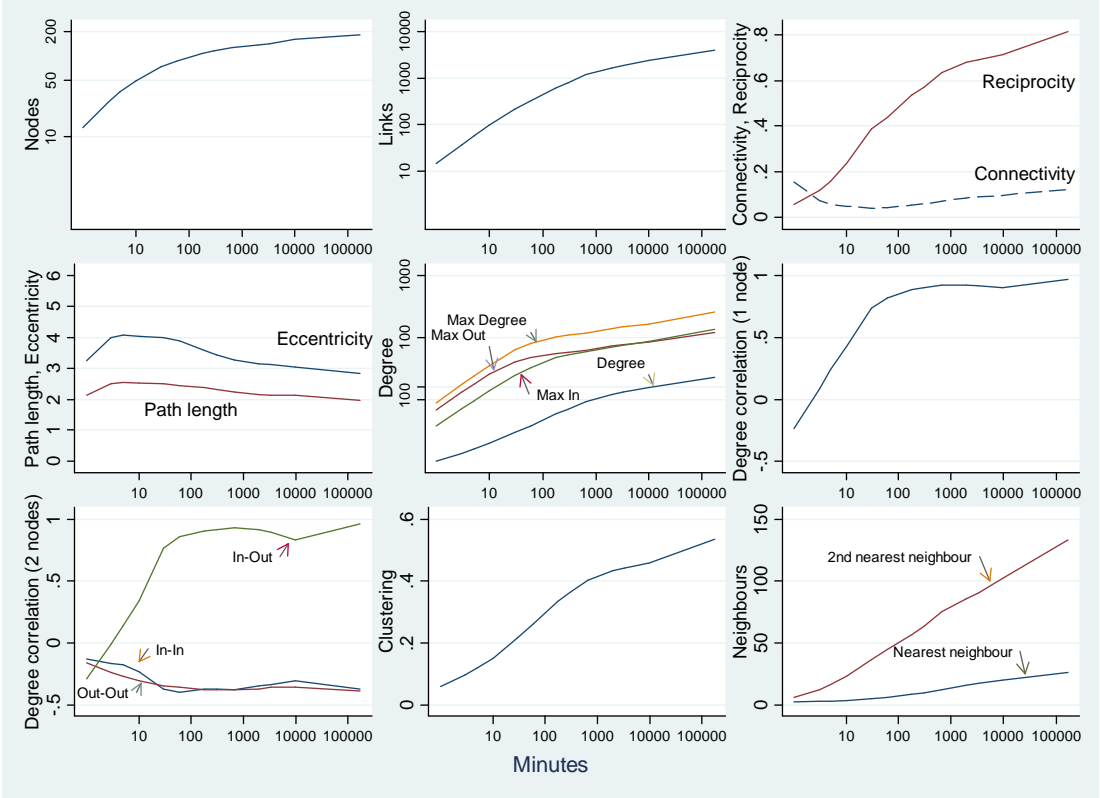
²⁰ After one (ten) minutes 36% (68%) of the data samples already consists of a single network component (Pröpper, et al. (2007)). In the case of the ten minute time frame, in the overwhelming majority of the 32% remaining cases there are only one or two minimally sized other components of two, three nodes.

²¹ And in for example Dorogovtsev and Mendes (2003) and Soramäki, et al. (2007).

²² Payments through TARGET to and from different banks in the same EU country are all recorded under the same country code and therefore belong to the same node. This leads to a downward bias.

The fraction of actual to possible links, *connectivity*, gives better insight in the relative growth of nodes and links based on the proper, quadratic relation between them (top right). The values show that the network remains very sparse over all time periods. Connectivity rapidly declines from 0.16 ± 0.12 after one minute to a minimum of 0.04 ± 0.01 after approximately 30 to 60 minutes, to increase thereafter at a lower pace to 0.07 ± 0.00 after one day and 0.12 after 257 days. The explosion of nodes in the first hour suppresses connectivity, because the growth of links does not keep up with the growth of nodes but after one hour the situation reverses. At all times, however, the network keeps its low connectivity and remains far from connected. Even after 257 days 88% of all theoretically possible links have not been used for a single transaction. *Reciprocity*, the fraction of links with a link in the opposite direction, displays a rapid increase in the first hour to 0.44 ± 0.12 and increases at a lower rate to 0.63 ± 0.02 after one day. It means that a link in one direction implies a high probability of a link in the opposite direction. Payments often take place in two directions. This, however, gives no information on the intensity of activity in both directions.

Figure 4: Development of network properties over time (in minutes): nodes, links, connectivity, reciprocity, path length, eccentricity, degree (also max degree, max in-degree, min degree, min in-degree) and degree correlation (for one node and for two nodes), clustering, nearest and second nearest neighbours



Note: panes displaying nodes, links and degrees use a logarithmic y-axis.

The average *path length* between two randomly selected nodes forms another way to measure the size of the network (middle left). This distance peaks at 2.5 ± 0.3 nodes at the beginning of network formation (5 min) to decline gradually thereafter to 2.0 nodes after 257 days. The latter is of the order of the logarithm of the size of the network (number of nodes), a feature predicted both for the classical

equilibrium and fat-tailed, scale-free non-equilibrium networks (Dorogovtsev and Mendes (2003)). The numbers indicate that every node on average connects to another node through only one intermediate node. The maximum distance, *eccentricity* (or diameter), amounts to 4.1 ± 0.7 after 5 minutes and declines gradually to 2.8. This means that over time the maximum number of steps participants have to take to reach the other participants decreases. Concentric, spherical connections gain strength in comparison to linear, radial connections. The network gets more structure and on a local level it becomes less 'tree-like'. The results again emphasize the small size of the network (in 'length' this time) and raise the question whether the intermediary node is also random in general, or that a core of central nodes exists through which other nodes connect.

Node *degree*, the number of links connected to a node, forms an essential measure for the description of the direct surroundings of a node (middle right). The degree measure can be split in in-degree and out-degree on the basis of the number of in- and outgoing links. The concept of a degree can easily be extended and generalized to concentric circles of neighbouring nodes with length 1, 2, ..., n ($n <$ network size). A close relationship therefore exists between degree and length of the network (Dorogovtsev and Mendes (2003)). Here the focus will only be on direct neighbours of length 1. From an initial value of 1 the network degree increases to 3.7 ± 0.8 links per node after one hour and at a somewhat slower pace to 9.2 ± 0.4 links per node after one day. It takes nearly the rest of the 256 days to (more than) double to 22.3 links per node.

These outcomes deviate significantly from many theoretical models of growing networks which assume a fixed degree (linear growth). In these models each added node is accompanied by a fixed number of new links (e.g. Barabási and Albert (1999); see also discussion in Dorogovtsev and Mendes (2003)). The payment network clearly exhibits a form of accelerated growth, because the degree increases during network growth. The model of network growth would also differ from theoretical models due to the upper boundary in the number of participants. Growth in nodes inevitably declines over time, since fewer nodes can be added. The theoretical model of the payment network, including accelerated growth in links and a declining growth in nodes due to the limited number of participants, is a subject for further study.

The time development of the *maximum degree* shows node degree covers a large range of values across the network (middle right). The maximum degree increases from a level of about 9 times the average degree after 1 minute, to a level of around 20 for periods between ten minutes and three hours, slowly declining over time to a level of 11 afterwards. Concretely, it means that after one hour the average node may hold 3.7 ± 0.8 links, but the maximally linked node actually holds 79 ± 13 links. *Maximum out-degree* surpasses *maximum in-degree* for periods up to a day. These maxima reflect the presence of one or more highly connected nodes. The presence of large differences in degree values also reflects large differences in the local network structure, hinting at a structure of many low-degree

and some high-degree nodes. The actual degree distribution of nodes across the network, discussed at the end of this section, therefore contains indispensable information about the local structure.

The *degree correlation* (centre) between in-degree and out-degree of individual nodes starts off negatively but becomes very strongly positive after just thirty minutes ([70%-100%]). This means that above (below) average in-degree has a high chance of being accompanied by above (below) average out-degree. Nodes that make payments to many counterparties also receive payments from many counterparties. The results on reciprocity already showed nodes are often counterparties in both directions. Degree correlations between in-degree and out-degree of two connected nodes largely follow the same pattern (bottom left). Degree correlations between in-degree respectively out-degree of two connected nodes prove negative. The results on degree correlations again suggest the existence of a few strongly connected nodes linking to several weakly connected nodes.

The *clustering coefficient* (bottom middle) measures the probability of two neighbours of a node sharing a link among themselves, too. Where distance measures length, clustering measures density of the network structure at a local level. It gives information about the direct surroundings of the nodes. As expected, the development of the clustering coefficient over time confirms that formation of connections across neighbours takes more time to develop than sheer growth of the network. Still, the rate of increase in clustering is relatively high in the first hour and somewhat lower afterwards. The average clustering coefficient increases from 0.26 ± 0.09 after one hour to 0.40 ± 0.02 after one day. After 257 days, average clustering amounts to 0.53. It means that, on average, in half of the cases the neighbours of a randomly chosen node are connected among themselves, too. When comparing this to classical equilibrium networks, the numbers indicate relatively high correlations in the form of clustering exist on a local level. As in Soramäki, et al. (2007) the number of nodes with a clustering coefficient 0, however, is very high: 49% after one hour, 27% after one day. These can mostly be attributed to nodes with only one or just a few links, because the probability of at least a single link across neighbours increases rapidly with the number of links (more precise: the number of combinations of neighbours increases with the number of links *squared*). In any case, local density proves absent for a significant part of the network on a local level.

The last pane (bottom right) displays the number of *nearest neighbours* z_1 and *second nearest neighbours* z_2 of a node. In determining z_1 and z_2 the direction of the links has been ignored,²³ since it is the presence of a connection that matters here (not the direction of it). The number of nearest and second nearest neighbours amounts to 5.7 ± 0.1 , respectively 45 ± 4 after one hour. These numbers increase to 12.5 ± 0.5 , respectively 75 ± 3 after one day. They emphasize that the number of direct contacts z_1 strongly increases over time and that the second line of contacts z_2 through z_1 (logically) is a multiple of the first line (the number of contacts grows strongly between the first and second line).

²³ This distinguishes z_1 from the degree.

Like in the case of the clustering coefficient, the results confirm that local structure takes more time to develop than the size of the network.

We now turn to *degree distributions*. As mentioned earlier, these distributions give indispensable information about the relative “popularity” of participants in the system. Large banks are obvious examples of popular (highly connected) participants, but also the specific clearing institution which settles many, relatively small, customer driven payments. As such payments go to and from most institutions, this clearing institution will be a highly connected node many participants attach to.

The concept of degree can easily be extended from nearest neighbours to second nearest neighbours, to third nearest neighbours, etc, but here the focus will be on nearest neighbours only. In Figure 5 the degree is plotted on the x-axis with the associated probability on the y-axis (both on a logarithmic scale).²⁴ The darker the dot associated with each degree bucket, the larger the size of the firm(s) in that bucket. We measure firm size by total annual transaction value.

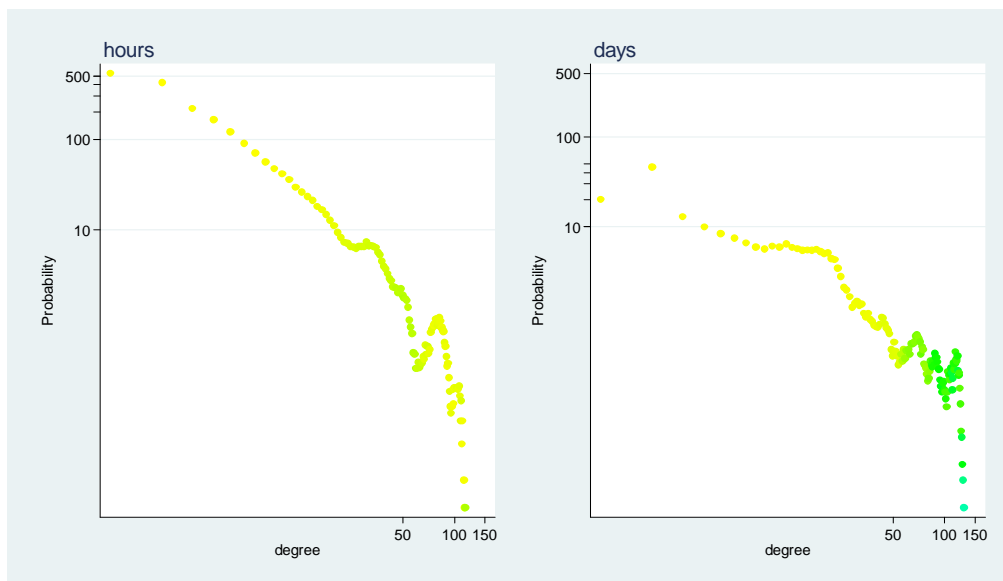
For one hour time snapshots the distribution already steeply declines at very low degree values; for one day time snapshots this rapid decline starts off from a degree value of about twenty. The largest observed size of a one hour (day) network amounted to 116 (130) nodes. The networks are characterized by a high number of nodes with relatively few connections and a small number of (relatively) highly connected nodes. Each of the smaller humps at the high degree end of the x-axis accounts for an individual node or a small group of individual nodes. These humps are basically distributions of individual (groups of) nodes.

Further, in going from one hour to one day time snapshots the frequency of highly connected nodes increases at the expense of weakly connected nodes. During the treatment of the clustering coefficient and degree correlations it was already mentioned that connections across neighbours take more time to develop than growth of the network. The local structure becomes stronger with time. Also, the degree distributions for the payment network cover too few order of degree (≈ 2), to perform any fitting to a power-law distribution (to test whether the distribution is a scale-free, fat-tailed distribution).²⁵

²⁴ Averaging has taken place over all snapshots in the dataset.

²⁵ See §5.6 and footnote 11 of Dorogovtsev and Mendes (2003) for a critical note on the empirical ‘observations’ of power-law distributions. Scale-free networks have a size-dependent cut-off, which sets strong restrictions for such observations over 2 or 3 orders of degree.

Figure 5: Degree distributions for time snapshots of one hour and one day, respectively



The shading of the dots in the graph tells us that for shorter time periods, in this case the one hour snapshots in the left pane, the firms with the highest degrees are not necessarily the largest firms (in terms of value transferred). The left pane clearly shows the importance of participants that handle batches of consumer payments; the value of these payments are not high but they do entail many connections. In comparison, the one day snapshots in the right pane show us that high degrees are associated with large turnover.

VULNERABILITY OF THE NETWORK STRUCTURE

In the introduction we noted that the study of the topological structure of the payment network (or any other network) is not a goal in itself, but a means for understanding the processes that make use of the structure. A particularly interesting topic of research is the vulnerability (or resilience) of the network to random or directed failures. The impact of a failure of a single node may remain confined locally or cause a shockwave that propagates through the system (systemic risk). The purpose here is to show that network theory provides tools for studying this risk.

We analyse the impact on network properties of removal, one by one, of the most highly connected nodes (cf. Albert, et al. (1999)). Risks to the system may surface upon showing the importance of specific nodes to the topological structure. The removal procedure is equivalent to building the network from the raw transaction data, but leaving out all individual transactions that involve the specific 'removed' nodes. It is a static procedure with shortcomings, like for instance the absence of any adaptive behaviour. The topology after removing n specific nodes is always the same, but the path in getting there will differ upon changing the order of removal of the nodes. It means the procedure will not identify a unique dependence of the topology on any of the individual, removed nodes. What

the procedure does bring to light is the dependence of the topology on removal of a limited number of highly connected nodes.

Following the literature we choose one day as the timescale to measure the network properties. Figure 6 shows the impact of the removal procedure on several selected network properties. The x-axes show network properties in the initial situation ('0') and, in going to the right, those properties after removing the most highly connected node ('-1'), the second most highly connected node ('-2'), etc., until ten nodes have been removed ('-10'). The initial situation ('0') for all properties is the same as in Figure 6 (one day time period).

The network becomes smaller and even more sparse as for instance shown by the *degree* values (top left corner) which decrease steadily, except for point '-7', from 9.2 to 4.3. This results from the number of nodes having declined from 129 to 89 (-31%) and the number of links from 1182 to 378 (-68%). Moreover, *connectivity* decreases from 0,072 to 0,049. The network loses more nodes than the 10 deliberately removed nodes, because on average 30 neighbouring nodes with a single link will lose their last connection during this procedure. The seventh node ('-7') is a good example, since the degree actually increases upon removing this node.

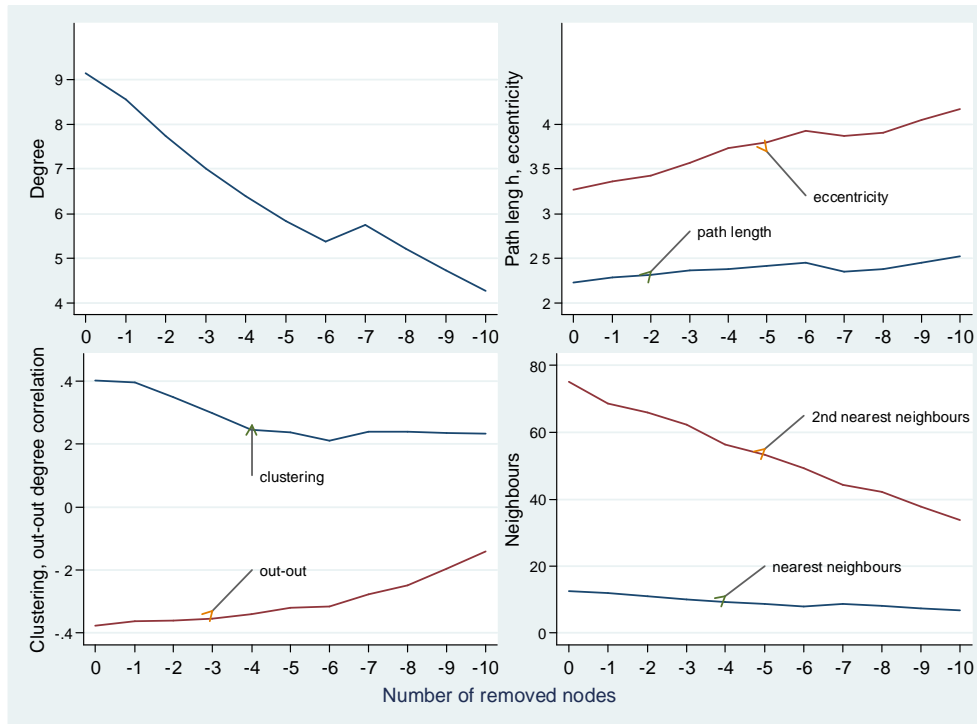
The removal of central, highly connected nodes increases the path lengths between the remaining nodes. In the removal of the seventh node this phenomenon is outweighed by the accompanying loss of the single link nodes and the shortest paths between them and all other nodes. Specifically, *path length* and maximum path length (top right corner), or *eccentricity*, increase from 2.2 to 2.5 and from 3.3 to 4.2, respectively.

The outcomes for clustering and correlations both show the local structure starts to break down (bottom left corner). Clustering, or density of connections on a local scale, decreases from 0.40 to 0.23. The removal of nodes two to four has an unevenly negative impact on clustering in comparison to the other nodes. The out-out degree correlation increases more steadily from -0.38 to -0.14 (= loss of correlation).²⁶ The outcomes for nearest neighbours and second nearest neighbours confirm this breakdown in structure (bottom right corner).

The impact of removing the ten most highly connected nodes on the key aggregate figures of the payment network is severe. Value transferred and number of transactions decline steeply to respectively only 6% and 12% of the initial situation (not shown here). This marks these nodes as essential to the core function of the payment network. This holds especially for the first 4 nodes, since by their removal value transferred and the number of transactions have already declined to 27% and 30%, respectively, of the initial situation.

²⁶ In-in degree correlation increases from -0.38 to -0.10. In-out degree correlation decreases from 0.93 to 0.59.

Figure 6: Impact of node removal on network properties: degree, path length, eccentricity, clustering, out-out degree correlation, nearest and second nearest neighbours (z_1 and z_2)



It should be clear that random removal of ten nodes would not have caused the same impact on the network structure and key aggregate figures. In this sense the results are comparable to those in Albert, et al. (1999) in that the system is vulnerable to a directed failure (here: removal of a highly connected node) due to the importance of the relatively highly connected nodes in the tail of the degree distribution. In addition, the discussed procedure of node removal convincingly shows network theory provides tools for analysing distortions to the network.

THE NETWORK STRUCTURE DURING RECENT TURMOIL IN CREDIT MARKETS

In the previous section the vulnerability of the network has been illustrated principally by a static, hypothetical procedure of node removal. The dependence of network properties on the most highly connected nodes proved to be strong. Other, more realistic events can also affect the proper functioning of the payment system. A prime example is a possible loss in confidence between banks which would reduce the liquidity of funds in the markets. If banks delay or stop making payments to (some) other banks, this will have its effect on the functioning of the payment system if the scale of such change in behaviour or the scale of banks involved is large enough.

In 2007 problems in the US 'sub prime' mortgage market have caused worldwide turmoil in credit markets. Initially, US homeowners with risky sub prime mortgages were confronted with strongly rising housing costs due to rising variable mortgage rates. Their subsequent payment troubles led to

heavy valuation losses on mortgage portfolios. These portfolios had often been securitised and sold to third-party investors. Previously liquid credit markets quickly dried up due to a loss of confidence between counterparties. Credit spreads increased strongly and central banks (ECB, FED, BoE) intervened to supply liquidity. In this section we investigate the impact on the functioning of the high value payment system. To this end we compare in Figure 7 the payment system between June and September 2007 (the part of the turmoil period we investigate) with the same period in 2006. This procedure ensures seasonal effects do not get misinterpreted.

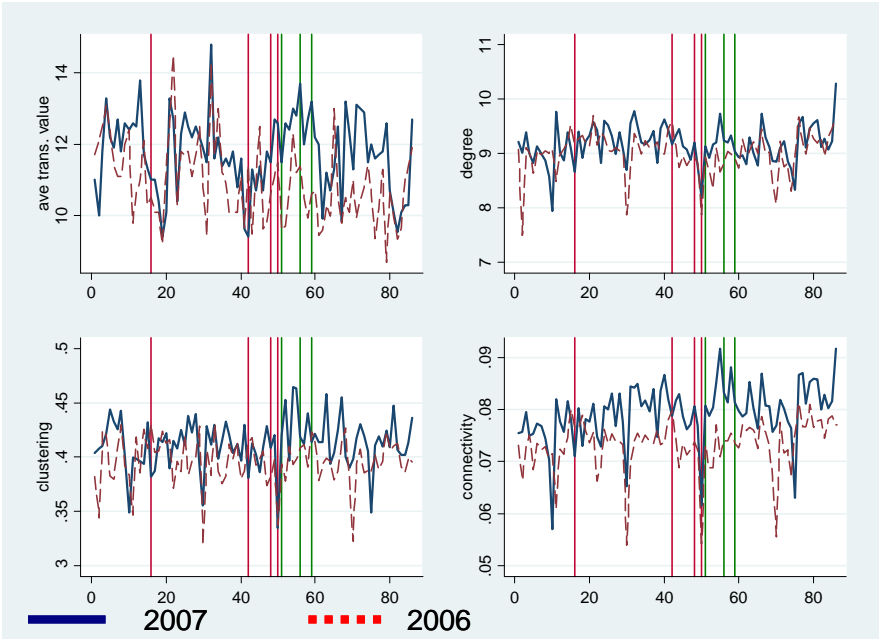
Figure 7 displays the functioning of the payment system by plotting a selection of previously discussed measures: average transaction value, degree, clustering and connectivity. If, as a result of a loss in confidence, participants would be more conservative in choosing their counterparties or in transferring large sums of money, it would show up in these measures. Some important events during the period of turmoil have been indicated by the vertical bars.²⁷ It's difficult to exactly pinpoint when the broad market turmoil started, but the first event ("Bear Sterns") is a good candidate.²⁸ This event is indicated in the figure by the most left vertical line. A closer view of the individual measures points to the importance of comparing the relevant period in 2007 with the previous year. The top left pane of Figure 7 shows the average transaction value of the investigated periods in 2006 and 2007. This figure shows that the average transaction value is higher in 2007 for most days compared to the same period in 2006. However there are no remarkable deviations, either positive or negative, on one of the highlighted events, nor on the other business days. The degree, shown in the top right of the figure, is slightly higher for the 2007 period but follows a similar trend. The clustering on the bottom left shows a similar result as the degree. The bottom right pane shows a higher connectivity during the turbulent period in 2007. In general it can be concluded that the payment activity was higher during the investigated turmoil period in 2007 than in the corresponding period in 2006. In both series we see similar periodic and non-periodic developments causing daily variability in the displayed measures. Unlike other markets where for example credit spreads steeply rose and activity declined, the payment system seems to have functioned properly and not to have been disrupted noticeable outside the bandwidth of regular daily fluctuations. Such disruptions should clearly show up in the discussed measures. These outcomes are representative for the other, non-plotted measures and have been tested thoroughly for their validity.²⁹

²⁷ The vertical bars from left to right indicate market-moving news involving Bear Sterns (22-06, working day (wd) 16), IKB (30-07, wd 42), Goldman Sachs (07-08, wd 48), BNP (09-08, wd 50), FED and ECB liquidity support (09-08, wd 51), FED cuts rates (17-08, wd 56) and Bank of America (22-08, wd 59).

²⁸ Earlier events before June 2007 such as the bankruptcy of a US mortgage bank did not receive the same worldwide attention.

²⁹ For example by filtering out transactions below different threshold values and by filtering out all TARGET transactions to and from other countries. Also, other time frames shorter than 1 day have been tested.

Figure 7: Development of a selection of traditional system measures and network properties over time: transaction value, degree, clustering, and connectivity.



Note: To make the two sets of data comparable we start both series on the same day of the week. Further, we have dropped all days corresponding to Dutch public holidays.

The results of this section demonstrate the usefulness of monitoring both traditional system measures and network properties in order to assess the functioning of the payment system.

CONCLUSIONS

Recently, interest in the topological structure of networks has risen significantly. Application of the ideas of network theory to payment systems is still limited, though. This study adds to literature in showing the measures available to characterize such networks. The application of these measures illustrates the influence of the chosen time frame on the properties of the payment network (discussed in more detail in Pröpper, et al. (2007)) and the central role of highly connected banks in the functioning of the payment network.

In this study we use data from the Dutch large value payment system TOP (part of TARGET) as TOP takes on a middle position in comparison to other (European) systems in terms of value transferred and number of transactions processed. It is a natural candidate to analyse, although naturally each country system will have its unique characteristics.

We first gave a description of the network in terms of ‘traditional’ measures such as turnover and average payment size. An international comparison with other payment systems revealed that the analyzed network is midsized and relatively active. Looking at intraday developments activity proves high right after opening, mainly due to queued orders entered the previous day. With a brief rise in activity right before lunch, the all day high in value transferred takes place between four and five o’clock.

Then we turned our attention to the question whether structural payment flows exist between (groups of) participants for two types of such flows: net value transfer and circular flows. Such dependencies might make the system more vulnerable to failures and would emphasize the need to include information on the direction of payments in the description of network properties. We examined this issue both for the full sample, as well as for the three largest banks in the sample.

The distributions of net value rankings of individual participants remain remarkably similar for different time periods. In addition, the relative net displacement of value (net value/gross value) of individual participants strongly decreases as the time period increases. These findings led us to conclude that the system does not contain a group of significant structural net receivers or payers. For periods of one day and beyond a dominant group of active participants exists that hardly transfers any net funds relative to gross value transferred. The outcomes suggest all significant participants actively manage their balance over time.

By focusing on the three largest banks in the system we could identify the existence of an important structural, circular flow in the payment system. It means the payment system is a directed network. In addition, a short analysis of total net cross-border payments through TARGET revealed The Netherlands are part of an international structural circular net flow (order of size of GDP). The observations of the presence of large circular *net* flows between TARGET countries add to the idea of

Rosati and Secola (2005) that significant and stable payment patterns between countries could entail dependencies resulting in possible channels of contagion of liquidity tensions. Further work could aim at identifying the most important, circular flows in the (international) system and at investigating the implications of structural flows on the stability of the payment system. Since the current study principally wants to reveal the time development of network measures, we have, however, chosen to proceed by analysing the system from the simplified perspective of an undirected network.

Then we proceeded with a presentation of the time development of several important network measures. As we have argued in Pröpper, et al. (2007), time is crucial in analysing networks with short lived links, like payment networks. The outcomes have made explicit to what extent fast development takes place in the early phase of network formation of about one hour and slower development afterwards. The payment network proves to be small (in nodes and links), compact (in path length and eccentricity) and sparse (in connectivity) for all time periods. Measurements of degree and degree correlations, clustering, and the number of nearest and second nearest neighbours describe the development of network structure at a local level. As expected, development of network structure takes more time than growth of the network in terms of size. The actual degree distribution contains indispensable information on the local structure. It proves that the network is characterized by a high number of nodes with relatively few connections and a low number of highly connected nodes. Of particular importance is the observation that the average degree increases during growth of the network. This contrasts with many theoretical models that assume node degree remains fixed. Further work could therefore be directed at modelling the payment network using a type of accelerated network growth in links.

Finally, we showed that the payment system is vulnerable to a directed failure and that recent market turmoil has not materially affected the network structure. The vulnerability of the network was tested by removing, one by one, the ten most highly connected nodes. Node removal had a strong impact on value transferred, number of transactions and network properties like degree, path length and eccentricity, clustering and degree correlation, and nearest and second nearest neighbours. These outcomes emphasized the central role the most highly connected banks, especially the top 4 of these, play in the payment system: they are essential to the core function of the payment network. We also investigated whether the recent 'subprime' turmoil in credit markets has led to changes in the structure of the payment network by comparing recent time-series of network measures with the same period a year earlier. We concluded that it has not materially affected the network structure of the payment system during the investigated period. Since severe disruptions in the payment system would inevitably show up in the discussed measures, it is useful to monitor for changes in traditional system measures and in network properties. A comparison with developments in previous years may ensure seasonal effects do not get misinterpreted.

The current study intends to show how various measures can be used in analysing payment networks. It is also an exploratory study: two clear directions for further research are analysing, first, the importance of link weights and, second, the role of collateral and available liquidity in absorbing shocks. Ultimately, knowledge of the functioning of the payment network is aimed at gaining a better understanding of the means to preserve its stability.

APPENDIX NETWORK PROPERTIES³⁰

Size

The most basic network properties are the number of nodes *nodes* (n) and *links* (l). The former is often referred to as the *size* of the system. The relative number of links l to the possible number of links determines the network *connectivity* (c). It represents the probability of two nodes sharing a link. For a directed network, with links between nodes in two directions, connectivity is given by $c = l / (n \cdot (n - 1))$. For a connected network (i.e. without disconnected components) $l \geq n - 1$. In the special circumstance $l = n - 1$ the network is a so called tree network with minimal connectivity $c = 1/n$. Connectivity reaches its maximum value $c = 1$ for a completely connected network. All possible links have then been realized. Reciprocity, finally, is the fraction of links with a link in the opposite direction (range from 0 to 1).

Path length

A *path* is an alternating sequence of connected nodes and links that starts and terminates at a node. If all links represent unit length, *path length* l_{ij} between nodes i and j is the length of the shortest path between the nodes. The average path length l_i for node i is the average distance to all other nodes. Although a directed network in principle consists of directed paths that are being traversed in the direction of the links, direction is not taken into account here. The path represents a connected sequence of contacts in the form of transactions rather than a sequence of directed flows of payments. Link weights in terms of value transferred may vary strongly over one path so that direction of flow, without explicitly taking into account link weights, not necessarily contains very valuable information. *Average network length* l_{avg} is the average of all path lengths l_i . It determines the average undirected shortest path. Network *eccentricity* (e) is defined as the largest of the observed path lengths: $e = \max_{i,j} (l_{i,j})$.

Degree

The number of links between one node i and other nodes determines the *node degree* (k_i). In a directed network these connections consist of incoming and outgoing links, which respectively determine the *in-degree* ($k_{in,i}$), the *out-degree* ($k_{out,i}$), and *node degree* (k_i) by $k_i = k_{in,i} + k_{out,i}$. Every link contributes exactly one unit to both the out-degree of the node at which it originates and to the in-degree of the node at which it terminates. The *average degree* (k_{avg}) of a network is the relative number of all links to all nodes: $k_{avg} = l/n = 1/2n \sum_i k_i = 1/n \sum_i k_{in,i} = 1/n \sum_i k_{out,i}$.

³⁰ Based on Dorogovtsev and Mendes (2003) and Soramäki, et al. (2007).

The *maximum in-degree* $k_{in,max} = \max_i(k_{in,i})$, *maximum out-degree* $k_{out,max} = \max_i(k_{out,i})$ and *maximum degree* $k_{max} = \max_i(k_i) = \max_i(k_{in,i} + k_{out,i})$ determine the maximum degree values and the maximum deviations (to the upside) from the respective average degree values. More informative and more elaborate to determine are the degree distributions $P(k_i)$, $P(k_{in,i})$ and $P(k_{out,i})$ for a specific node i . Summation over all nodes i and taking averages results in the total degree distributions $P(k)$, $P(k_{in})$ and $P(k_{out})$. Two examples of degree distributions are respectively a Poisson distribution and a power-law distribution. The former results when a fixed number of nodes is randomly connected on the basis of the fixed network degree k_{avg} . Larger networks asymptotically follow the Poisson

distribution $P(k) = \frac{e^{-k_{avg}} \cdot k_{avg}^k}{k!}$ (classical equilibrium network). In practice, however, many networks

have (relatively recently) been found to follow a power-law distribution $P(k) \propto k^{-\gamma}$. These non-equilibrium networks are characterized by fat-tails which mark the relatively high frequency of highly connected nodes in comparison to classical equilibrium networks. They originate from growing networks in which new nodes (linearly) preferentially attach to other nodes. They have no natural scale and are called scale-free networks. In recent years it has been demonstrated that many social, informational, technological and biological networks have fat-tailed, scale-free degree distributions (see for instance Amaral, et al. (2000), Newman (2003), and Dorogovtsev and Mendes (2003)).

Degree correlations

Degree correlations between neighbouring nodes provide additional information on the network structure. In an uncorrelated network the degree of one node is independent of its neighbouring nodes. Degree correlations therefore provide information on whether nodes are generally connected to nodes with comparable degree, to nodes of different degree, or if there is no relation at all. Classical random networks have no correlations. Fat-tailed, scale-free networks on the other hand may exhibit strong correlations.

Several measures exist for degree correlations. For example:

- Between k_{in} and k_{out} for individual nodes
- Between k_{in} and k_{out} , k_{in} and k_{in} , or k_{out} and k_{out} for two nodes

Clustering coefficient

Another concept to describe the correlation between nodes is the *clustering coefficient* (C_i), which gives the probability that two neighbours of a node share an undirected link among themselves. It marks the density of connections in the direct neighbourhood of a node (cliquishness). The clustering coefficient is determined by the number of actual undirected links between nearest neighbours ($l_{nn,i}$) of

a node i as a fraction of the number of possible undirected links: $C_i = \frac{2l_{nn,i}}{k_i \cdot (k_i - 1)}$. The average

clustering coefficient (C_{avg}) over all nodes determines the network clustering. The meaning of the coefficient becomes particularly clear in a social context where it is the extent of the mutual acquaintance of friends. The clustering coefficient ranges from 0 for a tree network to 1 for a completely connected network. The classical random network locally has a tree-like structure (loops cease to exist in the infinite network). Fat-tailed, scale-free networks may exhibit strong clustering.

REFERENCES

- ALBERT, R. (1999): "Diameter of the World Wide Web," *Nature*, 401, 130-131.
- ALBERT, R., and A.-L. BARABÁSI (2002): "Statistical Mechanics of Complex Networks," *Reviews of Modern Physics*, 74, 47-97.
- ALBERT, R., H. JEONG, and A.-L. BARABÁSI (1999): "Error and Attack Tolerance of Complex Networks," *Nature*, 406, 378-382.
- ALLEN, F., and D. GALE (2000): "Financial Contagion," *Journal of Political Economy*, 108, 1-33.
- AMARAL, L. A. N., A. SCALA, M. BARTHÉLÉMY, and H. E. STANLEY (2000): "Classes of Behaviour of Small-World Networks," XXXX.
- BARABÁSI, A.-L., and R. ALBERT (1999): "Emergence of Scaling in Random Networks," *Science*, 286, 509-512.
- BATAGELJ, V., and A. MRVAR (2006): "Pajek Manual: Program for Analysis and Visualisation of Large Networks," <http://vlado.fmf.uni-lj.si/pub/networks/pajek/>, Version 1.15.
- BECH, M. L., and R. GARRAT (2006): "Illiquidity in the Interbank Payment System Following Wide-Scale Disruptions," *Federal Reserve Bank of New York Staff Reports*, 239.
- BECH, M. L., and K. SORAMÄKI (2005): "Systemic Risk in a Netting System Revisited," in *Liquidity, Risks and Speed in Payment and Settlement Systems - a Simulation Approach*, ed. by H. Leinonen: Bank of Finland Studies, 151-178.
- BIS (2007): "Redbook: Statistics on Payment and Settlement Systems in Selected Countries."
- BOSS, M., H. ELSINGER, M. SUMMER, and S. THURNER (2004): "The Network Topology of the Interbank Market," *Quantitative Finance*, 4, 677-684.
- DOROGOVTSSEV, S. N., and J. F. F. MENDES (2003): *Evolution of Networks*. Oxford: Oxford University Press.
- ECB (2005): "Target Annual Report."
- ERDŐS, P., and A. RÉNYI (1959): "On Random Graphs," *Publicationes Mathematicae*, 6, 290-297.
- FALOUTSOS, M., P. FALOUTSOS, and C. FALOUTSOS (1999): "On Power-Law Relationships in Internet Topology," *Computer Communications Review*, 29, 251-262.
- FURFINE, C. H. (1999): "The Microstructure of the Federal Funds Market," *Financial Markets, Institutions and Instruments*, 8, 24-44.
- GOODRICH, M. T., and R. TAMASSIA (2006): *Data Structures and Algorithms in Java*. London: John Wiley & Sons.
- INAOKA, H., T. NINOMIYA, K. TANIGUCHI, T. SHIMIZU, and H. TAKAYASU (2004): "Fractal Network Derived from Banking Transaction - an Analysis of Network Structures Formed by Financial Institutions," *Bank of Japan Working Paper*, 04-E-04.
- LEDRUT, E. (2006): "A Tale of the Water-Supplying Plumber: Intraday Liquidity Provision in Payment Systems," *DNB Working Paper Series*, 99.

- LELYVELD, I. P. P. VAN, and F. R. LIEDORP (2006): "Interbank Contagion in the Dutch Banking Sector," *International Journal of Central Banking*, 2, 99-134.
- LUBLÓY, Á. (2006): "Topology of the Hungarian Large-Value Transfer System," *Magyar Nemzeti Bank Occasional paper*, 57.
- NEWMAN, M. E. J. (2003): "The Structure and Function of Complex Networks," *SIAM Review*, 45, 167-256.
- OORD, A. V., and H. LIN (2005): "Modelling Inter- and Intraday Payment Flows," *DNB Working Paper Series*, 74.
- PERICOLI, M., and M. SBRACIA (2003): "A Primer on Financial Contagion," *Journal of Economic Surveys*, 17, 571-608.
- PRÖPPER, M. H., I. P. P. VAN LELYVELD, and R. H. HEIJMANS (2007): "Time and Interbank Payment Networks," *Mimeo*.
- SORAMÄKI, K., M. L. BECH, J. ARNOLD, R. J. GLASS, and W. E. BEYELER (2007): "The Topology of Interbank Payment Flows," *Physica A*, 379, 317-333.
- Rosati, S. and S. Secola (2005), "Explaining cross-border large-value payment flows – Evidence from TARGET and EURO I data," ECB Working Paper Series, No. 443