DNB Working Paper
No. 683 / April 2020

Is there anybody out there? Detecting operational outages from LVTS transaction data

Neville Arjani and Ronald Heijmans

DeNederlandscheBank
EUROSYSTEEM

Is there anybody out there? Detecting operational outages from LVTS transaction data

Neville Arjani and Ronald Heijmans*

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

Working Paper No. 683

April 2020

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

Is there anybody out there? Detecting operational outages from LVTS transaction data *

Neville Arjani[a] and Ronald Heijmans[b,c]

[a]*Canada Deposit Insurance Corporation*
[b]*De Nederlandsche Bank*
[c]*Payments Canada*

April 2020

Abstract

This paper develops a method to identify operational outages of participants in the Canadian Large Value Transfer System. We define an operational outage as either no or unusually low activity. We test our algorithm against a database of outages by participants reported in order to reduce false negatives. The false positives can be reduced by excluding "outages found" by the algorithm if a participant historically has no payment in a given five minute time interval. Additionally, we can test whether participants do indeed report all their operational outages. The results show that our algorithm works best for the largest participants as they send in payments continuously. Our method can be used by LVTS system operators and overseers to identify sources of operational risks.

Keywords: Operational Risk, LVTS, Financial Market Infrastructures, outages.
JEL classifications: E42, E58, G21.

*Arjani worked as research director at Payments Canada at the time this research project was carried out. Heijmans is the corresponding author and can be reached at ronald.heijmans@dnb.nl. The authors would like to thank Jan Paulick and Marc Glowka for providing useful comments.

1 Introduction

Financial market infrastructures (FMIs) play a crucial role in the economy. FMIs represent vast financial networks that connect financial intermediaries (and their clients) by means of technology, rules, standards and procedures to facilitate the transfer of value between agents in an economy. Large value payment systems (LVPSs), including real-time gross settlement (RTGS) systems, are a type of FMI that facilitate the exchange, clearing and final settlement of typically large-value and time-sensitive monetary transfers between participating financial intermediaries on their own behalf and that of their clients.[1] They are a critical component of a country's financial system and economy, as they serve as the platform for the implementation of monetary policy, as well as supporting hundreds of billions of dollars in real and financial transaction activity each day. Needless to say, if these systems do not function properly, they can seriously hinder economic activity. Therefore, FMIs have to live up to high international standards and are generally subject to a strong degree of regulatory and operational oversight, see CPSS (2012).

As LVPSs connect financial intermediaries, interdependencies exist where the actions of one participant in the network will have an impact on the others. For instance, if one participant in the network suffers a technical outage that prevents it from sending payment instructions to other participants in the network as planned, this could result in other participants only being able to fulfill their payment obligations at a higher cost as they must seek intraday funding somewhere else to meet these obligations, e.g., through a central bank intraday credit facility.[2] During the technical outage, the stricken bank serves as a liquidity sink in which incoming liquidity stays trapped on its account and cannot be recycled to the rest of the system. This problem is exacerbated when other participants, either unaware or dismissive of the fact that the stricken bank is unable to send payments, continue to funnel transfers to the stricken bank through the net-

[1] LVPSs around the world include the US Fedwire and US CHIPS systems, BOJ-NET in Japan, the TARGET2 system in the Euro Area, UK CHAPS, and the Canadian LVTS. Virtually all major economies around the world have adopted a real-time gross settlement (RTGS) model for their LVPS implementation (Bech and Hobijn, 2007).

[2] Incoming payments represent the cheapest source of intraday liquidity for participants in these systems.

work.[3] In the worst-case scenario, a technical outage at a participant could result in a gridlock in which no participant is willing or able to make any more payments and waits for others to make the first payment.[4] To mitigate this possibility, timely indication of participant technical outages is key. Therefore, operators of an FMI diligently monitor the activity of their system's participants.[5]

The question we address in this paper is how can operational outages of individual participants in the Canadian Large Value Transfer System (LVTS) can be detected in a more timely manner. Even though LVTS participants are required by the LVTS Rules to inform the system operator within 15 minutes of detecting a technical disruption in their ability to send and/or receive payments, anecdotal evidence suggests that this is often not the case, for a variety of reasons.[6] For instance, a participant that is busy trying to restore connectivity may may miss the 15-minute deadline, or forget to contact Payments Canada altogether. In economics parlance, this generates an externality by increasing the risk of financial loss to the system as a whole through this single outage. In all cases, this increases the risk of financial loss associated with operational outages. Given the strong seasonal patterns observed in the LVTS on an intraday, daily, monthly and even quarterly basis, coupled with the prospective real-time availability of these data to Payments Canada and in the development of machine learning algorithms focused on anomaly detection, it seems that more timely and accurate detection of operational outages in the LVTS is certainly possible. This has the potential to reduce the cost associated with such events, thereby reducing the risk profile of the system as a whole.

To detect operational outages, we build on the approach by Glowka et al. (2018) and Klee (2010), who have developed an algorithm for the European RTGS, TARGET2, and

[3] This could be due to a healthy participant thinking that the stricken bank will recover quickly.

[4] Liquidity injections are still possible of course (e.g., central bank lending facility) in this scenario, but they too carry a financial cost.

[5] Moreover, some system operators may choose to impose strict operational availability requirements on participants in the network (as is the case in Canada), which are subject to compliance protocols if not adhered to. As well, given the importance of these systems, there are often manual or other special back-up facilities that might be available for an affected participant to continue to send value over the network, usually to facilitate any time-sensitive transfers. While these back-up facilities are typically not designed to handle the regular daily flow of participants, they can at least be used to help circulate some liquidity from the affected participant back to the network to help avoid a gridlock scenario. Before backup facilities can be started it has to be known that there is an (operational) problem with a participant.

[6] Non-compliance with this rule garners further follow-up and potential penalty.

the American RTGS, Fedwire, respectively. Glowka et al. (2018) looks in contrast to Klee (2010) at periods of unusual low activity. They argue that, in the event that a participant is not able to send in payment instructions via the standard channels, there may be back-up functionalities that allow at least a few (very urgent) payments to be made. We modify their approach to make it suitable for the Canadian LVTS. This type of algorithms inevitably leads to potential Type I and Type II errors. A Type I error would be identifying an outage when non has occurred , because the direct clearing member had no (or few) payments to send in. A Type II error would be not identifying an operational outage when in fact there was one. This could be the case if the activity during the outage period exceeds the 1 percentile threshold used. Unlike Glowka et al. (2018), we have a data base containing the participants' reported outages over a two year period, which we used to test our algorithm.

Our paper adds to the growing literature on identifying risks for FMIs. Benos and Zimmerman (2012) develop risk indicators measuring the impact of liquidity risk due to operational outages. Their findings indicate that the impact of operational outages on system liquidity has increased since the collapse of Lehman Brothers, compared with the previous observation period. Berndsen and Heijmans (2020) develop risk indicators based on TARGET2 transaction data that are linked to the principles of financial market infrastructures, can be used by FMI operators and overseers. Clarke and Hanckock (2013) study how the design of the payment system can impact operational disruptions. Diehl and Müller (2015) study the use of bilateral limits in TARGET2. They find that limits are not actively used and if they are used they are relatively constant over time. Both papers find that liquidity saving mechanisms in an LVPS can reduce the impact of an operational outage. Triepels et al. (2018) implemented a neural network to identify intraday liquidity outliers of an individual bank for TARGET2. Their method can detect starting bank runs from RTGS data. One frequently investigated aspect of large value payment systems is timing. There may be timing incentives in the payment system, as described theoretically by the game theoretical model by Bech and Garratt (2003, 2006). Their model has been run in an experimental real life game by Abbink et al. (2017) and by Heemeijer and Heijmans (2015). Heijmans and Heuver (2014) study several liquidity aspects that can be derived from RTGS data including payments' timing.

The main sources of liquidity they identify from LVPS data are incoming payments, interbank money market loans (a special subset of interbank payments), monetary loans (transactions with the central bank) and intraday credit (which is backed by collateral), which allows them to make payments even though their account balance is insufficient). Kaliontzoglou and Müller (2015) build on this work to derive a payments delay indicator, which can be used for all banks in the Eurosystem. Diehl (2013) explains that free-riding incentives may be behind the timing of payment transactions, although they find no free riders in the German part of TARGET2.

This paper is organized as follows. Section 2 describes the main features of LVTS and their data. Section 3 explains the set up of the algorithm used to identify the operational outages. Section 4 provides the results of the algorithm and the check against the reported operational outages. Section 5 concludes.

2 Large Value Transfer System

2.1 The system

The Large Value Transfer System (LVTS) is used to facilitate interbank settlement in Canada. Owned and operated by Payments Canada, the LVTS has been in place since 1999 and clears roughly CA$C200 billion in transaction value each day. This translates to clearing roughly the equivalent of Canadian nominal GDP every nine business days.

The LVTS is best described as a real-time net settlement system. At the heart of the LVTS settlement model stands novation netting, where bilateral settlement obligations stemming from intraday exchange of payment instructions between pairs of participants are immediately extinguished upon LVTS processing and replaced with a multilateral settlement obligation on the part of the sender to the rest of the system. Settlement of multilateral obligations of participants should occur at 18:30 EST on each business day. Settlement takes place in central bank money through daily transfer of value over participants' settlement accounts maintained at the Bank of Canada. Settlement risk stemming from this model is managed via a dual-stream risk model.

The participants can use two different streams in LVTS. It is their choice which stream

they send the payments to. The two streams (or Tranches) differ with regard to the way settlement risk exposure generated by a participant is controlled. In the first stream, or Tranche 1, any intraday credit exposure posed by a participant to the rest of the system is supported on a dollar-for-dollar basis with financial collateral posted to the central bank by the participant ('defaulter-pays'). In the second stream, intraday credit exposure posed by a participant to the rest of the system is supported by a joint ('survivors-pay') collateral pool and further backed by a central bank commitment. Through a combination of financial collateral, and corresponding bilateral and multilateral net debit limits, the LVTS risk model ensures that, at a minimum, there will always be sufficient collateral value posted to cover the largest possible multilateral net settlement obligation of any single participant in the system. The central bank commitment helps to further ensure that the LVTS will settle under all circumstances. Additionally, a participant may use either Tranche to send payment instructions, subject to applicable bilateral and multilateral net debit limits.[7] There were 17 participants in the LVTS (see Payments Canada) including domestic and international financial intermediaries. There is a high degree of concentration, whereby the largest six LVTS participants are typically responsible for over 90% of daily LVTS transfer value. The five smallest banks are only responsible for only 1% of the turnover and fewer than 1% of the transactions. This means that the participants in LVTS are highly heterogeneous.

2.2 Transaction data

We have transaction data available for Tranches 1 and 2 for the period 2002- 2018. For each transaction we have used the settlement date, settlement time, sending and receiving participant, and amount. The number of participants varies between 14 in 2002 and 17 in 2018.

The data has been cleaned for the following transactions. First, we only kept transactions between 8.00 AM to 6.00 PM (EST). These are the normal opening hours in which banks are active. LVTS is also open besides these times, but activity is normally limited. During the night, for example, the payments to CLS are made;[8] These CLS

[7] see Arjani and McVanel (2006) for a more in-depth discussion of the LVTS model.

[8] CLS stands for continuous linked settlement and takes care of foreign exchange transactions. It oper-

payments are highly urgent but very limited in number. Second, we removed all outgoing transactions from the central bank as we are interested in outages of commercial banks only. Payments from a participant to the central banks remain in the sample as this constitutes payment activity of a commercial bank (participant). Third, we removed all Canadian public holidays from the transaction data set. On most public holidays the system will be closed, which means that be no payments at all are executed in. There are also 'regional' public holidays on which the system is open, but the activity can be much lower than usual. Moreover, activity may also slow down around US public holidays. This can cause false positives in which outages are detected by our algorithm, but in fact thees are just the reduced activity of a normal 'slow' day. Some LVTS participants are foreign bank branches or subsidiaries operating in Canada, and thus have their headquarters outside the country, e.g. in France, the United Kingdom or the United States. For these participants, we also removed the transactions on days on which their home country has a public holiday.

Figure 1 gives insight into the number of transactions settled in LVTS. Figure 1a shows that the monthly daily average number of transactions has increased from roughly 26,000 in 2012 to 36,000 in 2018. This rise is mainly the result of the larger banks increasing their activity in LVTS. Figure 1b shows the percentage share of the transactions settled in a given hour. Transactions listed under '8' are those settled between 8.00 AM and 9 AM, etc. It is clear from this figure that the number of transactions is at its lowest the last two hours of the day. The last hour, especially, has a very low transaction volume. This is mainly due to the fact that banks use this hour primarily to make sure they have sufficient liquidity available in their account by the system's closing time.

2.3 Reported operational outages

Besides LVTS transaction data, we have a list of operational outages reported by LVTS's participants. This states when they were unable to make any payments (Severity 1 outage) or only a few transactions (Severity 2 outage) in the years 2016 and 2017. As well as the participant name and severity type, the list contains the dates and the start and end times of the incidents. The total numbers of reported outages in 2016 and 2017

ates in many different time zones.

Figure 1: Number of transactions in LVTS between 2012 and 2018.

(a) Monthly daily averages.

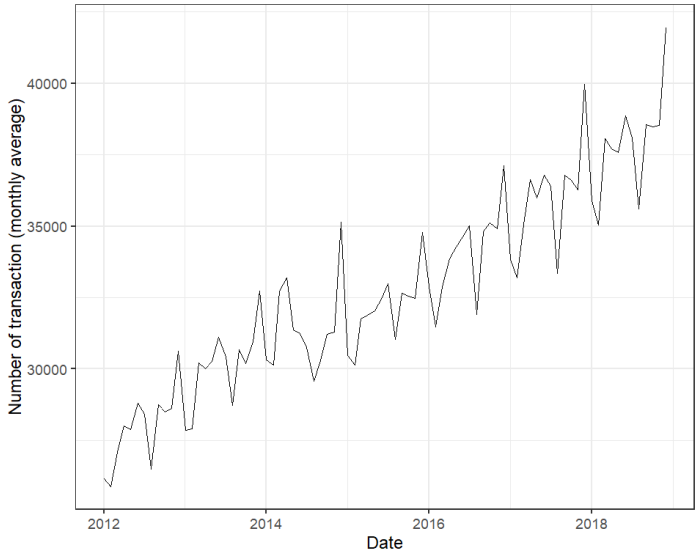

(b) Hourly averages.

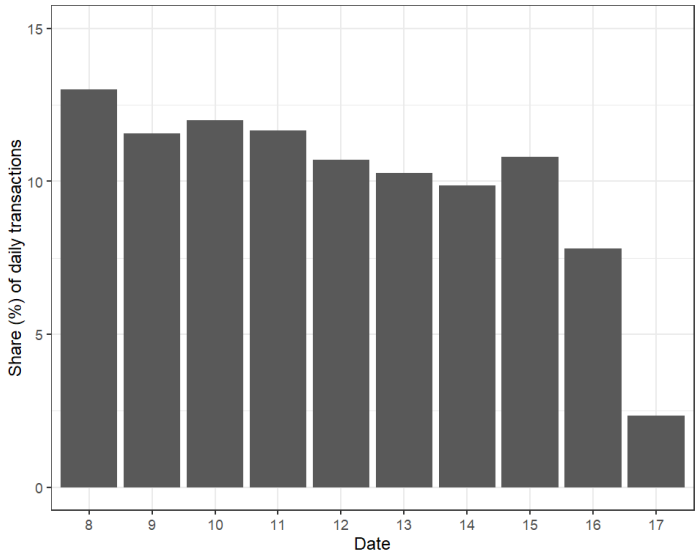

were 23 (16 Severity 1 and 7 Severity 2) and 26 (19 Severity 1 and 7 Severity 2), respectively. We have only included outages that took place between 8 AM and 6 PM. The average outage duration in those two years were 02:05 and 02:25 hours, respectively. The longest operational problem recorded in 2016 lasted 12:40 hours and in 2017 12:46 hours.

2.4 Special arrangements or features

Procedures to follow in the case of an operational outage are governed by LVTS Rule 12. Rule 12 and the accompanying procedures distinguish between different types of outages. A Severity 1 outage occurs when a participant is unable to send/receive any payments via the LVTS. A Severity 2 outage is when the participant is able to send payments, but not at a normal processing pace; its normal payment processing operations are still disrupted. For Severity 1 and Severity 2 outages alike, section 12.14 of Rule 12 states that, if a participant encounters a technical, site or external problem (e.g., SWIFT connection problem, payments environment problem, site problem, or weather or civil problem) that may impact its ability to send or receive LVTS payment messages, the participant must notify the LVTS Help Desk within 15 minutes of becoming aware of the problem, identifying itself, the nature of the problem and, if possible, the expected time by which the situation will be resolved. In the case of time-sensitive payments or given risk of a material liquidity trap, contingency procedures are available to mitigate risk of further disruption to other participants and FMIs. For instance, if a participant is experiencing a Severity 1 Incident and cannot send payment messages via its SWIFT interface because of a technical, site or external problem that results in trapped liquidity or the inability of that participant to meet critical time-sensitive obligations (e.g. CLS pay-ins, funding for ACSS settlement, and CDSX settlement), the participant may use the direct network to send payment messages in accordance with the procedures contained within Rule 12. Note that the direct network can only be used to foster interbank transfers (SWIFT MT 205 messages).[9]

3 The algorithm

In line with Glowka et al. (2018), we also look at two different types of potential outages. First, if there are no transactions at all during a particular five minute time interval, this interval is referred to as one of "no payment activity (NoPA).[10] Second, if activity

[9] SWIFT message type 205 is often referred to as 202 in other jurisdictions.

[10] They look at intervals of 10 minutes, being one-third of the length of an outage that should be reported by a participant in TARGET2. In the LVTS, participants must inform Payments Canada within 15 minutes of encountering a disruption. We Therefore look at five minute intervals, this also being one-third of the minimum report interval length. However, this period can be set to any desired value by the operator.

is much lower than usual (e.g. below 1 percentile), meaning that only a very limited number of transactions is taking place, we refer to this interval as one of "low payment activity (LowPA)". The reasoning for looking at periods with lower transaction volumes, as well as none at all, is that financial institutions have a manual contingency procedure whereby they can still process a limited amount of very urgent payments during an outage. In this way we also limit susceptibility to false negatives, as these periods with low activity should be seen as potential outages, but are not picked up by the NoPA method.

Unlike Glowka et al. (2018), we do not have to exclude extremely small participants. After all, LVTS only has 17 participants. By contrast, TARGET2 has approximately 1000 direct participants, many of them very small ones in terms of payment activity. Nevertheless, the difference in activity between the largest and smallest participants in LVTS is also quite wide.

If our algorithm were to see a period of low or no activity by a small participant, that would not necessarily mean that it had indeed faced a technical outage. That could instead be the result of not having any payment obligations to settle for some time, or of an intentional delay (with no technical problems). Besides, with the available data, we are not be able to prove in real-time that there was indeed an outage. Our algorithm is intended to inform the payment operator that a participant has had low or no activity over say for 15 minutes. If this is too long a period of inactivity from the operator's perspective, it can then contact the participant to see what is going on. Lastly, participants only have to report to the operator when they have had a technical outage of at least 15 minutes, which is three consecutive monitoring intervals in our algorithm.

3.1 No payment activity

The "no payment activity algorithm" (NoPA) identifies periods during which a participant has sent no payment instructions into the LVTS. We look at intervals with no

transactions of five minutes.

$$NoPA = \begin{cases} yes, \text{if} \quad TNR_{fi,t} = 0 \\ no, \text{if} \quad TNR_{fi,t} > 0 \end{cases} \tag{1}$$

where $TNR_{fi,t}$ is the number of transactions financial institution fi has at five minute time interval t.

Figure 2 illustrates the absolute (2a and 2b) and relative (2c and 2d) number of five minute periods without transactions in LVTS for all banks (2a and 2c) and for the largest five banks (2b and 2d) in LVTS, respectively.[11]

Figure 2c shows that the number of five minute periods without transactions is just below 50%. The reason for so many periods without any transactions is that quite a few banks are not very active in the system throughout the whole day. These banks have periods in which they do not send in any payments. If we zoom in on the five largest financial institutions, which tend to be the most active banks, this number decreases to below 20%, see Figure 2d. The number of periods without transactions decrease over the years, to close to 15%. This remains a relatively large number as there are still no transactions in 1 out of every 5-6 periods of five minutes.

3.2 Low payment outages

By contrast with the NoPA outage setup, in the low payment activity (LowPA) setup we do accept a few payments being sent in to the LVTS. We base this low activity on the behaviour of the individual participant. Low activity is defined as the number of transactions settled in a given 5 minute interval over the whole data sample below 1 percentile. We zoom into the payment behaviour of that particular participant during that particular time interval. First, participants often send in more payment instructions at given times, for example early in the morning and in the afternoon, and have fewer payments at other times, see e.g. Massarenti et al. (2012) or Heijmans and Heuver (2014). Second, we wish to exclude periods of no activity when they are 'normal' for that

[11] As the number of financial institution increased from 14 in 2002 to 17 in 2017, the absolute number of five minute intervals increased accordingly. As the number of business days varies per year, the total number of five minutes intervals also varies.

Figure 2: Five minute intervals with and without transactions.

(a) Absolute number all banks.

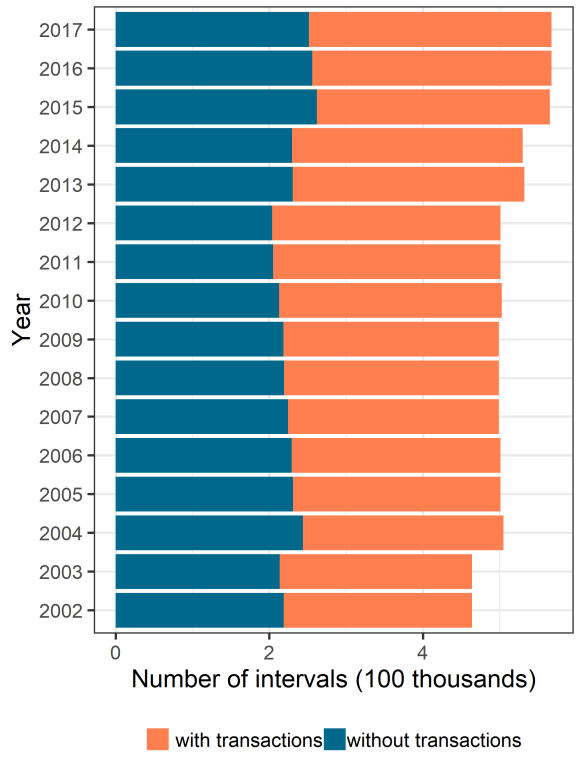

(b) Absolute five largest banks.

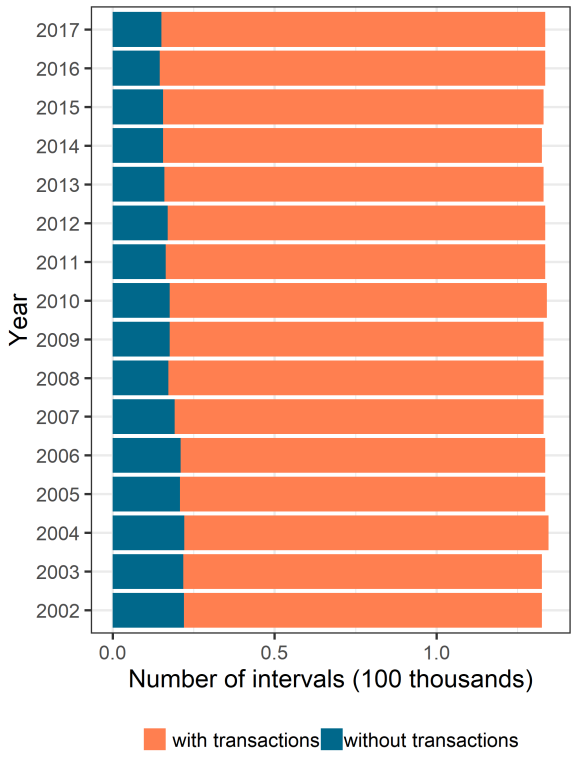

(c) Relative share all banks.

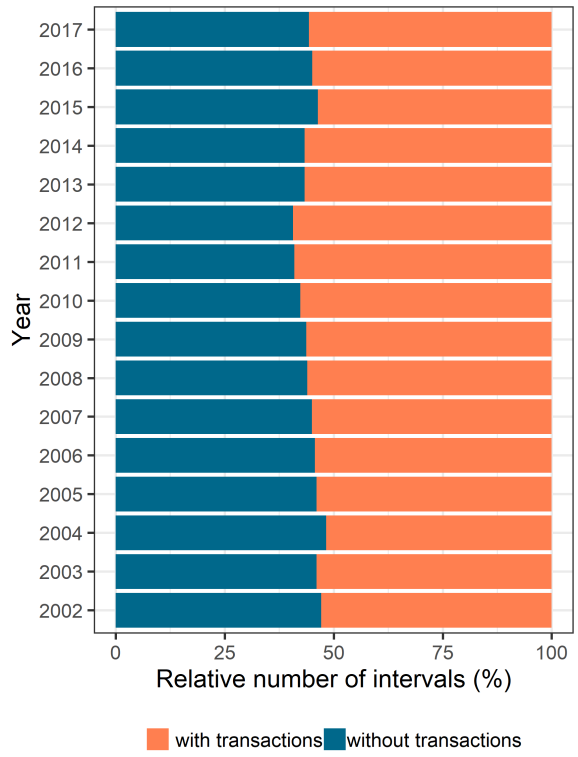

(d) Relative share five largest banks.

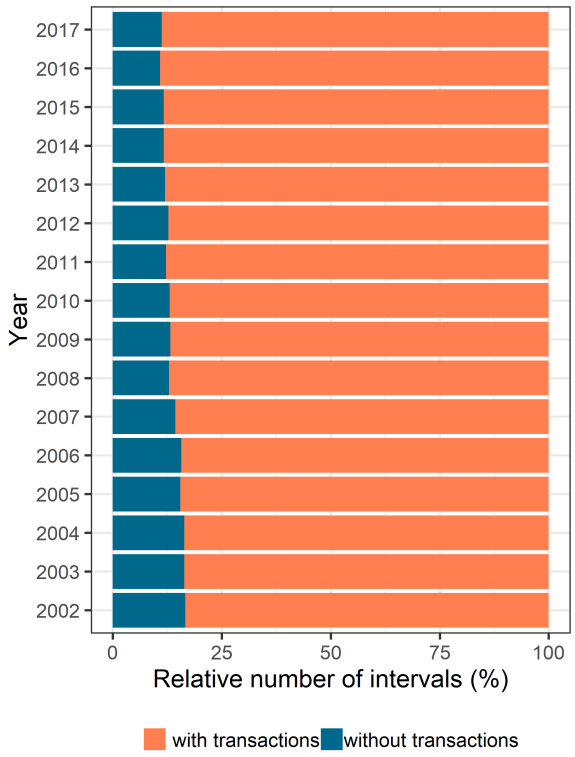

particular participant. To be able to look at unusually low payment activity we must correct for this natural difference in payment activity at the participant level:

$$\frac{TNR_{fi,t}}{\overline{TNR_{fi,t,y}}} < 1\,percentile \tag{2a}$$

$$TNR_{fi,t} > 5 \tag{2b}$$

where $TNR_{fi,t}$ is again the number of transactions a participant fi has at five minute time interval t $\overline{TNR_{fi,t,y}}$ is the mean number of transactions of participant fi at five minute interval t over the year y. The $\overline{TNR_{fi,t,y}}$ is the mean per year as the numbers may change over the years.

In our LowPA model setup, an event is considered an outage if the payment activity is below the 1 percentile for that particular participant and time interval (equation 2a) and at least five transactions have been sent in (equation 2b).

4 Results

This section describes the outcome of our outage detection method compared to the reported outages. Section 4.1 shows the detected outages for the period ranging from 2002 to 2017. Section 4.2 describes the intra week and intraday patterns of the outages. Unlike Glowka et al. (2018), we have a list of outage events reported by the participants. This allows us to check whether they do indeed report their outages, see section 4.3. These reported outages presents an opportunity to assess the validity of our model and the empirical likelihood of false positives. There may be an incentive for participants to not report their operational outages. First, per the LVTS Rules participants need to have up-time of at least 98% during each calendar month. Every time they report an outage, their up-time decreases. Second, they do not want to be known by the operator (and their peers) as a bad performer in terms of their technical infrastructure.

4.1 Outages detected by the NoPA and LowPA algorithms

Figure 3 shows the number of outage intervals in the NoPA and LowPA models for all banks and for the largest five banks. As expected, the number of outages decreases in the low payment activity (LowPA) setup we do. As Glowka et al. (2018) state, the longer an outage, the less likely that it is just a coincidence.

Figures 3a and 3b show the outages for the differentiated model with low activity. The number of outages found by this model setup for the largest five banks in the system is close to the numbers found for all banks (including the largest five). The reason for this is that outages are only included if the number of transactions found is lower than the 1 percentile. For some of the smaller banks, the 1 percentile is equal to zero transactions. As lower than zero transactions is not possible, these periods will not be picked up by this second algorithm (although they will be picked up by the first).

4.2 Intraweek and intraday Outages.

Figure 4 shows the number of outages (lasting at least four consecutive five minute periods) found by the LowPA algorithm, differentiated by day of the week. We use a minimum of four consecutive intervals in order to exceed the reporting obligation time interval of 15 minutes. The figure clearly shows that there are more potential outages detected on Monday than over the rest of the week combined. According to the system operator, this can be explained by the fact that software updates and new releases of the LVTS system and the participant's internal systems are usually introduced over the weekend, and any resulting operational problems occur on the next working day. But these are usually solved on the same day. As the number of outages detected on Tuesdays is much smaller than that on Mondays, and is in line with the rest of the week, this seems plausible.

Figure 5 depicts the outages found by the LowPA model by the hour of the day. The last two hours monitored, 4-6PM, see many more outages found than the rest of the day. This is because the participants typically use that period to make sure that their end of day balance is adequate, resulting in payments which are relatively high in value but low in number. The other hours of the day do not show a large variation.

Figure 3: Number of successive single, double, triple and quadruple intervals without transaction.

(a) LowPA, all banks.

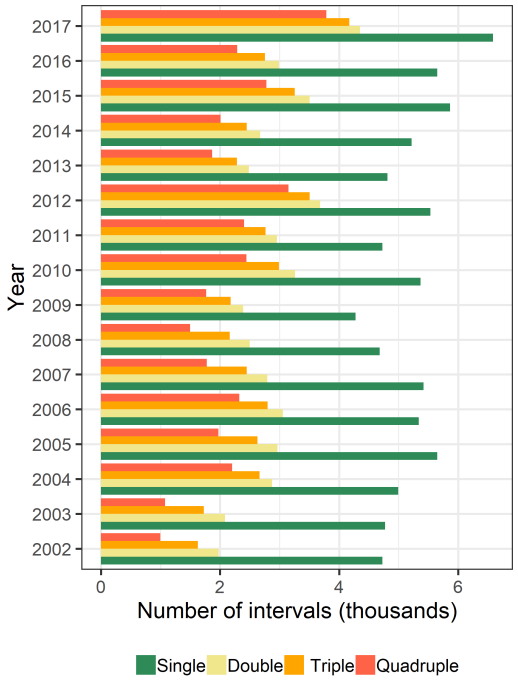

(b) LowPA, five largest banks.

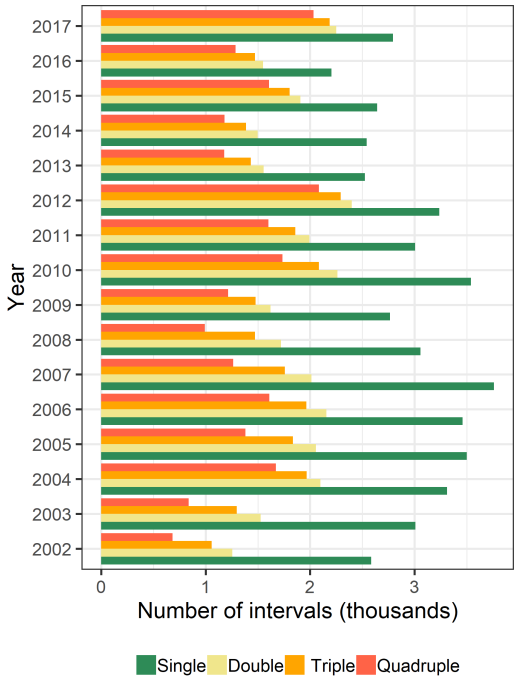

(c) NoPA all banks.

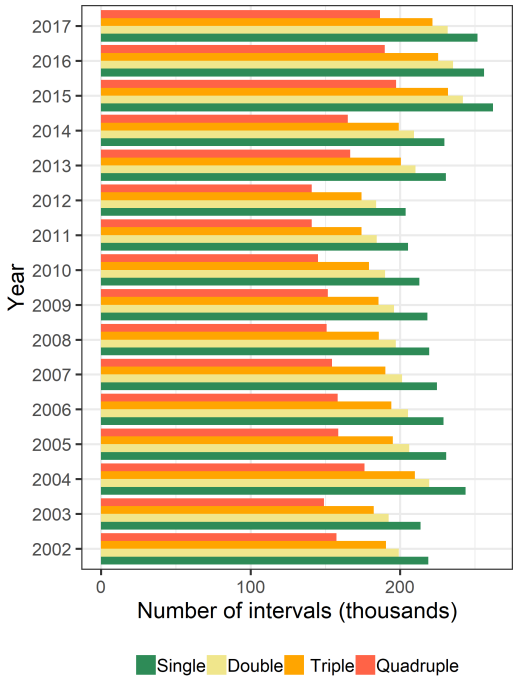

(d) NoPA five large banks.

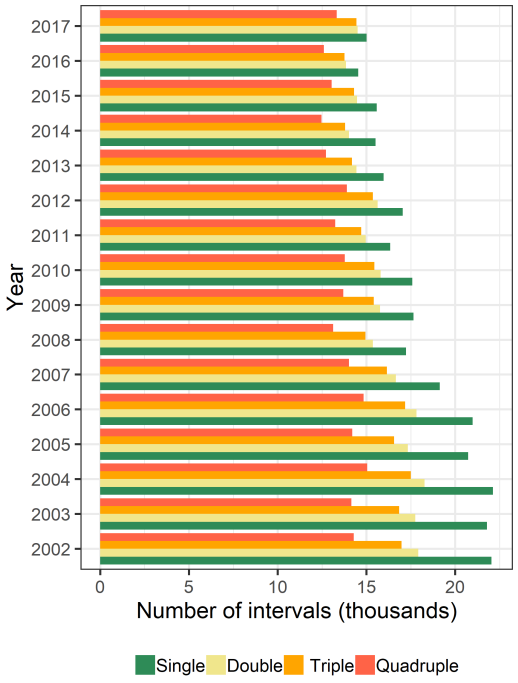

Figure 4: Outages found by LowPA algorithm by day of the week.

(a) all banks

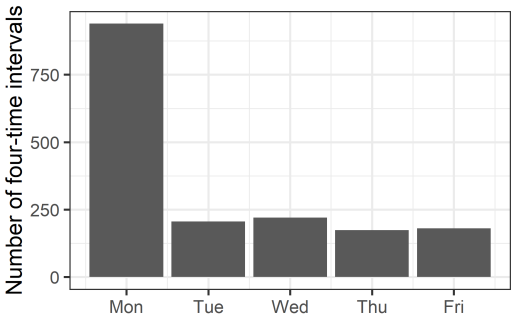

(b) 5 large banks

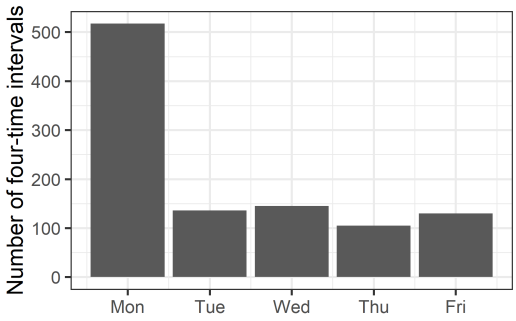

Figure 5: Outages by LowPA algorithm by hour of the day.

(a) All banks.

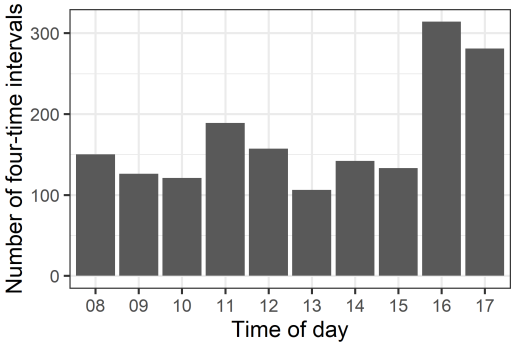

(b) Five large banks.

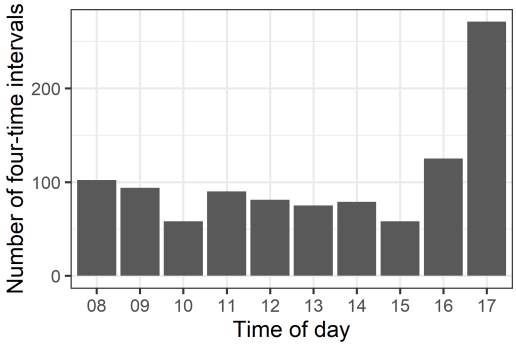

4.3 Validation of the algorithm: detected versus reported outages

We now move to the list with outages reported by the participants, see section 2.3. We check whether the algorithm picks up those periods as outages, which is basically a check of the false negative error. To check the algorithm we only include: 1) those participants with a reported outages; and 2) time interval of the reported outage. The time interval has to be between 8.00 AM and 6.00 PM as this is the period for which we studied the data. As the algorithm works best for the largest five participants in the system, we focus on them in this section. https://www.overleaf.com/project/5e15f17fa8c03300016d9875

Figure 6 shows the relative number of periods with and without transactions for the largest five banks. During 75% of the reported outage periods, no transactions are detected at all. In the other 25%, some payments are detected. Most of these are related

Figure 6: Periods of outages reported by the five large banks.

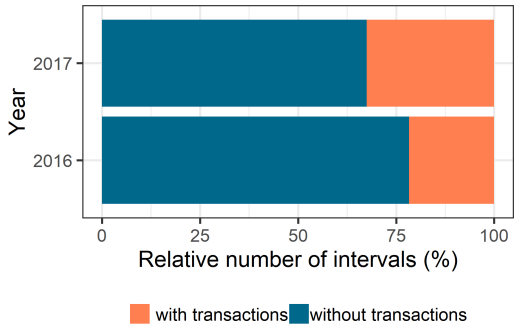

to Severity 2 incidents in which participants are still able to make some payments. The number of payments is still higher than the 1 percentile and therefore not excluded by the algorithm. Ideally, an operator may want our algorithm to pick up these Severity outages. In the event of a Severity 2 error, however, it is aware of the problem anyway and so does not need to be informed by the algorithm. Some of the payments are within the same five minute interval investigated, but occur (just) before the outage started. This means that if the outage started at 10.04 AM and a few payments had already been between 10.00-10.03 AM the five minute interval (10.00-10.05 AM) has been counted towards the intervals with transactions.

Figure 7 shows the number of successive periods with low (7a) and no payment activity (7b) at the time of outages reported by the five large banks. The number of outage intervals without transactions is much larger than the number with transactions. From Figure 7b we can see that the number of quadruple period outages is almost the same as the number of single period ones. As all outages except one lasted longer than three five minute intervals, we can state that all reported outages are picked up by the algorithm.

Looking at the LowPA output (7a) we see that there is a few cases where there is still some activity when an outage is reported. Some of these detected time intervals with a low number of transactions will be at the beginning or the end of an outage with no transactions at all. In the event that there are four consecutive periods with low activity, however, we can assume that this is linked to Severity2 outages.

Figure 7: The number of single, double, triple and quadruple outage intervals.

(a) LowPA five large banks.

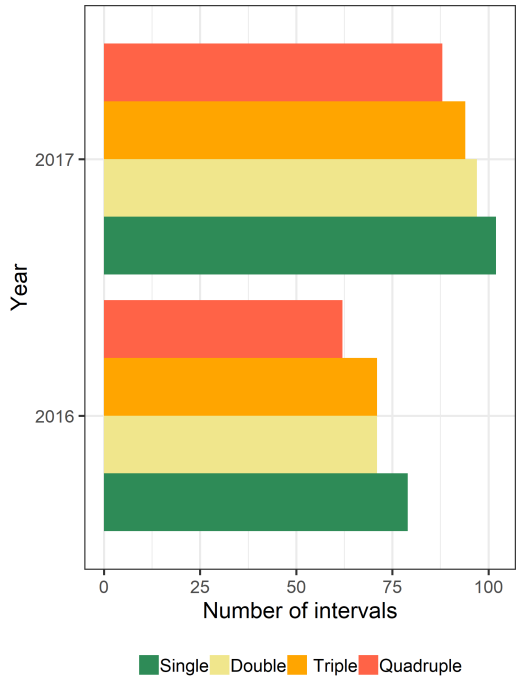

(b) NoPA five large banks.

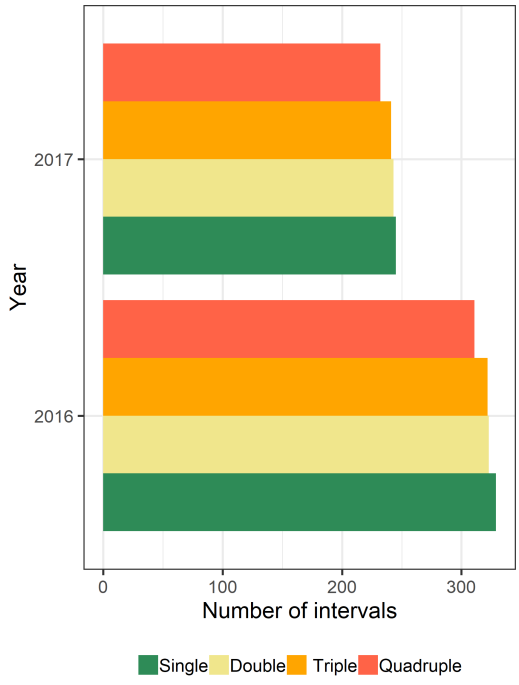

5 Conclusions

This paper develop an algorithm to detect potential operational outages at individual participants in LVTS. We distinguish two different setups of the algorithm. First, we identify periods in which no payments at all are made by a participant. A bank that is not able to send in payment instructions will become visible in the data if no payments have been observed for a certain period of time. Second, we identify periods in which there is an unusually low number of payment instructions sent in by the participant. The rationale behind this second set-up is that a bank that is unable to send in payments in the normal way may have a back up functionality for very urgent payments. The results show that a participants sends in not payments in roughly 50% of the five minute periods. Zooming in on the five largest financial institutions, this percentage decreases to less than 20%. As expected, the number of consecutive periods of five minutes decreases as their total duration increases from five minutes (one period) to 20 minutes (four periods). For the five large participants, however, this decrease is much lower than for all participants. Basically, this is because Large banks make payments during most

five minute intervals and, therefore, it is unlikely that they do not have at least some payments in two or three consecutive periods. When these large financial institutions have more than one period without payments, this is thus more likely to be due to a technical problem. And these problems may take some time to resolve (more than 20 minutes).

Most of the outages detected occur on Mondays. This is due to the fact that new software releases and updates usually take place over the weekend, when the system is closed. Problems with the new functionality and therefore experienced on the next working day, Monday. These are usually solved on the same day, though, after which participants can send in payment instructions normally again. Reports of operational outages by the participants to the operator show that our algorithm picks up these incidents. During the reported outages, for the most part the participants are unable to make any payments. In a minority of cases, however, they are still able to send in some payments.

Overall, we can conclude that our algorithm works well for the five large banks. As the small banks do not send in enough payments during the day and have too many consecutive periods with no payment activity, it will lead to too many false positives or negatives. The impact of the low activity on the part of the smaller participants is limited as the other do not expect to receive large amounts of liquidity from them.

In order to implement a "real time potential outage detection method" as a monitoring tool for the payment operator, it is crucial to have real-time data. The method presented in this paper is only able to detect potential outages, at this stage, and not predict them. As this type of algorithm will produce false positives (and negatives), a potential outage detected by it does not mean with 100% certainty that the participant concerned is actually suffering a technical outage. It may simply have no payment instructions to undertake or it may be delaying payments intentionally. It is therefore advisable to look at payment delay indicators to check whether a participant is potentially delaying its payments.

References

Abbink, K., Bosman, R., Heijmans, R., and van Winden, F. (2017). Disruptions in Large Value Payment Systems: An Experimental Approach. *International Journal of Cen-*

tral Banking, 13(4):63–95.

Arjani, N. and McVanel, D. (2006). A primer on canada's large value transfer system. *Bank of Canada*.

Bech, M. and Garratt, R. (2003). The Intraday Liquidity Management Game. *Journal of Economic Theory*, 109:198–210.

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.

Bech, M. and Hobijn, B. (2007). Technology diffusion within central banking: The case of real-time gross settlement. *International Journal of Central Banking*.

Benos, E. and, G. R. and Zimmerman, P. (2012). Bank behaviour and risks in chaps following the collapse of lehman brothers. *Bank of England Working Paper*, (451).

Berndsen, R. and Heijmans, R. (2020). Near real-time monitoring in a real-time gross settlement system: A traffic light approach. *Journal of Risk*, 22:39–64.

Clarke, A. and Hanckock, J. (2013). Payment system design and participant operational disruptions. *Journal of Financial Market Infrastructures*.

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

Diehl, M. (2013). Measuring Free Riding in Large-Value Payment Systems: The Case of TARGET2. *Journal of Financial Market Infrastructures*, 1(3):31–53.

Diehl, M. and Müller (2015). Analysis of the use and impact of limits. In Laine, T., editor, *Quantitative Analysis of Financial Market Infrastructures: Further Perspectives on Financial s Stability*, volume E:50, pages 14–42. Bank of Finland.

Glowka, M., Paulick, J., and Schultze, I. (2018). The absence of evidence and the evidence of absence: an algorithmic approach for identifying operational outages in target2. *Journal of Financial Market Infrastructures*, 6(2/3):63–91.

Heemeijer, P. and Heijmans, R. (2015). Central bank intervention in large value payment systems: An experimental approach. *Journal of Financial Market Infrastructures*, 3(3):17–49.

Heijmans, R. and Heuver, R. (2014). Is this Bank Ill? The Diagnosis of Doctor TARGET2. *Journal of Financial Market Infrastructures*, 2(3):3–36.

Kaliontzoglou, A. and Müller, A. (2015). Implementation aspects of indicators related to payments timing. In Diehl, M., Alexandrova-Kabadjova, B., Heuver, R., and Martinez-Jaramillo, S., editors, *Analyzing the Economics of Financial Market Infrastructures*, pages 169–190.

Klee, E. (2010). Operational outages and aggregate uncertainty in the federal funds market. *Journal of Banking and Finance*, 34(10):2386–2402.

Massarenti, M., Petriconi, S., and Lindner, J. (2012). Intraday Patterns and Timing of TARGET2 Interbank Payments. *Journal of Financial Market Infrastructures*, 1(2):3–24.

Triepels, R., Daniels, H., and Heijmans, R. (2018). Detection and explanation of anomalous payment behaviour in real-time gross settlement systems. In *Enterprise Information Systems. ICEIS 2017. Lecture Notes in Business Information Processing*, volume 321. Springer, Cham.

Previous DNB Working Papers in 2020

No. 662 **Carin van der Cruijsen-Knoben, Jakob de Haan and Ria Roerink**, Financial knowledge and trust in financial institutions
No. 663 **Jon Frost,** Economic Forces Driving FinTech Adoption
No. 664 **Anna Samarina and Nikos Apokoritis,** Evolution of monetary policy frameworks in the post-crisis environment
No. 665 **Christian König-Kersting, Stefan T. Trautmann and Razvan Vlahu,** Bank instability: Interbank linkages and the role of disclosure
No. 666 **Claus Brand, Gavin Goy, Wolfgang Lemke,** Natural Rate Chimera and Bond Pricing Reality
No. 667 **Joost Bats,** Corporates' dependence on banks: The impact of ECB corporate sector Purchases
No. 668 **Marc van Kralingen, Diego Garlaschelli, Karolina Scholtus and Iman van Lelyveld,** Crowded trades, market clustering, and price instability
No. 669 **Mark Mink, Rodney Ramcharan and Iman van Lelyveld**, How Banks Respond to Distress: Shifting Risks in Europe's Banking Union
No. 670 **Jasmira Wiersma, Rob Alessie, Adriaan Kalwij, Annamaria Lusardi and Maarten van Rooij,** Skating on thin ice: New evidence on financial fragility
No. 671 **Michiel Bijlsma, Carin van der Cruijsen and Nicole Jonker,** Consumer propensity to adopt PSD2 services: trust for sale?
No. 672 **Duncan van Limbergen and Robert Vermeulen**, The importance of value chains for euro area trade: a time series perspective
No. 673 **Martijn Boermans en Bram van der Kroft,** Inflated credit ratings, regulatory arbitrage and capital requirements: Do investors strategically allocate bond portfolios?
No. 674 **Andras Lengyel and Massimo Giuliodori**, Demand Shocks for Public Debt in the Eurozone
No. 675 **Raymond Chaudron, Leo de Haan and Marco Hoeberichts,** Banks' net interest margins and interest rate risk: communicating vessels?
No. 676 **Martijn Boermans and John Burger**, Global and local currency effects on euro area investment in emerging market bonds
No. 677 **Patty Duijm and Ilke van Beveren**, Product diversification as a performance boosting strategy? Drivers and impact of diversification strategies in the property-liability insurance industry
No. 678 **Richard Heuver and Ron Berndsen,** Liquidity Coverage Ratio in a Payments Network: Uncovering Contagion Paths
No. 679 **Gabriele Galati, Jan Kakes and Richhild Moessner,** Effects of credit restrictions in the Netherlands and lessons for macroprudential policy (not published yet)
No. 680 **Frank van der Horst, Jelle Miedema, Joshua Snell and Jan Theeuwes,** Banknote verification relies on vision, feel and a single second
No. 681 **Leonard Sabetti and Ronald Heijmans,** Shallow or deep? Detecting anomalous flows in the Canadian Automated Clearing and Settlement System using an autoencoder
No. 682 **Shaun Byck and Ronald Heijmans,** How much liquidity would a liquidity-saving mechanism save if a liquidity-saving mechanism could save liquidity? A simulation approach for Canada's large-value payment system

DeNederlandscheBank

EUROSYSTEEM

De Nederlandsche Bank N.V.
Postbus 98, 1000 AB Amsterdam
020 524 91 11
dnb.nl