

A liquidity black hole: What is the impact of a failing participant in a large value payment system and does time matter?



Ronald Heijmans and Ellen van der Woerd | De Nederlandsche Bank

Keywords: Financial market infrastructures, TARGET2, liquidity risk, operational risk, systemic risk, financial stability

JEL codes: E42, E50, E58, E59

Abstract

This paper presents a methodology to detect potential failing banks in large value payment systems and measure the intraday impact of outages, considering Liquidity, Systemic, and Receiver Impacts. Medium and high risk thresholds are established to create a combined risk indicator. Outages of large banks can be detected within 10 minutes, while smaller banks may take over 30 minutes. Impact and risk levels vary by the size of the bank and the start time of the outage. Large banks can reach high-risk levels in 30 minutes, highlighting the need for timely detection, whereas smaller banks rarely reach high-risk levels.

Disclaimer: This policy brief is based on research by [Heijmans and Van der Woerd \(2025\)](#). Views expressed here are those of the authors and not necessarily those of De Nederlandsche Bank.

Introduction

Financial systems depend on the smooth functioning of Financial Market Infrastructure (FMI) networks. One critical type of FMI is a Large-Value Payment System (LVPS), a system used by banks to transfer large sums, in central bank money, typically in real-time and with immediate finality. TARGET2, operated by the Eurosystem, is an example of such an LVPS. TARGET2 processes around EUR 2000 billion per day (roughly the weekly GDP of the euro area). About a thousand banks participate, but a small number of large banks account for a large share of the volume.

When one major bank in such a system fails to send payments, for example, due to an operational outage or cyberattack, it can become a “liquidity black hole”. In such cases, the liquidity that other banks were expecting from the affected institution remains trapped, preventing them from fulfilling their own payment obligations. This can trigger a chain reaction and potentially disrupt the broader financial system.

International standards, such as the Principles for Financial Market Infrastructures (PFMIs), emphasize the importance of closely monitoring participants and responding swiftly to disruptions CPSS-IOSCO (2012). However, in practice, large banks in LVPSs are typically required to report an outage only after a certain threshold, such as 30 minutes in the case of TARGET2. Unfortunately, if a critical bank halts payments, waiting this long can be too late. By then, billions of euros in expected outgoing payments can be trapped in the failing bank, placing significant stress on the bank’s counterparties and increasing systemic risk.

This policy brief introduces a new early warning and impact assessment toolkit to address this problem. Our toolkit consists of three parts:

- Near real-time detection of an operational outage of a bank in minutes, by leveraging the typical payment pattern of that bank.
- Near real-time measurement of systemic impact through intuitive indicators that track how an outage affects liquidity and other banks as time passes.
- Clear communication of risk through a simple traffic light signal (green, orange, and red) to inform the system operator when an incident is minor, serious, or critical.

The goal is to help payment system operators act quickly and effectively, preventing a localized technical failure from snowballing into a system-wide crisis. As an example, we apply this toolkit to the 2022 TARGET2 data.¹

The remainder of this note contains three parts. First, we describe the method and results. Second, we explain how the operator could use the early warning dashboard. Lastly, we conclude and provide policy recommendations.

Near real-time detection of an outage

Each bank tends to have a characteristic timing and frequency of sending payments. Some large banks send transactions every few minutes or even seconds, whereas smaller banks may exhibit longer intervals between payments. We use these characteristic payment patterns to detect outages near real-time. We define the **Minimum Outage Time Interval (MOTI)** for each bank as the longest period that it usually lasts without sending a payment during a given hour, using a methodology based on historical data. If a bank exceeds this interval without payments, it

¹ However, this method can be used in any large value payment system that handle payments individually (gross). In addition, our method can also be used in future versions of TARGET2.

may indicate an operational problem.² Identifying such deviations early can help reduce the impact of disruptions by enabling faster responses.

For example, if Bank A almost never goes more than 10 minutes without a payment around mid-morning, then 10 minutes is its MOTI threshold for that time. If Bank A does not initiate payments for a longer period of time, an automatic alert is triggered. By customizing the MOTI for every bank individually (and for different times of day), we can catch outages of active banks within minutes, while avoiding false alarms for less active banks.

We find that banks with the shortest MOTI, between 1 and 10 minutes, tend to be high volume banks that process a large number of payments each day. In contrast, banks with longer MOTIs tend to process fewer payments. In addition, we observe that the MOTI varies during the day for large banks. More specifically, MOTI tends to increase during the first two hours after the system is opened and during the last hour of the day. In contrast, banks with lower activity levels show a more consistent MOTI throughout the day.

Determining the impact of an outage

Detecting an outage is only the first step. We also need to understand its impact on the system in (near)real-time. We introduce three metrics:

- **Liquidity Impact (LI):** the total value of the payments the affected bank has not sent due to the outage. This indicates how much liquidity “is missing from the system” compared to normal.
- **Systemic Impact (SI):** the number of other banks that have not received payments they normally would have from the affected bank. This shows how widely the outage spreads through the network.
- **Receiver Impact (RI):** the number of banks put in a vulnerable position because they miss a large portion of the expected funds from the affected bank. If a bank relies heavily on payments from the failed participant, it could struggle to meet its own obligations. RI counts how many banks face that situation.

These metrics start at zero when the outage begins and grow throughout the duration of the outage. We combine them into a single Combined Risk Indicator (CRI) with a traffic light interpretation:

- **Green (Low risk):** The impact is minimal (all metrics are low).
- **Orange (Medium risk):** The outage is causing noticeable strain (two or more metrics at a medium level or at least one at a high level).
- **Red (High risk):** The situation is critical (multiple metrics at high levels, indicating a potential systemic crisis).

The CRI gives operators an immediate sense of severity. For example, if a large bank experiences an outage and after 20 minutes has missed payments worth, say, EUR 5 billion affecting 80 other banks (with a couple of those banks in trouble), the CRI would likely signal orange and trend toward red, indicating that urgent intervention may be needed.

We use these metrics to simulate the outages of different banks in TARGET2. We find the following results: Outages of major banks escalate quickly and affect many others. For a large bank (one that processes thousands of payments a day), even a short outage can withhold a large amount of liquidity and impact a wide network of counterparties. Our analysis shows that if one of the largest banks experiences an outage at the start of the day, within 30 minutes it could withhold *several billion euros* of payments (high LI) and leave more than a hundred other banks waiting for incoming liquidity (high SI). Some of these banks would soon find themselves missing a significant share of their liquidity, risking

² It is also possible in exceptional circumstances that a bank exceeds this limit for other reasons. Using historical data we can determine that this does not happen more than a couple of times a year.

a domino effect (RI rising). In contrast, if a smaller bank (say one with a few hundred payments a day) has an outage, the impact builds much more slowly and often stays low.

Outages early in the day are far more dangerous than late-day outages. Timing makes a big difference. If a bank fails in the morning, it misses the whole day's payments; if it fails late in the afternoon, most of its important payments are likely to be made earlier in the day. In our simulations, a large bank outage starting at 9:00 AM could breach high-risk levels by noon, while the same outage starting at 3:00 PM might only reach medium-risk by the end of the day (since fewer payments remain unsent).

The MOTI early-warning system buys valuable time. In almost every severe case that we examined, the bank's MOTI alert would ring before the situation became critical. For example, a major bank might trigger an alert after 5–10 minutes of no payments. Without MOTI alert system, operators might only notice the outage after 30 minutes when a manual report arrives; by then, the liquidity shortfall and network effects could already be at dangerous levels. Early detection allows operators to initiate contingency plans or assist affected banks early, greatly increasing the chances of containing the problem.

Operational toolkit: Near real-time monitoring in practice

A key innovation of our approach is to provide actionable information to operators and management through a visual dashboard. Once a MOTI-based outage alert is triggered for a bank, the dashboard can display the status of that bank in terms of the different metrics LI, SI and RI and the CRI level, in near real-time.

Figure 1 illustrates the concept. The dashboard for an affected bank would show, for example:


- A time-line of the CRI (with segments colored green, orange, red as risk rises).
- The growing total value of missed payments (LI over time).
- The count of banks affected (SI over time).
- The count of banks in trouble due to missing liquidity (RI over time).

When the alarm is triggered the operator can see what the expected evolution is and when a medium or high risk will be triggered, and what the main underlying cause is, e.g., the euros not sent or the number of banks affected. This concrete information helps in deciding the next steps. First, an operator would call the bank in case an alert is raised to check whether there is an operational problem. Follow-up steps could be the issue of an alert to the market, providing emergency liquidity to certain banks, or extending operating hours to clear backlogs.

We suggest implementing our method in a large value payment system, as such a tool transforms system monitoring from reactive to proactive. Instead of waiting and guessing, operators have a clear early signal and a quantified view of stress, allowing timely intervention. It also creates a common language for the response to a crisis: for example, management could establish a rule that “if the CRI goes red, convene the crisis team immediately”.

Furthermore, capturing these metrics for each incident helps in after-action reviews and improvements. Operators can identify which banks repeatedly appear as highly impacted receivers, examine whether MOTI thresholds need adjustment, and refine risk thresholds as the system evolves.

Figure 1.



Schematic of the near real-time risk monitoring dashboard for an outage. As soon as a bank is detected to be inactive beyond its MOTI, the dashboard tracks the LI, SI and RI. The combined indicator (CRI) provides an overall risk level (green = low, orange = medium, red = high) to guide operators. In this illustrative example, the bank experiences an outage at 7:00 AM leading to a medium just before 1:00 PM, evolving into high risk at the end of the day.

Conclusion and policy recommendations

Operational failures in payment systems are inevitable, but not their systemic consequence. The insights and tools of this study show that with the right monitoring, a potentially catastrophic liquidity black hole can be identified and mitigated in near real-time.

Policy makers, overseers, and system operators should consider the following actions:

- **Adopt early-warning indicators:** Implement bank-specific no-payment timers (like MOTI) in large-value payment systems to catch outages within minutes. Rapid detection is the first line of defense.
- **Near real-time dashboards:** Use combined risk indicators (CRI) to monitor and communicate the severity of outages. A traffic light-style dashboard can streamline decision-making and escalation during incidents.
- **Focus on critical participants:** Ensure that the largest banks (those whose failure would have systemic impact) maintain robust operational continuity plans. Overseers might impose higher resilience standards on these banks and conduct regular outage drills.
- **Plan for peak risk periods:** Recognize that outages at the start of the day pose the greatest risk. Be especially vigilant during opening hours. If a major outage occurs early, deploy contingency measures (like liquidity support or payment rescheduling) promptly.
- **Continuously update the framework:** Regularly recalibrate MOTI thresholds and risk limits using fresh data, and learn from each incident. Payment patterns change over time; the monitoring system should evolve with them to remain effective.

By taking these steps, financial authorities can significantly reduce the chance that a single technical failure triggers a wider financial panic. In essence, we are equipping the “central plumbing” of finance with better sensors and circuit breakers. This strengthens trust in the system: even if one pipe causes a leak, we can quickly contain it before the whole network is flooded.

References

CPSS-IOSCO (2012). Principles for financial market infrastructures, bank of international settlements.

Heijmans, R. and Van der Woerd, E. (2025). A liquidity black hole: What is the impact of a failing participant in a large value payment system and does time matter? *DNB Working Paper*, 836.

About the author(s)

Ronald Heijmans earned his PhD in Economics from the University of Groningen, with a thesis focused on large value payment systems. He previously obtained an MSc in astrophysics from Leiden University. He has over 19 years of experience at DNB, where he works for the Market Infrastructure and Innovation department as a senior policy advisor and data scientist. He has made significant contributions to the literature on payment systems and financial market infrastructure, covering topics such as machine learning, risk indicators, liquidity, network topology, and game theory on TARGET2, and crypto and stablecoin data.

Ellen van der Woerd earned her PhD in theoretical physics from the KU Leuven, with a focus on string theory and cosmology. She has six years of experience at DNB, where she works for the Market Infrastructure and Innovation department as a senior data scientist. Her main areas of expertise are the analysis of large financial data sets such as EMIR, SFTR, and TARGET. Using these data, she reduces complex problems into valuable insights for research and to support data-driven policy making.

SUERF Policy Notes and Briefs disseminate SUERF Members' economic research, policy-oriented analyses, and views. They analyze relevant developments, address challenges and propose solutions to current monetary, financial and macroeconomic themes. The style is analytical yet non-technical, facilitating interaction and the exchange of ideas between researchers, policy makers and financial practitioners.

SUERF Policy Notes and Briefs are accessible to the public free of charge at <https://www.suerf.org/publications/suerf-policy-notes-and-briefs/>.

The views expressed are those of the authors and not necessarily those of the institutions the authors are affiliated with.

© SUERF – The European Money and Finance Forum. Reproduction or translation for educational and non-commercial purposes is permitted provided that the source is acknowledged.

Editorial Board: Ernest Gnan, David T. Llewellyn, Donato Masciandaro, Natacha Valla

Designed by the Information Management and Services Division of the Oesterreichische Nationalbank (OeNB)

SUERF Secretariat

c/o OeNB, Otto-Wagner-Platz 3A-1090 Vienna, Austria

Phone: +43 1 40 420 7206

E-Mail: suerf@oenb.at

Website: <https://www.suerf.org/>