JULY 30, 2024

# Protecting Distributed Financial Networks

By Corey Lofdahl (Leidos), Amanda Gentzel (Leidos), Chris Wheeler (Mastercard), Michael Dewar (Mastercard), Rafael Alonso (Alonso Associates)

*Distributed financial networks are a feature of the international financial system of payments. Still, they are also increasingly vulnerable to disruption as new technologies create unexpected opportunities for surprises, threats, and shocks. These vulnerabilities arise due to current economic and technical trends, including the increasing velocity and digitalization of individual economic activity, as well as the growing interconnectedness of the global economy. In this brief, we discuss these financial system challenges through the lens of a credit card payment system. We present a range of integrated tools and procedures tailored to meet the needs of the financial firm, network, and system as no single "silver bullet" solution exists. Instead, protecting networks requires multiple, integrated solutions that work together to reduce system fraud and errors.*

## INTRODUCTION

Distributed financial networks are a feature of the international financial system of payments. Still, they are also increasingly vulnerable to disruption as new technologies create unexpected opportunities for surprises, threats, and shocks. These vulnerabilities arise due to current economic and technical trends, including the increasing velocity and digitalization of individual economic activity as well as the growing interconnectedness of the global economy. Due to these expanding economic relationships and evolving forms of financial intermediation, new challenges emerge over addressing vulnerabilities such as jurisdictional regulatory differences, cyber-attacks, and transnational fraud. These factors have taxed financial and regulatory organizations' ability to ensure consistency across distributed financial systems and maintain their stability due to potential losses of both data integrity and user confidence.

Addressing these vulnerabilities at the institutional and transaction level is challenging, but so is identifying and recovering from coordinated systemic shocks across multiple institutions. Financial institutions typically ensure their system resiliency through individual backups, [12] but when securing a distributed Financial Market Infrastructure (FMI), similar practices require far more coordination, since data storage and their backup are uncoordinated across institutions [3].

This brief addresses the challenges presented by distributed FMIs in four parts. First, we investigate the processing of consumer payment transactions as an example to determine industry best practices. This is done using anonymized Mastercard payment

Note: This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

data [7] to examine transaction relationships among firms rather than at a specific firm. Second, we look at the networks that result from these transactions among firms, which allows us to explore an empirical example, albeit limited, of a distributed financial network.

Third, we review ways to identify significant changes to the distributed system. The data can reveal multiple insights, but doing so requires analyzing and synthesizing many transactions. We demonstrate a model that can perform this automatically. Fourth, we look at institutional processes that protect the financial system when problems occur. To maintain the resilience of financial systems, informed actions must be undertaken to address identified problems. Using Mastercard credit card payments, we focus on transaction fraud. Using the demonstrated detection models, the number of problems and system features can be expanded to include other problematic dynamics including system attacks. The ability to detect and address financial systems threats beyond transaction fraud can be expanded to ensure financial network integrity and maintain user trust.

## PROCESSING PAYMENT TRANSACTIONS

Mastercard credit card data proves an excellent example of a distributed financial network because it provides worldwide transaction coverage to more than 210 countries [8]. The credit card transaction data we received from Mastercard was comprised of two transaction components, authorization and clearing, as shown in **Figure 1**. When a customer makes a purchase at a merchant, the customer initiates a transaction authorized through a call from the acquiring (merchant's) bank to the issuing (customer's) bank. This authorization allows the customer to make the purchase with the understanding that the item will be paid for later through the clearing process in which the balance between the issuing and acquiring banks is settled.

The transaction data is anonymized[1] from customer information and contains information on authorizing and clearing transaction records. Specifically, the data comes from Canadian Mastercard transactions and are for March 2020 and 2021 representing approximately 500 thousand transactions per day.[2]

We generate a set of consistency metrics to assess if a problem exists between the issuing and acquiring banks in the distributed financial network. For example, high-level balances between banks are calculated from each bank's perspective, and these are then checked against each other periodically in a periodically generated consistency report. Significant discrepancies then provide the impetus for more detailed analysis and investigation into the source of the discrepancy, which contributes to the confidence that the financial network is functioning properly. In this case, calculation and comparison of these reported balances did not yield any significant discrepancies,

Figure 1. Mastercard Use Case

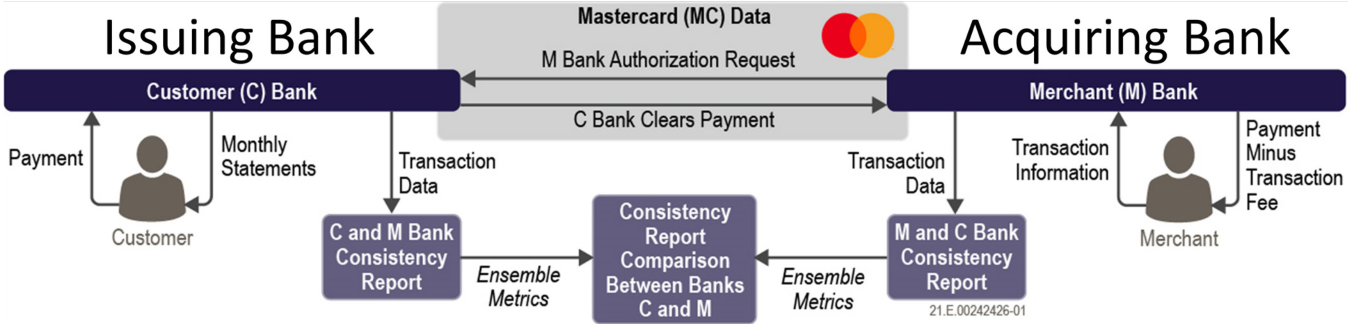

Note: Mastercard use case shows an example of authorization and clearing for a transaction between the customer/issuing (C) and merchant/acquiring (M) banks.

Source: Authors' creation

which was expected as Mastercard spends considerable effort ensuring that bank balances are accurate.

We calculated six variables to support the consistency analysis: (1) total transactions; (2) total cleared transactions; (3) total incomplete transactions; (4) total corrupt transactions; (5) transaction minimum and maximum; (6) transaction expected value. Credit card transaction types are categorized by Merchant Category Codes (MCCs) [4], of which there are approximately 500. Four code categories were selected given their prominence and potential vulnerability: (1) restaurants; (2) airfare; (3) gambling; (4) crypto/quasi-cash, which are shown in **Table 1**.

Table 1. Canadian Transactions by Merchant Category Codes

| Merchant Category Code (MCC) | Transactions Per Day | Mean Transaction Value |
|---|---|---|
| Restaurants | 500,000 | $50 CAD |
| Airfare | 30,000 | $500 CAD |
| Gambling | 3,000 | $70 CAD |
| Crypto/quasi-cash | 300 | $1000 CAD |

Note: Canadian transactions characterized broken out by Merchant Category Codes (MCCs) to show significant differences among Transactions per day and Mean transaction value.

Source: Authors' analysis [7]

**Table 1** shows that most credit card transactions are restaurant transactions with the number reducing by an order of magnitude from airlines to gambling and crypto / quasi-cash. However, in terms of transaction value, the airline transactions are about an order of magnitude greater than restaurants, and crypto / quasi-cash is significantly higher than the other three categories. This shows that MCC sub-categories can have very different characteristics, reflecting each category's particular dynamics. In terms of raw numbers, over 90% of transactions are restaurant transactions, which tend to have lower costs and often exhibit a difference between clearing and authorizing amounts due to tips. These look very different from crypto transactions, which are much lower volume but significantly higher value.

**Table 2** shows a subset of transactions broken out by MCC and categorized as *cleared, incomplete,* and *corrupt*. These categories are determined by transaction response codes, which are defined by the International Organization of Standards (ISO) [9]. Cleared transactions have a response code denoting, "Approved or completed." Such transactions have a corresponding clearing transaction that normally occurs within 48 hours of the authorization. However, we observed that restaurants exhibit more delay than the other MCCs, sometimes appearing more than a week after the initial authorization. Incomplete transactions, of which there are approximately 50, denote various transaction status conditions and problems. For example, Incomplete Transactions are those that are blocked because of a simple problem that the cardholder can easily correct, such as "insufficient funds, over credit limit," "transaction not permitted to issuer/cardholder," and "not declined, valid for zero amount transactions" ( i.e., transaction performed not for a purchase but to determine if the card is active). A subset of the Incomplete Transaction types prevents potentially Corrupt Transactions such

Table 2. Canadian Transactions for March 1, 2020

| Merchant Category Code (MCC) | Cleared Transactions | Incomplete Transactions | Corrupt Transactions |
|---|---|---|---|
| Total | 98.0% | 1.9% | 1.4% |
| Restaurants | 98.9% | 1.1% | 0.7% |
| Airfare | 91.7% | 8.1% | 7.0% |
| Gambling | 83.0% | 16.1% | 11.8% |
| Crypto/quasi-cash | 77.8% | 20.7% | 16.6% |

Note: Canadian transactions for March 1, 2020, characterized by transaction response codes [9] and categorized as cleared, incomplete, and corrupt.

Source: Authors' analysis [7]

as “capture card,” “do not honor,” and “invalid card number.”

The transaction analysis in **Table 2** is over a single day (March 1, 2020) and constrained geography (Canada). More extensive analysis, both temporally and geographically, would provide additional results that could be analyzed and tracked over time. Still this limited MCC and response code-based analyses gives a sense of the character and richness of credit card transaction data.

## EXPLORING PAYMENT NETWORKS

Credit card transactions combine to create network relationships among acquiring and issuing banks. Looking beyond the set of bilateral bank relationships allows us to consider how the larger payments system creates the distributed financial network. To give a sense of the size and scale of the Mastercard network of banks in Canada, there are 42 acquiring banks and 96 issuing banks in the dataset. Of the 4,032 possible combinations of issuing and acquiring bank transactions, we find 1,146 (28%) have transactional relationships.

Note that these relationships reveal significant clustering. That is, there are few banks with large numbers of relationships and many banks with few relationships, which is described as a *power law relationship* and is common in dynamic networks [1]. For example, **Figure 2** shows that acquiring banks are highly centralized, especially with 186 at the upper left, while the issuing banks are more evenly distributed. Note also bank 389 at the bottom center is both an acquiring and issuing bank, which is another form of centralization. Note that these graphics are created from one day of transaction data, so further analysis is required to determine if and how these relationships evolve over time. This visualization shows a dense network of edges, each representing a relationship between two financial institutions, which, in aggregate, compose a highly complex and interconnected financial network, highlighting the need for protection, not just of the firms themselves, but of the system as a whole.

Figure 2. Sankey Diagram Showing Transaction Flow

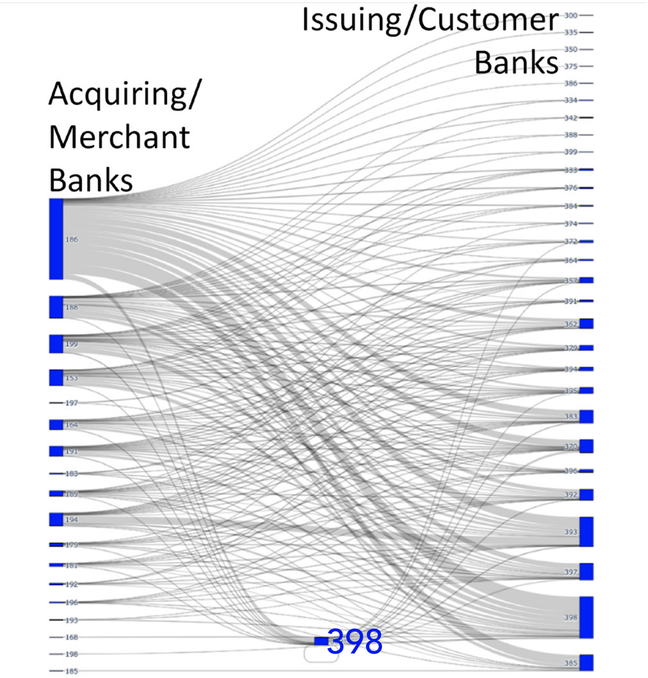

Note: Sankey diagram showing transaction flow from acquiring/merchant (left) to issuing/customer (right) bank, where width of connection indicates the number of transactions between bank pairs, filtered to top 25% of relationships by volume. Note bank 389 (bottom center) operates as both an authorizing and issuing bank.

Source: Authors’ analysis [7]

## IDENTIFYING SYSTEMIC CHANGES

The collected transaction data can be used to identify shifts and changes in the underlying financial system. For example, we find two significant systemic changes when we compare March 2020 to March 2021 as shown in **Figure 3**. First, daily transaction counts for March 2020, shown in orange, reveal a significant decrease in transaction volume as measured by the number of authorizations in the second half of the month. Second, for March 2021, the transaction volume was significantly decreased compared to the previous year, March 2020, but was relatively stable across the month. This leads to two questions: first, what caused these changes, and second, can such changes be identified automatically?

Before offering a possible answer to the first question, note that this is typically the kind of puzzle that confronts those who notice unusual trends in transaction data. The change may be easily noticed, but identifying and understanding the reasons underlying it requires additional research and access to other data sources. Addressing the first question, we postulate that Canada's COVID-19 quarantine and distancing policies added in mid-March 2020 resulted in a measurable decrease in credit card transactions. Measurements taken the following year, in March 2021, reveal a relatively diminished number of transactions due to the ongoing COVID control measures. Still they are comparatively consistent within the month (apart from a clear weekly cycle), indicating a long-term change from the earlier systemic shock.

Answering the second question—Can these changes be identified automatically?—our analysis of **Figure 3** required a significant level of human-directed computation, visualization, and analysis to identify the likely COVID-19 impacts, which is not feasible for operational datasets at scale. Our approach to automatically identifying inconsistencies is based on learning causal models of normal behavior from the Mastercard data and using these models to infer inconsistencies. Specifically, we determined a set of transaction features that provide information regarding normal transaction behavior. This consists of a combination of straightforward transaction information (e.g., the transaction amount and the merchant category) and more complex features (e.g., the average historic transaction amount for this merchant). Using causal structure learning algorithms, we can then use these features to learn causal models of transaction behavior.

Figure 3. Authorization Counts for March 2020 and 2021 Compared

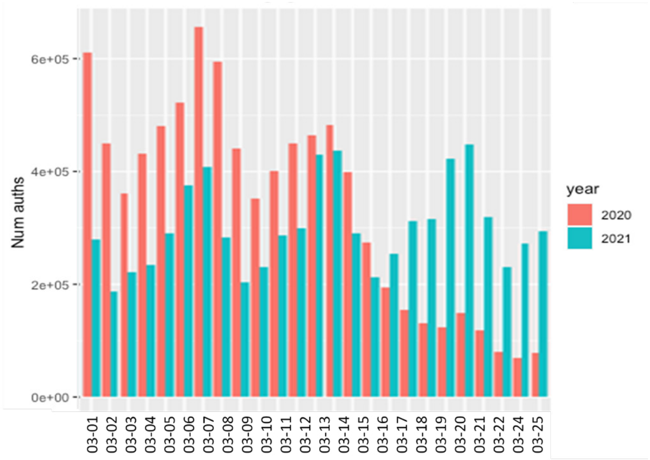

Note: March 2020 transaction volume decreases significantly in the second half due to COVID sanctions impacting economic activity, while March 2021 transactions remain comparatively stable.

Source: Authors' analysis [7]

Causal models in this setting consist of Bayesian networks, where an edge from variable A to variable B is interpreted as "A causes B" [10, 13]. The network represents the joint probability distribution over the transactions in an interpretable way. Such models are typically constructed in two main steps—structure learning and parameter fitting—with structure learning algorithms falling broadly into score-based and constraint-based categories. Score-based methods attempt to find the network structure that optimizes a score function (such as the Bayesian Information Criterion, or BIC). In contrast, constraint-based methods perform a series of conditional independence tests on the data, adjusting the edges in the network to reflect the results of those tests. Once a structure is learned, a conditional probability distribution is fitted at each node, denoting the probability distribution of each variable conditioned on the values of its direct causes—that is, the variables that point to it—in the network. The product of these conditional probability distributions composes a joint distribution over the data.

**Figure 4** provides an example of a causal structure learned using Mastercard data from the first week of March 2020. The network structure can provide insight into the behavior of transactions, and the probability distributions can be used to detect significant system anomalies. For example, the edge "crypto -> term attend" suggests that the probability distribution of a card terminal being human-attended differs for crypto and non-crypto transactions, and the edge "S. Amt -> B. Amt" suggests that the transaction settlement amount affects the billing amount. Some background information of temporal precedence was provided to the algorithm in the form of logical rules that cannot be violated as an *edge blacklist* (e.g., 'mean historic merchant transaction amount' cannot occur after 'transaction amount'), and, when an undirected edge (i.e., an arrow without a causal interpretation)

Figure 4. Causal structure of transactions during the first half of March 2020[3]

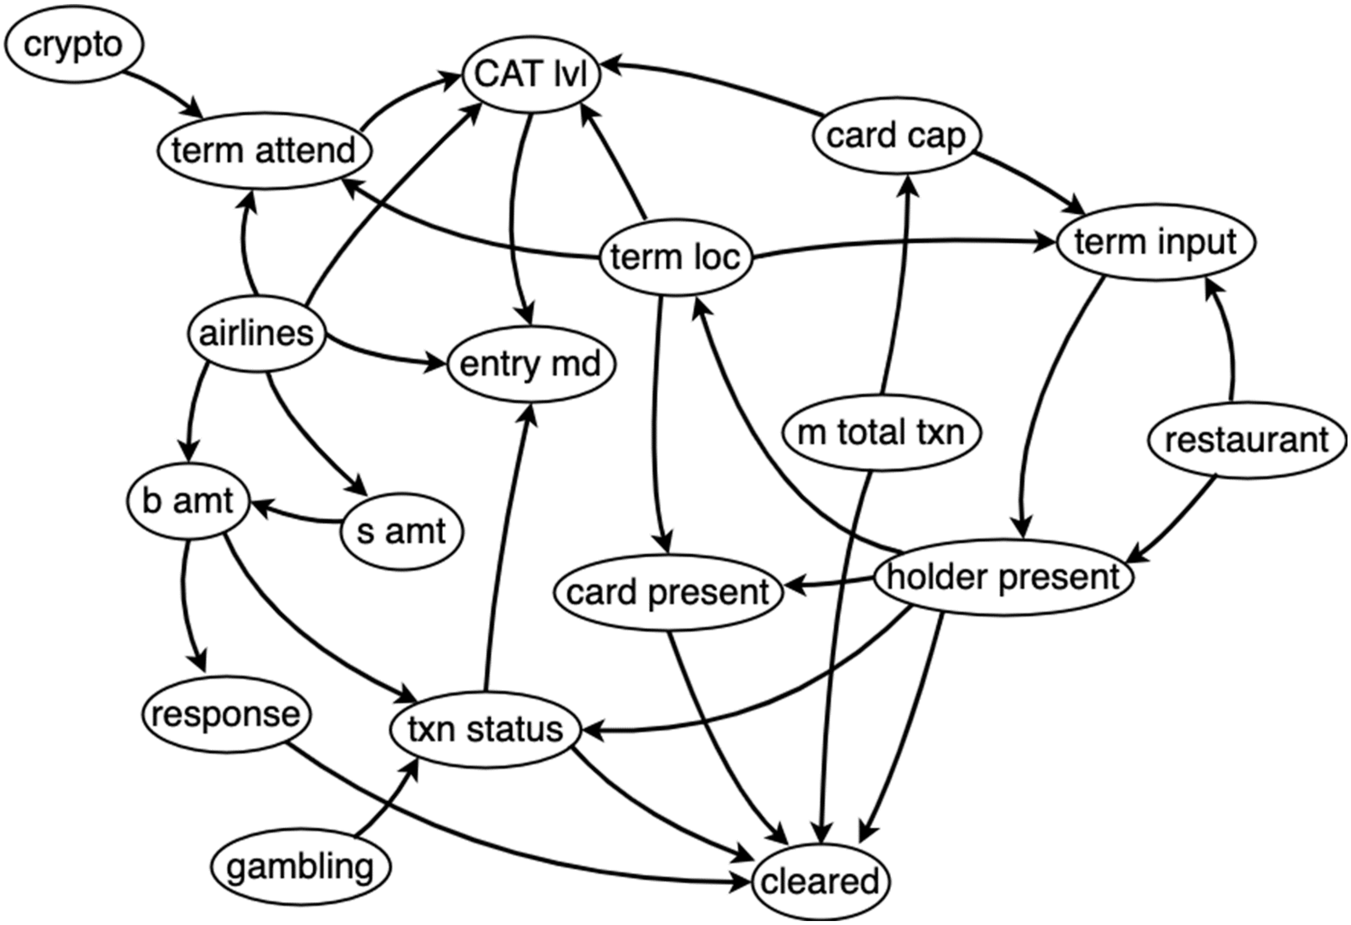

Note: A causal graph with nodes representing features of a transaction (as described in footnote 3) and directed edges representing the causal dependence structure among those features.

Source: Authors' analysis [7]

was returned by the structure learning algorithm, it was modified by hand using knowledge about the problem domain.

These models can be used to detect systemic shifts in the transaction population by calculating the joint probability of the features of any individual transaction. By calculating this on a sample of transactions, we get a distribution of transaction probabilities. Assuming the transactions in the sample are drawn from the same underlying distribution as the transactions in the training data, we expect the distribution of transaction probabilities to be static over time. If the distribution starts to deviate, this suggests a systemic change in how transactions function, which may indicate a large-scale event. As an initial test of this, we trained separate models on the first week of March 2020 and the first week of March 2021. We then sampled transactions from days throughout March 2020 and 2021 and calculated their probabilities according to their respective models. As we can see in **Figure 5**, March 2020 sees a steep drop in probability—as calculated by multiplying the conditional probabilities of the observed values for each variable given the observed values of its parents in the network—that persists in the second half of March, which is consistent with the large-scale changes brought about by COVID lockdowns at that time.

**Figure 6** shows a return to a more consistent transaction pattern in March 2021, like in the first half of March 2020. Because these joint probabilities are

Figure 5. Joint Probability of a Sample of Transactions in March 2020

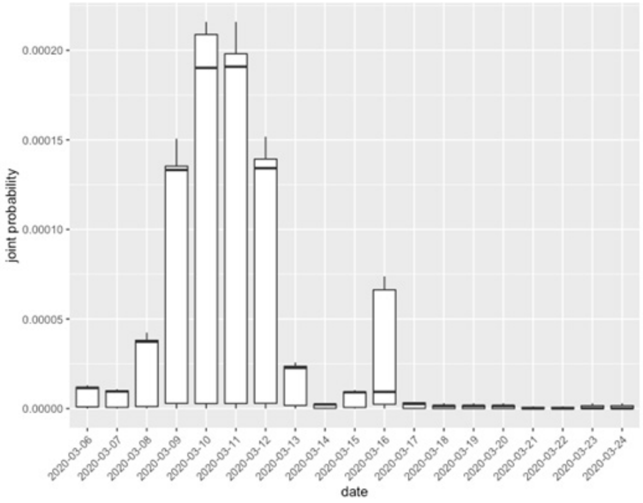

Note: Joint probability of a sample of transactions based on a model learned on data from March 6-12, 2020, which shows a clear reduction in transaction volume after March 14.

Source: Authors' analysis [7]

Figure 6. Joint Probability of a Sample of Transactions in March 2021

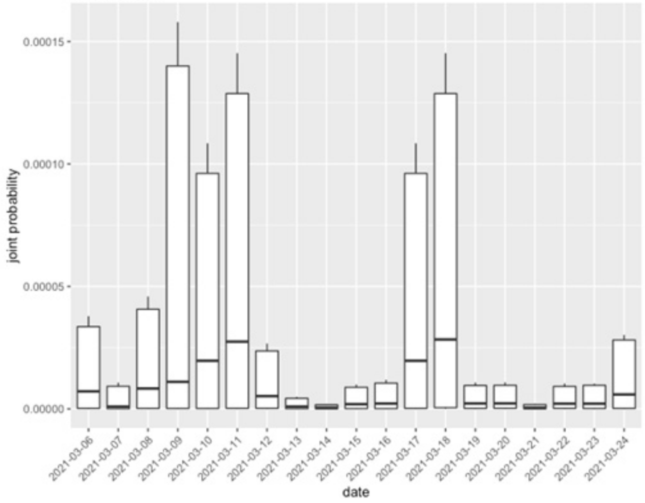

Note: Joint probability of a sample of transactions based on a model learned on data from March 6-12, 2021, which shows a comparatively consistent transaction pattern during this period.

Source: Authors' analysis [7]

composed of conditional probabilities at each node, we can further investigate which variables experienced a probability drop, further helping guide the investigation to identify potential causes of a systemic change.

**Figures 5 and 6** were created using transactions from all MCCs during the training period. However, models can be trained over any subset of the data, providing more focused and tailored insights. For example, a model could be trained on transactions from only a single issuing bank, a particular industry or economic sector, or a selected country or region over a focused timeframe. As shown above, we can learn a model of typical transaction behavior for that use case and calculate the distribution of transaction probabilities over time. A shift in these probabilities would signify a change in behavior for that focused subsystem, rather than the entire financial system. We can apply such analysis to any subset of the data for which we have sufficient transactions. While the model's causal structure alone can be informative, the model also contains a set of conditional probabilities that can be used to detect anomalies—automatically identify system changes. Such models encode what behavior is typical for payment transactions, and this can be used to identify automatically when transactions shift away from that behavior.

## PROTECTING FINANCIAL NETWORKS

Computational tools and processes provide vital protection to financial networks and systems. However, by themselves, they are insufficient as personnel must be trained to know when and how to use them through preparation, coordination, and practice [11]. For example, Mastercard employs a Safety Net system that automatically applies rule-based models to detect transaction fraud [6]. Given the importance of fraud detection to a credit card's business, this capability is fundamental to supporting the integrity of the authorization process. Mastercard emphasizes fraud detection during the authorization step to avoid the cost of unwinding cleared transactions through the *chargeback process*, which can be significant. Credit card companies therefore err on the side of caution with the understanding that transactions denied during authorization can always be re-initiated. These types of checks can be expanded using the types of causal models described in Section 3 to address other types of relationships revealed by consumer transactions including transnational fraud [2]; cyber-attacks [3,5]; and money laundering as well as other financial crimes [14].

While causal models can be applied to identify a range of systemic threats and dynamics, several challenges exist when considering applications to transnational fraud. First, while the Mastercard payments dataset spans multiple regulatory regimes (210 countries and territories [8]) that allow for transnational dynamics to be detected, the volume is quite high – on the order of 60k transactions/sec – so automatic detection models must be computationally efficient. Prioritizing model application based on risk—a function of problem frequency and impact—is therefore required [3].

Second, the transaction environment can be counter-intuitive and hard to interpret, as what appears at first glance to be problematic often is not, and what seems normal can be problematic. This complicates detection and causal identification, so care must be given to the human interactions that support these systems. New models will need to be added and others removed, requiring a model testing process to ensure that they perform effectively and as expected.

Third, the transaction payments system itself, rooted in human behavior, is dynamic, changing and evolving over time. Therefore, identifying and addressing problems such as fraud doesn't fully eliminate them so much as cause the system to change and evolve. This presents an opportunity to apply artificial intelligence to create models that can learn, adapt, and update to track the new "normal." Additional human analysis, processes, checks, and institutional infrastructure will always be required to identify and address newly emergent transaction problems, introducing additional delays, costs, and complexities into automated detection systems. However, automatic models provide an important capability to help identify financial network changes and threats.

## CONCLUSION

In conclusion, a range of capabilities are available to improve the resilience and performance of distributed financial networks. This study demonstrates several features of a complete solution. First, transactions among firms, like the credit card transactions captured in the Mastercard data [7], produce complex networks. Data backups at the firm level provide an important resilience capability, but consideration must also be given to the network structure and the relationships among banks. We provide a set of metrics to check the balances among banks and identify problems early and efficiently.

Second, we find Bayesian network models can automatically recognize changes in the financial network. Moreover, the causal relationships that comprise these models can be queried to provide additional information about how the system has changed. While such models are used to explore fraud identification in our setting, they can also be adapted and focused to address various other system changes and threats.

Third these models can be used in an operational system to address threats quickly when time is of the essence. The detection and response to financial network threats require ongoing diligence to ensure personnel are trained to respond effectively when threats are identified. Addressing them requires a range of integrated tools and procedures tailored to meet the needs of the financial firm, network, and system as no single, "silver bullet" solution exists. Instead, improving systems and protecting networks requires multiple, integrated solutions that reduce errors and address problems from different perspectives.

## REFERENCES

1. Bak, Per, 2013. *How Nature Works: The science of self-organized criticality.* Springer Science & Business Media.

2. Conley, Heather A., James A. Lewis, Eugenia Lostri, and Donatienne Ruy, 2022. “Strategic Competition in the Financial Gray Zone.” White Paper, Center for Strategic and International Studies (CSIS), Washington, DC (April).

3. DTCC, 2021 September. “Cyber Threats and Data Recovery Challenges for Financial Market Infrastructures (FMIs).” White Paper, Depository Trust & Clearing Corporation (DTCC), Jersey City, NJ.

4. Mastercard, 2018. “Quick Reference Booklet—Merchant Edition.” Technical Report, Merchant Category Code (MCC) list, Mastercard, Purchase, NY (Nov. 15).

5. Mastercard, 2020. “Mastercard Launches AI-Powered Solution to Protect the Digital Ecosystem.” Press Release, Mastercard, Purchase, NY.

6. Mastercard, 2021. “Securing Trust in Central Bank Digital Currencies.” White Paper, Mastercard, Purchase, NY.

7. Mastercard, 2022. “Anonymized Canadian Authorization and Clearing Transaction Data: March 2020 and 2021.” Mastercard, Purchase, NY. Provided as part of DARPA Ensuring Consistency of Systemic Information (ECoSystemic) Artificial Intelligence Exploration (AIE) Leidos Distributed-database Analysis and Consistency-Checking System (DACCS) Project, Agreement No. HR00112290007.

8. Mastercard, 2024. “Who We Are.” Website, Mastercard, Purchase, NY. Accessed at https://www.mastercard.us/en-us/vision/who-we-are.html on Feb 24.

9. Mastercard, 2024. “Network Response Codes.” ISO Codes from Website, Mastercard, Purchase, NY. Assessed at https://developer.mastercard.com/mastercard-send-funding/documentation/response-error-codes/network-response-codes/ on Mar. 24.

10. Pearl, J., 2009. *Causality.* Cambridge University Press.

11. Raskin, Sarah. 2016. “Remarks by Deputy Secretary Sarah Bloom Raskin at the Cybersecurity Docket’s Incident Response Forum 2016.” Archived Content, U.S. Department of Treasury, Washington, DC (March 31).

12. Sheltered Harbor 2023. “Operating Rules.” Technical Report, Sheltered Harbor, New York (Jan.).

13. Spirtes, P., Glymour, C.N. and Scheines, R., 2000. *Causation, Prediction, and Search.* Cambridge, MA: MIT Press.

14. Vocalink, 2017. “The Rise of the Mule.” White Paper, Vocalink (proprietary to Mastercard), London.

# ENDNOTES

1 The data we worked with was anonymized through a *hashing* process, in which bank and Personally Identifiable Information (PII) are transformed into arbitrary numerical values, thereby ensuring privacy.

2 As Canada's GDP is about one-tenth the size of the US (1.8T PPP (2020) vs. 22T PPP (2019) respectively), US transaction data is anticipated to be an order of magnitude larger.

3 The following variables are used to construct the Figure 4 causal model: b amt – billing amount; s amt – settlement amount; m total txn – total transaction amount for this merchant in the training period (proxy for merchant size); txn status – transaction status or purpose (e.g., normal request, account status inquiry, preauthorized request); cleared – whether or not the transaction cleared within a week; term attend – is the terminal used for the transaction human-attended?; term input – input method for the terminal; entry md – terminal entry mode (PIN entry capability); term loc – terminal location (on or off premise, or no terminal used); card cap – does the terminal used have card capture capabilities?; holder present – is the cardholder present for the transaction? (if not, order method, such as phone or electronic); card present – is the card present for the transaction?; response – response code (approval or decline with reason); CAT lvl – if relevant, Cardholder Activated Terminal (CAT) security level of transaction; [crypto, gambling, restaurants, airlines] – merchant category for this transaction.