Fraud Risk Mitigation in Real-Time Payments: A Strategic Agent-Based Analysis

Katherine Mayo , **Nicholas Grabill** and **Michael P. Wellman**
University of Michigan
{kamayo, grabilln, wellman}@umich.edu

Abstract

Whereas standard financial mechanisms for payment may take days to finalize, *real-time payments* (RTPs) provide immediate processing and final receipt of funds. The speed of settlement benefits customers, but raises vulnerability to fraud. We seek to understand how bank nodes may strategically mitigate fraud risk in RTPs, through investment in fraud detection and restricting payments eligible for real-time processing. To study this, we introduce an agent-based model of the payment network supporting both real-time and standard payments, and define a game among banks and fraudsters. Using empirical game-theoretic analysis, we identify Nash equilibria in nine game configurations defined by network attributes. Our analysis finds that as banks become more liable for fraud, they continue to allow RTPs but are more likely to employ both restrictions and a high level of fraud detection. Fraudsters, in response, switch from targeting only RTPs to attempting fraud with any type of payment and tend to exploit banks where they have historically been most successful. We also conduct a *strategic feature gains assessment* to further understand the benefit offered by each of the bank's risk mitigation measures, which confirms the importance of selective RTP restrictions. Finally, we find that in equilibrium bank strategic decisions negatively affect fraudsters while minimally impacting customers.

1 Introduction

A payment begins with the sender declaring its value and intended recipient to its bank. From there, payment processing unfolds over a series of defined steps, referred to as the clearing and settlement, ending with the final receipt of the funds. Customarily, these steps take a day or more to complete as a result of batch processing and requirements for communication between banks. However, recent advancements in technology have allowed for increasingly faster processing and the introduction of new, faster payment types.

Real-time payments (RTPs) execute the clearing and settlement steps directly upon initiation of the payment, allowing the payment's receiver immediate access to the funds [Committee on Payments and Market Infrastructures, 2016]. The total processing time for a RTP is about 10 seconds [Diadiushkin *et al.*, 2019; European Central Bank, 2023]. Importantly, this convenience is available 24/7/365 rather than subject to business hours as with traditional payments. Many RTP systems are in use around the world today, such as Zelle and FedNow in the United States.

The immediacy of RTPs benefits both senders and receivers of payments. Benefits may also extend to malicious actors (***fraudsters***), who can exploit the limited ability of fraud detection systems to handle this speed. A survey of payment service providers found that current manual review processes average 5 to 10 minutes [Diadiushkin *et al.*, 2019], too slow for RTPs. Thus, fraud detection for RTPs necessitates reliance on fully automated quick checks, which tend to be less accurate. A UK study found a 132% increase in fraud the year the *Faster Payment Service* (FPS) was introduced, though numbers for FPS alone were not available. RTPs also attract a difficult-to-detect attack known as *authorized push payments* (APP) fraud, wherein a customer is tricked into authorizing a fraudulent payment. APP fraud was the second largest type of payments fraud in the UK in 2018 [Taylor and Galica, 2020].

Given the risk of fraud in RTPs, we seek to understand how banks may mitigate fraud risk while providing these beneficial payments to customers. We study this question using an agent-based model of a payments network with banks, fraudsters, and customers modeled as nodes in the network, and supporting the modeling of both real-time and standard payments. Within this model, we define a RTP fraud game played by the bank and fraudster nodes. To mitigate risks—including fraud risk—of RTPs, maximum threshold values above which a payment cannot be sent in real time have been found to be useful tools for banks [Weyman, 2016]. Thus, bank nodes in our game make strategic choices regarding threshold value setting and investment in RTP fraud detection. Fraudsters in the game make strategic decisions determining targets of their fraudulent payment attempts.

The game is analyzed using *empirical game-theoretic analysis* (EGTA) to identify Nash equilibria for various game configurations. To evaluate the benefit to banks for inclusion of each risk mitigation technique, we introduce a process of ***strategic feature gains assessment***, which generalizes an ap-

proach to assessing strategic factors previously employed in an EGTA study by Mayo and Wellman [2021]. Finally, we study outcomes in the network in strategic equilibrium to understand the impact on network participants.

We find that as bank liability for RTP fraud increases, they adopt both threshold values and fraud detection with higher probability. When banks utilize both mitigation techniques, fraudsters switch from attempting fraud only in RTPs to considering any payment type. In selecting banks to target, fraudsters tend to rely heavily on information on historical success. Our strategic feature gains assessment highlights the importance of controlling the bank threshold value for mitigating initial fraud risk. By analyzing equilibrium outcomes, we find that bank strategies are effective at disrupting fraudsters with minimal negative impact to customers.

The main contributions of this work are:

- an agent-based model of a payment system consisting of standard and real-time payments supporting the modeling of fraudulent activity;
- a game-theoretic analysis of RTPs fraud risk mitigation; and
- a process for *strategic feature gains assessment* to analyze the impact of available strategy options.

2 Related Work

An early use of directed graphs to model trust in the context of transactions was an auction application by Ghosh et al. [2007]. Related models were applied to informal borrowing networks [Karlan *et al.*, 2009] and other applications, and later unified by Dandekar et al. [2011] as a general formalism for *credit networks*. Subsequent work explored strategies for issuing credit (i.e., determining trust levels) in such networks [Dandekar *et al.*, 2015]. Cheng et al. [2016] expanded the formalism to including interest rates, referring to the augmented model as *financial credit networks* (FCNs).

General discussions of RTPs describe their distinction from standard payments, and potential benefits and risks for participants [Committee on Payments and Market Infrastructures, 2016; Hartmann *et al.*, 2019]. There are several studies of the design of RTP systems [Guo *et al.*, 2015; Kulk, 2021; Guo and Ma, 2023], including blockchain implementations [Arshadi, 2019; Zhong *et al.*, 2019]. Additional work focuses on risks posed by immediacy and 24/7 availability of RTPs. This includes difficulties with performing maintenance and repair, bank runs in times of distress [Weyman, 2016], fraud detection [Diadiushkin *et al.*, 2019], and liquidity needs [Hellqvist and Korpinen, 2021].

Studies of existing RTP systems analyzed their effects on the broader payment ecosystem. A study of the FPS in the United Kingdom found heavy use of FPS for standing orders and credit card bills, but almost no effect on point-of-sale transactions [Greene *et al.*, 2014]. Another study compared the adoption of RTP systems in six countries identifying the most important factors for successful adoption [Hartmann *et al.*, 2019]. It also found the greatest impact of RTPs was on traditional credit transfers and the use of checks.

To better understand the adoption of RTPs, two prior works analyzed related proxy systems. Bech et al. [2017] identified parallels in the adoption of real-time gross settlement systems for processing whole sale payments and RTPs, particularly noting both experienced a 15-year gap between first introduction and prominent worldwide up-take. Hayashi and Toh [2020] noted the importance of mobile banking as a gateway for accessing RTPs and identified gaps in usage that will need to be addressed for wide adoption of RTPs.

Lastly, there exists a large body of research related to the use of machine learning based techniques for fraud detection [West and Bhattacharya, 2016; Ryman-Tubb *et al.*, 2018], including deep learning [Lin *et al.*, 2021]. Recent work also explores faster methods for fraud detection, some of which is motivated by RTPs [S.N. John *et al.*, 2016; Said and Hajami, 2021; Madhuri *et al.*, 2023].

3 Payments Network Model

We describe the elements of our model of the payment system, building on FCNs [Cheng *et al.*, 2016], which implements the payment network infrastructure on which the RTP fraud game is played.

3.1 Bank and Customer Nodes

Banks and customers are represented by nodes in the network. There are b bank nodes, $B = \{B_1, \ldots, B_b\}$, and n customer nodes, $C = \{C_1, \ldots, C_n\}$. Each customer node is connected to a bank node by edges as shown in Figure 1. The dashed ***debt edge*** represents customer deposits held by the bank, with value denoting the amount of deposits held (i.e., what the bank owes the customer). The customer's willingness to hold additional deposits in its account is represented by a solid ***credit edge***, essentially a credit line bounding what the bank can owe the customer. Such bounds may reflect deposit insurance limits or other customer preferences. Figure 1 shows customer node C_1 with 100 units currently in its bank account at bank node B_1 and a willingness to hold 100 additional units.

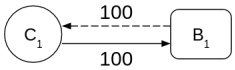

Figure 1: An example of customer node C_1 with 100 units deposited at bank node B_1 and a willingness to hold an additional 100 units in the account for a maximum of 200 units at any given time.

Banks are responsible for routing payments in the network on behalf of their customers through the ***interbank network***. We model this as a fully connected network of credit edges, representing all banks' willingness to interact with one another. Bank nodes are capable of routing both standard and real-time payments through this network. An example of the interbank network connecting two bank nodes, B_1 and B_2, is shown in Figure 2. In this simple network, bank node B_1 has one customer C_1 with 100 units currently in its account and bank B_2 has a customer C_2 with 50 units currently deposited. Each bank node has a willingness to route up to 1,000 units to the other.

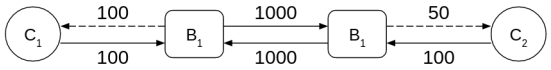

Figure 2: A small example network with two bank nodes, each with a single customer node, who are connected by the interbank network.

3.2 Modeling Payments

Payments are implemented by a series of updates to the edges between parties involved: the sender, receiver, and their respective banks. The payment type is differentiated by the timing of these updates.

Real-Time Payments

When a customer node (***sender***) initiates a payment, it is reflected by a decrease in the value on the debt edge by the amount of the payment. This represents the customer drawing on their deposits to make the payment. The sender now holds fewer deposits in its account, so its willingness to hold additional deposits must increase by the amount of the payment captured by an increase on the value of its credit edge. The sender's bank node then routes the payment to the receiver's bank, creating a debt edge from the sender's bank to the receiver's bank of the payment's amount. The action of routing a payment decreases the receiver bank's willingness to interact with the sender bank in the future by the amount of the payment, decreasing the value on the credit edge from the receiver bank to the sender bank. Finally, we model the receiver obtaining the funds in their bank account by an increase of the value on the debt edge from the receiver's bank to the receiver equal to the payment amount. With an increase in the amount of deposits in its account, the receiver's willingness to hold additional deposits must decrease by the amount of the payment reflected by a decrease in the value on the receiver's credit edge.

In an RTP scenario, these processing steps all occur instantaneously and thus all edges will be updated as described in the same time step the payment was initiated. An example of customer node C_1 making a RTP of 20 units to C_2 in time step $t = 1$ is shown in Figure 3 with the initial state of the network shown in Figure 2. As can be seen, all edges on the path between the sender and receiver node have been updated.

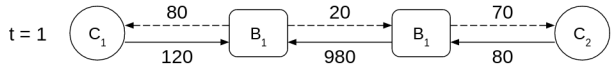

Figure 3: An example of C_1 sending a RTP of 20 units to C_2 in time step $t = 1$, resulting in all updates to the edges occurring at $t = 1$. The initial network configuration is shown in Figure 2.

.

Standard Payments

To model the delay between payment initiation and completion of payment processing for a standard payment, we simply delay some of the edge updates. Specifically, when a payment is initiated, only the credit and debt edges between the sender and its bank node are updated. The remaining steps go into the sender bank's payment queue, representing the delayed processing. At a later time period, called the ***clearing period***, the unprocessed payments in the bank queues are removed and applied to the network. At this time, the edges of the interbank network and between the receiver and its bank will be updated as described for RTPs.

The same payment of 20 units from customer node C_1 to C_2 is shown now as a standard payment in Figure 4. While the payment updates unfold over several time steps, the final impact of the payment on the network remains the same.

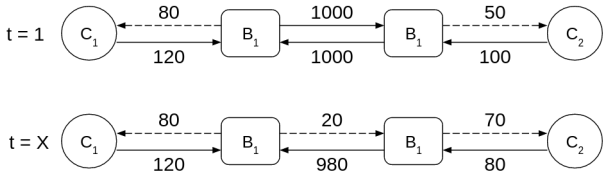

Figure 4: Customer node C_1 makes a standard payment of 20 units to C_2 in time step $t = 1$, which completes in the next clearing period X.

3.3 Modeling Fraud

Fraud in this network is committed by a ***fraudster node*** drawing on customer node's deposits to make payments within the network, referred to as ***impersonating*** a customer node. When a fraudster node impersonates a given customer, it creates a ***fraud edge*** connecting the two nodes. The value on the fraud edge is equal to the amount of fraud committed by the fraudster while impersonating the customer. Since the other nodes are unaware the payment is truly fraudulent, the steps for each type of payment outlined above remain the same. Figure 5 illustrates an example of fraudster node F_1 impersonating C_1 to make a fraudulent RTP of 20 units to C_2. As shown, this creates a fraud edge with 20 units connecting F_1 and C_1, while the remaining edge updates appear the same as the non-fraudulent RTP shown in Figure 3. A standard payment of 20 units would look similar to Figure 4 with the addition of the fraud edge.

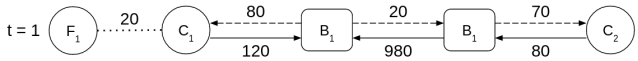

Figure 5: Fraudster node F_1 impersonates customer C_1 to make a fraudulent RTP of 20 units to customer C_2.

4 RTP Fraud Game Initialization

We describe the formation of the initial payments model in which the RTP fraud game will be played, as well as the strategy sets available to the strategic agents. As the game is played, the initial network will dynamically update as a result of behavior described in Section 5.

4.1 Customer Nodes

We initialize $n = 200$ customer nodes with initial deposits drawn from an exponential distribution with an average value of 10,000. Each customer node is also initialized with its

own preferences for RTPs usage as both a sender and receiver. A given customer will be willing to only receive RTPs with probability λ and with probability $1 - \lambda$ be willing to accept any type of payment. Furthermore, some customers may have an aversion to utilizing a new, unfamiliar technology to send payments. To capture this, each customer node has its own personal threshold, $CT_i \in \{200, 400, 600, 800, 1000\}$, uniformly randomly chosen during initialization. A customer node in the network is unwilling to send a payment in real time if its value exceeds the customer's personal threshold. In our model, the maximum payment value is 1,000, so a customer with $CT_i = 1,000$ is always willing to send RTPs.

Lastly, we note that customers may not necessarily always desire to send an eligible payment in real time due to personal preference and payment requirements. We model this with the ***urgency*** parameter such that with probability $u = 0.5$ an eligible payment is sent in real time and with probability $1 - u$ it is sent as a standard payment. The urgency parameter universally applies to all customer nodes.

4.2 Bank Nodes

A set of $b = 4$ banks nodes is initialized to form the interbank network and each bank is initialized with its own standard payments queue, Q_i. In our game, banks have an infinite willingness to route payments on behalf of their customers. Bank nodes have access to two fraud detectors, one for each type of payment. The focus of our work is to understand the use rather than emphasize a particular method of fraud detection, and thus we abstract away implementation details by modeling both detectors as black boxes. Each detector is characterized only by the probability (γ) it correctly labels a payment as fraudulent or not fraudulent relative to its true label. Bank node B_i's standard payments fraud detector γ_i^S is initialized with a random draw from U $[0.8, 0.9]$ and its RTPs detector γ_i^R by the bank node's chosen strategy as described below.

Bank Strategy Set

During initialization, each bank node will select from a set of two-part strategies that control RTPs usage and fraud detection. The first part of the strategy sets a bank threshold $BT_i \in \{0, 200, 400, 600, 800, 1000\}$ which is the maximum payment value bank B_i allows to be sent in real time. A threshold *BT = 0* does not allow any RTPs, threshold *BT = 1000* does not restrict RTPs, and the remaining choices offer varying degrees of limitation. The threshold choices of bank nodes is assumed to be public information that attracts customers to the bank. The second part of the strategy is the bank's investment level in RTPs fraud detection, which sets the bank node's RTP fraud detection capabilities for the duration of the game and incurs a one-time cost. A higher investment allows for a more accurate fraud detector (higher γ^R), but at a higher cost to the bank. To capture the challenges of speed, fraud detection accuracy for RTPs at every investment level is strictly less than that of standard payments. Table 1 shows the mapping of investment level choices to one-time cost and γ^R setting for the game, derived from the bank's γ^S setting. The strategy set for bank nodes contains all 24 combinations of threshold choices and investment levels.

investment level	cost (units)	$\gamma^R =$
none	0	—
low	1,000	$0.65 \times \gamma^S$
high	3,000	$0.9 \times \gamma^S$

Table 1: A mapping of the investment level of a bank node in RTPs fraud detection to the one-time incurred cost and RTPs fraud detection capability (γ^R) for the game.

4.3 Fraudster Nodes

The network is initialized with $m = 1$ fraudster node. The fraudster node maintains a history of its success committing fraud at each bank node to inform its strategic behavior. At the start of the game, the fraudster is assigned its first customer node to impersonate (***victim***) based on its selected strategy.

The strategies of the fraudster node determine its bank and payment type targeting rules. We assume given the opportunity, the fraudster node will always prefer to send RTPs to exploit the limits of fraud detection in the quicker payment scheme. If banks severely restrict RTPs, however, it may be beneficial for the fraudster to gamble on the imperfect nature of standard payments fraud detection rather than completely miss an opportunity to attempt fraud. Thus, the fraudster will choose between one of two rules: attempt RTPs only (RTP-only) or be willing to attempt a standard payment if a RTP is not possible (Any).

The fraudster in the game targets bank nodes, but does not differentiate between customer nodes of a targeted bank. We assume information regarding customer and bank relationships is accessible to the fraudster at negligible cost. The fraudster node selects the bank node to target using its strategy, and then randomly selects a victim from customers of the target bank. When a payment attempt is caught, the fraudster will select a new victim, potentially of a different target bank depending on the strategy being followed.

There are two main bank targeting rules the fraudster chooses from: threshold and success. The threshold rule simply targets the bank with the highest RTPs threshold. The success rule targets the bank where the fraudster has been most successful in the past, modeled as a *multi-armed bandit problem* and implemented using the *Upper Confidence Bound* algorithm [Auer *et al.*, 2002]. Specifically we calculate the target bank (b^*) as:

$$argmax_b \frac{\sum_{s=1}^{t} I_{b(s)=b} X_b(s)}{N_b(t)} + \rho \times \sqrt{\frac{log(t)}{N_b(t)}},$$

where t is the current time step, $N_b(t)$ is the number of times bank b has been targeted, and X_b is 1 to indicate a successful attempt at b and 0 otherwise. A larger value for ρ indicates greater exploration than exploitation. We test three versions of this strategy defined by $\rho \in \{0.2, 1, 2\}$ denoted success_ρ. For both rules, if multiple banks fit the criteria, one bank is uniformly randomly chosen as the target each time.

The last targeting rule, random, randomizes over threshold and success each time a new victim is selected. We test two versions of this strategy defined by $\rho \in \{0.2, 2\}$ used for the success rule which we refer to as random_ρ. With two

payment targeting rules and six bank targeting rules, there are 12 possible fraudster node strategies.

4.4 Network Formation

The network is formed by assigning customer nodes to bank nodes using bank and customer RTP thresholds. We assume customers prefer a bank that allows them to send RTPs with values at least equal to their personal threshold and if no such bank exists, would prefer a bank offering the highest possible threshold. Customers are assigned to their preferred bank and in the event multiple fit the criteria, one such bank will be randomly selected. Upon assignment to a bank, customers deposit all their initial funds creating a debt edge from the bank node to the customer node as outlined in Section 3.1. Customer nodes in our game have an infinite willingness to hold additional deposits in their account.

5 RTPs Fraud Game

The RTPs fraud game unfolds over $T = 2,880$ time steps with each time step containing payment creation, fraud detection, and possible network updates. Every 96 time steps (the clearing period), the bank queues are cleared to update the network as described in Section 3.2 before new payments are created. At the end of the game, the payment queues are cleared for the final time and bank and fraudster nodes receive a payoff for the performance of their selected strategies.

5.1 Payment Creation

During the payment creation phase, $\eta = 20$ customer nodes are randomly selected, without replacement, to make a payment. With probability $\mu = 0.7$ a fraudster node is one of the η selected. The value of the payment is drawn from $v = U\ (0, min(1000, d)]$ where d is the amount of deposits currently in the sender's bank account. The payment receiver is randomly selected from the customer nodes in the network. The payment type is determined by the sender, its bank, and v as outlined below.

Customer Node Payments

Consider a payment sent by customer node C_1 belonging to bank B_1. If the receiver only accepts RTPs, the payment must be sent in real time or the attempt will be canceled due to incompatibility and considered a missed opportunity for a transaction within the network. The payment can be sent in real time only if $v \leq CT_1$ *and* $v \leq BT_1$. If the receiver is willing to accept any type of payment, it will be sent as a standard payment if the sender is unwilling ($v > CT_1$) and/or unable ($v > BT_1$) to send it in real time. Otherwise, we consider the *urgency* the customer has for the timely processing of this payment as described in Section 4.1.

Fraudster Node Payment Type

The fraudster node is constrained only by the sending bank's threshold and as mentioned, always prefers RTPs when possible. Thus, the fraudster with a victim belonging to bank node B_1 will attempt a RTP if $v \leq BT_1$. Otherwise, the payment type is dictated by the receiver's preferences and the fraudster's strategy. If the receiver is willing to accept any payment and the fraudster is following the Any strategy, a standard payment attempt will be made. Otherwise, the payment attempt is canceled triggering the fraudster's selection of a new victim.

5.2 Fraud Detection and Network Updates

Payments not canceled in the creation phase go through the sender bank's black box fraud detectors. Bank node B_i's fraud detector labels the payment correctly with respect to the true label with probability γ_i and incorrectly with probability $1 - \gamma_i$. If a payment is a RTP γ_i^R is applied and if it is a standard payment, γ_i^S. Payments labeled fraudulent are canceled regardless of the true label and remaining payments are used to update the network according to payment type as described in Section 3.

5.3 Payoffs

The payoff to bank nodes is composed of the effects of services offered and the effects of invoking fraud detection. Bank node B_i receives the following payoff:

$$\delta_1 ID_i + \delta_2 VA_i - \alpha VRF_i - VSF_i - RFI_i - \beta VFP_i.$$

By offering RTPs, a bank attracts some amount of customer deposits (*ID*). Though the deposits themselves are a bank liability, they represent partial assets in the form of continued business from customers which we capture with $\delta_1 = 0.5$. Banks in the network route some total amount of payments in the network on behalf of customers, *VA*, on which banks may be able to charge small fees represented by $\delta_2 = 0.01$. When fraud occurs, bank nodes are liable for the full amount of standard payment fraud (*VSF*) and may hold some liability for RTPs fraud (*VRF*) as described in Section 6.1. Choosing to invest in fraud detection results in a one-time cost to banks (Table 1) denoted here by *RFI*. Lastly, imperfect fraud detectors may lead to some payments erroneously being labeled fraudulent and subsequently being canceled. So called false positives represent missed opportunities for routing and may lead to frustration and less frequent business by customers. To account for this potential cost we define *VFP* as the total value of false positive payments and define $\beta = 0.2$.

The fraudster node seeks only to maximize the fraud it successfully commits, reflected by a payoff equal to the total value of fraud it commits across all banks and payment types over the course of the game.

6 Experiments

6.1 EGTA

We analyze the RTPs fraud game using EGTA [Wellman, 2016; Tuyls *et al.*, 2020], a method employing extensive simulation of strategy profiles in a game to identify Nash equilibria. A strategy profile is selected for simulation, and bank and fraudster nodes in our model are assigned to play the strategies according to the profile in the game described in Section 5, after which nodes receive a payoff. This process iterates many times such that each profile is simulated in 3,000 random generations of the game and payoffs for the strategies in the profile are calculated from the sample average of payoffs observed over the many simulation runs. EGTA employs an iterative procedure to select profiles for simulation,

seeking to identify symmetric mixed-strategy Nash equilibria while avoiding an exhaustive search of the profile space [Cassell and Wellman, 2013; Fearnley *et al.*, 2015]. To understand how the Nash equilibria are affected by attributes of the network, we perform the EGTA analysis on nine different configurations of the RTPs fraud game defined by customer demand ($\lambda \in \{0.25, 0.5, 0.75\}$) and bank liability for RTPs fraud ($\alpha \in \{0, 0.5, 1\}$).

6.2 Strategic Feature Gains Assessment

To better assess the individual benefit of each of the fraud risk mitigation techniques to bank nodes, we introduce the *strategic feature gains assessment*. The assessment aims to understand the benefit to an agent in a game provided by gaining access to a set of strategies united by commonalities such as the basis for decision making or their effect on agent behavior. For a game with strategy set S, we define this set of strategies, the ***deviation set***, as $\Delta \subset S$. We also define the set of initial strategies agents may access, the ***base set***, as $\Omega \subset S$ which is disjoint from Δ. Finally, we define σ as a strategy profile in the game, σ^* as denoting a Nash equilibrium, and $u_i(\sigma)$ as the payoff obtained by agent i for playing strategy profile σ. The payoff to agent i for deviating to strategy s while the other agents play according to σ is denoted $u_i(\sigma_{-i}, s_i)$. The steps of the assessment are defined as follows:

1. Define Δ and Ω
2. Obtain $\sigma^*(\Omega)$ using EGTA
3. Gain from $\Delta = \max_{s \in \Delta} u_i(\sigma^*_{-i}, s_i) - u_i(\sigma^*)$

When a game contains agents belonging to multiple roles, such as banks and fraudsters, we may only wish to perform the assessment on one role. In this case, agents of the non-assessed role (here fraudsters) retain use of their full strategy set. Performing this assessment with different Ω and Δ sets can provide a point of comparison between different scenarios or different information sets.

We perform the four assessments listed in Table 2 on the nine configurations of the RTPs fraud game. Assessments A1 and A2 seek to understand the gains to a bank node with access to both mitigation techniques after initially starting with just one: either setting a threshold (*BT*) or investing in fraud detection (*FDI*). Assessments A3 and A4 focus on the scenario where banks begin without any mitigation measures, in both the no RTPs and unrestricted RTPs cases, to see which technique they employ first when forced to choose only one.

assessment	base set	deviation set
A1	BT only	BT *and* FDI
A2	FDI only	FDI *and* BT
A3	no RTPs	FDI *or* BT
A4	BT=1000, FDI=none	FDI *or* BT

Table 2: A description of each assessment performed describing the strategies in the base and deviation sets, with *BT* denoting threshold setting and *FDI* denoting fraud detection investment setting.

The assessment introduced here is a generalization of an approach employed by Mayo and Wellman [2021] in their EGTA study of eliminating debt cycles among agents in a financial network. In the model studied in that work, agents must decide to agree or disagree to a potential cycle elimination using a strategy from a set of strategies that make the strategic decision using information available to the agent, for example debts owed. For analysis, the strategies were grouped by information employed for the decision and the groups were referred to as *features*. For example, the *assets* included all strategies using agent asset holdings to make the elimination decision. The feature analysis sought to understand the importance of each feature for decision making. Similar to our assessment, gains were measured by the maximum payoff gain to an agent for deviating from the equilibrium in a base set to a strategy in a feature group which functioned as the deviation set. However, the feature analysis defined the base and deviation sets more stringently such that $\Omega \cup \Delta$ was the full strategy set S. Furthermore, if two strategies employed the same information, they must both be in the same feature group. Our generalization allows some strategies to be left out entirely from the assessment as in A3 and A4 and provides more flexibility for defining the membership of each set.

6.3 Effects on the Network

Finally, we analyze the effects of the equilibria on the network by measuring outcomes for various network participants over 1,000 runs of the game. In the event of a mixed-strategy equilibrium, nodes are assigned to play a pure strategy from a weighted draw according to the equilibrium distribution. We compare these outcomes to those in the two settings: where banks do not offer RTPs and where bank nodes do not restrict RTPs, nor invest in fraud detection.

7 Experiment Results

7.1 Nash Equilibria

We find that bank nodes in our game select strategies with higher value thresholds, but often balance this with the use of fraud detection when banks are liable for RTP fraud. Specifically, when banks are not liable for RTP fraud, they do not restrict customer access to RTPs, nor do they employ fraud detection. With partial liability, bank nodes begin to set thresholds and invoke fraud detection. Interestingly, the equilibrium show a strategic trade-off between customer access to RTPs and the cost of fraud detection in this setting. Bank nodes tend to either invoke fraud detection to avoid restricting customer RTP usage *or* set a threshold value instead of incurring the cost for fraud detection. When bank nodes become liable for the full amount of RTPs fraud, they exhibit a higher probability of both setting a threshold and investing in good fraud detection. In this case, the trade-off between fraud detection costs and customer RTP restriction is only seen when customer demand for RTPs is low. We find that customer demand for RTPs has a smaller effect on bank equilibria than RTP fraud liability, with greater customer demand resulting in only slight increases in threshold values selected.

In response to the banks' strategic decisions, the fraudster in our network tends to target banks based on historical success, either with success or random strategies, with a

high degree of exploitation. The fraudsters prefers to follow RTP-only unless bank nodes implement both fraud mitigation measures, invoking fraud detection *and* setting a threshold value. In this case, more barriers to successful RTPs fraud forces the fraudster to adopt the Any strategy.

7.2 Strategic Feature Gains Assessment

While both mitigation techniques appear in the equilibria of bank nodes, the strategic feature gains assessment highlights the particular importance of control over customer RTP use. The results from A1 and A2, shown in Figure 6, demonstrate greater gains to bank nodes when starting by setting a threshold (*BT*) and gaining the ability to invest in RTP fraud detection (*FDI*) compared to the reverse scenario. When thresholds are set first, the addition of FDI allows banks to deviate to a strategy with a higher threshold value, expanding RTP usage. Conversely, starting with FDI results in a restriction to RTP usage and a lower payoff gain for banks. Though it should be noted bank nodes almost always benefit from the ability to use both mitigation techniques with the exception of the cases where the equilibrium is to not restrict RTPs nor offer fraud detection. The importance of initial threshold setting is further supported by the results of A3 and A4. In both cases we find a bank node in all game configurations will choose to deviate to a strategy controlling the threshold value when given the option to choose only one mitigation technique.

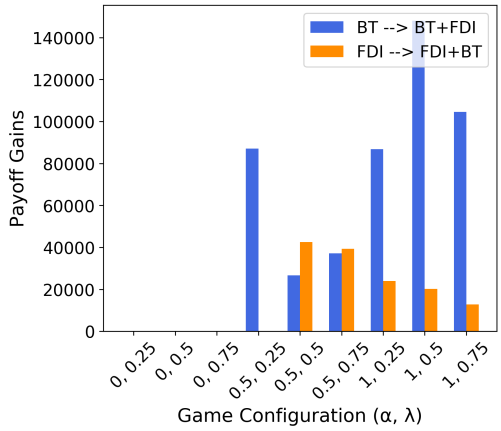

Figure 6: The results from A1 and A2 show that a bank node benefits the most by gaining the ability to control fraud detection investment level (*FDI*) *after* initially controlling only the threshold value (*BT*).

7.3 Network Effects

Under the Nash equilibria, all bank nodes will be targeted by the fraudster over the course of the game and at least 1/3 of customer nodes experience impersonation. The number of victims increases as bank nodes become more concerned with fraud and implement stricter mitigation measures, leading to less success by the fraudster. This forces the fraudster to impersonate more customers to accomplish its goal resulting in up to 80% of customer nodes experiencing impersonation.

We explore the effects of banks employing fraud risk mitigation techniques on customer nodes by analyzing their rate of successful payment attempts under the Nash equilibria compared to when bank nodes do not restrict RTPs nor use fraud detection. Figure 7 shows only a small gain in success for customer nodes when bank nodes switch to the latter case and a much larger gain for the fraudster. This indicates the strategic choices of banks greatly reduce fraudster success with minimal interruptions to customers.

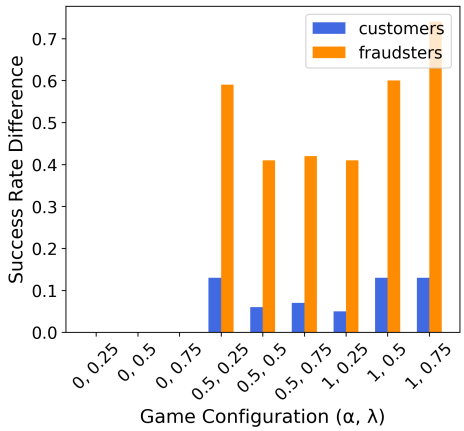

Figure 7: Gain in success rate for fraudsters and customers when banks allow unrestricted RTPs, in equilibrium comparisons.

8 Conclusion

We explore banks' strategic mitigation of fraud risk in RTPs, balancing benefit to customers with vulnerability to fraud. To maintain this balance, banks often supplement the use of fraud detection with restrictions to RTPs use. We investigate how banks may strategically trade-off between the use of these two mitigation methods and the strategic response of fraudsters.

We study this problem in an agent-based model of a payment system in which a RTP fraud game is played among banks and a fraudster. The game is analyzed using EGTA to identify Nash equilibria in nine game configurations defined by bank liability for RTPs fraud and customer demand for RTPs. Our results indicate that with no liability, bank nodes are willing to allow unrestricted access to RTPs and do not invoke fraud detection. However, as banks take on more liability, they become more likely to employ higher fraud detection measures and restrict RTPs. In response, the fraudster node must adjust from targeting only RTPs to a willingness to use any payment type. We further study the strategic decision of bank nodes by introducing the generalized strategic feature gains assessment to gauge the relative and individual importance of each risk mitigation technique. The results of this assessment identify the importance of initial restriction to RTPs for mitigating fraud risk. Lastly, we find bank strategic decisions in equilibrium effectively lower fraud risk with little impact on customer outcomes.

Our current work defines customer preferences stochastically, however, exploring the effect of strategic customers on banks decisions may be an interesting avenue for future research.

Ethical Statement

There are no ethical issues.

References

[Arshadi, 2019] Nasser Arshadi. Blockchain platform for real-time payments: A less costly and more secure alternative to ACH. *Technology & Innovation*, 21(1):3–9, October 2019.

[Auer *et al.*, 2002] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, May 2002.

[Bech *et al.*, 2017] Morten L. Bech, Yuuki Shimizu, and Paul Wong. The quest for speed in payments. *BIS Quarterly Review*, March 2017.

[Cassell and Wellman, 2013] Ben-Alexander Cassell and Michael P. Wellman. EGTAOnline: An experiment manager for simulation-based game studies. In *Multi-Agent Based Simulation XIII*, volume 7838 of *Lecture Notes in Artificial Intelligence*, pages 85–100. Springer, 2013.

[Cheng *et al.*, 2016] Frank Cheng, Junming Liu, Kareem Amin, and Michael P. Wellman. Strategic payment routing in financial credit networks. In *17th ACM Conference on Economics and Computation*, pages 721–738, 2016.

[Committee on Payments and Market Infrastructures, 2016] Committee on Payments and Market Infrastructures. Fast payments: Enhancing the speed and availability of retail payments. Technical Report 154, Bank for International Settlements, November 2016.

[Dandekar *et al.*, 2011] Pranav Dandekar, Ashish Goel, and Ramesh Govindan. Liquidity in credit networks: A little trust goes a long way. In *12th ACM Conference on Electronic Commerce*, pages 147–156, 2011.

[Dandekar *et al.*, 2015] Pranav Dandekar, Ashish Goel, Michael P. Wellman, and Bryce Wiedenbeck. Strategic formation of credit networks. In *ACM Transactions on Internet Technology*, volume 15, pages 1–41, March 2015.

[Diadiushkin *et al.*, 2019] Alexander Diadiushkin, Kurt Sandkuhl, and Alexander Maiaitin. Fraud detection in payments transactions: Overview of the existing approaches and usage for instant payments. *Complex Systems Informatics and Modeling Quarterly*, pages 72–88, 2019.

[European Central Bank, 2023] European Central Bank. What are instant payments?, 2023.

[Fearnley *et al.*, 2015] John Fearnley, Martin Gairing, Paul Goldberg, and Rahul Savani. Learning equilibria of games via payoff queries. *Journal of Machine Learning Research*, 16:1305–1344, 2015.

[Ghosh *et al.*, 2007] Arpita Ghosh, Mohammad Mahdian, Daniel M. Reeves, David M. Pennock, and Ryan Fugger. Mechanism design on trust networks. In *Third International Workshop on Internet and Network Economics*, pages 257–268, 2007.

[Greene *et al.*, 2014] Claire Greene, Marc Rysman, Scott D. Schuh, and Oz Shy. Costs and benefits of building faster payment systems: The UK experience and implications for the United States. Technical report 14-5, Federal Researve Bank of Boston Research Paper Series Current Policy Perspectives, October 2014.

[Guo and Ma, 2023] Zhiling Guo and Dan Ma. Catching the fast payments trend: Optimal designs and leadership strategies of retail payment and settlement systems. *MIS Quarterly*, 47(2):669–704, 2023.

[Guo *et al.*, 2015] Zhiling Guo, Rob Kauffman, Mei Lin, and Dan Ma. Near real-time retail payment and settlement systems mechanism design. Working paper, SWIFT Institute, September 2015.

[Hartmann *et al.*, 2019] Monika Hartmann, Lola Hernandez van Gijsel, Mirjam Plooij, and Quentin Vandeweyer. Are instant payments becoming the new normal? A comparative study. Occasional Paper Series 229, European Central Bank, August 2019.

[Hayashi and Toh, 2020] Rumiko Hayashi and Ying Lei Toh. Mobile banking use and consumer readiness to benefit from faster payments. *Federal Reserve Bank of Kansas City Economic Review*, 105:1–5, April 2020.

[Hellqvist and Korpinen, 2021] Matti Hellqvist and Kasperi Korpinen. Instant payments as a new normal: Case study of liquidity impacts for the Finnish market. Technical report 7, BoF Economics Review, 2021.

[Karlan *et al.*, 2009] Dean Karlan, Markus Mobius, Tanya Rosenblat, and Adam Szeidl. Trust and social collateral. *Quarterly Journal of Economics*, 124(3):1307–1361, 2009.

[Kulk, 2021] Erwin Kulk. Request to pay: Monetising the instant payments investment. *Journal of Digital Banking*, 5(3):193–203, January 2021.

[Lin *et al.*, 2021] Wangli Lin, Li Sun Qiwei Zhong, Can Liu, Jinghua Feng, Xiang Ao, and Hao Yang. Online credit payment fraud detection via structure-aware hierarchical recurrent neural network. In *30th International Joint Conference on Artificial Intelligence*, pages 3670–3676, 2021.

[Madhuri *et al.*, 2023] T. Sirisha Madhuri, E. Ramesh Babu, B. Uma, and B. Muni Lakshmi. Big-data driven approaches in materials science for real-time detection and prevention of fraud. *Materials Today*, 81:969–976, 2023.

[Mayo and Wellman, 2021] Katherine Mayo and Michael P. Wellman. A strategic analysis of portfolio compression. In *2nd ACM International Conference on AI in Finance*, 2021.

[Ryman-Tubb *et al.*, 2018] Nick F. Ryman-Tubb, Paul Krause, and Wolfgame Garn. How artifical intelligence and machine learning research impacts payment card fraud detection: A survey and industry benchmark. *Engineering Applications of Artifical Intelligence*, 76:130–157, 2018.

[Said and Hajami, 2021] Marouane Ait Said and Abdelmajid Hajami. Ai methods used for real-time clean fraud detection in instant payments. In *13th International Conference on Soft Computing and Pattern Recognition*, 2021.

[S.N. John *et al.*, 2016] Okokpujie Kennedy O. S.N. John, C. Anele, F. Olajde, and Chinyere Grace Kennedy. Real-time fraud detection in the banking sector using data mining techniques/algorithm. In *International Conference on Computational Science and Computational Intelligence*, 2016.

[Taylor and Galica, 2020] John L. Taylor and Tony Galica. A new code to protect victims in the UK from authorized push payments fraud. *Banking Finance Law Review*, 35:327–332, May 2020.

[Tuyls *et al.*, 2020] Karl Tuyls, Julien Perolat, Marc Lanctot, Edward Hughes, Richard Everett, Joel Z. Leibo, Csaba Szepesvari, and Thore Graepel. Bounds and dynamics for empirical game-theoretic analysis. *Autonomous Agents and Multi-Agent Systems*, 34(7), 2020.

[Wellman, 2016] Michael P. Wellman. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems*, 30:1175–1189, 2016.

[West and Bhattacharya, 2016] Jarrod West and Maumita Bhattacharya. Intelligent financial fraud detection: A comprehensive review. *Computers Security*, 57:47–66, 2016.

[Weyman, 2016] Julius Weyman. Risks in faster payments. Working paper, Retail Payments Risk Forum, May 2016.

[Zhong *et al.*, 2019] Lin Zhong, Qianhong Wu, Jan Xie, Zhenyu Guan, and Bo Qin. A secure large-scale instant payment system based on blockchain. *Computers & Security*, 84:349–364, 2019.