RZ 3787 (# Z1009-001) 09/09/2010
Computer Science 14 pages

Research Report

Risk-Based Payment Fraud Detection

K. Julisch

IBM Research – Zurich
8803 Rüschlikon
Switzerland

Research
Almaden • Austin • Beijing • Delhi • Haifa • T.J. Watson • Tokyo • Zurich

Risk-Based Payment Fraud Detection

Klaus Julisch

IBM Research
Zurich Research Laboratory
Switzerland
kju@zurich.ibm.com

Abstract. Payment frauds are a form of intrusions where money is transferred from a victim's bank account to the bank account of a fraudster. This paper reviews why in practice, it is neither possible nor economical to prevent 100 percent of payment frauds. It therefore becomes important to detect payment frauds and to control their cost. The paper then combines two unique insights into the nature of payment fraud to propose a new risk-based payment fraud detection method. This method does not try to detect individual fraudulent payments but rather seeks to quantify the expected loss from frauds over a given time period. This expected loss is the risk posed by frauds, and fraud managers only need to intervene if this risk exceeds a maximum acceptable loss threshold. Below this threshold, payment fraud related losses represent a contained risk, which should be viewed as an operating cost just like shoplifting is an operating cost in retail. The paper critically appraises the risk-based method and discusses its applicability in practice.

Key words: Payment fraud, fraud detection, risk management.

1 Introduction

Fraud is a deception deliberately practiced in order to secure unfair or unlawful gain [1]. Payment fraud occurs when deceptive methods are used to transfer money against the payor's will to an illegitimate beneficiary's account. The parties involved in electronic payment fraud are as follows (Figure 1) [2,3]:

- The *Originator* (a.k.a. *payor*) rightfully owns an account and uses it to issue payment transactions (credits or debits).
- The *Originating Financial Institution (OFI)* is a financial institution (bank, credit union, savings banks) that executes payment transactions on behalf of the originator.
- There are many *channels* that the originator can use to submit payment transactions to the OFI. Examples of such channels include online banking, fax or phone orders, standing orders, debit card payments, and checks. Each of these channels has been targeted by fraudsters.
- The *Fraudster* is a malicious attacker who seeks to illegitimately appropriate funds from any of the other parties shown in Figure 1.

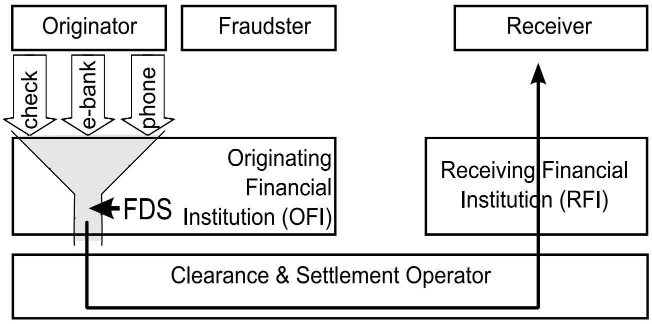
Originator
Fraudster
Receiver
check
e-bank
phone
Originating Financial Institution (OFI)
FDS
Receiving Financial Institution (RFI)
Clearance & Settlement Operator

Fig. 1. Overview of electronic payment systems.

- The *Clearance & Settlement Operator* runs an electronic network that transmits, clears, and settles payments between financial institutions.
- The *Receiving Financial Institution (RFI)* is a financial institution that receives and executes payment transactions from OFIs. In general, each RFI is also an OFI and vice versa.
- The *Receiver* (a.k.a. *payee*) is the person or institution that receives the payment transactions (credit or debit) from the originator. *Credit transactions* increase the receiver's account balance, while *debit transactions* decrease it.

In this paper, I take the position of the OFI who seeks to detect and prevent payment frauds. In general, OFIs want to detect *any* fraud, even when they are not legally or financially liable for them. For example, Swiss banks are not liable for frauds that are committed using the e-banking channel. In practice, however, Swiss banks (and banks, in general) are very anxious to prevent and detect e-banking frauds as a service to their customers and to protect their reputations. I therefore defined frauds as payments that the payor did not *intend* to happen. This is clearly a "soft" definition as it is impossible to discern the payor's intentions. Nonetheless, this definition is appropriate as it captures the problem as it presents itself in practice.

As shown in Figure 1, payment transactions are submitted via many channels. All those channels are merged into a single payment format (as indicated by the funnel), which is then submitted to the Clearance & Settlement Operator. To cover frauds from all channels as well as frauds that involve multiple channels (a.k.a. *multi-channel frauds*) it is advisable to place a Fraud Detection System (FDS) at the stem of the funnel, as indicated in the figure by the "FDS" symbol. A typical example of a multi-channel fraud consists in a fraudster breaching an account via the online channel to steal account balances, check images, and signature blocks; this and other information is subsequently used to commit wire, check other other off-line frauds, which never get linked to the original online breach [4]. In a variant of the architecture shown in Figure 1, the FDS at the stem of the funnel is replace or complemented by specialized fraud detection systems that are located within the channels; the alerts from these systems can be correlated to detect multi-channel frauds. This paper assumes the architectural variant shown in Figure 1.

The remainder of this paper is structured as follows: Section 2, classifies the methods used by fraudsters to commit payment frauds and it explains why it is neither possible nor economical to prevent 100% of payment frauds. This emphasizes the importance of payment fraud detection. Section 3 reviews today's state of the art in payment fraud detection. Section 4 presents a new fraud detection method, which breaks with the convention of detecting fraudulent transactions and instead, detects unacceptably high payment risks. Section 5 discusses the pros and cons of this method as well as its applicability in practice. Section 6 summarizes and concludes the paper.

The novel contribution of this paper is to provide a detailed introduction to payment fraud (Sections 1-3), a topic that has been treated sparingly in the literature due to its sensitive nature. Moreover, the paper proposes a new payment fraud detection method (Sections 4 and 5), which detects unacceptably high risks rather than detecting individual fraudulent payments.

2 Understanding Payment Fraud

Fraudsters use various methods in an attempt to appropriate other parties' funds. These methods are *impersonation*, *deception*, and *server-side attacks*. I will consider each of these methods in turn and then address why fraud detection must complement any preventive measures that an OFI may take.

Impersonation (a.k.a. Identity Theft): A fraudster impersonates an account holder and initiates fraudulent payments using the account. Phishing as well as man-in-the-middle attacks are classic methods used to obtain the credentials needed to impersonate account holders [5,6]. Further impersonation techniques include social engineering where a fraudster persuades account holders to surrender their credentials. Check forgery also falls into this category and so does a fraud [7] where the fraudsters rent space in the same building as a company and then apply for corporate credit cards using the company's name. The application passes the credit check because the company name and address match, and the fraudsters receive the cards in their mailbox. In another impersonation fraud, criminals tried to impersonate the National Bank of Ethiopia in order to transfer money out of their accounts [8]. These examples represent a small sample because criminals constantly invent new impersonation frauds.

Deception: In deception, a fraudster does not try to impersonate a victim, but rather tries to deceive the victim into initiating payments to the fraudster's benefit. These frauds exploit the credulity of people. Pyramid and Ponzi schemes fall into this category. Another example is a fraudster who calls a (typically small) company and says something like: ”This is your supplier XYZ; we are updating our accounts. For all future payments, please use our new account number #xyz”, which, of course, is the fraudster's account number [6]. In another scam, the victim receives a letter informing him or her about some unexpected windfall. The letter includes a check with the money, but also asks the victim to reimburse taxes or processing fees. The enclosed check is a fraud and will bounce, but most victims pay the tax/processing fee before they notice this. In bust-out frauds, the fraudster assumes a fake identity and engages in business transaction with other parties so as to build credibility and a credit history. Then, the fraudster ”maxes out” his credit from the other parties, takes the money and runs [6]. Again, these are mere examples and fraudsters regularly invent new ways to exploit the credulity of people.

Server-Side Attacks: The fraudster compromises the information systems used by financial institutions to process payments and then triggers illegitimate payments. Such server-side attacks are different from malware, spyware, or other client-side attacks that steal login credentials in support of impersonation frauds. Examples of server-side attacks include insider fraud by bank employees and the infamous logic bomb that shaved fractions of pennies off each payment transaction [9].

The difficulties associated with preventing payment fraud now become apparent: Fraud is at least partially a ”human problem” as evidenced by deception frauds or insider frauds. Even if we assume that eventually, nobody will click on phishing frauds anymore, fraudsters constantly work on new ways to deceive people. This makes it very difficult to prevent frauds.

A second problem with preventative methods is that there is a large and growing number channels. Each of these channels has its own weaknesses and vulnerabilities, but each channel must be protected or it will become the ”weakest link” that criminals choose to perpetrate their frauds. For example, Javelin Strategy & Research has found important weaknesses in the security of the still popular phone banking channel [10]. Currently, paying bills with smart phones and other mobile banking services is becoming popular [11,12]. With each new channel, however, new vulnerabilities and security challenges arise, hence increasing the complexity of protecting them all.

The final point to bear in mind is that 100% fraud prevention is generally not economical. A key difference between payment frauds and computer intrusions is that the cost of payment frauds is always known: It is the amount of money lost. As such, it becomes possible to conduct cost-benefit analyses. When the expected incremental savings from more fraud prevention are smaller than the cost of implementing the preventive measures, then further spending on fraud prevention is no longer economically justified. This argument obviously also limits the amount of money one would spend on fraud detection, but at least a basic level of fraud detection is always required to manage the residual risk of imperfect fraud prevention.

3 State of the Art in Payment Fraud Detection

Just like in intrusion detection methods, fraud detection methods can be classified into *knowledge-based* and *behavior-based* ones [13]. Knowledge-based methods use knowledge of past frauds to detect future instances of the same frauds. Behavior-based methods build a model of normal (i.e. non-fraudulent) payment activity and then detect deviations from this model. Commercial payment fraud detection systems use both techniques in combination. Examples of knowledge-based fraud detection rules that the author has seen in practice include the following:

- Alert on unusually large international payments that are made to beneficiaries that the payor had never sent money to before.
- Alert on payments to hot-listed accounts that have been associated with fraud.
- Alert on multiple payments from one account to another account if these payments occur within a given (short) time interval and in aggregate exceed a given dollar amount.
- Alert on address changes that are followed by large numbers of payments or payments over large dollar amounts.
- Alert on changes in the IP address, operating system, or user agent used for online banking.
- Alert when the same IP address is used to access several different accounts within a given time interval.
- Use automatically learned fraud detection rules that were generated using supervised learning techniques [14,15].

Behavior-based techniques compare observed payments to their ”expected values”, which are derived from historic payments. Various statistical methods have been used to derive values for expected payments, including link analysis (which identifies hidden relationships between entities), clustering (which groups payments based on similarity), and outlier detection [14,15].

Commercial fraud detection systems are extremely complex software packages that cannot be evaluated or compared simply. Fraud managers who have to select and implement payment fraud detection solutions are therefore recommended to take the following factors into consideration:

1. Quality of the case management tools that analysts use to evaluate fraud alerts;
2. Ease of integration with the core banking system of the financial institution;
3. Scalability to very large payment volumes (up to dozens of millions of payments per day);
4. Support for synchronous real-time fraud detection (as opposed to batch fraud detection);
5. Support in automating responses to frauds (e.g. by blocking payments);
6. Ability to anticipate frauds rather than detect them after the fact;
7. Detection rate at a given level of false positives;
8. Ability to detect special fraud classes such as multi-channel frauds;
9. Ability to adapt and evolve over time as fraudsters change their behavior and respond to new prevention or detection methods;
10. Total cost of ownership (including annual operating costs).

In the author's experience, today's payment fraud detection systems offer very mature solutions in the areas (1)-(5). Factors (6)-(9) are increasingly available in commercial products but offer more room for improvement. Factor (10), i.e. the total cost of ownership, has been the most important hurdle in the fraud detection projects that the author was involved with. As a very rough rule of thumb, a mid-sized payment fraud detection solution for several million transactions per day will cost $800,000 in software licensing fees, $800,000 in implementation costs, $150,000 in annual software service fees to the fraud detection vendor, and $100'000 in annual in-house labor for fraud investigations. That adds up to an upfront investment of $1,600,000 and running costs of $250,000. While these costs are not excessive, fraud managers compare them to recent fraud losses and may conclude that there simply is no business case.

There has been a significant amount of work on other types of fraud such as credit card fraud, money laundering, telecommunications fraud, insurance fraud, or accounting fraud. These topics are outside the scope of this paper, but very readable reviews of these topics can be found in [14,15,16].

4 Risk-Based Payment Fraud Detection

Motivated by the relatively high cost of commercial payment fraud detection systems, I now present a fraud detection method that is less powerful but cheaper than today's leading products. As such, it is a suitable solution for fraud managers that seek a sound payment fraud detection system but shy away from the high cost of many commercial packages.

The new fraud detection method is based on two observations: *Observation I* is that payment frauds generally become known fairly early (i.e. days or weeks) after they are committed because the fraud victims complain. *Observation II* is that the objective of fraud detection is *not* to detect fraudulent payment transactions. This just happens to be the common operationalization of fraud detection. In practice, however, the objective of fraud detection is to manage risk and to keep fraud-related losses acceptably small. Margalit and Fine, for example, beat commercial fraud detection systems by optimizing their system to detect the largest possible fraction of *monetary fraud losses*, rather than trying to detect as many *fraud instances* as possible [17]. This illustrates that fraud detection performance can be improved in a meaningful way by abandoning the conventional focus on the detection of individual fraudulent payments.

The first observation has important implications: The moment a financial institution learns about a payment fraud, they start investigating it. This investigation generally reveals some vulnerabilities in the information systems and those vulnerabilities are remediated instantaneously. As a consequence, the fraud stops working and the fraudster's ability to make money disappears. Fraudsters are aware of this and they adapt accordingly. Thus, after a new fraud has been discovered, fraudsters have an incentive to exploit it as fast and aggressively as possible to extract the maximum gain before the fraud is discovered and stops working. This leads to so-called ”Tsunami Frauds”, i.e. very aggressive frauds that lead to large losses in short periods of time. Many of these Tsunami frauds never get published; one of the few published Tsunami Frauds occurred in November 2008, when criminals stole $9 million in a few hours using cloned debit cards [18]. Tsunami frauds are a major threat to financial institutions.

In this respect, payment fraud is fundamentally different from classic intrusion detection where successful attacks – and the underlying vulnerabilities – can remain undetected for years. Payment frauds are almost always reported, investigated, and mitigated. So, while a software vulnerability can work for a very long time, payment frauds hardly ever do. John Boyd's warfare theory therefore also becomes the theoretical basis for winning the war of fraud [19,20]: Financial institutions must detect and respond to frauds *faster* than criminals can execute them, hence changing the environment so frauds no longer work. Conversely, criminals seek to outrun the defenders so they can steal as much money as possible before the environment changes and the frauds stop working. This property of observing, deciding, and acting faster than one's opponent is referred to as ”being inside the enemy’s decision loop”. It is the key to successful payment fraud management.

Observation II reminds us that financial institutions want to manage the risk of fraud-related losses. For any given time interval T, it is possible to estimate this risk rather accurately. Let $\mathbb{P}_T$ be the payments that occurred during interval T and let $p \in \mathbb{P}_T$ be a payment in $\mathbb{P}_T$. The risk R_f of fraud-related losses during interval T then follows as:

$$R_f = \sum_{p \in \mathbb{P}_T} amount(p) \times P_f(p) \tag{1}$$

with $amount(p)$ being the nominal dollar amount of the payment p and $P_f(p)$ being the probability that payment p is fraudulent. The value $amount(p)$ is known as it is contained in each payment transaction and the probability $P_f(p)$ can be estimated from historical payment data (see Section 4.1). For any sliding time window T, a financial institution can therefore calculate the payment fraud risk R_f that it is exposed to. If this value exceeds a threshold value *maxLoss* then all transactions in this time window must be blocked and investigated. Otherwise, the risk is contained and no action needs to be taken. Figure 2 summarizes the resulting fraud detection system.

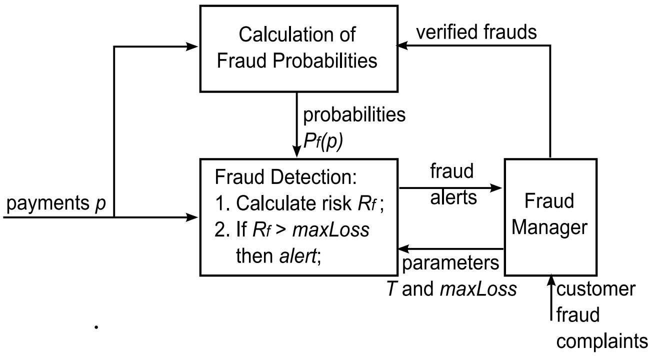

Fig. 2. Risk-based payment fraud detection system.

Let us consider the effect that this risk-based fraud detection system has on fraudsters. Basically, fraudsters are faced with two poor choices (from their perspective). If they want to be aggressive and monetize their frauds before they become known and ineffective then they will be flagged by the risk-based fraud detection system as soon as the risk R_f of fraud-related losses exceeds the threshold value. If, on the other hand side, the fraudster wants to ”stay below the radar” of the risk-based fraud detection system then he will succeed a few times before customers complain and the bank mitigates the vulnerability as explained above. In this case, the loss to the bank remains acceptably small, and the payoff to the fraudster may be too small to make such an approach viable. Section 5 evaluates the strengths and shortcomings of the risk-based fraud detection system more fully.

4.1 Estimation of the Probabilities $P_f(p)$

The probability $P_f(p)$ of payment p being a fraud can be expressed as:

$$\begin{aligned} P_f(p) &:= P[p \text{ is fraudulent}|p \text{ occurs}] && (2)\\ &= \#[p \text{ is fraudulent}]/\#[p \text{ occurs}] && (3) \end{aligned}$$

Equation (2) defines $P_f(p)$ as the conditional probability that payment p is fraudulent provided it occurs in the historic data. Equation (3) expresses this conditional probability as the number of times that, in the historic data, payment p was fraudulent divided by the number of times the payment occurred in the historic data. Equation (3) gives us a point estimate of the binomial fraud probability. The problem with this point estimate is that it is inaccurate (i.e. it has noticeable errors that are a result of the random sample used to calculate it) unless #[p occurs] is at least 250 and ideally larger [21]. This is a problem for infrequent payments p, and we next discuss how to deal with this.

A payment p is an n-tuples of attribute values such as the account numbers, payor and payee names, payment time and date, reference number, payment type, dollar amount etc. [22]. As such, payment tuples define a high-dimensional space that in parts, is too sparse for the estimator of Equation (3) to be accurate. A good way to deal with this problem is to map the payment tuples onto a lower-dimensional space where the mapped payments p' are "packed" more densely so the estimator can be calculated. Let $M()$ be the function that performs this mapping, i.e. $p' = M(p)$. We then evaluate Equation (3) in the mapped tuple space of p'. To further increase the amount of payment data available for the estimation process, it is advisable to use a full six months of historic payment data. Using older data is not recommended as it tends to be unrepresentative of current fraud patterns. (Six months is an experience value in the industry for the time it takes fraudsters to change their fraud patterns.)

The details of defining the mapping function $M()$ is the subject of another paper, but a few key points are worth making: First, $M()$ should eliminate redundant attributes such as "account number" and "name of account holder", which basically contain the same information. Second, $M()$ should eliminate attributes that are only marginally relevant for distinguishing normal payments from fraudulent ones. Automatically generated transaction numbers are an example of such irrelevant attributes. Third, $M()$ should construct additional attributes or modify existing ones so they become more indicative of frauds. For example, it is a good idea to replace the payee's account number by the type of payee (e.g. private person, business, type of business, etc.) and the payee's geographic location. Moreover, attributes of high cardinality should be mapped to attributes of lower cardinality. For example, the timestamp of a transaction should be mapped to discrete values such as ”morning”, ”afternoon”, ”evening”, and ”night”.

Calculating Equation (3) in the mapped tuple space p' increases the accuracy of estimates because the mapped space has fewer dimensions and the dimensions have lower cardinality. If this is still not enough to obtain an accurate estimate, then we set $P_f(p) := x$ if in the $M()$-mapped space, x percent of historic transactions occur more frequently than $M(p)$. For example, if $p_1 = M(p)$ and the historic transactions in the $M()$-mapped space are given by $\{p_1, p_1, p_2, p_2, p_2, p_3, p_3, p_3\}$, then $P_f(p) = 6/8$. This rule is a bit arbitrary, but it captures the intuition that payment p is likely to be a fraud if it is a type of payment that the payor rarely ever submits.

5 Discussion

Evaluating fraud or intrusion detection methods is in many ways very difficult. Synthetically generated test data tends to be biased [23], and evaluations in real-world deployments are subject to environmental factors such as the types of frauds that (coincidentally) occurred during the evaluation period. Moreover, real-world experiments only offer individual generally highly anonymized (due to the sensitivity of fraud data) and non-replicable data points that have limited generalizability. Also, while evaluations of detection methods tend to rely on the Receiver Operating Characteristic (ROC) curve, this is insufficient from a practical point of view. Specifically, Section 3 alone listed ten factors that affect the quality of fraud detection systems. Ideally, these factors should be reflected in any evaluation of a fraud detection system. To deal with these difficulties in a pragmatic way, I will assess the strengths and weaknesses of the proposed fraud detection system and I will offer evidence that in practice, it is ”useful”.

The claim of the presented risk-based fraud detection system is that it offers good fraud detection performance at a low cost. The low cost is a result of its simplicity: The system only requires two parameters to be set (T and $maxLoss$) and it requires no maintenance or upkeep. Most commercial payment fraud detection systems, by contrast, require regular updates in order to adapt to the changing behavior of fraudsters.

The practical usefulness of the fraud detection system is evidenced by a fraud case that the author was involved with. A major financial institution had incurred significant losses because of massive cash withdrawals at ATMs in a different geography. The cash withdrawals were performed using cloned debit cards and the entire fraud played out within a few days. The size and frequency of cash withdrawals combined with the fact that they occurred in a geography in which the debit cards had never previously been used would have triggered the risk threshold and flagged the fraud (the system was not used at the time of the fraud). As a further proof point, Figure 3 shows the stylized and anonymized fraud frequencies and monetary fraud losses at a typical financial institution over a two year time period. The figure shows very clearly the ebb and flow of fraud Tsunamis that our risk-based system has been optimized to detect. Again, the practical usefulness of the proposed risk-based fraud detection system is apparent.

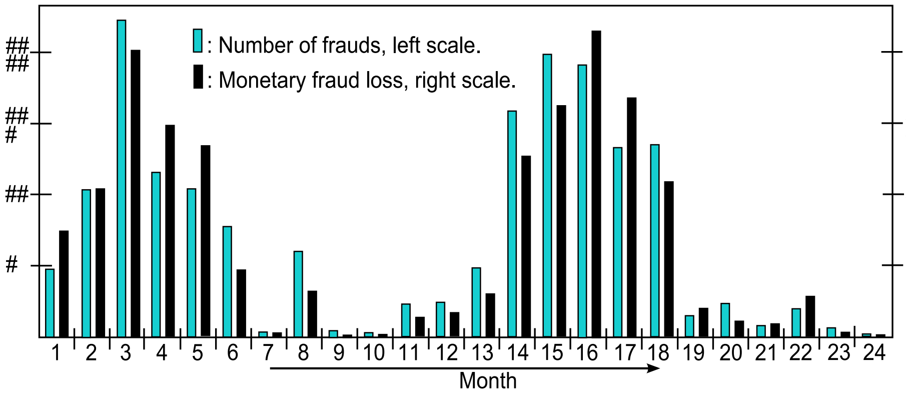

Fig. 3. Fraud frequencies and fraud losses at a typical financial institution.

Having discussed the strengths of the proposed fraud detection system, I next consider its weaknesses. To begin with, even though I have not done a one-on-one comparison with commercial fraud detection systems, I fully expect that commercial systems detect more frauds. Moreover – just like any other fraud detection system – the system presented in this paper will miss certain frauds (e.g. certain slow and low-profile frauds). That is OK because this paper's objective was not to build the highest-performing fraud detection system. Rather, the objective was to build a system that is suitable for financial institutions that have low fraud losses, consequently shy away from the relatively high cost of many commercial systems, and yet, want to implement a sound fraud detection solution.

A relevant limitation of the presented risk-based system is that it operates in near-real-time. This is because the system has to observe enough payments before it can raise an alert. Therefore, frauds are detected in near-real-time with a delay of up to T time units. This delayed detection also implies that real-time fraud response is only possible within limits. To understand why, just note that by the time a fraud alert has been triggered some of the payments in the suspicious time interval T may have moved on to a another processing stage where they can no longer be stopped. The risk-based method is also prone to false positives and false negatives in case the fraud probabilities are inaccurate. This can happen due to estimation errors in Equation (3) or because fraud probabilities are calculated from historic payments that are no longer representative of current fraud probabilities [14].

It is also worth noting that the presented risk-based system constitutes a natural extension of the classic banking practice of screening high-dollar-value payments for fraud [24]: This paper's risk-based system weights the dollar-values of payments with their probability of being fraudulent and puts them into the context of other transactions within the same time window.

In summary, the presented risk-based method is a cost-effective and practically ”useful” way to detect frauds that exceed an institution's risk limit. It is best suited where near-real-time detection is sufficient.

6 Summary and Conclusion

This paper defines payment fraud as the use of deceptive methods to transfer money against the payor's will to an illegitimate beneficiary's account. Frauds enter financial institutions through various channels (checks, telephone, online, debit card payments, etc.) and fraudsters use various methods (impersonation, deception, and server-side attacks) in an attempt to appropriate other parties' funds. The paper reviews these channels and fraud methods and concludes that given their multitude, it is not possible in practice to prevent all frauds. This emphasizes the importance of fraud detection.

Next the paper reviews the technologies used by today's state of the art payment fraud detection systems. One key point here is that even though the detection performance of a fraud detection system is important, this is only one of many factors that fraud managers need to consider when choosing a system for their institutions. The paper specifically lists ten factors that affect purchasing decisions and points out that the relatively high cost of fraud detection systems is still a factor that slows down the adoption of payment fraud detection systems in practice.

Motivated by the need for lower-cost fraud detection methods, the paper presents a new risk-based method for detecting payment frauds. This method exploits two key observations: Firstly, payment frauds generally become known fairly early after they are committed because the fraud victims file complaints; secondly, the objective of fraud detection is not to detect fraudulent payment transactions but rather to manage risk and to keep fraud-related losses acceptably small. These observations form the basis for the new fraud detection method presented in this paper. This method quantifies the expected loss from frauds over a given time period and raises an alert if this risk exceeds a threshold value.

The presented risk-based fraud detection method works because it forces fraudsters to execute their frauds slowly if they want to remain undetected. As a consequence, financial institutions incur small but acceptable losses; eventually, however, they will stop the fraud because customers will detect and report the fraud. The impact on fraudsters is that frauds become less profitable, potentially to the point where they are no longer worthwhile.

The paper concludes with a critical appraisal of the risk-based fraud detection method. Most notably, it is emphasized that the method has a low total cost of ownership, that it is ”useful” (but by no means infallible), and that it is most suitable where real-time detection and fraud response are not required. As other fraud detection methods, the risk-based method of this paper is prone to false positives and false negatives if estimation errors occur or when historic probabilities are unrepresentative of present fraud patterns.

References

1. Editors of The American Heritage Dictionaries: The American Heritage Dictionary. 2nd edn. Houghton Mifflin (1982)
2. Humphrey, D.B.: Payment Systems: Principles, Practice, and Improvements. World Bank Publications (1995)
3. Grant, N.: ACH 102: Who We Are and What We Do (April 2008) `http://www.electran.org/docs/annual_meeting/2008/presentations/ACH_vs_Check21-Grant.pdf`, last accessed April 4, 2010.
4. Sausner, R.: Reconnaissance Leads To Multi-Channel Fraud. Bank Technology News **21**(10) (2008) 20–20
5. Weigold, T., Kramp, T., Baentsch, M.: Remote Client Authentication. IEEE Security $ Privacy **6**(4) (2008) 36–43
6. Anonymous: Payments Fraud. How it Happens. And What You Can o to Protect Your Organization. Technical report, J.P. Morgan Treasury Services (2009)
7. Tozzi, J.: Identity Theft: The 'Business Bust-Out'. BusinessWeek Online (July 2007)
8. Weiser, B.: Nigerian Accused in Scheme to Swindle Citibank . The New York Times (2009)
9. Hall, J.A.: Accounting Information Systems. 6th edn. South-Western College Publisher (2008)
10. Kim, R.: Telephone Banking Authentication: Practical Approaches to Securing a Popular yet Vulnerable Channel. Technical report, Javelin Strategy and Research (2007)
11. Vyas, C.: From Niche Play to Mainstream Delivery Channel: US Mobile Banking Forecast, 200813. Technical report, Tower Group (2009)
12. Wu, C., Zafar, S., Cloninger, J.: 2008 Mobile Financial Services Study. Technical report, Edgar, Dunn & Company (2008)
13. Debar, H., Dacier, M., Wespi, A.: Towards a taxonomy of intrusion-detection systems. Computer Networks: The International Journal of Computer and Telecommunications Networking **31**(9) (1999) 805–822
14. Bolton, R.J., Hand, D.J.: Statistical Fraud Detection: A Review. Statistical Science **17**(3) (2002) 235–255
15. Phua, C., Lee, V., Smith, K., Gayler, R.: A Comprehensive Survey of Data Mining-Based Fraud Detection Research. Artificial Intelligence Review (2005)
16. Singleton, T., Singleton, A., Bologna, J., Lindquist, R.: Fraud Auditing and Forensic Accounting. 3rd edn. John Wiley & Sons (2006)
17. Margalit, O., Fine, S.: Detecting Fraud in Financial Transactions. IBM Brochure (2009)
18. Poulsen, K.: Global ATM Caper Nets Hackers $9 Million in One Day (2009) `http://www.wired.com/threatlevel/2009/02/atm/`.
19. Coram, R.: Boyd: The Fighter Pilot Who Changed the Art of War. Back Bay Books (2004)
20. Hammond, G.T.: The Mind of War: John Boyd and American Security. Smithsonian Books (2004)
21. Lewis, E.E.: Introduction to Reliability Engineering. John Wiley & Sons (1996)
22. Westphal, C.: Data Mining for Intelligence, Fraud & Criminal Detection. CRC Press (2008)
23. John McHugh: Testing Intrusion Detection Systems: A Critique of the 1998 and 1999 DARPA Intrusion Detection System Evaluations as Performed by Lincoln Laboratory. ACM Transactions on Information and System Security (TISSEC) **3**(4) (2000) 262–294
24. Khan, I.: Recognizing and Managing Payment Fraud. J.P. Morgan Expos 2008 (2008) `https://host5.agsdc.net/jpmchase/files/attachments/PaymentFraud_Khan.pdf`.