

Fraud Prevention using Automated Audit Systems: A Strategic Imperative

Jamil Mulani¹ Zhang Xianzhi^{2*}

1. Doctoral Student, School of Accountancy, Dongbei University of Finance & Economics, Dalian Liaoning, China, 116025
2. Professor – School of Accountancy, Dongbei University of Finance & Economics, Dalian, Liaoning China, 116025

*Corresponding Author: Email: zxz@dufe.edu.cn

Abstract

This study examines a mathematical model to determine the timing and consequently volume of transactions to be audited in a continuous audit system to detect potentially fraudulent transactions. The interactions between the audit system and a potential fraudster are modeled as Continuous Time Markov Chain and the transition probabilities from one state to another are determined using game theoretic approach. We believe that such a model has the potential to be deployed in an audit system to detect potentially fraudulent or malicious transactions. In this research, the information system is modeled as a continuous time Markov chain (CTMC), where the transition from one state to another occurs due to actions of a person with malicious intent. The present state of the system depends only on the past state. At each state, the fraudster can either continue with the next step in the fraud or can cease and desist from the fraud. The interaction between the actions of the audit module and fraudster is modeled as a two-player simultaneous zero-sum game and the probability of transition from one state to another is derived from the payoff table. This payoff will be decided by the outcome of a game theoretic model. A sensitivity analysis showed that when an organization has strong anti-fraud controls, the probability of fraud decreases and the need for frequent audit decreases. The limitations of the model are that, the game theory model assumes a zero-sum game where the payoffs are known and certain.

Keywords: Continuous Audit Systems, Game Theory, Audit Timing, Fraud prevention, Sensitivity analysis.

1. Introduction

During the last decade, continuous audit systems (CAS) have become widely used in large transaction processing systems (PricewaterhouseCoopers, 2006; 2009). These systems are defined as “a continuous (concurrent) monitoring technique in which a set of programmed modules are directly integrated into the organization's application system” (Groomer and Murthy, 2003).

In a manual audit, the scope and scale of the audit procedures is limited by the amount of physical data that an auditor can process. CAS can process vastly more data than what can be handled by manual means. However, this increase in scale and scope of data poses its own set of problems. If the CAS scans too much data it will find a lot of spurious items of interest, or false positives. Scan too little data and the CAS finds only trivial issues that would just as easily have been detected by manual means. This key issue regarding the timing and volume of transactions to be examined is highlighted in the problem statement below.

Consider the problem of access controls in a database. Operating system objects, such as files, are unrelated items, whereas records, fields, and elements of a database are related. Although a user cannot determine the contents of one file by reading others, a user might be able to determine one data element just by reading others. The ability to determine data values from other related data values is called inference (Weber, 1998).

This could happen in two ways. One is by using “Context Dependent” queries. Using diverse set of queries, a user may be able to infer data that would not normally be available to him. In order to secure against such improper access, access is restricted to all contextual references to data with higher security. For example, even though a user may not have access to income data of customers, he may query the database and obtain a list of customer making high value purchases, which allows him to infer the income of the customer. The second way is by using “History Dependent” queries. Here a user may construct a series of queries in such a way that each query and its answer does not violate the security policy. But all the answers of these queries together would give her information, for which she is not authorized (Weber, 1998).

A similar problem arises when a malicious user attempts a fraud in an accounting system. Usually, in order to perpetrate the fraud, the user must succeed in creating a series of false entries, over a period of time that ultimately results in financial loss to an organization. Each fraudulent entry in itself does not necessarily cause a loss, and may not trigger an alert from the CAS in the early stages. However, at a later stage, by analyzing the series of transactions entered by that user, the audit module is in a better position to flag a transaction as a potentially fraudulent.

To prevent these leakages of information or loss, the CAS must keep track of the queries or transactions that each user has run in the past. However, the key question is how far back the audit module has to scan to

determine if an attempt has been made to violate information security. Or to put it differently, how much and how soon must the audit module scan the history of a user's actions to determine potential fraud.

In this paper, we use game theoretic and stochastic modeling tools to determine the length of time for which the audit module must scan the past history of transactions. To determine this, one needs a metric which tells the module the amount of time it takes to violate an information system. We model the interactions between the user ("the fraudster") who wishes to obtain data to which she is not privileged and the audit module using a stochastic model, where the decision to commit the fraud is based on the payoffs determined by a game theory model.

2. Literature Review

Reliability of a system is usually measured by estimating how long a system can operate without failing, usually termed as mean time to failure, or mean time between failures. This enables one to measure reliability in terms of a behavior that can be perceived by the user. Users are more interested in the measures of reliability in terms of failure rate rather than knowing it in terms of its static features, properties, or mode of construction (Littlewood et al., 1993). Though such static features, properties or mode of construction do influence reliability, they are not adequate in themselves to determine reliability. In a similar vein, feature set, design philosophy of firewalls, intrusion detection systems and access controls systems are not enough to determine whether a particular system is capable of maintaining integrity or confidentiality of data. What is required is a measure of actual achieved information integrity, analogous to how mean time to failure is used, as a yardstick for measuring reliability (Littlewood et al., 1993).

One of the earliest attempts to quantify security was taken by Littlewood et al., (1993), who commented on the scarcity of literature on the probabilistic treatment of system security. Littlewood et al., (1993) drew an analogy between the problem of reliability and problem of security, and elucidated on the similarities and differences. They pointed out that the key difference between reliability and security is the presence of malicious threats in the security domain.

Another early attempt at quantification of information security metrics was undertaken by Ortalo, Deswarte, and Kaâniche (1999). They attempted to provide a quantitative assessment system security level at a point in time. They modeled system vulnerabilities as a privilege graph, where each arc in the graph is assigned a weight that estimates the degree of effort required by a potential attacker to make a move from a less privileged state to another with greater privileges. They setup an experimental system to capture data on actual system vulnerabilities and attempts to circumvent controls.

These early researchers concentrated on measuring static measures that measure actual security at a point in time. They fail to take into account the fact that once a vulnerability has been identified, administrators take steps to fix that vulnerability and concurrently, attackers modify their behavior and seek other vulnerabilities. To address this dynamic nature of the security modeling, researchers have resorted to game theory.

The use of game theory to model information security issues is fairly recent. Hamilton, Miller, and Saydjari (2002) pointed out that it is possible to formulate attack and defense scenarios as a game of moves and counter moves. Game theory already has algorithms to predict the likelihood an action would be selected in such scenarios. Alpcan and Başar (2003) pointed out that game theory provides a rich set of tools for modeling and analyzing information security issues. These game theoretic tools can also be used to develop practical and cost effective solutions that can be implemented in the real world. Alpcan and Başar (2003) modeled an Intrusion Detection System (IDS) as a network of sensors. They modeled the behavior of an attacker and the IDS as a two-person, nonzero-sum, noncooperative game. They applied game theoretic concepts to develop a formal decision and control framework for a platform-independent IDS.

Lye and Wing (2005) constructed a two-player stochastic game to represent the interaction between a system administrator and an attacker. They computed the Nash equilibria or best response strategies for the attacker and administrator. They used a two-player game model and pointed out that a team of attackers or team of administrators can be modeled as a single omnipresent attacker or defender. Thus, a two player game is sufficient for the modeling problems of information security.

Liu, Zang, and Yu (2005) used a game-theoretic approach to inferring attacker intent, objectives, and strategies (AIOS) and captured the inherent interdependency between AIOS and defender objectives. The authors were primarily interested in examining the characteristics of attackers rather than attacks themselves. The attacker's intent is modeled using game theoretic models. They make the case that the attacker's intent can be inferred from the nature of the attack and this in turn influences the nature of defense strategy to be adopted.

Cavusoglu et al. (2005; 2008) demonstrated that firms incur lower costs when they use the game theory as opposed to the decision theory. According to them, the firm's payoff is maximized when the firm credibly commits and communicates its strategy to the hacker. Even if the communication of strategy is not credible, the firm enjoys a higher payoff if the firm and attacker play a sequential game. Their research focused on the overall economics of investment in information security.

Sallhammar and Knapkog (2004) built upon the work of Alpcan and Başar, 2003; Liu, Zang, and Yu, 2005; Lye

and Wing, 2005 and modeled a security breach as a series of state changes from an initially secure state to compromised state. They used game theory to model and compute the probabilities of malicious user actions.

Loss of integrity implies that purity of an information system has been compromised in some way. This loss of integrity could be accidental or malicious. While accidental events are usually not recurring, once a malicious event occurs, the perpetrator is usually emboldened to carry out further attacks. These events carry huge financial and reputational risks for the organization concerned. In a state space model, these malicious and intentional events can viewed as a series of changes from an initially uncompromised state to a state that represents a severe loss of integrity.

In order to successfully carry out a history dependent attack, the fraudster must often carry out a series of actions that each breaches the integrity of the system, but does not necessarily cause a monetary loss. For example, in order to obtain fraudulent disbursement, a fraudulent invoice, payroll slip, expense report or another document must be created. Some of these may require additional documents. In order to support a fraudulent pay slip, it may be necessary to substantiate it with a forged timecard.

In this paper, we are not concerned with the actual nature of fraud. Any action that undermines the integrity and/or confidentiality of the system would eventually cause a loss, if undetected would constitute a fraud. Given a measure of how long it takes for a fraudster to compromise the system, it would be easy to configure a CAS to scan that length of time and corresponding volume of transactions to detect potential frauds. If it scans for a shorter period of time, it is likely to miss the fraud. We use the approach of Sallhammar and Knapskog (2004) and demonstrate how their model can be used to determine the timing of audit in an continuous audit system. From a programming standpoint, the audit modules are subroutines that are executed when certain criteria set by the auditor are met. In this environment, the auditor needs to specify the extent and timing of transactions that will be subject to audit once a trigger event occurs. A trigger event could be, for example, queries or transactions from a particular user, which is far in excess of the usual number of transactions from that user based on historical usage patterns. While Sallhammar and Knapskog (2004) proposed a model to quantify security, in this paper we look at how the model can be adapted to improve the effectiveness of continuous audit system. In particular, this paper attempts to determine that extent of transactions to scan in order to detect a potentially malicious attempt.

The problem of determining audit timing and volume has been studied analytically by two other researchers, Hughes (1977), and Morey and Dittman (1986). Hughes (1977) modeled the audit timing problem as an infinite state time varying Markov decision process and solved it using dynamic programming. Morey and Dittman (1986) attempted to determine optimal elapsed time between audits or the number of financial transactions that should occur before an audit should be initiated. They used two models, one for protecting against overstatement errors and the other for both overstatement and understatement. They provided a closed form solution technique for determining the optimal timing of audit. Both these researchers did not consider the possibility of deliberate fraud. Instead, they assumed that errors occur at a finite rate and that these errors are caused by non-malicious factors.

The key contribution of current study is the determination of audit timing and volume in an continuous audit system. In traditional audits, the timing and periodicity is usually determined by statute (in most countries, statutory audits are annual and internal audits are usually conducted on a quarterly basis). These strategies are inadequate in an CAS. To the best of our knowledge, this paper represents one of the first attempts to determine the timing and frequency of audits through game theoretic methods.

3. Mathematical model and empirical test

In the state transition diagram as shown in figure 1 (appendix), the present state of the system depends only on the past state. And the time spent in any given state is stochastic. In case the fraudster fails at an intermediate state (state i) and resigns, the system continues to remain in a compromised state until the audit system restores it back to a pristine state.

An information system's fraud-free operating time is modeled as having an exponential distribution. The exponential distribution is widely used in modeling system reliability (Littlewood et al., 1993; Singh and Chander, 2008) due to its memory less property. Experimental data from an intrusion detection experiment also supports the view that an exponential distribution can be applied to modeling information security (Goseva-Popstojanova et al., 2001). The time between any two successive fraudulent entries is modeled as an exponential distribution (whose cumulative distribution function is given by $P(t)$) with parameter λ_{fraud}

(where λ_{fraud} is the average rate at which fraudulent transactions are initiated). Thus,

$$P(t) = 1 - e^{-\lambda_{\text{fraud}} t}$$

Once a fraudulent transaction is initiated, the system enters state $i = 1$. The time for the next fraudulent entry

$(1/\lambda_i)$ and the time needed for the audit subsystem to detect $(1/\mu_i)$ and restore the system to the pristine state 0 are also modeled by exponential distributions.

$$P_{\text{fraud}(i)}(t) = 1 - e^{-\lambda_i t} \quad \text{and} \quad P_{\text{audit}(i)}(t) = 1 - e^{-\mu_i t}, \text{ respectively.}$$

$P_{\text{fraud}(i)}(t)$ and $P_{\text{audit}(i)}(t)$ are the cumulative distribution function of fraud process and the audit process

respectively. Thus, $\frac{1}{\lambda_i}$ and $\frac{1}{\mu_i}$ will be the respective mean time that a fraudster and the system spend in state i , of

the model. State i represent an intermediate situation, where the system has been compromised by the entry of fraudulent transactions, but no monetary loss has occurred. The two processes, λ_i and μ_i , representing the

fraudster and the audit process can be merged into one Poisson process. Now due to the memoryless property of the exponential distribution, the state transition model can be transformed into a continuous time Markov chain (CTMC) with discrete space (Ross, 2006), formally described as,

$$\{X(t): t \geq 0\}, X_s = \{0, 1, 2, \dots, n\},$$

where, the subscript s refers to the state the system is in. Thus, if $X(t)=2$, it implies that the system is in state 2 at time t . The model as described here is shown in Figure 2 in appendix.

At each intermediate state i , the fraudster has two possible courses of action. She can either continue with the next step in the fraud or she can cease and desist from the fraud. The probability of continuing with the fraud is $p_i(f)$, and the probability that she would cease is $1 - p_i(f)$. At each stage i , the probability that she continue

with the fraud has to be computed, since at each stage she will evaluate the payoff of continuing with the fraud. Also, in order for the fraud to have financial repercussions, at each step the fraudster must succeed. The probability of success in each case must also be computed and this is indicated by $p_i(s)$ and has been

incorporated into the model.

The fraudsters and the audit module's action rates and the transition probabilities allow the computation of the instantaneous transition rates between state i and state $i+1$:

$$q_{i,i+1} = p_i(f)p_i(s) \cdot \lambda_i \text{ and the state between state } i \text{ and state } 0 \text{ as } v_{i,0} = \mu_i \quad \forall i = 1, \dots, n.$$

Since the stochastic model is a continuous time Markov chain, it is possible to compute the limiting probabilities of each state in the model. The transition probabilities are also shown in Figure 2. In order for Markov chain to have a steady state solution, the rate at which the transition to state i occurs should match the rate of transitions out of state i . This results in a series of balance equations. Solving these equations, one can obtain the stationary probability of being in fraudulent state, P_n .

$$\begin{cases} P_0 \cdot \lambda_{\text{fraud}} = P_1 (q_{12} + v_{10}) \\ P_0 \cdot q_{12} = P_2 (q_{23} + v_{20}) \\ \vdots \\ P_{n-1} \cdot q_{n-1,n} = P_n \cdot v_{n0} \\ \sum_{i=1, \dots, n} P_i = 1 \end{cases}$$

With this information one can compute the mean time taken for the fraudster to commit the first fraud with the following equation.

$$\text{Mean time to first loss from fraud} = \frac{1 - P_n}{P_n \cdot v_{n0}}.$$

For a system that starts in an initially secure state, this metric provides a quantitative measure of the system's ability to maintain integrity.

To illustrate, consider a transaction processing environment, where transactions are examined to ensure compliance with business rules, and are concurrently checked for abnormalities to ensure that they are not

erroneous. A transaction that is compliant with normal business rules may still be abnormal. For example, assume that a system checks the orders entered by salespersons to ensure that they are not erroneous or fraudulent. The system could compare the order entered with previous history of orders entered by a particular salesperson to ensure that orders are free from error or fraud. Suppose, at a particular point in time, the system detects that an order entry by a salesperson is four times the average order for that salesperson. Such transactions may need further validation or investigation. In the above scenario, a single abnormal transaction may not constitute an attempt at fraud. The CAS may need to scan several similar transactions before it can assess the probability that the transaction may be fraudulent. The above metric provides an estimate of how much time must elapse for someone to commit fraud on an initially pristine system. This time estimate, in turn, provides an estimate of the volume of transactions that must be scanned to have a reasonable chance of detecting the fraud.

One issue to be addressed in the above model is the probability of continuing with the fraud at each state. One of the ways is to run an experiment on a real system and observe the actual actions of persons (Ortalo, Deswarte, and Kaâniche, 1999). This is an extremely resource-intensive approach. Another less resource intensive way is the use of game theory to estimate this probability, $p_i(f)$. One can view each transition in the stochastic model as a game, where a fraudster's choice of action is based on his utility derived from a rational considerations of benefits and costs of the fraud. The actions of the audit module and the fraudster can be modeled as a two player simultaneous game. In each state the fraudster has two possible courses of action, either to continue the fraud or withdraw.

The fraudster and the audit module have two actions each:

f_i , represents the action that changes the state of the system from i to $i+1$. This is when a fraudster upon

discovering a weakness in the control system, initiates or continues a fraud by entering a fraudulent transaction.

r_i , represents the action, wherein the fraudster decides to resign and desists from carrying out any other actions.

d_i , represents the action where the audit module detects the fraud.

Φ_i , represents the action (or situation) where the audit module fails to detect the fraud.

A game model can be constructed for each state i in the stochastic model. The game model $G(i)$ would be defined by

$N(\text{number of players}) = \{1, 2\} = \{\text{fraudster, audit module}\}$

$A_i(\text{Action Set}) = \{f_i, r_i, d_i, \Phi_i\}$, and

$\gamma_i = \{\gamma_{i1}, \gamma_{i2}, \gamma_{i3}, \gamma_{i4}\}$ is the payoff received by the fraudster for each possible combination of action and

response by the audit module. Here, we are assuming a zero sum game, where the payoff received by the fraudster is exactly equal loss suffered by the audit system.

$$\gamma_i = \begin{array}{|c|c|c|} \hline & d_i & \Phi_i \\ \hline f_i & \gamma_{i1} & \gamma_{i2} \\ \hline r_i & \gamma_{i3} & \gamma_{i4} \\ \hline \end{array}$$

In the above case, the fraudster's expected payoff for a course of action is,

$$v_i(f_i) = p(d_i) \cdot \gamma_{i1} + p(\Phi_i) \cdot \gamma_{i2}$$

$$v_i(r_i) = p(d_i) \cdot \gamma_{i3} + p(\Phi_i) \cdot \gamma_{i4}$$

Given that a fraudster is unlikely to know the exact probability of being detected, he can assume that the audit module is rational and seeks to minimize the expected payoff of the fraudster. If one assumes a zero sum game, where one player gain is another's loss, the minimax solution of the game is the Nash Equilibrium of the game (Gibbons, 1992).

To summarize, we have a stochastic model that depicts a fraud as a series of state transitions from a initially un compromised state to a compromised state through one or more intermediate states. At each state the fraudster has to evaluate the risk vs. reward trade off of continuing with the fraud or desisting from continuing with the fraud. The game theory model determines this risk vs. reward tradeoff and provides an input to the stochastic model. The stochastic model outputs the time required to commit the fraud, assuming the system starts of in clean state. The estimate of time indirectly provides an estimate of the volume of transactions to be examined.

4. Model testing

Consider a scenario wherein an employee creates a fraudulent invoice and submits it to his employer and the employer makes a payment based on the fraudulent invoice. In order to successfully carry out the fraud, the perpetrator must have generated a fraudulent purchase order and a fraudulent goods received note, before the

fraudulent invoice is paid. In this scenario,

Initially, the system has no fraudulent entries and is in a pristine state. We call this State 0.

The fraudster realizes that the system has vulnerabilities and decides to commit fraud (State 1).

She initiates the fraud by first entering a fraudulent purchase order, circumventing any controls. The system enters State 2, an intermediate state.

In case the fraudster's transaction continues undetected, she finally enters the fraudulent invoice which later gets paid in due course. The fraud would then be complete.

At each stage(except state 0 and 1), the audit module will have an opportunity to detect the fraud and take corrective action to rollback the transaction and restore the integrity of the system. In a database, in order to prevent an unwarranted inference only the last query has to be prevented but the audit module must remember to scan an unspecified number of previous transactions.

The stochastic model for this is depicted in Figure 3. To make the exposition easier, the fraudster's payoff function at each state is assumed to be as given below.

$$\gamma_i =$$

	<i>detected</i>	<i>undetected</i>
<i>Fraud</i>	-1	2
<i>Resign</i>	0	-1

This implies that if the fraudster enters a fraudulent transaction and remains undetected she will receive a positive payoff ($\gamma_{i2} = 2$). In case the fraud is detected or if the fraudster resigns, the payoff would be negative. If she gives up after being detected, the payoff is zero.

The minmax strategy which will provide the highest payoff to the fraudster is given by,

$$p_i(f) = 0.25 \quad \forall i = 1, \dots, n.$$

The other numerical values are as given as given below:

- The probability of fraud ($1/\lambda_{\text{fraud}}$) is 1 in 1,000,000. Note that since we are using Poisson distribution, the probability that a single event will occur during the time interval is proportional to the size of the interval.
- Probability of committing the first step in a fraud ($1/\lambda_1$) is 1 in 20,000.
- Probability of committing the second step in the fraud ($1/\lambda_2$) is 1 in 20,000.
- Probability of that the first step ($p_1(s)$) will succeed is 90%.
- Probability of that the second step ($p_2(s)$) will succeed is 70%.
- Probability of audit module detecting the fraud in the first stage ($1/\mu_2$) is 1 in 80,000.
- Probability of audit module detecting the fraud in the second stage ($1/\mu_3$) is 1 in 4,000.

By inserting the values from the above table into the following set of equations, we can calculate the mean time before the first loss from fraud occurs.

$$\begin{cases} P_0 \cdot \lambda_{\text{fraud}} = P_1 \cdot (p_1(f)p_1(s) \cdot \lambda_1) \\ P_1 \cdot (p_1(f)p_1(s) \cdot \lambda_1) = P_2 \cdot (p_2(f)p_2(s) \cdot \lambda_2 + \mu_2) \\ P_2 \cdot (p_2(f)p_2(s) \cdot \lambda_2) = P_3 \cdot \mu_3 \\ P_0 + P_1 + P_2 + P_3 = 1 \end{cases}$$

$$\text{Mean time to first loss from fraud} = \frac{1 - P_3}{P_3 \cdot \mu_3} = 17.07 \text{ days.}$$

We can now answer the question that is posed in the problem statement. Give the assumptions above the CAS will need to examine the data or transactions relating to the last 17 days in order to detect the loss of confidentiality or integrity.

In order to examine the internal validity of the model, we performed a simulation by varying key parameters and examining the results. The first parameter, λ_{fraud} , is the initial probability of fraud and the second parameter,

$p_i(f)$, is the probability of continuing with the fraud at each stage, which is the output from the game theory

model. The results are presented in Table 1.

We see from Table 1, that when the initial probability of fraud increases from 1 in 1,000,000 to 1 in 100, the audit interval decreases from 17 days to 13 days. The increase in the audit frequency corresponds with the increased probability of fraud, but it is not a linear relationship. While the probability of fraud increases 100-fold, the increase in audit frequency is only about 30%. Similarly, when $p_i(f)$, the probability of continuing with the

fraud at each stage decreases, there is a corresponding decrease in audit frequency. Again the relationship is non-linear. For example, when the probability decreases from 20% to 10%, the decrease in audit frequency is of the order of 17%. Note that when the audit frequency decreases, the volume of transactions that are to be audited also correspondingly decreases.

Note that any change in the payoffs only affect $p_i(f)$. If the payoff from the fraud were such that the

reward from fraud outweighed the risk from detection, then the probability of continuing with the fraud goes up and this in turn reduces the audit frequency.

We also look at the effect of varying the probability of success at each stage. The probability of success in the first stage is $p_1(s)$ and the probability of success in the second stage is $p_2(s)$. The effect of varying these

probabilities from 90% to 10% independently, while keeping the other variables constant, is shown in Table 2.

It can be seen that as the probability of success reduces, the audit frequency decreases. This makes intuitive sense as when an organization has strong anti-fraud controls, the probability of fraud decreases and this implies that the need for frequent audit decreases. Again note that when the frequency of audit decreases, the volume of transactions to be audited during each audit goes up.

5. Conclusions

In current study, we present a simple model that probes the expected behavior of a fraudster in an accounting or transaction processing system. We have used a stochastic model to look at the behavior of the system and used game theory for modeling the expected behavior of a fraudster. This allows us to compute how far back a CAS must scan transactions in order to detect a potential breach of system security. A pressing concern while applying traditional external audit techniques to large transaction systems is that the audit data is gathered long after the economic events are recorded and is too late to prevent corrective action. In online transaction systems, even when a detailed audit trail exists, examining the data after the fact does not prevent losses from occurring. This necessitates the evaluation of controls and data in real time or as close to real time as possible. Such systems are called continuous audit systems. There is a dire need to determine the volume, timing and frequency of the audit in continuous audit systems as it can materially affect the conclusions drawn from the data. Given that a fraudster is unlikely to know the exact probability of being detected and can assume that the audit module is rational and seeks to maximize audit systems payoff, which is equivalent to minimizing the expected payoff of the fraudster in a zero sum game. In this scenario, the minimax solution of the game is the Nash Equilibrium of the game.

This work presents some limitations as well. We have assumed one single profile for a potential perpetrator of fraud. However, realistically, there are many kinds of perpetrators of fraud with different skills, motives and payoffs. An extension of this work would be extrapolating this model to include more than one type of threat to integrity. A further limitation of this model is that we are assuming that all information necessary for the analysis is known. This may not hold in practice. However, this limitation could be addressed in future studies by looking at game theoretic tools which can incorporate incomplete information as well.

References

1. Alpcan, T., & Başar, T. (2003). A game theoretic approach to decision and analysis in network intrusion detection. In Proceedings. 42nd IEEE Conference on Decision and Control, 2003 (Vol. 3, pp. 2595-2600).
2. Cavusoglu, H., Mishra, B., & Raghunathan, S. (2005). The value of intrusion detection systems in information technology security architecture. *Information Systems Research*, 16(1), 28-46.
3. Cavusoglu, H., Raghunathan, S., & Yue, W. (2008). Decision-theoretic and game-theoretic approaches to IT

- security investment. *Journal of Management Information Systems*, 25(2), 281-304.
4. Gibbons, R. *Game Theory for Applied Economists*. Princeton University Press, Princeton, New Jersey, 1992.
 5. Goseva-Popstojanova, K., Vaidyanathan, K., Trivedi, K., Wang, F., Wang, R., Gong, F., & Muthusamy, B. (2001). Characterizing Intrusion Tolerant Systems Using A State Transition Model. In *DARPA Information Survivability Conference and Exposition (DISCEX II'01) Volume II-Volume 2*. Presented at the DARPA Information Survivability Conference and Exposition.
 6. Groomer, S., & Murthy, U. (2003). Monitoring High Volume On-line Transaction Processing Systems Using a Continuous Sampling Approach. *International Journal of Auditing*, 7(1), 3-19.
 7. Hamilton, S. N., Miller, W. L., & Saydjari, A. O. O. S. (2002). The role of game theory in information warfare. In *4th Information Survivability Workshop (ISW-2001/2002)*.
 8. Hughes, J. S. (1977). Optimal Internal Audit Timing. *The Accounting Review*, 52(1), 56-68.
 9. Littlewood, B., Brocklehurst, S., Fenton, N., Mellor, P., Page, S., Wright, D., Dobson, J., McDermid, J., and Gollmann, D. (1993). Towards operational measures of computer security. *Journal of Computer Security*, 2(2-3), 211-230.
 10. Liu, P., Zang, W., & Yu, M. (2005). Incentive-based modeling and inference of attacker intent, objectives, and strategies. *ACM Transactions on Information and System Security*, 8(1), 78-118.
 11. Lye, K., & Wing, J. M. (2005). Game strategies in network security. *International Journal of Information Security*, 4(1), 71-86.
 12. Morey, R. C., & Dittman, D. A. (1986). Optimal Timing of Account Audits in Internal Control. *Management Science*, 32(3), 272-282.
 13. Ortalo, R., Deswarte, Y., & Kaâniche, M. (1999). Experimenting with quantitative evaluation tools for monitoring operational security. *IEEE Transactions on Software Engineering*, 25(5), 633-650.
 14. PricewaterhouseCoopers. (2006). State of the internal audit profession study: Continuous auditing gains momentum. Retrieved on September 1, 2010 from <http://www.pwc.com/us/en/internal-audit/publications/2009-study-internal-audit-profession.jhtml>
 15. PricewaterhouseCoopers. (2009). State of the internal audit profession study: Business upheaval: Internal audit weighs its role amid the recession and evolving enterprise risks. Retrieved on September 1, 2010 from <http://www.pwc.com/us/en/internal-audit/publications/2009-study-internal-audit-profession.jhtml>
 16. Ross, S. M. (2006). *Introduction to Probability Models*, Ninth Edition (9th ed.). Burlington, MA: Academic Press.
 17. Sallhammar, K., & Knapskog, S. J. (2004). Using Game Theory In Stochastic Models For Quantifying Security. In *Proceedings of the 9th Nordic Workshop on Secure IT-systems*. Presented at the Nordsec 2004, Espoo, Finland, November 4-5, 2004.
 18. Sallhammar, K., Knapskog, S. J., & Helvik, B. E. (2005). Using stochastic game theory to compute the expected behavior of attackers (pp. 102-105). Presented at the Applications and the Internet Workshops, 2005. Saint Workshops 2005. The 2005 Symposium on.
 19. Singh, H., & Chander, V. (2008). Estimating the Variance of an Exponential Distribution in the Presence of Large True Observations. *Austrian Journal of Statistics*, 37(2), 207-216.
 20. Weber, R. A. (1998). *Information Systems Control and Audit*. Upper Saddle River, NJ: Prentice Hall.

Appendix

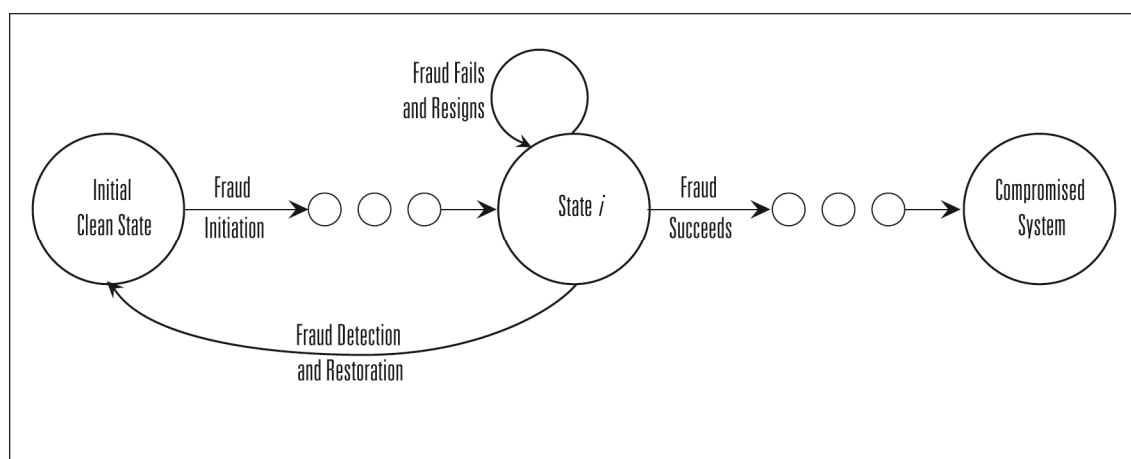


Figure 1: State Transition Diagram for Fraud




Figure 2: State Transition Rates




Figure 3: Stochastic Model of Fraud

$1/\lambda_{\text{fraud}}$	Days
0.000001	17.08
0.00001	14.49
0.0001	13.19
0.001	13.06
0.01	13.05

$p_i(f)$	Days
0.25	17.08
0.20	17.73
0.15	18.82
0.10	20.99
0.05	27.51

Table 1: Simulation Results

$p_1(s)$	Days
0.9	17.08
0.8	17.24
0.7	17.44
0.6	17.72
0.5	18.11
0.4	18.68
0.3	19.65
0.2	21.59
0.1	27.36

$p_2(s)$	Days
0.9	16.78
0.8	16.91
0.7	17.08
0.6	17.30
0.5	17.61
0.4	18.07
0.3	18.84
0.2	20.38
0.1	25.01

Table 2: Simulation Results

The IISTE is a pioneer in the Open-Access hosting service and academic event management. The aim of the firm is Accelerating Global Knowledge Sharing.

More information about the firm can be found on the homepage:
<http://www.iiste.org>

CALL FOR JOURNAL PAPERS

There are more than 30 peer-reviewed academic journals hosted under the hosting platform.

Prospective authors of journals can find the submission instruction on the following page: <http://www.iiste.org/journals/> All the journals articles are available online to the readers all over the world without financial, legal, or technical barriers other than those inseparable from gaining access to the internet itself. Paper version of the journals is also available upon request of readers and authors.

MORE RESOURCES

Book publication information: <http://www.iiste.org/book/>

IISTE Knowledge Sharing Partners

EBSCO, Index Copernicus, Ulrich's Periodicals Directory, JournalTOCS, PKP Open Archives Harvester, Bielefeld Academic Search Engine, Elektronische Zeitschriftenbibliothek EZB, Open J-Gate, OCLC WorldCat, Universe Digital Library, NewJour, Google Scholar

