IIMB Management Review (2012) 24, 79–84

available at www.sciencedirect.com

SciVerse ScienceDirect

journal homepage: www.elsevier.com/locate/iimb

IIMB Management Review

Determining expected behaviour of fraudsters for a continuous audit system

Mathew A. Thomas*, Rahul R. Marathe [1]

Department of Management Studies, Indian Institute of Technology Madras, I.I.T. Post Office, Chennai 600 036, India
Available online 3 April 2012

KEYWORDS
Continuous audit system;
Stochastic game theory;
Fraud prevention

Abstract This research attempts to determine the behaviour of fraudsters in a continuous audit system where the fraudsters have multiple options for committing fraud. The system is modelled as a Continuous Time Markov Chain where the state changes are caused by the fraudster's actions. The model uses a dynamic game with probabilistic transitions to determine the expected behaviour of the fraudster.

Introduction

After a string of high profile business failures (most notably the Enron scandal) caused by non-transparent financial statements and outright fraud by senior management, which were not detected by auditors, the US Congress passed the Sarbanes–Oxley Act in 2002. This Act mandated a set of internal controls and required management and external auditors to report on the adequacy of internal controls. This required the management and external auditors to document and test manual and automated controls. The Act specifically required the management to perform a fraud risk assessment and evaluate controls designed to prevent or detect fraud. As part of risk assessment, an organisation would like to find out which components of its system are more susceptible to certain threats than others. In this research, we attempt to find the components of a transaction system that are susceptible to fraud by analysing the behaviour of a potential fraudster. This will assist a security administrator in taking steps to mitigate fraud risks and fine tune the parameters of a continuous audit system. It will also assist the auditor in focussing on risk prone areas during an audit.

One of the key problems when applying traditional external audit techniques to large transaction systems is that audit data is gathered long after the economic events are recorded and it is often too late to prevent corrective action (Vasarhelyi & Halper, 1991). To mitigate this problem, large transaction systems use continuous audit systems. Continuous audit is defined as 'a comprehensive electronic audit process that allows auditors to provide some degree of assurance on continuous information simultaneous with, or shortly after, the disclosure of the information' (Rezaee, Sharbatoghlie, Elam, & McMickle, 2002). However, internal auditors and more recently system auditors have increasingly been required to examine

* Corresponding author. Tel.: +91 9884701002.
E-mail addresses: mathewdoms@gmail.com (M.A. Thomas), rrmarathe@iitm.ac.in (R.R. Marathe).
[1] Tel.: +91 44 257 4579.
Peer-review under responsibility of Indian Institute of Management Bangalore.

Production and hosting by Elsevier

0970-3896
http://dx.doi.org/10.1016/j.iimb.2012.02.002

organisational data more frequently in order to detect errors and/or malicious activity. This also implies analysing the controls in these systems. In an online transaction system, 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 (Vasarhelyi & Halper, 1991) or as close to real time as possible. In a continuous audit system the data stream is monitored and analysed at relatively short intervals (hourly, daily, etc.) using a set of auditor-defined rules (Vasarhelyi & Halper, 1991). When the system detects an exception to these rules, it triggers an alarm notifying the auditor about the occurrence of the exception.

In this paper we examine how the stochastic game theory can be used to model the behaviour of a malicious person or a potential fraudster. The output of the model will be used to determine the time taken by a potential fraudster to complete a fraudulent action. We assume that the fraudster is rational and will factor in the consequences of his actions before committing fraud.

Literature review

Previous studies in the area of audit timing (Boritz & Broca, 1986; Hughes, 1977; Morey & Dittman, 1986; Rossi, Tarim, Hnich, Prestwich, & Karacaer, 2010; Wilson & Ranson, 1971) have been concerned with fixing the time interval for internal audits. All these papers assume that in the absence of audit, errors in the accounting system will occur and that the number of such errors will accumulate with time until the audit detects and corrects them. The errors are modelled as increasing linearly with time (Hughes, 1977; Wilson & Ranson, 1971) or based on the auditor's subjective judgment (Boritz & Broca, 1986) or occurring randomly (Morey & Dittman, 1986). Some researchers (Dodin & Elimam, 1997; Dodin, Elimam, & Rolland, 1998) do not consider errors at all. In cases where the researchers do not consider errors, the authors assume that the auditors have to complete certain audits within a certain time, subject to certain constraints, irrespective of whether the transaction system has errors or not. In the previous research it is generally assumed that the errors are non-malicious in nature. In this paper, we focus only on malicious threats. (In a later work, we develop a model that combines both malicious and non-malicious threats.)

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 defence scenarios as a game of moves and counter moves. Game theory already has algorithms to predict the likelihood of an action being selected in such scenarios. Alpcan and Başar (2003) pointed out that game theory provides a rich set of tools for modelling and analysing 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) modelled an Intrusion Detection System (IDS) as a network of sensors. They modelled the behaviour of an attacker and the IDS as a two-person, nonzero-sum, non-cooperative game and developed 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 the administrator. Working with the two-player game model, they pointed out that a team of attackers or team of administrators can be modelled as a single omnipresent attacker or defender. Thus, a two-player game is sufficient for the modelling 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 interdependence between AIOS and defender objectives. The authors were primarily interested in examining the characteristics of attackers rather than the attacks and the attacker's intent was modelled 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 the defence strategy to be adopted.

Sallhammar and Knapskog (2004) built upon the work of Alpcan and Başar (2003), Liu et al. (2005), Lye and Wing (2005) and modelled an intrusion as a series of state changes from an initially secure state to a compromised state. They were primarily interested in devising a metric for an information system's 'trustworthiness' in terms of its dependability in resisting intrusions. In a later paper, Sallhammar, Knapskog, and Helvik (2005) expanded their model to include attacker intent, objectives, and strategies. In this paper, we use the approach of Sallhammar and Knapskog (2004) and demonstrate how their model can be used to determine the timing of audit in a continuous audit system. Prior research (Cavusoglu, Mishra, & Raghunathan, 2005; Cavusoglu, Raghunathan, & Yue, 2008) has 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 maximised when the firm credibly commits and communicates its strategy to the attacker. 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.

In this paper we are explicitly concerned with the behaviour of malicious agents or fraudsters who attempt to circumvent the controls in a financial system in order to perpetuate a fraud.

Fraudster's expected behaviour

The integrity of a transaction system is modelled as a continuous time Markov chain (CTMC), where the first state is the initial secure state. An attempt at fraud causes an intentional state change from one state to another. Each state, except the first, represents a situation where the integrity of the information system is breached. However, this model ignores the effects of unintentional errors.

Systems, in general, have weaknesses and a potential fraudster can take recourse to several courses of action to exploit them. In order to commit a successful fraud, a fraudster has to perform several atomic actions and in order for the fraud to succeed, all atomic actions must be successfully completed. The fraudster also has the option

of either not initiating the fraud or not continuing with it after initiating it. Also, at each possible stage, a fraudster may have multiple actions available to her to compromise the system. For example, a person who intends to defraud a payroll system may create a fake attendance record, or fake an unauthorised perk or bonus. In either case, the fraudster could end up earning more than what she would be entitled to.

In each state s:$s = 1,\ldots,z$, a fraudster can take m_s actions, which can broadly be categorised as follows:

1. Continue the fraud by choosing an action f_i^s, where $i = 1,\ldots, m_s - 1$. In case the action is unsuccessful, there is no reward and if the action is successful, she obtains a reward. There is also a possibility that the fraud is detected at this stage, in which case she has to pay a penalty (cost).
2. Cease and desist from the fraud action, $f_{m_s}^s$. In this case, the fraudster has to bear a cost depending on how far the fraud has progressed.

The probability that the fraudster will choose action i in state s is denoted by $p_{\text{fraud}}(f_i^s)$. Hence, for each state s in the state transition model, the fraudster's expected choice of action can be represented by a stochastic vector,

$$\overrightarrow{p}_{\text{fraud}}(f^s) = \left(p_{\text{fraud}}\left(f_1^s\right), \ldots, p_{\text{fraud}}\left(f_{m_s}^s\right)\right)$$

where $\sum_{i=1,\ldots m_s} p_{\text{fraud}}(f_i^s) = 1$

The complete set of such stochastic vectors is given by

$$P_{\text{fraud}} = \left\{\overrightarrow{p}_{\text{fraud}}(f^s) | s = 1, \ldots, z\right\} \quad (1)$$

To continue the fraud from state s, the fraudster must also succeed in the action chosen by her $[p_{\text{success}}(f_i^s)]$. In order to find the stationary probability distribution, we consider the embedded discrete time Markov chain (DTMC) within the CTMC. An embedded DTMC is obtained by looking at only the transition instances. In this case, since the fraudster's reward and cost come only from the transition, we can ignore the time spent in any given state and only consider the probabilities of making a transition from one state to another. This embedded DTMC must be specified in terms of transition probabilities rather than transition rates. Thus, the probability that the fraudster causes a transition from state s (at time $_k$) to state t (at time $_{k+1}$) is given by:

$$P(X_{k+1} = t | X_k = s) = p_{\text{fraud}}\left(f_i^s\right) \cdot p_{\text{success}}\left(f_i^s\right) \quad (2)$$

If the fraudster's chosen action fails or she chooses $f_{m_s}^s$ the state of the system is unchanged. To model the state transition model the following information is necessary:

1. Enumerate all the possible choices that a potential fraudster has at each state.
2. Estimate the success probabilities to those choices. These could be based on subjective assessments.
3. Compute the probabilities of the actions. This paper deals with the manner of computing these probabilities.

This state of affairs can be represented as a dynamic game with probabilistic transitions. The audit system and the fraudster can be modelled as playing the game in stages. The players select actions and receive payoffs that depend on the current state and the chosen action. This results in a new state whose distribution depends on the previous state and the actions chosen by them. The game then moves to a new state and the play is repeated at the new state. Thus, there are z states or game elements Γ_s:$s = 1,\ldots,z$. Thus, the complete stochastic game model can be defined as:

$$\begin{aligned}
\text{Players} &: N = \{\text{fraudster, audit system}\}\\
\text{Game elements} &: \Gamma_s : s = 1, \ldots, z\\
\text{Fraudsters actions} &: F_s = \left\{f_1^s, \ldots, f_{m_s}^s\right\}\\
\text{Actions of the audit system} &: \text{As} = \{d^s, \phi^s\} = \{\text{detect, miss}\}
\end{aligned}$$

For each possible combination of actions F_s and A_s, there is a payoff, $\gamma_{f_i\phi}^s = p_{\text{success}}(f_i^s)(\gamma(f_i^s, \phi^s) + \gamma(\Gamma_t))$. This implies that in state s, if the fraudster succeeds with his ith strategy he get a payoff of $\gamma(f_i^s, \phi^s)$ and the expected payoff from playing the next game $\gamma(\Gamma_t)$. These payoffs $[\gamma(f_i^s, \phi^s) + \gamma(\Gamma_t)]$ are contingent upon success of the action. Thus, the payoffs are multiplied by the probability of success $[p_{\text{success}}(f_i^s)]$. In case the fraudster chooses action i but is detected by the audit system the payoff is $\gamma_{f_i d}^s = \gamma(f_i^s, d^s)$.

The game ends when the fraud is detected or if the fraudster opts not to continue. The fraudster assumes that the audit system is a rational player who seeks to minimise the fraudster's expected payoff. This enables the Nash Equilibrium of the complete stochastic game to be calculated and results in a set of minimax solution vectors. These solution vectors represent a complete attack strategy, which maximises the expected payoff of the fraudster and ensures that she has no ex-post regrets.

Nash Equilibrium of a stochastic game

A two player zero sum stochastic game (Owen, 1995; page 96) is a set of z game elements or states Γ_s:$s = 1, \ldots, z$. Each game element is represented by a $m_k \times n_k$ matrix (where the fraudster has m_k actions and the system has n_k actions), whose entries are of the form

$$\mu_{ij}^k = u_{ij}^k + \sum_{l=i,\ldots,z} q_{ij}^{kl} \Gamma_l$$

with $q_{ij}^{kl} \geq 0$ and $\sum_{l=i,\ldots,z} q_{ij}^{kl} < 1$.

This implies that in state k, if player I chooses his ith pure strategy and player II chooses his jth pure strategy, player I will receive a payoff of u_{ij}^k plus a possibility of future payoff, q_{ij}^{kl}, from playing the lth game. The condition $\sum_{l=i,\ldots,z} q_{ij}^{kl} < 1$ ensures that probability of infinite play is zero and that all expected payoffs are finite.

A strategy set for player I is the set x^{kt} of m_k vectors for $k = 1,\ldots,z$ of m_k vectors satisfying

$$\sum_{i=1,\ldots,m_k} x_i^{kt} = 1,\ x_i^{kt} \geq 0.$$

Here x_i^{kt} is the probability that player I will choose action i, assuming that she is playing the game element Γ_k at the tth stage of the game. Strategy for player II is a similar set of n_k vectors y^{kt}.

Given a pair of strategies, an expected payoff can be calculated for any $k = 1,\ldots,z$ on the assumption that the first stage of the game is Γ_k. Thus, the expected payoff for a pair of strategies is a z-vector (since there are z states, and for each state we have a probability vector). The lowest expected payoffs are the optimal strategies and the value of the game is the expected value of the game if the players play their optimal strategies. This expected value is a z-vector, $v = (v_1, v_2,\ldots,v_z)$. In order for the value vector to exist, the game element Γ_l must be replaced by the value component v_l. A sequence of vectors which converges to the desired vector can be constructed as follows.

$$v^0 = (0, 0, \ldots, 0),$$

$$x_{ij}^{kr} = \gamma_{ij} + \sum q_{ij}^{kl} v_i^r, \quad r = 1, 2\ldots$$

$$v_i^{r+1} = \text{value}\left(x_{ij}^{kr}\right).$$

The sequence of vectors will converge and the limit of the converged value vector is the optimal stationary strategy (Owen, 1995; pages 97–99) for the stochastic game. This vector provides us with a means to compute the probable behaviour of a fraudster at each state of the game.

Illustration

Consider a financial system that processes payroll, receivables and inventory. The notation used is the following,

$$X_s = \{(x, y, z) | x, y, z \in \{0, 1\}\}, \tag{3}$$

where for example, (1,0,1) means that payroll (x) and inventory (z) are compromised but the receivables (y) are not. The fraudster's payoff for compromising the inventory system is assumed to be 30, for compromising the receivables, 20 and payroll, 10. These payoffs also determine the priorities of the fraudster, and once he has compromised any one system, he will target a system with a higher payoff.

The game's action set can be defined as follows:

$$F_{(0,0,0)} = \{f_1, f_2, f_3, f_4\}, \tag{4}$$

$$F_{(1,0,0)} = \{f_2, f_3, f_4\}, \tag{5}$$

$$F_{(0,1,0)} = F_{(1,1,0)} = \{f_3, f_4\}, \tag{6}$$

$$A_{(0,0,0)} = A_{(1,0,0)} = A_{(0,1,0)} = A_{(1,1,0)} = \{a_1, a_2\}, \tag{7}$$

Where, the actions f_1, f_2, f_3 and f_4 represent 'defraud payroll', 'defraud receivables', 'defraud inventory' and 'do nothing' respectively, and the actions, a_1, a_2 represent 'detect' and 'miss' respectively for the audit system. The transition probabilities (shown in Table 1) are chosen arbitrarily and have no particular significance. The only assumption here is that the payroll system is the easiest to defraud, and the inventory system is the hardest. The payoffs are intended to reflect this assumption. Violating this particular assumption does not in any way invalidate the model. In the next section, we show the effects of varying the transition probabilities.

The game elements are as

$$\Gamma_{(0,0,0)} = \begin{pmatrix} -10 & 0.6\cdot(10 + \Gamma_{(1,0,0)}) \\ -20 & 0.4\cdot(20 + \Gamma_{(0,1,0)}) \\ -30 & 0.1\cdot 30 \\ 0 & -5 \end{pmatrix},$$

$$\Gamma_{(1,0,0)} = \begin{pmatrix} -20 & 0.5\cdot(20 + \Gamma_{(1,1,0)}) \\ -30 & 0.2\cdot 30 \\ 0 & -10 \end{pmatrix}, \tag{8}$$

$$\Gamma_{(0,1,0)} = \begin{pmatrix} -30 & 0.3\cdot 30 \\ 0 & -10 \end{pmatrix}, \ \Gamma_{(1,1,0)} = \begin{pmatrix} -30 & 0.4\cdot 30 \\ 0 & -15 \end{pmatrix} \tag{9}$$

$$\vec{p}_{\text{fraud}}\left(f^{(1,1,0)}\right) = (0.26, 0.74) \tag{10}$$

$$\vec{p}_{\text{fraud}}\left(f^{(0,1,0)}\right) = (0.20, 0.80) \tag{11}$$

$$\vec{p}_{\text{fraud}}\left(f^{(1,0,0)}\right) = (0.28, 0, 0.72) \tag{12}$$

$$\vec{p}_{\text{fraud}}\left(f^{(0,0,0)}\right) = (0.29, 0, 0, 0.71) \tag{13}$$

The vector $\vec{p}_{\text{fraud}}(f^{(1,1,0)}) = (0.26, 0.74)$, implies that if the fraudster has compromised the payroll and inventory systems ($f^{(1,1,0)}$), he will attempt to compromise the inventory system with a probability of 0.26 and cease to commit the fraud with a probability of 0.74. Plugging the resulting value of the game into equation (8) above, we work backwards and arrive at the probability that the fraudster will initiate the fraud on an uncompromised system $\vec{p}_{\text{fraud}}(f^{(0,0,0)})$. The results show that a rational fraudster will defraud the payroll system with a probability of 0.29.

Sensitivity analysis

In order to examine the behavioural validity (Bossel, 1994) of the model, we vary the transition probabilities from 10% to 90% and examine the result (Fig. 1). The transition probabilities represent the probability of succeeding in the fraud attempt at each stage.

As can be seen from the figure, transition probabilities (which serve as a proxy to indicate the strength of the internal control system) increase from 10% to 100%, but the probability of committing fraud decreases from about 33% to 8%. The lower the state transition probabilities, the better the internal control system and vice-versa. The above result implies that the weaker the internal control system, the lower will be the probability of committing fraud. This rather counterintuitive result is due to the adoption of minmax strategy. As the transition probabilities increase, the fraudster maximises her worst-case payoff by reducing the probability of committing fraud. However, by examining the changes in the payoff matrix, we come to more interesting results as shown in the following paragraphs.

Table 1 Transition probabilities.

	(1,0,0)	(0,1,0)	(0,0,1)	(1,1,0)	(1,0,1)	(0,1,1)	(1,1,1)
(0,0,0)	0.6	0.4	0.1	—	—	—	—
(1,0,0)	—	—	—	0.5	0.2	—	—
(0,1,0)	—	—	—	—	—	0.3	—
(1,1,0)	—	—	—	—	—	—	0.4

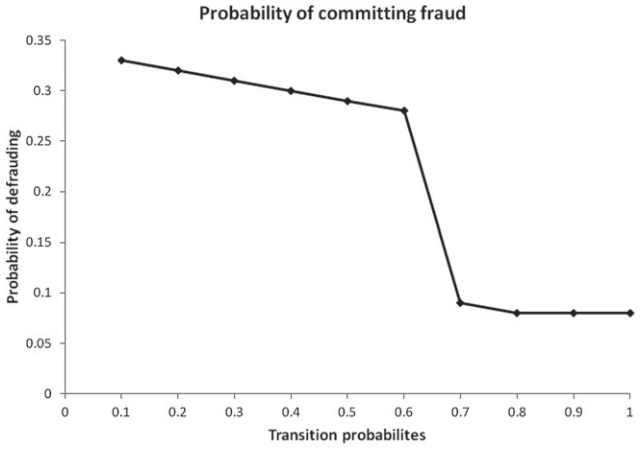

Figure 1 Probability of committing fraud.

We carried out another simulation to check whether the payoffs chosen have influenced the results. In each case, the inequality ('payoff of inventory' $>$ 'payoff of receivables' $>$ 'payoff of payroll') was maintained since that is a basic feature of the model. Altering the inequality will only alter the order in which the systems are compromised and will have no effect on the mathematical model. When the payoffs were scaled uniformly there was no significant change in the final result. We then changed the payoffs arbitrarily while maintaining the inequality between the payoffs. In such a case the final result (probability of committing fraud) changed by about $+$ or $-$ 3%. That is, the final probability of committing fraud varied from 67% to 72%.

There are three significant results that emerge from the sensitivity analysis. Firstly, the stochastic model dominates the game mode. In mathematical terms, it implies that the transition matrix determines the probability of committing fraud. Secondly the structure of the payoff matrix determines the priorities of the fraudster. Thirdly, the model is fairly robust and consistent across a range of payoff values. In other words, more than the payoffs it the ease with which a fraudster can commit fraud that determines whether fraud occurs or not. This has implications for security administrators and auditors. The transition matrix represents the internal control system and the payoff matrix, the incentive for committing fraud. The research indicates that the weakness in the internal control system is more likely to induce a fraud rather than the potential reward from committing fraud. Generally the tendency of an organisation is to place tight controls on high value transactions and be relatively lax with low value transactions. This research suggests that this might not be the correct approach as fraudsters are influenced more by ease of carrying out the fraud rather than the potential monetary reward from fraud.

Conclusion

To the best of our knowledge this paper represents one of the first attempts at using probability vectors to model state transitions resulting from fraud. The probability vectors allow the modelling of multiple actions by the fraudster and intentional state changes caused by fraud are computed by employing these probability vectors. This allows us to compute the expected behaviour of a fraudster in a transaction system and can be used to decide the timing and frequency of a continuous audit system that maximises the probability of detecting the fraud.

There are three key limitations of this model. Firstly, the game model uses a zero-sum game. The benefit of this approach is that a fraudster is required to know only the payoffs to herself and need not know the payoff to the audit system. The payoff to the audit system is simply the inverse of the payoff to the fraudster. This simplifies the evaluation of the model considerably. In reality, the game between the audit system and a putative fraudster is unlikely to be a zero-sum game. Secondly, we assume that the payoffs are known and certain. This assumption may not always be true. While it is relatively easy for a fraudster to estimate the potential reward for committing fraud, she may not be in a position to estimate the cost of being detected. Similarly, for an audit system the cost of fraud loss may not be easily measurable, as it involves intangible costs like loss of reputation, loss of brand image, investor disaffection, etc. Thirdly, this research assumes that the audit system's choices are either detect the fraud or fail to detect the fraud. In many circumstances, it may be more appropriate to give a probabilistic value that accounts for the uncertainty inherent in deciding whether a set of transactions constitutes a fraud or not.

The model presented here can be used in an embedded module within a continuous audit system. Instead of having the audit parameters set manually and subjectively, the embedded module automatically determines the audit parameter and timing of the continuous audit system. There are several benefits that accrue from this. Firstly, the audit system is likely to operate more effectively by increasing the probability of detecting frauds. Secondly, the organisation can minimise losses from frauds by detecting them early. Thirdly, since the audit parameters are set objectively, they provide greater legal protection in the event of a fraud being uncovered. For example, if a fraud occurs and the auditor/manager is being sued for

professional negligence, then the fact that audit parameters are set objectively would offer greater protection and provide defence against accusations of bias.

References

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, Vol. 3* (pp. 2595—2600).

Boritz, E., & Broca, D. S. (1986). Scheduling internal audit activities. *Auditing: A Journal of Practice & Theory, 6*(1), 1—19.

Bossel, H. (1994). *Modeling and simulation*. Wellesley, MA: A K Peters.

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.

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.

Dodin, B., & Elimam, A. A. (1997). Audit scheduling with overlapping activities and sequence-dependent setup costs. *European Journal of Operational Research, 97*(1), 22—33.

Dodin, B., Elimam, A. A., & Rolland, E. (1998). Tabu search in audit scheduling. *European Journal of Operational Research, 106*(2—3), 373—392.

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*, Vancouver, BC, Canada.

Hughes, J. S. (1977). Optimal internal audit timing. *The Accounting Review, 52*(1), 56—68.

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.

Lye, K., & Wing, J. M. (2005). Game strategies in network security. *International Journal of Information Security, 4*(1), 71—86.

Morey, R. C., & Dittman, D. A. (1986). Optimal timing of account audits in internal control. *Management Science, 32*(3), 272—282.

Owen, G. (1995). *Game theory* (3rd ed.). New York: Academic Press.

Rezaee, Z., Sharbatoghlie, A., Elam, R., & McMickle, P. L. (2002). Continuous auditing: building automated auditing capability. *Auditing, 21*(1), 147—164.

Rossi, R., Tarim, A., Hnich, B., Prestwich, S., & Karacaer, S. (2010). Scheduling internal audit activities: a stochastic combinatorial optimization problem. *Journal of Combinatorial Optimization, 19*(3), 325—346.

Sallhammar, K., & Knapskog, S. J. (November 4—5, 2004). Using game theory. In *Stochastic models for quantifying security. Proceedings of the 9th Nordic Workshop on secure IT-systems, Espoo, Finland*, .

Sallhammar, K., Knapskog, S., & Helvik, B. (2005). Using stochastic game theory to compute the expected behavior of attackers. In *Presented at the 2005 symposium on applications and the Internet workshops, 2005. Trento, Italy* (pp. 102—105).

Vasarhelyi, M., & Halper, F. (1991). The continuous audit of online systems. *Auditing: A Journal of Practice and Theory, 10*(1), 110—125.

Wilson, D., & Ranson, R. (July—August 1971). Internal audit scheduling—a mathematical model. *The Internal Auditor*, 42—50.