

Configuration of Detection Software: A Comparison of Decision and Game Theory Approaches

Huseyin Cavusoglu

A. B. Freeman School of Business, Tulane University, New Orleans, Louisiana 70118, huseyin@tulane.edu

Srinivasan Raghunathan

School of Management, The University of Texas at Dallas, Richardson, Texas 75083, sraghu@utdallas.edu

Firms are increasingly relying on software to detect fraud in domains such as security, financial services, tax, and auditing. A fundamental problem in using detection software for fraud detection is achieving the optimal balance between the detection and false-positive rates. Many firms use decision theory to address the configuration problem. Decision theory is based on the presumption that the firm's actions do not influence the behavior of fraudsters. Game theory recognizes the fact that fraudsters do modify their strategies in response to firms' actions. In this paper, we compare decision and game theory approaches to the detection software configuration problem when firms are faced with strategic users. We find that under most circumstances firms incur lower costs when they use the game theory as opposed to the decision theory because the decision theory approach frequently either over- or underconfigures the detection software. However, firms incur the same or lower cost under the decision theory approach compared with the game theory approach in a simultaneous-move game if configurations under decision theory and game theory are sufficiently close. A limitation of the game theory approach is that it requires user-specific utility parameters, which are difficult to estimate. Decision theory, in contrast to game theory, requires the fraud probability estimate, which is more easily obtained.

Key words: detection software; fraud detection; intrusion detection; false alarm rate; detection rate; ROC curve; decision theory; game theory

History: Received on April 21, 2003. Accepted by Robert T. Clemen and Don N. Kleinmuntz on March 30, 2004, after 3 revisions.

1. Introduction

In this paper we investigate the problem of optimal configuration of detection software using decision and game theory approaches. Firms are increasingly using software to detect fraudulent activities in a variety of domains such as security, financial services, tax, and auditing. Detection software uses pattern recognition to classify an event as a potentially normal or a potentially fraudulent event. In the IT security area, intrusion detection software (IDS) attempts to detect unauthorized use of computer resources by analyzing a user's audit trail. Physical security mechanisms such as authentication systems that allow people to gain entry or access (e.g., airport security) utilize user-profile data. Fraud detection software used by credit card companies analyzes transaction data such as the amount charged and the location of the transaction. Audit programs used by the U.S. Internal Revenue Service (IRS) attempt to catch tax evasion

by analyzing tax return data. The effectiveness of detection software in these settings is measured in terms of its classification errors, which consist of false-positive and false-negative rates. False positives occur when the detection software classifies a normal event as fraudulent. False negatives occur when the detection software classifies a fraudulent event as normal. Whereas it is desirable to have low false-positive and low false-negative rates in a detection program, a reduction in one type of error is often accompanied by an increase in the other type. The goal of detection software configuration is to balance the two error rates to minimize the firm's cost.

Decision theory is often used to analyze the risk associated with fraud and to configure the fraud detection software. Recently, Ulvila and Gaffney (2004) proposed a decision theory-based approach to configure an IDS. Although the decision theory-based approach can provide a useful starting point for managing risk

in settings where potential for fraud exists, we argue in this paper that this method is incomplete because of the problem's strategic nature. The reason for the limitation of the decision theory approach can be stated as one simple proposition: It does not allow the firm's decisions to influence the behavior of fraudsters. Researchers and practitioners have long recognized the behavioral influences of a firm's actions on fraudsters. For example, in the IT security area, it has been pointed out that security should be viewed as a "cat-and-mouse" game played by firms and hackers (Jajodia and Miller 1993). Hackers do not randomly select their targets. They rationally make their choices based on how much effort will be required to succeed in hacking, the probability of getting caught, and the possible penalty (NIST 800-30). In the context of credit card fraud, firms increasingly recognize that credit card fraud is organized crime, and fraudsters change their patterns to avoid detection (Alaric 2003). Such strategic interactions between a firm's decisions and fraudsters have to be captured in the model used to configure detection software. Because decision theory is designed to analyze decision making under uncertainty where "nature" is the only "opponent," it is fundamentally inadequate to address the fraud detection problem where firms deal with strategic adversaries (Fellingham and Newman 1985). Modeling the interaction between firm and user decisions requires game theory.

Traditional decision theory assumes that the firm exogenously estimates the probability of fraud before choosing the configuration. Although the firm can perform sensitivity analysis with respect to the estimated probability, the model still provides only partial solutions. In the game theory, both configuration and fraud probability are endogenously determined. We compare the decision theory approach with two types of games: simultaneous and sequential. We find that if the game between the firm and the user is sequential with the firm as the leader and the user as the follower then the firm incurs a lower cost when it uses the game theory than it does when it uses the decision theory. Even when the firm and the user play a simultaneous game—i.e., when neither party has knowledge of the other party's decision prior to making its decision—game theory results in a lower cost except when the configuration under decision theory

is sufficiently close to the configuration under game theory.

In spite of the superiority of the game theory approach to the decision theory approach, decision theory appears to be more popular for risk management that uses detection software. Perhaps the attractiveness of the decision theory approach lies in the fact that it (only) requires the estimation of fraud probability as opposed to game theory, which requires the user's utility, a more difficult estimate.

1.1. Related Work

The paper that comes closest to ours is Ulvila and Gaffney (2004). They present a method that is based on decision theory for comparing and evaluating the performance of intrusion detectors and for determining the best configuration point of a detector for a given operating environment. Their study integrates costs of dealing with false-positive and false-negative errors and the quality profile of the IDS as indicated by its receiving operating characteristics (ROC) curve that relates these two error rates. Our work complements Ulvila and Gaffney. Although Ulvila and Gaffney consider the hostility of operating environment as one of the factors that determine the configuration, they assume a fixed attack level. Our study assumes that the attack level is determined by the operating environment and the configuration. In contrast to Ulvila and Gaffney (2004), one of the goals of this paper is to develop a framework that explicitly models the strategic aspect related to the user behavior in response to software configuration in a cost-minimizing framework. Another goal is to compare the decision theory approach with the game theory approach for software configuration.

The broader issue addressed by this paper relates to configuration management and performance evaluation of software. Guidelines from commercial classification software manufacturers—e.g., Novell (Sriram 2002), as well as research institutes such as the Software Engineering Institute (SEI) at Carnegie Mellon University—emphasize the need for proper configuration of detection systems. For example, SEI's report on IDSs (Allen et al. 2000) cautions firms against accepting the default settings automatically and recommends appropriate configuration to balance security and operational requirements. Similar observations have been made for other detection software

such as explosives detection systems used by airports (NMAB 1998). According to a federal report, between November 2001 and February 2002 security screeners missed 70% of knives, 60% of simulated explosive devices, and 30% of guns. Commenting about this, the former FAA security chief Billie Vincent noted that if the metal detectors are tuned to detect that much metal, there would be an alarm on every person going through, which would imply hand screening of all bags and passengers and a tremendous increase in the cost of security at the airports (CNN 2002).

In order to support and manage software configuration, a whole industry that develops configuration management tools has evolved. The overriding goal of these efforts has been the performance of the software as measured by its classification accuracy. Evaluation of software such as IDS, machine learning systems, and other classification programs has relied on false-positive and false-negative rates (Sarkar and Sriram 2001, Durst et al. 1999). Modeling the accuracy of classification software is a well-established area with many known models and measures such as lift, response ratio, L-quality, and others (Provost and Fawcett 1997, Shapiro and Masand 1999, Shapiro and Steingold 2000, Steingold et al. 2001). All these models use the Bayesian theory. Our model also relies on the Bayesian theory to generate the quality profile of detection software. However, unlike previous models, our model for software configuration does not use classification accuracy exclusively. Firms are interested in not only raw performance measures, but also the overall cost of the detection process. Recently researchers have recognized the importance of costs of misclassification in measuring the detection software performance. Lee et al. (2002) developed a detection model that incorporates these costs in the classification algorithm itself in order to minimize costs. Chavez (2000) proposed a decision analytic rule that minimizes cost in order to decide whether to release or keep working on commercial software. Our model for software configuration also minimizes the firm's cost. The innovation in our analysis, absent in prior work, is the explicit modeling of user behavior in response to the firm's decisions in a cost-minimizing framework.

The need for incorporating user behavior in software configuration has been recognized in the

IT security context; IT security needs to develop better security breach prevention and detection software. Jonsson and Olovsson (1997, p. 235) pointed out that the common criteria used to evaluate security software "reflect static design properties and do not incorporate the interaction with the environment in a probabilistic way." Using an experiment, they modeled attacker behavior and concluded that the attacking process can be split into learning, standard attack, and innovative attack phases. They also found that the time between security breaches could be modeled using an exponential distribution. In a related work, Ortalo et al. (1999) proposed measures using a Markovian model that estimates the effort an attacker might expend to exploit system vulnerabilities. Although these models capture the dynamic aspects of user behavior, we use a static model that captures attacker behavior through the attack probability.

The rest of this paper is organized as follows. Section 2 summarizes the statistical theory that underlies most detection software. Section 3 introduces our model framework to address the configuration problem. Section 4 derives the optimal configurations under decision theory and game theory approaches. Section 5 compares the results under these two approaches. In §6, we discuss the limitations of the game theory approach. Section 7 concludes the paper.

2. Detection Software and ROC Curves

The principles underlying detection software are grounded in classical statistical decision theory. In the simplest case, there are two types of sources that generate inputs to the detection software: normal (H_0) and fraudulent (H_1). The normal source generates legal or authorized transactions. The fraudulent source generates illegal or fraudulent transactions. In a typical real-life detection scenario, a large percentage of transactions are legal. The skewed nature of the frequency distribution makes detection of illegal transactions difficult. The detection software observes the transaction but does not know whether it came from a normal or fraudulent source. The goal of the detection software is to classify each transaction as legal or fraudulent. Two types of errors can occur

in this classification: classification of a fraudulent transaction as a legal transaction (false negative) and classification of a legal transaction as a fraudulent transaction (false positive).

We define

Probability of detection = $P_D = \Pr(\text{classify into } H_1 | H_1 \text{ is true})$, or

Probability of false negative = $1 - P_D$

Probability of false positive = $P_F = \Pr(\text{classify into } H_1 | H_0 \text{ is true})$.

In general, we would like to have P_D as large and P_F as small as possible in detection software. However, it is not always possible to increase P_D and decrease P_F simultaneously. This is because of the variability of data associated with legal and illegal transactions and the imprecision of algorithms used by detection software. Many detection programs classify a transaction based on whether a numerical score computed from the transaction data exceeds a threshold value, or whether the transaction data satisfy a rule, or both. The quality parameters P_D and P_F of detection software can be adjusted, though not independently, by configuring its threshold value or rules. Consequently, the quality profile of configurable detection software is characterized by a curve that relates its P_D and P_F , known as the ROC curve (Trees 2001).

The ROC curve of detection software can be derived experimentally or analytically (Durst et al. 1999, Lippman et al. 2000, McHugh 2000). The analytical procedure proceeds as follows: Consider a detection software that uses a numerical score x computed from transaction data and a threshold value t to detect fraudulent transactions. Let the software classify a transaction as fraudulent if $x > t$ for that transaction. It follows that


$$P_D = \int_t^\infty f_F(x) dx$$

$$P_F = \int_t^\infty f_N(x) dx,$$

where $f_N(x)$ and $f_F(x)$ are the probability density functions of x for legal and fraudulent transactions, respectively. Figure 1 illustrates these probability calculations.

The shape of the ROC curve depends on the probability density functions of x_N and x_F . We assume that the numerical score used to distinguish normal

Figure 1 Computation of P_D and P_F



from fraudulent transactions follows an exponential distribution. Exponential distributions, besides being analytically tractable, capture the skewed nature of transaction data very well. Let the numerical scores for the normal and fraudulent transactions follow exponential distributions with parameters λ_N and λ_F , $\lambda_N > \lambda_F$, respectively. Then we can write P_D and P_F as

$$P_D = \int_t^\infty \lambda_F e^{-(\lambda_F x)} dx = e^{-\lambda_F t} \tag{1}$$

$$P_F = \int_t^\infty \lambda_N e^{-(\lambda_N x)} dx = e^{-\lambda_N t}. \tag{2}$$


Thus P_D can be expressed as a function of P_F as

$$P_D = P_F^r, \tag{3}$$

where $r = \lambda_F/\lambda_N$ is between 0 and 1. Lower values of r result in steeper ROC curves. Figure 2 shows sample ROC curves for various values of r .

Because detection programs are imperfect, i.e., they have false-positive and false-negative errors, in many

Figure 2 ROC Curves



organizations these programs are only used as decision support tools. A human expert might randomly investigate transactions classified to be fraudulent by the detection system. Similarly, the human expert might also randomly investigate transactions classified as legal by the detection system. The frequency of such random investigations depends on various costs, including the cost of performing the investigation and the cost of misclassification.

3. The Model Framework for the Configuration Problem

We consider a firm that uses a fraud detection software. The firm's purpose is to configure the fraud detection software by choosing a point on its ROC curve that minimizes the total cost. A user can generate legal or fraudulent transactions.¹ Users might commit fraud, depending on factors such as the benefit they derive from committing fraud, the penalty they will receive if they are caught, and the likelihood that they will be caught. We assume that a user committing the fraud derives a benefit of μ , if the fraud is undetected. If the fraud is detected, the user incurs a penalty of β for a net benefit of $\mu - \beta$. The penalty can take different forms depending on the nature of fraud. For instance, it can be the cost from legal prosecution or social humiliation. We denote the probability that a user commits fraud by ψ .

The detection program analyzes each of the transactions. If the program deems a transaction to be fraudulent, it generates a signal. The firm then decides whether or not to investigate the transaction. The firm makes decisions about whether or not to investigate based on the state (the signal or the lack of a signal) of the classification software. However, when the program generates a signal, the firm does not know with certainty whether there has been a true fraud or whether the program generated a false alarm. Similar

¹ We assume that every user is a potential fraudulent user. It might be more reasonable to assume that some users will not commit fraud at all. We can incorporate this aspect in our model by assuming two types of users: a user can be either honest with probability $1 - \lambda$ or dishonest (potentially fraudulent user) with probability λ . The results for this model are qualitatively identical to those discussed here. Here we are presenting only the case in which all users are potentially fraudulent, for brevity.

Table 1 List of Notations

Parameters	
d	Damage from an undetected fraud
c	Cost of manual investigation
β	Penalty to users when a fraud is detected
μ	Benefit to users when fraud is undetected
r	Proportion of the mean score of normal transactions to that of fraudulent transactions
Decision variables	
P_D	Probability of correct classification by the program when there is a fraud
$P_F (=P_D^c)$	Probability of incorrect classification by the program when there is no fraud
ψ	Probability of a user committing a fraud
ρ_1	Probability of manual investigation when the program signals a fraud
ρ_2	Probability of manual investigation when the program does not signal a fraud

uncertainty exists when the program does not generate a signal.

The firm supplements the program with manual investigation by a human expert. The human expert might investigate only a proportion (ρ_1) of transactions that generate signals. Furthermore, he might investigate a proportion (ρ_2) of transactions that did not generate signals. The firm incurs a cost of c each time the human expert performs a manual investigation. We assume that manual investigation always detects fraud.² If the firm detects the fraud, the firm does not incur any loss,³ other than the cost of manual investigation. When a fraud is undetected, the firm incurs a damage of d . Most companies estimate these possible damages in the risk-assessment phase before implementing and configuring the detection program.

The quality profile of the detector is modeled through its ROC curve. Table 1 summarizes the notations used in our model. We further assume that all parameters are common knowledge.

² We could extend the model easily to the case when manual investigation is not 100% effective. The results would not change qualitatively.

³ Again, our model could easily be extended to the case in which the firm recovers only a portion of the damage that has already been inflicted. The qualitative nature of results under this more general model is identical to that presented in this paper.

4. Optimal Configurations Under Decision Theory and Game Theory Approaches

Our analysis for deriving the optimal configuration uses backward induction, as follows: First, for a given configuration, i.e., P_D and P_F , we determine the optimal ρ_1 and ρ_2 as a function of P_D and P_F . Then, we determine the optimal P_D and P_F by minimizing the firm’s expected cost subject to the ROC curve constraint. The firm can use decision theory or game theory in determining the optimal decisions. We derive the optimal configuration for the decision theory approach first, followed by the game theory approach.

4.1. Decision Theory Approach

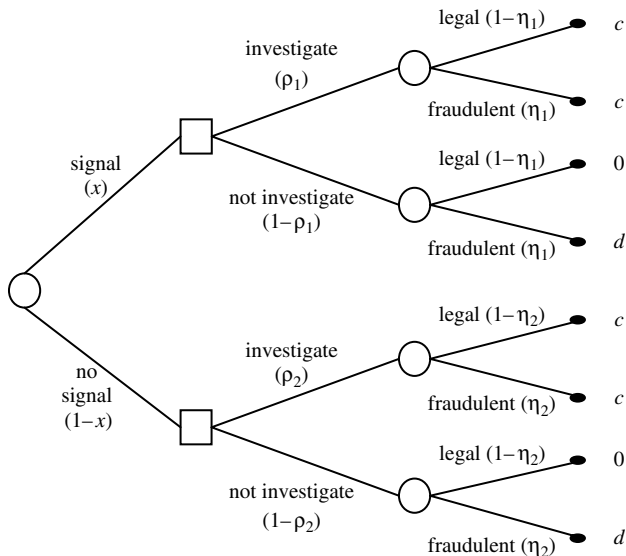
The decision theory approach determines the optimal configuration by minimizing its cost for a given risk environment. Firms assess the risk before applying a decision-theoretic model. The decision-theoretic model for a given configuration is shown in Figure 3. The payoff for each branch is given at the far right in the figure.

The probabilities of the firm being in the signal and no-signal states are

$$P(\text{signal}) = P_D\psi + P_F(1 - \psi)$$

$$P(\text{no-signal}) = (1 - P_D)\psi + (1 - P_F)(1 - \psi).$$

Figure 3 The Decision-Theoretic Model for Configuring a Detection Program



When the firm observes a signal or no signal from the detection program for a transaction, it updates its belief about the transaction type using Bayes’s rule. If the firm gets a signal, it determines the posterior probability of fraud as

$$\eta_1 = P(\text{fraud} | \text{signal}) = \frac{P(\text{signal} | \text{fraud})P(\text{fraud})}{P(\text{signal})}$$

$$= \frac{P_D\psi}{P_D\psi + P_F(1 - \psi)}.$$

Similarly, when the firm does not get any signal, it calculates the posterior probability as

$$\eta_2 = P(\text{fraud} | \text{no-signal})$$

$$= \frac{P(\text{no-signal} | \text{fraud})P(\text{fraud})}{P(\text{no-signal})}$$

$$= \frac{(1 - P_D)\psi}{(1 - P_D)\psi + (1 - P_F)(1 - \psi)}.$$

We can show that when $P_D \geq P_F$, as is the case in a ROC curve, $\eta_1 \geq \eta_2$.

The expected cost for each action in each of the decision nodes is computed to be the following:

$$\text{Cost}(\text{investigate} | \text{signal}) = \eta_1c + (1 - \eta_1)c = c$$

$$\text{Cost}(\text{not investigate} | \text{signal}) = \eta_1(d) + (1 - \eta_1)0 = d\eta_1$$

$$\text{Cost}(\text{investigate} | \text{no signal}) = \eta_2c + (1 - \eta_2)c = c$$

$$\text{Cost}(\text{not investigate} | \text{no signal}) = \eta_2(d) + (1 - \eta_2)0 = d\eta_2.$$

The firm decides to investigate or not investigate in a state (signal and no signal) by choosing the action that yields the lower expected cost. Result 1 shows the firm’s optimal strategies in the signal and no-signal states.

RESULT 1. The optimal manual investigation frequencies for a given detection software configuration in the decision theory framework are as follows:

$c/d < \eta_2$	$\eta_2 < c/d < \eta_1$	$c/d > \eta_1$
$\rho_1 = 1, \rho_2 = 1$ ①	$\rho_1 = 1, \rho_2 = 0$ ②	$\rho_1 = 0, \rho_2 = 0$ ③

Result 1 is consistent with our intuition. Choosing to investigate every transaction irrespective of whether or not the program generates a signal is the best strategy for the firm if the ratio (cost/benefit) of investigation is sufficiently low. However, it is optimal

Table 2 Firm's Expected Cost Under Decision Theory Approach

Region	Firm's expected cost
1	c
2	$cP_F(1 - \psi) + d\psi + P_D(c - d)\psi$
3	$d\psi$

for the firm not to investigate any transaction if the ratio is sufficiently high. If the ratio is moderate, the firm should investigate all and only those transactions that generate signals. Substitution of the above optimal investigation strategies in the firm's expected cost expressions gives the firm's expected cost for different parameter regions, as given in Table 2.

An examination of the expected cost expressions given in Table 2 reveals that configuration is relevant only in Region 2 because the expected costs in other regions are independent of P_D and P_F .

REMARK 1. Detection software is valuable and configuration is important only if (cost/benefit) ratio of manual investigation, (c/d) , is greater than η_2 and less than η_1 .

Since configuration is relevant only in Region 2, firms only in that region will use detection software to supplement their manual investigations. We focus only on this region to determine the optimal configuration. That is, we assume that the firm's cost parameters are such that it operates in Region 2. Writing P_F as a function of P_D (using Equation 3), we get the firm's expected cost in Region 2 as $c\sqrt{P_D}(1 - \psi) + d\psi + P_D(c - d)\psi$. Minimizing this over P_D gives

$$P_D^* = \left[\frac{c(1 - \psi)}{r\psi(d - c)} \right]^{r/(r-1)} \quad (6)$$

$$P_F^* = \left[\frac{c(1 - \psi)}{r\psi(d - c)} \right]^{1/(r-1)} \quad (7)$$

Substituting these optimal configuration points into the expected cost expression gives the expected cost of

$$d\psi - (1 - r)(d - c)\psi \left[\frac{c(1 - \psi)}{r\psi(d - c)} \right]^{r/(r-1)}.$$

If P_D^* computed from Equation (6) is greater than one, it will be set equal to one. When $P_D^* = 1$, the expected

cost is c . Result 2 summarizes our findings for the configuration under decision theory:

RESULT 2 (DECISION THEORY CONFIGURATION). Given that detection software is valuable, a firm using decision theory will configure the software such that

$$P_D^* = \min \left\{ \left[\frac{c(1 - \psi)}{r\psi(d - c)} \right]^{r/(r-1)}, 1 \right\}$$

and

$$P_F^* = \min \left\{ \left[\frac{c(1 - \psi)}{r\psi(d - c)} \right]^{1/(r-1)}, 1 \right\}$$

and realize an expected cost of

$$\min \left\{ d\psi - (1 - r)(d - c)\psi \left[\frac{c(1 - \psi)}{r\psi(d - c)} \right]^{r/(r-1)}, c \right\}.$$

The decision theory approach uses an estimate of the probability of fraud while solving the configuration problem. However, in reality, the probability that a user commits fraud need not necessarily be the same as the firm's estimated probability. The realized probability of fraud depends critically on how the user determines his strategy.

4.2. Game Theory Approach

The game theory approach differs from the decision theory approach in one important respect: In the game theory approach, the firm makes its decisions by anticipating the behavior of the user in response to its actions. Thus, while the decision theory approach assumes that the probability of fraud ψ is unaffected by P_D , P_F , ρ_1 , and ρ_2 , the game theory approach uses the fact that the user will change his strategy based on the firm's decisions. The nature of the game depends on the timings of the user's actions relating to fraud and the firm's actions relating to the investigation strategies. We consider two scenarios. In the *simultaneous* game, the user and the firm, respectively, make their fraud and investigation decisions simultaneously. In the *sequential* game, the firm decides on its investigation strategies first and then the user learns of them and makes his decision. For instance, the IRS routinely makes public the percentages of income tax returns it will audit for different income categories. Although the sequential game might appear to be more plausible in a real-life setting, the solution for the sequential game converges to that of the simultaneous game when several periods are considered

in which a player makes his decisions based on the decision made by the opponent in the previous period. In both games, we identify a Nash equilibrium. In a Nash equilibrium, neither player has an incentive to deviate from the equilibrium as long as the other player does not deviate.

4.2.1. Simultaneous Game. In order to analyze the strategic interactions between the firm employing the detection software and the user, we enhance the decision theory model to include the user's and firm's strategies. We characterize the game in strategic (normal) form in Table 3. The user's strategies are to commit fraud, F , or not commit fraud, NF , i.e., $S^U \in \{F, NF\}$. The firm might choose to investigate, I , or not investigate, NI , when the detector generates a signal. The firm has the same two strategies even when the detector does not generate a signal. Thus, the strategy space for the firm is the Cartesian product of the actions available at each of these two information sets. That is, $S^F \in \{(I, I), (I, NI), (NI, I), (NI, NI)\}$, where the first element in each pair specifies the action to be taken when the firm observes a signal, and the second element is the action when the firm does not observe any signal. For example (I, NI) implies that the firm investigates the transaction if it causes the detector to generate a signal and does not monitor the transaction if it does not cause the detector to generate a signal.

The solution to the game depends critically on parameter values. We use Nash equilibrium as the solution concept. It is a vector of strategies (one for the firm and one for the user) such that no player can increase its payoff by unilaterally changing strategies. There is always at least one solution, although this solution might not be in pure strategies. In other words, each player could play a mixed strategy by

randomly choosing from his pure strategies according to a probability distribution. However mixed strategies are not optimal in decision theory approach, as we saw in §4.1.

We derive the Nash equilibrium in behavioral strategies. As shown in Fudenberg and Tirole (1993, pp. 89–90), in a game of perfect recall such as ours, mixed and behavioral strategies are equivalent, and we use the mixed and behavioral formulations interchangeably for notational convenience. That is, we solve the game as if the strategy for the user is $\psi \in [0, 1]$ and the strategy space for the firm is $(\rho_1, \rho_2) \in [0, 1] \times [0, 1]$. The firm's expected cost for the signal and the no-signal states respectively are as follows:

$$F_S(\rho_1, \psi) = \rho_1 c + \eta_1(1 - \rho_1)d \quad (8)$$

$$F_N(\rho_2, \psi) = \rho_2 c + \eta_2(1 - \rho_2)d. \quad (9)$$

The firm's overall expected cost for a transaction is given by

$$F(\rho_1, \rho_2, \psi) = (P_F + \psi(P_D - P_F))F_S(\rho_1, \psi) + (1 - P_F - \psi(P_D - P_F))F_N(\rho_2, \psi). \quad (10)$$

The user's expected benefit is

$$H(\rho_1, \rho_2, \psi) = \psi\mu - \psi\beta(\rho_1 P_D + \rho_2(1 - P_D)). \quad (11)$$

The firm minimizes $F_S(\rho_1, \psi)$ when it gets a signal from the detection software and $F_N(\rho_2, \psi)$ when it does not get a signal from the detection software. The user maximizes $H(\rho_1, \rho_2, \psi)$. In the simultaneous game, both optimizations are done simultaneously. In the sequential game, the firm optimizes first using the fact that the user will determine his optimal strategy next after observing the firm's decision.

Result 3 holds in our model (the proof is given in the appendix).

RESULT 3. Assuming that $P_D \geq P_F$, the manual investigation probability when there is a signal is greater than or equal to the manual investigation probability when there is no signal (i.e., $\rho_1 \geq \rho_2$). In addition, both probabilities cannot be positive and less than one at the same time.

Result 3 is consistent with the intuition that the firm will investigate a larger fraction of the cases that generate signals from the detection software compared

Table 3 Game in Strategic Form

Firm strategies	User strategies	
	F	NF
(I, I)	$(c, \mu - \beta)^*$	$(c, 0)$
(I, NI)	$(d + (c - d)P_D, \mu - P_D\beta)$	$(cP_F, 0)$
(NI, I)	$(c - (c - d)P_D, \mu - (1 - P_D)\beta)$	$(c(1 - P_F), 0)$
(NI, NI)	(d, μ)	$(0, 0)$

Note. The firm's cost, the user's benefit.

with those that do not generate signals. This result is obtained because $P_D \geq P_F$. Now we can present the optimal frequency of manual investigations. Again, we present the proof in the appendix.

RESULT 4. The optimal frequencies for manual investigation for a given configuration of the detection software when using a simultaneous game theoretic framework are as follows.

Region	$c/d < 1$	$c/d > 1$
$\mu/\beta > 1$	$\psi = 1$ $\rho_1 = 1, \quad \rho_2 = 1$	③
$P_D < \mu/\beta < 1$	$\psi = \frac{c(1 - P_F)}{c(P_D - P_F) + (1 - P_D)d}$ $\rho_1 = 1, \quad \rho_2 = \frac{\mu - P_D\beta}{(1 - P_D)\beta}$	②
$\mu/\beta < P_D$	$\psi = \frac{cP_F}{P_Dd - c(P_D - P_F)}$ $\rho_1 = \mu/(P_D\beta), \quad \rho_2 = 0$	①
		④

The equilibrium policies given in Result 4 are characterized by the ratios c/d and μ/β . If the cost of manual investigation is greater than the benefit from it ($c/d > 1$), then the firm chooses not to investigate at all. Knowing that the firm will not investigate, the user always commits fraud in this case. However if $c/d < 1$ and if the benefit from fraud is higher than its cost to the user ($\mu/\beta > 1$), then the user commits fraud all the time, which, in turn, causes the firm to investigate all transactions. When both μ/β and c/d are less than one, there is no pure strategy equilibrium for the game. This is because if the firm does not investigate, the user strictly prefers committing fraud, and therefore the firm is better off investigating all transactions. If the firm investigates all the time (signal and no-signal cases), however, the user prefers not committing fraud, implying that the firm is better off not investigating at all. In this case both firm and user randomize their actions over their strategy spaces to make each other indifferent between any of their pure strategies. For instance, choosing to commit fraud with probability $(c(1 - P_F))/(c(P_D - P_F) + (1 - P_D)d)$ in Region 2 makes the expected gain from investigation equal to the expected cost of investigation for the firm. Similarly, the firm employs an investigation strategy in which it investigates every transaction that generates a signal and $(\mu - P_D\beta)/((1 - P_D)\beta)$ fraction

of transactions that do not generate a signal from the IDS in Region 2. This makes the expected gain from fraud to the user equal to his expected penalty if caught. This is the equilibrium, since neither the firm nor the user has an incentive to deviate unilaterally.

We can explain the equilibria, especially the mixed strategy, in the following way. A pure strategy equilibrium constitutes a rule that tells the players what action to choose, whereas a mixed strategy constitutes a rule that tells them what dice to throw in order to choose an action (Rasmusen 1998). Mixed strategies occur frequently in real life. In American football games, for example, the offensive team has to decide whether to pass or to run. Passing generally gains more yards, but what is most important is to choose an action that is not expected by the other team. Teams decide to run part of the time and pass part of the time in a way that seems random to observers but that is rational to game theorists. The IRS randomly selects which tax returns to audit. Telephone companies randomly monitor their operators' conversations to discover whether or not they are polite. In our context, the user randomizes over his actions, fraud or no fraud, whereas the firm randomizes over its actions, investigate or not investigate, for each information set. There are two ways to interpret mixed strategy equilibria in our situation. First is to assume that there are many users, with identical tastes and payoff functions. Let's say that in a mixed strategy equilibrium each user chooses to commit fraud with probability 0.2, and the firm chooses to investigate with probability 0.6 when there is a signal and to not investigate when there is no signal. That means that 20% of the users follow the pure strategy of committing fraud and 80% choose the pure strategy of not committing fraud. Individual characteristics outside the model could determine who chooses what actions at an instance. Another interpretation of mixed strategies, which works even in the single-user case, is to assume that the user is drawn from a population of users and that the firm does not know his characteristics. The firm only knows that there are two types, honest user and fraudulent user. Then, there is a 20% chance that the user is a fraudulent user if the firm chooses to monitor 60% of the transactions that generate signals and none of the transactions that do not generate signals.

Table 4 Firm's Expected Cost

Region	Firm's expected cost
1	$cP_F d / (P_D d - c(P_D - P_F))$
2	c
3	c
4	d

Table 4 shows the firm's expected cost under the optimal manual investigation policy for different regions.

The expected cost is the same in Regions 2 and 3, although the fraud probability and investigation rate are higher in Region 3 than in Region 2. This occurs because in Region 3 the higher fraud probability is offset by the higher detection that results from a higher investigation rate. Result 5 is apparent from Table 4.

RESULT 5. Classification and detection software is valuable and, as a result, configuration is important only if benefit to cost ratio of fraud for the user (μ/β), and cost to benefit ratio of manual investigation for the firm (c/d) are both less than one.

Configuration is important only in Regions 1 and 2. Through configuration, the firm can choose which region to lie in by specifying P_D and P_F . Subtracting the expected cost in Region 1 from that in Region 2 gives

$$\frac{c(P_D - P_F)(d - c)(d(1 - P_D) + c(P_D - P_F))}{(d(1 - P_D) + c(P_D - P_F))(P_D(d - c) + cP_F)} \geq 0.$$

Hence, the firm will choose its configuration so that the equilibrium point lies in Region 1. Next, the firm should decide where to lie within Region 1. Writing the cost expression in Region 1 as a function of P_D gives

$$\frac{cd\sqrt{P_D}}{c\sqrt{P_D} + P_D(d - c)},$$

$$\frac{\partial(\cdot)}{\partial P_D} = \frac{cd\sqrt{P_D}(1 - r)(d - c)}{r[c\sqrt{P_D} + P_D(d - c)]^2} \geq 0. \quad (12)$$

This derivative implies that the firm will choose to set P_D as small as possible. Since the firm wants to be in Region 1, the firm sets P_D of its detection system to μ/β .⁴ That is, the optimal configuration for the

⁴ Actually, the firm will set P_D to $\mu/\beta + \varepsilon$, $\varepsilon > 0$, where ε is an infinitesimally small number.

detection software is

$$P_D^* = \frac{\mu}{\beta} \quad (13)$$

$$P_F^* = \left[\frac{\mu}{\beta} \right]^{1/r}. \quad (14)$$

Substituting the above optimal configuration point into the cost expression gives an expected cost for the firm of

$$\frac{d}{1 + (\mu/\beta)^{1-1/r}(d/c - 1)}.$$

It is interesting to note that under optimal configuration the firm will investigate all and only those users that generate signals. That is, the optimal configuration discriminates the users perfectly for manual investigation purposes.

4.2.2. Sequential Game. In this game the firm chooses the configuration point and inspection strategy first and then the user chooses his best strategy. Again, as in the earlier analysis, we assume that the firm configures its software first and then determines the optimal monitoring strategy for the given configuration. So the firm, as the leader of the game, conjectures the user's best response in its determination of configuration point and investigation strategy. We again use backward induction to solve the game. Let P_D and P_F be the configuration points and ρ_1 and ρ_2 be the firm's manual investigation probabilities in the signal and no-signal states respectively. The firm estimates the user's best response from his payoff function as

$$\begin{cases} \psi^* = 1 & \text{if } \frac{\mu}{\beta} > \rho_1 P_D + \rho_2 (1 - P_D) \\ \psi^* = 0 & \text{if } \frac{\mu}{\beta} < \rho_1 P_D + \rho_2 (1 - P_D) \\ \psi^* \in [0, 1] & \text{if } \frac{\mu}{\beta} = \rho_1 P_D + \rho_2 (1 - P_D). \end{cases}$$

When $\mu/\beta > \rho_1 P_D + \rho_2 (1 - P_D)$, the user commits fraud with a probability of one. In this case $\partial F / \partial \rho_1 |_{\psi \rightarrow 1} = -P_D(d - c) < 0$, $\partial F / \partial \rho_2 |_{\psi \rightarrow 1} = -(1 - P_D)(d - c) < 0$, which shows that the firm chooses $\rho_1 = 1$, $\rho_2 = 1$. However, the detection software is not valuable and configuration is not relevant under this condition. Therefore, we assume $\mu/\beta \leq \rho_1 P_D + \rho_2 (1 - P_D)$ for the configuration problem. If $\mu/\beta < \rho_1 P_D +$

$\rho_2(1 - P_D)$, then the user does not commit fraud at all. If $\mu/\beta = \rho_1 P_D + \rho_2(1 - P_D)$, then the user is indifferent between committing fraud and not committing fraud. We summarize this finding in Result 6.

RESULT 6. Classification and detection software is not valuable and as a result configuration is not important if benefit to cost ratio of fraud for the user (μ/β) is greater than one. If benefit to cost of fraud for the user (μ/β) is less than one, configuration is important only when $\mu/\beta \leq \rho_1 P_D + \rho_2(1 - P_D)$.

Result 6 implies that firm should select ρ_1 , ρ_2 , and P_D such that $\rho_1 P_D + \rho_2(1 - P_D) \geq \mu/\beta$. Now there are two cases: $\rho_1 P_D + \rho_2(1 - P_D) > \mu/\beta$ and $\rho_1 P_D + \rho_2(1 - P_D) = \mu/\beta$.

Case 1. $\rho_1 P_D + \rho_2(1 - P_D) > \mu/\beta$. If $\rho_1 P_D + \rho_2(1 - P_D) > \mu/\beta$ then $\psi^* = 0$. We can use (10) to write the firm's cost as

$$\begin{aligned} F(\rho_1, \rho_2; \psi^* = 0) &= (\rho_1 - \rho_2)cP_F + \rho_2c \\ &= c(\rho_2 + (\rho_1 - \rho_2)P_D^{1/r}). \end{aligned} \quad (15)$$

Then we can write firm's optimization problem for determining the monitoring strategy as

$$\begin{aligned} \text{Min}_{\rho_1, \rho_2} \quad & c(\rho_2 + (\rho_1 - \rho_2)P_D^{1/r}) \\ \text{s.t.} \quad & \rho_2 + (\rho_1 - \rho_2)P_D \geq \frac{\mu}{\beta} + \varepsilon > \frac{\mu}{\beta}. \end{aligned}$$

We can rewrite the objective function as $c\{(\rho_2 + (\rho_1 - \rho_2)P_D) - ((\rho_1 - \rho_2)(P_D - P_D^{1/r}))\}$. The minimum of this expression is achieved when ρ_1 and ρ_2 minimize $(\rho_2 + (\rho_1 - \rho_2)P_D)$ and maximize $(\rho_1 - \rho_2)(P_D - P_D^{1/r})$ simultaneously, of course subject to the constraint $\rho_2 + (\rho_1 - \rho_2)P_D > \mu/\beta$. The minimum feasible value of $(\rho_2 + (\rho_1 - \rho_2)P_D)$ is $\mu/\beta + \varepsilon$, where ε is an infinitesimally small positive number, and the maximum feasible value of $(\rho_1 - \rho_2)(P_D - P_D^{1/r})$ is $(P_D - P_D^{1/r})$, which is achieved when $(\rho_1 - \rho_2)$ has the maximum feasible value. Thus, the maximum value of $(\rho_1 - \rho_2)$ that satisfies $\rho_2 + (\rho_1 - \rho_2)P_D = \mu/\beta + \varepsilon$ results in the minimum value for the objective function. It is straightforward to see that

$$\rho_1 = \frac{\mu/\beta + \varepsilon}{P_D}; \quad \rho_2 = 0$$

is thus the optimal solution. Substituting these values in the firm's cost expression gives a cost

of $c\mu P_D^{(1-r)/r}/\beta$.⁵ Taking derivative with respect to P_D gives

$$\frac{\partial(\cdot)}{\partial P_D} = \frac{(1-r)c\mu P_D^{(1-2r)/r}}{r\beta}. \quad (16)$$

The derivative is always positive, so the firm will keep P_D as small as possible. However, $P_D \geq \mu/\beta$ must hold to satisfy $\rho_1 \leq 1$. Hence, the firm sets the value of P_D to μ/β . Thus, the equilibrium solution of the sequential game when $\rho_1 P_D + \rho_2(1 - P_D) > \mu/\beta$ is as follows:

$$\begin{cases} \rho_1^* = 1 \\ \rho_2^* = 0 \\ P_D^* = \mu/\beta \\ P_F^* = (\mu/\beta)^{1/r} \\ \psi^* = 0 \\ \text{MinimumCost} = c(\mu/\beta)^{1/r}. \end{cases}$$

Case 2. $\rho_1 P_D + \rho_2(1 - P_D) = \mu/\beta$. Assume that the firm sets ρ_1 , ρ_2 , and P_D such that the above constraint is satisfied. In this case, the user is indifferent between committing fraud and not committing fraud. We can show that the firm's cost is increasing in ψ , i.e., $(\partial F(\rho_1, \rho_2, \psi, P_D))/\partial \psi > 0$, when $\rho_1 \geq \rho_2$ (Result 3). Thus, for any given ρ_1 , ρ_2 ($\rho_1 \geq \rho_2$), and P_D , the firm achieves the minimum cost when $\psi = 0$. We also know that among all strategies that result in $\psi = 0$, the optimal strategies given for Case 1 offer the minimum cost to the firm. Thus, the optimal strategy for Case 2 cannot be better than that of Case 1.⁶ Thus, as the leader of the sequential game, the firm will never choose Case 2. That is, it will investigate all transactions that generate signals and will not investigate any transaction that does not generate a signal. The firm will set a detection probability very close to μ/β .

5. Comparison of Game Theory and Decision Theory Approaches

Having derived the optimal configurations under the decision theory and game theory approaches, we can

⁵Since $\varepsilon \cong 0$, we ignore this from the cost and subsequent expressions.

⁶Case 1 and Case 2 yield the same cost to the firm if the user does not commit fraud in Case 2 even though he is indifferent between committing fraud and not committing fraud. Thus, as long as there is a nonzero probability of the user committing fraud, Case 1 will be strictly better than Case 2 for the firm.

now compare the firm’s realized costs under these approaches. We use the following definitions for our comparisons.

$P_D^D \equiv$ optimal probability of detection under decision theory approach

$P_F^D \equiv$ optimal probability of false alarm under decision theory approach

$P_D^G \equiv$ optimal probability of detection under the game theory approach, simultaneous or sequential, used for comparison

$P_F^G \equiv$ optimal probability of false alarm under game theory approach, simultaneous or sequential, used for comparison

$\psi_R \equiv$ realized probability of fraud

$\psi_D \equiv$ estimated probability of fraud by the firm under decision theory approach

The realized cost at the optimal configuration under the decision theory approach is given by $d\psi_R + (c - d) \cdot \psi_R P_D^D + c(1 - \psi_R)P_F^D$. We need to determine the realized probability of fraud in order to compute the realized cost under the decision theory approach. Since the user is strategic, he will adjust his strategy depending on the strategy used by the firm. That is, the user will choose ψ_R based on his utility. From the user’s utility function, we find that the user will commit fraud, i.e., $\psi_R = 1$, if $P_D^D < \mu/\beta$, and will not commit fraud, i.e., $\psi_R = 0$, otherwise. Thus, the realized cost under decision theory approach is computed to be the following:⁷

$$\left\{ \begin{array}{l} cP_F^D \text{ if } \frac{\mu}{\beta} < P_D^D = \min \left\{ \left(\frac{c(1 - \psi_D)}{r\psi_D(d - c)} \right)^{r/(r-1)}, 1 \right\} \\ \text{(i.e., } \psi_R = 0) \\ cP_D^D + d(1 - P_D^D) \\ \text{if } \frac{\mu}{\beta} > P_D^D = \min \left\{ \left(\frac{c(1 - \psi_D)}{r\psi_D(d - c)} \right)^{r/(r-1)}, 1 \right\} \\ \text{(i.e., } \psi_R = 1). \end{array} \right.$$

⁷ The costs given in these expressions are for the case when the firm’s estimate of fraud probability is ψ_D . It is possible that the firm learns and updates its estimate of the fraud probability over time. In that case, the realized cost will be different in each period. Our results hold for any given period in which the fraud probability is estimated to be ψ_D . We explore the implication of learning by the firm for our results in §6.

Under game theory, the realized costs are, respectively, $d/(1 + (\mu/\beta)^{1-1/r}(d/c - 1))$ and $c(\mu/\beta)^{1/r}$ in the simultaneous and sequential games.

A comparison of the realized costs under the decision theory and game theory approaches gives Result 7 (the proof is in the appendix).

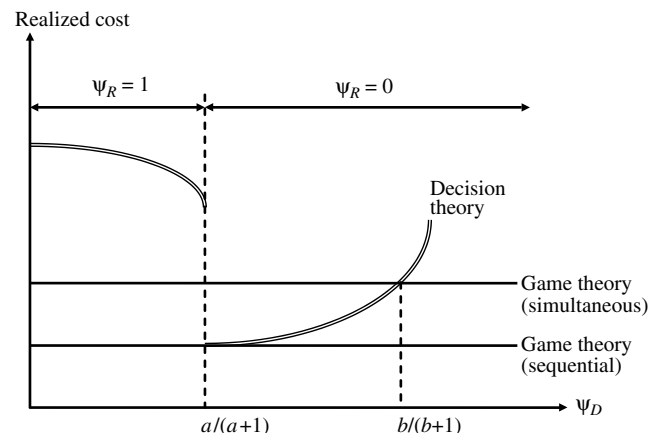
RESULT 7. (i) Under the assumptions of simultaneous game, configuration using game theory results in lower cost than configuration using decision theory unless

$$\frac{c}{r(d - c)} \left(\frac{\mu}{\beta} \right)^{(1-r)/r} < \frac{\psi_D}{1 - \psi_D} < \frac{c}{r(d - c)} \left(\frac{\mu}{\beta} \right)^{(1-r)/r} \left(\left(\frac{c}{d} \right) \left(\frac{\mu}{\beta} \right)^{1/r} + \left(\frac{\mu}{\beta} \right) \left(1 - \frac{c}{d} \right) \right)^{r-1}.$$

(ii) Under the assumptions of sequential game, configuration using game theory always results in same or lower cost than the configuration using decision theory.

The above result can be illustrated using Figure 4, which plots the realized costs under the three approaches for different values of ψ_D . It also shows the realized values of fraud probability in different regions. Game theory results in a lower cost to the firm in all cases except when a simultaneous game is played, and the firm uses a fraud probability that lies in the interval $(a/(a + 1), b/(b + 1))$, where a and b are, respectively, the lower and upper limits for $\psi_D/(1 - \psi_D)$ given in Result 7, under the decision theory approach. It is intuitive that the firm realizes a

Figure 4 Comparison of Realized Costs for Configuring a Detection Program



lower cost when it uses game theory instead of decision theory under a sequential game. The firm is the leader and the user is the follower in the sequential game. That is, the user reacts to the decisions made by the firm. In this scenario, the firm will always do better if it anticipates what the user's reaction for its action will be as opposed to how the firm will do if it does not take user's reaction into consideration while making its decisions. By using game theory, the firm exploits its first-mover advantage in the sequential game scenario.

If the simultaneous game is played, then decision theory performs worse than game theory if the firm does not deter the user from committing fraud. In other words, if the firm estimates a low probability of fraud and configures the program with a low probability of detection, the user ends up committing fraud. In this scenario, the firm incurs a higher cost when it uses decision theory than when it uses game theory in which the probability of fraud is less than one. If the firm is successful in deterring the user from committing fraud when it uses the decision theory approach, then its cost under decision theory is lower than under game theory only when it does not *overdeter* the user, i.e., when it does not configure the program with a very high probability of detection. When the firm's estimate of the fraud probability is very high, it will configure the program to have a high probability of detection, which will deter the user. Although the firm is successful in deterring the user, it also incurs a high cost from the high rate of false signals from the software. Consequently, the firm realizes a lower cost under game theory even though the game theory solution does not deter the user from committing fraud.

The result that decision theory might perform better than game theory when the simultaneous game is played (though only when the condition stated in Result 7 is satisfied) is counterintuitive. One would expect that the firm would be better off when it makes the decisions by anticipating the reaction of the user, as in game theory, as compared to when it does not take the user reaction into consideration, as in decision theory. The counterintuitive result can be explained as follows: We can show that when $\psi_D = a/(a+1)$, i.e., when $\psi_D/(1-\psi_D)$ is equal to the lower limit of the condition stated in Result 7, the firm uses

the same configuration under decision theory as game theory. Under this scenario, although the final outcomes in terms of configuration are identical under the decision theory and game theory methods, the models used to arrive at these outcomes are different. If the firm is lucky enough to configure the software correctly under the decision theory approach, then the firm will incur the lowest possible cost because, essentially, the firm behaves as though it knows the realized fraud probability, which leads to a first-mover advantage for the firm. Recall that when the firm uses the decision theory approach it makes its decisions once, based on its estimate of the fraud probability. The strategic user adjusts his strategy based on the firm's decision. Thus, the firm acts as a leader and the user acts as a follower under the decision theory approach. This is evident from Figure 4, which shows that the realized costs under decision theory and sequential game theory approaches are equal when $\psi_D = a/(a+1)$. Neither party has a first-mover advantage in the simultaneous game. Hence, the firm incurs a higher cost under simultaneous game than under sequential game or decision theory approach. This first-mover advantage to the firm under the decision theory approach exists as long as the firm does not overconfigure (i.e., set a high detection or false-positive rate). If the firm overconfigures, the first-mover advantage is offset by the higher costs associated with more frequent false alarms.

We illustrate the results of our comparison of decision theory and game theory solutions using the following numerical example. We use the following parameter values: $d = 650$, $c = 200$, $\beta = 1,000$, $\mu = 800$, and $r = 0.1$. The estimated probability of fraud was varied between 0.1 and 0.8.

As can be seen from Table 5, the firm configures the detection software to have a high detection rate as the assessed risk of fraud increases in the decision theory approach. For example, the firm chooses to set the detection rate at 0.6637 if the assessed fraud probability is equal to 0.1. However this rate is set at 0.9884 if the estimated probability increases to 0.8. Higher estimated fraud probability also leads to higher expected cost, as expected.

When we compare expected and realized costs, we see that the expected cost exceeds the realized cost if the realized fraud probability turns out to

Table 5 A Numerical Example for the Configuration Problem

ψ_R	ψ_D	P_F	P_D	Expected cost	Realized cost
Panel A: Decision theory approach					
1	0.1000	0.0166	0.6637	38.1188	351.3202
1	0.1500	0.0277	0.6987	55.0514	335.5658
1	0.2000	0.0409	0.7263	71.1685	323.1585
1	0.2500	0.0562	0.7499	86.5720	312.5423
1	0.3000	0.0744	0.7711	101.3063	302.9864
1	0.3500	0.0958	0.7909	115.3830	294.0729
0	0.3736	0.1074	0.8000	121.8063	21.4748
0	0.4000	0.1215	0.8099	128.7892	24.2983
0	0.4500	0.1525	0.8286	141.4904	30.5070
0	0.4908	0.1830	0.8438	151.2811	36.5913
0	0.5000	0.1906	0.8473	153.4284	38.1270
0	0.5500	0.2383	0.8664	164.5160	47.6504
0	0.6000	0.2991	0.8863	174.6265	59.8260
0	0.6500	0.3792	0.9076	183.5757	75.8490
0	0.7000	0.4887	0.9309	191.0877	97.7453
0	0.7500	0.6462	0.9573	196.7297	129.2312
0	0.8000	0.8895	0.9884	199.7709	177.9051
		ψ_R	P_F	P_D	Game cost
Panel B: Game theory approach					
Simultaneous		0.056294	0.1074	0.8	36.5913
Sequential		0	0.1074	0.8	21.4748

be higher than the estimated probability. Thus, if the firm underestimates the fraud probability, then it incurs a higher-than-expected cost. Similarly, an over-estimation of fraud probability has an opposite effect.

Panel B in Table 5 shows the solutions under game theory. Although the firm configures the detection program at the same level in both simultaneous and sequential games, the firm incurs a lower cost under the sequential game because of the first-mover advantage the firm enjoys in the sequential game. Comparing with the decision analysis results shown in Panel A, we see that game theory results in lower cost than decision theory when realized fraud probability is equal to one, irrespective of how the game is played. However if the realized fraud probability under decision theory analysis is equal to zero—i.e., when configuration based on decision theory is successful in deterring the user from fraudulent activity—decision theory gives better results than simultaneous game theory (simultaneous case) if the estimated fraud probability is less than 0.4908.

6. Limitations of the Game Theory Approach

Our results clearly show that under most circumstances a firm, when faced with a strategic adversary, realizes a lower cost when it uses the game theory as opposed to the decision theory to configure detection programs. However, there seems to be a dichotomy between our results and current business practices; this dichotomy seems to favor a decision theory approach. We hypothesize two reasons for the dichotomy. One reason could be that firms are truly unaware of the potential benefits of applying the game theory approach for the fraud detection problem. We believe that this paper provides insights to firms on how and why game theory performs better than does decision theory. Another reason could be that firms view decision theory as a simplification of the more complex game theory approach. The decision theory and game theory models require estimation of several parameters. In some sense, the game theory model requires deeper user-specific parameters that are more difficult to obtain. For instance, game theory requires that the firm know the user utility and penalty parameters, whereas decision theory requires only the final outcome of users’ decisions in terms of fraud probabilities. We believe that fraud probabilities are easier than utility to obtain because firms can use historical log records to estimate fraud probability. We conjecture that the difficulty in estimation of user-specific parameters is one reason why firms prefer to use decision theory instead of game theory.

The game theory model requires the firm to estimate two user-specific parameters: benefit from undetected fraud and penalty if fraud is detected. The estimation methodology as well as the difficulty in estimation depends on the problem context. With respect to the user’s benefit from undetected fraud, it is relatively easy to estimate the benefit to the user in contexts where the user benefit is equal to damage incurred by the firm. For instance, in the case of credit card fraud, one could approximate the user’s benefit to be the credit limit on the account. Similarly, in the context of tax fraud, the user’s benefit can be equated to the maximum tax loss incurred by IRS. In other contexts, cost of fraud to the firm might not be the same as benefit to the user. In such cases, the

firm might not have any data to estimate the user benefit. For instance, in the information security context, the value of information for the firm and the hacker can be different. Gordon and Loeb (2001) state that IT security frauds are often committed to gather competitive intelligence data. In such cases, the dollar amount competitors are willing to pay for such information can be used as a basis for estimating the user's benefit from hacking. Estimation of penalty incurred by detected fraudsters is easier than the user's benefit. The penalties for different kinds of fraud are documented in laws and regulations. For example, the Computer Crime and Intellectual Property Section of the Criminal Division of the U.S Department of Justice (2004) provides punishment given to hackers in terms of sentencing, fine, forfeiture, and restitution. Estimation of user-specific parameters can be difficult initially, but could become accurate over time. If good estimates of values of these parameters are unavailable, our model can still be used to perform sensitivity analysis. We should recognize that quantitative models such as ours are often best judged on the usefulness and the validity of implications they generate. It is encouraging to note that our analytical results suggest that the model's results are consistent with our intuitions, an aspect that increases our confidence in the validity of the model to the configuration problem.

A related question is what happens when the firm under- or overestimates the benefit-to-cost ratio for the firm. This might occur because of the difficulties in estimation as discussed previously. If the firm underestimates the benefit-to-cost ratio and applies the optimal configuration strategy consistent with the lower estimate, then our model shows that the user will commit fraud with a probability of one. This leads to the firm incurring a higher cost because of the higher-than-expected damage. If the firm overestimates the benefit-to-cost ratio, then again it incurs more cost because of the higher level of monitoring. However, if the firm monitors the fraud activity and updates its estimate, it could adjust its strategy in subsequent periods. Eventually, the game will reach the derived equilibrium from which neither will have any incentive to deviate. Another reason for error in the firm's estimation of user's benefit-to-cost ratio could be the user's incentive to make the firm believe that

the ratio is low. However, for the same reasons cited above, although the user might trick the firm and make himself better off in the first period, this cannot be sustained in the long run.

In this paper we assumed that the model parameters are common knowledge to the firm and users. One area that seems particularly interesting for configuration of detection systems is the area of games with incomplete information, in which either the firm or the user is uncertain about the other's payoffs. This perspective allows incorporation of uncertainty about the nature of the game being played. Many of these common knowledge-related assumptions have been analyzed by game theorists. Harsanyi (1967; 1968a, b) has shown that the lack of knowledge about payoffs does not alter the basic representation of a game and the qualitative nature of the results. If the firm is uncertain about the utility of fraud to the user, it can use its prior knowledge to come up with a monitoring strategy. For instance, assume that the firm believes that the utility can be high or low. We can represent the firm's beliefs with a subjective probability distribution, p_H and p_L . In other words, the firm places a $100p_H\%$ chance that it is dealing with a high-utility user and a $100p_L\%$ chance that it is dealing with a low-utility user. However the user knows his true utility. We can analyze this setting using Harsanyi transformation, in which nature makes the first move and chooses the user type in accordance with the firm's subjective probabilities. We leave the detailed analysis of this model to future research. However, we can speculate that the firm's uncertainty about the user's utility from fraud will make game theory solution less attractive than the one presented here. A higher level of uncertainty will reduce the firm's profit in the game theory setting, and the advantage of game theory over decision theory will diminish.

The advantages of the game theory approach over the decision theory approach might also diminish over time if the players learn about the opponent's likely strategies. In this paper we assumed that the user is always strategic and changes his strategies in response to the firm's strategies. The firm is nonstrategic when it uses the decision theory model. In our decision theory model, we assumed a one-shot setting in which the firm makes its configuration and investigation probabilities based on its estimate of the

hacker's probability of hacking. However, it might be more realistic to consider a multiperiod model in which the firm revises its estimates every period based on its observations of the hacker's strategy in previous periods. This learning has been analyzed in game theory. Fudenberg and Levine (1998, pp. 34–35, Proposition 2.2) show that when players learn but use a myopic approach every period, and if the empirical distribution over each player's choices converges, then the strategy profile is a Nash equilibrium. If convergence is achieved, the game theory-based configuration and decision theory-based configuration yield identical outcomes in the equilibrium. However, it should be noted that until the convergence is reached, our results apply and the game theory-based approach would result in lower cost than the decision theory-based approach, except when the condition in Result 7 is satisfied. Fudenberg and Levine also state that the empirical distributions need not always converge.

In addition, the type of learning model used also has an effect on the convergence. For instance, an open issue is what type of learning model is appropriate for our context. Some of the questions include "Is learning based only on the most recent move or is it based on the history of all moves?", "What relative weights should be assigned to different move?", and "How are the probabilities updated based on the history?" Answers to these questions will provide valuable additional insights into the tradeoff between game theory- and decision theory-based approaches to the detection software configuration problem. A valuable extension of our research is to analyze a dynamic model that incorporates learning within the context of the decision theory approach and compares the results of this model with the game theory approach.

7. Conclusions

Business firms have been automating a variety of business processes that involve critical real-time decision making. Many of these decisions involve classification or detection of some sort. False-positive and false-negative errors in detection could affect the value of these systems significantly. Consequently, firms need to configure these programs carefully. In this paper we presented two models—the first

based on decision theory and the second based on game theory—to assist firms in the configuration process of detection software. In the decision theory approach, the firm estimates the fraud probability exogenously and assumes that its actions do not alter users' behavior. In the game theory approach, the firm makes its decisions by assuming that the firm's decisions alter user behavior. We considered both simultaneous and sequential games. We found that firms incur lower cost under most situations when they use game theory as opposed to decision theory. The decision theory approach results in a lower cost only under the simultaneous game assumptions and when the firm neither underconfigures nor overconfigures the detection software by a significant amount.

The results of our model suggest that using the decision theory approach to manage fraud, which ignores the reactions of strategic adversaries, can cause significant harm to firms. Traditionally, software configuration is viewed as a firm's internal problem that affects only the firm. Although this is true of software that deals with operational problems such as transaction processing systems, strategic applications require modeling the strategic interactions. Consequently, the design and configuration of such software need to take into account the effect of software configuration on user behavior.

Appendix

PROOF OF RESULT 3. To prove that $\rho_1 \geq \rho_2$, all we need to show is $\partial F_S / \partial \rho_1 \leq \partial F_N / \partial \rho_2$, since we try to maximize F . Taking the partial derivatives we get

$$\frac{\partial F_S}{\partial \rho_1} = c - \eta_1 d \quad (\text{A1})$$

$$\frac{\partial F_N}{\partial \rho_1} = c - \eta_2 d. \quad (\text{A2})$$

Algebraic manipulations show that $\eta_1 \geq \eta_2$ when $P_D \geq P_F$, and hence $\partial F_S / \partial \rho_1 \leq \partial F_N / \partial \rho_2$.

To show the second part, suppose that $0 < \rho_1 \leq \rho_2 < 1$ is possible. Then the first-order conditions for the firm's optimization in the signal and no-signal states must be satisfied for the same value of ψ . Equating (A1) and (A2) to zero and solving for ψ gives

$$\psi = \frac{c P_F}{P_D d - c(P_D - P_F)} \quad (\text{A3})$$

$$\psi = \frac{c(1 - P_F)}{c(P_D - P_F) + (1 - P_D)d}. \quad (\text{A4})$$

ψ given by (A3) and (A4) are equal to each other iff $c = d$. However when $c = d$, $\rho_1 = \rho_2 = 0$, which contradicts our supposition. \square

PROOF OF RESULT 4. A user's optimization condition is

$$\frac{\partial H}{\partial \psi} = \mu - \beta(\rho_1 P_D + \rho_2(1 - P_D)) = 0. \quad (\text{A5})$$

The firm's optimization conditions are

$$\frac{\partial F}{\partial \rho_1} = -P_D d \psi + c(P_F + (P_D - P_F)\psi) = 0 \quad (\text{A6})$$

$$\frac{\partial F}{\partial \rho_2} = -(1 - P_D)d \psi + c((1 - P_F) - (P_D - P_F)\psi) = 0. \quad (\text{A7})$$

There are parameter values for which all three conditions are not satisfied simultaneously. Consequently, we also need to consider the corner solutions. We analyze all possible solutions below.

(a) $\psi = 1$, $\rho_1 = 0$, $\rho_2 = 0$ is an equilibrium iff $\partial F / \partial \rho_i |_{\psi=1} > 0$ for $i = 1, 2$ and $\partial H / \partial \psi |_{\rho_1, \rho_2=0} > 0$. These conditions are both satisfied when $c > d$.

(b) $\psi = 1$, $\rho_1 = 1$, $\rho_2 = 1$ is an equilibrium iff $\partial F / \partial \rho_i |_{\psi=1} < 0$ for $i = 1, 2$ and $\partial H / \partial \psi |_{\rho_1, \rho_2=1} > 0$. These conditions are satisfied only if (i) $c < d$ and (ii) $\mu > \beta$.

(c) $\psi = 1$, $\rho_1 = 1$, $\rho_2 = 0$ is an equilibrium iff $\partial F / \partial \rho_1 |_{\psi=1} < 0$ and $\partial F / \partial \rho_2 |_{\psi=1} > 0$ and $\partial H / \partial \psi |_{\rho_1=1, \rho_2=0} > 0$. These conditions are satisfied only if (i) $c < d$, and (ii) $c > d$, and (iii) $\mu > P_D \beta$. Since (i) and (ii) contradict each other, $\psi = 1$, $\rho_1 = 1$, $\rho_2 = 0$ is not a feasible equilibrium.

Since ρ_1 and ρ_2 cannot both be less than 1 and $\rho_1 \geq \rho_2$ (Result 3), the other possible equilibriums are $(0 < \psi < 1, \rho_1 = 1, 0 < \rho_2 < 1)$ and $(0 < \psi < 1, 0 < \rho_1 < 1, \rho_2 = 0)$.

(d) If $(0 < \psi < 1, \rho_1 = 1, 0 < \rho_2 < 1)$ is an equilibrium, first-order condition for the firm with respect to ρ_2 and first-order condition for the user must be satisfied. Equating (A5) to zero gives the relationship between ρ_1 and ρ_2 as follows:

$$\frac{\mu}{\beta} = P_D(\rho_1 - \rho_2) + \rho_2. \quad (\text{A8})$$

Plugging the equilibrium value of $\rho_1 = 1$ into the above equation and solving for ρ_2 gives

$$\rho_2 = \frac{\mu}{\beta(1 - P_D)} - \frac{P_D}{1 - P_D}. \quad (\text{A9})$$

Equating Equation (A7) to zero and solving for ψ we get

$$\psi = \frac{c(1 - P_F)}{c(P_D - P_F) + (1 - P_D)d}. \quad (\text{A10})$$

There are two constraints for this equilibrium to exist. First, equilibrium values found in (A9) and (A10) must be between zero and one. Second, equilibrium values must make the derivative of the payoff for the firm with respect to ρ_1 positive (since $\rho_1 = 1$). These constraints yield

$$P_D < \frac{\mu}{\beta} < 1 \quad \text{and} \quad c < d.$$

(e) If $(0 < \psi < 1, 0 < \rho_1 < 1, \rho_2 = 0)$ is an equilibrium, first-order condition for the firm with respect to ρ_1 and first-order condition for the user must be satisfied. Equating (A5) to zero gives the relationship between ρ_1 and ρ_2 as given in (A8). Plugging the equilibrium value of $\rho_2 = 0$ into Equation (A5) and solving for ρ_1 gives

$$\rho_1 = \frac{\mu}{P_D \beta}. \quad (\text{A11})$$

Equating (A6) to zero and solving for ψ we get

$$\psi = \frac{c P_F}{P_D d - c(P_D - P_F)}. \quad (\text{A12})$$

There are also two constraints for this equilibrium. First, equilibrium values found in (A11) and (A12) must be between zero and one. Second, equilibrium values must make the derivative of payoff for the firm with respect to ρ_2 negative (since $\rho_2 = 0$). These constraints yield $0 < \mu/\beta < P_D$ and $c < d$. \square

PROOF OF RESULT 7. Assume that $\mu/\beta > P_D^D$ holds for the simultaneous game (i.e., $\psi_R = 1$). Realized cost under decision theory < Realized cost under game theory iff

$$c P_D^D + d(1 - P_D^D) < \frac{d}{1 + (\mu/\beta)^{1-1/r}(d/c - 1)}.$$

This inequality can be written as

$$P_D^D > \frac{\mu d}{c\beta(\mu/\beta)^{1/r} + \mu(d - c)}.$$

Since

$$\frac{\mu d}{c\beta(\mu/\beta)^{1/r} + \mu(d - c)} > \frac{\mu}{\beta},$$

it contradicts $\mu/\beta > P_D^D$. Hence, the game theory solution always dominates the decision theory solution when $\mu/\beta > P_D^D$.

Assume that $\mu/\beta < P_D^D$ holds for the simultaneous game (i.e., $\psi_R = 0$). Realized cost under decision theory < Realized cost under game theory iff

$$c P_F^D < \frac{d}{1 + (\mu/\beta)^{1-1/r}(d/c - 1)}.$$

Plugging the equilibrium value of P_F^D and solving for ψ_D gives

$$\begin{aligned} \frac{\psi_D}{1 - \psi_D} < \frac{c}{r(d - c)} \left(\frac{\mu}{\beta} \right)^{(1-r)/r} \\ \cdot \left(\left(\frac{c}{d} \right) \left(\frac{\mu}{\beta} \right)^{1/r} + \left(\frac{\mu}{\beta} \right) \left(1 - \frac{c}{d} \right) \right)^{r-1}. \end{aligned}$$

We can write $\mu/\beta < P_D^D$ in terms of ψ_D as

$$\frac{\psi_D}{1 - \psi_D} > \frac{c}{r(d - c)} \left(\frac{\mu}{\beta} \right)^{(1-r)/r}.$$

Hence, the decision theory solution dominates the game theory solution when

$$\frac{c}{r(d-c)} \left(\frac{\mu}{\beta}\right)^{(1-r)/r} < \frac{\psi_D}{1-\psi_D} < \frac{c}{r(d-c)} \left(\frac{\mu}{\beta}\right)^{(1-r)/r} \cdot \left(\left(\frac{c}{d}\right) \left(\frac{\mu}{\beta}\right)^{1/r} + \left(\frac{\mu}{\beta}\right) \left(1 - \frac{c}{d}\right) \right)^{r-1}.$$

Assume that $\mu/\beta > P_D^D$ holds for the sequential game (i.e., $\psi_R = 1$). Realized cost under decision theory < Realized cost under game theory iff $cP_D^D + d(1 - P_D^D) < c(\mu/\beta)^{1/r}$. We can rewrite this condition as $d[1 - P_D^D] < c[(\mu/\beta)^{1/r} - P_D^D]$. Since $d > c$ and $1 > \mu/\beta$. This inequality never holds. So, the game theory solution always dominates the decision theory solution when $\mu > \beta P_D^D$.

Assume that $\mu/\beta < P_D^D$ holds for the sequential game (i.e., $\psi_R = 0$). Realized cost under decision theory < Realized cost under game theory iff $cP_F^D < c(\mu/\beta)^{1/r}$. Since $\mu/\beta < P_D^D$ implies that $(\mu/\beta)^r < P_F^D$, $cP_F^D < c(\mu/\beta)^{1/r}$ never holds. So, the game theory solution always dominates the decision theory solution when $\mu/\beta < P_D^D$. \square

References

- Alaric. 2003. The card fraud detection problem. Alaric Systems Ltd., http://www.alaric-systems.co.uk/fractals_problems.htm/.
- Allen, J., A. Christie, W. Fithen, J. McHugh, J. Pickel, E. Stoner. 2000. State of the practice of intrusion detection technologies. Technical report CMU/SEI-99-TR-028 ESC-99-028, Carnegie Mellon Software Engineering Institute, Pittsburgh, PA.
- Chavez, T. 2000. A decision-analytic stopping rule for validation of commercial software systems. *IEEE Trans. Software Engrg.* **26**(9) 907–918.
- CNN. 2002. Knives, guns, fake bombs elude airport security. CNN.com (March 26).
- Computer Crime and Intellectual Property Section. 2004. Criminal Division of the U.S. Department of Justice, <http://www.usdoj.gov/criminal/cybercriminal>.
- Durst, R., T. Champion, B. Witten, E. Miller, L. Spagnuolo. 1999. Testing and evaluating computer intrusion detection systems. *Comm. ACM* **42**(7) 53–61.
- Fellingham, J., P. Newman. 1985. Strategic considerations in auditing. *Accounting Rev.* **60**(October) 634–650.
- Fudenberg, D., D. Levine. 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Fudenberg, D., J. Tirole. 1993. *Game Theory*. MIT Press, Cambridge, MA.
- Gordon, L. A., M. P. Loeb. 2001. Using information security as a response to competitor analysis systems. *Comm. ACM* **44**(9) 70–75.
- Harsanyi, J. C. 1967. Games with incomplete information played by Bayesian players, I: Basic model. *Management Sci.* **14**(3) 159–182.
- Harsanyi, J. C. 1968a. Games with incomplete information played by Bayesian players, II: Bayesian equilibrium points. *Management Sci.* **14**(5) 320–334.
- Harsanyi, J. C. 1968b. Games with incomplete information played by Bayesian players, III: The basic probability distribution of the game. *Management Sci.* **14**(7) 486–502.
- Jajodia, S., J. Miller. 1993. Editor's preface. *J. Comput. Security* **16**(4) 43–53.
- Jonsson, E., T. Olovsson. 1997. A quantitative model of security intrusion process based on attacker behavior. *IEEE Trans. Software Engrg.* **23**(4) 235–245.
- Lee, W., W. Fan, M. Miller, S. Stolfo, E. Zadok. 2002. Toward cost-sensitive modeling for intrusion detection and response. *J. Comput. Security* **10**(1/2) 5–22.
- Lippmann, R., J. W. Haines, D. J. Fried, J. Korba, K. Das. 2000. The 1999 DARPA off-line intrusion detection evaluation. *Comput. Networks* **34**(4) 579–595.
- McHugh, J. 2000. Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory. *ACM Trans. Inform. System Security* **3**(4) 262–294.
- NIST 800-30. Risk management guide for information technology systems. National Institute of Standards and Technology Special Publication, Gaithersburg, MD.
- NMAB. 1998. Configuration management and performance verification of explosives-detection systems. Publication NMAB-482-3, National Academy Press, Washington, D.C.
- Ortalo, R., Y. Deswarte, M. Kaaniche. 1999. Experimenting with quantitative evaluation tools for monitoring operational security. *IEEE Trans. Software Engrg.* **25**(5) 633–650.
- Provost, F., T. Fawcett. 1997. Analysis and visualization of classifier performance: Comparison under imprecise class and cost distributions. *Proc. KDD-97*. AAAI Press, Newport Beach, CA, 43–48.
- Rasmusen, E. 1998. *Games and Information*, 2nd ed. Blackwell, Cambridge, MA.
- Sarkar, S., R. Sriram. 2001. Bayesian models for early warnings of bank failures. *Management Sci.* **47**(11) 1457–1475.
- Shapiro, G. P., B. Masand. 1999. Estimating campaign benefits and modeling lift. *Proc. KDD-99*. ACM Press, San Diego, CA, 185–193.
- Shapiro, G. P., S. Steingold. 2000. Measuring lift quality in database marketing. *SIGKDD Explorations* **2**(2) 81–86.
- Sriram, T. 2002. Blocking virus requests in Novell BorderManager's HTTP accelerator. Feature article. Novell Appnotes, <http://developer.novell.com/research/appnotes/>.
- Steingold, S., R. Wherry, G. P. Shapiro. 2001. Measuring real-time predictive models. *Proc. IEEE Internat. Conf. Data Mining*. San Jose, CA, 649–650.
- Trees, H. V. 2001. *Detection, Estimation and Modulation Theory-Part I*. John Wiley, New York.
- Ulvila, J. W., J. E. Gaffney. 2004. A decision analysis method for evaluating computer intrusion detection systems. *Decision Anal.* **1**(1) 35–50.