U.S. DEPARTMENT OF ENERGY

PNNL-19665

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Identifying at-Risk Employees: A Behavioral Model for Predicting Potential Insider Threats

FL Greitzer
LJ Kangas
CF Noonan
AC Dalton

September 30, 2010

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

DISCLAIMER

PACIFIC NORTHWEST NATIONAL
LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF
ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America
Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical
Information Service,
U.S. Department of Commerce, 5285 Port Royal Rd.,
Springfield, VA 22161
ph: (800) 553-6847
fax: (703) 605-6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/ordering.htm

PNNL-19665

Identifying at-Risk Employees: A Behavioral Model for Predicting Potential Insider Threats

FL Greitzer
LJ Kangas
CF Noonan
AC Dalton

September 2010

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Summary

A model was developed to assess employees' behavioral manifestations of a number of psychological and personality predispositions that are hypothesized to indicate an increased risk of insider abuse. This psychosocial model is based on case studies and research literature on factors and correlates associated with behavioral precursors of individuals committing insider crimes. In many of these crimes, managers and other coworkers observed that the offenders had exhibited signs of stress, disgruntlement, or other issues, but no alarms were raised. Barriers to using such psychosocial indicators include the inability to recognize the signs and the failure to record the behaviors so that they can be assessed.

The model has been implemented as a Bayesian belief network, designed with the help of human resources staff experienced in evaluating workplace behaviors. We conducted an experiment to assess the agreement of the model's risk assessment output with judgments of human resources and management professionals on the relative insider threat risks of a collection of sample scenarios. The model exhibited strong agreement with judgments of the human experts, suggesting that it has potential as a tool to raise an alarm about employees who pose higher insider threat risks. While additional testing is needed, we suggest that combining this type of analysis with more traditional cyber/workstation monitoring tools can ease the processing burden and improve performance of computer-assisted insider threat monitoring and detection.

Acknowledgments

The authors wish to sincerely thank a number of individuals from different parts of our organization who contributed in various ways to this work, including Deborah A. Frincke, Lead for the Information and Infrastructure Integrity Initiative at PNNL, and team members Tom Carroll, Ryan Hohimer, Duane Klotz, Patrick Paulson, Christine Ulibarri, and former colleague Mariah Zabriskie. This work was supported by the Information and Infrastructure Integrity Initiative of the Pacific Northwest National Laboratory. The Pacific Northwest National Laboratory is operated by Battelle for the U.S. Depart of Energy under Contract DE-AC06-76RL01830.

Acronyms and Abbreviations

ACAMP	Adaptive Cyberdefense using an Auto-associative Memory Paradigm
ANN	Artificial Neural Network
BN	Bayesian network
CERT	Computer Emergency Readiness Team
CSO	Chief Security Officer
CWB	counterproductive work behavior
ECPA	Electronic Communications Privacy Act
EU	European Union
FFM	Five Factor Model
HR	Human Resources
LR	Linear Regression
NIAC	National Infrastructure Advisory Council
PIPEDA	Personal Information Protection and Electronic Documents Act
PNNL	Pacific Northwest National Laboratory
RMSE	Root Mean Squared Error

Contents

Figures

Tables

1.0 Introduction

Espionage and sabotage involving computers and computer networks are among the most pressing cyber security challenges that threaten government and private sector information infrastructures. The annual e-Crime Watch Survey conducted by the Chief Security Officer (CSO) Magazine in conjunction with other institutions (CSO, U.S. Secret Service, Software Engineering Institute, CERT Program at Carnegie Mellon University and Deloitte, 2010) reveals that for both the government and commercial sectors, the most costly or damaging cybercrime attacks are caused by insiders such as current or former employees and contractors.

The insider threat refers to harmful acts that trusted insiders might carry out. For example, something that causes harm to the organization, or an unauthorized act that benefits the individual. The insider threat is manifested when human behaviors depart from established policies, regardless of whether it results from malice or disregard for security policies. The types of crimes and abuses associated with insider threats are significant; the most serious include espionage, sabotage, terrorism, embezzlement, extortion, bribery, and corruption. Malicious activities include an even broader range of exploits, such as copyright violations, negligent use of classified data, fraud, unauthorized access to sensitive information, and illicit communications with unauthorized recipients.

Surveys and studies conducted over the last decade and a half have consistently shown the critical nature of the problem in both government and private sectors. A 1997 Department of Defense Inspector General report (Department of Defense, 1997) found that 87 % of identified intruders into Department of Defense (DoD) information systems were either employees or others internal to the organization. The 2010 e-Crime survey showed that most insiders target proprietary information including intellectual property and customer or financial information (CSO, et al., 2010).

More generally, recent studies of cybercrime such as the e-Crime Watch Survey (see also Keeney, et al., 2005) in both government and commercial sectors reveal that the financial impact and operating losses due to insider intrusions are increasing. Among those companies experiencing security events, the majority (55%) report at least one insider event, which was an alarming increase from 39% in 2005. A recent report covering over 143 million data records collected by Verizon and the U.S. Secret Service analyzed a set of 141 confirmed breach cases in 2009 and found that 46% of data breaches were attributed to the work of insiders (Verizon and the U.S. Secret Service, 2010). Of these, 90% were the result of deliberate, malicious acts; six percent were attributed to inappropriate actions such as policy violations and other questionable behavior, and four percent to unintentional acts (Verizon and the U.S. Secret Service, 2010).

2.0 Relevant Research

We conducted a broad review of the literature examining the related topics of workplace aggression (Ambrose, Seabright, and Schminke, 2002; Andersson and Pearson, 1999; Beugre, 2005; Hershcovis and Barling, 2010; LeBlanc and Barling, 2005), entitlement (Harvey, 2009), and counterproductive work behavior (Beugre, 2005; Folger and Skarlicki, 2005; Fox and Spector, 2005; Kelloway, et al., 2010; Pearson, Andersson, and Porath, 2005; Spector and Fox, 2005; Tripp and Bies, 2009). This was expanded to also include computer deviance in the workplace (Mastrangelo, Everton and Jolton, 2006; Robinson and Bennett, 1995; Weatherbee, 2010), information security (Coles-Kemp and Theoharidou, 2010; Colwill, 2010), and criminal profiling (Gudaitis, 1998; Nykodym, Taylor and Vilela, 2006). Mastrangelo, Everton and Jolton (2006) report that 5-10% of employees engage in "antagonistic forms of deviant computer use" (p. 739) —socially undesirable behaviors such as gambling at work, downloading pornography, asking coworkers for dates, and violating confidentiality—and that the most common forms of deviant computer use involved personal email and chat sessions. The focus on counterproductive computer use in the workplace is of particular importance due to the risks of the insider threat (Bishop, Engle, Peisart, Whalen, and Gates, 2008; Pfleeger, Predd, Hunker and Bulford, 2010; Probst, Hunker, Gollman and Bishop, 2010; Schultz, 2002; Vasiu and Vasiu, 2004; Weiland, Moore, Cappelli, Trzeciak, and Spooner, 2010). This disparate literature has far-reaching application in a variety of fields of inquiry including information security studies, computer science, criminology, psychology, organizational behavior, and many more. The interdisciplinary and somewhat fragmented nature of the topic has lead to theoretical isolation and the lack of a unifying vocabulary to describe behaviors surrounding information technology misuse (Fox and Spector, 2005; Weatherbee, 2010).

Burroughs and James (2005) review personality research to identify individual differences that may account for counterproductive work behavior. People with certain dispositions are found more likely to engage in antisocial behaviors or to direct harmful actions against others. This may include exhibiting traditional workplace retaliation behaviors in order to right a wrong, in response to organizational upheaval or organizational injustice, and in response to breach of contract (e.g., psychological contract breach) (Ambrose, et al., 2002; Folger and Skarlicki, 2005; Pearson and Andersson, 2005; Rosen, Chang, Johnson and Levy, 2009; Tripp and Bies, 2009). Retaliatory behaviors may include but are not limited to calling in sick when not ill, gossiping about one's boss or coworkers, wasting company materials, damaging equipment or work processes. Shropshire (2009) recently conducted a canonical analysis of sixty-two intentional security breaches by insiders. His study indicated a positive correlation between four general variables and predictions of insider threat, each of which is observable by conscientious managers and/or supervisors. Financial changes correlated positively with information technology espionage while relationship strains, substance abuse, and job changes all positively correlated with information technology sabotage. This work is similar to studies conducted by Verizon and the U.S. Secret Service (2010), Cappelli, Moore, Trzeciak and Shimeall (2009), and the NIAC (2008) that assessed the relationship between insiders' backgrounds and motivations and their resulting deviant behaviors.

Despite a growing body of research into the psychology and motivation of insiders, it is difficult to predict who will commit security fraud (Kramer, Heuer and Crawford, 2005). Shaw and Fischer (2005) noted that most of the threats in their study could have been prevented by timely and effective action to address the anger, pain, anxiety, or psychological impairment of perpetrators, who exhibited signs of vulnerability or risk well in advance of the crime:

> *Combined with the strong finding above that nine of the 10 subjects were engaged in serious employment crises, this finding on the use of operations security to hide their system abuse reinforces the high value of personnel problems as a predictor of insider risk. This conclusion was further reinforced by findings on the occurrence in nearly every case studied of subject disgruntlement and serious personnel problems months prior to an attack. These subjects reacted to off-line personal conflicts, stresses, and disappointments through electronic behavior. The data from these subjects also indicated that the post termination window for an attack can range from hours to up to 2 months. One of the most important findings of this research was that there was a window of opportunity for dealing with the personnel problems affecting these subjects. These individuals were reportedly disgruntled in some cases for over a year prior to their attacks, and management was aware of these personnel problems weeks, if not months, prior to the attack. Yet there were consistent intervention problems. In fact, in many cases ill considered management actions exacerbated the problem. This finding indicates the need for improved management training and procedures covering interventions with at-risk individuals. [Shaw and Fischer, 2005, pp. 41-43]*

To date, no systematic methods have been used to evaluate psychosocial behaviors that can predict increased risk for insider threats. To fill this research gap, the present work follows recommendations by Schultz (2002) to develop a "new framework" for insider threat detection, which is based on multiple indicators that not only address workstation and network activity logs but also include preparatory behavior, verbal behavior and personality traits. Gudaitis (1998) makes a similar argument that "human based data gathering, assessment and profiling" must be synthesized and integrated with information security techniques to achieve an effective overall security package (p. 322). Nevertheless, Gudaitis argues effectively that traditional means of assessing psychological profiles and predispositions are problematic, in part because of existing laws, such as the Americans with Disabilities Act of 1992 and the Rehabilitation Act of 1973 Sections 503 and 504, against using clinical testing to indicate mental disabilities both during a pre-hiring phase for screening and during employment. Another issue undermining the usefulness of the traditional psychological profile assessment approach is due to the lack of valid data for "good" employees versus hackers or disgruntled employees. Instead, he recommends consideration of nonclinical instruments that measure personality and behavioral characteristics—tests that "not only focus on job suitability and skills, but they do not contain the obvious psychiatric questions that are easily picked out and answered 'appropriately'" (Guidatis, 1998, p. 327). Gudaitis advocates these types of tests as part of employee selection, but he allows that even this approach has drawbacks largely due to the unpredictability of the employees' life and work circumstances in relation to their work place behavior after being employed. To be specific, while a prospective employee may enter the workforce as a "good employee," he may turn into a disgruntled employee over time because of certain life circumstances and workplace experiences. For these and other reasons, Gudaitis concludes that a more integrated approach is required that involves assessment from multiple perspectives, and a common thread among each of these perspectives concerns the behaviors, motivations and expectations of employees and changes over time that may constitute threat indicators.

More recently, Phelps, Cappelli, Moore, Shaw, and Trzeciak (2007) reported wide support for relationships between dimensions of personality defined by the Five Factor Model (FFM) (Goldberg, 1993) and counterproductive work behaviors. The FFM describes personality factors of openness to experience, extraversion, conscientiousness, agreeableness, and neuroticism or emotional stability. Significant correlations have been reported between elements of the FFM (openness and agreeableness)

and irresponsible or counterproductive behaviors such as absenteeism, disciplinary issues, and drug or alcohol abuse. On a practical level, an efficient approach to insider threat mitigation that takes these factors into account without administering personality test instruments is to conduct audits aimed at identifying manifestations of these personality factors such as absenteeism, disgruntlement, disciplinary issues, and so forth. Along these lines, a Defense Personnel Security Research Center report advocates design of insider risk mitigation plans for individual employees. These plans take into account employees who display one or more concerning behaviors indicative of increased risk, such as an IT security violation or an altercation with his supervisor or coworkers (Shaw, Fischer and Rose, 2009). These behavioral manifestations may then justify the organization to conduct an insider risk evaluation.

Workplace disgruntlement and employee dissatisfaction are identified as two key underlying causes of deviance in the workplace and organizational crime (Willison, 2009; Moore, Cappelli and Trzeciak, 2008). When an individual has unsatisfied expectations of the organization, he or she might be motivated to address the expectations through malicious action against the organization (Moore, Cappelli, and Trzeciak, 2008). Unmet expectations might include organizational factors such as the level of compensation, promotional potential, and the organization's grievance and conflict resolution policy. Extra-organizational factors such as marital and familial problems, personal finances, and addictions can also influence the intensity of disgruntlement and dissatisfaction. The findings by Keeney, et al. (2005) reveal that 85% of the insiders identified in their study experienced grievances before carrying out attacks and in 92% of the sabotage cases the grievance was related to employment.

Employee attitudes—especially strong negative affect—are precursors of intentional counterproductive (and even subversive) behaviors ranging from absenteeism to various forms of retaliation (Workman, 2009b; Workman and Gathegi, 2007). Wells (2001) points out that employee dissatisfaction with the work organization is a powerful predictor of workplace fraud. Hollinger and Clark (1982) report a positive relationship between employee dissatisfaction and employee theft. Drawing from the organizational justice literature, Willison (2009) analyzed how distributive, procedural, interactional, interpersonal, and informational injustice within the organization can trigger insider computer crime and argued that there is a substantial relationship between employees' perception of injustice in the workplace and their deviant behavior such as theft, violence, and sabotage.

In the case of an employee trying to gain financially by exploiting a corporation's intellectual property, a desire for revenge may be driven by the satisfaction of causing costly damage to the corporation but it can be also include a motive of financial gain. In either case, the employee may have exhibited stress or some form of dissatisfaction about his or her circumstances. These factors, if properly evaluated in a timely manner, could alert an organization about a developing insider crime.

Identifying employees who show elevated risk of insider threat has two benefits: preventing an unnecessary cost to the employer, and helping the employee before a bad situation turns worse. A properly administered intervention will help find a solution that benefits both parties. In some cases, it can involve moving an employee to a more suitable position within the corporation. Thus, a psychosocial model benefits both the employees and the employers if the model is incorporated as a tool in regular staff evaluations, and if appropriate action is taken in accordance with the predictions of the tool.

3.0 Privacy and Ethical Issues

The privacy and ethics debate is clearly a contentious issue that deserves more discussion. There is a fine line between what the organization "needs to know" and what is firmly in the realm of the employee's expectation of privacy. Indeed, empirical evidence demonstrates that only a minute percentage of employees actually engage in activities that constitute insider threat; the rest of the population comprises honest, hard-working staff who would be highly offended to learn they were monitored. To gain a better understanding of the important legal and ethical context for our research, in the following section, we focus on how privacy and electronic monitoring are defined, regulated in the U.S. and abroad, and their implications for employment satisfaction, trust and insider threats.

Privacy. There has been a very long-running debate over privacy issues. Warren and Brandeis (1890) defined privacy as "the right to be let alone." Lasprogata, King, and Pillay (2004) refine this definition by describing historically protected areas of informational privacy (personal information), physical privacy, and decisional privacy (the right to be let alone regarding personal decisions). They note that American law on privacy stands apart from most of the world. In Europe, for example, the right to privacy is considered a fundamental right in that the Treaty of the European Union[1] requires member States to respect the fundamental rights as set forth in the Charter of Fundamental Rights of the European Union[2] that everyone "has the right to respect for his or her private and family life, home and communications." In 1995, the EU adopted the EU Privacy Directive that established national data protection laws administered by strong legal regimes to protect personal data privacy. In the EU, informational privacy is defined broadly to cover personal information processed in employment context, including electronic monitoring. Canada's Personal Information Protection and Electronic Documents Act (PIPEDA), which became widely effective in 2004, follow the EU Privacy Directive. However, in the United States, there is no comparable right to privacy of personal information (Lasprogata, King, and Pillay, 2004).

Electronic Monitoring. Lasprogata, King and Pillay (2004) describe three usages of the term "electronic monitoring": (a) it includes an employer's use of electronic devices to review and evaluate employee performance; (b) it includes surveillance of employee when not performing work tasks; and (c) it includes use of computer forensics by the employer to recover and reconstruct electronic data after an "exploit" such as deletion, concealment, or attempted destruction of the data. Vasterman, Yzermans, and Dirkwager show an increasing trend in these surveillance practices, growing from 67% in 1999 to 92% by 2003 in the United States (cited in Workman, 2009a). Surveillance, in the security literature, is defined as the physical or electronic observation of people's activities and behavior (Ball and Webster, 2003). According to the 2007 Electronic Monitoring & Surveillance Survey (American Management Association, 2008), 43% of companies monitor employee e-mail, 66% monitor Internet connections, and 45% of employers track content, keystrokes, and time spent at the keyboard. Over 58% of the managers surveyed had fired workers for email or Internet misuse. Of the employees dismissed, 64% violated company policy and 22% breached confidentiality rules. Each of the employer responses above has psychological and financial costs for the employer.

In the United States, electronic workplace monitoring is generally unrestricted except in circumstances relating to disability or health information. The Electronic Communications Privacy Act

[1] Treaty Establishing the European Community, Feb. 7, 1992, O.J. (C224) 1 (1992).
[2] Chapter of Fundamental Rights of the European Union, art. 7, Dec. 7, 2000, O.J. (C364) 1(2000).

(ECPA) of 1986, which is intended to provide individuals with some privacy protection in their electronic communications, provides only limited protection to private sector employees (General Accounting Office, 2002). Some U.S. laws provide privacy protection only for the contents of employees' electronic communications, while other U.S. laws protect the privacy of only personal data relating to medical/health information (Lasprogata, King, and Pillay, 2004). Under current judicial interpretation of federal and state privacy statutes, a U.S. employer can implement an electronic employee monitoring policy so long as it has obtained employee consent. There is no such thing as "covert" monitoring when the employee has agreed to be monitored at the discretion of the employer (Lasprogata, King and Pillay, 2004), and, therefore, there is no expectation of privacy by the employee. In contrast to the status of workplace monitoring in the United States, the EU requires that monitoring of employee electronic communications be subject to a "contract" with the employee or in compliance with a legal obligation of the employer or need to conduct legitimate business, so long as the monitoring does not violate the employee's fundamental rights. Canada has a similar human rights/ethical position in restricting such surveillance.

U.S. Federal and state laws and judicial decisions have generally given private sector companies wide discretion in their monitoring and review of employee computer transmissions, including the Internet and e-mail (General Accounting Office, 2002). Employer use of electronic monitoring of employee cyber activities has been growing and, in the United States, it is widely acknowledged that employers have the right to monitor employee cyber activities (Lasprogata, King and Pillay, 2004). The 9/11 attacks on the United States ushered in an expansion and legitimization of surveillance trends. The premise is that if threats can be identified and neutralized, stability and control will be assured. Many do not object to this. In fact, technologies such as mobile phones and debit cards, which are ubiquitous tools enabling flexibility and freedom in one's mobility, are accompanied by surveillance measures to ensure connectivity, correct billing, and other precautionary measures.

The relationship between workplace monitoring and trust is delicate and complex. On the one hand, opponents to workplace monitoring claim that an employer's use of monitoring devices, whether covert or overt, threatens both employee privacy and morale. Among the arguments cited against employee electronic monitoring are deleterious effects on employee morale, on the trust relationship between the employee and the employer, and on the work product itself (Ariss, 2002). Advocates of privacy rights seek to ensure that employees will not suffer unwanted intrusions and that potentially harmful information will not be acquired about them. Critics note that monitoring can increase employee stress, reduce commitment, and lower productivity (Brown, 1996; 2000). Monitoring perceived as invasive with an implied lack of trust may contribute to employee job dissatisfaction. Workman (2009b) observes that "When attitudes are negative about surveillance, employees are less likely to be committed to their work, they display lower organizational citizenship behaviors, and in some cases resort to furtive means of retaliation against management and the organization" (p. 348). In some cases management intervention on suspected employee disgruntlement issues may actually increase an employee's frustration level (Shaw and Fischer, 2005). On the other hand, lack of monitoring due to inflated or unjustified trust can also produce adverse effects. It has been observed that inadequate attention and action by an employer can increase insider activity. Such influences on trust have been described as the "trust trap" (e.g., Band, et al., 2006).

The ramifications of workplace monitoring and the use of such information in terms of employee job satisfaction and public relations can be severe. Trust is a fundamental concept underlying the issue of privacy and workplace monitoring. Tabak and Smith (2005) assert that the initiation of trust and

subsequent trust formation affects managerial implementation of electronic monitoring policies, and these policies have implications for workplace privacy rights. Similarly, employee perception of management practices influences employee trust in and commitment to the organization. Brown (2000) concludes that "workers are driven to semi-schizoid responses to the power and authority" of technologically mediated supervision (p. 65). He posits that this resulting loss of privacy can create feelings of vulnerability that may give rise to deep-seated feelings of alienation, and psychological patterns of behavior often observed in extreme adversity, including selective apathy, emotional disengagement, and narcissistic survivalism. However, from the employer's perspective, there are many good business reasons to electronically monitor employees in the workplace, including assessing worker productivity, protecting company assets from misappropriation, and ensuring compliance with workplace policies and nondiscrimination laws, as well as national security concerns (King, 2003). Monitoring promotes productivity and affords better control over counterproductive employees. The justification of such a practice is based on the argument that employers "own" or pay for employee time and resources such as computer equipment and network connections.

To be sure, employment is founded upon trust, and even though the organization typically asserts and society acknowledges its right to conduct electronic workplace monitoring, there is the potential for reduced trust. However, if the process is disclosed fully, explained, and managed equitably across employees, it may not be considered as unfair by employees, and the mutual trust relationship required for a healthy organization may remain intact. Workman (2009b) argues that such "procedural justice" is a crucial element in mitigating negative attitudes among employees and provides a set of guidelines to help management structure employee monitoring and surveillance to facilitate perceptions of procedural justice by employees. Thus, the data monitoring needed to inform a predictive psychosocial model should be done openly with proper privacy safeguards, and based on actual behaviors and events that are identified as part of the normal performance assessment process.

4.0 Psychosocial Model Description

In keeping with an approach that attempts to reflect relationships between certain personality characteristics and counterproductive work behavior (CWB) or higher-risk employees, we conducted discussions with HR professionals and managers at our organization to identify behavioral "proxies" for such characteristics that may, to varying degrees, produce a heightened concern about possible insider threat risks. Informed by the Five Factor Model (FFM) and previous research and case studies documenting personality disorders and factors of concern, these discussions focused on the kinds of behaviors that would likely be observed and "known" by managers and HR staff because of the level of concern that they bring about. The model that evolved from these discussions was therefore highly observation-based, i.e., focusing on observable behaviors that could be recorded and audited. Therefore, although the model is based on behavioral observables, it can support making inferences about the possible psychological/personality/social state of an employee; hence we refer to our model as a "psychosocial" model to capture the wide spectrum of inferences it is capable of producing. The implementation of the psychosocial reasoning used a data-driven approach based on personnel data that are likely to be available (see Greitzer, Frincke and Zabriskie, 2010 for a discussion). The indicators used in the model, such as disgruntlement, anger management issues, and disregard for authority, are listed and defined in Table 1. It is worth noting that these psychosocial indicators contribute differentially to the judged level of psychosocial risk with disgruntlement, difficulty accepting feedback, anger management issues, disengagement, and disregard for authority having higher weights than other indicators, for example.

Table 1. Psychosocial Indicators.

Indicator	Description
Disgruntlement	Employee observed to be dissatisfied in current position; chronic indications of discontent, such as strong negative feelings about being passed over for a promotion or being underpaid, undervalued; may have a poor fit with current job.
Not Accepting Feedback	The employee is observed to have a difficult time accepting criticism, tends to take criticism personally or becomes defensive when message is delivered. Employee has been observed being unwilling to acknowledge errors; or admitting to mistakes; may attempt to cover up errors through lying or deceit.
Anger Management Issues	The employee often allows anger to get pent up inside; employee has trouble managing lingering emotional feelings of anger or rage. Holds strong grudges.
Disengagement	The employee keeps to self, is detached, withdrawn and tends not to interact with individuals or groups; avoids meetings.
Disregard for Authority	The employee disregards rules, authority or policies. Employee feels above the rules or that they only apply to others.
Performance	The employee has received a corrective action (below expectation performance review, verbal warning, written reprimand, suspension, termination) based on poor performance.
Stress	The employee appears to be under physical, mental, or emotional strain or tension that he/she has difficulty handling.
Confrontational Behavior	Employee exhibits argumentative or aggressive behavior or is involved in bullying or intimidation.
Personal Issues	Employee has difficulty keeping personal issues separate from work, and these issues interfere with work.
Self-Centeredness	The employee disregards needs or wishes of others, concerned primarily with own interests and welfare.
Lack of Dependability	Employee is unable to keep commitments /promises; unworthy of trust.
Absenteeism	Employee has exhibited chronic unexplained absenteeism.

It should also be noted that the judgments based on observations will necessarily be subjective since there is no expectation that an objective test instrument will emerge from this research. Nevertheless, with appropriate training, we believe that management and HR personnel would better understand the nature of the threat and the likely precursors or threat indicators that may be usefully reported to cyber security officers.[3] Most importantly, the approach in predictive modeling is to provide "leads" for cyber security officers to pursue in advance of actual crimes, without which they would likely have little or no insight to select higher-risk "persons of interest" and focus analyses.

For security analysis purposes, only individuals about whom a manager is "highly concerned" would be considered for further analysis in the insider threat model. As the model's assignment of risk level increases for an individual, so too would be the level of monitoring and analysis of that individual.

4.1 Bayesian Network Development

The psychosocial indicators and the psychosocial risk were implemented as binary variable nodes in a Bayesian network model (Heckerman, 2008; Pearl, 1985) using GeNIe [GeNIe 2.0.3006.0, Decision Systems Laboratory, University of Pittsburgh] as shown in Figure 1.

The development of a Bayesian network requires several steps. First, the network is constructed with linked conditionally dependent random variables that each takes on values from a domain. In our model, these values are *True* or *False*, corresponding to whether the indicators, the behaviors, were observed severe enough to be a concern. Second, prior probabilities (*priors*) are assigned to each random variable. These priors, which were estimated by HR experts, reflect the frequencies at which random variables take on values from their domains. For example, the prior probability that an employee is observed to exhibit severe stress is denoted P_{Stress}, and the complementary case (employee does not exhibit severe stress) is $1 - P_{Stress}$.

It is interesting to note that when the priors were solicited from the HR experts, the experts were initially asked to provide the priors as probabilities. An examination of these priors and discussion with our experts suggested that these initial estimates were inflated. Recognizing that there may be certain biases associated with probability estimation, particularly for rare events with negative consequences or utilities (e.g., Harris, Corner and Hahn, 2009), we, therefore, asked the HR experts to estimate the *number* of cases that occur per year in which an employee exhibits a given indicator. We assumed a baseline context of about 4000 employees, which is consistent with their experiences at our institution. The rephrased question format appeared to have provided better estimates of the priors.

[3] Training on recognition of threat indicators, as well as on consistent reporting and effective mitigation strategies are essential for successful application of these concepts. An ongoing R&D program funded through the Office of the Secretary of Defense focuses on accelerated learning through serious game technology concerning behaviors and indicators of potential insider threat (see Andrews, 2010 and http://www.acq.osd.mil/osbp/sbir/solicitations/sbir083/osd083.htm).

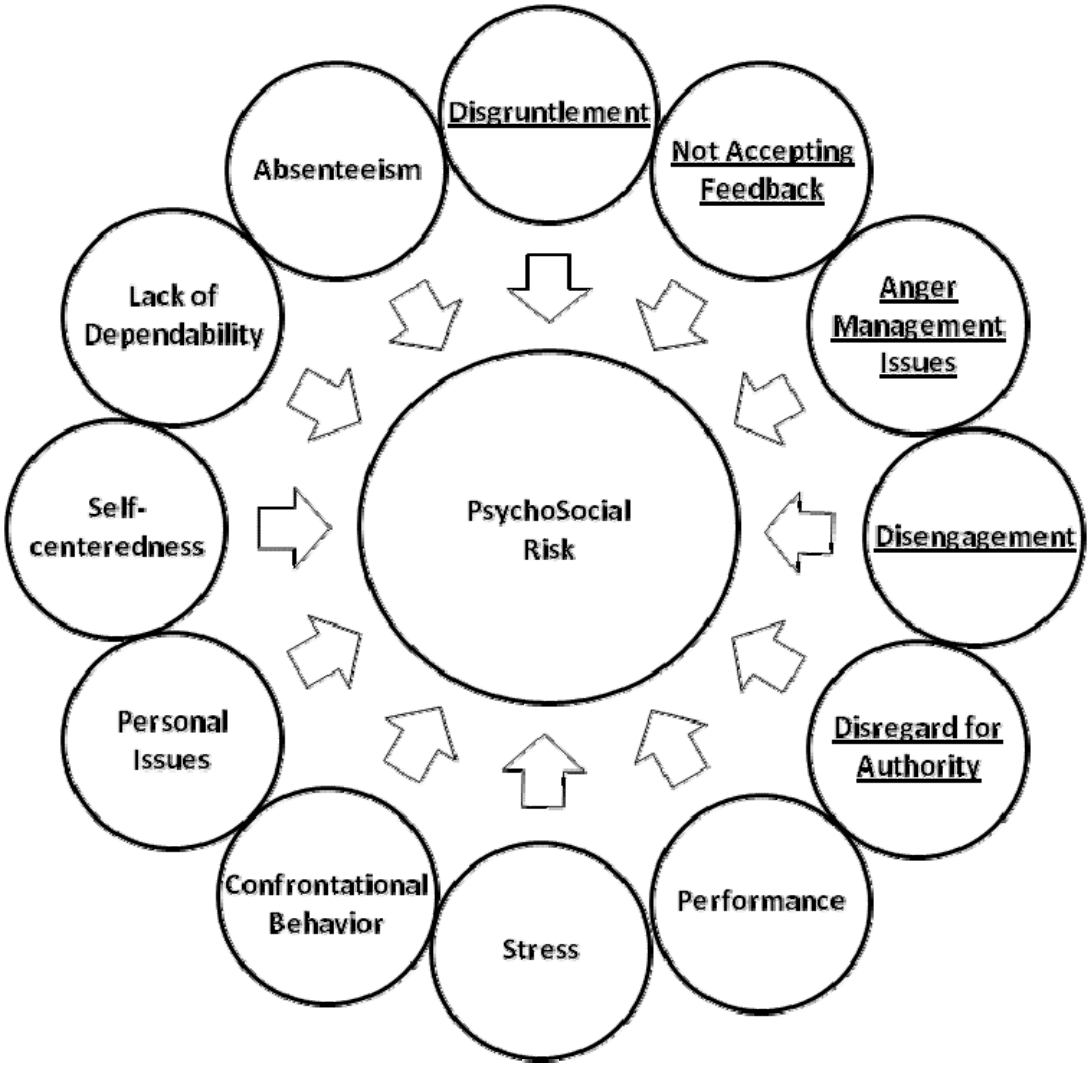

Figure 1. The Psychosocial Model with Behavioral Indicators. Indicators determine the relative "risk level" of an individual. The five underlined indicators are considered to be higher risks than the other indicators.

Table 2 shows the priors for observing the employee behaviors in a year as estimated from our HR experts. The table also shows relative judgments obtained from the HR experts of the weight of each indicator in influencing ones risk assessment for insider threat when the indicator is observed alone. We discuss below how we used a different method to assess the risk from combinations of indicators.

The table shows that (extreme) Disgruntled behavior occurs relatively seldom (0.025) but has a high influence on the associated insider risk (0.400), while (extreme) Self-centeredness occurred relatively often (0.100) but has a lower influence (0.180). Intuitively, one can conclude that extreme self-centeredness, when observed alone in 10% of employees, should not cause alarm for an insider risk; otherwise an employer would have to conduct comparatively high levels of insider threat monitoring on 10% of its workforce. (The priors differ for different labor forces and corporations—the numbers shown in Table 2 are estimated by the HR experts for research scientists.)

Table 2. Priors and weights for the model's random variables.

Parameter	Prior	Weight
Disgruntled	0.025	0.400
Accepting Criticism	0.060	0.280
Anger Management	0.019	0.260
Engagement	0.040	0.310
Disregards Rules	0.075	0.340
Performance	0.020	0.160
Stress	0.030	0.200
Confrontational	0.063	0.120
PersonalIssues	0.080	0.140
SelfCentered	0.100	0.180
Dependability	0.038	0.060
Absenteeism	0.010	0.060

The third step in developing a Bayesian network is to determine the influence of the random variables on the risk output to be encoded. One way to do this is to consult HR experts in order to enter numeric values directly in the conditional probability table (CPT) of the Bayesian network's Psychosocial Risk random variable node. Clearly, this involves a large number of complicated judgments in which various numbers of factors are combined.[4] Because this method was highly impractical, we used an alternative approach to derive the conditional probabilities through a training methodology that acquires expert judgments for only about 3% of the total number of cases and that builds on the judgments of individual priors shown in Table 2. We constructed 110 different scenarios where employees had zero to five indicators set to TRUE, and then asked the HR experts to assign insider threat risks to employees who would exhibit those behaviors. These scenarios with the experts' assignment of risks were then used in the expectation maximization algorithm in GeNIe to set the conditional probabilities in the CPT for the risk variable. In essence, this step thus transferred the expertise of our HR experts to the Bayesian network with the expectation that the network should predict the same insider risk in employees as the HR experts would for novel employee evaluations.

[4] Formally, the total number of possible combinations is referred to as the power set *P(S)*, the set of all possible subsets of the set *S*, which is in fact 2^S. In the present case, there are 4096 possible cases.

4.2 Verification Experiment

The Bayesian network model depicted in Figure 1 was developed from judgments of two HR experts in several knowledge engineering meetings. In a verification study, we asked these HR experts to judge the severity of 61 cases on a 0-10 scale and compared their averaged scores with the output of the model (the Bayes probabilities were normalized to a 0-10 risk scale). The results, shown in Figure 2, indicate a high level of agreement between the expert judgments and the psychosocial model (R^2 = 0.920).

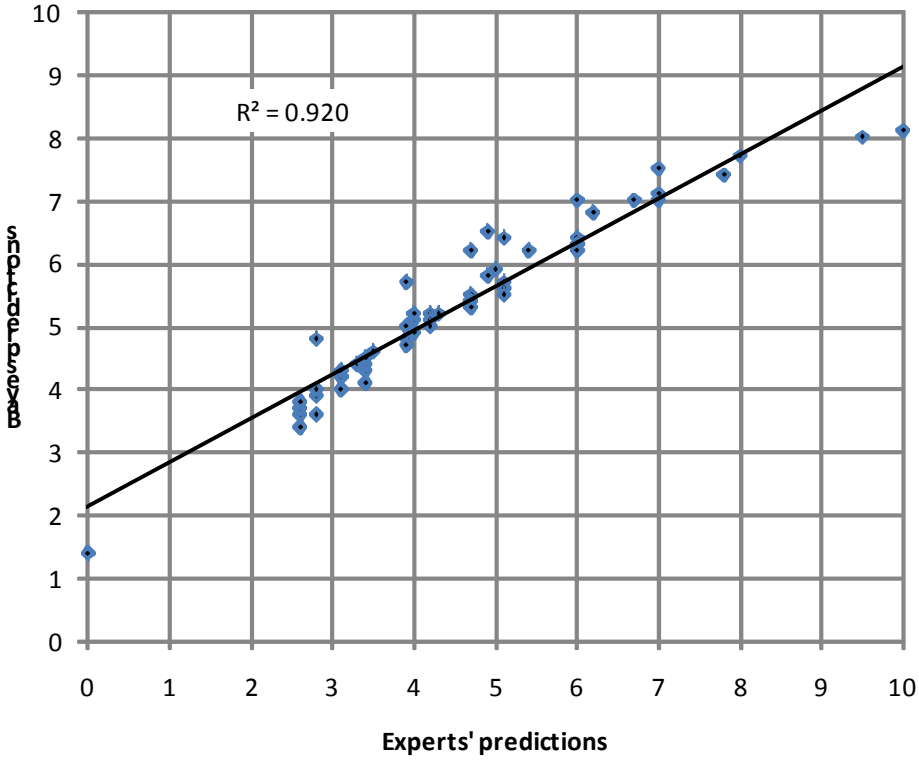

Figure 2. Verification Test of Psychosocial Model During Development. Shows the Fit of the Psychosocial Model to Expert Judgments used During Development of the Model.

5.0 Formal Study

The analysis described in the previous section provided a degree of verification that the psychosocial model represented the judgments of experts whose assessments were elicited to construct the model. A more rigorous assessment was obtained by examining the degree to which the model fits judgments of another set of experts.[5]

5.1 Participants and Procedure

Ten staff members from the Pacific Northwest National Laboratory were recommended by HR management to participate in the experiment due to their breadth and depth of experience. Nine of the participants were human resources professionals and one was a line manager. Informed consent was obtained from these volunteers.

Each participant attended one of three 2-hour sessions conducted over a day and a half, during normal business hours. Three or four participants attended each session (for a total of ten participants). Each session was conducted according to a scripted procedure. After welcoming the participants and collecting consent forms, the experimenter spent ten minutes describing the general insider threat problem and explained that the focus of this study was on behavioral factors, acknowledging that a complete response to the insider threat problem would also require integration of behavioral analyses with workstation/electronic monitoring approaches. We also explained that the focus of this study was on malicious insiders, not on individuals who inadvertently propagate malicious exploits by others from the outside, such as phishing attacks and the like. After this discussion, the participants were asked to read a one-page description of the problem that included a table identifying and describing twelve behavioral factors that were studied. The experimenter explained that these factors were of interest, and the opinions of the participants were sought to help validate the importance of each of the factors. Thus, it was noted that there was no expectation that these factors contribute equally to identifying and predicting potential insider threats. With no additional discussions or questions by the participants, the experimenter described the procedure that the participants were asked to follow.

Each participant was given a collection of 24 cases; each typed on a separate page with the case presented both in tabular and narrative form. Figure 3 shows an example of a case that exhibited three indicators, Disgruntlement, Anger Management Issues, and Disregard for Authority. All 24 cases are summarized in Table 3.

Each set of 24 cases was shuffled prior to the experiment session so there was no consistent order of the cases in the sets given to the participants. The participants were asked to sort the 24 cases into up to ten categories, ranging from “no concern” on the left to “highest concern” on the right. The sorting

[5] We recognize that the strongest form of validation would entail the ability of the model to predict actual insider exploits rather than to generate predictions that are consistent with expert judgments. A longitudinal study is required in which data are collected over a period of time and then predictions of the model are compared to actual observed events. This was beyond the scope of the present study.

categories were not labeled, and participants were instructed that they must use the first and last category, but they were not required to use every category in between.

Case 39567.

Indicator	Observed?	
Disgruntlement	Yes	
Not Accepting Feedback		No
Anger Management Issues	Yes	
Disengagement		No
Disregard for Authority	Yes	
Performance		No
Stress		No
Confrontational Behavior		No
Personal Issues		No
Self-Centeredness		No
Lack of Dependability		No
Absenteeism		No

Adam was sure he would be picked for a one-year offsite assignment that he wanted very much. Not receiving the assignment made Adam disappointed and angry at management. After Adam's manager observed that he continued to hold a grudge and exhibit the anger (**Disgruntlement, Anger Management Issues**), the group manager entered these observations into Adam's personnel folder. Following this incident, the group manager received word of Adam breaking company rules, possibly in defiance of its management (**Disregard for Authority**). No other risk indicators have been observed or recorded in his personnel folder. Other than these indicators, no additional issues have been observed or documented.

Figure 2. Example of case in tabular and narrative form presented to test participants.

Participants worked on their own and at their own pace. When the sorting task was finished, the experimenter asked the participants to further rank-order any cases that appeared in the same category. At the conclusion of this task, the participants were asked to rank-order the twelve factors, from highest to least importance, as lone indicators or predictors of potential insider threat risks.

Table 3. Twenty-four cases with subsets of indicators true (1) and false (0).

Case / Indicator	Disgruntlement	Not Accepting Feedback	Anger Management Issues	Disengagement	Disregard for Authority	Performance	Stress	Confrontational Behavior	Personal Issues	Self-Centeredness	Lack of Dependability	Absenteeism
1	1	0	0	0	0	1	0	0	0	1	0	0
2	1	0	0	0	1	0	0	0	0	0	0	0
3	1	0	0	1	0	0	0	0	0	0	0	0
4	1	1	0	0	0	0	0	0	0	0	0	0
5	1	0	1	0	0	0	0	0	0	0	0	0
6	0	1	0	1	0	0	0	0	0	0	0	0
7	0	0	0	1	0	0	1	0	0	0	0	0
8	0	0	0	1	1	0	0	0	0	0	0	0
9	0	0	0	1	0	0	0	0	1	1	0	0
10	0	0	0	0	1	1	0	0	0	1	0	0
11	0	0	0	0	1	1	0	1	0	0	0	0
12	1	1	1	0	0	0	0	1	0	1	0	0
13	1	0	0	0	1	0	0	0	0	0	0	0
14	0	0	0	0	0	0	0	0	0	1	1	0
15	0	0	0	0	0	1	0	0	0	0	0	1
16	0	0	0	0	0	0	1	0	1	0	0	0
17	1	0	1	0	1	0	0	0	0	0	0	0
18	0	0	0	0	0	0	0	1	0	0	1	0
19	0	0	0	0	0	1	1	0	0	0	0	0
20	0	0	1	1	0	0	0	0	0	0	0	0
21	0	0	0	0	0	0	0	0	0	0	0	1
22	0	0	0	0	0	0	0	0	0	1	0	0
23	0	0	0	0	1	0	0	0	0	0	0	0
24	1	0	0	0	0	0	0	0	0	0	0	0

5.2 Measures

For measurement/identification purposes, we labeled the ten sorting categories using the numbers 0-9. Each case was coded into its respective sorting category (0-9), with additional discrimination using the rankings within categories, as follows: The first case in category 0 (no concern) was assigned a score = 01; the second case in that category was given score = 02; and so on. For category 2, the cases were assigned scores of 21, 22 …, depending on the number of cases contained in that category. Similarly, scores were assigned for all 24 cases.

For the indicator ranking task, the twelve indicators were assigned numeric/integer ranks from 1 to 12, based on the rankings assigned by the participants.

5.3 Results

Table 4 shows the rank orders of the 12 indicators as given by the participants. The average standard deviation for the twelve indicators is 1.69, indicating a good consensus among the participants for the indicators. A few indicators stand out as having a large range among the participants: e.g., the ranks of Disengagement range from 1 to 12. As can be seen in Table 4, eight of the ten participants seemed to agree on a middling ranking (corresponding to the average rank of 5.2), while one participant considered the indicator to be the most important and another participant considered it the least important. Although the participants in the tests were given some guidelines for the definitions of the indicators, some subjective interpretation seemed to persist at the time of testing. Variations in the participants' rankings are expected based on their different experiences and perceptions.

Table 4 shows that Self-Centeredness has the highest standard deviation (3.44). This indicator also has the highest prior (0.100) in Table 2. This means that in a given year an estimated 10% of the staff will be observed to have extreme self-centeredness (i.e., severe enough to justify recording and concern). Although our participants for testing were selected for their many years of HR and/or managerial experience, some had worked in different environments from a research laboratory. Within the Laboratory, self-centeredness is believed to be frequent as staff increase in seniority. These phenomena could also be observed in other environments where the accumulation of knowledge and experience put staff at considerable advantage in performing high quality work.

Table 4. Indicator rank orders by ten experts.

Indicators\Expert	1	2	3	4	5	6	7	8	9	10	Average	StdDev
Disgruntlement	1	1	6	6	3	1	5	2	1	1	2.7	2.16
Accepting Feedback	6	6	8	5	5	6	6	8	6	5	6.1	1.10
Anger Management	5	3	2	3	4	2	2	4	4	3	3.2	1.03
Disengagement	4	4	5	7	1	12	7	5	3	4	5.2	2.97
Disregard for Authority	3	2	1	1	2	3	3	1	2	2	2.0	0.82
Performance	8	11	12	8	7	8	10	9	10	6	8.9	1.85
Stress	7	7	9	10	9	10	8	7	5	7	7.9	1.60
Confrontational	2	5	3	4	5	4	4	3	7	8	4.5	1.84
Personal Issues	10	10	10	11	8	9	9	12	11	9	9.9	1.20
Self-Centeredness	9	9	4	2	11	5	1	6	8	10	6.5	3.44
Lack of Dependability	11	8	7	9	10	7	12	10	9	11	9.4	1.71
Absenteeism	12	12	11	12	12	11	11	11	12	12	11.6	0.52
											Average:	**1.69**

Table 5 shows the insider threat risks assigned by the participants to our 24 test cases. The assigned risk levels range from 1 to 93, and were computed as explained above. The table shows that the average standard deviation is a low (14.62) for the 24 cases, and suggests that the ten participants were relatively consistent in judging the risk for the individual cases. The lowest standard deviation is 6.77 for Case 21 and the highest is 27.33 for Case 7. The table shows that for Case 7, two participants assigned risks of 1 and 2, respectively, and two other participants assigned this case risks of 73 and 74, respectively.

Table 5. Ten experts' scores given to cases.

Case\Expert	1	2	3	4	5	6	7	8	9	10	Average	StdDev
1	82	73	55	81	61	85	71	64	82	13	66.7	21.41
2	71	41	34	94	91	73	41	91	61	61	65.8	22.19
3	72	61	51	61	82	72	51	51	64	63	62.8	10.35
4	74	72	52	51	72	61	23	81	44	62	59.2	17.30
5	73	71	54	62	54	51	61	72	52	41	59.1	10.59
6	31	42	61	32	81	42	22	31	53	51	44.6	17.51
7	11	74	31	11	73	1	33	2	45	11	29.2	27.33
8	61	43	53	91	93	32	21	73	63	44	57.4	23.71
9	32	62	57	71	41	71	42	41	81	22	52.0	19.24
10	41	51	91	82	52	91	81	53	91	82	71.5	19.79
11	63	83	56	95	62	82	52	63	71	81	70.8	13.89
12	91	92	93	92	71	81	82	62	93	91	84.8	10.79
13	3	22	21	22	22	52	12	12	2	1	16.9	15.07
14	12	33	24	64	21	21	14	22	43	2	25.6	17.53
15	4	11	1	34	13	12	32	11	31	3	15.2	12.52
16	13	21	2	1	1	13	11	1	11	12	8.6	6.93
17	81	82	92	96	92	92	91	82	92	92	89.2	5.37
18	21	81	32	63	23	62	34	52	62	43	47.3	19.86
19	2	23	5	33	11	14	13	23	42	21	18.7	12.34
20	62	91	33	41	83	31	44	42	72	52	55.1	21.01
21	1	2	3	21	3	11	1	13	1	4	6.0	6.77
22	5	1	4	31	2	41	31	21	21	5	16.2	14.65
23	14	31	23	93	51	83	45	61	41	31	47.3	25.47
24	51	32	22	42	53	84	43	71	51	42	49.1	17.90
											Average:	**14.62**

The wide range among the participants in a few cases is assumed, as explained above with the indicator scores, to result from the individual participants having had different experiences. The observed indicators in Case 7 are Disengagement and Stress. Table 4 showed that two participants ranked Disengagement at the opposite extremes at how much it predicts insider threat risk, which may explain some of the wide range for the case.

For the test results in Table 5, the inter-rater agreement on the 24 scenarios was high (pair wise mean Pearson correlation coefficient = 0.684, standard deviation = 0.095; intra-class correlation coefficient = 0.651 with 1.0 being perfect agreement; inter-rater reliability coefficient with Spearman-Brown correction = 0.949; the nonparametric, Kendall's w, i.e., Kendall's coefficient of concordance, is 0.707 ($p < .001$) with 0 being no agreement and 1, perfect agreement). The coefficient of concordance suggests there is a high level of agreement among the raters and the agreement is statistically significant.

6.0 Test of Model

The Bayesian network was tested using a round robin procedure, leaving out the 24 cases from one rater for testing while the 24 cases from each of the other nine raters were used in GeNIe's expectation maximization algorithm to learn the probabilities in the conditional probability tables in the network. Figure 4 below shows the Bayesian network predictions for the 240 cases left out in testing. The scatter plot shows a clear vertical division between predictions ~0.3-0.4. Above this boundary there is at least one behavioral indicator observed from the group of the five most important indicators shown in Figure 1 with underlined labels. And below the boundary there are only behaviors that appear less indicative of insider threat risk. Although the experts' ratings of the cases do not show this separation explicitly, the Bayesian network learned to deduce this pattern from the experts' predictions [R^2 = 0.598, Root Mean Square Error (RMSE) = 0.188].

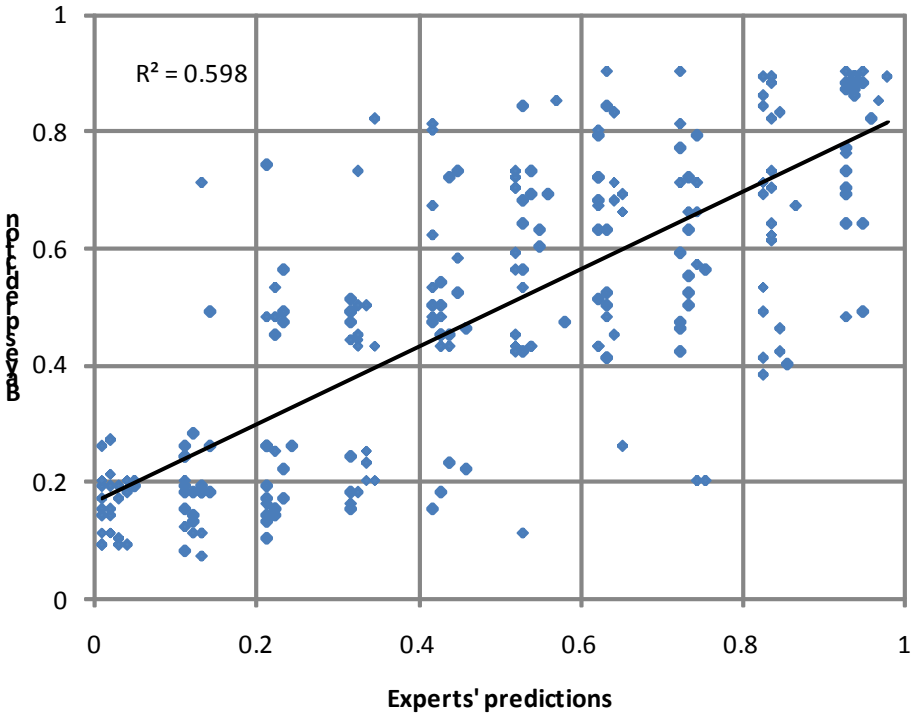

Figure 3. Bayesian Model's Predictions of Expert Judgments. Shows prediction of 24 unique cases for a total of 240 test cases from Round-Robin testing.

6.1 Comparing to Alternative Models

Besides the psychosocial model, three other models were developed and tested to predict the risk that a staff member would pose if subsets of the indicators had been observed as shown in Table 3. Like the Bayesian network (BN) model, a linear regression (LR) and a feedforward artificial neural network (ANN) were trained and tested using round robin test. A Counting model was simply required to count the number of indicators observed and it did not require estimates of weights for the individual indicators and thus no training or round robin was needed.

Figure 5 shows that the counting model yielded a Pearson R2 = 0.253 (RMSE = 0.260). The relatively poor performance is attributed to the fact that all indicators were weighted equally in this model while the experts clearly considered some indicators more important than others. This confirms the consensus in the ranking of indicators shown in Table 4—the experts considered the indicators as having different weights in predicting the risk.

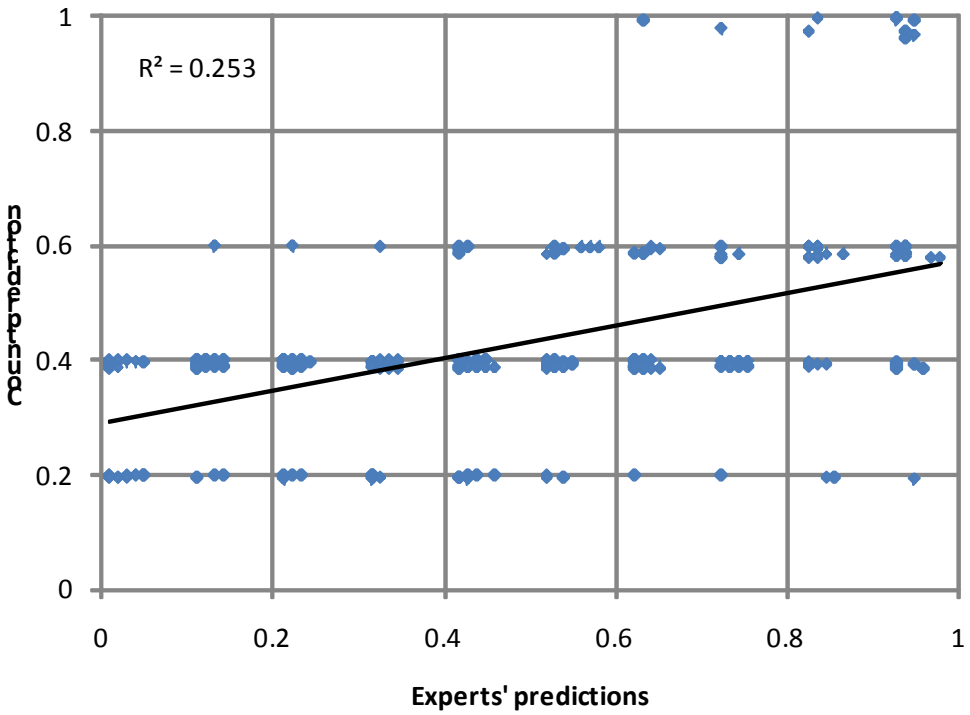

Figure 4. Counting Model's Prediction of Expert Judgments. The model counts each behavioral indicator unweighted for a prediction.

The LR model had a specific weight for each indicator. Two methods were tried when developing this model. First, a genetic algorithm (Goldberg, 1989; Holland, 1992) was configured to optimize the indicator weights to produce the risk measures assigned by the experts for the cases. Second, a feed-forward artificial neural network (ANN) was configured without a hidden layer (12 inputs; 1 output) and trained with the backpropagation algorithm (Rumelhart, Hinton and Williams, 1986; Werbos, 1974, 1994). The fully-trained ANN yields the weight for each indicator. As both methods optimize the weights, they produced the same weights. (Observe that round robin was used for both methods as explained above.) As shown in Figure 6, the performance of the LR model was R^2 = 0.592 (RMSE = 0.191), a substantial improvement over the counting model, further confirming that the individual indicators should be weighted differently.

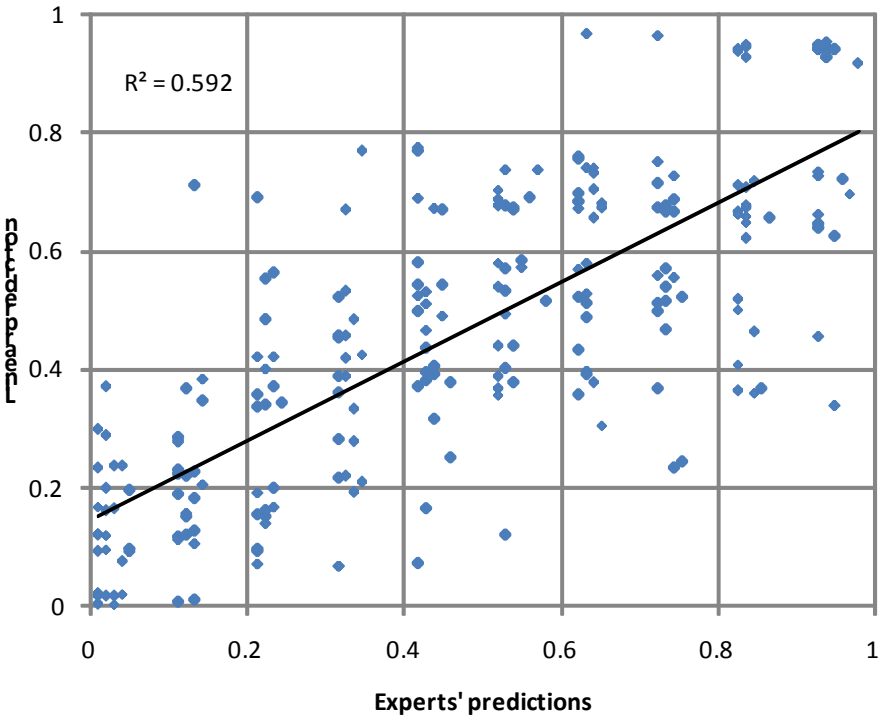

Figure 5. Linear Regression Model's Prediction of Expert Judgments. Each behavioral indicator weight was learned from the experts in round robin testing**.**

The ANN with one hidden layer (12 inputs, 2 hidden nodes, and 1 output) was tested to discern if the problem had nonlinear properties. The performance increased only slightly to an R^2 = 0.606 (RMSE = 0.185), shown in Figure 7, suggesting the problem is mostly nonlinear. As shown in Figure 3, the Bayesian network (BN) model, gave an R^2 value of 0.598 (RMSE = 0.188), a value only slightly below that of the (nonlinear) ANN. Additionally, variations of the number of hidden nodes were tested and our results show that more than two hidden nodes did not improve performance suggesting the problem has a low complexity (Observe that one hidden node is the same as a linear model.

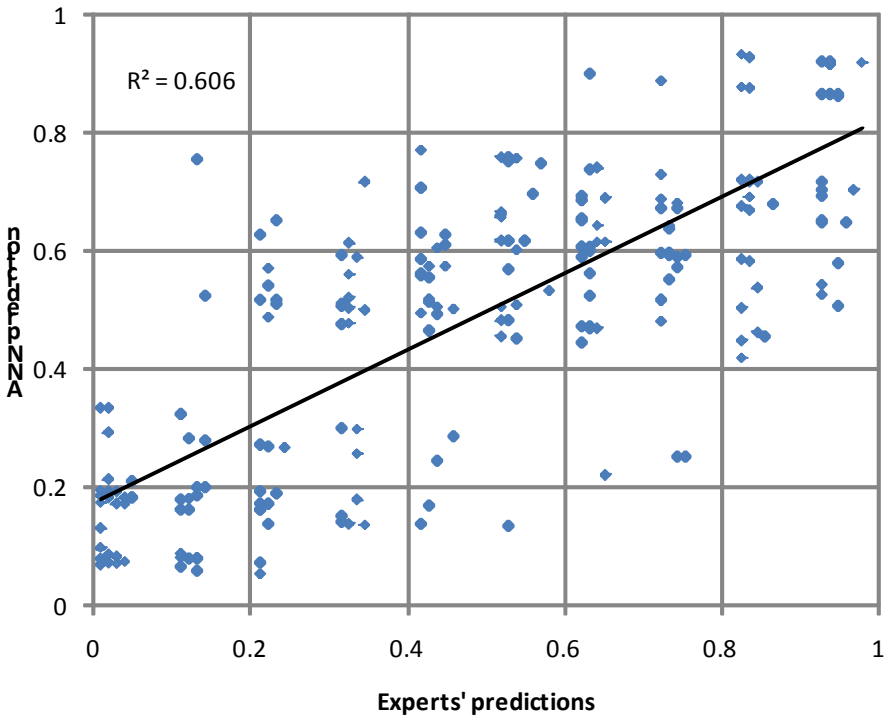

Figure 6. Nonlinear Artificial Neural Network Model's Prediction of Expert Judgments. Nonlinear regression model where each behavioral indicator weight was learned from the experts in round robin testing.

6.2 Discussion of Models

The ANN and the BN scatter plots both show significant gaps in how these models assigned risks. These gaps coincide with low and high risk indicators. When the indicators were developed, certain indicators were deemed more significant than others. The low risk indicators (e.g., lack of dependability and absenteeism) according to the experts in the development team were those that can be observed alone in staff members without them posing a risk. Only when these are observed together with high risk indicators (e.g., disregard for authority, disgruntlement, and anger management) do they increase the risk of committing an insider crime. Below the gaps in the graphs, there are only combinations of the low risk indicators. Above the gap, there is one or more high risk indicators combined with the low risk indicators. None of these relationships was explicitly discussed with the test participants; the models revealed this relationship without it being observed in the scores given by the participants. One possible interpretation for the gap is that it provides a natural threshold for deciding when staff members should be monitored more closely.

Table 6 summarizes the results obtained for the individual models. Clearly, the Bayesian network, the Linear Regression, and the ANN models that differentially weigh the individual behavioral indicators perform best. Both the Bayesian network and the ANN capture some nonlinearity in the data that the linear regression model fails to capture.

Table 6. Performance of the Models.

Model	R^2	**RMSE**
Bayesian Network	0.598	0.188
Counting Model	0.253	0.260
Linear Regression	0.592	0.191
Artificial Neural Network	0.606	0.185

Even though the performance of the Bayesian network is a little below that of the ANN model and only a little above the linear regression model, it has three important advantages. First, a regression model, either linear or nonlinear like the artificial neural network, typically requires that each indicator is set to true or false for binary values. The Bayesian network is better suited to work with these missing values. If behavioral indicators are neither observed/confirmed as true or false, then the Bayesian network will use the prior probabilities of those indicators to make the best predictions. Second, the Bayesian network also gives probability estimates, assuming true risk rates and priors were available during model development, an advantage that the other two models do not have. Finally, compared to an artificial neural network, the Bayesian network is typically more acceptable to users because it provides simpler explanations of why specific risks are assigned.[6]

The modeling results show that an R2 of about 0.6 is achievable in an expert system that simulates the consensus of ten experts. No judgment is made here as to the accuracy of the experts' opinions in predicting threat, but we observe that their consolidated judgments represent many years of experience in managing human resources. We suggest that the "average" risk predictions generated by a model representing these experts' consolidated wisdom is better than the prediction that an individual expert can provide due to possible information processing limitations, individual biases, or varying experiences. An expert system model also enables the automatic screening of staff members, which is consistent and independent of the experiences an individual human resources staff may have.

[6] Also, in their typical application, artificial neural networks only estimate probabilities.

7.0 Discussion

We have described research that suggests that any attempt to seriously address the insider threat, particularly through proactive means, must consider behavioral indicators in the workplace in addition to the more traditional workstation monitoring methods. The timeline for data breaches, as assessed in the Verizon and the U.S. Secret Service (2010) report covering breach cases in 2009, is instructive. While 31% of data breaches took on the order of minutes, a very large percentage (60%) took days to months, a phenomena also reported by Shaw and Fischer (2005). The Verizon report observed: "If victims truly have days or more before an attack causes serious harm, then this is actually pretty good news. It means defenders can take heart that they will likely get more than one chance at detection" (p. 47). The report goes on to say: "The bad news is that organizations tend not to take advantage of this second window of opportunity. The telltale signs are all too often missed and the attacker has all the time they need to locate and compromise data" (p. 47). Recognizing behavioral indicators is difficult and requires training, but we suggest that raising managers' and HR staff's awareness and skills in recognizing potential risks can only improve their effectiveness in dealing with everyday workplace challenges as well as severe insider threat risks. For a very large organization, it is difficult and costly to train sufficiently many experts so that all employees' risks are regularly and consistently analyzed. Transferring HR expertise to a computer model-based decision aid will help ensure that the "system" is applied consistently and fairly. The model automates this process as long as the organization creates employee evaluation processes that enable appropriate data to be recorded into a database of personnel files either at regular intervals such as during employee performance evaluations or when these behaviors are observed.

From the limited set of HR experts and managers that we have interviewed, we gained an understanding that these experts are not at all "clueless" with regard to insider threat detection. Good managers and HR staff are well aware of incidents and issues relating to "concerning behaviors" such as increasing complaints to supervisors regarding salary, increased cell phone use at the office, refusal to work with new supervisors, increased outbursts directed at coworkers, and isolation from coworkers (Randazzo, Keeney, Kowalski, Cappelli and Moore, 2004; Shaw and Fischer, 2005; Cole and Ring, 2006; Phelps, et al., 2007; Shaw, Fischer and Rose, 2009). For the most serious occurrences, which are the focus of our model, there will be communications and discussions among HR staff and management. Management is not only aware of the most egregious behaviors, but indicators of concerning behaviors may appear many months before an actual attack. Shaw and Fisher (2005) observe that signs of disgruntlement may appear from 1 to 48 months before the attack. This provides a window of opportunity during which employers' awareness of risk linked to effective interventions could reduce the threat of an attack (Band, et al., 2006; Shaw and Fischer, 2005). Randazzo, et al., (2004) reported that eighty percent of insider cases in their study raised official attention for concerning behaviors such as tardiness, truancy, arguments with coworkers, and poor job performance; in 97% of those cases, supervisors, coworkers, and subordinates were aware of these issues. However, typically there is no formal infrastructure for recording and tracking such behaviors, except when they become critical to the point where disciplinary action is taken. We are advocating the establishment of a system for collecting and tracking concerning behaviors so they may be taken into account by an objective system that integrates these psychosocial indicators with physical and cyber monitoring data to derive a more complete picture of potential "problem employees" and insider threats (Greitzer, et al., 2009).

It has been argued that insider threat assessment based on screening of personal characteristics may be inadequate because there may be considerable characteristic commonalities shared by both "good" and

"bad" actors and malicious insiders do not share a common profile, as Pfleeger, Predd, Hunker and Bulford (2010) point out, "Because the set of malicious insiders is small and diverse, no single personal characteristic or set of characteristics can act as a reliable predictor of future misbehavior" (p. 174). We are not advocating a model based only on personal characteristics, but rather a model that integrates multiple sources of data, which is consistent with the relevant literature. For instance, Gudaitis (1998) advanced the multidimensional profiling concept; Schultz (2002) advocated for systems that monitor and analyze numerous clues of diverse types including personal characteristics and suspicious cyber activities; Weatherbee (2010) developed the cyberdeviancy model that aims to link individual, situational, and organizational variables to events and processes, and Nykodym, Taylor and Vilela (2005) categorized computer crime committed by insiders as a step towards cyber criminal profiling. It is in the area of multidimensional, integrative modeling of the insider threat that our model contributes to the extant research on insider threat detection and interdiction.

We envision that a psychosocial model will be helpful to an HR expert with access to a database containing information and judgments associated with the psychosocial factors (judgments obtained from managers based on observed concerning behaviors). Because consistency and objectivity are of paramount importance in providing this type of input, managers who supply this information must be given guidelines and effective training on recognizing psychosocial indicators. A system developed to record these judgments would only take data for cases for which there are grave concerns about the employee's behavior. These judgments are combined with other behavioral data that may be available (such as disciplinary actions) and with cyber/workstation monitoring data to produce a composite picture that a security analyst can examine. Analysis of outputs from a psychosocial model and other more conventional workstation activity monitoring would be used in informed decisions of a multidisciplinary team comprising management, HR, security, cybersecurity personnel as well as counterintelligence officer for the most serious transgressions. It warrants noting, however, the automated decision aid should be used only to inform and advise decision makers rather than acting as the sole foundation for invoking unilateral sanctions or responses. The impact of the deployment of a psychosocial model-based decision aid would be to help managers and HR staffs identify employees who are at greater risks of slipping into a state that puts the organization or its employees at risk. Rather than being regarded as a Machiavellian/punitive system, the use of such tools can be considered to offer a fair and consistent application of behavioral monitoring, and with proper privacy safeguards in place, it benefits both employees and employers.

8.0 Conclusions

The insider threat poses a very hard detection problem (Band et al., 2006; Colwill, 2010; Probst, Hunker, Gollman and Bishop, 2010; Schultz, 2002; Weatherbee, 2010; Weiland, Moore, Cappelli, Trzeciak and Spooner, 2010) and an even harder prediction problem. It is therefore not surprising that the insider threat ranked second of the eight problems on the list of hard problems identified by the 2005 Information Security (INFOSEC) Research Council (INFOSEC 2005). The potential for harm and cost to both employers and employees increases as sensitive information becomes more accessible in cyber space. To protect all parties, systematic methods are needed to reduce the risk of deliberate attempts to harm organizational interests or individuals. Research shows that in the preponderance of cases the malicious intent of the perpetrator was "observable" prior to the actual exploit (Randazzo, Keeney, Kowalski, Cappelli and Moore, 2004; Shaw and Fischer, 2005; Cole and Ring, 2006; Phelps, Cappelli, Moore, Shaw and Trzeciak, 2007; Shaw, Fischer and Rose, 2009). Considering this research literature along with input from human resources staff, we developed a prototype psychosocial model that uses twelve behavioral indicators to predict the level of risks of insider threats. A set of twelve indicators was deemed important in predicting unstable individuals who may, under the right circumstances, decide to seek out avenues for activities that violate policies or break laws in anticipation of financial gain or revenge. This set of twelve indicators was also selected to be easy to monitor and record regularly. We envision the indicators to be recorded by employees' managers as the behaviors are observed throughout the year.

Assuming that these behavioral indicators are available for collection, we have shown that a predictive model of insider threat risk can be developed. A test of a prototype Bayesian network model was conducted with ten experienced HR staff. The model was compared to three other models, demonstrating that the proposed behavioral indicators have their own differential predictive weight and combinations of indicators do not automatically add to a risk. Thus, a good model needs to weigh each indicator separately and combine indicators intelligently.

The Bayesian network model is attractive in that it is based on Bayes probabilities that take into account experts' subjective belief and assessments, and, since it uses prior probabilities as defaults, it is robust against missing data (in the present context, missing data are indicators that have not been observed as either true or false). The performance of the Bayes network was close to the performance of a nonlinear feedforward artificial neural network known in theory to be able to learn any function to any degree of accuracy (Hornik, 1989).

In the validation test, round robin test results showed that using the twelve indicators and a good model, the insider threat risk among employees can be assessed to be highly correlated with expert HR judgments. We believe that if the developed model is incorporated to monitor employees with proper recording of the behavioral indicators, and combined with detection and classification of cyber data from employees' computer/network use, the integrated system will empower a HR/cyber/insider threat team with enhanced situation awareness to facilitate the detection and prevention of insider crimes.

9.0 Ongoing and Future Research

Research ongoing at our Laboratory is addressing several questions and challenges. One research question that was briefly discussed in the present paper concerns the extent to which HR or management personnel in an organization are informed about the kinds of behavioral indicators posited in our model. Informal discussions with other researchers and stakeholders produced mixed opinions. To help inform this question, we are planning to conduct a relatively informal survey to identify the level of awareness about concerning behaviors in the workplace. A substantial survey study is likely to better address this research question by providing statistics broken down by type of organization, size of organization, position within organization (HR, management), and so on.

To be sure, regardless of the present level of awareness by HR or managers about insider threat, there is a need to develop effective training and awareness programs for managers and HR staff. Ongoing research funded by the Air Force Research Laboratory (Andrews, 2010) seeks to develop approaches to accelerating the learning of management/supervisory personnel about behavioral precursors and indicators that suggest employee issues that relate to insider threat.

There is a need for more thorough validation of insider threat models and tools. Additional validation testing is needed before deploying the tool that we developed. We have verified that the model derived from HR experts in our study adequately represents the judgments of these experts as well as judgments of experts in the same organization who were not initially consulted in developing the model. A broader population of HR/management experts would certainly strengthen the reliability of the model. Further, a rigorous validation of the model is required. To conduct such a test, a longitudinal study is required. To that end, designated input data should be collected (but not used) over a period of time (e.g., several years), as well as recording of outcomes. After the data collection period, statistical analyses may be conducted to assess the strength of the relationship between the predictor variables (the model outputs) and the recorded outcomes (ground truth).

Another active area of research in combating the insider threat concerns the development of methods and models for analyzing cyber data streams for evidence of suspicious computer activities that may portend insider crimes. There is a vast research literature focusing on development of detection tools (see, e.g., Gabrielson, et al., 2008 for a review), but recent assessments of research and practice conducted in our organization identifies a continuing need for integrated solutions that transcend anomaly/signature detection methods (Greitzer and Frincke, 2010). The PsyberSleuth prototype developed at PNNL (Greitzer, et al., 2009) aims to integrate more traditional cyber security audit data monitoring with psychosocial data collection in an effort to develop a more effective system with potential to anticipate and prevent insider crimes. While an initial prototype utilized fuzzy probability and finite state automata models implemented within a general Bayesian network modeling framework, current research[3] and model development in our Laboratory seeks to re-frame the architecture of the model using an hierarchical network of knowledge-intense case based reasoners that conduct sophisticated pattern classification processing of input data to analyze and recognize invariant forms of stored patterns. This dynamic/adaptive approach focuses on the semantic content (not just syntax) of monitored data. The new

[3] This internal Laboratory-Directed R&D project, called Adaptive Cyberdefense using an Auto-associative Memory Paradigm (ACAMP), is being carried out under the Information and Infrastructure Integrity Initiative. See www.i4.pnl.gov.

architecture will continue to take input from a psychosocial component of the system that will support integration of behavioral/psychosocial indicators with traditional cyber data. It is hoped that this more advanced architecture will provide greater adaptability while also improving scalability of the system.

More generally, a set of research topics suggested in the literature and by the present study includes: (a) research to further characterize data and indicators to support behavioral profiles that help to differentiate between accidental misuse and true malicious insider activities; (b) research to establish a greater understanding of the motivations of malicious insiders and related precursors or indicators of imminent insider attacks; (c) research on testing/evaluation methods to verify and validate proposed insider threat detection or prediction methods; (d) collection of useful data sets to support evaluation studies; and (e) a continuing responsibility to examine social issues, privacy and ethical issues, and legal considerations surrounding the deployment of insider threat monitoring, analysis, and mitigation systems that balance these public/human rights issues with the needs and responsibilities of organizations, enterprises, and governments in conducting business productively, safely, and securely.

10.0 References

Ambrose, ML, MA Seabright, and M Schminke. 2002, "Sabotage in the Workplace: The Role of Organizational Injustice." Organizational Behavior and Human Decision Processes 89:947-65.

American Management Association. 2008, "The Latest on Workplace Monitoring and Surveillance."
Andersson, LM, and CM Pearson. 1999, "Tit for Tat? The Spiraling Effect of Incivility in the Workplace." Academy of Management. The Academy of Management Review 24:452-71.

Andrews, D. 2010, "Cyber Training Research in the Air Force Research Laboratory, Panel Session." in Annual Computer Security Applications Conference (ACSAC), Honolulu, Hawaii.

Ariss, SS. 2002, "Computer Monitoring: Benefits and Pitfalls Facing Management." Information & Management 39:553-58.

Ball, K, and F Webster, eds. 2003. The Intensification of Surveillance: Crime, Terrorism and Warfare in the Information Age. Pluto Press, London.

Band, DR, et al. 2006, Comparing Insider It Sabotage and Espionage: A Model-Based Analysis. Technical Rpt. CMU/SEI-2006-TR-026, Carnegie-Mellon University. Software Engineering Institute. CERT Coordination Center.

Beugre, CD. 2005, "Reacting Aggressively to Injustice at Work: A Cognitive Stage Model." Journal of Business Psychology 20:291-301.

Bishop, M, et al. 2008, "We Have Met the Enemy and He Is Us." in Proceedings of the 2008 workshop on New security paradigms, pp. 1-12. ACM, Lake Tahoe, California, USA.

Brown, WS. 2000, "Ontological Security, Existential Anxiety and Workplace Privacy." Journal of Business Ethics 23:61-65.

---. 1996, "Technology, Workplace Privacy, and Personhood." Journal of Business Ethics 15:1237-48.

Burroughs, SM, and LR James. 2005, "Advancing the Assessment of Dispositional Aggressiveness through Conditional Reasoning." in Counterproductive Work Behavior: Investigations of Actors and Targets, eds. S Fox and PE Spector, pp. 127-50. American Psychological Association, Washington, DC.

Cappelli, D, et al. 2009, Common Sense Guide to Prevention and Detection of Insider Threats 3rd Edition – Version 3.1. Technical, Carnegie-Mellon University. Software Engineering Institute. CERT Coordination Center.

Coles-Kemp, L, and M Theoharidou. 2010, "Insider Threat and Information Security Management." in Insider Threats in Cyber Security, eds. CW Probst, et al., Vol 49, pp. 45-71. Springer US.

Colwill, C. 2010, in press, "Human Factors in Information Security: The Insider Threat - Who Can You Trust These Days?" Information Security Technical Report.

CSO Magazine, et al. 2010, 2010 Cybersecurity Watch Survey - Survey Results. Technical.

Folger, R, and DP Skarlicki. 2005, "Beyond Counterproductive Work Behavior: Moral Emotions and Deontic Retaliation Versus Reconciliation." in Counterproductive Work Behavior: Investigations of Actors and Targets, eds. S Fox and PE Spector, pp. 83-105. American Psychological Association, Washington, DC.

Fox, S, and PE Spector, eds. 2005. Counterproductive Work Behavior: Investigations of Actors and Targets. American Psychological Association, Washington, DC.

Gabrielson, B, et al. 2008, State-of-the-Art Report (Soar): The Insider Threat of Information Systems (Unclassified, for Official Use Only). Technical, Information Assurance Technology Analysis Center (IATAC), Herndon, VA.

General Accounting Office. 2002, Report to the Ranking Minority Member, Subcommittee on 21st Century Competitiveness, Committee on Education and the Workforce, House of Representatives: Employee Privacy—Computer-Use Monitoring Practices and Policies of Selected Companies. Technical Rpt. GAO-02-717, U.S. General Accounting Office, Washington, D.C.

Goldberg, DE. 1989. Genetic Algorithms in Search, Optimization and Machine Learning. Addison-Wesley, Reading, MS.

Goldberg, LR. 1993, "The Structure of Phenotypic Personality Traits." American Psychologist 48:26-34.

Greitzer, FL, and B Endicott-Popovsky. 2008, "Security and Privacy in an Expanding Cyber World, Panel Session." in Annual Computer Security Applications Conference (ACSAC), Anaheim, CA.

Greitzer, FL, and DA Frincke. 2010, "Combining Traditional Cyber Security Audit Data with Psychosocial Data: Towards Predictive Modeling for Insider Threat." in Insider Threats in Cyber Security, eds. CW Probst, et al., Vol 49, pp. 85-114. Springer US.

Greitzer, FL, DA Frincke, and MM Zabriskie. 2010, "Social/Ethical Issues in Predictive Insider Threat Monitoring." in Information Assurance and Security Ethics in Complex Systems: Interdisciplinary Perspectives, ed. MJ Dark, ed. IGI Global.

Greitzer, FL, et al. 2009, Predictive Modeling for Insider Threat Mitigation. Technical Rpt. PNNL-SA-65024, Pacific Northwest National Laboratory, Richland, WA.

Gudaitis, TM. 1998, "The Missing Link in Information Security: Three Dimensional Profiling." CyberPsychology & Behavior 1:321-40.

Harris, AJ, AL Corner, and U Hahn. 2009, "Estimating the Probability of Negative Events." Cognition 110:51-64.

Harvey, P. 2009, "Understanding and Managing Workplace Entitlement." in 8th Industrial & Organisational Psychology Conference. Australian Psychological Society.

Heckerman, D. 2008, "A Tutorial on Learning with Bayesian Networks." in Innovations in Bayesian Networks, eds. D Holmes and L Jain, Vol 156, pp. 33-82. Springer Berlin / Heidelberg.

Hershcovis, MS, and J Barling. 2010, "Towards a Multi-Foci Approach to Workplace Aggression: A Meta-Analytic Review of Outcomes from Different Perpetrators." Journal of Organizational Behavior 31:24-44.

Holland, J. 1992, "Genetic Algorithms." Scientific American 267:66-72.

Hollinger, R, and J Clark. 1982, "Formal and Informal Social Controls of Employee Deviance." Sociological Quarterly 23:333-43.

Hornik, K, M Stinchcombe, and H White. 1989, "Multilayer Feedforward Networks Are Universal Approximators." Neural Networks 2:359-66.

INFOSEC. 2005, National Scale Infosec Research Hard Problems List. Technical, INFOSEC Research Council.

Keeney, M, et al. 2005, Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors. Technical, U.S. Secret Service and Carnegie-Mellon University, Software Engineering Institute, CERT Coordination Center.

Kelloway, EK, et al. 2010, "Counterproductive Work Behavior as Protest." Human Resource Management Review 20:18-25.

King, NJ. 2003, "Electronic Monitoring to Promote National Security Impacts Workplace Privacy." Employee Responsibilities and Rights Journal 15:127-47.

Kramer, LA, RJ Heuer, Jr., and KS Crawford. 2005, Technological, Social, and Economic Trends That Are Increasing U.S. Vulnerability to Insider Espionage. Technical Rpt. TR 05-10, Defense Personnel Security Research Center, Monterey, CA.

Lasprogata, G, NJ King, and S Pillay. 2004, "Regulation of Electronic Employee Monitoring: Identifying Fundamental Principles of Employee Privacy through a Comparative Study of Data Privacy Legislation in the European Union, United States, and Canada." Stanford Technology Law Review 4:1-46.

LeBlanc, MM, and J Barling. 2005, "Understanding the Many Faces of Workplace Violence." in Counterproductive Work Behavior: Investigations of Actors and Targets, eds. S Fox and PE Spector, pp. 41-63. American Psychological Association, Washington, DC.

Mastrangelo, PM, W Everton, and JA Jolton. 2006, "Personal Use of Work Computers: Distraction Versus Destruction." CyberPsychology & Behavior 9:730-41.

Moore, AP, DM Cappelli, and RF Trzeciak. 2008, "The “Big Picture” of Insider It Sabotage across U.S. Critical Infrastructures." in Insider Attack and Cyber Security, eds. SJ Stolfo, et al., Vol 39, pp. 17-52. Springer US.

NIAC. 2008, Final Report and Recommendations: The Insider Threat to National Infrastructures. Technical, U.S. Department of Homeland Security, National Infrastructure Advisory Council.

Nykodym, N, R Taylor, and J Vilela. 2005, "Criminal Profiling and Insider Cyber Crime." Computer Law & Security Report 21:408-14.

Pearl, J. 1985, "Bayesian Networks: A Model of Self-Activated Memory for Evidential Reasoning (UCLA Technical Report CSD-850017)." in 7th Conference of the Cognitive Science Society, University of California, pp. 329-34, Irvine, CA.

Pearson, CM, LM Andersson, and CL Porath. 2005, "Workplace Incivility." in Counterproductive Work Behavior: Investigations of Actors and Targets, eds. S Fox and PE Spector, pp. 177-200. American Psychological Association, Washington, DC.

Pfleeger, SL, et al. 2010, "Insiders Behaving Badly: Addressing Bad Actors and Their Actions." IEEE Transactions on Information Forensics and Security 5:169-79.

Probst, CW, et al. 2010, "Aspects of Insider Threats." in Insider Threats in Cyber Security, eds. CW Probst, et al., Vol 49, pp. 1-15. Springer US.

Randazzo, MR, et al. 2004, Insider Threat Study: Illicit Cyber Activity in the Banking and Financial Sector. Technical Rpt. CMU/SEI-2004-TR-021, Carnegie-Mellon University. Software Engineering Institute.

Robinson, SL, and RJ Bennett. 1995, "A Typology of Deviant Workplace Behaviors: A Multidimensional Scaling Study." Academy of Management Journal 38:555-72.

Rosen, CC, et al. 2009, "Perceptions of the Organizational Context and Psychological Contract Breach: Assessing Competing Perspectives." Organizational Behavior and Human Decision Processes 108:202-17.

Rumelhart, DE, GE Hinton, and RJ Williams. 1986, "Learning Internal Representations by Error Propagation." in Parallel Distributed Processing: Explorations in the Microstructures of Cognition. 1: Foundations, eds. DE Rumelhart and JL McClellands, pp. 318-62. MIT Press, Cambridge, MA.

Schultz, EE. 2002, "A Framework for Understanding and Predicting Insider Attacks." Computers & Security 21:526-31.

Shaw, ED, and LF Fischer. 2005, Ten Tales of Betrayal: The Threat to Corporate Infrastructures by Information Technology Insiders. Report 1 - Overview and General Observations. Technical Rpt. TR 05-04, Defense Personnel Security Research Center, Monterey, CA.

Shaw, ED, LF Fischer, and AE Rose. 2009, Insider Risk Evaluation and Audit. Technical Rpt. TR 09-02, Defense Personnel Security Research Center, Monterey, CA.

Shropshire, J. 2009, "A Canonical Analysis of Intentional Information Security Breaches by Insiders." Information Management and Computer Security 17:296-310.

Spector, PE, and S Fox. 2005, "The Stressor-Emotion Model of Counterproductive Work Behavior." in Counterproductive Work Behavior: Investigations of Actors and Targets, eds. S Fox and PE Spector, pp. 151-74. American Psychological Association, Washington, DC.

Tabak, F, and WP Smith. 2005, "Privacy and Electronic Monitoring in the Workplace: A Model of Managerial Cognition and Relational Trust Development." Employee Responsibilities and Rights Journal 17:173-89.

Tripp, TM, and RJ Bies. 2009. Getting Even: The Truth About Workplace Revenge and How to Stop It. Jossey-Bass, San Francisco, CA.

Vasiu, L, and I Vasiu. 2004, "Dissecting Computer Fraud: From Definitional Issues to a Taxonomy." in 37th Annual Hawaii International Conference on System Sciences, pp. 8 pp.

Verizon and U.S. Secret Service. 2010, "2010 Data Breach Investigations Report."

Warren, SD, and LD Brandeis. 1890, "The Right to Privacy." Harvard Law Review 4:193-220.

Weatherbee, TG. 2010, "Counterproductive Use of Technology at Work: Information & Communications Technologies and Cyberdeviancy." Human Resource Management Review 20:35-44.

Weiland, RM, et al. 2010, Spotlight On: Insider Threat from Trusted Business Partners. Technical, Software Engineering Institute, Carnegie Mellon University.

Wells, JT. 2001, "Enemies Within." Journal of Accountancy 192:31-35.

Werbos, PJ. 1974, "Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences." in Department of Applied Mathematics, Vol PhD. Harvard University, Cambridge, MA.

---. 1994. The Roots of Backpropagation. John Wiley & Sons, Inc., New York.

Willison, R. 2009, Motivations for Employee Computer Crime: Understanding and Addressing Workplace Disgruntlement through the Application of Organisational Justice. Technical Rpt. Working Paper No. 1, Copenhagen Business School, Department of Informatics, Copenhagen, Denmark.

Workman, M. 2009, "A Field Study of Corporate Employee Monitoring: Attitudes, Absenteeism, and the Moderating Influences of Procedural Justice Perceptions." Information and Organization 19:218-32.

---. 2009, "How Perceptions of Justice Affect Security Attitudes: Suggestions for Practitioners and Researchers." Information Management and Computer Security 17:341-53.

Workman, M, and J Gathegi. 2007, "Punishment and Ethics Deterrents: A Study of Insider Security Contravention." Journal of the American Society for Information Science and Technology 58:212-22.

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by **Battelle** *Since 1965*

U.S. DEPARTMENT OF
ENERGY

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)
www.pnl.gov