

Running Head: Behavioral Information Security

Behavioral Information Security: Defining the Criterion Space

Jeffrey M. Stanton

Cavinda Caldera

Ashley Isaac

Kathryn R. Stam

Slawomir J. Marcinkowski

Syracuse University

Authors' note: Jeffrey Stanton, School of Information Studies, Syracuse University, Syracuse, NY 13244-4100 (jmstanto@syr.edu). This research was supported in part by a small grant from the SIOP research foundation and by award SES0196415 from the National Science Foundation. Neither SIOP nor the National Science Foundation necessarily endorse the results or conclusions of this study.

Please cite freely:

Stanton, J. M., Caldera, C., Isaac, A., Stam, K. R., & Marcinkowski, S. J. (2003, April).

Behavioral information security: Defining the criterion space. In P. M. Mastrangelo & W. J. Everton, The Internet at work or not: Preventing computer deviance. Symposium presentation at the 2003 meeting of the Society for Industrial and Organizational Psychology, Orlando, FL.

Behavioral Information Security: Defining the Criterion Space

Over recent decades most work organizations have come to depend on information technology for internal operations such as record-keeping, external transactions such as financial transfers, and mediated communications of all types (e.g., email). As connectivity among devices has increased, so has the likelihood of intrusion, theft, defacement, and other forms of information resource loss. Surprisingly, although organizations tend to be more concerned about vulnerability to external attack than internal, recent industry research by Ernst and Young (2002) suggests that more than three-quarters of security breaches result from activity within organizations. At the low end, losses from security breaches have been estimated at approximately \$20 billion per year across all U.S. organizations (Security Wire Digest, 2000). These losses have spurred increased spending on information security specialists and technology: According to a 2002 industry survey by Information Security Magazine, very large organizations spend an average of \$6 million per year apiece on information security. Smaller organizations spend on average nearly 20% of their overall information technology budgets on security related products. This ocean of cash has spawned a large new sub-industry dedicated to the design, development, and marketing of security-related devices such as firewalls, biometrics, and security scanners.

Product development in this new sub-industry has received ample intellectual backing from a rich array of academic research programs on cryptography, public key infrastructure, watermarking, access control, intrusion detection, and related topics. The CiteSeer automated indexing facility (<http://citeseer.nj.nec.com>) lists more than 10,000 academic science and engineering articles related to information security. Computer scientists, network engineers, information technology specialists and others have worked diligently over past decades to

develop technological solutions for fundamental information security problems: how to restrict information resource access to authorized individuals, how to transmit and receive information privately, how to keep public information accurate and available in the face of malicious intrusion, and so forth (e.g., Won, 2001). The list of accomplishments in these areas is long, and many of these developments have resulted in positive business, economic, and societal outcomes (Dhillon, 2001).

A constraint appears in all these efforts, however, in the form of the behaviors of the human agents who access, use, administer, and maintain information resources. The success of information security appears to depend in part upon the effective behavior of the individuals involved in its use. Appropriate and constructive behavior by end users, system administrators, and others can enhance the effectiveness of information security while inappropriate and destructive behaviors can substantially inhibit its effectiveness. Human behavior is complex and multi-faceted, and this complexity defies the expectations for control and predictability that developers routinely assume for the technology with which they work. As the Organisation for Economic Co-Operation and Development's (OECD, 2002) Guidelines for the Security of Information Systems states, "The diversity of system users—employees, consultants, customers, competitors or the general public—and their various levels of awareness, training and interest compound the potential difficulties of providing security." Even this statement belies a certain system-centric view of security: Security is something that is provided (ostensibly by technology) rather than something that is enacted by users and system administrators.

The present research takes a different perspective on information security by focusing on "behavioral information security" which is defined as the complexes of human action that influence the availability, confidentiality, and integrity of information systems. Because research

in this area is so new, in the present in study we focused on delineating and understanding the behavioral domain. Our goal for this study was to construct and test a taxonomy of information security behaviors. We expect that this knowledge can support later research efforts that focus on understanding the antecedents and consequences of information security behavior.

Behavioral Information Security: An Overview

Most research on information security focuses on algorithms, methods, and standards that support the three basic functions of information security: confidentiality, integrity, and availability. In addition to this basic research in computer science and mathematics, human factors experts have worked to simplify and rationalize the user interfaces of security-related systems. Likewise, management experts have analyzed business risks associated with information systems and have drafted organizational policies to cope with these risks. We believe that an important missing layer in this assortment of approaches lies between the human-computer interface and the business-level concerns of management. In particular, we believe that information security research presently gives too little attention to the antecedents of behavior in organizations.

As an example, despite the ready availability of encrypted email products as well as ubiquitous organizational policies decreeing the importance of secure communications, few individuals and few organizations use such products. Each research camp might offer a plausible explanation: Technologists might lament the lack of a widely accepted industry standard, human factors scientists might criticize the user interfaces for securing email as too complex and counterintuitive, and management scholars might say that the risk of costly disaster has historically been too low to bother enforcing the relevant security policies.

The behavioral information security perspective would offer a different scenario:

Workers find the use of encrypted email very inconvenient, particularly in light of the fact that they are under serious pressure to get a lot of work accomplished without delays. Additionally, the workers see little information of value in their routine correspondence and in those rare cases when there is a sensitive message to pass, a phone call or face-to-face meeting will suffice commendably. Finally, the worker sees that even top management never uses the secure email function, never mentions it as a high priority to the organizational mission, never offers training on its use, and never rewards those few workers who use the feature diligently. In *Secrets and Lies*, Bruce Schneier (2002) says, “Mathematics is logical; people are erratic, capricious, and barely comprehensible.” On the contrary, this example suggests that behavior is understandable, organized, and laden with meaning both for those who enact it and those who work at making sense of it.

Researchers have long used this foundational assumption as the basis for developing theory and practice for understanding and influencing behavior in organizations; a few have even begun to take tentative steps toward applying research in organizational behavior to information security problems. For example, Straub (1990) investigated the impact of sanctions and other forms of obtaining compliance in organizations to ascertain the extent to which the severity and certainty of sanctions would influence “computer abuse.” This early effort preceded a whole new line of research on counterproductive computer usage that has included projects by Loch & Conger (1996); Young (1998); Armstrong, Phillips and Saling (2000); Stanton (2002a); Morahan-Martin and Schumacher (2001); and others.

Interestingly, these projects and related work on the “insider threat” to information security (e.g., Anderson et al. 1999; Schultz, 2002; Shaw, Post, & Ruby, 2002) have all tended to

focus on the intentionally disruptive behavior enacted by a small proportion of the workers in any given organization. The few exceptions to this spotlight on troublesome actions have included some examinations of the importance of user awareness and training (e.g., Spurling, 1995; Thompson & von Solms, 1997, 1998), and analyses of the ethical guides that may influence security related behavior (e.g., Siponen, 2001; Trompeters & Eloff, 2001). We believe that projects like these hold substantial promise for helping to shift research away from the common assumption that workers are wrongdoers whose behavior must be carefully circumscribed. In contrast to that common assumption, our own research, as described below, explored both positive and negative behaviors.

Method

We conducted 110 interviews with managers, information technology professionals, and regular employees during which we asked respondents to describe both beneficial and detrimental behaviors that employees within organizations enact that affect information security. From the transcripts of these interviews we compiled a raw list of security related behaviors. Next, we prepared a card deck listing the 82 resultant behaviors. Ten subject matter experts (graduate students and faculty in information technology) sorted the cards into self-generated categories. By collapsing across the many similarities among these expert-generated categories, we developed a six-element taxonomy of security behavior that varied along two dimensions: intentionality and technical expertise. The intentionality dimension appeared to capture whether the behavior described on the card was intentionally malicious, intentionally beneficial, or perhaps somewhere in between (i.e., absent explicit intention to help or harm). The technical expertise dimension focused on the degree of computer or information technology knowledge and skill that the actor needed to have in order to perform the behavior described on the card.

Table 1 and Figure 1 below depict the six categories arranged on these two dimensions. To illustrate with contrasting categories, “aware assurance” is the positive practice of information security carried out by well-trained personnel, while “detrimental vexation” refers to the inappropriate and intentional behaviors of non-expert individuals who misuse information resources. An example of aware assurance would be when a well-trained junior systems administrator vigilantly updates her systems with the latest security patches. An example of detrimental vexation would be when a regular worker uses the company’s email to send sales pitches for his sideline business. In Figure 1, each category shows a slight degree of overlap with its neighbors in recognition of the existence of behaviors near the borderlines. For instance, forging an email header to make it seem like the boss has distributed a rude joke requires some expertise and a bit of malicious intent. Thus, this behavior may lie at the intersection of detrimental vexation and intentional destruction.

The central “dark cloud” of unintentional (in)security suggests that sometimes individuals act without explicit intentions either to harm or help information security behavior, even though the outcome may suggest otherwise. Our list contained many examples of “naïve mistakes” (e.g., using one’s social security number as a password) that suggested a lack of awareness of basic information security principles rather than an intention to cause harm. Similarly, dangerous tinkering suggests that an individual with a higher degree of technical expertise might affect information security as an unintended consequence of his or her ability to setup a complex technology configuration with unintentional properties (e.g., by deploying a wireless network gateway that allowed non-company personnel to use the company’s network). We tested this tentative six-category taxonomy to ascertain whether a larger panel of subject matter experts could agree on the classification of behaviors into the categories we had defined, to obtain

ratings of the level of expertise required, and to ascertain the apparent intent associated with each behavior.

Respondents. Respondents for this second phase of the study were 49 students in information technology bachelor's and master's degree programs. Because of the preliminary nature of this study, we did not capture demographic information about individual respondents, but as a group all respondents were between 20 and 35 years of age, were either undergraduate, Masters, or Ph.D. students in an information related field (e.g., information management and technology majors), and comprised both males and females from multiple ethnic backgrounds.

Stimulus Materials. For the second phase of the study, we developed a nine-page paper and pencil instrument that included instructions and rating anchors on the first page (which the respondents could tear off and use as a reference) and eight pages listing out each of the behaviors. Based on feedback we had received in the first phase of the study, a small set of cards had their respective descriptions divided into two discrete behaviors. Also based on feedback from the initial study, we reworded each item so that it consistently referred to the actor by the impersonal pronoun he or she. We presented these items in alternation to avoid bias based on the gender of the portrayed actor. The result was a series of 94 statements such as, "He brought a wireless gateway device into his office, and installed it on the network without authorization."

Procedure. For each of the 94 behaviors listed on the paper and pencil instrument, we asked respondents to assign the behavior to one of the six categories appearing in Figure 1 and described in Table 1. We also asked the respondent to rate the apparent intentionality of the behavior on a scale ranging from 1 ("Highly malicious intentions to compromise resources") to 5 ("Highly benevolent intentions to preserve resources"), and the apparent degree of necessary technical know-how on a scale ranging from 1 ("No special expertise or training required") to 5

(“A lot of special expertise and/or training required.”). Note that the former scale was bipolar while the latter was unipolar. Because of the large number of behaviors on the list, we asked respondents to work at their own pace and spread out the work over several sessions if necessary. We provided an addressed, stamped envelope so that respondents could return the surveys to us when complete. We distributed a total of 75 surveys and received 49 completed surveys for a response rate of 65 percent.

Results

Analytical Strategy. Because our primary interest lay in the list of behaviors rather than in individuals’ responses, we began by aggregating category assignments and ratings across all 49 respondents. To aggregate the categorical variable that designated the six categories of behavior we took the modal response. In the case of three behaviors, there were ties (i.e., bimodality): In these cases, we chose one of the two modal categories at random. To aggregate the two rating variables, we calculated both the mean and the standard deviation across all 49 responses. The mean provided a single index of each behavior’s standing on the intentionality scale and the expertise scale. The standard deviation provided an index of disagreement among raters about the status of intentionality and the degree of required expertise for each behavior.

Descriptive Results. Tables 2, 3, and 4 provide a descriptive overview of our item-by-item findings. Tables 2 and 3 focus on the means of the items, whereas Table 4 focuses on the standard deviations of the items. Notably, Table 2, which shows the ten most extreme behaviors with respect to expertise, and Table 3 which shows the ten most extreme behaviors with respect to intentionality, are suggestive of the success of our two-dimensional taxonomy. Behaviors listed in these two tables and apparently fit our definitions of high and low expertise as well as our definitions for beneficial and detrimental intentionality. Possibly more interesting are the

contents of Table 4. This table provides a view of the items that generated the most disagreements, with the top five items focusing on disagreements on the expertise ratings and the bottom five items focusing on disagreements on the intentionality ratings. Note that the smallest observed standard deviation for an item was .47, so the values in this table (up to three times as large) suggest substantial disagreements on the items highlighted here. With respect to expertise, it is interesting to note that the three out of five items pertain to training. It seems possible that in the case of these items, raters were unsure whether they should rate the actor's expertise prior to the training or following it. For intentionality, it is interesting to note that the stated behaviors may in some cases be legitimate for incumbents in certain roles in the organization (e.g., security specialists) while being inappropriate for other roles (e.g., regular employee), and the absence of information about the actor's position in the organization may have made judging intentionality much more difficult in these cases. Another note of interest about the final five behaviors in table 4 is that they are all actions with serious import: I.e., their ultimate results or implications could have substantial negative effects on the actor's organization.

Inferential Results. The results described above provided interesting overview of our respondents' ratings of expertise and intentionality, but our primary goal was to ascertain whether our six-category taxonomy provided a reasonably good fit to the data. To accomplish this, we used the modal category designators to code each behavior's membership in an intentionality factor and an expertise factor. The first two columns of table 1 document the six cells that created by crossing the three levels of intentionality with two levels of expertise. As an example, 39 out of 49 raters assigned item 11 ("She did not change her password for over two years") to the "Naïve Mistakes" category, and we therefore placed this item in the low expertise condition (expertise factor) and the neutral intentionality condition (intentionality factor).

Tests of Means. Next we ran a two factor MANOVA, using individual items as cases and mean scores on expertise and intentionality as the dependent variables. Figure 2 shows the results for expertise and Figure 3 shows the results for intentionality. The multivariate omnibus tests were statistically significant for the expertise categorization factor, $F(2,87)=92.5$, $p<.001$, the intentionality categorization factor, $F(4,174)=98.8$, $p<.001$, and the interaction of the two factors, $F(4,174)=4.89$, $p=.001$. The alert reader will recognize that the F-values for the main effects signify quite substantial effect sizes. Indeed, the eta-squared value for the main effect of expertise categorization was .68 and the eta-squared value for the main effect of intentionality was .69: Both values can be interpreted as large effect sizes (Cohen, 1992). These large effect sizes support the success of the categorization scheme by demonstrating that the marginal means for the two conditions of expertise appeared at substantially different points in our five-point continuum (see Figure 2) and that the three conditions of intentionality did likewise (see Figure 3). Further, the ordering of the means was, in both cases, as one would expect, with the means for low expertise and malicious intent near the bottom of their respective scales and the means for high expertise and benevolent intent near the top of their respective scales.

As the multivariate results suggest, the univariate effect of the expertise categorization on the mean item expertise score was statistically significant, $F(1,88)=183.1$, $p<.001$. Likewise, the univariate effect of the intentionality categorization on the mean item intentionality score was also statistically significant, $F(2,88)=337.0$, $p<.001$. Of greater interest, however, the univariate effect of the intentionality factor on expertise ratings was also statistically significant, $F(2,88)=125$, $p<.001$, as was the interaction effect on expertise ratings, $F(2,88)=3.39$, $p<.05$. An examination of the pattern of means in Figure 2 reveals that raters perceived the difference between levels of expertise in the neutral intent category to be larger than the differences in the

other two intentionality categories. In particular, behaviors in the naïve mistakes category received the lowest expertise ratings of any of the six categories. Additionally, post-hoc least significant difference tests showed that behaviors in the malicious category generally received higher expertise ratings than the other two intentionality categories. These results hint at the possibility of an interesting attributional bias: that these raters saw “bad actors” as intrinsically more expert than actors whose intentions were not apparently malicious.

In addition to this interesting finding for expertise, there was an interaction effect on intentionality ratings as well, $F(2,88)=6.54$, $p<.01$. In this case, inspection of the pattern of means in Figure 3 suggests that minimal differences appeared in ratings based on level of expertise in either the neutral intention or the benevolent intention condition. In the malicious intentions condition, however, the intentionality of those behaviors that had been classified as signifying high expertise were rated as more malicious than those behaviors that had been classified as signifying low expertise. As above, these results suggest that raters believed in a contingency between expertise and malevolence: Other things being equal, a detrimental behavior requiring more expertise was also seen as more malevolent.

Tests of Standard Deviations. As a final analytical strategy we ran a two factor MANOVA, using individual items as cases and aggregated standard deviations on expertise and intentionality as the dependent variables. Because the purpose of our rating task was to establish normative levels of expertise and intentionality on these behaviors, aggregated standard deviations can be interpreted as an index of disagreement among raters about the apparent expertise needed to enact the behavior and the apparent intentionality of the behavior. Figure 4 shows the results for expertise and Figure 5 shows the results for intentionality. The multivariate omnibus tests were statistically significant for the expertise categorization factor, $F(2,87)=6.57$,

$p < .01$, the intentionality categorization factor, $F(4,174)=12.4$, $p < .001$, and the interaction of the two factors, $F(4,174)=3.83$, $p = .01$. The eta-squared value for the main effect of expertise categorization was .13 and the eta-squared value for the main effect of intentionality was .22: Both values can be interpreted as small effect sizes (Cohen, 1992). Despite the small magnitude of these effect sizes these results demonstrating that the level of disagreement depended both upon the classification of expertise (see Figure 4) and the classification of intentionality (see Figure 5).

First, examining the results for disagreement about expertise ratings, the univariate effect of the expertise classification was not statistically significant, $F(1,88)=1.11$, N.S. in contrast, the univariate effect of the intentionality classification was statistically significant, $F(2,88)=11.7$, $p < .001$. Inspection of the pattern of means in Figure 4 shows that there was little disagreement about expertise ratings for behaviors that were classified as having neutral intent. For behaviors that were classified as having either beneficial or malicious intent, however, post hoc tests showed significantly more disagreement on expertise ratings. In other words, it seemed to be more difficult for raters to agree on the level of expertise required for a particular behavior if that behavior was apparently intentional rather than neutral in intent.

Next, examining the results for disagreement about intentionality ratings, the univariate effect of the expertise classification was statistically significant, $F(1,88)=12.6$, $p = .001$, as were the effect of intentionality classification, $F(2,88)=13.3$, $p < .001$, and the interaction of the two factors, $F(2, 88)=6.7$, $p < .01$. Inspection of the pattern of means in figure 5, suggests that it was relatively easy to agree about the intentionality of the behavior for behaviors that were classified as positive (and regardless of the classification of expertise). In other words, virtually everyone was able to agree that a beneficial behavior was positive in intent. In contrast, there were high

levels of disagreement about of the intentionality of behaviors that seemed to require high expertise, in both the neutral category and the malicious category. This was particularly true for behaviors that had been classified as malicious: Raters had great difficulty agreeing on the level of intentionality and this difficulty was especially pronounced for behaviors that had been classified as malicious and requiring high expertise. Interestingly, this result dovetails with the pattern of interactions ascertained for mean ratings. While the mean ratings suggested that a substantial subset of raters considered behaviors requiring more expertise as also signifying more malicious intent, the pattern of disagreement suggests that this very same category was also the one that was most difficult to rate.

Discussion

In the primary purpose of this study was to transform a raw list of security related behaviors into a more manageable taxonomy with recognizable dimensions that had logical and definitional appeal. Our results suggest that we achieved a degree of success in this goal. While three of the 94 behaviors did not generate a clear majority vote, for the remaining 91 behaviors a clear consensus emerged on where that behavior belonged in our six-element taxonomy. Further, when we used this consensus as a basis for comparing ratings of expertise and intentionality (the two dimensions of our taxonomy) statistical analysis of mean ratings clearly showed that as a group our raters assigned normative levels of expertise and intentionality consistent with the classification scheme. We feel confident at this point that virtually any security related behavior that occurs in organizations can be positioned at some point within our six category taxonomy.

From a practical perspective, this confirmation of our taxonomy simplifies the task of developing criterion measures that assess security related behavior in organizations. We recommend that such criterion measures attempt to capture, through observations, self ratings, or

other ratings, the occurrence of behaviors that require high and low expertise and the occurrence of behaviors whose apparent intentionality ranges from malevolent to benevolent. We also believe that one important characteristic of this taxonomy is that it provides clear indications of paths that organizations can take towards improving their security status. In general, any interventions that shift intentionality towards the benevolent end of the continuum ought to improve the organization's security status. Likewise, with the exception of those employees who may have malevolent intentions towards the organization, providing training and other forms of expertise development appear to have the potential to benefit the organization's information security. Of course, demonstrating the validity of these assertions requires further research.

From a scientific perspective, these results provide several points of interest. First, a pattern of ratings and disagreements appears to imply the possibility of an attributional bias. From a theoretical perspective, it would be interesting to know whether employees judge other people's intentions differently based on whether they see the behavior as requiring a lot of expertise or a little expertise. This question ties in with the notion of trust as a fundamental component of security in organizations. Knowing who to trust and under what circumstances to trust can make an important difference in an employee's ability to detect and respond to potential security problems. If an employee misjudges the trustworthiness of another actor based on irrelevant cues, this may 1) make it more difficult for certain employees to conduct a legitimate activities, and/or; 2) make it easier for individuals with malevolent intentions to carry out detrimental tasks.

Note that there are several weaknesses in the methodology of the present study that should temper interpretation of the results. First, the kinds of inferences that we wanted to make about the quality of our taxonomy required a representative sample of security related behaviors

in organizations. Although we used over a hundred interviews with individuals who occupied different organizational roles, this strategy in itself is not a guarantee that we obtained such a sample. In particular, to our knowledge the baseline distributions of positive and negative security behaviors and low expertise and high expertise security behaviors is unknown. While it is likely that behaviors in each of the six categories of our taxonomy occur sometimes in some organizations, it is difficult to know whether we have adequately covered the behavioral domain in each of our six categories or indeed whether additional categories exist for which none of our respondents described a representative behavior. In short, some security related behaviors may exist that do not fit our categorization scheme.

Additionally, the results that we obtained suggesting the possibility of attribution bias may have been an artifact of our sample of raters. Using student raters, who generally lack lengthy work experience, may have resulted in some anomalies in our ratings. Future research should attempt to replicate our results using various samples of incumbent employees.

References

- Anderson, R. H., Feldman, P. M., Gerwehr, S., Houghton, B., Mesic, R., Pinder, J. D.,
Rothenberg, J., & Chiesa, J. (1999). Securing the U.S. defense information infrastructure:
A proposed approach. Washington, DC: Rand.
- Armstrong, L., Phillips, J. G., & Saling, L. L. (2000). Potential determinants of heavier Internet
usage. International Journal of Human-Computer Studies, *53* (4), 537-550.
- Cohen, J. (1992). A power primer. Psychological Bulletin, *112* (1), 155-159.
- Dhillon, G. (Ed.) (2001). Information security management: Global challenges in the new
millennium. Hershey, PA: Idea Group Publishing.
- Ernst and Young LLP. (2002) Global Information Security Survey. Published in the UK by
Presentation Services.
- Loch, K. D., & Conger, S. (1996). Evaluating ethical decision-making and computer use.
Communications of the ACM, *39*(7), 74-83.
- Morahan-Martin, J., & Schumacher, P. (2000). Incidence and correlates of pathological Internet
use among college students. Computers in Human Behavior, *16* (1), 13-29.
- OECD (2002). OECD Guidelines for the Security of Information Systems and Networks:
Towards a Culture of Security. Organisation For Economic Co-Operation And
Development. Available at: <http://www.oecd.org/pdf/M00033000/M00033182.pdf>.
- Schultz, E. E. (2002). A framework for understanding and predicting insider attacks. Computers
and Security, *21* (6), 526-531.
- Schneier, B. (1995). Applied Cryptography: Protocols, Algorithms, and Source Code in C (2nd
edition). New York: Wiley.
- Schneier, B. (2000). Secrets and Lies. New York: Wiley.

Security Wire Digest. (2000, March 27) CSI/FBI study says: Security breaches on the rise.

Author. Available at: http://www.lexias.com/1.0/securitywiredigest_27MAR2000.html

Shaw, E. D., Post, J. M., & Ruby, K. G. (2002). Inside the Mind of the Insider. Available at:

<http://www.securitymanagement.com/library/000762.html>

Siponen, M. T. (2001). On the role of human morality in information systems security.

Information Resources Management Journal, 14 (4), 15-23.

Spurling, P. (1995). Promoting security awareness and commitment. Information Management &

Computer Security, 3 (2), 20-26.

Stanton, J. M. (2002a). Company profile of the frequent Internet user: Web addict or happy

employee? Communications of the Association for Computing Machinery, 45 (1), 55-59.

Straub, D.W. (1990). Effective IS security: an empirical study. Information System Research, 1

(2), 255-77.

Thomson, M.E. and von Solms, R. (1997). An effective information security awareness program

for industry. Proceedings of the WG 11.2 and WG 11.1 of the TC11 IFIP.

Trompeters, C. M., & Eloff, J. H. P. (2001) A Framework for the Implementation of Socio-

ethical Controls in Information Security. Computers & Security, 20, 384-391.

Won, D. (Ed.) (2001). Proceedings of the Third International Conference on Information security

and cryptology (ICISC 2000), Seoul, Korea, December 8-9, 2000. Berlin: Springer.

Table 1:
Two Factor Taxonomy of Security Behaviors

Expertise	Intentions	Title	Description
High	Malicious	Intentional Destruction	Use this category when a behavior requires quite a lot of technical expertise together with a strong intention to do harm to the organization's information technology and resources. One example would be when an employee breaks into an employer's protected files in order to steal a trade secret.
Low	Malicious	Detrimental Vexation	Use this category when a behavior requires minimal technical expertise but nonetheless includes some intention to do harm through annoyance, harassment, rule breaking, and other misbehaviors. One example would be using company email facilities as a method of sending out SPAM messages marketing one's hobby business.
High	Neutral	Dangerous Tinkering	Use this category when a behavior requires quite a lot of technical expertise but no clear intention to do harm to the organization's information technology and resources. One example would be when a technically savvy employee sets up a wireless gateway that inadvertently allows wireless access to the company's network by people in passing cars.
Low	Neutral	Naïve Mistakes	Use this category when a behavior requires minimal technical expertise and no clear intention to do harm to the organization's information technology and resources. One example would be when an unaware and untrained employee chooses a bad password such as "password" that could be easily guessed by someone who was trying to break into the organization's systems.
High	Beneficial	Aware Assurance	Use this category when a behavior requires quite a lot of technical expertise together with a strong intention to do good by preserving and protecting the organization's information technology and resources. One example would be when a well-trained and motivated system administrator diligently and regularly applies security patches to all relevant software.
Low	Beneficial	Basic Hygiene	Use this category when a behavior requires minimal technical expertise but nonetheless includes some clear intention to do good by preserving and protecting the organization's information technology and resources. One example would be when a trained and aware employee resists an attempt at social engineering by telling a caller that security policies affirm that passwords can never be revealed to a system administrator or anyone else.

Table 2:
Ten Most Extreme Behaviors on Expertise

Item#	Mean	Behavior Description
32.00	4.29	He created a denial of service attack on a competitor's website using the company's computers.
31.00	4.24	She set up a packet spoofing application just to test out her programming ability.
34.00	4.20	He set up a network monitoring device which intercepted data not intended for his system to assess how well the network was running.
67.00	4.13	He built a special script that disabled other users' terminal sessions.
33.00	4.04	She forged routing information to make it seem like someone else had sent some packets.
4.00	1.68	She wrote her password on a sticky tape and put it on her monitor.
66.00	1.61	She chose a password that was "1234."
3.00	1.60	He used his social security number as a password.
13.00	1.55	She wrote her password on a slip of paper and taped it under her keyboard.
2.00	1.53	She shared her account information with a friend.

Table 3:
Ten Most Extreme Behaviors on Intentionality

Item#	Mean	Behavior Description
75.00	4.75	He did a training program to learn about the sensitivity/criticality of special company files so that he could apply appropriate protective measures when handling the information.
76.00	4.64	She did a training program to become familiar with indicators of virus infection and learn how to report operational anomalies to resource administrators.
88.00	4.62	She reported a discovered security vulnerability to the appropriate authorities.
85.00	4.60	He would not release non-public company data/information to a reporter.
94.00	4.57	She used excellent access codes (passwords and usernames) and changed them periodically.
58.00	1.83	She forged her email header information to make it look like her boss had sent a message.
33.00	1.82	She forged routing information to make it seem like someone else had sent some packets.
20.00	1.80	He transmitted a harassing message using the company's email.
24.00	1.75	He used a file decryption program to discover the contents of a file containing trade secrets.
40.00	1.63	He intentionally introduced a Trojan horse program into the network.

Note that the five-point intentionality scale was bipolar and ranged from 1 (“Highly malicious intentions to compromise resources”) to 5 (“Highly benevolent intentions to preserve resources”).

Table 4:
Behaviors Generating Greatest Disagreement

Item#	SD Expertise	SD Intentionality	Behavior Description
35.00	1.60	1.10	She mailed an encryption software CD to a foreign country in violation of international or regional export control laws.
76.00	1.49	.65	She did a training program to become familiar with indicators of virus infection and learn how to report operational anomalies to resource administrators.
81.00	1.44	.86	He participated in advanced security training designated by the organization.
75.00	1.42	.47	He did a training program to learn about the sensitivity/criticality of special company files so that he could apply appropriate protective measures when handling the information.
8.00	1.38	1.17	She constructively criticized organizational security policies to her boss.
69.00	1.29	1.42	He found and saved trade secret information about other companies using the Internet.
34.00	.84	1.34	He set up a network monitoring device which intercepted data not intended for his system to assess how well the network was running.
59.00	1.03	1.33	He used an intrusion detection program on the company's network even though that was not part of his job.
63.00	1.11	1.28	He used unsolicited email to advertise a service offered by the organization.
37.00	1.20	1.21	She deleted a colleague's account information so that he would not be able to access his files.

Note: Bolded values indicate which variable served as the sorting key: The first five behaviors were the most disagreed upon for the expertise variable, the last five were the most disagreed upon for the intentionality variable.

Figure 1: Six-Category Behavioral Taxonomy

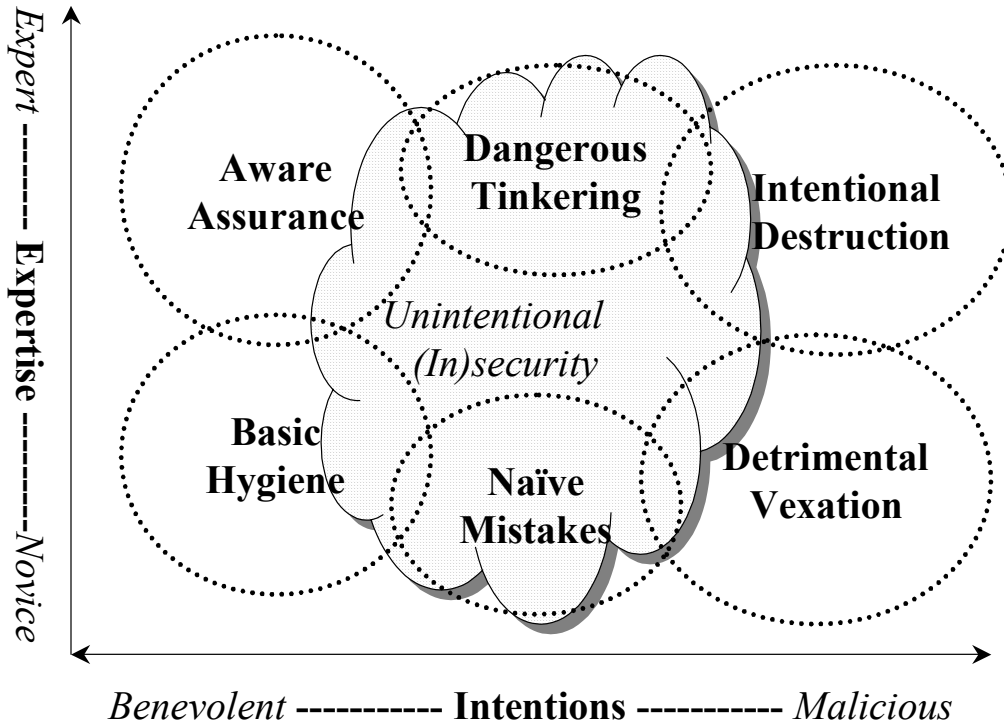


Figure 2: Plot of Means for Expertise

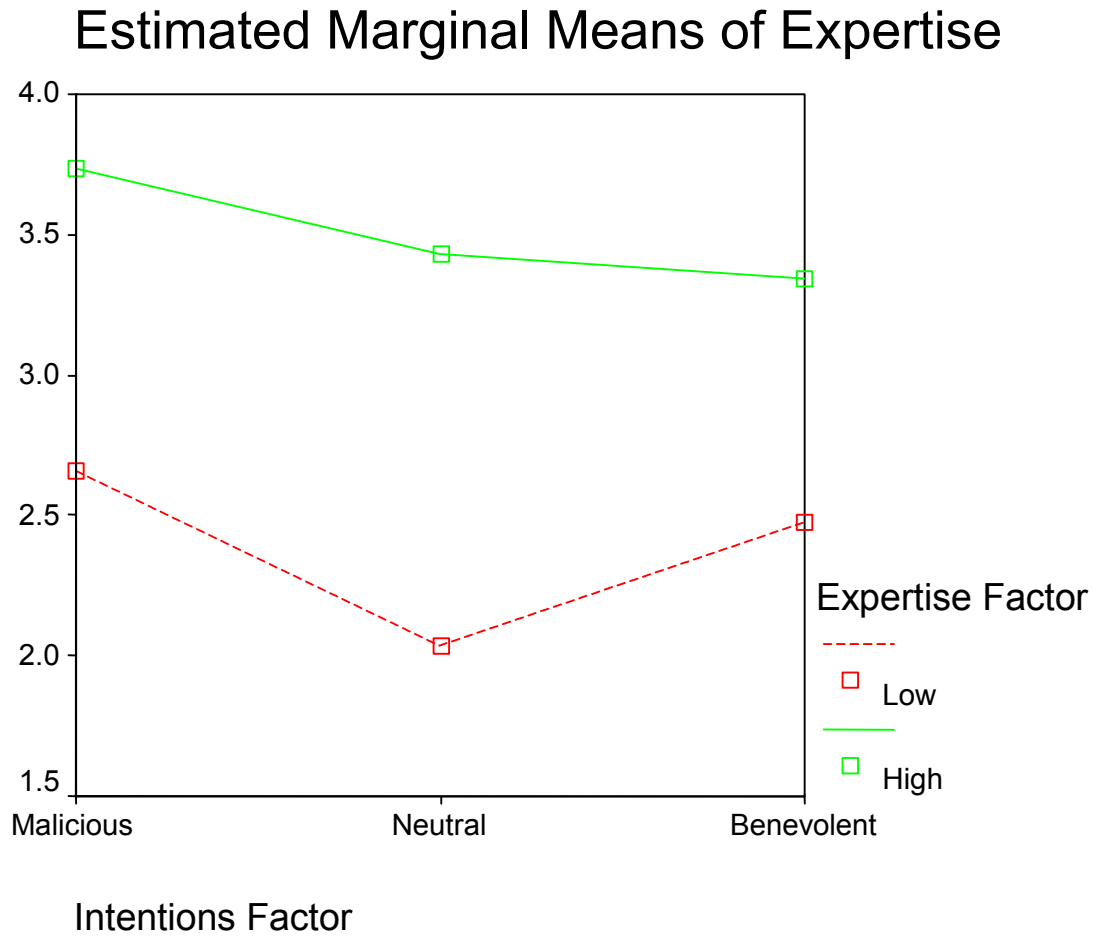


Figure 3: Plot of Means for Intentionality




Figure 4: Plot of Standard Deviations for Expertise




Figure 5: Plot of Standard Deviations for Intentionality

