

CyberTrust: Models for Security Behavior and Dynamic Security Policy

Douglas C. Sicker (PI), Donna Caccamise, Clayton Lewis, Tom Lookabaugh
University of Colorado, Boulder 80309-0530

1 Prior Support

Donna Caccamise: Currently Dr. Caccamise serves as a co-principal investigator and program manager for an NSF IERI grant to develop, implement and evaluate scientifically-based literacy tools in several school districts throughout Colorado. This effort includes the development of intelligent agents who interact with students as they engage with these computer-driven literacy tutors. Prior to moving to academia 5 years ago, Dr. Caccamise spent nearly 20 years as a human factors engineer and manager of security related projects for the U.S. government (DOD and DOE). These projects, which investigated and provided solutions for security readiness for high value assets, involved a contextual approach where security systems were assessed in terms of the integration of equipment, people and policy.

Clayton Lewis: Most recently Lewis served as co-investigator for The Digital Commonsense, an NSF research infrastructure award. His role was to advise on procurement and deployment of computing and communications equipment for a broad research program in information integration. Earlier NSF support for Lewis and collaborators led to the development of the Cognitive Walkthrough, now one of the most widely-used user evaluation methods in industry worldwide, as well as theoretical advances in understanding user interactions with systems.

Tom Lookabaugh: Tom Lookabaugh has supervised students for two summers under an NSF REU grant, focusing on multimedia security and security of Voice-over-IP under potential extensions of the Communications Assistance for Law Enforcement Act. He has served as Co-PI on a NETS Institute grant focused on the evolution of E911 emergency services. This work has supported papers focused both on the technology of security and on policy and regulatory concerns around security.

Douglas Sicker: Douglas C Sicker has been supported by various sponsors including NSF, Internet2, NETS Institute, Colorado Institute of Technology and the state of Colorado. He has also received equipment donations from Cisco, Agilent, Avaya and Finisar. In 2004, Dr. Sicker was co-PI on an NSF ITR grant on agile spectrum management. This work seeks to improve spectrum efficiency by developing and applying secure, cooperative techniques among the radio devices composing a wireless network. Several papers have been submitted along the lines of this work. Dr. Sicker has also received funding for NSF REU sponsored undergraduate research, which focused on the development of agents for enhancing privacy.

2 Introduction

Security and trustworthiness are not properties of technology alone, but also of the social system in which information technology is embedded. We believe that we can make the most rapid progress in changing both the effectiveness and efficiency of security by tackling the problem using coordinated research in psychology and technology. To this end, we have assembled an academic research team with skills in experimental psychology and organizational behavior, cognitive science and human interface design, security policy, security protocols and technology, and management, supplemented by a practicing industry expert security technologist. We propose a two-pronged research program consisting of cross-validating efforts in developing a new psychometric instrument to measure and understand security behavior with new technology based tools to monitor and affect security behavior. We expect our research to help with security technology and human interface design, human resources and personnel management, and adaptive and flexible security

policies. The result should be organizations that are both more secure for a given level of investment in security and more capable of managing to an appropriate tradeoff between security benefits and security costs.

The dominant contemporary approach to securing information technology in organizations is based on risk analysis, policy development, and enforcement. The paradigm is at serious risk, though, when it comes into conflict with economics or with individual and social psychology. The cost of security – both direct and indirect costs – may be either too high or too low, often simply because the resulting policy is not flexible and adaptive. And security rules may be thwarted by individuals, even friendly ones, for a variety of reasons: they are making what they believe are responsible tradeoffs on the organization's behalf (and they may be right), they are not entirely aligned with the organization's goals or don't understand exactly how they relate to a particular security related behavior (and it may be cognitively unreasonable to expect them to), or they are responding to basic social norms or fundamental cognitive limits. A techno-social system that suffers from misalignment between the security aspects of its technology and the behavior of its people is intrinsically weaker both economically and against malicious attackers.

To best affect this, our two pronged strategy first uses psychological and social-psychological research methods to build on our own and others' qualitative research in order to develop a new, validated psychometric instrument for measuring security behavior (a *security behavior inventory*) and a model for predicting it. Our second prong extends our current experimentation in personal computer based security monitoring to embed specific security oriented information, persuasion, and monitoring components in information systems. The two research efforts are designed to reinforce and validate each other. To guide our research, we augment the traditional concept of relating security technology to users through *enforcement* of policy with strategies based on (1) *persuasion* when behaviors are desired but the cost of enforcement outweighs its benefits and (2) *alignment* when it is appropriate to delegate decisions about security behavior to individuals providing we have aligned their interests and capabilities with those desired by the organization. We expect to:

- Improve the design and implementation of IT security policies, procedures and guidelines, by creating adaptive and flexible ones that better fit each application;
- Provide concrete guidance to the management of organizations, including organizational design, culture, training, and compensation; and
- Offer specific tools and processes to guide the future design of IT security technology, including monitoring, analysis, decision making, reporting, and information presentation.

We expect an important benefit to society in the form of more effective and more efficient security. This includes the development of tools that industry and government can employ to achieve unprecedented levels of IT security. We expect the tools we develop to not just monitor security compliance, but to implicitly educate users in the affected sections of the workforce about security and how to interact with it, and to do so in a timely and efficient manner by being embedded in their normal work routines. And we are also pursuing other benefits by using the project to train graduate student researchers in the rare skill of effective interdisciplinary research, and to generate new and important curricular materials for our institution and others in the areas of psychology, computer science, and management. To insure the broadest possible impact of our work, we plan a dissemination campaign based on a combination of traditional research conference, publication, web distribution, and researcher training with the development of materials and their presentation to classes, trade shows and industry conferences and periodicals, consultancies, and government agencies.

3 Existing Work in Security Behavior, Policy, Technology, and Work

3.1 *User behavior, productivity and motivation*

Emphasis on application of social sciences tools to information system security may be recent, but the problem of the “human factor” in security is an established one. As Bruce Schneier points out, “people often represent the weakest link in the security chain and are chronically responsible for the failure of security systems,” ([SC00], p. 255); historically technically strong military security systems have consistently been weakened by spies, clerical mistakes, and simple human foibles. And, the notorious hacker Kevin Mitnick’s describes social engineering (the process of eliciting human indiscretions to hack a computer system) as follows: “A lot of companies are clueless, because they spend most or all of their security budget on high-tech security like fire walls and biometric authentication – which are important and needed – but then they don’t train their people. All it takes is for a hacker to place one phone call to an unsuspecting person at the company and pull out the information from them.” [KO02]

We take the perspective that such incidents can be better understood not simply as evidence of human frailty but rather as instances where humans are making choices, whether consciously or not, between behaviors that promote security and those that weaken them. In other words, we should ask: how are humans making security behavior tradeoffs in the context of the socio-technical system in which they are embedded?

The evolution of security thinking towards this socio-technical perspective can be thought of as a sequence of genres – James [JA96] describes audit and checklist, risk analysis, cost justification, and finally “high level conceptual management models,” while Siponen [SI01] describes information/data base modeling, responsibility modeling, and “security-modified IS development” – but it may be useful instead to consider some of the major concepts developed to date.

It is well known that users don’t necessarily understand the reasons for security or their role in security [AD99, GA00, BR02, SA01]; this makes it difficult to expect intrinsic motivation. Further, as security needs increase, policies become more intrusive and conflicts with other goals increase. Users can respond to such conflicts by emphasizing their other goals (e.g., productivity) and circumventing security policies; moreover, once adopted such behavior can become self-reinforcing as once a user has circumvented the policy, he or she concludes that the policy is not that important anyhow [AD99, BR02, WH01].

At an organizational level, this can be addressed by “responsibility modeling,” [ST93, BA96], which tries to capture how an organization actually gets work done in order to influence security requirements; in this sense it attempts to transcend the more prevalent strategy of generating security requirements from threat analysis and apparent work processes (e.g., [MC99, AL03]). The challenge with responsibility modeling is the necessity of a detailed and complete specification of real work processes; this is both laborious and inflexible. An alternative is to focus on the relationship between security and culture (admittedly with less predictable results): Dhillon and Backhouse, for example, call for instilling responsibility, integrity, trust, and ethicality [DB00] while Weirich and Sasse propose using principles of social marketing and fear appeals to elicit a security conscious culture [WE01]. But, whether strategies aimed at influencing organizational culture are sufficiently effective is a major open question in socio-technical approaches to security.

A number of researchers have examined trends in creativity, innovation, and productivity [DR99, LA00, AR98, BA99a, VE90, BA99b]. Much of this work looks at the role of

motivation and empowerment and its impact on productivity. These trends in modern work practices can run strongly counter to security practices that come more from the military “command and control” perspective. In particular, the notion of increasing the role of knowledge, distributed decision making, and risk taking is antithetical to directly circumscribing freedom of action by users in pursuit of higher security. We can summarize tensions between current security practice and evolving practices below.

	Contemporary Security Practice	Evolving Work Practices
Communication	Need-to-Know	Open Book Management
Work Processes	Documented and controlled in order to meet security requirements	Flexible to promote innovation, subject to change by empowered users
Roles	Need to be specified in order to assign information access rights	Flexible to promote intrinsic motivation
Work Culture	Command and control; responsibility for security centralized	Security responsibility diffused.

Table 1. Conflicts between traditional security practices and evolving work practices.

3.2 Security Policy, Practice and Technology

In 2002, the National Research Council published its most recent report on the state of computer security and in it concluded that, “not much has changed with respect to security as it is practiced,” since its first report in 1991. [NRC02] While researchers are contributing considerable effort to improving the practice of security, the effective implementation of this work lags. This is particularly true when we consider the process of implementing security policies and practices. This gap is well understood by the security community and often cited in journals and the popular press. [EC02] An interesting and somewhat ironic historical note can be seen in the title of an article published 13 years ago, “On the Buzzword ‘Security Policy’”. [ST91] Security policy continues to be a buzzword, if for no other reason than that it is still not being done well. In the following, we review previous work in the areas of security policy, its impact on users and role of technology in security policy.

3.2.1 Security Policy

A security policy defines the rules for securing a system. [CC98] Generally, these are high-level management specifications that capture the security goals of an organization. These high-level goals are brought closer to the security process by guidelines and procedures that translate the policy into something more specific, such as step-by-step processes or user requirements. These processes can be captured in frameworks for security, such the OCTAVE methodology from CMU. [AL99] To track the evolving nature of computer security, such policies and practices are described in terms of a life cycle. The life cycle concept, as described in [CC98], is a fairly slow moving process. The US government (and particularly, the Department of Defense) exerted a considerable amount of effort in developing security policy specifications over the last three decades, see [DoD79, DoD83, DoD85, NCS87, CC94]. More recently, a number of researchers have examined ways of improving the composition and enforcement. Bauer, Walker and Ligatti have contributed significantly to designing enforceable policies through the use of security monitors. [BA02a, BA02b, BA02c, [BA03] Erlingsson and Schneider are also working in this area and have developed an inline security reference monitor. [ER99, ER00]

3.2.2 Policy and User Research:

Much of the efforts in security policy research and design focus on developing comprehensive security policies. Walton [WA02] developed a framework for the process of adopting a more comprehensive security policy. Outlined in this research is the necessity of

the security policy implementation team to achieve a certain level of ‘buy-in’ by users regarding the usability of the procedures and technologies manifested by such a policy implementation. In addition, the author stresses that the policy must maintain a flexible and dynamic structure, and be able to incorporate information obtained regarding the secured systems and the behaviors of users within those systems. Sasse and Adams [SA99] question the notion that users are not interested in behaving securely. They note that lack of information regarding proper password security behaviors might not best be addressed through harsh enforcement. “While we advise against “punishing” users who circumvent security mechanisms, such behavior needs to be detected and challenged in a constructive manner ... At the same time, an environment giving the impression that its security mechanisms are invincible is likely to foster careless behavior among users, since the level of perceived threats to security is low.” Sasse, Brostoff, and Weirich [SA01] in their research involving security problems that arise out of the password requirements contend, “Security designers must identify the causes of undesirable user behaviour, and address these to design effective security systems”. Dourish [DO03] discuss the notion of developing security measures within software and systems that better reflect the “relationship between security facilities and everyday work”. They state:

“Since security requirements depend on the specific circumstances of action are subject to continual reflection and revision, it is necessary to provide people with the means to understand the security implications of the current configuration of technologies at their disposal. Rather than being “transparent”, then, security technologies need to be highly visible – available for inspection and examination seamlessly as a part of work.”

Fogg coins the word “captology” to describe the use of computers as persuasive technology (hence “capt”). [Fog03] Captology extends theories of human persuasion to persuasion of humans by machines, and a number of the concepts presented by Fogg directly inform possible mechanisms for affecting security behavior, including notions of persuasion by reducing complexity, guiding interactions, customizing interactions, providing timely suggestions, facilitating self-monitoring, and using appropriate surveillance and conditioning. Weirich and Sasse [WE02] discuss the use of persuasive tactics embedded in “the mechanisms themselves, policies, tutorials, training and the general discourse...” related to organizational security. Two of the conclusions that they reach are, “Users' willingness to make the extra effort that security conscious behavior requires is a vital variable influencing the effectiveness of this system.” and “Users cannot be forced to behave in a proper fashion, but an effort to *persuade* them to do so has to be made.” According to de Souza, Basaveswara and Redmiles [DE02], the theme of user awareness “involves information about the activity of software systems as well as the activity of human collaborators.” These researchers identify visualization as a tool to promote user awareness with the use of gauges. A continuation of this research from Dourish and Redmiles [DO02] takes the concept of visualization to promote awareness with respect to system security. Specifically, the architecture they propose is “designed to gather, integrate, and interpret information about security, which is distributed across a large number of systems and components: and then, to present this information as a set of real-time visual displays.”

3.2.3 Technology Policy Research

The captology concept proposed by Fogg parallels much of the work done in the area of *computer agents*. In simply terms, an agent helps the user perform some task, but can also be used to influence decisions that the user makes. Below, we describe how various researchers have applied this agent concept to security.

Jendricke and Markotten [JE00] attempt to assist users in better utilizing the security functionality of applications with the use of an “Identity Manager” that allows the user to configure security settings and preferences within the manager, which then dictates these preferences to all applications accessed through this identity. This addresses the issues that users avoid security functionality within applications because of the inconvenience involved with configuring this functionality within each application utilized. Ackerman and Cranor [AC99] examine the feasibility of applying privacy critics (semi-autonomous agents) to “help people protect [users’] online privacy by offering suggestions and warnings”. The authors note that these agents do not initiate actions themselves; they simply inform the user of the potential consequences of a given behavior, and allow the user to decide how to proceed. In [VI98] several authors discuss similar work, particularly within mobile agent security; this work mainly focused on plug-ins and Java Applets for web browser security. The development of security monitoring tools represents another area closely related to our research. [JE00] [DE02] We plan to extend some of the prior work done in this area.

4 A Proposed Framework for Security and Behavior and Initial Work

As our understanding of security in human and organizational contexts expands, we can augment the core concept of “enforcement” with two different perspectives about how we can expect security technology to interact with individuals: “persuasion” and “alignment”.

Enforcement. Enforcement dominates traditional conceptions of the relationship between individuals, the organization, and technology appropriate to managing security. If the expected risk to an organization dominates any direct or indirect costs associated with proscribing individuals’ behaviors, then enforcement is the correct paradigm. In this case, our focus becomes making enforcement efficient and effective by developing and using the more detailed understanding of human interactions with security technology that we hope to develop through our research.

Persuasion. Enforcement may be inappropriate, though, because it is impractical or too costly relative to its benefit; in other words, while we may not specifically be interested in enabling a range of behaviors and individual discretion in making choices regarding security, the direct and indirect costs of completely effective enforcement exceeds the benefit of risk reduction to the organization. For these cases, we try to influence behavior, expecting that through persuasion we can increase compliance and lower risks and losses, although we accept that we cannot guarantee compliance. Persuasion plays a central role in the work of Sasse and colleagues [WE01, SA01] and we can link this to implementation ideas drawn from Fogg’s work on computer-based persuasion [FO03].

Alignment. Rather than trying to prevent or discourage active trade-offs in security behavior, we may expect them and work instead to maximize the likelihood that the trade-offs made are optimal from the overall institution’s perspective. This is the process of alignment that is the core of much contemporary work in human resource management [BA99b]. Alignment involves a combination of education, dynamic information, and motivation that influences individuals in the instant at which they make trade-offs, but respects the ability of intelligent individuals “on the spot” to make reasonable decisions. We would like to help the individual to accurately estimate risks and costs, to do so from the same perspective that would make sense to the organization, and to make a rational decision based on these estimates. The concept of alignment as a mode of interaction between individuals and organizations regarding security is novel to the best of our knowledge, although consistent with general trends in job and work design.

4.1 Initial Qualitative Research and Model Building

We initiated some limited qualitative research to inform a model of security behavior, building on the literature cited in Section 2.1, and in particular, the work of Dourish et al. [DO03] and a series of efforts by Sasse and colleagues [SA01, WE01, BR02]. Semi-structured interviews were conducted with employees of two medium sized high-tech companies, Company A and Company B. Company A has roughly 180 employees and 11 were interviewed, all being engineers with the exceptions of the CFO, Network Administrator, and Director of IS. Company B has roughly 350 employees and 7 were interviewed, all being engineers with the exception of the CFO and Director of IS. The intent of interviewing two companies was to gain a better understanding of the data through comparing the results from each data set. The focus on engineers was based on the premise that these individuals are subject to variation in work content (e.g., critical project deadlines, urgent customer interactions) that would highlight conflicts with static security policies and that these individuals have more than average ability to circumvent information technology security; hence they represent a population that is likely to demonstrate the effects we are seeking to study. The interview questions pertained to how each user interacts with security, what affects security has on their work (positive or negative) and how they value security. The particular interview based methodology used was grounded theory [CO98].

Several key themes emerged from our analysis of these interviews. In particular,

- Individuals have a strong preference for security to be invisible
- Users are unsure of proper procedures when dealing with security – there is a distinct variation in “common sense” security practice
- Perceptions vary by position held by employee
- The structure of the organization affects security openness

Our qualitative research combined with the literature already suggests several components of a potential model for security behavior, although we stress that it would be premature to develop a detailed model at this point; such an effort will be driven by the substantial data that will be collected as part of the proposed research via interviews, surveys, and measurements in both laboratory and open settings. Nonetheless, we expect to explore in particular the following relationships:

- The dependence of average security behavior across on organization on perceptions among organization members of financial slack, trust, exposure of the organization to threats, necessary and desirable flexibility, and innovativeness.
- The dependence of an individual’s average security behavior on organization level influences, demographics, individual experience, and exposure to training and education on security issues.
- The dependence of an individual’s instantaneous security behavior on average security behavior, on task context, and on individual characteristics such as interpersonal trust, innovativeness, conscientiousness, frustration, and cognitive limits.

4.2 Initial Technological Research

Motivated by our preliminary model for security behavior, we propose two linked models for designing security interventions. The first model aims to make security policy responsive to data about actual security behavior. The second model describes appropriate security interventions, based on technical, management and policy considerations. We then describe an agent-based intervention tool that we have implemented based on these ideas. This agent

takes the form of a software program running on a machine and can be represented to the user as a small icon on the desktop. This agent monitors the way an individual conforms to or deviates from the expected security requirements and provides feedback to the user. The agent also reports to a central network monitoring software that runs on the organization's network. The output of this agent-based system provides the input for a dynamic security policy life cycle.

4.2.1 Model for Dynamic Security Policies

We begin by suggesting a more dynamic approach to security policy and practice¹. We move away from a model of enforcement of rigid rules towards a model based on persuasion and ultimately alignment between the organization's preferred decisions and the individual's tradeoffs between security behaviors and other goals. The new model changes the interface between individuals and information technology to encourage better security behavior decisions through techniques such as monitoring and feedback, timely provision of relevant information, personification, and other contemporary persuasion oriented techniques from the human-computer interaction community. We incorporate a new sub-component, which will influence each phase of the corporate security policy life cycle: the behavior agent. As shown in Figure 2, the subcomponent extends traditional security models [CC94] to add a dynamic to each phase of the security policy life cycle. The behavior agent monitors actual security behavior (assessment), provides data to management, and influences both education and deployment practices in the organization. The fundamental goal is to allow the user and the organization to consider security in a less rigid and static manner.

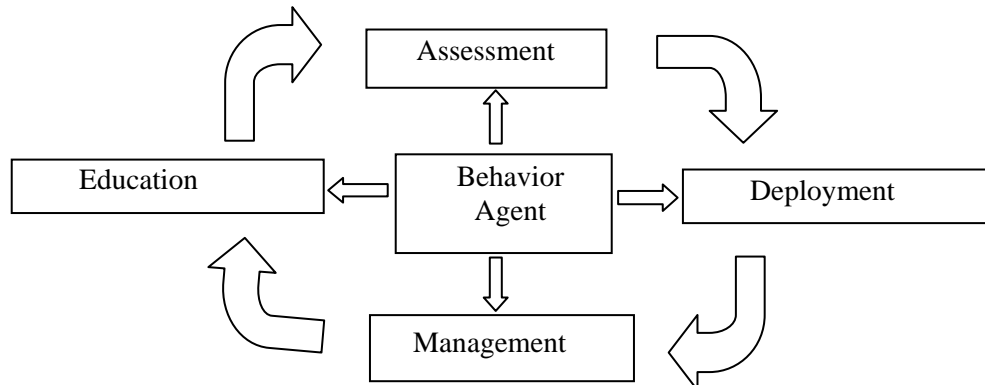


Figure 1. Model for Dynamic Security Life Cycle

4.2.2 Intervention Model

We envision three types of intervention that can promote better security outcomes for organizations (i.e., outcomes that better represent the tradeoffs that one would want made²):

- Technology interventions– new technologies or modifications to existing technologies.
- Policy interventions– changes to the content and even the character of security policies. In particular we envision moving from static policies not just to flexible policies but to dynamic policies, policies that change in response to data.
- Management interventions– changes in the way organizations relate to employees (e.g., through information dissemination and motivation).

¹ We use the term *Security Policy* in a generic manner to represent (1) the life cycle of security policy, practices and deployment and (2) the actual managerial guidelines for securing an organization.

² Note that we are careful here to avoid saying “more security” as there is a point at which security becomes too expensive to be justified (either in its direct costs or costs it creates for other activities).

In realizing these ideas, we will develop a software agent that will allow us to experiment with all three modes of interacting with security: enforcement, persuasion and alignment, and drive all three kinds of interventions. The agent will engage in three canonical tasks:

- Observation – measuring behavior and recording events that are observable to the agent. Research issues here include maximizing the set of relevant observations (including such things, for example, as settings that affect security in a user’s local environment, behavior with respect to accesses control) and explicitly considering tradeoffs with privacy.
- Data Analysis and Decision-making – determining frequencies or trends, comparing to templates or limits, etc., and deciding what actions are required. Research issues here include linking observable behavior to an analysis tool that can be used to discover trends and determine appropriate actions.
- Action – taking action on the base of decisions. Actions may span the space of alignment, persuasion and enforcement. For instance, for alignment purposes, we might simply provide information intended to insure that a user is correctly estimating risks taken. For persuasion, we might provide specific suggestions and indicate that we are sharing observations with others (e.g., the individual’s supervisor). For enforcement, we might advertise and then take action designed to punish a breach or we might proactively deny access to computer or network resources. The same framework should be able to effectively scale through these paradigms depending on an organizational analysis of the relative cost and risk of behaviors. Moreover, the system will be designed to be dynamic and require limited training on the part of the user; the feedback, suggestions, and actions are intended to be self-explaining. Other actions might include altering technology in ways to make deviation more difficult, providing education to the users about the intention of the policy or altering a policy no longer useful or effective.

In this effort we will build on pilot work already carried out using two simple prototype security agents. These prototypes have allowed us to explore more concretely the ideas in our framework, and how they can be implemented. Within these agents, we have incorporated aspects of technology, policy and management intervention. One prototype was built on a Linux platform and the other is a C# program, running under Windows, and have the following functions.

Observation and Monitoring. The agents examine a wide range of security settings, including for example Microsoft Security Configuration settings, the security audit log for the machine (which includes records of successful and unsuccessful login attempts) and the list of currently running processes. It also detects software update (patch) requests. A central server collects data acquired by the agent and subsequently serves as a source for propagating this information. We have developed (and tested) other similar capabilities in our agents.

Data Analysis and Decision-making. The agent detects particular events, such as unsuccessful login attempts, patch failures, weak passwords (detected through an automated password cracking tool), and so on. In total, we now can analysis 20 such events. Analysis of data at the server is planned but only partially implemented.

Actions. The agent responds to events in ways that reflect the alignment, persuasion and enforcement aspects of our framework. An *alignment* action would be to notify the user of a risk associated with a behavior, such as ignoring a patch request. Carrying on with the patch example, the agent can also perform actions intended to *persuade* the user to respond, for example by sending a message that warns the user that a note will soon be sent to their supervisor if the patch is not installed. Finally, the agent can *enforce* a response by the user by blocking network access if a patch is not applied.

The information gathered by this agent can be incorporated back into the security cycle. The feedback can be done statically (through a typical security policies/procedures review) or dynamically (through agents altering security procedures described by certain policies). We can then expand this to further technology interventions, security policy modifications and recommendations for management. This will allow us to validate our model of human behavior relative to security and our interventions in order to influence that behavior.

5 Research Plan

Our research plan builds on our initial work in qualitative research on security behavior and model building and on technological tools to measure security behavior. Our specific goals consists of four elements and their cross validation as shown in Figure 2.

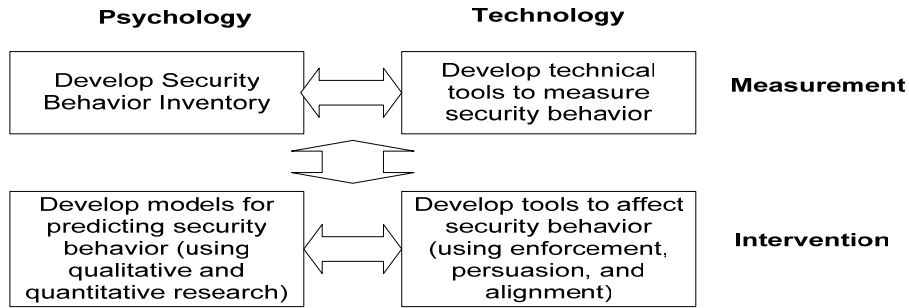


Figure 2. Overall Structure of the Research

Our research design consists of overlapping phases. The first phase builds on work we have already initiated and focuses on measuring security behavior. The second phase focuses on intervening to affect security behavior. The first and second phase overlap as we expect to use results from experiments in intervention to revisit the models developed for measurement. Each phase includes complementary psychology and technical research. Figure 3 lays out a proposed time line for our research. Details on each of the four elements (plan, methods, and resources) follow. Each element has a lead researcher team assigned. Doug Sicker will provide overall project management. We will develop four reports, each accompanied by dissemination to the academic and user communities as described in Section 5 at the times shown in the Figure. We anticipate publishing in multiple academic disciplines, including computer science, psychology, information systems, management, and policy. A website devoted to this effort will provide the current status of our efforts, provide a collaborative opportunity for other researchers to engage in dialog about the effort, and general serve as a portal to help disseminate our tools and finding.

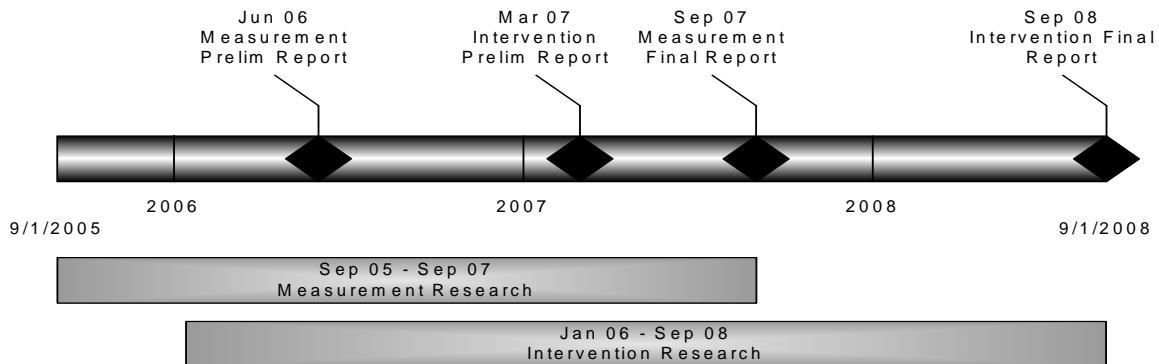


Figure 3. Research Timeline

5.1 Develop Security Behavior Inventory

We propose to develop a *security behavior inventory*, a new, validated psychometric instrument for measuring the behavior of an individual with respect to information security. Caccamise and Lookabaugh will lead this effort.

5.1.1 Plan

Developing a psychometrically reliable and valid instrument involves several key steps. We select candidate elements of the instrument based on related work, in our case the substantial work in the traditional security community to construct checklists of desirable and undesirable security behaviors, and combine this with pre-existing work in related validated psychometric instruments, such as those relating to safety. We have already constructed a preliminary, not yet validated list of behaviors. This is the basis for further development that will result in a statistically reliable test instrument that will also have concurrent validity with other measures that are related to subscales of component behaviors. We will also establish predictive validity for reported security behavior and observed security behavior. During this process we will simultaneously ask individuals to complete the instrument and monitor behavior of these individuals using the other system-driven measurement tools being developed in parallel.

5.1.2 Resources

We will recruit and supervise a Ph.D. candidate graduate student in psychology or educational psychology with a particular interest in the development and validation of psychometric instruments. We will also have minor expenses in executing surveys.

5.1.3 Broader Impact of Security Behavior Inventory

Development of such an instrument is not only central to our research, but should have two important broader impacts:

- Validated psychometric instruments are a building block for all researchers in a field and we would expect our work on a security behavior instrument to assist other researchers working on many aspects of psychology and security.
- Such an instrument or its derivatives can be very useful in personnel and human resource management, particularly in situations in which the expected cost of security breaches is very high (e.g., national or homeland defense). Examples where previously developed psychometric instruments have become central to human resource management include personality inventories such as Meyer–Briggs. These tools will be developed with scaling and sustainability in mind so has to have a national impact on IT security in industry and government enterprises.

5.2 Developing Technical Tools to Measure Security Behavior

We will develop a family of tools intended to measure security behavior. Such tools have three important roles: (1) most appreciably, they provide a method for examining how users interact with security policies and procedures, (2) they form an important building block for tools we will later develop to affect security behavior and (3) we use them along with our survey-based and empirically validated security behavior inventory to develop and subsequently cross-validate each other. Sicker and Lewis will lead this effort supported by Escamilla.

5.2.1 Plan

We envision three types of capabilities for such tools:

- Client based tools reside on end user devices and monitor behaviors of the user in interacting with the device that could affect security. Such tools might be automatic or involve user action (facilitated self-monitoring). We include strong assumptions on privacy including, notably, full transparency of monitoring. We have already prototyped a Windows-based PC and a Linux version of such a tool as detailed in Section 3.2, although we expect to expand and refine this tool throughout this phase of the project.
- Network based tools operate across a network to monitor behaviors and interact with the client tools. As with existing monitoring tools, a network-based solution will allow us to correlate meta-events for trending, which will provide feedback into the inventory and behavior models. Here again, we assume a strong transparency model.
- External monitors measure security behaviors not directly visible from a client or across a network; for example, physical access to a computing facility (e.g., a machine room) can often be monitored by a card reader access log or by the use of cameras. We expect to target external monitoring strategies for elements of our proposed security behavior index that are difficult to validate via client or network based tools. In limited cases, we may use manual monitoring or logging for validation purposes only. Here again, we assume a strong transparency model to manage the impact on privacy.

5.2.2 Objectives

- Develop a set of security behavior criteria that can be monitored by technology tools.
- Develop a set of monitoring tools to observe security behavior.
- Incorporate these tools into refining our security behavior inventory.

5.2.3 Resources

We will recruit and supervise a Ph.D. candidate graduate student in computer science with a particular interest in the development, integration and testing of a family of security behavior measurement tools. We expect that this student will (1) develop and test the tools, (2) assess user interaction and (3) cross-validate the inventory model.

5.2.4 Broader Impact of Behavior Monitoring Tools

Development of such tools is not only central to our research, but should have three important broader impacts:

- Such tools could be used to discover aspects of how users interact with technology.
- Such tools could be used as educational instruments, to disseminate understanding in such areas such as privacy and monitoring.
- As part of an existing REU grant, we will sponsor undergrad research projects in this category through the NSF security REU project that our department currently supports.

5.3 *Develop Models for Predicting Security Behavior*

We propose to develop and validate a model for predicting security behavior based on organizational and individual factors. Caccamise and Lookabaugh will lead this.

5.3.1 Plan

We will continue to execute targeted qualitative research (in line with the “theoretical sampling” concept of grounded theory [CO98]) to augment our model building. We will select from previously validated psychometric instruments to measure proposed causal factors. We then use surveys and observed and measured behavior in usability laboratory settings and in open environments to validate relationships between proposed causal factors and our security behavior inventory. Because instruments used to measure causal factors can each require multiple survey questions, we will need to develop and validate our model in

phases. We also expect to cross-validate the model with specific technology based tools: where a presumed causal relationship exists and we have an opportunity to affect the cause via a technology tool, we should be able to verify or not the strength of the relationship from the measured security behavior that results when the tool is invoked. This type of experimental cross-validation will continue throughout this phase.

5.3.2 Resources

We will recruit and supervise a Ph.D. candidate graduate student in psychology, educational psychology, sociology, or communication with a particular interest in the development of social and psychological models and their validation using triangulated (qualitative and quantitative) research. We will also have minor expenses in executing interview and survey based research.

5.3.3 Broader Impact of Predictive Models

This element of our research is expected to have direct impact on how security policies are implemented in organizations and on management practice. To accelerate these impacts, we expect to develop a new model for security policy design and implementation and disseminate it through the avenues described in Section 5. This model will become a tool both for educators and security practitioners in organizations. We also expect to develop accessible articles for the general management press that help define our proposed “best practices” for security management and the relationship with organizational design and personnel management.

5.4 *Developing Tools to Affect Security Behavior*

We propose to build on our work in tools specifically designed to affect security behavior. Sicker and Lewis will lead this effort, supported by Escamilla.

5.4.1 Plan

These tools will have the same fundamental types described in Section 4.2 (client, network, and external tools) but will draw on a number of techniques to determine how they interact with users, grouped roughly on our categories of enforcement, persuasion, and alignment. *Enforcement* tools build on our measurement tools to either directly preclude a particular undesired behavior or provide a predictable, anticipated, immediate and severe response to an undesirable behavior, in analogy to enforcement in the physical world. *Persuasion* tools invoke a variety of methods to influence behavior without the cost of attempting to enforce it, including reducing complexity of desirable security behaviors, guiding users through desirable behaviors, customizing the user interaction, providing timely suggestions, facilitating user self-monitoring, exposing results of user surveillance, and conditioning user behavior [Fog03]. We build in full user awareness of both the function and intention of persuasion tools to manage privacy and other ethical issues. *Alignment* tools anticipate and encourage user decision making about which security behaviors to adopt in a particular circumstance but attempt to align choices with the organization’s overall interests. These include information tools designed to provide cognitively manageable information about security tradeoffs at the time they occur, security management tools designed to assist users in minimizing the risk entailed in chosen behaviors (e.g., automatically closing an opened firewall port after an agreed time), and motivational alignment tools, such as maintenance of a “balance” of security behavior “costs” accruing to a particular individual based on their behavior, and visible in either a private, semi-private (e.g., visible to supervisor), or public (e.g., visible to peers) mode.

A number of inquiries will be used as the basis in assessing each approach. We will evaluate an individual's understanding of the policies and the organization's understanding of the policies (as assessed from trend data about the users as a group). We will assess the effectiveness of dynamic policies and practices. Lastly, we will compare the effectiveness of private, peered and supervisory monitoring, as well as the overall effectiveness of enforcement versus persuasion versus alignment.

5.4.2 Objectives

- Further develop our set of tools to influence security behavior.
- Combine these tools with the security behavior measurement techniques to create a dynamic security life cycle tool.
- Develop of set of criteria to assess the individual's understanding of policies, influence of education, and the effectiveness of persuasion, alignment and enforcement.
- Develop and test the benefit and detriments of incorporating dynamic security practices.

5.4.3 Resources

We will recruit and supervise a Ph.D. candidate graduate student in computer science with a particular interest in the development and testing of a family of security behavior tools. Additionally, to supplement our existing security research facility, we will require the hardware and software necessary to support this research. This includes four personal computers and a limited set of software development tools.

5.4.4 Broader Impact of Technical Tools

We expect the tools we develop to be directly useful in monitoring and influencing security behavior across a wide range of environments. Similarly, we expect that a new understanding of dynamic security policy could influence policy development in many institutions. Both of these elements also lend themselves directly to inclusion in education and training activities, both in standard academic courses and in tutorials and workshops targeted at security practitioners and users.

5.5 *Summary of Broader Impacts*

Our work will have impact on the security research community and on the practitioner community. Validated psychometric instruments, like the Security Behavior Inventory we will develop, become building block for all researchers in a field. This instrument will also be widely used in personnel and human resource management, particularly in situations in which the cost of security breaches is very high (e.g., national or homeland defense.)

Other results from our project will also find wide use in research and practice. Our predictive model of security behavior, security behavior measurement tools, and tools that influence security behavior will all be made available on our project website. We expect these products to shape widely used commercial tools and practices.

To enhance the impact of our results, we will develop an inclusive model for security policy design and implementation and disseminate it through multiple channels as described in our dissemination plan. This model will become a tool both for educators and security practitioners in organizations.

As educators, we will also apply the research directly in the several security, management, and human computer interface oriented classes and certificates that we teach. As we hone the application of this material in teaching and training settings, we will provide it for use by other educators through our project website.

6 Dissemination Plan

We expect our results will be useful to both researchers and security users. As shown in Figure 4, we propose to take our three major research products - (1) a new, validated security behavior inventory, (2) a set of technical tools to predict, measure and modify behavior, and (3) a new model for security policy design - and make them available in the form of standard research papers and data, monograph and tutorial information meant for users and security consultants, and participation by our research team in workshops, panels, and presentations accessible to the general user community.

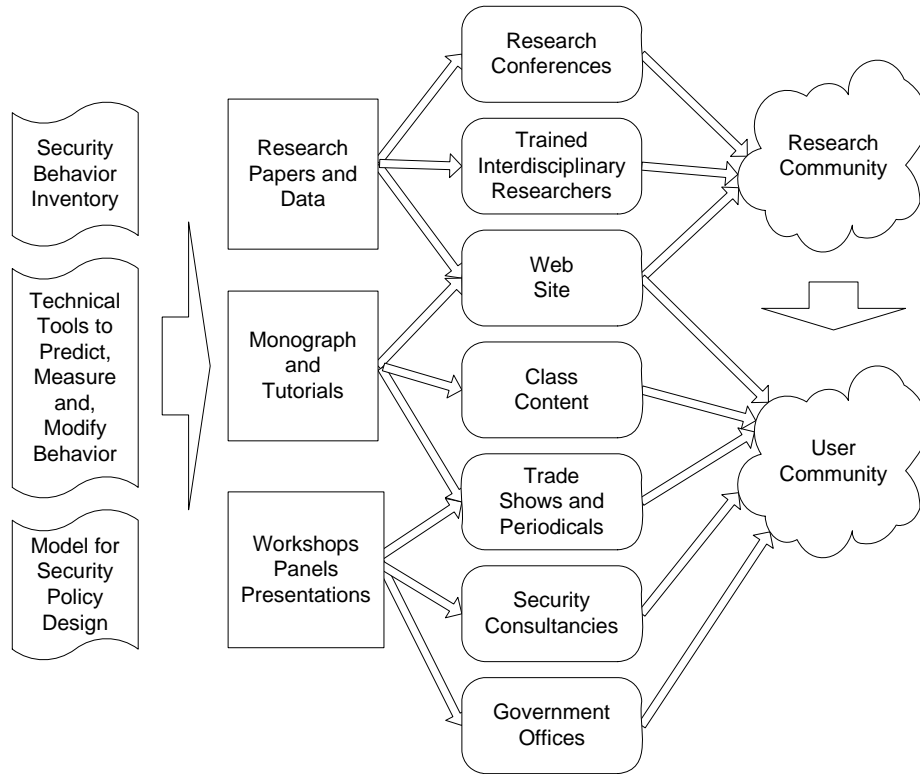


Figure 4. Dissemination Plan

For the research community, we plan to present in traditional research conferences and journals, through a dedicated project web site, and by graduating trained interdisciplinary researchers with a background in security. For the user community, we will use multiple channels: a web site with material understandable by users, new curricular content for classes on security, participation in high visibility trade shows and periodicals (e.g., Communications of the ACM), direct interaction with security consultancies that influence and educate end users (e.g., we have existing relationships with IBM and Coalfire security consultancies), and presentations to influential government offices that inform both government internal and public practice (e.g., the Personnel Security Research Center in the Department of Defense - PERSEREC, the National Threat Assessment Center in the Department of Homeland Security, the Office of Security in the Department of Energy).