

Death by a Thousand Facts: Criticising the Technocratic Approach to Information Security Awareness

G. Stewart¹ and D. Lacey²

¹Risk Intelligence Ltd, United Kingdom

²David Lacey Consulting, United Kingdom

e-mail: geordie@risk-intelligence.co.uk, rabbitscottage@tiscali.co.uk

Abstract

Mainstream information security awareness techniques are failing to evolve at the same rate as automated technical security controls. Humans are increasingly seen as the weak link in information security defences and attackers are starting to prefer exploiting human factors such as greed, curiosity and respect for authority. Problems with human behaviour in an information security context are assumed to be caused by a lack of facts available to the audience. Awareness therefore is largely treated as the broadcast of facts to an audience in the hope that behaviour improves. There is a tendency for technical experts in the field of information security to tell people what they think they ought to know (and may in fact already know). This “technocratic” view of risk communication is fundamentally flawed and has been strongly criticised by experts in safety risk communications as ineffective and inefficient. To improve the effectiveness and efficiency of security awareness techniques this paper leverages safety risk communications which is a mature discipline with common objectives. A critical feature of safety risk communications which is missing from the information security approach is a set of methodologies to systematically evaluate audience requirements. Accordingly, this paper explores the concepts of bounded rationality, mental models and the Extended Parallel Processing Model in an information security context.

Keywords

Information Security Awareness, Human Vulnerabilities, Risk Communication, Safety, Mental Models, Extended Parallel Processing Model, NIST 800-50, Bounded Rationality, Satisfying

1. Introduction

Over the last twenty years there have been enormous advances in the sophistication and maturity of automated technical information security controls. Advanced, automated technical controls such as client based firewalls, anti-virus and real time patching are now common.

Despite the presence of advanced technical controls, information systems remain vulnerable because of human behaviour (Lacey 2009). There is growing evidence to suggest that human vulnerabilities are increasingly being seen as an easier option to exploit information systems (Deloitte 2009). There are a number of reasons why this is the case. Researchers have noted that there are problems with the usability of information systems (Parkin et al 2010), unreasonable risk trade-off decisions

expected of users (Herley 2009) and limits to human tolerance to comply with instructions (Beautement et al 2008).

However, the authors of this paper propose that another key problem is the approach to information security awareness which is used to help control human vulnerabilities. In contrast to technical controls, awareness techniques are stale with little sign of innovation in the methodologies or approach. The approach used today is much the same as it was ten or twenty years ago (McIlwraith 2006).

If the current trend continues then human vulnerabilities will be an increasing target of attack. It is therefore of paramount importance to improve the effectiveness and efficiency of awareness techniques. Effectiveness, in that the objective is being successfully achieved and efficiency in that the objective is being achieved at an acceptable cost.

Safety risk communications is a parallel discipline that offers an opportunity to leverage successful methodologies for use in security awareness. The safety field is significantly more mature than information security and safety theory is supported by a large body of academic research. Many of the techniques of safety risk communication are applicable to information security awareness because of the common goal to help people recognise risk, advise them what they can do to control it and motivate them to take action (Stewart 2010). Collectively, safety communication techniques offer the means to provide a more structured approach for what users are told, how they are told it and how often they are told it.

An abundance of empirical data is available for the effectiveness and efficiency of safety risk communications because safety reporting is often a matter of public record which is enforced by legislation. Good examples are the results for public information campaigns on issues such as drink driving and smoking (Delaney et al 2004). It is this abundance of data and the life or death relevance of safety communications techniques that has most likely driven the maturity of safety risk communications as a discipline. In contrast, information security can be difficult to measure which has probably contributed to its immaturity as a discipline. Security failures are not always recognised and some organisations have an interest in not disclosing information security incidents.

This paper discusses three key concepts which can be applied to information security awareness to improve effectiveness and efficiency:

- i. The psychological concept of bounded rationality which helps predict the limitations of human decision making
- ii. The mental models approach which is used by safety experts to conceptualise beliefs relevant to the human perception of risk

- iii. The Extended Parallel Processing Model which is used by safety experts to predict the likely outcomes of human behaviour when confronted with risk stimuli

2. The Rise of the Technocrats

To improve the effectiveness and efficiency of awareness methods it is necessary to identify the existing problems in security risk communications, many of which are caused by a technocratic approach to managing an audience. Historically, the information security function has usually been part of an information technology department where skills in communication or influencing users were not necessarily recognised as important. While the technocrats involved may have had an excellent understanding of the technical issues, they have had little concern for the existing beliefs, abilities or learning styles of their audiences. The technocrats assumed that problems with human behaviour in an information security context are caused by a lack of facts available to the audience and or adequate threats of consequences (Slovic 2000). Awareness therefore is largely treated as the broadcast of facts to an audience in the hope that behaviour improves. There is a tendency for technical experts in the field of information security to tell people what they think they ought to know (and may in fact already know). The problem is that technical specialists are venturing outside of their technical expertise when deciding what audiences will be told, how they will be told and how often they will be told. Technical expertise in understanding a risk does not automatically mean expertise in communicating that risk to others.

The “broadcast of facts” approach has been discredited by experts in safety risk communications:

“An effective communication must focus on the things people need to know but do not already. Rather than conduct a systematic analysis of what the public believes and what information they need to make the decisions they face, communicators typically ask technical experts what they think people should be told.” (Morgan et al 2002)

One of the origins of the “broadcast of facts” mentality can be found in the NIST SP800-50 model which defines Awareness, Training and Education:

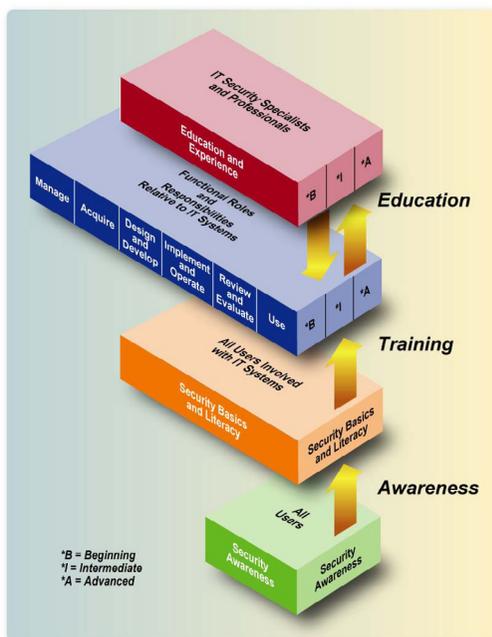


Figure 1: NIST Model for Awareness, Training and Education (NIST 2003)

The NIST model has been hugely influential in framing the perception of information security awareness activities. The NIST distinctions of awareness, training and education are pervasive through awareness literature (Herold 2005). The problem is that the NIST model focuses on technical competencies. For example the suggested needs analysis template provided in appendix A of NIST 800-50 reveals a focus on profiling the “rules of behaviour” that need to be transmitted to a given individual. Awareness needs are seen as a function of the duties and tasks performed. Other important factors are largely ignored. For example, there is no acknowledgement that people’s security awareness needs might also be based on their pre-existing beliefs about information security, some of which may not be correct or helpful. Similarly, there is no reference to other important considerations such as the culture and demographics of the organisation. Only focusing on the tasks that an audience needs to perform is a very narrow view of requirements.

The view that the contents of awareness communications is limited to increasing competencies for performing security tasks is also found in the relevant British and international information security standards:

“...All employees of the organization and, where relevant, contractors and third party users shall receive appropriate awareness training and regular updates in organizational policies and procedures, as relevant for their job function.” (BS/ISO Standard 27001:2005)

In contrast, the safety view of risk awareness acknowledges that all communications happen in a context of a person's existing perceptions about the topic and the choices that the individual perceives as available to them. The problem with the information security approach is that it is fact focused, not audience focused. This disconnection between the information security awareness approach and the needs of the audience has serious repercussions for information security governance. There are numerous examples of information security professionals blaming information system users for causing information security incidents. Words like "lazy", "stupid" or "ignorant" are often employed (Adams et al 1999) to supposedly "explain" human behaviour. Such labels ignore the limitations placed on users in terms of time and resources and actually reflect the ignorance of the observer in failing to understand the choices and opportunity costs perceived by the user.

Part of the problem appears to be a mistaken confidence in "best practice" approaches which have evolved in the information security profession with limited evidence to support their effectiveness or efficiency. For example, information security literature commonly refers to awareness activities being carried out on an annual basis. There is no evidence that such a frequency is optimal for an audience. In any case, since all audiences are unique and therefore have different requirements any arbitrary frequency would only be optimal for some audiences. The adoption of an annual program is probably more to do with a compliance target placed on the technocrats rather than the needs of the audience.

3. Bounded Rationality

There is significant debate within psychology literature as to the extent to which humans can be described as rational (Gross 2005). Rationality is described as the ability for individuals to select the "best" option when confronted with a set of choices. The best option is also referred to as a "value maximising" option when the most benefit is obtained for the least expenditure of resources or exposure to risk.

The problem is that people routinely fail to select a "value maximising" option in an information security context and exhibit apparently illogical behaviour. Commonly, an option mathematically modelled as the best choice by the technical experts isn't the choice chosen by information system users when responding to risk.

However, humans can be considered rational in so far as they attempt to make the best choice they can in a given situation. Research on passwords has demonstrated that behaviour which initially appears irrational such as the writing down of passwords becomes logical and therefore predictable once the constraints of the user are understood (Adams 1999).

Herbert Simon, an American psychologist proposed the concept of Bounded Rationality (Simon 1957) to explain why apparently logical people can make seemingly irrational decisions (as perceived by an independent, objective observer). It identifies a set of limitations on people's ability to make optimal decisions:

- a) That individuals make decisions based on their existing beliefs and attitudes. This is an important factor to help explain apparently illogical behaviour. Every individual approaching a decision will have a set of experiences in the form of beliefs and attitudes which frame their perception of the issue. Since all individuals have a unique set of experiences there will be an infinite range of beliefs and attitudes. Some will be helpful to information security risk management and others will be a hindrance.
- b) That individuals make decisions based on the limitations of their own actual and perceived cognitive ability. This helps explain why individuals may fail to attempt a task if they perceive it is beyond their ability to complete. This is particularly relevant when an individual is confronted with technological complexity or a poor user interface.
- c) That individuals make decisions based on time and resource constraints in consideration of other tasks and objectives. Experimental settings differ from “real life” decisions that individuals may face. Competing stimuli and time constraints in a real world scenario are likely to increase the likelihood of individuals relying on “rules of thumb” or heuristics when making decisions. While rules of thumb and heuristics are helpful and usually result in a reasonable outcome for the individual, they are unlikely to achieve an optimal one.
- d) That individuals learn to be content with a satisfactory outcome rather than an “optimal” one. “Satisficing” is a concept combining satisfactory and suffice which was proposed to explain why individuals were content with sub-optimal decisions. Simon demonstrated that in many situations, there was too much information that could realistically be processed by an individual in any meaningful way within acceptable time frames. The excessive investment of resources into optimising one decision would result in a reduction of resources available for other decisions. This constraint leads to a search for satisfactory solutions, rather than optimal ones.

Bounded rationality has important implications for information security. It provides an important reason why relying on best practice or topical subjects is unlikely to be effective or efficient for selecting awareness content. In any risk communication situation it is important to consider the limitations of the audience who are not looking for the perfect option as modelled by the experts, merely a satisfactory one given their other constraints. Without an understanding of these constraints in a given audience, any communications will be unreliable. Bounded rationality offers a way to recognise and predict the likely limitations of audiences so that the effectiveness of communications can be improved.

4. Mental Models

A problem identified in safety literature (Morgan et al 2002) is that technical experts approach risk communications with a different “mental model” of the risk than the

audience does. A mental model is a pattern of understanding held by an individual which provides context to their perception of rational choices. It consists of what beliefs they hold, the strength of those beliefs and the connections between beliefs. Safety experts note that when risk communication takes place the audience will have some degree of pre-existing knowledge which helps form their mental model:

“...for most risks, people have at least some relevant beliefs, which they will use in interpreting the communication. They may have heard some things about the risk in question. It may remind them of related phenomena.”
(Morgan et al 2002)

The safety approach to risk communication involves first understanding the mental models of the audience. By understanding the mental model of an audience it is possible to identify the specific beliefs which are causing the behaviour of concern. Communicating generic “facts” concerning a risk is unlikely to be effective or efficient. It may not even be feasible given the attention span of the audience. Using mental models, a communicator has the opportunity to identify what specific communications component would be most likely to influence the behaviour. Targeted communications components to influence a specific belief are much more likely to be effective and efficient. This could be for example a belief about the motivation of the threat actor, the value of the assets being protected, the likelihood of the threat or the likelihood of personal consequences.

The usefulness of the mental models approach has already been demonstrated in an information security context. Home users were surveyed to understand why they were vulnerable to bot-nets (Wash 2010). It was found that one of the mental models prevalent in the audience surveyed was a belief that the threat actors were mischievous, not malicious. This misunderstanding about the nature of the threat then helped explain the resultant behaviour of the audience in not taking sufficient preventative measures. To change the behaviour, the most effective and efficient way would be to target communications on the nature of the threat rather than generically reiterate the risks of internet security. With a general presentation of the facts, a communicator is less likely to communicate key information within the attention span of the audience. Instead, the focus should be on identifying and influencing key beliefs which act as a fulcrum for risk taking behaviour.

The difference in mental models between technical experts and their audiences are not only caused by differences in beliefs and their connections, but also by problems with terminology. False fluency is cited as an example (McIlwraith 2006) where key terms are misunderstood by audiences. The risk is that the choice of words used in awareness communications invokes the wrong mental model in the audience. If an audience has a fundamental misunderstanding about the meaning of a key word such as “virus” or “password” then unless reliable, agreed definitions are established then it is likely that any security awareness on the subject will fail to be effective or efficient.

Although the methods described here often refer to an individual it is recognised that it is not cost effective for most organisations to map the perceptions of all individuals. Instead, a marketing approach would be appropriate where audiences are sampled to identify common beliefs and the demographics to which they apply.

Focusing on specific elements of a mental model offers the opportunity for communications to be more effective and more efficient. Effective, in that key messages are communicated within the limited attention span of an audience and efficient in that the communication of key messages is less resource intensive than the general reiteration of facts as perceived by the technocrats.

5. Extended Parallel Processing Model

Safety researchers have noted that successful public information campaigns which influence risk taking behaviours tend to have been associated with the use of behavioural models to plan the communications approach (Delaney et al 2004). Numerous behavioural models exist such as the Theory of Reasoned Action, General Deterrence Theory and the Theory of Planned Behaviour. The models present different focuses on behaviour and allow for varying inputs and considerations for an individual making a decision. Information security normally attempts to prevent something negative from happening and usually relies on the threat of sanctions for compliance. Therefore the model most relevant is probably the Extended Parallel Processing Model because it focuses on the coping response of an individual when confronted with a perceived threat or risk.

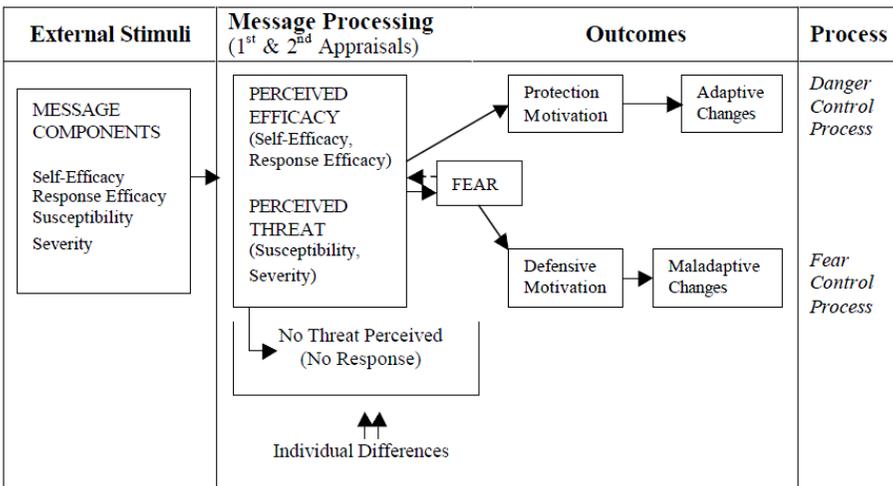


Figure 3: Extended Parallel Processing Model (Delaney 2004)

The Extended Parallel Processing Model identifies four key components in risk communications that shape how an individual responds.

- i. Self-Efficacy: an individual's perspective on their ability to perform a task competently. In an information security context this would include a user's perception of the complexity of a system versus their own perceived competency as a result of experience or training.
- ii. Response-Efficacy: an individual's perspective on the degree to which a choice of action has the ability to influence an outcome. In an information security context this is the degree to which the user has confidence that taking a proposed action will prevent a risk from occurring or reduce the potential damage if it does.
- iii. Susceptibility: an individual's judgement on how likely a risk is to impact them. In an information security context people's perceptions will also be shaped by the news stories that have been exposed to and the experiences of people they know.
- iv. Severity: an individual's judgement on the magnitude of the potential impact. Similar to susceptibility above, people's perceptions will be influenced by personal experiences.

Based on these inputs, the Extended Parallel Processing Model predicts three possible outcomes from risk communications. Safety experts note that if any one of the four key message components above fails then the likely consequence is that risk mitigation will fail with an individual either adopting a fear controlling response or ignoring the threat:

- i. Fear controlling response: the risk is perceived as significant but the individual perceives their ability to control the risk as low. A fear coping response often involves the use of a cognitive coping mechanisms to shield the individual from stress or worry about the possible consequences where a compensatory belief is adopted such as:
 - It will happen to me no matter what I do
 - I'm lucky so it won't happen to me
 - I'm an expert and I know what I'm doing so it won't happen to me

Fear controlling instead of risk controlling is a maladaptive response. It should be noted that a maladaptive response is a particular concern for information security where individuals are being challenged by technical complexity. For fear controlling responses it is likely that the individual's perception of Self-Efficacy or Response-Efficacy is the problem and a mental model should be used to map the individual's relevant perceptions of efficacy. For example, do subjects have an underlying belief (rightly or wrongly) that security software is complex and there is no point even trying to install or use it? This would explain a problem with self-efficacy. Or, do subjects have an underlying belief (rightly or wrongly) that anti virus doesn't work? This would contribute to a problem with response-efficacy.

- ii. No response: the risk is perceived as insignificant and is therefore ignored by the individual. If the risk was trivial then this is a successful outcome. If not, this is a negative outcome and a failure from the point of view of the risk communicators. For individuals inappropriately ignoring the risk it is likely that the perception of susceptibility or severity is the problem in which case mental models should be used to explore the perceived magnitude and likelihood of threats.
- iii. Risk controlling response: the risk is perceived as significant and the individual perceives their ability to control the risk is high. This is an adaptive response and a successful outcome from the point of view of the risk communicators.

If individuals are adopting a maladaptive response or inappropriately ignoring a risk then it is important to identify which of the four message components is contributing to the response. Traditionally when encountering non-compliant behaviour, technocrats would resort to reiterating general facts about the issue and increasing the threat of sanctions. An improved understanding of the beliefs supporting the perceptions of self-efficacy, response-efficacy, susceptibility and severity offers a far more efficient way of pinpointing specific problematic beliefs which are causing the undesired behaviour. It is this targeted approach to communications which is the main enabler for risk communications to be more effective and efficient.

6. Conclusions

Traditional information security awareness techniques expect that behaviour will improve by communicating the facts to a given audience and providing sufficient motivation to comply with instructions. Perspectives from safety risk communication have been presented to explain why the general presentation of facts is a narrow view of risk management that results in ineffective and inefficient communications. The tendency for technical experts in the field of information security to tell people what they think they ought to know (and may in fact already know) needs to be recognised as a failure.

A common theme from the safety approach is that to achieve effective and efficient communications it is critical to understand the relevant beliefs of the audience. It is not enough to know *what* behaviours exist that are causing information security risk. Communicators must understand *why* the behaviour is occurring which requires an understanding of an audience's constraints and supporting beliefs.

The use of conceptual frameworks such as bounded rationality, mental models and the Extended Parallel Processing Model offer an opportunity for a more formal and consistent approach to planning information security awareness. Bounded rationality helps explain why logical people do apparently illogical things. Mental models show the importance of existing beliefs and how they can be used to identify requirements for specific items of awareness content. The Extended Parallel Processing Model

shows how risk outcomes can be traced back to specific problems with one of four message components.

There is an urgent need to develop an improved process for conducting an information security awareness needs assessment. The audience's existing beliefs and constraints must be taken into account. For information security awareness techniques to improve in effectiveness and efficiency it is clear that information security awareness content can't be created from what the technical experts want to tell people, the contents of a technical standard or the prevailing best practice topics but rather the audience themselves that it seeks to influence and protect.

7. References

Adams, A. and Sasse, A.: Users Are Not the Enemy - Why users compromise computer security mechanisms and how to take remedial measures ACM 1999

Beautement, A Sasse, A Wonham, M: The Compliance Budget: Managing Security Behaviour in Organisations New Security Paradigms Workshop 2008

Delaney, A, Lough, B, Whelan, M and Cameron, M: A Review of Mass Media Campaigns in Road Safety Monash University Accident Research Centre 2004

Deloitte: Consumer Business Security Survey Deloitte 2009

Gross, R: Psychology: The Science of Mind and Behaviour Hodder Arnold 2005

Herley, C: So Long and No Thanks For The Externalities: The Rational Rejection of Security Advice by Users 2009

Herold, R: Managing an Information Security and Privacy Awareness and Training Program Auerbach Publications 2005

Lacey, D: Managing the Human Factor in Information Security John Wiley and Sons Ltd 2009

McIlwraith, A: Information Security and Employee Behaviour, How to Reduce Risk Through Employee Education, Training and Awareness. Gower Publishing 2006

Morgan, M. G., Fischhoff, B., Bostrom, A. and Atman, C.J.: Risk Communication: A Mental Models Approach. Cambridge University Press 2002

Parkin, S Moorsel, A Inglesant, P Sasse, A. A Stealth Approach to Usable Security: Helping IT Security Managers to Identify Workable Security Solutions. New Security Paradigms Workshop 2010

Simon, H. A. Models of man: Social and rational. New York: Wiley 1957

Slovic, P: The Perception of Risk Earthscan Publications 2000

Stewart, G: A Safety Approach to Information Security Communications Information Security Technical Report, Volume 14, Issue 4 2010

Wash, R: Folk Models of Home Computer Security 2010