



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Chris Garrett
July 23 2004
GSEC Practical
Assignment Version 1.4b
Option 1

Developing a Security-Awareness Culture – Improving Security Decision Making

© SANS Institute 2004, Author retains full rights.

Table of Contents

Abstract	2
Introduction	3
Cultural change.....	4
Individual Responsibility	5
Understanding the Decision-Making Process	7
What Decisions?	8
Some hypothetical decisions scenarios.	8
Judgmental Heuristics	11
The Availability Heuristic	11
The Representativeness Heuristic	13
Anchoring and Adjustment	13
Making Decisions in uncertain conditions	16
Improving Security Decision Making	18
References	21

Abstract

CIOs, managers and staff are faced with ever increasing levels of complexity in managing the security of their organizations and in preventing attacks that are increasingly sophisticated. As individuals we are subjected to enormous amounts of information across broad ranges of subjects, for example, security policies, new technologies, new patches, new threats, new sources of information, the list is endless. To fulfill the function of our role in the organization whether at a strategic or tactical level we make many decisions each day in the context of this information. As the environment continues to become more dynamic the process of making good security decisions is becoming more and more challenging. The answer lies in creating security-aware cultures in our organizations. This paper proposes that creating security aware cultures is dependent on improving how individuals make security decisions. Awareness of our decision-making processes as security practitioners can help us make better decisions in these uncertain conditions and help promote security-aware cultures in our organizations. Key to doing this is in understanding the process of how we really make decisions and what factors in the process may impair our abilities to make good security decisions for our organizations. This paper examines important facets of individual and group decision-making and provides prescriptive guidance on how we may improve the quality of our decision-making processes, leading to better security decisions.

Introduction

The Information Security Breaches survey 2004 conducted by PricewaterhouseCoopers for the Department of Industry (DTI) in the UK provides a useful starting point for this paper. The survey captures a snapshot of progress made by business and industry in dealing with information security threats and breaches particularly in the wake of 9/11. The survey reveals many interesting facts about the current state of security in the UK and four findings in particular are relevant to this paper:

- Information is regarded as the lifeblood of business and UK companies are increasingly reliant on the confidentiality, availability and integrity of their data
- Security is a high priority and firmly on the agenda of senior managers and executives
- Only a third of UK businesses have a security policy in place and only one in eight makes their staff aware of their security obligations
- It is widely accepted that the vast majority of security breaches are the result of a human error rather than technology flaws¹.

In the 2003 Global Security Survey of financial Institutions conducted by Deloitte Touche Tohmatsu, 80% of respondents reported that they have a formal information security strategy. However when asked if line and functional leaders led and embraced the strategy only 47% agreed that they do.² Ironically the financial services industry is one of the better industry sectors in addressing information security and many other industry sectors lag significantly behind.

Clearly, many organizations are cognizant of the importance of information and information assurance to the value of their business and organizations are moving to respond to the threat. Some by creating or elevating CIO positions to senior planning roles, removing the position from the IT department and integrating security planning and policy development into the strategic management process. Though the trend varies from industry-to-industry, security planning and policy development is clearly becoming an activity that is increasingly undertaken at the higher echelons of the organization either by, or in close consultation with senior management. Organizations also appear to understand that the vast majority of security breaches are the result of human error and the development of security policies and of security awareness training is seen as key to addressing the problems. However, there remain two key concerns. Firstly, and despite the increased emphasis placed on information and a greater awareness of the impact of security breaches, business is still unwilling

¹ DTI Information Security Breaches Survey 2004. Technical Report - PricewaterhouseCoopers

² 2003 Global Security Survey (Financial Services Industry) Deloitte Touche Tohmatsu

to invest to the degree necessary to improve security awareness among its employees. Secondly, policy development and security awareness classes are not enough. In this new age where there are numerous dangers to the critical data that drives our world, only a significant change in organizational culture can really reduce the number of security breaches experienced.


Cultural change

“So the challenge to many organizations is to create a security-aware culture. Making staff aware of the risks and their responsibilities helps them act in a sensible and secure manner”³

Understanding Organizational Culture and Cultural Change

It is useful at this point to define exactly what we mean by culture. The American Heritage English dictionary defines culture as “The totality of socially transmitted behavior patterns, arts, beliefs, institutions, and all other products of human work and thought.” Johnson and Scholes (1989)⁴ refer to the ‘Cultural Web,’ which supports the cultural paradigm of the organization.

Fig 1.



³ DTI Information Security Breaches Survey 2004. Technical Report - PricewaterhouseCoopers

⁴ Johnson J., Scholes K. Exploring Corporate Strategy. London, Prentice Hall (1989).

To shift the cultural paradigms of organizations into a mode where security becomes inherent requires changes in the structure of the cultural Web. Control systems, policies, organizational structures may all need to be adjusted to encourage cultural change, and there is some evidence that this is happening with the increases in the numbers of organizations that have implemented security policies and procedures. However, a critical additional factor in achieving real cultural change is in changing the behaviors of individual throughout the organization to support new policies, procedures and structures.

“The routine ways that members of the organization behave towards each other, and that link different parts of the organizations, comprise ‘the way we do things around here’, which at best lubricate the working of the organization and may provide a distinctive and beneficial organizational competency⁵ However, they can also represent a take-for-grantedness about how things should happen which is extremely difficult to change.”⁶

As the 2003 survey notes, even when the security related structures are changed, policies are developed and training given, that the behaviors of managers on a day to day basis reflect a current failure in integrating security into the cultural fabric of the organization. If the policies and practices are to be successful in preventing security breaches, they need to be adopted and practiced by every member of the organization on a daily basis. This is a significant cultural change in our organizational thinking and requires individuals at all levels in the organization to take the responsibility for making a decision that either safeguards, or exposes the organization to potential harm. The implication is obvious, unless we develop appropriate policies at a strategic level that are capable of shaping the future cultural fabric of the organization, and make available to each member of the organization the necessary information, training and guidance in making informed security decisions, we run the risk of exposing the organization to significant harm. We have to develop organizational cultures that inherently promote good security practices and behaviors. We have to evolve from the notion that if we bolt policies and practices onto the existing framework of the organization that cultural change will occur and information security will be addressed. In reality the policies and practices have to become an integrated part of the culture and need to be reflected in, and re-enforced by the behaviors of all members of the organization.

Individual Responsibility

So how do we encourage different behavior? As the research indicates, the vast majority of security breaches originate from human actions. There are a number of potential reasons for this:

⁵ RR Nelson and SG Winter, An Evolutionary Theory of Economic Change, Harvard University Press (1982)

⁶ Johnson J., Scholes K. Exploring Corporate Strategy. London, Prentice Hall (1989).

- People are poorly trained and have poor security awareness
- People are not motivated to perform at the required level
- People are malicious and deliberately expose the organization to risk
- People are aware of the problem of security but as managers and employees make poor decisions.

There is much in the literature about security awareness training and how to develop policies and training programs. Much of this work can be used to address the first three reasons that individuals may be responsible for security breaches. However, the implications of the research discussed so far alludes to the need to give individuals responsibility for their own behaviors and decision-making. The reality is that even if we develop policies, checklists, train people and weed out the insiders that are malicious, individuals are for a wide variety of reasons more than capable of making poor decisions that expose the organization to risk. This paper presumes that the structural aspects of the 'cultural Web' of an organization have been adjusted to address security threats and policies and practices are in place to some degree and that a level of security awareness exists. The paper will consider the underlying factor that is critical in developing security aware cultures in our organizations – Behavior and in particular how we make critical security-related decisions. If we are going to develop the kinds of security-aware cultures described we are going to rely on individuals to make good decisions. Unless we can get people thinking about the decisions they make we will always struggle to create a security aware culture in our organizations and will always be vulnerable to those who wish to do us harm. Therefore, we need to develop a greater understanding of what motivates our individual behaviors and what may impair our ability as individuals to make good security decisions. We have to create a culture of greater self-awareness in our employees as to the factors that can influence their judgment and lead them to make bad decisions. In this context this paper aims to do two things:

- Increase awareness by describing factors that may influence our security decision-making and judgment in specific situations.
- Prescribe ways in which we may use this awareness to improve our decision-making skills in securing our organizations.

The goal of the paper is to give individuals a greater sense of awareness about their own decision-making skills, knowledge that can be applied in their own organizations to improve judgment about security planning and security related decisions and help develop behaviors that support a security aware culture. It takes research from psychology and management science and how we can learn from these disciplines to develop security-aware managers and individuals who are capable of making good decisions at all levels in our organization. Decisions that may be highly complex, take weeks of contemplation and may influence security policies and infrastructure at the very highest level of the organization, to those decisions that we make many times each day, sometimes in a split second

(do I hold the door for the next guy, do I give a password over the phone etc.) and that occasionally have dire consequences. In order to do this, the paper focuses on research that shows how humans typically act in certain situations and how a number of factors can positively or adversely affect our perception of a certain situations. The paper will broaden self-awareness by:

1. Distinguishing between decision and outcome.
2. Examining how we actually make decisions, and how we deviate from the rational (ideal) model of decision-making.
3. Describing how we use heuristics (rules of thumb) to help us make decisions and how the biases associated with these heuristics can impact the quality of our security decision-making
4. Examine how we make decisions in uncertain conditions and in particular how we deal with risk

Each section will also suggest possible impacts on our decision-making as security practitioners by suggesting scenarios when our awareness of the decision-making process may help us make better decisions.

Understanding the Decision-Making Process

To help increase our self awareness and use good decision making processes it is important to firstly distinguish between what is an outcome and what is a good decision. The following provides a useful definition:

1. A good outcome is a future state of the world that we prize relative to other possibilities
2. A good decision is an action we take that is logically consistent with the alternatives we perceive, the information we have, and the preferences we feel
3. The quality of a decision should be judged by the knowledge and information available at the time the decision was made.⁷

We tend to judge the quality of a decision by its outcome, however this paper is focused on improving the decision-making process that is described in the second bullet rather than assessing individual outcomes.

Depending on our security role in the organization, these decisions may relate to many different things:

- Defining security policy and procedures
- IT Investment decisions
- Security staffing decisions

⁷ Reidar B. Bratvold, Would you Know a Good Decision if You Saw One – Psychological and Judgmental Aspects in Decision-Making –University of Adelaide, Australian School of Petroleum.
<http://www.spe-pb.org/attachments/articles/12/Bratvold%20-%20Mod%20for%20PDF.pdf>

- Data security issues
- Network Security
- Virus prevention
- Intrusion detection policy
- Help desk protocols
- Key card use and building access
- Back-up procedures and disaster recovery

What Decisions?

Some hypothetical decisions scenarios.

“There are some serious flaws in Internet Explorer's security model. Microsoft's latest IE patch removes one way of exploiting the flaw but it doesn't fix the root problem. Implementing an open source browser on all company machines would solve the security problem. What are the implications? What will I decide?”⁸

“I'm a CIO and money is tight. The IT budget is the same as last year and is to include provisions for system security. I have to update hardware and software, roll-out a wireless network, develop security policy and disaster recovery procedures, plan for security awareness training and pay for clearances for a number of my staff. What are my priorities?”

“One of my network admin staff has been acting a little out of character lately. He's been uncooperative with his colleagues, has been consistently late, taken a number of sick days. I know last year he was very unhappy at being passed over for promotion. I've let it slide a little because usually he's a good guy, but he has access to network passwords and a number of databases that contain sensitive data. Is it time to do something...what should I do?”

“I'm two days behind schedule for the release of a new Web system but the client has marketed the launch date and wants it launched on time. The site is dynamic and collects personal and credit card information. I haven't had time to check that all data inputs are validated correctly and that there is no threat from cross-site scripting⁹. Should I tell the client and ask for time, or launch the site and check the validations later?. I can always update the live site”

So how do we make decision about all of these important subjects?

“Strategic decision makers are rational actors who seek to maximize outcomes, first seeking all available information, then weighing up the various alternatives in order to select the best course of action”.¹⁰

⁸ US-CERT. Multiple Vulnerabilities in Microsoft Windows Components and Outlook Express. July 14 2004. <http://www.us-cert.gov/cas/techalerts/TA04-196A.html>

⁹ Allaire Security Bulletin (ASB00-05). Cross-Site Scripting Vulnerability Information for Allaire Customers. http://www.macromedia.com/devnet/security/security_zone/asb00-05.html

¹⁰ Hodgkinson G. P., Sparrow P. The Competent Organization: A Psychological Analysis of the Strategic Management Process (Managing Work and Organizations) London, McGraw Hill (2002)

So in other words we define the problem, identify and weight the criteria, generate and rate all of the alternatives by criterion and finally compute the optimal decision. In an ideal world yes, but in reality we rarely make decisions in this way. It is true that as decision makers for the most part we attempt to make rational decisions but there are many factors that limit our ability to truly take this approach and in our security planning and our security related activities it is extremely unlikely that we will use this all encompassing normative model that describes how we *should* make decisions as opposed to how we *actually* do make decisions. This is the first revelation in developing a sense of awareness of how we really make decisions that potentially safeguard our organizations. In reality there are good reasons why we rarely are able to use the normative approach. As early as 1975, Minzberg observed that managers engaged in a different activity every nine minutes and tended to avoid hard analytical, systematic data and relied heavily on intuitive judgment.¹¹ It is clear that time is a factor in limiting our ability to analyze all of the factors involved in a decision scenario. Quite often the quality and quantity of the data is insufficient, the problem poorly defined or the relevant criteria unclear. In today's world the business environment changes rapidly and the rate of change and complexity continually increases, further hampering our ability to make effective decisions based on the rational model. There are other factors also that limit our ability to use the rational model and In his Nobel prize winning work, Simon (1957)¹² and then March and Simon (1958)¹³ proposed that individual judgment is bounded in its rationality.¹⁴ The bounded rationality framework suggests that decision makers attempt to make rational decisions but are limited by:

- Missing or limited information
- Limited capacity to store and process information in usable memory
- Limitations on intelligence
- Their perceptions

These limitations prevent decision-makers from making optimal decision choices assumed in the rational model. Instead they choose options that appear to be satisfactory (satisficing) and meet a certain level of performance. In other words we analyze the information we have and use judgment in choosing an outcome based on our knowledge, perceptions and limited ability to store and process data. We also make decisions under the influence of factors from our external environment and in making security related decisions and the impact of

¹¹ Minzberg H. The nature of managerial work. New York Harper and Row (1975)

¹² Simon H. A., (1957) Models of Man: New York: Wiley. (1957)

¹³ March J. G., and Simon H.A., Organizations. New York. Wiley. (1958)

¹⁴ Bazerman M. Judgment in Managerial Decision Making. New York, Wiley (1988)

environmental factors is important and individual to each manager and employee. These may include:

- Financial and competitive issues and constraints
- Multiple reference points for security related information
- Complex interrelated security problems
- Continually changing technologies
- Continually evolving organizations – mergers buy-outs etc.
- Politics and legislation
- Time

Given all of the factors, in our security planning and policy development we are likely to choose options that satisfy and reach a certain level of performance. One classic example is the risk vs. IT investment equation where we try to balance the threat against the need to remain competitive while assuring the confidentiality, integrity and availability of information. The rational model would assume we analyze and weight every possible security threat and the means to prevent it and implement policies, processes and software accordingly. The problem in this scenario is that we are attempting to analyze a moving target as the threats and technologies evolve continually, so we have limited or missing information. The sheer complexity of analyzing, weighting and predicting the probability of every threat is realistically beyond the resources of the vast majority of organizations, so our ability to process and store such data is limited. So we tend to look for options that meet a more general level of acceptable performance.

Lets step back for a moment and consider the implications of bounded rationality in terms of security and information assurance and ponder the idea that in technology we often deal with absolutes, ones and zeros, definitions etc. Yet bounded rationality and the security environment dictate that we are going to have to make decisions that deviate from the rational model, decisions that are based on judgment. Even more disturbing, is the prospect of having to empower other people to make decisions based on their judgment. This is why we can develop all of the policy documents and checklists we like..... and build shelves for them to gather dust on. Unless we can get people thinking about the decisions they make we will always struggle to create a security aware cultures in our organizations and always be vulnerable to those who wish to do us harm. If we want to deal with the threats of social engineering we have to acknowledge that we are up against people who are clever and motivated. The people that exploit organizations through social engineering frequently exploit weaknesses in human judgment and decision making to gain access to organizational resources. It's not enough to teach people the policies, we have to get them aware of how their judgments may be flawed, get them engaged and out-thinking those that pose the threat. Therefore, it is vital that we understand how we use judgment in making rational decisions, and the factors that commonly impair judgment and lead to bad decisions.

Ready for some good news? Humans have developed cognitive tools called heuristics that help us deal with bounded rationality and they are generally effective in helping us make rational decisions. It is by understanding these tools and their potential pitfalls that we drill down to the next level of awareness of how we really use judgment.

Judgmental Heuristics

Knowing that we deviate from the rational model of decision making itself is interesting but fairly limited in improving our awareness of our decision-making processes and in turn the quality of our decisions. What is much more useful is to know how our decision making may be biased. Tversky and Kahneman (1974) continued what March and Simon had begun and developed an understanding of specific systematic biases that influence our judgment. Their research indicated that people rely on simplifying strategies which they termed heuristics or 'rules of thumb' and that these heuristics serve as mechanisms for coping with complex decision making scenarios. These tools have a direct effect on our judgment and understanding their associated biases is key in improving individual judgment and decision-making. We shall examine three key heuristics and some of their associated biases.

- The Availability Heuristic – Research has shown that managers tend to assess the frequency, probability or likely cause of an event by the degree to which instances or occurrences of that event are readily available in memory. (Tversky and Kahneman, 1973). An emotional, vivid or easily imagined event is more available to memory than a dull, vague or boring event.
- The Representativeness Heuristic – Managers also tend to assess the likelihood of an event's occurrence by the similarity of that occurrence to their stereotypes of similar occurrences.
- Anchoring and Adjustment – Managers make judgments by starting from a certain fixed point and adjusting their position in relation to the original point. The original value maybe suggested from historical, from the way a problem is presented, or from some arbitrary piece of information.¹⁵

The heuristics and their associated biases we are interested in terms of security decision-making are as follows.

The Availability Heuristic

The ease of recall bias - we have a tendency to judge the frequency of an event by the availability of its instances based on vividness and recency in memory rather than analysis of its actual frequency. Russo and Shoemaker (1989)

¹⁵ Bazerman M. Judgment in Managerial Decision Making. New York, Wiley (1988)

demonstrated that the availability of information in the media biases our perception of the frequency of an event¹⁶ This bias clearly has implications in the way we prioritize our security decisions. Are we overly influenced by the continual reporting and hype in the media about viruses and the cost to business generally? Or are we looking at all of the information we have regarding risk to our organization on an equally weighted footing regardless of recency and vividness? Consider the following scenario.

An organization devotes a good deal of time and IT security budget on ensuring that the organizations network is protected and that software patches and virus detection software are kept up to date. An update is made too slowly to the virus detection software and a virus is spread to a number of machines. The machines are cleaned up in a matter of a few hours and business returns to normal by the end of the day. Meanwhile, an employee working on the IT helpdesk, under pressure with the consequences of the virus outbreak gives a password to a caller on the phone who he assumes is a company employee. The caller is not an employee and gains access to a database of medical record information. Which of these is potentially more damaging to the business and which scenario did we pay most attention to during the development of security policies and risk assessment?

The truth is that the media tends to look at industry and business as a whole and the billions of dollars lost because of a virus or worm makes big news and is a vivid event. For example, it's estimated that the May 2000 "Love Letter" virus did an estimated \$2.6 billion in damages. At the level of the individual organization the reality maybe that the threat from viruses is relatively small and in applying defense-in-depth we must address issues that are seemingly much more mundane but that potentially may be much more destructive to our organization. In other words we have to balance the risk and the attention we pay to different issues based on the business model, infrastructure and culture of the organization and we have to base our decisions on data we know is objective and sound and not biased by descriptive influences.

Presumed associations – We tend to overestimate the likelihood of two events co-occurring. The implication of this bias in information assurance is the potential for us to overestimate the likelihood of two simultaneous events happening because in our cognitive processing of information we make associations that aren't born out by real data or actual facts, but by descriptive information or information that is more easily remembered. We also tend to ignore the basic statistical fact that if we are going to assess the probability of two associated factors, we need to analyze four different conditions.

In summary we need to ensure that our decision-making and judgment benefits from the availability heuristic and that we are not biased by erroneous factors that limit or misrepresent the true probability or characteristics of an event.

¹⁶ Russo J. E., and Shoemaker, P. J. H. Decision traps. New York: Doubleday (1989)

The Representativeness Heuristic

- Insensitivity to base rates – We tend to ignore base rates when descriptive information is available – even if it is irrelevant
- Insensitivity to sample size – we tend to ignore the importance of sample size in interpreting the significance of data.

The basic tenant in dealing with biases that emanate from the representative heuristic is, don't forget statistics 101! And more importantly don't be distracted by erroneous or descriptive information and apply the basics of statistics to solving problems or estimating risk. "Judgmental biases of this kind frequently occur when individuals cognitively ask the wrong question."¹⁷ It is interesting to consider this bias when we are making strategic security decisions or in assessing risks. This is typically done by some form of executive or strategic planning team that is comprised of individuals who have different interests, agendas and views on information assurance. Financial and performance factors may also play a part, as will personality and politics. As we know from experience this can often lead to conflict and disagreement. Facts can quickly become scarce and be replaced by opinion and descriptive or anecdotal evidence to support certain viewpoints. We need to begin such a process by identifying objective data that relates to the base of the problem we are dealing with. We also need to check that the data is weighted in its relevance by factors such as sample size etc. before the team considers descriptive or anecdotal information.

Anchoring and Adjustment

- Insufficient anchor adjustment
- Overconfidence

Now we come to potentially the most powerful biases in influencing our security judgments. Research has shown that we take information, form an estimate and develop an 'anchor' point (point of reference) based on this information. Kahnemann and Tversky (1974) provided empirical evidence to show the anchoring effect and how we typically don't adjust sufficiently away from that point even if the data the point is based on is irrelevant.

Things have been crazy on the help desk today and I just got a call from a woman who forgot her password. Didn't know her but she said she worked with a guy who I do know does work in that division so I gave her the password. I'm sure it'll be ok. Phone rings again....another forgotten password....what do I do?"

Arguably one of the most striking examples of anchoring and adjustment was the lack of preparedness before the terrorist attack of 9/11. We have had for many years as a nation and within the policy-making institutions of Washington DC a

¹⁷ Bazerman M. Judgment in Managerial Decision Making. New York, Wiley (1988)

collective anchor in terms of national security and of the threat to the continental US. This anchor was based on a large-scale nuclear attack by another superpower and despite the fact that there were a number of indications of a threat from other means, we were unable to adjust our anchor sufficiently from the status quo of the past 50 years to envisage that a small group of fanatical individuals with \$450,000 could kill 3000 people, cause millions of dollars of damage and negatively affect the economy for nearly two years. The major lesson for security practitioners is that we must continually question the validity of the anchor points we develop for their relevance. Take the help desk example. Will the anchor point of the individual, which is presumably the security policy for handling a forgotten password, now change if there are no negative consequences as a result of the initial decision?

Technology and the characteristics of the battle to keep information and organizations secure changes on a daily basis and we must always question our assumptions given the latest information. We must also be always aware that the anchor point of yesterday (policies, practices, threat, vulnerabilities) will overly influence and anchor these things today and we should always question whether we are overly biased in our judgment by the influences of the past and if those past reference points are still valid.

Overconfidence is a major concern in the information security world at the present time. The DTI survey of security breaches in the UK noted a worrying contradiction in the confidence of organizations in their ability to deal with security breaches and the actual number of security breaches taking place.

"The mismatch between the level of confidence organisations have and the number of incidents they are experiencing is worrying. There is no evidence to show that confidence is justified," said Andrew Beard, security consultant at PricewaterhouseCoopers.¹⁸

It is useful to look at research into the confidence bias to try and understand what is going on. Two significant research findings show that people tend to be most overconfident in situations that are of moderate to extreme difficulty (Lichtenstein, Fischhoff and Philips, 1982)¹⁹, and as peoples knowledge of a situation decreases they do not correspondingly reduce their level of confidence (Pitz, 1974)²⁰. These findings are extremely worrying in an environment where security practitioners are dealing with threats that are becoming more and more complex and we are making decisions under conditions of more and more uncertainty. In fact the DTI survey showed that respondents are increasingly more worried about the

¹⁸ Bill Goodwin, Businesses pay for over-confidence in firewall protection. Computer Weekly. <http://www.computerweekly.com/Article129375.htm>

¹⁹ Lichtenstein, S., Fischhoff, B., and Philips, L. D. (1982) Calibration of Probabilities. State of the art to 1980. In D Kahneman, P Slovic, and A. Tversky (Eds), Judgement under uncertain conditions: Heuristics and Biases. New York: Cambridge University Press.

²⁰ Pitz G. F. Subjective Probability distributions for imperfectly known quantities. In I.w. Gregg (Ed), Knowledge and Cognition, pp. 35-41. New York: Wiley (1974)

increased sophistication of threats against their computer systems.²¹ We clearly need to develop ways to improve how we make judgments on how confident we are that computer systems are secure. The section on improving decision-making later in this paper describes a number of methods for doing this.

Three more biases we should consider:

- The confirmation trap
- Irrational belief persistence

The confirmation trap basically means that we tend to look for evidence to validate a decision we are about to take before committing to it and that we tend to discount information that may challenge the decision. The confirmation trap was demonstrated by the following experiment. (Watson, 1960)²²

Consider the following set of numbers

2-4-6

This series conforms to some rule and the challenge was to discover the rule. Participants were allowed to propose any new sequence of numbers and were told whether or not the sequence conformed. What is the rule and how will you know when you have enough evidence to guess the rule? Most common guesses were; numbers that go up by 2; the difference between the first two numbers equals the difference between the last two numbers.

The actual answer is, three ascending numbers²³

Clearly, the solution requires that we accumulate disconfirming evidence and not confirming evidence. The implication is that to make better decisions, we need to attach equal weight to information that both supports and questions our tentative choices and overcome our natural tendency to look for and attach greater weight to evidence that supports the initial attractive choice.

Irrational belief persistence also describes the tendency that all of us have to look for evidence to positively support our viewpoints driven by our beliefs. These beliefs maybe religious, moral, social even related to a form of particular technology, but the effect of the irrational belief bias can be strong enough to overweight evidence consistent with that belief and prevent us from searching impartially for any evidence to question our beliefs. In technology and in

²¹ DTI Information Security Breaches Survey 2004. Technical Report - PricewaterhouseCoopers

²² Watson P. C. (1960) On the failure to eliminate hypothesis in a conceptual task. Quarterly Journal of Experimental Psychology 12, 129-140

²³ Reidar B. Bratvold, Would you Know a Good Decision if You Saw One – Psychological and Judgmental Aspects in Decision-Making –University of Adelaide, Australian School of Petroleum. <http://www.spe-pb.org/attachments/articles/12/Bratvold%20-%20Mod%20for%20PDF.pdf>

information security we all have our own particular beliefs. We should configure our network in a particular way, we should use a particular firewall, we should use Microsoft software, we should adopt Linux, we should move to open source, and our off-site back-ups should be located at location X etc. There are literally thousands of decisions we make to secure our networks. Has anybody actually stepped back and asked if our approach to securing our systems is based on the weighting of impartial evidence from reliable sources, configured to provide optimal protection to our systems? Any personal opinions, preferences, technological beliefs, or biases creeping in? It's worth at least thinking about.

Making Decisions in uncertain conditions

One of the most important areas in which we make security decisions is in determining risk. In an attempt to make this a more tangible, objective and scientific exercise, many risk assessment methodologies have been developed and many of them analyze risk based on variations of the following formula.

$$\text{Risk} = \text{Threat} \times \text{Vulnerability}$$

The clear limitation of this approach is that weights attached to threat and vulnerability needs to be figured out by individuals or groups and human decision-making is again involved. Add to this the major uncertainties of the security environment and it becomes a daunting task. One particularly relevant characteristic of human decision making that we should be aware of in terms of assessing risk in this situation is our tendency to react differently depending on how the scenario is framed, particularly in terms of perceived losses or perceived gains.²⁴ It is important to be aware of the phenomena as we conduct risk assessment so that we can check against its affects and make better risk assessments. Prospect Theory²⁵ attempts to describe how we make decisions under uncertain conditions like this. Like expected utility theory, prospect theory assumes that the value of an option or alternative is calculated as the summed products over specified outcomes. Each product consists of a utility and a weight attached to the objective and the probability of obtaining the outcome.

Prospect theory differs from expected utility theory in a number of important respects. Firstly, it handles the probabilities attached to particular outcomes differently. Prospect theory also treats the preferences as a function of "decision weights", and it assumes that these weights do not always correspond to probabilities. Specifically, prospect theory postulates that decision weights tend

²⁴ Stephanie Barclay McKeown. Framing Effects, UBC <http://epse501.freesevers.com/Framing-Effects.htm>

²⁵ Kahnemann D., and Tverskey, A. (1979). Prospect Theory: An analysis of decision under risk. *Econometrica* 47, 263-291

to overweight small probabilities and underweight moderate and high probabilities. Prospect theory also replaces the notion of "utility" with "value".²⁶

Tversky and Kahnemans' (1981)²⁷ Asian disease problem is one of the best examples of the effects of framing.

Problem 1

Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows.

Program A: If Program A is adopted, 200 people will be saved [72%].

Program B: If Program B is adopted, there is 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved [28%].

Which of the two programs would you favor?

Tversky and Kahneman (1981) found that the majority choice in this problem was risk averse: the prospect of saving 200 lives with certainty was more promising than the probability of a one-in-three chance of saving 600 lives. This risky prospect B was of equal expected value as the first prospect A.

A second group of respondents were given the same story of the Asian disease problem, but were provided with different program options.

Problem 2

Program C: If Program C is adopted 400 people will die [22%].

Program D: If Program D is adopted there is 1/3 probability that nobody will die, and 2/3 probability that 600 people will die [78%].

Which of the two programs would you favor?

The majority of respondents in the second problem chose risk taking: the certain death of 400 people is less acceptable than the two-in-three chance that 600 people will die.

This example shows how the framing of a problem in terms of gains and losses can affect the choices we make even though the value of the outcomes remain the same. Prospect theory demonstrates that we tend to be:

²⁶ Sussane Haberstroth, Prospect Theory, 1999. <http://mailhost.sfb504.uni-mannheim.de/glossary/prospect.htm>

²⁷ Tversky, A., and Kahneman, D. (1981). The framing of decisions and the psychology of choice. Science 211, 453-463

- Risk-seeking with low probability gains
- Risk-averse over low probability losses
- Risk-averse over high probability gains
- Risk seeking over high probability losses

We also need to appreciate that our experiences, perceptions and cognitive processes all affect the framing of a problem. Prospect theory and subsequent research has also shown that we have a tendency to overweight the probability of low probability events and underweight the probability of moderate and high probability events. This has serious implications in security planning and risk assessment and we must, through careful analysis ensure that we determine the real probability of an event regardless of how it is framed, through the data that is available and be resistant to using our instinct or relying on descriptive information.

Improving Security Decision Making

This paper is really about improving our self-awareness as security practitioners and key to doing this is in understanding the process of how we really make decisions and what factors in the process may impair our abilities to make good security decisions for our organizations. If we take a moment to step back from our decisions and consider some of the issues discussed will hopefully enable us to more aware of the context of the decisions and so improve our ability to prevent the effects of biases and framing. Equally as important as the understanding of what we really know as opposed to what we think we know, is an appreciation of the limits of our knowledge and what we don't know or have a poor understanding of. We also need to make sure we are asking the right questions. The following checklist may help improve some aspects of our decision-making.

Bazerman (1989) recommends the following questions in approaching uncertain situations. Some of these questions may be useful in determining how judgment is affected by the frame of a decision and may help in guiding security planning and conducting risk assessments:

1. How are you're decisions affected by the framing of choices?
2. How are you're decisions affected by the framing of outcomes?
3. How are you're decisions affected by the framed pseudocertainty and certainty of choices?
4. How do you differentially respond to the framing of "paying premiums" versus accepting sure losses?
5. How is your evaluation of the quality of a transaction affected by the frame in which it is presented?
6. How are you're decisions affected by summing gains and losses?
7. How does the frame of the problems affect how much you're time is worth?

8. How does ownership change your value of a commodity?
9. How rational are your inter-temporal choices?²⁸

Other Questions to consider in making security related decisions

1. Judgmental Biases –
 - Is my judgment biased?
 - I subject to any of the biases associated with judgmental heuristics.
 - Am I neglecting objective or relevant data in favor of descriptive or anecdotal information?
 - Am I ignoring the basics of statistics, and the statistical validity of the information (likelihood of conjunctive events, sample size)
 - Am I basing my judgments on objective facts or am I influenced by the recency or vividness of an event, particularly in the media
 - Am I influenced by an 'anchor point'. Is my anchor point valid and based on good data?
 - Am I moving too far or not adjusting enough from my anchor point.
 - Am I overestimating the likelihood of non-related events?
2. Overconfidence - Am I overconfident? It is striking how many times overconfidence or over optimism in both individuals and groups appears in the research described by this paper. Self-reflective questions to consider.
 - What is my confidence based on and is that data objective.
 - Have I considered all of the risks?
 - Have I truly identified and weighted the greatest risks and prioritized my decisions accordingly

Lichtenstein, Fischhoff and Philips (1982)²⁹ suggest two strategies for reducing overconfidence. Giving people feedback based on their judgments and asking people to explain why their answers may be wrong is useful in helping them see the contradictions in their judgment.

3. Overcoming the confirmation trap – We need to actively look for evidence that both promotes or challenges a proposed our proposed choice. It maybe that we have to bring in a third party who has no vested interest to objectively analyze the decision criteria and research.

- Am I finding evidence to support a belief while neglecting objective evidence that questions the validity of that belief?

²⁸ Bazerman M. Judgment in Managerial Decision Making. New York, Wiley (1988)

²⁹ Lichtenstein, S., Fischhoff, B., and Philips, L. D. (1982) Calibration of Probabilities. State of the art to 1980. In D Kahneman, P Slovic, and A. Tversky (Eds), Judgement under uncertain uncertainty: Heuristics and Biases. New York: Cambridge University Press.

- As I look at decisions in retrospective, am I selecting only data that supports the decision while neglecting data that cast doubt on its wisdom?

Group Think

Groups often undertake the process of defining security policy and taking strategic security decisions, so it is important to mention briefly a well-known phenomena that can occur in group decision-making that can be potentially disastrous so that it can be avoided. Group Think³⁰ was identified as the cause of the first space shuttle disaster in that the NASA mission team indulged in behaviors that led to critical errors in judgment. Group Think typically occurs when the group adopts the following behaviors:

- Illusion of Invulnerability: Members ignore obvious danger, take extreme risk, and are overly optimistic.
- Collective Rationalization: Members discredit and explain away warning contrary to group thinking.
- Illusion of Morality: Members believe their decisions are morally correct, ignoring the ethical consequences of their decisions.
- Excessive Stereotyping: The group constructs negative stereotypes of rivals outside the group.
- Pressure for Conformity: Members pressure any in the group who express arguments against the group's stereotypes, illusions, or commitments, viewing such opposition as disloyalty.
- Self-Censorship: Members withhold their dissenting views and counter-arguments.
- Illusion of Unanimity: Members perceive falsely that everyone agrees with the group's decision; silence is seen as consent.
- Mindguards: Some members appoint themselves to the role of protecting the group from adverse information that might threaten group complacency.³¹³²

Clearly some of the characteristics of groupthink have root is the individual biases described earlier. However the dynamics of the group help promote behaviors that reinforce collective thinking whether the thinking is likely to have a positive or negative outcome. Beware groupthink, and when engaged in a security decision-making process that requires group decision-making be aware of the common symptoms of groupthink and check that the group is not exhibiting

³⁰ Janis, I. L. & Mann, L. (1977). Decision making: A psychological analysis of conflict, choice, and commitment. New York: Free Press.

³¹ Janis, I. L. & Mann, L. (1977). Decision making: A psychological analysis of conflict, choice, and commitment. New York: Free Press.

³² Connie L. Fulmer, Group Think. *College of Education, Northern Illinois University Website* <http://www.cedu.niu.edu/~fulmer/groupthink.htm>

these types of behaviors. Awareness and monitoring is key to avoiding group-think.

References

- Bazerman M. Judgment in Managerial Decision Making. New York, Wiley (1988)
- Lichtenstein, S., Fischhoff, B., and Philips, L. D. (1982) Calibration of Probabilities. State of the art to 1980. In D Kahneman, P Slovic, and A. Tversky (Eds), Judgement under uncertain uncertainty: Heuristics and Biases. New York: Cambridge University Press.
- Janis, I. L. & Mann, L.. Decision making: A psychological analysis of conflict, choice, and commitment. New York: Free Press. (1977)
- Hodgkinson G. P., Sparrow P. The Competent Organization: A Psychological Analysis of the Strategic Management Process (Managing Work and Organizations) London, McGraw Hill (2002)
- Johnson J., Scholes K. Exploring Corporate Strategy. London, Prentice Hall (1989).
- Nelson R. R., Winter S.G. An Evolutionary Theory of Economic Change, Harvard University Press (1982)
- Minzberg H. The nature of managerial work. New York Harper and Row (1975)
- Simon H. A., Models of Man: New York: Wiley. (1957)
- Pitz G. F. (1974. Subjective Probability distributions for imperfectly known quantities. In I.w. Gregg (Ed), Knowledge and Cognition, pp. 35-41. New York: Wiley
- March J. G., and Simon H.A., Organizations. New York. Wiley. (1958)
- Russo J. E., and Shoemaker, P. J. H. Decision traps. New York: Doubleday (1989)
- National Security Agency. Defense in Depth – A practical Guide for achieving Information Assurance in today's highly networked environments. <http://nsa1.www.conxion.com/support/guides/sd-1.pdf>
- Reidar B. Bratvold, Would you Know a Good Decision if You Saw One – Psychological and Judgmental Aspects in Decision-Making –University of

- Adelaide, Australian School of Petroleum. <http://www.spe-pb.org/attachments/articles/12/Bratvold%20-%20Mod%20for%20PDF.pdf>
- US-CERT. Multiple Vulnerabilities in Microsoft Windows Components and Outlook Express. July 14 2004. <http://www.us-cert.gov/cas/techalerts/TA04-196A.html>
 - Allaire Security Bulletin (ASB00-05)
Cross-Site Scripting Vulnerability Information for Allaire Customers. http://www.macromedia.com/devnet/security/security_zone/asb00-05.html
 - Stephanie Barclay McKeown. Framing Effects, UBC <http://epse501.freeservers.com/Framing-Effects.htm>
 - Sussane Haberstroth, Prospect Theory, 1999. <http://mailhost.sfb504.uni-mannheim.de/glossary/prospect.htm>
 - Connie L. Fulmer, Group Think. *College of Education, Northern Illinois University Website* <http://www.cedu.niu.edu/~fulmer/groupthink.htm>
 - Bill Goodwin, Businesses pay for over-confidence in firewall protection. Computer Weekly. <http://www.computerweekly.com/Article129375.htm>
 - Watson P. C. (1960) On the failure to eliminate hypothesis in a conceptual task. Quarterly Journal of Experimental Psychology 12, 129-140
 - Tversky, A., and Kahneman, D. (1981). The framing of decisions and the psychology of choice. Science 211, 453-463
 - Kahnemann D., and Tverskey, A. (1979). Prospect Theory: An analysis of decision under risk. Econometrica 47, 263-291
 - DTI Information Security Breaches Survey 2004. Technical Report – PricewaterhouseCoopers
 - 2003 Global Security Survey (Financial Services Industry) Deloitte Touche Tohmatsu
 -