

A Call to Arms: It Is Time to Learn Like Experts

Join the Discussion
Connect

By Jay Jacobs – ISSA member, Minnesota, USA Chapter

This article is a call to action to break the reliance on unvalidated expert opinions by raising awareness of our decision environment and the development of context-specific feedback loops.

Abstract

Businesses and organizations rely almost exclusively on the opinions of security experts for direction and priority of risk-based resource allocation. This reliance on the subjective experience of experts begs the question, “Under what conditions are the opinions of security professionals worthy of trust?” This complicated question has a relatively simple answer: unvalidated expert judgment is not a reliable foundation for making risk-based information security decisions. The call to action is to break the reliance on opinions by raising awareness of our decision environment and the development of context-specific feedback loops.

“When we take action on the basis of an [untested] belief, we destroy the chance to discover whether that belief is appropriate.” – Robin M. Hogarth

Introduction

The report titled “Operation Black Tulip” from Fox-IT described the compromised servers for DigiNotar as being “physically very securely placed in a tempest-proof environment.”¹ For those not familiar with the technology, a tempest enclosure is intended to prevent eavesdropping of the electromagnetic frequencies emitted from electronic equipment. While nobody would argue about the effectiveness of tempest enclosures, many would argue about the importance of allocating those resources relative to other security controls. Yet, the architects at DigiNotar believed building a tempest-proof environment was more of a priority than designing an air-gap (off-network systems) into the design of their root certification authority. That decision contributed to a major breach of their systems. Unfortunately,

DigiNotar is not an isolated case of questionable risk-based prioritization of security resources. Given the lack of scientific studies and decision aids, untested beliefs, opinions, gut-feel, and emotions are the main components in information security decisions in practice.

There is a paradox here because relying on beliefs and intuition is often a successful approach to problem solving. Over the past 30 years or so, research in cognitive science and behavior economics has shed light on the human ability to leverage intuition and heuristics in decision-making, and why they are generally so successful in everyday circumstances: unconscious pattern recognition capabilities can help us see the order amidst complex patterns. In other words, intuitive judgment is improved by accumulating accurate patterns in our memory. “Intuitions based on a simple rule of thumb can be more accurate than complex calculations.”² “Skilled decision makers know that they can depend on their intuition, but at the same time they may feel uncomfortable trusting a source of power that seems so accidental.”³ We use heuristics and intuition so often in everyday life because it works most of the time and it saves time and effort. But what about the times when it does not work? Kahneman and Tversky, who demonstrated that our use of heuristics leads to deviations from economic rationality, stated, “People rely on a limited number of heuristic principles... In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors.”⁴ When can information security practitioners trust their intuition and when would that trust lead to severe and systematic errors? To answer those questions we

1 J. R. Prins, “Interim Report: DigiNotar Certificate Authority breach ‘Operation Black Tulip’” 9 5, 2011. <http://www.rijksoverheid.nl/ministeries/bzk/documenten-en-publicaties/rapporten/2011/09/05/diginotar-public-report-version-1.html> (accessed 9 10, 2011).

2 Gerd Gigerenzer, *Gut Feelings: The Intelligence of the Unconscious*. New York, NY: Penguin Group, 2007.

3 Gary Klein, *Sources of Power: How People Make Decisions*. Cambridge, MA: MIT Press, 1999.

4 Amos Tversky and Daniel Kahneman, *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press, 1982.

must first understand exactly what intuition is and how it is developed.

What is intuition?

Some people see intuition as a mystical thing, a gift that some people have and some people do not. Sigmund Freud referred to it as the unconscious or the primary process as opposed to the logic and reasoning of the secondary process. Seymour Epstein argues that this cognitive unconscious is an adaptive system that automatically and effortlessly learns through natural human experience. Current research defines intuition as simply “the recognition of patterns in memory,” which makes sense because as we experience things we store away the perceived lessons we learn to be recalled later. This leads to two very important implications: 1) intuition can be developed and improved since memory and experience are attainable, and 2) intuition can be developed wrong since perception, interpretation of outcomes, and memory are fallible. These are two very important points and set the stage for understanding the trustworthiness of intuition in risk-based decisions.

Many researchers talk about two different modes of decision making: first is “intuitive” (also called “experiential” or “tacit”), and second is “deliberate” (also called “rational” or “explicit”).⁵ In the intuitive mode, decisions are made with very little thought or effort and are generally made without any conscious awareness or attention. In the deliberate mode decisions are made with mental effort and over time, and often involve explicit analysis or logic. Most decisions in complex fields such as information security do not fall squarely into one mode or the other but instead fall somewhere between the two extremes. However, even a purely analytic approach is not free from the influence of intuition. “While tacit knowledge can be possessed by itself, explicit knowledge must rely on being tacitly understood and applied. Hence all knowledge is either tacit or rooted in tacit knowledge. A wholly explicit knowledge is unthinkable.”⁶

Multiple factors contribute to the development of intuition (creating the patterns in our memory), yet two factors are repeated in the literature as being the most critical: the ability of the environment to produce feedback and our opportunity to learn from that feedback.⁷ Situations or environments with clear feedback will improve both intuitive and deliberate processes, while having poor and misleading feedback will degrade both modes of decision-making. Robin Hogarth introduces an important concept to describe the relationship between the environment and our ability to learn:

*In general, you can think of learning structures as being favorable or unfavorable to accurate learning – that is, kind or wicked... The quality of your intuitions depends on the kinds of learning structures that prevailed when these were acquired. Kind learning structures lead to good intuitions; wicked ones do not.*⁸

Information security is a wicked environment

Hogarth describes feedback as the prime discriminator between a kind and wicked environment and consequently the quality of our intuition. A kind environment will offer unambiguous, timely, and accurate feedback. For example, most sports are a kind environment. When a tennis ball is struck, the feedback on performance is immediate and unambiguous. If the ball hits the net, it was aimed too low, etc. When the golf ball hooks off into the woods, the performance feedback to the golfer is obvious and immediate. However, if we focus on the feedback within information security decisions, we see feedback that is not timely, extremely ambiguous, and often misperceived or inaccurate. Years may pass between an information security decision and any evidence that the decision was poor. When information security does fail, proper attribution to the decision(s) is unlikely, and the correct lessons may not be learned, if lessons are learned at all. Because of this untimely, ambiguous, and inaccurate feedback, decision makers do not have the opportunity to learn from the environment in which the risk-based decisions are being made. It is safe to say that these decisions are being made in a wicked environment.

James Shanteau lists environmental characteristics associated with kind vs. wicked environments;⁹ comments for the decisions about risk-based prioritizations within information security are added (Table 1).

KIND ENVIRONMENT	WICKED ENVIRONMENT	INFORMATION SECURITY RISK-BASED DECISIONS
Decisions about things	Decisions about behavior	Adversary is adaptive and intelligent, behavior driven
Static stimuli	Dynamic stimuli	The speed of technology and complexity of systems make the stimuli very dynamic
Experts agree on stimuli	Experts disagree on stimuli	No agreement - based on the broad variation in risk analysis
Decision aids common	Decision aids rare	Decisions are based on opinion with aids
Objective analysis available	Subject analysis only	Subjective ordinal scoring techniques dominate the industry
Feedback is timely and accurate	Feedback is untimely and misleading	See text, feedback is nonexistent, ambiguous, or not timely

Table 1 – Characteristics of kind vs. wicked environments

5 Seymour Epstein, “Integration of the cognitive and the psychodynamic unconscious.” *American Psychologist*, no. 49 (1994): 709-724; Robin M. Hogarth, *Educating Intuition*. Chicago: University of Chicago Press, 2001; Michael Polanyi, *Knowing and Being: Essays by Michael Polanyi*. Chicago, University of Chicago Press, 1969.

6 Michael Polanyi, 1969.

7 Danial Kahneman and Gary Klein. “Conditions for Intuitive Expertise: A Failure to Disagree.” *American Psychologist*, no. 64 (2009): 515-526; James Shanteau, “Competence in experts: The role of task characteristics.” *Organizational Behavior and Human Decision Processes*, no. 53 (1992): 252-262; Robin M. Hogarth, *Educating Intuition*. Chicago: University of Chicago Press, 2001.

8 Robin M. Hogarth, 2001.

9 James Shanteau, 1992.

At this point, we can be completely confident in making the statement that risk-based decisions in information security exist in a wicked environment and developing a trustworthy intuition cannot be accomplished through mere exposure or tacit experience. In other words, Unvalidated expert judgment is not a reliable foundation for making risk-based information security decisions.

A mix of kind and wicked

Saying that expert judgment is not reliable across all types of information security decisions is not justified. Information security is a complex mixture of different environments depending on the work being performed. For example, security researchers who seek vulnerabilities (along with attackers and penetration testers) work in a kind environment. More often than not, they receive clear and timely feedback on their actions. They know when attacks succeed or fail and can adjust their activities based on the feedback. They have good opportunities to learn from their environment and may develop good intuition about *how* systems fail. However, researchers and penetration testers have no direct feedback about *when* attacks occur and will be successful, that is, they receive poor or no feedback about the overall risk. Researchers receive ambiguous feedback about the frequency of attacks and how probable one attack path may be over another. Receiving feedback on how a security system fails does not provide the opportunity to learn about risk-based prioritizing remediation efforts. The result of penetration testing produces a list of security weaknesses; any attempt at applying finite resources towards risk-based prioritizing remediation efforts must leverage alternative data or alternative feedback loops. Realizing what lessons can be learned from the feedback is as important as the feedback itself.

It is possible to shift a wicked environment towards a kind one.

A useful example is application software development. Developing secure code is currently done in a wicked environment. Developers rarely get timely and accurate feedback on their work. Most of the time they do not receive any feedback and overconfidence builds in developers. Some developers may go through their whole career blissfully unaware of their lacking skill to develop secure code. However, in looking at a possible new trend (as exemplified by Darren Meyer's talk for OWASP AppSec USA 2011¹⁰), building a feedback loop into the development cycle will break the tradition of poor and missing feedback. By putting static code analysis and other resources into the hands of developers – at the time decisions are being made (code is being written) – the feedback loop becomes timely and over time it may be more and more accurate as technology advances. Conversely, if that feedback is provided at the end of the development cycle (code review during the testing phase), the feedback moves towards untimely and will not provide the same opportunity for developers to learn. The good news here is there is already work

being done to shift secure coding from a wicked environment to a kind one through the introduction of timely feedback and decision aids.

A Call to Arms: Learn like experts

This is not to say that practitioner's opinions are always wrong, far from it. It simply shows that an unvalidated expert judgment is not a reliable method to make risk-based information security decisions. Currently, the guidance for practitioners comes in the form of best practices, which attempt to define how "experts" act. The unfortunate part of this is that expertise is not easy to establish when operating in a wicked environment. Mere exposure to a wicked environment is not guaranteed to generate expertise since opportunities to learn are not prevalent. When expertise is not easily identifiable (in a wicked environment) copying their unvalidated actions and thinking lead to epistemic fallacies. Gary Klein offers an alternative, "Instead of teaching people to *think* like experts, we can try to teach them to *learn* like experts. We can provide tools for helping people gain expertise on their own, without trying to predefine the nature of that expertise."¹¹ By taking a path of active learning we can begin a path towards improving the consistency and accuracy of our risk-based decisions. It all begins with two simple questions put forth by Hogarth: Why do I think that? and How would I know if I am wrong?

Why do I think that?

Insert breakpoints into the decision process and consider this first question in order to determine how information was obtained. By asking and answering this question we can identify when decisions are being based on feedback that is untimely, ambiguous, or misleading. It has been the author's experience that security practitioners who are used to operating unchallenged do not appreciate that question because they struggle to justify their intuitive guesses (making the question seem unfair). They typically do not have an answer or the answer is thinly veiled in fear, uncertainty, and doubt. Some answers may make an appeal to the authority of mislabeled experts by citing "best practice" for risk-based decisions. It is okay that we cannot cite metrics, research, or studies, but it is an indication of unverified feedback in our wicked environment. Either way, we should always ask the second question.

How would I know if I am wrong?

We should be wary of pursuing solutions that cannot answer this question because we will perpetuate a wicked environ-

Mere exposure to the environment is not guaranteed to generate expertise since opportunities to learn are not prevalent.

¹⁰ <http://www.appsecusa.org/> and <http://www.ustream.tv/recorded/17252001>.

¹¹ Gary Klein, "Developing Expertise in Decision Making." *Thinking and Reasoning* 3, no. 4 (1997): 337-352.

ment without any opportunity to learn from it. “Without the opportunities to learn, a valid intuition can only be due to a lucky accident or to magic – and we do not believe in magic.”¹² Seeking disconfirming evidence (rather than supporting evidence) forces an alternative perspective and forces thinking in terms of feedback loops. Perhaps this question cannot be answered now; perhaps the solution could be modified to provide useful metrics next month, next quarter, or even next year. It may be more cost-effective to study the effectiveness of security controls rather than putting in controls that are expensive or complex or generate inconvenience, especially if the control is intended to scale across an enterprise. Creating and improving on feedback loops may be just as important (if not more so in some cases) as the projects themselves, especially if both the cost and uncertainty (of control efficacy) are high. It is this building of feedback loops that we must be focusing on; otherwise our learning environment will remain wicked.

Sources of feedback

Plenty of sources exist for feedback. Multiple organizations in the industry are creating “data breach” or “state of” reports that communicate lessons learned from feedback at their disposal. While there is some variation in trustworthiness, they all offer a form of feedback and the opportunity to learn from their environment. Internally to organizations, every device, appliance, server, and application produces raw feedback in the form of logs and alerts; internal security events can provide a wealth of feedback if recorded, gathered, and analyzed.

There is a caveat to developing sources of feedback and that is they should be gathered in context. Most attempts at gathering feedback (a.k.a. metrics) are initiated for the purpose of gathering feedback, but the result is a collection of information with little applicability to any actionable context. Through asking and answering the questions presented here,

¹² Daniel Kahneman and Gary Klein, 2009.

seeking feedback gains context. For example, if the purpose of a remediation effort is to reduce the frequency and impact of a security breach on a specific platform, how many breaches or incidents have been recorded? Has a similar asset been targeted by attackers? Can the specific remediation efforts be identified as contributing to other internal breaches or in industry reports? These are just a few examples; the possibility for creating feedback exists both directly and indirectly. There is no wrong feedback, just more or less applicability to actionable context in this wicked environment.

Concluding comments

This article posited the question about risk-based decisions in information security, “Under what conditions are the opinions of security professionals worthy of trust?” The conclusion is that risk-based decisions in information security are made in a wicked environment. Simply accumulating experience does not justify trust in the opinions of security professionals where feedback is untimely, ambiguous, or misleading. The current approach to risk-based decisions warrants a change. It is time we learn like experts. The new approach begins with two simple questions, “Why do I think that?” and “How would I know if I am wrong?” By asking these two questions we can raise awareness of the problem and identify where feedback loops could be leveraged, improved, or created. It is only through this improvement in feedback that we’ll improve our opportunities to learn and break this self-perpetuating wicked environment.

About the Author

Jay Jacobs, CISSP, has worked in technology and information security for over 15 years. His areas of expertise include cryptography, risk management and mitigation, and security policy development. Jay currently works on key management and cryptography for a Fortune 500 company and serves on the board of directors for the Minnesota ISSA chapter. He can be reached at jay@beechplane.com.

