Security on Autopilot: Why Current Security Theories Hijack our Thinking and Lead Us Astray

Alan R. Dennis
Indiana University

Randall K. Minas
University of Hawaii at Manoa

Abstract

Most current information systems security theories assume a rational actor making deliberate decisions, yet recent research in psychology suggests that such deliberate thinking is not as common as we would expect. Much of human behavior is controlled by nonconscious automatic cognition (called System 1 cognition). The deliberate rational cognition of System 2 is triggered when System 1 detects something that is not normal; otherwise we often operate on autopilot. When we do engage System 2 cognition, it is influenced by the System 1 cognition that preceded it. In this paper we present an alternative theoretical approach to information security that is based on the nonconscious automatic cognition of System 1. In a System 1 world, cognition is a sub-second process of pattern-matching a stimulus to an existing person-context heuristic. These person-context heuristics are influenced by personality characteristics and a lifetime of experiences in the context. Thus System 1 theories are closely tied to individuals and the specific security context of interest. Methods to improve security compliance take on a very new form; the traditional approaches to security education and training that provide guidelines and ways to think about security have no effect when behavior is controlled by System 1, because System 1 cognition is instant pattern matching not deliberative. Thus in a System 1 world, we improve security by changing the heuristics used by System 1's pattern matching and/or by changing what System 1 sees as "normal" so that it triggers the deliberate cognition of System 2. In this article, we examine System 1 and System 2 cognition, while calling for increased research to develop theories of System 1 cognition in the cybersecurity literature.

Keywords: Information Security; Cybersecurity, Theory; Dual Process Cognition; System 1 Cognition; System 2 Cognition.

Introduction

Think about the last time you fell for a phishing email. Did you deliberately assess the situation (e.g., check for typos, check URL in the link) and conclude it was a valid email, or did you click on it without much deliberate thought?

Now think about the last time you received a badly crafted phishing email – one that was obviously fake – and later discovered that many people fell for it. Do you think they took time to examine the typos and consider an obviously fake URL, before they concluded that it was a valid email and clicked on the link?

Although technical solutions to security are important, it is becoming increasingly clear that employees, especially non-technical employees, are a key part of information security (Anderson & Agarwal, 2010; Bulgurcu, Cavusoglu, & Benbasat, 2010; Johnston & Warkentin, 2010; Liang & Xue, 2009, 2010; Warkentin, Johnston, & Shropshire, 2011). About 50 percent of all security incidents can be traced back to nonmalicious behavior of insiders not complying with organizational security policies (Ernst & Young, 2017; PricewaterhouseCoopers, 2017). Compliance has gained increased attention after a plethora of high-profile attacks (e.g., phishing), leading to multi-million dollar losses and damage to brand image (Kosner, 2014; Perlroth, 2014).

Much research has examined behavioral compliance (Johnston, Warkentin, McBride, & Carter, 2016; Johnston, Warkentin, & Siponen, 2015; Moody, Siponen, & Pahnila, 2018; Siponen & Vance, 2010), and many theories have been useful in explaining compliance behavior (e.g., Straub & Collins, 1990; Workman, Bommer, & Straub, 2008). Like most theories in information systems research, these theories quite naturally assume a rational actor thinking about and planning his or her behavior. But when a user falls for a phishing attack, was he or she actively thinking through the issues our theories suggest that rational actors consider before making a deliberate and considered decision?

Recent research in psychology would say no (Kahneman, 2011). This research concludes that humans have two distinct approaches to cognition: System 1, the nonconscious automatic cognition that is immediate and reactive, making decisions in less than a second, and System 2 that is the deliberate, thoughtful process we see in theories. Most researchers conclude that much human behavior is driven by System 1, with some believing that most behavior is controlled by System 1 (Kahneman, 2011; Stanovich & West, 2000). The exact amount of behavior controlled by System 1 is debatable and lies beyond the scope of this paper. However, the key point from this research is that most researchers studying dual process cognition agree that *a meaningful amount* of behavior is controlled by System 1 cognition.

System 1 cognition matters in the real world, yet most of our current information security theories assume System 2 cognition and do not consider System 1 cognition. Thus, our current theories are missing a key component to explain and predict behavior. This is not to suggest that our current theories are wrong; we believe they are useful in explaining behavior that is controlled by the deliberate and rational cognition of System 2. In the case of our hypothetical user above, it is possible that he or she engaged in System 2 cognition and was fooled by the badly crafted phishing email; on the other hand, it is also possible that System 1 ruled.

Our central thesis is that we need a different theoretical understanding to explain behavior that is controlled by the reactive and instant cognition of System 1. In a System 1 world, traditional security training is useless because System 1 does not consider what we have been taught (cf. Zhang, 2016). Instead, our behavior is heavily influenced by our past experience, the current context, and the specific stimuli we receive. Framing (Guo, Trueblood, & Diederich, 2017) and priming (Bargh & Chartrand, 2000) have very strong effects, so to an outside observer, our behavior might appear “irrational.” Yet in the words of Ariely (2009), we are “predictably irrational.” We argue that by understanding the fundamental nature of System 1 cognition, we can develop theories to predict this “irrational” behavior and produce better forms of training that are more effective in preventing security breaches that are due to nonmalicious behavior by insiders.

In this paper, we present an alternative theoretical approach to information security that is based on System 1 cognition. We begin by examining the nature of cognition, with a focus on System 1 cognition. We then turn to the topic of information security and use this research to partially develop a System 1-based theory to explain and predict how users respond to phishing emails as an example of how System 1 can be used to produce information security theories. We conclude by discussing the implications for research and practice that are raised by a System 1 view of information security.

The Nature of Cognition

Researchers have long argued that there are two fundamentally different forms of cognition. Many dual process models have been posed under a host of different names: see Evans (2008) for an analysis. In this paper, we adopt the commonly used terminology of Keith Stanovich (1999) and Daniel Kahneman (2011) who call these System 1 (automatic cognition) and System 2 (deliberate cognition). System 1 runs continuously, and delivers conclusions automatically and involuntarily; “it cannot be turned off” (Kahneman, 2011, p. 25). “A large part of what is commonly understood as ‘intuition’ … can be broadly conceived as the collection of automatic [System 1] processes” (Achtziger & Alós-Ferrer, 2013, p. 924).

System 2 runs much slower and most of the time adopts the conclusions of System 1 without thought (Kahneman, 2011). System 2 is activated when a surprise or error triggers us to focus our attention on a situation, but otherwise it typically idles and lets System 1 drive (Kahneman, 2011). “The arrangement works well most of the time because System 1 is generally very good at what it does; its models of familiar situations are accurate, its short-term predictions are usually accurate as well, and its initial reactions to challenges are swift and generally appropriate" (Kahneman, 2011, p. 25).

System 2 is familiar from prior information systems research: see de Guinea and Markus (2009) who call for a change to this. System 2 is the commonly assumed, deliberate, rational, effortful thinking that plays a central role in most information security research. Any theory that argues that deliberate attention is applied to some information and a conclusion is drawn based on a thoughtful process that balances the issues present in the situation has assumed that System 2 is operating. For examples, see Boss, Galletta, Lowry, Moody, and Polak (2015), Bulgurcu et al. (2010), D'Arcy, Hovav, and Galletta (2009), Johnston and Warkentin (2010), and Siponen and Vance (2010).

System 2 is well understood and well-represented in theories of information security. System 1 is not. Therefore, we focus this section on System 1 cognition, and its implications for information systems security theory. We want to foreshadow one major theme that runs through this paper: System 1 cognition is closely tied to a person-context relationship; that is, it is highly dependent upon the specific person (and the set of heuristics that guide his or her impression formation and behavior) and the very specific context which provide the stimuli that trigger these heuristics. Change the person or the context, and System 1 can produce different results. Nonetheless, we argue that there are predictable patterns to System 1 cognition, patterns that can be used to build theories of System 1 cognition. We will discuss this in more detail at the end of this section.

System 1 Cognition

System 1 exists so we can react faster than we can think (Kahneman, 2011). This may have developed as an evolutionary response; humans needed to be able to react to danger before the conscious mind could process stimuli, so System 1 cognition evolved to act faster than System 2 cognition (Kahneman, 2011). When a predator jumps out from behind a tree, we must run before we have time to deliberately assess the situation and plan our response. Neuroscience research shows that there is direct pathway from our sensory systems to the amygdala (which plays a key role in fear responses) in addition to the pathway that runs through the neocortex (which is responsible for the deeper cognition of System 2) and that this direct pathway is about twice as fast as the pathway through the neocortex (Loewenstein, O'Donoghue, & Bhatia, 2015).

When presented with a stimulus, our System 1 cognition automatically generates a response in less than one second (Bargh & Ferguson, 2000; Carlston & Skowronski, 1994; Fazio, Sanbonmatsu, Powell, & Kardes, 1986), often as quickly as 300 milliseconds (Todorov & Uleman, 2003). The most fundamental automatic cognition is to categorize a new stimulus as "good" or "bad," which triggers either an appetitive response (a.k.a. approach response) or an avoidance response (Chen & Bargh, 1999; Duckworth, Bargh, Garcia, & Chaiken, 2002; Lang, Bradley, & Cuthbert, 1990), although stimuli can also trigger other cognitions (e.g., the smell of apple pie could trigger hunger, or memories of childhood).

This process is nonconscious and unavoidable; we cannot prevent it (Kahneman, 2011). System 1 does not require deliberate "attention" (Evans & Stanovich, 2013). Instead, our System 1 runs continuously and supplies these assessments to our System 2, even though they are not asked for (Kahneman, 2011). For example, individuals will automatically categorize made-up words as warranting an appetitive or avoidance response even though there is no value in doing so (Duckworth et al., 2002). To borrow an illustration from Kahneman (2011),

> …look at the following words: *banana vomit*.
>
> A lot happened to you during the last second or two. Your face twisted slightly in an expression of disgust … and your sweat glands were activated. In short, you responded to the disgusting word with an attenuated version of how you would react to the actual event. All of this was completely automatic, beyond your control. (p. 50)

System 1 is not simply a faster version of System 2; it is fundamentally different (Bellini-Leite, 2013; Kahneman, 2011). System 1 is a pattern-matching system driven by the stimuli available in the current context (Fiske & Neuberg, 1990; Kahneman, 2011). A stimulus in the context automatically triggers an individual's heuristic(s) that matches the stimulus (hence person-context). The stimuli can be visual, auditory, or any of the other senses (touch, taste, or smell), although humans often give preference to sight and sound over the other senses (Kahneman, 2011). System 1 matches stimuli to the heuristics it has stored in memory, and these heuristics produce the matching response.

System 1 is a set of subsystems that simultaneously run in parallel triggered by the stimuli in the context (See Figure 1) (Bellini-Leite, 2013; Evans, 2008; Stanovich, 2004; Thompson, 2013). Because it has many subsystems operating in parallel, it has extremely high capacity to consider a wide range of

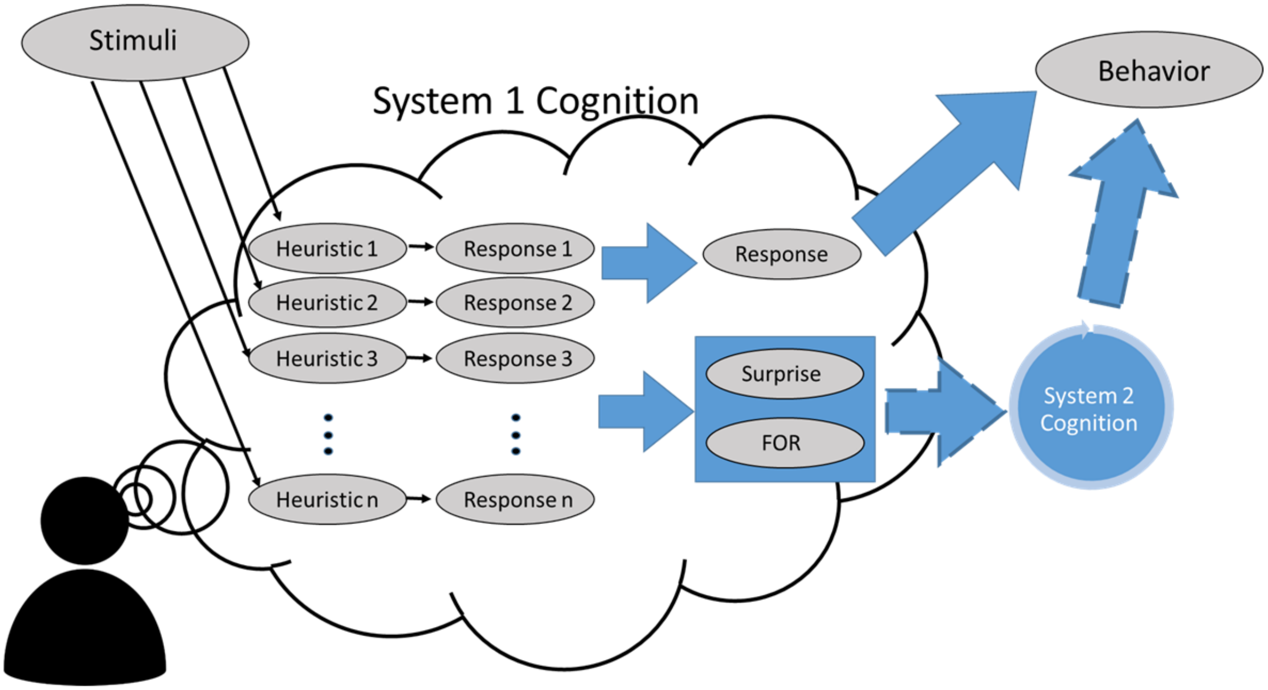

Figure 1. The parallel nature of System 1 and System 2 cognition

A stimulus triggers System 1 cognition, wherein heuristics are utilized to provide a response that spurs behavior or is classified as a "surprise" or triggers a low "FOR" that leads to System 2 processing of the information. System 2 cognition then controls subsequent behavior.

stimuli (Evans, 2008, 2014). In contrast, System 2 cognition is single-threaded and thus has much less processing capacity (Evans, 2014).

There are a variety of different stimuli present simultaneously in most contexts (e.g., sights, sounds, smells). Each of these can and do trigger different System 1 responses, as the different subsystems run in parallel (Bellini-Leite, 2013). There are likely different subsystems within the same sensory set. "Consider the situation in which one encounters a bizarrely dressed person in the mail room. A number of Type 1 processes are initiated, and their contents delivered to WM [working memory]: a feeling of surprise, an assessment of familiarity (is this someone I know?), an assessment of threat, and so on." (V.A. Thompson, 2013, p. 255). This is why a given stimulus can bring different thoughts to mind (e.g., the apple example above).

Evolutionary biology points to a set of parallel subsystems: there is some argument that System 1 is shared with animals (and thus may drive behavior seen as instinctive), while System 2 is limited to humans and some other cognitively advanced species (Bellini-Leite, 2013; Evans, 2014). Bacteria exhibit behavior but lack the brain needed for System 2 cognition, so behavior can come without System 2 cognition (van Dijk et al. 2008). "In other words, the brain is best viewed not as a commander or director of behavior, but rather as only one of the players among equally important others (i.e., the body and the world)." (van Dijk, Kerkhofs, Rooij, & Haselager, 2008, p. 298). System 1 may also be conceived of as a "traffic facilitator" among the different subsystems that propose behavioral responses to stimuli (Bellini-Leite, 2013; van Dijk et al., 2008).

The Nature of Heuristics

The System 1 responses to a stimulus are driven by the heuristics that that the stimulus automatically triggers (Kahneman, 2011). System 1's automatic processing uses these heuristics via category-based processing (Fiske & Neuberg, 1990). The stimulus is first matched to one or more of the individual's existing classifications and then the heuristic associated with that classification(s) is applied to form a response.

An individual's response to a stimulus depends on the general heuristics he or she has developed over time (Kahneman, 2011). People's personality, culture, habits, and experiences lead to different heuristics that have developed over a lifetime of experiences (Chen & Bargh, 1999; Maheswaran & Chaiken, 1991). Nonetheless, there is some commonality in the response to stimuli; many individuals respond similarly to the same stimuli (Chen & Bargh, 1999; Duckworth et al., 2002). Surprise, for example, triggers avoidance (Chen & Bargh, 1999). This suggests that some heuristics are instincts, shared by most humans (Evans, 2014). For example, how does a newborn baby know to cry whan it needs something? It cannot have learned that crying triggers adults to act, but must have been born with a heuristic for crying.

Other heuristics are learned (Evans, 2014; Kahneman, 2011). For example, consider this problem: At an event for circus people, you see someone who is funny; is the person more likely to be a clown or an acrobat? Most people would respond clown, because our System 1 has a heuristic that links clowns to funny, and no heuristic that links acrobats to funny. We are not born with heuristics about circus people, but learn them over time from our experiences. So, there are likely some common heuristics shared by individuals who have had similar experiences.

System 1 heuristics are individual and separate from each other. Thus System 1 cognition is characterized by an elimination by aspects process that considers one and only one problem attribute at a time that is used to produce an answer (Achtziger & Alós-Ferrer, 2013), rather than a more optimal multi-criteria decision making process. Each attribute recognized by System 1 could trigger a different heuristic focused on it (Bellini-Leite, 2013).

This is one reason why individuals often arrive at erroneous conclusions to simple reasoning problems that present two different types of information (Bago & De Neys, 2017). Consider this classic problem: At an event that has 5 clowns and 95 accountants, you see someone who is funny; is the person more likely to be a clown or an accountant? Many individuals will answer clown, although the Bayesian logic of our System 2 would suggest the correct answer is accountant, given the highly skewed base rate. This is because the problem as stated triggers two heuristics in most people: one about clowns being funny and one about accountants not being funny. System 1 triggers both heuristics in parallel and both produce an answer: clown and not-accountant. The System 1 traffic facilitator integrates these two answers, which are complementary, to produce the answer: clown. The base rate information is not linked to a heuristic, so our System 1 ignores it.

Of course, this is not to suggest that our System 1 heuristics are accurate. Not all accountants are not funny (we know several exceptions), but this is a widespread heuristic. Suppose for a moment that the problem had used auditors and IRS agents instead of clowns and accountants. Would our System 1 have produced an answer? For most people, the heuristics attached to tax accountants and IRS agents do not include the word "funny" so System 1 would have a hard time finding a heuristic to quickly produce a result. We might then have to use System 2 cognition to think about the base rate and draw a conclusion.

In any situation, there are usually many heuristics that could be triggered. System 1 prioritizes speed over accuracy (Kahneman, 2011), so the heuristics it uses are those likely to produce a fast answer. It is also likely that the most functional heuristic(s) in terms of the situation get priority (Fazio & Olson, 2003) – if I am surprised by something, a heuristic that produces a fight or flight response will get priority. Simple-to-use heuristics, such as those related to easily recognized classifications (e.g., gender, age, and race), tend to be activated first, rather than more complex classifications (S.T. Fiske, Lin, & Neuberg, 1999). For example (from Kahneman, 2011), consider whether the following statements are true: cows have three legs; chickens have four legs. Chances are, your answer to the first question was faster that the second. You have a heuristic that states no normal animal has three legs, so your System 1 could quickly match the statement to this heuristic and produce an answer. The second statement was more difficult because your System 1 knows there are some animals with four legs and some with two legs, and you had to sort out which group chickens fell into.

The affect-driven part of System 1 tends to place more priority on proximity (in time and/or space) or the vividness of stimuli, so that a snack placed in front of dieter is more likely to be accepted than a snack offered at a future time using words (Loewenstein et al., 2015). The vivid and proximate physical snack triggers a stronger System 1 response than the abstract nature of a future snack, so System 1 is more likely to drive behavior than the System 2's desire for weight control.

System 1 goes for the easy, most available, heuristics when it produces a result (Kahneman, 2011). In other words, the heuristics that are the most accessible to working memory are more likely to be triggered (Fiske & Neuberg, 1990; Srull & Wyer, 1979, 1980). This has two profound implications.

First, accessibility matters, which explains framing effects (Guo et al., 2017). It is well known that framing a problem in positive terms leads to different decisions than framing the same problem in negative terms (e.g., a problem that says there is a 75 percent chance of success versus one that says a 25 percent chance of failure leads to different decisions)

(Kahneman & Tversky, 1979; Tversky & Kahneman, 1981). While positive framing and negative framing appear identical to our System 2 cognition, they appear very different to our System 1 cognition. System 1 uses the heuristics triggered by the stimulus and a stimulus that contains the word “success” triggers different heuristics than a stimulus that contains the word “failure.” Findings in neuroscience indicate that the amygdala shows increased activation when subjects displayed framing effects (suggesting greater System 1 processing), while parts of the cortex showed increased activation when subjects avoided framing effects (suggesting System 2 cognition) (Benedetto, Kumaran, Seymour, & Dolan, 2006). Thus our System 1 cognition is particularly susceptible to the very specific nature of a stimulus (e.g., words that are synonyms may trigger different heuristics), unless our System 2 cognition steps into to override System 1. We will return to the role of System 2 in overriding System 1’s decisions in a little bit.

The second profound implication of availability is that the decisions produced by System 1 depend upon the contents of working memory (Fiske & Neuberg, 1990) – in other words, what we were thinking about immediately before the stimulus appeared – which explains priming. Research shows that we respond to and interpret the same stimulus in different ways depending upon the context and what occurred immediately before (Fazio & Olson, 2003). Priming is an intervention intended to activate desired concepts in working memory with the goal of changing subsequent interpretations and behavior (Bargh & Chartrand, 2000). The scope of priming research is expansive, with researchers using it to influence a wide variety of behaviors such as purchase intentions (Yi, 1990), walking speed (Bargh, Chen, & Burrows, 1996), rude behavior (Shariff & Norenzayan, 2007), snack food consumption (Minas, Poor, Dennis, & Bartelt, 2016), and the brainstorming productivity of virtual teams (Dennis, Minas, & Bhagwatwar, 2013). In other words, the decisions produced by System 1 are unstable because the heuristics that a stimulus triggers depend upon what occurred immediately prior to the stimulus, whether due to intentional or unintentional priming effects; give a stimulus to System 1 and it might produce different behavior today than the same stimulus produced yesterday. This one reason why humans often do not present a stable set of well-ordered preferences (Bellini-Leite, 2013).

Recall the earlier “banana vomit” example, and now consider how you might react if someone offered you banana bread. Chances are, you have a slight aversion to bananas right now because your mind assumed a temporal sequence and causal connection between the two words in working memory forming a sketchy scenario in which eating bananas caused the sickness, even though there was no reason to do so (Kahneman, 2011). Your System 1 reaction to being offered banana bread would likely be negative because we primed this aversion. If we ask about banana bread later in this paper after the two-word combination has faded from working memory, your System 1 reaction will likely be different. Information that is in working memory – even random information – has a strong effect on the results produced by System 1, and information stored in long-term memory not presently in working memory might as well not exist (Kahneman, 2011). Priming has a very uneven history of working, likely because it has strong effects on System 1 and weaker effects on System 2 (Stafford, 1996).

The Role of System 1 in System 2 Cognition

Individuals can deliberately choose to invoke System 2 cognition at any time, but they often rely on the automatic processing of System 1 when it produces a result and they do not perceive anything unusual in the environment to suggest the result is not appropriate (John A Bargh & Chartrand, 1999; John A Bargh & Ferguson, 2000; Schwarz & Clore, 2007; Smith, Cacioppo, Larsen, & Chartrand, 2003) So why do we not engage in System 2 thinking? Do you remember learning to drive and how difficult it was because you had to pay attention to and think about so many different things? Once you gained enough experience, driving became simpler, because you could rely on your System 1 to manage much of it and not need to expend as much System 2 cognition. Today, you probably use your System 2 to think about what awaits you at work as you drive to the office and for the most part, leave the driving to System 1, something you could never do as you were learning to drive.

System 2 cognition is effortful, and most humans are “cognitive misers” in that they attempt to minimize the effort needed to meet their goals (Taylor & Fiske, 1978). We do not have the capacity to deliberately think about everything, so if we are not motivated, we do not exert the effort required for the deliberate assessment of System 2 (de Guinea & Markus, 2009; Gersick & Hackman, 1990; Jasperson, Carter, & Zmud, 2005; Louis & Sutton, 1991). We have to deliberately decide that some issue is important enough to engage System 2 cognition, or our System 1 has to signal that something is important enough to warrant investing System 2 cognition, before we engage it. So what triggers System 1 to signal us to engage System 2? There are two commonly accepted types of triggers.

One trigger for invoking System 2 is a negative stimulus (Kahneman, 2011) or a discrepancy from normal expectations (Gersick & Hackman, 1990; Louis & Sutton, 1991). Negative stimuli draw attention (Smith et al., 2003) and are likely to trigger the conscious processing of System 2 because humans are loss averse (Schwarz & Clore, 2007). Discrepancies from what our System 1 considers normal usually need to be significant to trigger a departure from System 1's automatic process (Kahneman, 2011; Louis & Sutton, 1991), because once established, the invocation of the System 1 routine drops from awareness (Gersick & Hackman, 1990). Nonetheless, for some individuals, even a minor discrepancy is sufficient to trigger the use of System 2 (Gersick & Hackman, 1990; Louis & Sutton, 1991).

What is "normal" may differ between novices and experts, because experts have a greater pool of past experiences to draw upon (Evans, 2008; Kahneman, 2011). So a critical abnormal stimulus may be overlooked by a novice and recognized by an expert; for example, Klein (1999) describes a situation when an expert firefighter ordered a team to evacuate a building immediately before it collapsed because his System 1 detected an abnormality. Conversely, novices may perceive something as abnormal when an expert does not, because they lack experience; one of the authors was sent to an expert ophthalmologist by a novice optometrist because he misidentified an unusual pigmentation as the start of detached retina.

The other trigger is conflicting System 1 results. The multiple System 1 subsystems work in concert and lead to behavior when System 1's traffic facilitator resolves the results they produce. However, sometimes the System 1 subsystems will produce results that are in conflict – the results from the different subsystems do not agree. System 1 also produces a "Feeling of Rightness" (FOR) that is a measure of this conflict (De Neys, 2014; Thompson, Prowse Turner, & Pennycook, 2011). When the System 1 traffic facilitator integrates results, it produces this FOR as a byproduct. When the System 1 subsystems produce complementary results, FOR suggests that all is well (Bago & De Neys, 2017). When there is conflict among System 1 results, the FOR creates a sense that something is not right (Bago & De Neys, 2017). For example, consider this classic problem: Suppose a bat and ball together costs $1.10 and the bat costs $1.00 more than the ball; how much does the ball cost? Most people's System 1 immediately produces the answer of $.10. But did you have the sense that something was not right? That was your FOR trying to signal that there was conflict among your System 1 subsystems, because $.10 is the wrong answer.

Like other System 1 processes, the processing of FOR is automatic and does not depend on System 2 (De Neys, 2014; Johnson, Tubau, & De Neys, 2016). Empirical evidence shows that the longer it takes to produce a System 1 answer, the lower the FOR (Thompson et al., 2011). It is unclear whether the longer time is taken by the System 1 subsystems to produce their various results, or whether the longer time is taken by the System 1 traffic facilitator in resolving the results of the different subsystems; theory would suggest the longer time is due to the traffic facilitator.

Thus a low FOR is a second common trigger for System 2 cognition (De Neys, 2014; Thompson & Morsanyi, 2012; Thompson et al., 2011). Different individuals may choose to respond to FOR in different ways. Some individuals may have a high need for cognition and therefore engage System 2 in situations where the FOR is only slightly low, while other individuals may be "miserly" or reluctant to engage in System 2 cognition even when the FOR is very low (Evans & Stanovich, 2013; Johnson et al., 2016). Empirical evidence suggests that when System 1 produces an erroneous response, most people have a low FOR and whether System 2 is triggered or not depends on the person's personality (Johnson et al., 2016); thus it is not a differential ability to produce reasonable FOR values but rather how one chooses to respond to the FOR that drives whether System 2 will be invoked or not (Johnson et al., 2016). "Put differently, people do not fail to detect that they need to think harder, they fail to complete the effortful, hard thinking." (Johnson et al., 2016, p. 61).

If System 2 is triggered, then the individual engages in deliberate cognition about the situation. System 2 can override the automatic response of System 1, and in fact it often does. Suppose you are on a diet to lose weight and are feeling hungry when you see someone eating a snack; your System 1 likely generates an immediate appetitive response. Whether you choose to actually eat a snack likely depends on what your System 2 decides (Loewenstein et al., 2015). If you are like us, it probably depends upon how much willpower you have to overcome the System 1 desire (cf. Loewenstein, et al. 2015). Thus behavior can be a product of System 1 and System 2 working together – in concert or in conflict.

We have treated System 1 cognition as separate and distinct from System 2 cognition, which it is (Kahneman, 2011). However, System 2 cognition is influenced, sometimes very strongly, by the results of System 1 (Kahneman, 2011). There are three distinct theoretical routes by which System 2 cognition may be influenced by the fast and often inaccurate assessment of System 1 (remember that System 2 is often triggered by a low FOR when System 1 produces a result it suspects is not right).

The first theoretical route is confirmation bias (Nickerson, 1998; Tsohou, Karyda, & Kokolakis, 2015). We tend to focus on information that supports our initial beliefs and discount information that opposes them (Lord, Scott, Pugh, & Desforges, 1997; Wood, Quinn, & Kashy, 2002). If we choose to override our System 1 response and devote cognitive effort to deliberately applying System 2 cognition, the initial System 1 conclusion colors our System 2 cognition (Tsohou et al., 2015). We see the available information, but tend to discount information that does not support the initial – and potentially inaccurate – System 1 result.

The second theoretical route is by the judgments that System 1 automatically attaches to the information (Loewenstein et al., 2015). In many cases, the information in a situation is ambiguous and could be interpreted in different ways (Srull & Wyer, 1979). For example, suppose someone promises to email you a report on Thursday and it arrives on Thursday at 11:59pm; is this person "reliable"? You could conclude yes (because the deadline was met) or maybe not (because it just barely made it). Yet our System 1 performs this judgment automatically in less than a second, and is influenced by whatever heuristics are most accessible in working memory at the time (Srull & Wyer, 1979, 1980), meaning your judgment is influenced by the exact words used (framing effects: Guo et al. (2017)) and your judgment today might be different than your judgment tomorrow (priming effects: Srull and Wyer (1980)). This judgment influences how subsequent information is processed by System 2.

The third theoretical route is by the way System 1 and working memory treats the System 1 results. The results created by System 1 are stored in working memory alongside the "facts" that produced them (Srull & Wyer, 1980). "The outcome of autonomous [System 1] processes automatically become part of the representation of the problem space." (Thompson, 2013, p. 254). System 2 then has equal access to the "facts" and System 1's unreliable result, and can and does use both in its deliberations. In the brain, working memory works closely with the anterior cingulate cortex to provide quick error detection that affects the assessment of the "facts" and influence one's FOR. There is some evidence that System 2 gives more weight to System 1's result than to the "facts" that produced it, and that this weight increases over time (Srull & Wyer, 1980, 1983). As time elapses from the initial System 1 result, we have to decide how much information about the situation we transfer from short-term memory into long-term memory. The System 1 summary result is easier to remember that the host of "facts" that produced it, so we tend to remember the System 1 result and forget the underlying "facts" (Srull & Wyer, 1980, 1983). Thus an erroneous System 1 result is more likely to linger than the "facts" and thus have greater influence on our subsequent System 2 cognitions and behavior.

The Formation of System 1 Heuristics

To this point we have focused on how System 1 operates using heuristics. One important question remaining is how are an individual's heuristics formed? Some heuristics are instinct and some are learned based on experience (Bellini-Leite, 2013; Evans, 2014; Kahneman, 2011). Learned heuristics are based on past experience. We repeat behaviors that have been successful in the past and avoid those that have not (Achtziger & Alós-Ferrer, 2013; Evans, 2014). The essence of learned heuristics is reinforcement learning (Holroyd & Coles, 2002; Sutton & Barto, 1998). Given a certain situation, we perform a behavior and assess the result. This assessment is often heavily effect-laden (Holroyd & Coles, 2002) and heavily anchored in outcomes, even when the outcomes have a large random component (Achtziger & Alós-Ferrer, 2013).

Nonetheless, it is this series of learning by doing that creates our System 1 heuristics (Achtziger & Alós-Ferrer, 2013; Sutton & Barto, 1998). When we are successful, we form the integrated situation-response in a heuristic to guide future behavior (often an appetitive or avoidance reaction, although stimuli can trigger emotions and memories). Learning is facilitated when the behavior and resulting assessment occur close in time (Sutton & Barto, 1998), and when there is repetition that reinforces the newly emerging heuristic (Achtziger & Alós-Ferrer, 2013; Kahneman, 2011; Sutton & Barto, 1998). Although it may take many repetitions to build – and especially change – heuristics (Fazio & Olson, 2003).

However, the experiential memory used to create heuristics is unusual. Kahneman (2011) describes an experiment in which participants received two treatments (60 seconds of moderate pain, and 60 seconds of the same moderate pain followed by 30 seconds of minor pain) and then were given a choice of which they would prefer for a third treatment. To an outside observer, logic says participants should choose the first treatment because it has less pain, yet 80 percent of participants choose the longer duration, preferring to endure 30 seconds more pain than they needed to. A dispassionate outside observer uses System 2's logic and reasoning and focuses on the total pain experienced (i.e., the area under the curve). In contrast, a participant uses System 1's memory of individual slices of representative time periods, heavily biased by the last time period experienced.

Thus one important take-away is that the System 1 heuristics are experiential patterns (Kahneman, 2011). They are not learned from training or the advice from others. Such training and advice augments the knowledge used by System 2, but does not influence the heuristics used by System 1. Realistic training (e.g., flight simulators), may lead to the development of System 1 heuristics because this type of training is close to the actual experience. Perhaps the best analogy is that expertise reflects System 2 knowledge, while skills reflect System 1 heuristics (Kahneman, 2011).

System 1's "nonconscious behavioral guidance Systems" (Bargh & Morsella, 2008) are the building blocks for habitual behavior using well-learned patterns (Bargh & Chartrand, 1999; Triandis, 1971). When an individual repeatedly encounters a situation and forms the same attitude or intention in response, that situation and attitude/intention become matched in memory. When that situation is encountered again, the matching attitude/intention is formed automatically (Bargh & Chartrand, 1999; Bargh & Ferguson, 2000). System 1 attitude formation does not require conscious control (Wood et al., 2002). One might even say that "the environment selects appropriate behaviors from the behavioral repertoire, without any internally computed behavioral plans." (van Dijk, et al., 2008, p. 307).

In a stable and predictable working environment where nothing unusual is perceived, individuals are not motivated to think (Schwarz & Clore, 2007; Smith et al., 2003), and System 1 runs efficiently with System 2 "in a comfortable low-effort mode, in which only a fraction of its capacity is engaged" (Kahneman, 2011, p. 24). System 1 generates responses, which System 2 adopts with little thought (Kahneman, 2011).

This offers one explanation for the difference between novice and expert problem solving. It is well known that novices tend to consider more information than experts and spend much more time and effort to explore it (Kahneman, 2011). In contrast, experts quickly focus on a small part of the available information and produce fast answers (Kahneman 2011, Evans 2008). For example, Klein's (1999) studies of expert decision making by fire fighters and paramedics reveals very little of the rational decision making we associate with System 2. Instead, it is replete with examples in which the expert matches the situation to one encountered previously and rapidly retrieved a solution schema – the processes we associate with System 1 cognition. Kahneman (2011, pp. 11-12) argues that we should be no more surprised by this fast expert judgment in complex situations than we are when a 2-yearold sees a dog and exclaims "doggie!"; the processes are the same System 1 at work, just using different heuristics.

Summary

Figure 1 summarizes our arguments about System 1 and System 2 cognition. We argue that the stimuli in the context individually trigger a set of heuristics that that are simultaneously processed by our numerous System 1 subsystems in parallel. These heuristics produce responses that are integrated by our System 1 traffic facilitator into a response (or a set of complementary responses). System 1 also produces a Feeling of Rightness (FOR) and may issue an alert if it detects surprise or an anomaly. We begin acting on the System 1 response(s) while the response(s) and accompanying FOR and surprise alert is passed along to System 2, which decides whether to continue acting on the System 1 response(s) or to override it. In most normal situations, the FOR is high and we detect no anomalies, so System 2 passively lets System 1 guide our behavior. In other situations when FOR is low, when we detect an anomaly, or when we are motivated to think about the situation, we may engage System 2 cognition, which may let the System 1 behavior continue or step in and override it (See Figure 1).

Implications for Security Research

As we examine past security research, it becomes immediately clear that past theorizing and empirical research is dominated by System 2 thinking (Hui, Vance, & Zhdanov, 2016). We see little discussion of pattern matching, categorization, and heuristics that are the hallmarks of System 1 thinking. Instead, we see important factors (e.g., response efficacy, response cost) balanced against each other in a struggle for influence -- which is the hallmark of System 2 cognition. The few security theories that do consider System 1 cognition typically approach it as bias, and are concerned with how to mitigate the effects of such biased heuristics on the subsequent System 2 cognition (Tsohou et al., 2015). Such theories are good first steps, yet they still focus on System 2 cognition as the driver of behavior.

In hindsight, the dominance of System 2 thinking in our theories should have been obvious. "When we think of ourselves, we identify with System 2, the conscious reasoning self that has beliefs, makes choices, and decides what to do. Although System 2 believes itself to be where the action is, the automatic System 1 is the hero…" (Kahneman, 2011, p. 21). We want to believe we make conscious, rational decisions using System 2. "Most often, however, you are just an observer to a global evaluation that your System 1 delivers." (Kahneman, 2011, p. 310).

In other words, while System 2 cognition dominates our theories, System 1 cognition drives a meaningful portion of our behavior. No wonder security is still a major problem.

We are not arguing that System 2 theories should be discarded; there is an important place for System 2 theories because System 2 cognition controls a substantial portion of information security behavior. Instead, we are arguing that there is a fundamental flaw in our understanding of information security behavior because our System 2 theories and empirical research do not address the meaningful proportion of behavior guided by System 1 cognition.

In order to develop System 1 theories, we need to understand the key aspects of System 1 cognition that make it different from System 2 cognition. Let us summarize the key elements of System 1 cognition. System 1 cognition is, by design, fast (usually less than a second) but not necessarily accurate. It works by matching stimuli in the current context to readily accessible heuristics that are instinctive or learned from past experience. The many different stimuli trigger many different heuristics that fire simultaneously and produce a set of responses that are integrated by a traffic facilitator.

Thus the primary building block for a System 1 theory is the heuristic, a matching situation-response. Heuristics are created by individual experience so they are person-context specific: each person has his or her own set of heuristics that are triggered by stimuli in the context. The first element of a System 1 theory is the context. Each context is different. Thus a System 1 theory cannot be about "security policy compliance" because this is not a single context. Instead System 1 theories will focus on very specific activities, such as phishing or logging out of a computer when not using it, because the context for phishing (e.g., reading email) is fundamentally different than that for logging out (e.g., taking a break).

The second important element of a System 1 theory is the person. Some heuristics may be instinctive and shared by all humans, and likewise some experiences are widely shared by some groups of people so they have common heuristics. As we think about creating a System 1 theory, we need to theorize about the experiences and resulting heuristics of the people interacting in contexts we are theorizing about. What experiences are these people likely to have in common that are relevant for the context we are theorizing about?

The heuristics that are triggered by stimuli are highly dependent on the specific stimuli encountered and heuristics most accessible in working memory. We should expect framing effects because the words "failure" and "success" trigger very different heuristics. Likewise, we should expect priming effects because the contents of working memory are affected by the events immediately prior to the stimulus. Thus the actions of System 1 may appear unstable over time and "irrational" to an outside observer who is grounded in System 2 thinking. Nonetheless, the actions are very "rational" to System 1.

Although System 1 cognition is common, sometimes the results of System 1 trigger the use of System 2 cognition. A negative stimulus, a discrepancy, or a low FOR inform us that System 2 cognition is called for. Initial empirical evidence suggests that most people's System 1 have uniformly good abilities to detect the need for System 2 cognition, but there are large individual differences in choosing whether or not to actually engage in the more effortful thinking of System 2. The implication for System 1 theory development is that personality differences are likely to play a key role in the persistence in using System 1 cognition (in the face of erroneous System 1 conclusions) or a switch to System 2 cognition.

Information Security in a System 1 World

It is difficult to estimate the extent to which information security tasks are left to automatic System 1 cognition versus tackled by deliberate System 2 cognition. Much of what we do is under the control of System 1 (Kahneman, 2011), so it is likely that System 1 controls a meaningful amount of security behavior. In this paper, we will illustrate how to develop a System 1 theory to explain and predict user response to phishing emails. Falling for phishing emails is usually unintentional and nonmalicious, and research shows that such unintentional nonmalicious behavior is a major cause of security breaches (Ernst & Young, 2017; PricewaterhouseCoopers, 2017). We have quite deliberately picked a context in which we believe we can persuade most readers that behavior is likely under the control of System 1 (Vishwanath, Herath, Chen, Wang, & Rao, 2011). One natural objection is that we have "cherry-picked" the context and that our thesis is limited to a very narrow set of situations where we use System 1 cognition.

Nonetheless, if current, well-established, theories fail to explain behavior in even one narrow case, then we argue that this is prima facie evidence that they need to be improved. More importantly, though, for the many times when behavior is under the control of System 2, System 1 cognition has already completed and its results have biased the System 2 cognition that follows it (Kahneman, 2011). Most prior research would treat this bias as the unexplained variance we see in empirical studies of System 2 theories, but we argue that some of this unexplained variance (which is often substantial: see Moody et al. (2018)) is actually predictable using System 1 theories.

What influences information security behavior in a System 1 world? System 1 behavior is driven by rapid pattern matching, usually driven by a stimulus, which is typically something new in our environment (e.g., an email message, a web pop-up, a request in person or over the phone). In this world, heuristics built on past experience that are close at hand drive behavior (Kahneman, 2011). Thus theory becomes less about general factors (e.g., perceptions of threat severity) and more person-context specific because behavior results from contextual stimuli that trigger individual heuristics built over a lifetime of experiences. We begin by examining phishing and looking at how three existing security theories would explain user behavior. We then consider the three fundamental elements of a System 1 phishing theory (the phishing email, the context in which it read, and the person receiving it). We conclude this section by developing some propositions for a System 1 theory for phishing.

Phishing

Phishing and spear phishing emails are becoming increasingly common (Ernst & Young, 2017; PricewaterhouseCoopers, 2017). A phishing email designed to trick employees into revealing their passwords often works by providing a link to a phony URL which requests them to login using their company credentials. A phishing email designed to spread malware such as ransomware often just requires employees to click the link and the malware takes over.

Suppose we consider phishing through the lens of several commonly used System 2 theories. Bulgurcu et al. (2010) draw on the theory of planned behavior (TPB) to argue that an employee's intention to comply with an information security policy depends on the attitude towards compliance, which is formed using beliefs about the costs and benefits of compliance and noncompliance. Cox (2012) found that subjective norms and self-efficacy significantly influence attitudes about compliance. Other studies have examined online privacy protection strategies (Yao & Linz, 2008), ethical attitude development and its impact on behaviors in information security practice (Chatterjee, Sarker, & Valacich, 2015; Chiang & Lee, 2011), and attitudes towards security behavior for executives' use of biometric authentication (Seyal & Turner, 2013). When applied to phishing, a security theory based on this research would argue that an employee receiving a phishing email deliberately considers the costs and benefits of clicking the link, and subjective norms (e.g., would my colleagues think I should click this link?) to form an attitude toward clicking the link, which would then influence actual behavior.

Another long-used theory is deterrence theory (DT) (D'Arcy et al., 2009; Straub & Nance, 1990), which also proposes that an individual deliberately weighs the costs and benefits of an action (with a focus on sanctions) before choosing to do it or not (Pratt, Cullen, Blevins, Daigle, & Madensen, 2006). Herath and Rao (2009) found that severity and certainty of penalties, social pressures of normative beliefs, and peer behavior significantly affected an employee's extrinsic motivation to comply with an information security policy, while perceived effectiveness influenced the individual's intrinsic motivation to comply. Other studies have found similar effects, with recent studies adding a "rational choice calculus" construct that is influenced by an individual's propensity, moral beliefs, and perceived deterrence (Hu, Xu, Dinev, & Ling, 2011). DT would argue that when an employee receives a phishing email he or she would deliberately think about the severity of a penalty (e.g., what is the penalty for clicking a phishing link?), certainty of a penalty (e.g., how likely is it I will get caught?), social pressures (e.g., would it be embarrassing if others find out about my noncompliance?), and finally perceived effectiveness (would clicking this link improve my effectiveness?). These thoughts then inform a rational choice on whether to click the link.

Protection Motivation Theory (PMT) was developed to explain how individuals cope with threats (Rogers, 1983). PMT for information security was extended by Johnston and Warkentin (2010) who argued that it invokes two separate appraisals, the threat (severity and likelihood) and the response (efficacy, cost, and one's self-efficacy to enact it). Studies have examined habit and PMT (Vance, Siponen, & Pahnila, 2012), used PMT and DT to further examine compliance (Johnston et al., 2015), and various extensions to PMT and fear appeals (Crossler et al., 2013; Johnston et al., 2016; Warkentin et al., 2011; Willison & Warkentin, 2013). For our phishing example, PMT would argue that an employee receiving an email with a link would first assess the likelihood that it was a threat and the magnitude of threat, and then assess the response cost (and efficacy and self-efficacy) of not clicking the link. The decision to click or not click would be based on this balance of threat and response.

Hopefully at this point, you have begun to realize that TPB, DT, and PMT do not really fit phishing (Vishwanath et al., 2011). They are System 2 theories designed to explain behavior when a user deliberately thinks about the situation and makes an informed choice or how to proceed. Phishing is a situation where we often fail to engage our System 2 cognition – or at least when we are tricked by a phishing email (Vishwanath, et al., 2011). And as we argued above, it is possible that a user who falls for a phishing email has engaged his or her System 2 and thought deeply about the email – he or she was simply fooled. Our central thesis is that there are times when System 1 controls behavior (or shapes the System 2 cognition that follows it) and our current theories fail to capture this.

Elements of a System 1 Phishing Theory

We need to consider three separate elements as we build a System 1 theory for phishing: the phishing email and stimuli it contains, the context, and the person.

A phishing email contains a number of distinct elements, each of which can be a stimulus. For example, the name of the sender and the subject line may trigger heuristics. The body of the email can we written in a formal letter style (with a greeting and the recipient's name) or can just launch into the content. The style as well as the content can trigger different heuristics, and can then match between the sender, subject and style. The content of the email can be crafted in a number of different ways. It can appeal to an individual's sense of greed (e.g., lottery winnings, a paid survey), loss aversion (e.g., loss of email access), humanitarianism (e.g., charitable appeal), authority (e.g., IRS), affiliation (e.g., neighborhood group), prurient desires (e.g., naked celebrities), or a host of other motivations (Hong, 2012). A spear phishing email often masquerades as a legitimate email from a fellow employee and asks for some innocuous action, such as RSVPing to party or accessing a work-related document (Hong, 2012).[1] Many phishing emails attempt to create a sense of scarcity or urgency (Hong, 2012). Each of these elements will trigger a different set of System 1 heuristics; greed and loss aversion, for example, are processed in very different ways by our System 1 brain (Guo, et al., 2017; Lowenstein, et al., 2015).

The context is the situation in which the stimulus is received. The context is a user reading his or her email. Users read email in the office, at home, and on mobile devices in a variety of locations. Email use in an office context is often utilitarian, while email use at home or on mobile devices could be either utilitarian or hedonic (van der Heijden, 2004; Wakefield & Whitten, 2006). The utilitarian mindset is quite different than a hedonic one (Hirschman & Holbrook, 1982) with sharp differences in the contents of working memory that could influence what heuristics are most accessible.

A person can be viewed through several different lenses. Personality is an important lens because it influences how we react to stimuli and the interpretation of past experiences needed for the development of heuristics (Shaffer, 2008). Personality is a combination of innate and learned characteristics and is often stable over some moderate periods of time. For example, the anxiety about phishing was found to play role in the accuracy of detecting phishing emails (Wang, Li, & Rao, 2017). Another lens is to consider job roles and professions (e.g., police officer, doctor, accountant, sales representative) because experiences in each of these can lead to different heuristics. Culture is also an important lens, because people from the same culture often have common experiences that lead to similar heuristics. Chao and Moon (2005) argue that an individual's culture can be viewed as a mosaic which consists of "tiles" from three distinct categories (demographic, geographic, and associative), each of which can lead to its own heuristics. The demographic tile includes physical characteristics and social identities inherited from parents and ancestors, such as race and gender. The geographic tile includes the region or nation that influences identity (e.g., American, Yankee). The associative tile includes the formal or informal organizations that an individual belongs to or identifies with (e.g., company, sports team). With regard to the demographic tile, Oliveira et al. (2017) found that one element of demographics (age) influenced response to phishing emails, with older users more likely to fall for phishing emails using reciprocity or security, and younger users for emails using authority or discounts. In the next section, we use these different elements to develop part of a System 1 theory for phishing.

Developing a System 1 Phishing Theory

In building a System 1 theory for phishing, most researchers would try to start at the beginning and move logically through the conceptual flow: the phishing email, then the context in which it is read, and finally the user who reads it. This is quite appropriate, but there are other equally appropriate ways to proceed. We are strong believers in the garbage can model of science, which argues there is no orderly progression when we develop theory; instead key theoretical elements are tossed in an intellectual garbage can, mixed together, and when the can is emptied on the ground they are intermixed and can be picked up in any order (Dennis & Valacich, 2001; Martin, 1980).

We start with context, and focus on an employee reading email in the office. What would be the differences in response to a spear phishing email purportedly from a senior manager with information on an important project versus an email from a neighbor with a link for a 4-hour online sale at Home Depot? The context is such that the employee is primed to fall for the senior manager email, because his or her working memory contains work concepts and thus an email from a superior is not unexpected. Flip the context to an individual reading email on a mobile phone at home during a commercial break from a football game – or a favorite HGTV show – and the elements of the theory and predicted outcome might be different, because the context leads to different priming. We also note that since there are predictable patterns when people are working in the office versus engaging in leisure activities at home, one might take the next theoretical step and predict the times of the day/week when different types of phishing emails would be more or less successful. Thus we theorize one proposition that follows from this match between the phishing email and the context in which it is read:

> P1: An employee is more likely to fall for a phishing email when the purported author matches the context in which it is read.

In a System 2 world, the context in which an email is read is irrelevant, because where one reads an email does not affect whether it is or is not a phishing email. System 2 theories quite rightly propose broad generalized factors that would influence a thoughtful rational actor (e.g., costs, benefits, social norms, threat severity), so factors such as context would be error terms, or irrational oddities. For System 1 theories, they are an important theoretical focus.

Recognizing that System 1 not System 2 has made the decision is challenging. Self-reports are seldom useful, because individuals are typically unware of the influence of System 1 (Kahneman, 2011). Observation of the time taken is a better measure, as decisions made by System 1 are faster than those of System 2 (Kahneman, 2011), although one has to control for between-subject differences because some people read faster than others. Use of tools from neuroscience can identify what brain regions were involved in a decision, so these may help (Benedetto et al., 2006; Dimoka et al., 2012). In any case, these approaches will not uncover the effects of System 1 cognition on the System 2 cognition that follows. Perhaps the best measure is to examine the behavior of interest (response to phishing emails) to see how it is affected. Thus one way to test this proposition experimentally would be to prime subjects with either an office setting or a home setting and then test the effectiveness of phishing emails from different authors. If we see differences, then we can surmise that System 1 is at work. However, note our concerns with experimental research on System 1 in the Discussion.

Another important System 1 element is the match between the phishing email and the recipient's personality, which is also not often considered in System 2 security theories. For the purposes of example, let us select a context of an employee working in the office who receives a phishing email purportedly from a superior. There are many aspects of personality that could be important to in this situation. One of the most useful models of personality is the “Big Five” model (Barrick & Mount, 1991). We will focus on one personality trait as an example: agreeableness. Agreeableness is cooperation and harmony with others (Barrick & Mount, 1991). Individuals high in agreeableness are optimistic, trusting, and willing to compromise with others (Barrick & Mount, 1991). Because they are more trusting, they are more susceptible to phishing (Cho, Hasan, & Oltramari, 2016; Modic & Lea, 2011). In a System 1 world, individuals high in agreeableness have many easily accessible heuristics generating a fundamental appetitive response (Costa, McCrae, & Dye, 1991) that drives them to trust and comply with phishing email (Cho et al., 2016). When asked to comply with good security behaviors, such as changing their password or backing up data, they are more likely to comply. However, when exposed to phishing and asked to click on a link, they are also more likely to comply and compromise security.

Individuals low in agreeableness focus on their own self-interest and tend to be skeptical and suspicious of others (De Dreu & Nauta, 2009). Individuals low in agreeableness have many easily accessible heuristics generating resistance and noncompliance. They are more skeptical and less likely to trust an email (Cho et al., 2016; Modic & Lea, 2011). Their default answer is a fundamental avoidance reaction (Costa et al., 1991) that drives them away from a requested behavior. Thus when asked to comply with a security policy or a phishing email, their System 1 heuristics will generate avoidance. If the individual persists in considering the phishing email, and engages his or her System 2 cognition, it will be heavily shaded by System 1's immediate negative response. So individuals low in agreeableness are less likely to be tricked by phishing emails. Thus we could theorize a proposition that follows from this personality trait:

> P2: An employee whose personality is high in agreeableness is more likely to fall for a phishing email.

In a System 2 world, the reader's personality is less important, because the reader's personality does not affect whether email is a phishing email. Personality may influence the phishing assessment, but once again it is often treated as an error term, something to be avoided in the crisp clean world of rational thinking (Huber, 1983). In contrast, theorizing in a System 1 world embraces personality differences because they are important driver of behavior.

The inherent bias in the heuristics accessible to individuals who are high or low in agreeableness will influence their response to phishing and other information security behaviors. But, once again, we have to be very careful to focus on the joint person-context relationship. We cannot say that high agreeableness leads to better or worse security compliance, because the heuristics linked to the agreeableness personality trait are biased to an appetitive reaction -- compliance with the request, whether the request is to comply with security policy or respond to a phishing email.

Personality and experience work together. So, if I am high in agreeableness and I have had my identity stolen by clicking on a link in a phishing email, it is likely that heuristics associated with my bad experience dominate my fundamental personality when I am in context that triggers heuristics based on that experience. The heuristics that we use are often driven by vividness and proximity (Loewenstein et al., 2015), and it is likely that identity theft is a rather vivid experience; thus, an identity theft heuristic is likely to outweigh an agreeableness heuristic when System 1 comes to put them together. In general, avoidance responses dominate appetitive responses, perhaps because of the evolutionary need to avoid danger (Kahneman, 2011). Thus, one bad experience often replaces a set of good experiences with an avoidance heuristic. Likewise, if the identity theft was recent (temporal proximity) then the heuristic associated with it is likely to be even stronger. Thus we could theorize a proposition that follows from experience:

> P3: An employee with a prior bad experience due to a phishing email is less likely to fall for a phishing email.

In a System 2 world, past experience does not change the present phishing email, but it would be logical to argue that past experience with phishing could influence current behavior. System 2 theories could argue that past experience would make one more vigilant and more likely to compare the current email to the past phishing one, and so on. In contrast, System 1 theories would argue the same outcome, but a different theoretical mechanism: instant avoidance. Some people bitten by a dog have instant System 1 aversive reactions to dogs (i.e., fear); they do not compare the current dog or its behavior to the one that bit them; instead, the sight of a dog invokes a fear. The same could happen with phishing emails.

It is also useful to theorize about other effects of past experience, specially, situational normality. System 1 heuristics are triggered by the situation. If the phishing email appears normal, then normal heuristics will be triggered. If the individual is in an organization that routinely sends emails with embedded links and invites employees to click them, then the individual will have a strong heuristic to click the phishing link. Note that is separate and distinct from any training that advocates not clicking links in emails; training is not as readily available as past successful experience clicking links. Similarly, individuals often receive emails from stores, organizations, universities etc. that he or she interacts with; if the individual has a history of good experiences clicking email links from these, then there will also be a general heuristic that generates an appetitive response to a link in an email. Thus:

> P4: An individual who routinely receives emails containing links is more likely to fall for a phishing email purportedly from one of these routine senders.

We note that System 2 security theories have begun to include a role for habit (Moody et al., 2018), which they argue are behaviors that "have become automatic insofar as they are performed without mindful instruction to do so" (Moody et al., 2018). Habits are akin to System 1 cognition, although most prior research treats them as another factor in a variance model of System 2 cognition (Moody et al., 2018), rather than the separate and distinct process that psychology argues System 1 cognition to be (Kahneman, 2011).

As System 1 theories are generated on specific contexts, personalities, and individual differences in heuristics and past experiences, there are also widely prevalent cognitive heuristics that can be examined. As mentioned previously, one such heuristic is confirmation bias (Klayman, 1995). Confirmation bias is ubiquitous and is used to simplify information processing and reduce cognitive dissonance (Kunda, 1990). Individuals prefer information that conforms with their prior beliefs, rather than rely on information that contradicts it. Applying this to the phishing example, an individual's response to a phishing attempt will be dependent on whether the information contained in the phishing attempt conforms with their previously formed beliefs (e.g., an email from a supervisor with a link to a relevant news story). If the individual has not received a phishing email from a compromised account before, confirmation bias would indicate they would click the email because they have not experienced anything that would challenge the confirmation bias of “my supervisor occasionally sends me news articles relevant to my work.” However, if the individual has been regularly exposed to phishing attempts using the same methodology, confirmation bias would play in the opposite manner, leading them to ignore the email (System 1) or report the email (System 2). Therefore, those that have been exposed to legitimate emails before that are similar to the phishing attempt, the individual will have an affirmative confirmation bias that drives their System 1 reaction to the behavior, thus increasing their likelihood for falling for the phishing attempt. Thus:

> P5: When an employee has not been exposed to a phishing attempt before, but has been exposed to legitimate emails similar to the phishing attempt, they will be more likely to fall for the phishing attempt.

In summary, in a System 1 world, it is the person-context that drives behavior – the stimulus triggers heuristics built by personality traits and past experiences in the context (or related contexts). In many cases, System 1 makes the decision, and we act. In other cases, as we noted above, the System 1 answer triggers System 2 to act -- we realize we need to apply the deliberate cognition of System 2 to the situation, although our System 2 cognition is heavily influenced, for good or bad, by our System 1 answer.

As this section argues, the individual and the context matters. When we develop System 1 security theories, they tend to be more specific than their System 2 counterparts. System 1 cognition involves quickly matching (in less than one second) elements in the current situation to heuristics learned from past behaviors and producing a recommended action that is either immediately implemented, or an alternative course of action that often serves as an anchor for the detailed System 2 cognition that follows it. Because we theorize about specific people in the specific contexts, we need to know more about the individuals of interest (e.g., gender, culture) and how past behaviors have been experienced. We also need to pay particular attention to priming and framing effects that appear illogical to our thoughtful System 2, but strongly influence our fast and reactive System 1.

Discussion

Past information security research has focused on the deliberate cognition of System 2. These theories are useful, but have an important boundary condition: they cannot explain or predict security behavior that is controlled by the instant and reactive cognition of System 1. System 1 accounts for a substantial portion of human behavior (Kahneman, 2011), so this presents a sizeable hole in our understanding of security behavior, since a large proportion of security incidents can be traced back to unintentional nonmalicious employee behavior (Ernst & Young, 2017; Guo et al., 2017; PricewaterhouseCoopers, 2017). Even when System 2 cognition drives behavior, System 1 cognition biases the System 2 cognition that follows it (e.g., confirmation bias) (Kahneman, 2011).

In this paper, we have offered a different theoretical approach to information security, based on System 1 cognition. We view this System 1 approach as complementing prior research because it focuses on the one key boundary condition of most prior theories: most theories assume a rational individual making and thoughtful decisions under the guidance of System 2 cognition. Research on System 1 theories of information security is in its infancy, so it is difficult to draw data-based conclusions about its usefulness.

With the introduction of tools from cognitive neuroscience (i.e., EEG, fMRI, fNIRS), the information systems field has a greater ability to examine System 1 cognition, elucidating latent biases (see, for example, Minas, Potter, Dennis, Bartelt, & Bae (2014)). These tools provide the opportunity to create and strengthen System 1 theories in information security research. Therefore, we believe there are several important implications for research and practice.

Implications for Research

First and foremost, the ubiquity of System 1 cognition calls for more research on System 1 approaches to security. We expect that building System 1 theories will be difficult for most researchers at first, because they start in a very different place than our System 2 theories. Thinking in System 1 terms is not the normal way we approach theory development. We have provided some initial examples of theoretical reasoning based on a System 1 worldview, and we call on researchers to develop and test more well-developed System 1 security theories to fill the current hole in our current understanding of information security. Our focus was on nonmalicious employee behavior – in other words an insider who is not attempting to cause harm – which is a major cause of security breaches (Ernst & Young, 2017; PricewaterhouseCoopers, 2017). We believe there are important boundary conditions and that System 1 theories may be more limited in explaining deliberate malicious behavior – but this is an issue for future research.

There are many potentially interesting places to start creating System 1 security theories. One prime candidate is situations where System 2 theories fail to explain a large amount of variance, especially if the System 2 theory has been shown to be useful in other situations. After all, if a System 2 theory explains a lot of variance, a System 1 theory is not likely to add much. On the other hand, if a System 2 theory does not explain much variance, then it could be an inherent problem with the theory itself (i.e., a problem that could be addressed by adding more constructs), or it could be a situation where System 2 cognition is not widely used, so a System 2 theory has little opportunity to explain behavior.

Second, as we argued above, System 1 theories are likely to be person-context dependent (although there are some common biases and heuristics that warrant investigation as well). This means we will need to have a better understanding of the person. What aspects of personality are likely to translate into heuristics that favor an appetitive or avoidance response? What specific past experiences are both predictable and important? What low-probability experiences should we consider? This calls for theories that are much more grounded in the person, and we may have different theories for different types of people.

This also means that empirical research will need to be more cognizant of and careful to describe the participant population. Student participants, for example, may have fewer life experiences with negative security events such as identity theft or password theft because of their age and because they may be less likely to be targets. Thus students may have Systematic differences in the heuristics they draw upon than other populations. This does not mean they are less appropriate for research, but in theorizing for a specific study, researchers need to consider the life experiences of participants – building theories separate from a specific population may be difficult.

Third, the person-context nature of System 1 also suggests that we will need different theories for different contexts of security behavior. Unlike System 2 theories that generalize to a variety of security behaviors, such as building strong passwords and avoiding spear phishing (e.g., deterrence theory, protection motivation theory), System 1 theories will be closely tied to specific security behaviors. A System 1 theory for strong passwords is likely to have different elements than avoiding phishing because the past experiences that have created heuristics in these two contexts are likely different. Physical context may also matter, because the contextual heuristics we apply when working at the office may be different than those when we surf the net for pleasure at home. Likewise, the nature of the attack may influence the theory. A phishing request may get a different response if it is delivered via email, Facebook, or phone; or if it purports to be from a friend, a boss, or an unknown person; or because the heuristics triggered by different media and different requestors may be different.

Fourth, System 1 theories start with a very narrow focus because of the highly specific person-context that may lead to different outcomes. They are not as abstract and context independent as System 2 theories such as PMT that can apply to phishing, password sharing, failing to logout, and so on. As we learn more about System 1 theories, we may find larger elements that we can abstract to. For example, researchers may choose to bound their theories to a work context or a home context, thereby eliminating a huge swath of complexity. Likewise, higher levels of abstraction about the person may simplify the theories, either by drawing on causal logic or correlation. For example, transformational leaders are often high in agreeableness (Bono & Judge, 2004), so we may also be able to theorize that senior leaders would be very likely to be deceived by a phishing email. Alternately, researchers may deliberately choose to make a simplifying assumption about the person or context and remove either from the theory and propose a theory that is potentially applicable to multiple persons or contexts and let the empirical data suggest its boundary conditions. The specific nature of System 1 theories does not make them less valuable, it implies that there are likely to be more of them that develop initially to account for the person-context complexity.

Fifth, the person-context nature of System 1 thinking has one unfortunate implication for research. System 1 heuristics are triggered by normal situations (Bargh & Chartrand, 1999). When we move individuals out of normal situations, they are less likely to use System 1 thinking. This means that research done outside of the normal context may not be generalizable to the normal context. When users were asked to participate in a research study and self-report how likely they were to fall for different sample phishing emails, their responses did not match their actual behavior when exposed to phishing in the wild (Oliveira et al., 2017). If we study phishing by using a lab experiment or a survey to present participants with different email messages, they may be aware that they are in a research study and thus be more likely to engage in System 2 cognition, because a research study does not match a normal pattern that triggers System 1 thinking. Thus we may not find the same results in the “abnormal” research context as when we send a phishing email to participants unannounced and they encounter it as part of their normal day-to-day routine. Experimental and survey research may be less likely to find significant effects because System 2 cognition is more likely to override System 1 cognition in the unfamiliar contexts of experimental and survey research. This may mean that researchers will need to take a different approach to traditional laboratory and survey research, such as priming subjects with the context in which the phenomenon is theorized to occur.

Finally, there may be an opportunity to fit into System 1 theorizing into existing System 2 theories (Tsohou et al., 2015). We have treated System 1 cognition as separate and distinct from System 2 cognition, which it is (Kahneman, 2011). However, if we choose to override our System 1 response and devote cognitive effort to deliberately applying System 2 cognition, as current security theories argue, the initial System 1 response colors our System 2 cognition (Tsohou et al., 2015). Our current security theories (e.g., TPB, DT, PMT) might be improved by adding elements of System 1 cognition.

Implications for Practice

The implications for practice from System 1 theorizing call for profound changes for practice. The implications from System 2 theories would suggest that educating users about threat severity and likelihood, and/or sanction severity and likelihood, should affect security compliance. Thus security education, training, and awareness (SETA) initiatives should focus on these important levers. In contrast, theorizing based on System 1 cognition suggests that these types of SETA activities will have little effect on behavior because these factors are ignored by System 1 heuristics. System 1 does not use knowledge; knowledge is used by System 2. Recent empirical evidence suggests that these types of SETA activities have no effect on behavior after a few weeks (Zhang, 2016), offering some support for our arguments.

System 1 theorizing would suggest that there are three levers we can use to improve security compliance behavior. The first is to change the person -- that is, one's personality. Replacing people or altering personality is unlikely to be a useful lever to improve security in most organizations.

A second approach is to change the experience-based heuristics that System 1 uses. Heuristics are heavily context dependent, so this entails a lot of very detailed thinking about specific contexts. Suppose for example, we want to reduce the effectiveness of phishing. Individuals have a deep set of positive experiences with clicking on links in emails. Many organizations send emails with legitimate links, and we are used to clicking them to accomplish our goals with positive results. One approach would be to replace this heuristic that generates an appetitive response to email links with a heuristic that generates an avoidance response to email links. In other words, aversion training.

This may not be as nefarious as it sounds. For example, it could take the form of the organization *deliberatively* sending phishing emails regularly, then locking individuals that click on them out of their account for 15 minutes (or some other similar annoyance). Likewise, suppose we send a series of phishing messages with links that trigger a very loud alarm when they are clicked (e.g., air horn). After being tricked once or twice, most users would develop a very strong avoidance heuristic to clicking links. We could of course, merely display a warning message, but this would have a less powerful effect for building a System 1 heuristic because a warning message requires System 2 cognition to understand. An air horn, on the other hand, is very powerful negative feedback that requires no cognition to understand. The next time your hand moved the mouse over an email link, your System 1 would immediately begin shouting a strong avoidance response. This will help improve their heuristics and also, initially, trigger more System 2 cognition for clicking on email links.

A third approach is to attack situational normality. System 1 heuristics are triggered by situations that match well-learned patterns of normal behavior (Bargh & Chartrand, 1999; Triandis, 1971). If we change the situation so common attacks no longer match normal situations, then we break the automatic pattern matching of System 1 cognition. Consider our phishing example that requires the user to click an email link. One way to break the situational normality of clicking email links would be to send several phishing messages with links to all employees every day for several weeks. Employees who receive many such emails and are fooled by a few of them (even without an air horn to provide a negative response) will shift their definition of normal, so that any email with a link is "normally" a phishing attack; thus the System 1 heuristic when viewing an email with a link will automatically generate an avoidance response.

Another way to break this situational normality would be to prohibit organizational emails from containing a clickable link; instead all emails would instruct users to go to a common, well-known site (e.g., department home page) and provide detailed instructions to locate the page of interest (e.g., menu directions or search instructions). Therefore, any email that contains a clickable link is in violation of the policy and is not normal. Such a break in situational normality should break the cycle of System 1 cognition and trigger System 2 cognition. Of course, such a policy would impose greater costs on those sending and receiving emails, a cost that would have to be balanced against the cost of successful phishing attacks.

We also note that because System 2 cognition follows System 1 cognition, we need to understand and actively combat the way System 1's heuristics shapes System 2 cognition (Tsohou et al., 2015). SETA training can make users aware that their instinctive reactions to a stimulus that provokes an action with security consequences (such as avoiding social engineering (e.g., phishing), complying with policy (e.g., regular backups) or deciding not to break policy (e.g., not taking confidential information home), may not be the most appropriate reaction and so they should engage in System 2 cognition that recognizes the potential biases of System 1 heuristics.

Conclusion

Information Systems security continues to be of critical concern to both researchers and managers (Bulgurcu et al., 2010; Kappelman, McLean, Johnson, & Torres, 2016; Moody et al., 2018). We argue that our current research and theories are useful for behavior based on System 2 cognition, but therein lies the problem. We have focused our research efforts on one type of cognition. We have focused on the times that deliberate cognition flies the security plane, not the autopilot that controls much of the journey. Perhaps one reason why information security is such a major issue in organizations and why existing training programs seem to have little effect (Siponen & Vance, 2010), is because we have failed to focus on the second locus of control: System 1.

System 2 theories will continue to be important, because we do use deliberate cognition. However, because System 1 cognition is integral to an individual's behavior, we need new security research and theories based on System 1 cognition to complement our current understanding. Such research and theory are likely to produce a very different understanding of the factors that drive behavior, and are likely to produce very different recommendations for improving security in practice, recommendations that can complement existing SETA programs.

Notes

[1] A well-designed spear phishing attack at a Big Ten University tricked a substantial number of IT employees because it was sent using the email address of the CIO and purported to show office and cubicle assignments in the newly redesigned IT office building.

References

Achtziger, A., & Alós-Ferrer, C. (2013). Fast or rational? A response-times study of bayesian updating. *Management Science, 60*(4), 923-938.

Anderson, C. L., & Agarwal, R. (2010). Practicing safe computing: A multimedia empirical examination of home computer user security behavioral intentions. *MIS Quarterly, 34(*3), 613-643.

Ariely, D. (2009). Predictably irrational. New York: HarperCollins.

Bago, B., & De Neys, W. (2017). Fast logic: Examining the time course assumption of dual process theory*. Cognition, 158*, 90-109.

Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist, 54*(7), 462.

Bargh, J. A., & Chartrand, T. L. (2000). Studying the mind in the middle: A practical guide to priming and automaticity research. In H. Reis & C. Judd (Eds.), *Handbook of research methods in social psychology.* New York: Cambridge University.

Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology, 71*(2), 230-244.

Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin, 126*(6), 925.

Bargh, J. A., & Morsella, E. (2008). The unconscious mind. *Perspectives on Psychological Science, 3*(1), 73-79.

Barrick, M. R., & Mount, M. K. (1991). The big five personality dimensions and job performance: A meta-analysis*. Personnel Psychology, 44*(1), 1-26.

Bellini-Leite, S. d. C. (2013). The embodied embedded character of system 1 processing. *Mens Sana Monographs, 11*(1), 239-252.

Benedetto, D. M., Kumaran, D., Seymour, B., & Dolan, R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science, 313*(684), 684-687.

Bono, J. E., & Judge, T. A. (2004). Personality and transformational and transactional leadership: A meta-analysis. *Journal of Applied Psychology, 89*(5), 901-910.

Boss, S. R., Galletta, D. F., Lowry, P. B., Moody, G. D., & Polak, P. (2015). What do systems users have to fear? Using fear appeals to engender threats and fear that motivate protective security behaviors. *MIS Quarterly, 39*(4), 837-864.

Bulgurcu, B., Cavusoglu, H., & Benbasat, I. (2010). Information security policy compliance: An empirical study of rationality-based beliefs and information security awareness. *MIS Quarterly, 34*(3), 523-548.

Carlston, D. E., & Skowronski, J. J. (1994). Savings in the relearning of trait information as evidence for spontaneous inference generation. *Journal of Personality and Social Psychology, 66*(5), 840.

Chao, G. T., & Moon, H. (2005). The cultural mosaic: A metatheory for understanding the complexity of culture. *Journal of Applied Psychology, 90*(6), 1128.

Chatterjee, S., Sarker, S., & Valacich, J. S. (2015). The behavioral roots of information systems security: Exploring key factors related to unethical IT use. *Journal of Management Information Systems, 31*(4), 49-87.

Chen, M., & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin, 25*(2), 215-224.

Chiang, L., & Lee, B. (2011). Ethical attitude and behaviors regarding computer use. *Ethics & Behavior, 21*(6), 481-497.

Cho, J., Hasan, C., & Oltramari, A. (2016). Effect of personality traits on trust and risk to phishing vulnerability: Modeling and analysis. *IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA).* New York: IEEE.

Costa, P. T., McCrae, R. R., & Dye, D. A. (1991). Facet scales for agreeableness and conscientiousness: A revision of the neo personality inventory. *Personality and Individual Differences, 12*(9), 887-898.

Cox, J. (2012). Information systems user security: A structured model of the knowing-doing gap. *Computers in Human Behavior, 28*(5), 1849-1858.

Crossler, R. E., Johnston, A. C., Lowry, P. B., Hu, Q., Warkentin, M., & Baskerville, R. (2013). Future directions for behavioral information security research. *Computers & Security, 32*, 90-101.

D'Arcy, J., Hovav, A., & Galletta, D. (2009). User awareness of security countermeasures and its impact on information systems misuse: A deterrence approach. *Information Systems Research, 20*(1), 79-98.

De Dreu, C. K. W., & Nauta, A. (2009). Self-interest and other-orientation in organizational behavior: Implications for job performance, prosocial behavior, and personal initiative. *Journal of Applied Psychology, 94*(4), 913-926.

de Guinea, A. O., & Markus, M. L. (2009). Why break the habit of a lifetime? Rethinking the roles of intention, habit, and emotion in continuing information technology use. *MIS Quarterly, 33*(3), 433-444.

De Neys, W. (2014). Conflict detection, dual processes, and logical intuitions: Some clarifications. *Thinking & Reasoning, 20*(2), 169-187.

Dennis, A. R., Minas, R. K., & Bhagwatwar, A. P. (2013). Sparking creativity: Improving electronic brainstorming with individual cognitive priming. *Journal of Management Information Systems, 29*(4), 195-216

Dennis, A. R., & Valacich, J. S. (2001). Conducting experimental research in information systems. *Communications of the Association for Information Systems, 7*(1), 5.

Dijk, J. v., Kerkhofs, R., Rooij, I. v., & Haselager, P. (2008). Special section: Can there be such a thing as embodied embedded cognitive neuroscience? *Theory & Psychology, 18*(3), 297-316.

Dimoka, A., Banker, R. D., Benbasat, I., Davis, F. D., Dennis, A. R., Gefen, D., . . . Pavlou, P. A. (2012). On the use of neurophysiological tools in is research: Developing a research agenda for NeuroIS. *MIS Quarterly, 36*(3), 679-702.

Duckworth, K. L., Bargh, J. A., Garcia, M., & Chaiken, S. (2002). The Automatic evaluation of novel stimuli. *Psychological Science, 13*(6), 513-519.

Ernst, & Young. (2017). *EY's 19th global information security survey 2016-17*. http://www.ey.com/gl/en/services/advisory/ey-global-information-security-survey-2016. Current February, 2018.

Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology, 59*(1), 255-278.

Evans, J. S. B. T. (2014). Two minds rationality. *Thinking & Reasoning, 20*(2), 129-146.

Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science, 8*(3), 223-241.

Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*(1), 297-327.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50*(2), 229-238.

Fiske, S. T., Lin, M., & Neuberg, S. L. (1999). The Continuum Model: Ten Years Later. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social and cognitive psychology* (pp. 231-254). New York: Guilford.

Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (pp. 1-74): Cambridge, MA: Academic Press.

Gersick, C. J. G., & Hackman, J. R. (1990). Habitual routines in task-performing groups. *Organizational Behavior and Human Decision Processes, 47*(1), 65-97.

Guo, L., Trueblood, J. S., & Diederich, A. (2017). Thinking fast increases framing effects in risky decision making. *Psychological Science, 28*(4), 530-543.

Herath, T., & Rao, H. R. (2009). Protection motivation and deterrence: a framework for security policy compliance in organizations. *European Journal of Information Systems, 18*(2), 106-125.

Hirschman, E. C., & Holbrook, M. B. (1982). Hedonic consumption: emerging concepts, methods and propositions. *The Journal of Marketing, 46*(3), 92-101.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review, 109*(4), 679-709.

Hong, J. (2012). The state of phishing attacks. *Communications of the ACM, 55*(1), 74-81.

Hu, Q., Xu, Z., Dinev, T., & Ling, H. (2011). Does deterrence work in reducing information security policy abuse by employees? *Communications of the ACM, 54*(6), 54-60.

Huber, G. P. (1983). Cognitive style as a basis for MIS and DSS designs: Much ado about nothing? *Management Science, 29*(5), 567-579.

Hui, K. L., Vance, A., & Zhdanov, D. (2016). *Securing digital assets*. https://www.misqresearchcurations.org/. Current February, 2018

Jasperson, J., Carter, P. E., & Zmud, R. W. (2005). A comprehensive conceptualization of post-adoptive behaviors associated with information technology enabled work systems. *MIS Quarterly, 29*(3), 525-557.

Johnson, E. D., Tubau, E., & De Neys, W. (2016). The doubting system 1: Evidence for automatic substitution sensitivity. *Acta Psychologica, 164*, 56-64.

Johnston, A. C., & Warkentin, M. (2010). Fear appeals and information security behaviors: an empirical study. *MIS Quarterly, 34*(3), 549-566.

Johnston, A. C., Warkentin, M., McBride, M., & Carter, L. (2016). Dispositional and situational factors: influences on information security policy violations. *European Journal of Information Systems, 25*(3), 231-251.

Johnston, A. C., Warkentin, M., & Siponen, M. T. (2015). An enhanced fear appeal rhetorical framework: Leveraging threats to the human asset through sanctioning rhetoric. *MIS Quarterly, 39*(1), 113-134.

Kahneman, D. (2011). *Thinking, fast and slow*. New York: Macmillan.

Kahneman, D., & Tversky, A. (1979). On the interpretation of intuitive probability: A reply to Jonathan Cohen. *Cognition, 7*(4), 409-411.

Kappelman, L., McLean, E., Johnson, V., & Torres, R. (2016). The 2015 SIM IT issues and trends study. *MIS Quarterly Executive, 15*(1), 55-83.

Klayman, J. (1995). Varieties of confirmation bias. *Psychology of Learning and Motivation, 32*, 385-418.

Klein, G. (1999). Applied decision making. In P.A. Hancock (Ed.), *Human performance and ergonomics* (pp. 87-107). San Diego: Academic Press.

Kosner, A. W. (2014). *Actually two attacks in one, target breach affected 70 to 110 million customers*. https://www.forbes.com/sites/anthonykosner/2014/01/17/actually-two-attacks-in-one-target-breach-affected-70-to-110-million-customers/#187b23b45482. Current February, 2018.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*(3), 480-498.

Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (1990). Emotion, attention, and the startle reflex. *Psychological Review, 97*(3), 377-395.

Liang, H., & Xue, Y. (2009). Avoidance of information technology threats: A theoretical perspective. *MIS Quarterly, 33*(1), 71-90.

Liang, H., & Xue, Y. (2010). Understanding security behaviors in personal computer usage: A threat avoidance perspective. *Journal of the Association for Information Systems, 11*(7), 394.

Loewenstein, G., O'Donoghue, T., & Bhatia, S. (2015). Modeling the interplay between affect and deliberation. *Decision, 2*(2), 55-81.

Lord, C. G., Scott, K. O., Pugh, M. A., & Desforges, D. M. (1997). Leakage beliefs and the correspondence bias. Personality and Social *Psychology Bulletin, 23*(8), 824-836.

Louis, M. R., & Sutton, R. I. (1991). Switching cognitive gears: from habits of mind to active thinking. *Human Relations, 44*(1), 55-76.

Maheswaran, D., & Chaiken, S. (1991). Promoting systematic processing in low-motivation settings: Effect of incongruent information on processing and judgment. *Journal of Personality and Social Psychology, 61*(1), 13-25.

Martin, J. M. (1980). A garbage can model of the research process. In J. McGrath, J. Marting & R. Kulka (Eds.), *Judgment calls in research*. Beverly Hills, CA: Sage.

Minas, R. K., Poor, M., Dennis, A. R., & Bartelt, V. L. (2016). A prime a day keeps calories away: The effects of supraliminal priming on food consumption and the moderating role of gender and eating restraint. *Appetite, 105*, 494-499.

Minas, R. K., Potter, R. F., Dennis, A. R., Bartelt, V., & Bae, S. (2014). Putting on the Thinking Cap: Using NeuroIS to Understand Information Processing Biases in Virtual Teams. *Journal of Management Information Systems, 30*(4), 49-82.

Modic, D., & Lea, S. E. *How neurotic are scam victims, really?* The big five and Internet scams. https://ssrn.com/abstract=2448130. Current February, 2018.

Moody, G. D., Siponen, M., & Pahnila, S. (2018). Toward a unified model of information security policy compliance. *MIS Quarterly, 42*(1), 285-A222.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175.

Oliveira, D., Rocha, H., Yang, H., Ellis, D., Dommaraju, S., Muradoglu, M., . . . Ebner, N. (2017). Dissecting spear phishing emails for older vs young adults: On the interplay of weapons of influence and life domains in predicting susceptibility to phishing. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems.* New York: Association for Computing Machinery.

Perlroth, N. (2014). *Home Depot data breach could be the largest yet.* https://bits.blogs.nytimes.com/2014/09/08/home-depot-confirms-that-it-was-hacked/. Current February, 2018.

Pratt, T. C., Cullen, F. T., Blevins, K. R., Daigle, L. E., & Madensen, T. D. (2006). The empirical status of deterrence theory: A meta-analysis. In F. Cullen, J. Wright & K. Blevins (Eds.), *Taking stock: The status of criminological theory, Vol. 15* (pp. 367-396). New York: Rutledge.

PricewaterhouseCoopers. (2017). *The global state of information security survey 2017*. Retrieved from http://www.pwc.com/gsiss2015. Current February, 2018.

Rogers, R. W. (1983). Cognitive and physiological processes in fear appeals and attitude change: A revised theory of protection motivation. In J. Cacioppo & R. Petty (Eds.), *Social psychophysiology: A sourcebook* (pp. 153-176). New York: Guilford Press.

Schwarz, N., & Clore, G. L. (2007). Feelings and phenomenal experiences. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: A handbook of basic principles.* New York: Guilford Press.

Seyal, A. H., & Turner, R. (2013). A study of executives' use of biometrics: an application of theory of planned behavior. *Behavior & Information Technology, 32*(12), 1242-1256.

Shaffer, D. (2008). *Social and personality development, 6th edition*. Toronto, ON: Nelson Education.

Shariff, A. F., & Norenzayan, A. (2007). God Is watching you: Priming God concepts increases prosocial behavior in an anonymous economic game. *Psychological Science, 18*(9), 803-809.

Siponen, M., & Vance, A. (2010). Neutralization: New insights into the problem of employee information systems security policy violations. *MIS Quarterly, 34*(3), 487-502.

Smith, N. K., Cacioppo, J. T., Larsen, J. T., & Chartrand, T. L. (2003). May I have your attention, please: Electrocortical responses to positive and negative stimuli. *Neuropsychologia, 41*(2), 171-183.

Srull, T. K., & Wyer, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology, 37*(10),

Srull, T. K., & Wyer, R. S. (1980). Category accessibility and social perception: Some implications for the study of person memory and interpersonal judgments. *Journal of Personality and Social Psychology, 38*(6), 841-856.

Srull, T. K., & Wyer, R. S. (1983). The role of control processes and structural constraints in models of memory and social judgment. *Journal of Experimental Social Psychology, 19*(6), 497-521.

Stafford, T. F. (1996). Conscious and unconscious processing of priming cues in selling encounters. *Journal of Personal Selling & Sales Management, 16*(2), 37-44.

Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. New York: Psychology Press.

Stanovich, K. E. (2004). Metarepresentation and the great cognitive divide: A commentary on Henriques' "Psychology Defined." *Journal of Clinical Psychology, 60*(12), 1263-1266.

Stanovich, K. E., & West, R. F. (2000). Advancing the rationality debate*. Behavioral and Brain Sciences, 23*(5), 701-717.

Straub, D. W., & Collins, R. W. (1990). Key information liability issues facing managers: Software piracy, proprietary databases, and individual rights to privacy. *MIS Quarterly, 14*(2), 143-156.

Straub, D. W., & Nance, W. D. (1990). Discovering and disciplining computer abuse in organizations: A field study. *MIS Quarterly, 14*(1), 45-60.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction, Vol. 1*. Cambridge, MA: MIT press.

Taylor, S. E., & Fiske, S. T. (1978). Salience, attention, and attribution: Top of the head phenomena. In L. Berkowitz (Ed.), *Advances in experimental social psychology, Vol. 11* (pp. 249-288). Cambridge, MA: Academic Press.

Thompson, V., & Morsanyi, K. (2012). Analytic thinking: Do you feel like it? *Mind & Society, 11*(1), 93-105.

Thompson, V. A. (2013). Why it matters: The Implications of autonomous processes for dual process theories - Commentary on Evans & Stanovich (2013). *Perspectives on Psychological Science, 8*(3), 253-256.

Thompson, V. A., Prowse Turner, J. A., & Pennycook, G. (2011). Intuition, reason, and metacognition. *Cognitive Psychology, 63*(3), 107-140.

Todorov, A., & Uleman, J. S. (2003). The efficiency of binding spontaneous trait inferences to actors' faces. *Journal of Experimental Social Psychology, 39*(6), 549-562.

Triandis, H. C. (1971). *Attitude and attitude change*. New York: Wiley.

Tsohou, A., Karyda, M., & Kokolakis, S. (2015). Analyzing the role of cognitive and cultural biases in the internalization of information security policies: Recommendations for information security awareness programs. *Computers & Security, 52*, 128-141.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*(4481), 453-458.

van der Heijden, H. (2004). User acceptance of hedonic information systems. *MIS Quarterly, 28*(4), 695-704.

Vance, A., Siponen, M., & Pahnila, S. (2012). Motivating IS security compliance: Insights from habit and protection motivation theory. *Information & Management, 49*(3-4), 190-198.

Vishwanath, A., Herath, T., Chen, R., Wang, J., & Rao, H. R. (2011). Why do people get phished? Testing individual differences in phishing vulnerability within an integrated, information processing model. *Decision Support Systems, 51*(3), 576-586.

Wakefield, R. L., & Whitten, D. (2006). Mobile computing: a user study on hedonic/utilitarian mobile device usage. *European Journal of Information Systems, 15*(3), 292-300.

Wang, J., Li, Y., & Rao, H. R. (2017). Coping responses in phishing detection: An investigation of antecedents and consequences. *Information Systems Research, 28*(2), 378-396.

Warkentin, M., Johnston, A. C., & Shropshire, J. (2011). The influence of the informal social learning environment on information privacy policy compliance efficacy and intention. *European Journal of Information Systems, 20*(3), 267-284.

Willison, R., & Warkentin, M. (2013). Beyond deterrence: An expanded view of employee computer abuse. *MIS Quarterly, 37*(1), 1-20.

Wood, W., Quinn, J. M., & Kashy, D. A. (2002). Habits in everyday life: Thought, emotion, and action. *Journal of Personality and Social Psychology, 83*(6), 1281-1297.

Workman, M., Bommer, W. H., & Straub, D. (2008). Security lapses and the omission of information security measures: A threat control model and empirical test. *Computers in Human Behavior, 24*(6), 2799-2816.

Yao, M. Z., & Linz, D. G. (2008). Predicting self-protections of online privacy. *CyberPsychology & Behavior, 11*(5), 615-617.

Yi, Y. (1990). Cognitive and affective priming effects of the context for print advertisements. *Journal of Advertising, 19*(2), 40-48.

Zhang, T. (2016). Training decrement in security awareness training. *KSU Conference on Cybersecurity Education, Research and Practice.*

About the Authors

Alan R. Dennis is Professor of Information Systems and holds the John T. Chambers Chair of Internet Systems in the Kelley School of Business at Indiana University. He was named a Fellow of the Association for Information Systems in 2012. Professor Dennis has written more than 150 research papers, and has won numerous awards for his theoretical and applied research. His research focuses on three main themes: team collaboration; IT for the subconscious; and information security. He is co-Editor-in-Chief of AIS Transactions on Replication Research and Vice President for Conferences for the Association for Information Systems. He also has written four books (two on data communications and networking, and two on systems analysis and design).

Randall K. Minas is an assistant professor and the Hon Kau and Alice Lee Faculty Fellow at the Shidler College of Business at University of Hawai'i at Manoa. His research interests include cognitive neuroscience, cognitive and affective responses to information systems, cognitive biases, and collaboration research. He has been published in journals such as the *Journal of Management Information Systems*, *Journal of the Association of Information Systems*, *AIS Transactions on Human Computer Interaction*, and the *Journal of Applied and Preventive Psychology*. He has received Best Paper at the conferences Hawaiian International Conference on System Sciences (HICSS) and Human-Computer Interaction International (HCII). He also served as the managing editor of *MIS Quarterly Executive* from 2010-2014, a journal directed to executives with the goal of improving practice. He is co-founder of the Hawaii Interdisciplinary Neurobehavioral and Technology (HINT) Lab at the University of Hawaii at Manoa. Professor Minas received a Master's of Business Administration from Indiana State University and Bachelors of Science in Psychology with Neuroscience from Vanderbilt University.