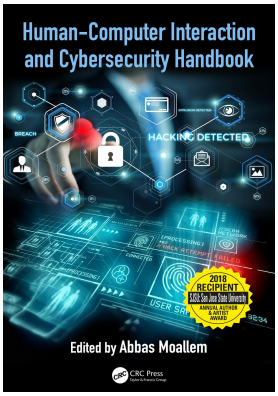

Human-Computer Interaction and Cybersecurity Handbook

Abbas Moallem

Trust

Publication details
https://test.routledgehandbooks.com/doi/10.1201/b22142-5
David Schuster, Alexander Scott
Published online on: 24 Oct 2018

chapter five

Trust

David Schuster and Alexander Scott

Contents

5.1 Introduction

"Trust is good, but control is much better," attributed to Vladimir Lenin [1], suggests that being trusting means being vulnerable. Rather, trust is an adaptation to an uncertain, risky situation; humans apply trust to make decisions and minimize risk. At present, cybersecurity occurs in a context characterized by high risk, uncertainty, time pressure, and an almost inconceivable number of agents potentially affecting the security of a network. This environment challenges cybersecurity professionals, who defend organizations

against threats, and consumer users of the Internet, who are willing or unwilling participants in their own personal and organizational security. These challenges will only increase as computer networks, and attacks on computer networks, grow in their size and sophistication.

Through their decision-making, users can mitigate or exacerbate threats. Human decision-making is critical to cybersecurity across roles and contexts. Poor decisions that affect security outcomes are the core reason why users are cited as the weakest link in security [2] or why social engineering is often an easier vector for attacking an organization than through electronic means [3]. Symantec's report on cyberthreat trends [3] found that malicious actors are increasingly utilizing social engineering tactics, citing effectiveness and ease of use. In this chapter, we focus on how trust affects decision-making, and ultimately, information security. By seeing trust as an adaptive process that can affect decision-making, we can better understand why social engineering works. The problem is not that humans trust, but that trust can be misplaced; we argue that designers should encourage the development of *appropriate* trust, which facilitates good decision-making. Further, we will describe how trust can be incorporated into a user-centered design process.

While human–computer interaction (HCI) in cybersecurity applies to an operator and a computer, interactions among members of a team, or between an individual and an attacker, also affect security outcomes. Complicating matters, many interactions between individuals are computer mediated. Sometimes, these interactions take the form of computer-mediated communication, as in sending an e-mail. Other times, an interaction could be one person's observation of another's behavior observable only via the computer network. For example, a cyberdefender may watch for changes in the dashboard of a security tool to determine whether an attacker was successful at infiltrating the network.

We offer these examples to distinguish cybersecurity, in its size and complexity, from other domains of interest to HCI. Methods traditionally used to understand less complex interactions, such as a heuristic evaluation of an application, are still used in cybersecurity domains. However, the increasing capability of automated tools in other domains has blurred the lines between people and computers in some ways. While continuing increases in computer processing power, network bandwidth, and storage space drive some of these improvements, cybersecurity is a challenging societal problem of unprecedented scale. The expected cost of data breaches by 2019 is $2.1 trillion [4]. The need for improved cyberdefenses by organizations will also drive the increasing sophistication of cybersecurity tools targeting cybersecurity professionals and other users at work and in the home. As the complexity of these tools grows, the interactions between users and the sociotechnical system may resemble interpersonal interactions in some ways.

5.1.1 *Defining trust*

Trust has been defined differently across disciplines. In this chapter, we take a user-centric definition of *trust* appropriate for HCI. From our perspective, *trust* is the "willingness to be vulnerable to the actions of another party" [5]. Our daily lives, professional and personal, are characterized by many decisions that are based on trust in others. We trust cashiers to accurately total our purchases, firewalls to remain active, and other vehicle drivers to also follow traffic laws. Trust is one way we adapt to an inherently risky and uncertain world that requires time-critical decision-making. Under these circumstances, we could not use our computers if we required certainty to proceed. Trust allows us to accept and limit risk in these situations despite uncertainty and dependence on others [5].

The application of trust is not universal across all people, computers, and situations. Trust in another person is called interpersonal trust while trust in a nonhuman agent is called trust in automation [6]. While the word *trust* dates to the year 1200 describing a general concept of reliance [7], trust has been applied to technology relatively recently. The application of trust to technology has led to a new construct in the literature. *Trust in automation* has been defined by Lee and See [8] as "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability." Trust in automation is particularly relevant to HCI in cybersecurity because untrusted technology might not be used; the recommendations of trusted technology are followed more than untrusted technology [8]. Empirical research has consistently shown that trust in automation predicts decisions to use tools [9–11] such that people are more likely to use and follow the guidance of technology that they trust [12].

To date, research has identified both differences and similarities between interpersonal trust and trust in automation. A theoretical explanation for differences in trust in automation versus people is that automation lacks intentionality, a quality of thinking being directed toward a known goal [8]. Machines execute instructions but do not consider goals. Our interactions with other people can benefit from social cues, such as symmetry of trust, which occurs when each partner understands how they will be perceived by the other member [13].

Interpersonal trust and trust in automation both suffer when the trustee is unreliable [13]. However, unreliability in machines is perceived differently than unreliability in humans [14]. Humans tend to expect machines to perform predictably, leading to more rapid declines in trust when unreliability is encountered [15]. According to Madhavan and Wiegmann [14], automation trust develops from a schema of perfect automation, making users more aware of errors and heavily basing trust on the perceived performance of the automation. Further, not all automation errors affect trust in the same way; the easy-errors hypothesis says that when people observe automation making mistakes on tasks they consider to be easy, trust falls even if the aid is otherwise quite reliable [16]. For example, if users cannot find a file known to exist using a desktop search feature and then find it manually, they may be disinclined to use the search feature in the future when searching for other files. In contrast, trust in humans develops from a schema of imperfection (i.e., to err is human), making users more forgiving of errors and basing trust on knowledge of the trustee.

An important consideration in HCI in cybersecurity, then, is whether the trustee is a human or a machine. In the cybersecurity context, interpersonal communication is often mediated by a computer instead of taking place face to face; that is, words written by a person are transmitted by computer. For example, crowd-sourced ratings of a product are written by individual people but are aggregated, filtered, and sorted by machines. Consequently, the distinction between trust in a human or machine is likely to be driven by the perception that an agent's behavior is dictated by a human or determined by an algorithm. One present author and colleagues explored this distinction in an experiment manipulating the origin of app security ratings that were presented as being algorithmically generated or human generated. They found that users were more likely to follow algorithmically generated ratings than human-generated ones when both provided similar levels of risk [17].

Despite differences in people's trust of humans versus machines, people tend to anthropomorphize, or apply human qualities to machines [18]. In other words, people sometimes treat technological agents the same way as they would treat human beings [19]. For example, people prefer technological agents that they feel are like themselves [18].

Designers may capitalize on anthropomorphism so that trust in automation more closely resembles interpersonal trust. Automation that supports more naturalistic interactions with people, and possibly attitudes resembling interpersonal trust, is called human–automation teaming.

Human–automation teaming [20] can be thought of as a style of interaction that describes how the nature of HCI will change as automation becomes more capable of participating in decision-making processes and interacting with humans in naturalistic ways. Human–automation teaming includes automation capable of participating in team process, the dynamic coordinating behaviors that humans engage in as part of human teamwork [21]. Human–automation teaming has been applied to the development of military robotics (called human–robot teaming; [22]) and in airspace operations research [23]. In both domains, the technological capability that would support human–automation teaming is nascent. This is similarly true in cybersecurity, but advances in machine learning and artificial intelligence applied to big data will lead to new methods to detect security threats. The continuous and rapid changes to the threat landscape pressures solution developers to leverage big data so that previously unseen or emerging threats can be detected. It is unlikely that such tools will be able to operate independent of human intervention. Cybersecurity professionals will still need to make decisions about identified threats and bridge between security outcomes and organizational impact to mitigate risk. For example, if a cybersecurity tool based on machine learning provided recommendations to a human operator for mitigation actions while describing the available evidence for the suggested decision with a chat interface, the operator would be able to hold a conversation in natural language to further explore the evidence for the threat along with possible mitigation strategies. Subsequently, with a more holistic view of the situation augmented by the security tool, a human operator would synthesize the machine-derived intelligence with their own knowledge of the context, such as whether the threat affects a system critical to business operations, to make an action decision. Interactions that support human–automation teaming further blur the line between interpersonal trust and trust in automation. If human interactions with automation begin to resemble human interactions with humans, the differences between trust in humans and trust in automation may decrease. However, there is a need for more research; research on trust in automation has so far focused on relatively simple interactions, decision selections, and information analysis in domains such as combat identification, decision aids, monitoring, visual inspection, route planning, and collision warning [24].

Trust is an intuitive concept and still emerging as a measurable construct of interest in HCI. Following a review of current theoretical perspectives on trust, we apply trust constructs to two user populations: cyberdefenders and consumer users of the Internet.

5.1.2 Models of trust

Trust reflects a relationship between an entity doing the trusting (the trustor) and the entity being trusted (the trustee). Trust implies that the trustor is invested in an outcome or goal and that there is a possibility of failure [25], which provides risk. In organizational contexts, goals are layered and may reflect individual, team, or higher-level goals [5]. Trust has been distinguished from similar constructs such as prediction, which also reduces risk; cooperation, which can occur with or without trust; and confidence, which can occur independently of decision-making [5].

Schaefer, Chen, Szalma, and Hancock [18] described trust as a relationship with three components. The first is propensity to trust, a relatively stable trait of the trustor. Propensity

to trust is the initial likelihood to trust in an entity before having any experience with it [12]. Propensity to trust is a "generalized expectation about the trustworthiness of others" (p. 715) and reflects a willingness to trust [5]. Individuals with a higher propensity to trust may be more likely to maintain trust than individuals with less propensity to trust [26].

The second property is a situationally defined state of trusting. The third property reflects that trust changes over time, and the change over time is a property of the interaction between the trustor and the trustee. Both the trustor and trustee affect the bidirectional trust relationship. Trust is adaptive when it is justified, when there is a match between the trustor's trusting and the trustee's trustworthiness [6].

The bidirectional nature of the trust relationship presents challenges to our understanding of trust. Trust is both a predictor of successful human–technology interaction and an outcome of human–technology interaction. Trust is determined by properties of the trustee, suggesting terms such as *trustworthiness*, *dependability*, and *reliability*. But trust is also determined by properties of the trustor, such as propensity to trust and the trustor's understanding of the technology [24].

Another continued challenge frequent in the scientific study of HCI is the complex nature of cognition in human–technology interactions in operational environments. It is difficult to identify the difference between, and relationship between, dynamic processes and outcomes. In operational environments, trust is employed within complex and parallel decision-making processes with no clear start or end [6]. History-based trust changes over time as a user gains experience with automation and varies across entities being trusted [12]. Another research question is whether the loss of trust is distinct from the building of trust, with some evidence that these are separate constructs [27].

Although a research need remains, it is established that trust is a dynamic attitude held by the trustor, and the impacts of trust are observable in the behavior of the trustor over time. In HCI, we are primarily concerned with the trustor as a user or customer. Trust impacts the quality of an interaction across performance, safety, and user satisfaction outcomes. At the most surface level, trust may provide an explanation for user outcomes.

5.2 How trust affects decision-making

Users affect cybersecurity through their decision-making. *Decision-making* is defined as the selection of an option in a situation with some ambiguity or risk, provided the decision takes longer than 1 s to make [28]. This definition distinguishes decision-making from faster reactions to perceptual stimuli. Decision-making is critical to successful security outcomes, both for cybersecurity professionals and consumer users.

As a process, decision-making varies depending on the decision maker's strategies, experience, and resolution of conflicts, which reflect uncertainty [29]. In their integrative decision-making model, Lehto, Nah, and Yi [29] identified the conflicts resolved through decision-making (listed in Table 5.1).

Generally, decision-making can be understood as a process that seeks to resolve one or more of these conflicts. For example, conflicting objectives are resolved through judgments about the importance of objectives, leading to prioritization as an outcome. How human decision makers perform this process has been extensively studied, resulting in two broad approaches to modeling the cognition behind human decision-making: (1) normative models of decision-making and (2) descriptive models of decision-making. While both are useful and informative, neither provides a comprehensive model of how external information predicts decision-making.

Table 5.1 Conflicts resolved through decision-making

Lack of consensus
Uncertain consequences
Uncertain preferences
Conflicting objectives (bad consequences)
Uncertain aspirations
Need to compare alternatives
Unidentified conditions, alternatives, or consequences

Source: Lehto, Mark R., Fiona Fui-Hoon Nah, and Ji Soo Yi. *Handbook of Human Factors and Ergonomics*, John Wiley & Sons, Hoboken, NJ, 191–242, 2006.

Normative models, with origins in behavioral economics, reflect an information processing approach that maximizes utility. That is, people acquire information relevant to a decision; weigh the quality, importance, and reliability of this information; and select the option that maximizes their utility. Decades of research have demonstrated that normative models do not sufficiently predict how people make decisions in operational contexts. Decision-making in a naturalistic environment is not a linear process of weighting and synthesizing complex decision criteria [30].

Descriptive models of decision-making explain how decisions are made in context. A major feature of descriptive decision-making is satisficing, which is choosing the option that is good enough, even if it is not the optimal decision. One reason that normative models do not hold across contexts is that the amount of information required to make a normative decision grows exponentially as more factors are considered and as more options become available. Humans are adapted to be efficient in their cognition [31]; that is, we take shortcuts in our mental processing of information. Decision heuristics are an example of a descriptive approach to decision-making [32]. Decision heuristics are shortcuts that reduce the information-processing load by selectively disregarding some information that could inform a decision [33]. Heuristics can be adaptive in that they support faster decision-making with acceptable outcomes much of the time, especially when employed by experts, who are able to focus on the most relevant cues [34]. For example, a cyberdefender may use the representativeness heuristic to match critical cues in an investigated event to an attack observed earlier rather than considering the probability of the attack, which may be rare or difficult to identify.

Together, normative and descriptive models of decision-making show that humans are rational information processors willing and able to shortcut this process when some uncertainty is difficult or impossible to resolve. This process is adaptive to decision makers in a complex world, as is trust. Understanding the impacts of decision-making, and being able to predict user decisions, is critical to HCI in cybersecurity.

Cyberdefense, by both professionals and consumer users, is characterized by the use of multiple tools to defend against and respond to threats. Professionals and consumer users face a multitude of security-critical decisions about which tools should be used under what circumstances and what actions to take because of alerts from these tools. For example, a warning about an expired certificate suggests to a user that they should not continue to the intended site. From a normative decision-making perspective, the user must gather and appropriately weigh the relevant evidence to decide if this warning is spurious or consequential and then take appropriate action. However, consumer users often lack the expertise necessary to interpret security warnings [35,36], and therefore, a normative decision-making model is unlikely to predict the user's decision. Further, these

decisions are made quickly and are not part of the user's primary task. As a quick decision characterized by uncertainty, trust may determine whether the user will follow the recommendations of the warning or not.

Trust is useful for describing or predicting other decision-making factors, especially those that are difficult to include in a normative decision-making model. For example, while the reliability of a tool predicts trust in it, trust may be inertial [37]. That is, occasional false alarms or misses by an otherwise reliable tool might not reduce the use of the tool. Further, trust in a security tool can affect use in two ways. Dixon and Wickens [38] demonstrated that user compliance and reliance can be affected independently. Compliance is the operator's use of automation when it has presented a signal, such as when security software indicates that a threat is present. Reliance is the use of automation when no threat signal has been presented, as occurs when security software indicates no threats. People separately consider the potential for misses, when automation fails to detect a threat, and false alarms, when automation falsely suggests that a threat is present. Trust can affect compliance, reliance, or both.

5.3 How to incorporate trust into design for security in HCI

Trust can be incorporated into the design of cybersecurity tools in both research and practice contexts. Trust predicts and explains user decisions, making it a useful variable in research studies. Designers may also increase user satisfaction and security when they ensure that tools are appropriately trusted. In this section, we propose strategies that designers can employ to address challenges related to trust in cybersecurity tools.

5.3.1 Challenge 1: Either too much or too little trust can be problematic

A challenge in designing for trust is that the designer must design for *appropriate* trust. We have described how trust is an adaptive process to facilitate decision-making in cybersecurity; trust is something to be attended to in design, not something to be avoided. It is adaptive, and when trust is placed in trustees who deserve to be trusted, better decisions can be made with less time and effort. Trust is based, in part, on observations of the behavior of targets of trust. When trustors observe potential trustees acting reliably and transparently, their trust increases. Similarly, trust decreases when potential trustees act in unpredictable ways or seem to make mistakes. This is not a perfect process, however; trust can become miscalibrated when trustors' perceptions of reliability are inaccurate. When trustors award less trust than deserved to those who would be trustees, called distrust, decisions may be suboptimal. Trustees, both human and technological, can be underutilized if they are undertrusted. This could be one reason that a tool that provides an important warning or alert goes ignored.

More trust does not always lead to better cybersecurity outcomes. Overtrust is an inappropriately applied trust that leads to overreliance on automation. When users overrely on automation, they give too much weight to automation in their decision-making and do not sufficiently anticipate automation failures. Overreliance predicts complacency, a state of insufficient monitoring of the output of automation [39]. Complacency may go undetected if automation only sporadically fails or misses a threat that occurs rarely. Thus, even if overtrust seems to satisfy users in the short term, performance and trust will eventually be harmed when user expectancies are violated, possibly in a security-critical situation. Because of the complex nature of threat defense, cybersecurity tools often provide an incomplete picture or provide results with a limited level of confidence. Thus, overtrust

Table 5.2 Summary of challenges in leveraging trust for security in HCI

Challenge	Proposed solution
1. Either too much or too little trust can be problematic.	Strive for appropriate trust: Identify targets of trust, describe trustworthiness criteria, and ensure that users understand when and why they should trust.
2. Trust is context and user dependent.	Incorporate trust measurement in each iteration to understand the effect of design decisions on trust.
3. Poor usability leads to mistrust.	Identify and match user expectancies to improve usability.
4. Trustworthiness depends on reliability, which may be unknown.	Provide transparency about the reasons for automation failure to support accurate inferences about reliability.

can negatively impact security outcomes when users do not investigate or respond to an alert from a tool.

5.3.1.1 Proposed design strategy

Designers must strive for appropriate trust. To assess appropriate trust, first identify the entities being trusted and identify them as human or technological. Second, describe the criteria under which each entity should be trusted. As a simple example, a firewall may be an entity that requires no monitoring, and thus a high level of trust, under regular operational conditions. If the organization is the victim of an attack, however, verifying the configuration and performance of the firewall may be important. Finally, designers can measure trust and employ strategies to increase or decrease trust that reflect the nature of the entity (human or technological) and the circumstances under which users should rely on (Table 5.2).

5.3.2 Challenge 2: Trust is context and user dependent

Trust, especially in automation, is contingent on the goals of the task or user. Just as performance is the degree to which user goals are achieved, trust in automation is contingent upon a goal [6]. A user might trust a virus scanner to find a known virus but not to monitor their home for a break-in. Consequently, HCI methods that measure or manipulate trust must consider to what degree the trust is task specific and whether findings of trust in automation to do one task may hold as the task parameters vary, especially as changes are made to the design or functionality of the automation. For example, the circumstances under which users trust virus scanners to detect viruses may differ from the circumstances under which users trust password managers to protect their login information.

Designers must also consider how user characteristics affect trust. Factors related to the trustor include emotive factors, cognitive factors, traits, and states [18]. Emotive factors include subjective outcomes such as user satisfaction, comfort, and attitudes toward the trustee. Emotive factors closely tie to user satisfaction outcomes, with the implication that user satisfaction may predict trust (see the following challenge).

Cognitive factors relate to the trustor's understanding of the automation, the ability to use the automation, and expectancies of the automation. These factors are relevant in security contexts where there may be variability in users' expertise.

As user characteristics, states and traits are distinguished by how stable they are within one individual over time. States are dynamic individual differences such as mood. States could affect the trust development, although there is a need for more research specific to

trust. Traits are distinguished by their stability over time. Traits are relatively stable across task situations, including age, personality, and propensity to trust. Traits interact with task characteristics [40], preventing broad recommendations for design based on specific characteristics. However, designers should understand that user traits may interact with automation characteristics to impact trust. Most directly, propensity to trust is a characteristic that users bring to an interaction before trust is affected by the use of an automated tool.

5.3.2.1 Proposed design strategy

Because context matters, appropriate trust is not easily predicted a priori. As with other outcomes of interest to HCI practitioners, such as user satisfaction, an iterative, user-centered design process will lead to the most optimal solution as a product evolves. HCI practitioners can measure trust in each iteration using one or more of the methods described later in this chapter. Specifically, state measures of trust in a specific entity complement measures of propensity to trust.

5.3.3 Challenge 3: Poor usability leads to mistrust

The ability to use automation is closely related to usability, such that highly usable products may be trusted more [41,42]. Hoff and Bashir [24] presented several literature-based design recommendations to increase trust. They suggested that increasing anthropomorphism, creating usable interfaces, adopting a polite communication style, providing accurate feedback about and context for automation failure, and avoiding errors during early interactions or on easy tasks would increase trust. In a security context, this means that poor usability can lead to miscalibrated trust, especially distrust of automated systems.

5.3.3.1 Proposed design strategy

Automation that is compatible with users' expectancies is trusted more, especially when the consistency helps users to predict its behavior [43]. Users have expectancies of the automation, based on their understanding of the relationships among relevant system concepts [18,44]. When automation matches expectancies, it is trusted more. Trust is harmed when the behavior of automation is unexpected or incompatible with the internal representation held by the user. Assessments of user expectancies can be used to suggest modifications to automation behavior.

5.3.4 Challenge 4: Trustworthiness depends on reliability, which may be unknown

In security design, automation reliability and predictability are engineering challenges, especially in security technologies that use behavior-based methods to respond to novel threats. In contrast to signature-based methods, behavior-based methods identify anomalous user or software behavior that might indicate a threat or differ from expected behavior [45]. If security technology allows for the detection of previously unknown threats, then it may be impossible to quantify the probability of future threat detection. Consequently, HCI practitioners face a challenge when implementing an interface for an aid with limited reliability. Because trust is subjective, the users' subjective perceptions of automation reliability and predictability may be just as important as the true reliability of the tool. However, users make attributions of the reliability of technology in their interactions with it over time [12], so attempting to mislead users about reliability is unlikely to be effective. Not all errors have the same effect, however. High false alarm rates in cybersecurity tools

may lead to reduced compliance with alerts. While research has suggested that operators will devote additional attentional resources to automation with high miss rates [38], this may not be possible for cybersecurity professionals who perform network defense; the high workload due to the volume of network traffic that must be investigated could limit the attentional resources available to compensate for an aid that misses threats. Further, security solution vendors may be motivated to adopt a liberal criterion because of the criticality of catching threats.

5.3.4.1 Proposed design strategy

In professional environments, providing transparency about the reasons and types of automation failure is likely to be useful [21,46] if cybersecurity professionals understand the feedback provided. That is, more detailed feedback is helpful if it is compatible with users' understanding and does not put excessive cognitive demands on the user. For interfaces designed for consumer users, the minimal understanding of how a tool works may challenge the building of appropriate trust and limit the usefulness of transparency if such feedback is not understandable to users. This is a challenge in designing tools for widespread use, but so too is the technical problem of automation reliability.

5.4 Three ways to measure trust

Several measures have been published that can be used to measure trust in HCI. Trust is commonly measured using self-report. Jian, Bisantz, and Drury [47] developed a subjective trust in automation scale known as the Checklist for Trust between People and Automation. The survey asks users to rate the intensity of their feeling of trust across 12 items scaled from 1 to 7. Five items are reverse coded. This survey focuses on trust in a specific entity, such as a software application, so it is suitable as a subjective trust measure in usability testing. It can also be applied to a variety of automated tools without needing to adapt the measure. However, professionals who use this instrument should take care to ensure that participants understand what entity is being assessed. This is especially important in the usability testing of complex systems in which the trust of only one agent, such as an intelligent assistant, is being evaluated. Practitioners may also consider the contingent nature of trust in automation. The capability of automation affects trust, but this capability is relative to the users' goals. For example, automotive cruise control would have high capability to maintain a set speed, but it would have low capability to autonomously drive a car across an intersection.

Rotter's Interpersonal Trust Scale [48] has been used as a measure of *propensity to trust*, but its items are not specific to automation. The scale has 25 items; 13 are positive statements about trust, leaving 12 items as reverse-coded negative statements about trust. For example, one item asks participants to rate their agreement with the statement, "Parents usually can be relied upon to keep their promises" [48, p. 654]. As we described earlier, empirical research has supported theoretical differences in the way people build interpersonal trust and trust in automation [14]. However, these differences are at least related to automation capability and, at most, will lessen as human–automation teaming becomes a viable interaction paradigm.

An automation-specific measure for the propensity to trust trait is the Automation-Induced Complacency Potential Rating Scale [49]. This 12-item survey asks users to rate their agreement on statements of trust in automated devices generally. For example, one item asks whether participants agree that medical automation saves time and money in the diagnosis of disease [49]. The measure was demonstrated to have high internal consistency

(a = 0.90) and test–retest reliability measured after three months (a = 0.87) [49]. A present limitation of this instrument is that the automated systems encountered in daily life have changed since the measure was published. Items about making purchases electronically rather than over the phone, or recording television shows using a videocassette recorder, may not be diagnostic today. Modifying this measure to include current technology could provide a better measure of propensity to trust at the cost of losing the ability to compare scores to researchers using the original measure.

5.5 *Applications of trust in automation for cybersecurity professionals*

5.5.1 *What defines a cybersecurity professional?*

Cybersecurity professionals are a diverse group of individuals who are responsible for ensuring the ongoing security of the computer networks in their organization. As a whole, information security professionals are in great demand; in 2014, Cisco estimated a shortage of more than a million information security professionals worldwide [50]. Critically, commonalities of knowledge, skills, and attitudes across job roles and organizations are only starting to be defined. One reason for this is the rapid evolution of cyberthreats and cyberthreat actors. Of primary focus are cybersecurity professionals who "protect, monitor, analyze, detect and respond to unauthorized activity," a task called computer network defense (CND) [51]. Because of the large and growing volume of network activity, the unaided performance of this task is impossible in large organizations. To reduce the human information processing requirements, automated tools are used. One example is an intrusion detection system (IDS), which examines server log files to find patterns associated with anomalies. When such a pattern is found, cybersecurity professionals can be alerted to investigate. However, IDSs are limited in their sophistication and reliability; this has been true for most forms of automation for CND. Because of this, CND is a joint human–machine collaborative task in which people depend on automated tools to perform their jobs but must remain in the loop as an information processor and decision maker. Consequently, the cybersecurity professional is a critical line of defense in CND. Effective human decision-making is a determinant of successful cybersecurity.

To address the high and growing demand for cybersecurity professionals, the US National Institute of Standards and Technology has led the development of the National Initiative for Cybersecurity Education Cybersecurity Workforce Framework (NICE NCWF) to describe the work of cybersecurity professionals. Five core functions underlie cyberoperations across an organization: identify, protect, detect, respond, and recover. Seven high-level categories are defined that span cyberdefense, intelligence, and forensics across an organization and include two or more of the core functions. These categories are listed in Table 5.3 [52].

Applications of trust are most relevant to the job roles that most require automated tools: protect and defend, analyze, collect and operate, and investigate. These job roles all have an intelligence activity that relies on the synthesis of large amounts of data. At present, these activities typically require the use of a suite of tools, each with a limited purpose [53]. In some cases, information from multiple tools may be integrated into a dashboard interface. Increasingly, cyberdefense tools leverage the vast amounts of data collected across the network of the organization to participate to a greater degree in decision-making.

Table 5.3 NCWF workforce categories and their definitions

Secure provision	Conceptualizes, designs, and builds secure information technology (IT) systems, with responsibility for aspects of systems and/or networks development
Operate and maintain	Provides the support, administration, and maintenance necessary to ensure effective and efficient IT system performance and security
Oversee and govern	Provides leadership, management, direction, or development and advocacy so the organization may effectively conduct cybersecurity work
Protect and defend	Identifies, analyzes, and mitigates threats to internal IT systems and/or networks
Analyze	Performs highly specialized review and evaluation of incoming cybersecurity information to determine its usefulness for intelligence
Collect and operate	Provides specialized denial and deception operations and collection of cybersecurity information that may be used to develop intelligence
Investigate	Investigates cybersecurity events or crimes related to IT systems, networks, and digital evidence

Source: National Institute of Standards and Technology. *NICE Cybersecurity Workforce Framework*. Last modified July 7, 2017. https://www.nist.gov/itl/applied-cybersecurity/nice/resources/nice-cybersecurity-workforce-framework.

5.5.2 *Cybersecurity professional characteristics*

Cybersecurity professionals develop their expertise over time as they engage in goal-directed practice. For example, this knowledge might include specific threats and vulnerabilities or maintain awareness of the organization's local area network/wide area network pathways [54].

Because network attacks can occur within milliseconds, and cybersecurity professionals' decisions are critical to the security of the organization, their work is categorized by uncertainty and time pressure. Miscalibrated trust is problematic at both ends. If a tool is capable of informing security decisions but the decision maker does not trust it, the tool is distrusted. Distrusted tools may be underutilized, increasing threat vulnerability when automation that could detect or mitigate a threat is ignored or disabled. Further, the workload of defenders increases when automation is not given the authority to manage detection tasks it is better suited for than a person. On the other end, overtrusted tools may lead to complacency, an insufficient checking of the results of the tool. This could manifest in a variety of outcomes. For those responding to low-level alerts, there may be insufficient research before alerts are escalated, causing overload at higher levels, or the misses of a tool may go undetected when it is overtrusted to just work without intervention. Fortunately, designers can design for appropriate trust by understanding how their design changes are more or less trusted by their cybersecurity professional users. Because trust predicts compliance and complacency, it should be measured as an outcome as part of the user-centered design process. Doing so can identify problematic levels of trust, or unanticipated changes in trust, as a design evolves.

Designers of tools for cybersecurity professionals must ensure that trust is appropriately calibrated. Cybersecurity professionals must understand the true reliability of the tool. In many cases, the reliability cannot be expressed as a simple percentage (e.g., "this IDS misses 10% of threats"). The circumstances surrounding automation failure in cybersecurity may be both complex and unpredictable. For example, an attacker may directly target a specific automated tool to knock it offline or have it provide inaccurate output. Designers should communicate these circumstances to cybersecurity professionals by way of the interface or by training.

Table 5.4 Recommendations to improve appropriate trust in design for cybersecurity professionals

1. Measure trust as a user-centered design outcome.
2. Provide evidence of the known reliability of automated tools whenever available. Provide transparency about the reasons for automation failure when reliability is unknown.
3. Ensure that the functionality of tools is understood through mental model elicitation techniques.
4. Incorporate human–automation teaming to facilitate trust calibration.

Cybersecurity professionals have a greater capacity for understanding complex reasons behind automation failures than consumers. However, the cybersecurity professional's knowledge of the function of a tool, and the circumstances for its failure, must match reality to build appropriate trust. Mental models, which describe structural knowledge, are particularly relevant. Mental models are "the mechanisms whereby humans are able to generate descriptions of system purpose and form, explanations of system functioning and observed system states, and predictions of future states" [44, p. 351]. Originally described as an internal representation of external reality by Craik [55], mental models support trust by providing a framework for understanding, interpreting, and integrating environmental cues [56]. Mental model elicitation techniques, such as concept mapping, can be used to evaluate how a cybersecurity professional understands the relationships between a new tool, existing tools, and the threat landscape (see Crandall, Klein, and Hoffman [57]).

Finally, the concept of human–automation teaming offers promise to facilitate trust calibration by incorporating naturalistic communication and team process in a way that allows humans to better leverage metaphors of interpersonal interaction. For example, a future automated tool may be able to explain why it made an error, minimizing the loss of trust that is likely in present automation (Table 5.4).

5.6 Applications of trust among consumer users

An understanding of how trust impacts use and performance has implications in consumer cybersecurity applications. Researchers in the cybersecurity domain have begun to investigate factors influencing trust; their research has shown that consumers who are not cybersecurity professionals lack security knowledge and frequently overtrust the capabilities of their security software, which lessens their computer security [58]. An illustration of the state of users' cybersecurity knowledge can be found in a 2017 McAfee survey. The study found that 41% of their users were not aware of how to check whether their devices are compromised [59].

Considerations of trust should be used by cybersecurity product designers to design products that foster appropriate user automation trust. As with many HCI guidelines, we offer solutions to leverage changes in the design as a first-line intervention. Failing that, leveraging the characteristics of users, or modifying user behavior through training, could be an option. Compared to cybersecurity professionals, however, consumer user tool use is highly discretionary, occurs with low frequency, and is characterized by mental models that may not reflect deep or accurate knowledge of security technology.

5.6.1 Increase trust in features that promote cyberhygiene

By fostering appropriate trust through interface design, researchers have sought to promote security-conscious user behavior and subsequently bolster the overall security of their

systems. Through the manipulation of message framing, length, and the use of an anthropomorphic messenger, Rodríguez-Priego and van Bavel influenced user behavior on a consumer shopping website [60]. The researchers found that by manipulating the presentation of security warnings, they could manipulate the trust and the resulting behavior of site visitors. The researchers found that a lengthier message paired with a male anthropomorphic character led to safer online behavior. They also found that a loss-framed security message led users to engage in safer behavior. The loss-framed message emphasized what users stood to lose from cybersecurity threats rather than a gain-framed message emphasizing what users stood to gain by staying safe. This research demonstrates that design can lead to proactive behaviors, which can improve security for all through a reduction of attack surface. The concept of users proactively participating in their own security has been called cyberhygiene [61]. Cyberhygiene depends on the appropriate trust in tools that support cyberhygiene. Users may undermine their own security by underutilizing or circumventing security technologies that they do not trust to work toward their goals. For the security of users, it is imperative for the design to support trust in features that support cyberhygiene behaviors.

5.6.2 Design for user diversity

Defined by their stable nature, traits include dimensions such as age, gender, ethnicity, and personality [15,40,62,63]. As an illustration to the applicability of traits to the design of cybersecurity systems, we can consider a company that provides products to both Mexico and the United States. Hoff and Bashir [24] found that compared to their counterparts in the United States, Mexican users were more likely to trust automated systems. Thus, designers may need to design their product in a way that accounts for diversity in trust across various segments of the user population. More broadly, differences in traits and their resulting effects on trust should be represented in designs by increasing the individualization of system designs aimed toward different population segments. Indeed, designing around the needs and traits of individuals has been suggested by researchers as an effective way of improving system performance [64].

5.6.3 Mental model compatibility

Researchers have found that matching an interface with the mental model of their users increases the credibility and subsequent development of trust in the system [65,66]. Conversely, a mismatch between the mental model of the user and how technology operates can cause distrust in the system. Thus, practitioners should consider exploratory usability testing with the purpose of eliciting and describing the mental models of target users relevant to the product.

5.6.4 Incorporating elements of human–automation teaming

As with cybersecurity professionals, human–automation teaming may support appropriate trust in consumer users as interactions become more natural. Teaming may start to be incorporated in consumer-facing products in two ways: through naturalistic communication and adaptive coordination with people.

5.6.4.1 Communication

The reliability and consistency of a system over time have a significant effect on user trust. While interface designers might not be able to improve those metrics, implementing

Table 5.5 Recommendations to foster appropriate trust in consumer users

1. Identify features that support cyberhygiene and ensure they are usable and trustworthy.
2. Design for user diversity in trust.
3. Assess mental model compatibility to ensure that security technology is understandable.
4. Incorporate coordination and naturalistic communication features to facilitate trust calibration.

effective system–user communication through clear, appropriate, and accurate feedback to users can mitigate performance lost through poor reliability and inconsistent performance [67].

The appearance and sound of an interface have been correlated with likeability and subsequent trust. Specifically, more anthropomorphism has been shown to lead to greater trust [68]. Users tend to trust human speech over synthetic speech for the presentation of information [69]. Additionally, research by Parasuraman and Miller [70] found that a polite interface positively impacted trust development. These findings suggest that by utilizing more familiar and natural communication modes, and by increasing the politeness of communication, systems become more trustworthy. Thus, system designers should strive to design interfaces that communicate with their users in a polite manner that imitates human speech patterns and content.

5.6.4.2 Coordination with users: Adaptive automation

The level of the involvement of a tool in a task, and authority to act, is an antecedent to trust [26,71]. Automated tools can more effectively coordinate with users by providing the information users need at the time that they need it, called adaptive automation [72]. Users show more trust in systems that collaborate with them and account for their state [73]. In a study using a driving simulator, Cai and Lin [74] tested a form of driving automation that took control based on a function incorporating the criticality of the situation along with the user's present level of cognitive engagement as assessed using an eye tracker. They compared an adaptive condition to both strong involvement (i.e., automation taking action) and weak involvement (i.e., automation suggesting action) and found that the adaptive level of automation resulted in higher trust levels and was preferable to users than either fixed levels of involvement. Thus, interaction design for appropriate trust in security tools should facilitate collaboration; tools should adapt to the perceived state of the user. For example, security software should avoid alerts that are not time sensitive during periods of high user workload (Table 5.5).

5.7 Conclusions

Trust is an adaptation to uncertainty and risk. Cybersecurity is a domain with many participants who have diverse perspectives and interests, and participating in cybersecurity as a professional or consumer user means navigating a dynamic, risky, and uncertain environment. When trust is calibrated appropriately, it allows trustors to minimize risk and move forward with decision-making. When trust is not appropriately calibrated, however, security is likely to suffer. In describing trust and offering initial recommendations for its use in user-centered design, we hope that designers, researchers, and HCI practitioners will incorporate this construct as an outcome of interest. As cybersecurity tools, for professional and home uses, continue to evolve, additional research will be needed to clarify how trust is built and maintained in cybersecurity. We offer that trust is not something to

be maximized nor minimized, but a feature of human information processing that interacts with other characteristics to explain how people can use tools most effectively and, ultimately, make decisions that maximize security outcomes.

Acknowledgments

This work is based upon work supported by the National Science Foundation under Grant No. 1553018. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. The authors gratefully acknowledge Elizabeth Shallal and Steven Wu for their help in formatting citations.

References

1. Seligman, Adam B. Trust and sociability. *American Journal of Economics and Sociology* 57, no. 4 (1998): 391–404. doi:10.1111/j.1536-7150.1998.tb03372.x.
2. West, Ryan. The psychology of security. *Communications of the ACM* 51, no. 4 (2008): 34–40. doi:10.1145/1330311.1330320.
3. Symantec. *Symantec Internet Security Threat Report: Trends for 2016*. Mountain View, CA: Symantec Corporation, 2017. https://www.symantec.com/content/dam/symantec/docs/reports/istr-22-2017-en.pdf.
4. Juniper Research. *Cybercrime Will Cost Businesses Over $2 Trillion by 2019*. Press release, May 12, 2015, 2015. http://www.prnewswire.com/news-releases/cybercrime-will-cost-businesses-over-2-trillion-by-2019-finds-juniper-research-503449791.html.
5. Mayer, Roger C., James H. Davis, and F. David Schoorman. An integrative model of organizational trust. *Academy of Management Review* 20, no. 3 (1995): 709–734.
6. Hoffman, Robert R., Matthew Johnson, Jeffrey M. Bradshaw, and Al Underbrink. Trust in Automation. *IEEE Intelligent Systems* 28, no. 1 (2013): 84–88. doi:10.1109/MIS.2013.24.
7. Harper, Douglas. Etymology of trust. *Online Etymology Dictionary*. Retrieved on May 22, 2018. http://www.etymonline.com/index.php?term=trust&allowed_in_frame=0.
8. Lee, John D., and Katrina A. See. Trust in automation: Designing for appropriate reliance. *Human Factors* 46, no. 1 (2004): 50–80. doi:10.1518/hfes.46.1.50.3039.
9. Muir, Bonnie M. Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies* 27, no. 5–6 (1987): 527–539. doi:10.1016/S0020-7373(87)80013-5.
10. Parasuraman, Raja, Thomas B. Sheridan, and Christopher D. Wickens. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans* 30, no. 3 (2000): 286–297. doi:10.1109/3468.844354.
11. Sheridan, Thomas B. HCI in supervisory control: Twelve dilemmas. In *Human Error and System Design and Management*, edited by Peter F. Elzer, Rainer H. Kluwe, and Badi Boussoffara, 1–12. Vol. 253. London: Springer London, 2000.
12. Merritt, Stephanie M., and Daniel R. Ilgen. Not all trust is created equal: Dispositional and history-based trust in human-automation interactions. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 50, no. 2 (2008): 194–210. doi:10.1518/001872008X288574.
13. Lewandowsky, Stephan, Michael Mundy, and Gerard Tan. The dynamics of trust: Comparing humans to automation. *Journal of Experimental Psychology: Applied* 6, no. 2 (2000): 104–123. doi:10.1037//1076-898X.6.2.104.
14. Madhavan, Poornima, and Douglas A. Wiegmann. Similarities and differences between human-human and human-automation trust: An integrative review. *Theoretical Issues in Ergonomics Science* 8, no. 4 (2007): 277–301. doi:10.1080/14639220500337708.
15. Dzindolet, Mary T., Linda G. Pierce, Hall P. Beck, Lloyd A. Dawe, and B. Wayne Anderson. Predicting misuse and disuse of combat identification systems. *Military Psychology* 13, no. 3 (2001): 147–164. doi:10.1207/S15327876MP1303_2.

16. Madhavan, Poornima, Douglas A. Wiegmann, and Frank C. Lacson. Automation failures on tasks easily performed by operators undermine trust in automated aids. *Human Factors: The Journal of the Human factors and Ergonomics Society* 48, no. 2 (2006): 241–256. doi:10.1518/001872006777724408.
17. Schuster, David, Mary L. Still, Jeremiah D. Still, Ji Jung Lim, Cary S. Feria, and Christian P. Rohrer. Opinions or algorithms: An investigation of trust in people versus automation in app store security. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*, edited by Theo Tryfonas, 415–425. Springer, Cham, 2015. doi:10.1007/978-3-319-20376-8_37.
18. Schaefer, Kristin E., Jessie Y. C. Chen, James L. Szalma, and Peter A. Hancock. A meta-analysis of factors influencing the development of trust in automation: Implications for understanding autonomy in future systems. *Human Factors* 58, no. 3 (2016): 377–400. doi:10.1177/0018720816634228.
19. Nass, Clifford, and Youngme Moon. Machines and mindlessness: Social responses to computers. *Journal of Social Issues* 56, no. 1 (2000): 81–103. doi:10.1111/0022-4537.00153.
20. Christoffersen, Klaus, and David D. Woods. How to make automated systems team players. In *Advances in Human Performance and Cognitive Engineering Research: Automation*, edited by Eduardo Salas, 1–12. Columbus: Elsevier Science/JAI Press, 2002. doi:10.1016/S1479-3601(02)02003-9.
21. Cuevas, Haydee M., Stephen M. Fiore, Barrett S. Caldwell, and Laura Strater. Augmenting team cognition in human-automation teams performing in complex operational environments. *Aviation, Space, And Environmental Medicine* 78, no. 5 (2007): B63–B70.
22. Ososky, Scott, David Schuster, Elizabeth Phillips, and Florian G. Jentsch. Building appropriate trust in human–robot teams. In *AAAI Spring Symposium: Trust and Autonomous Systems, Palo Alto, CA, March 25–27, 2013*, 60–65.
23. Prevot, Thomas, Nancy Smith, Everett Palmer, Todd Callantine, Paul Lee, Joey Mercer, Jeff Homola, Lynne Martin, Connie Brasil, and Christopher Cabrall. An overview of current capabilities and research activities in the Airspace Operations Laboratory at NASA Ames Research Center. In *14th AIAA Aviation Technology, Integration, and Operations Conference (ATIO), National Harbor, MD, January 13–17, 2014*, 1–20.
24. Hoff, Kevin Anthony, and Masooda Bashir. Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors* 57, no. 3 (2015): 407–434. doi:10.1177/0018720814547570.
25. Ekman, Frederick, Mikael Johansson, and Jana Sochor. To see or not to see: The effect of object recognition on users' trust in automated vehicles. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction, Gothenburg, Sweden, October 23–27, 2016*, 1–4. New York: ACM.
26. Merritt, Stephanie M., Heather Heimbaugh, Jennifer LaChapell, and Deborah Lee. I trust it, but I don't know why: Effects of implicit attitudes toward automation on trust in an automated system. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 55, no. 3 (2013): 520–534. doi:10.1177/0018720812465081.
27. Yang, X. Jessie, Christopher D. Wickens, and Katja Hölttä-Otto. How users adjust trust in automation: Contrast effect and hindsight bias. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, Washington, D.C., September 19–23, 2016*, 196–200. Los Angeles, CA: SAGE.
28. Wickens, Christopher D., John D. Lee, Yili Liu, and Sallie Gordon-Becker. *Introduction to Human Factors Engineering (2nd Edition)*. Upper Saddle River, NJ: Pearson, 2003.
29. Lehto, Mark R., Fiona Fui-Hoon Nah, and Ji Soo Yi. Decision-making models and decision support. In *Handbook of Human Factors and Ergonomics*, edited by Gavriel Salvendy, 191–242. Hoboken, NJ: John Wiley & Sons, 2006. doi:10.1002/0470048204.ch8.
30. Klein, Gary. Naturalistic decision making. *Human Factors* 50, no. 3 (2008): 456–460. doi:10.1518/001872008X288385.
31. Fiske, Susan T., and Shelley E. Taylor. *Social Cognition: From Brains to Culture*. Thousand Oaks, CA: SAGE, 2013.
32. Tversky, Amos, and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Oregon Research Institute Research Bulletin* 13, no. 1 (1975): 141–162.
33. Gigerenzer, Gerd, and Wolfgang Gaissmaier. Heuristic decision making. *Annual Review of Psychology* 62 (2011): 451–482.

34. Garcia-Retamero, Rocio, and Mandeep K. Dhami. Take-the-best in expert-novice decision strategies for residential burglary. *Psychonomic Bulletin & Review* 16 (2009): 163–169.
35. Bravo-Lillo, Cristian, Lorrie Faith Cranor, and Julie Downs. Bridging the gap in computer security warnings: A mental model approach. *IEEE Security & Privacy* 9, no. 2 (2011): 18–26. doi:10.1109/MSP.2010.198.
36. Furnell, Steve M., Peter Bryant, and Andy Phippen. Assessing the security perceptions of personal Internet users. *Computer Security* 26, no. 5 (2007): 410–417. doi:10.1016/j.cose.2007.03.001.
37. Parasuraman, Raja, and Victor Riley. Humans and automation: Use, misuse, disuse, abuse. *Human Factors* 39, no. 2 (1997): 230–253. doi:10.1518/001872097778543886.
38. Dixon, Stephen R., and Christopher D. Wickens. Automation reliability in unmanned aerial vehicle control: A reliance–compliance model of automation dependence in high workload. *Human Factors* 48, no. 3 (2006): 474–486. doi:10.1518/001872006778606822.
39. Manzey, Dietrich, J. Elin Bahner, and Anke-Dorothea Hueper. Misuse of automated aids in process control: Complacency, automation bias and possible training interventions. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, San Francisco, CA, October 16–20, 2006*, 220–224. Los Angeles, CA: SAGE.
40. Szalma, James L., and Grant S. Taylor. Individual differences in response to automation: The five factor model of personality. *Journal of Experimental Psychology: Applied* 17, no. 2 (2011): 71–96. doi:10.1037/a0024170.
41. Atoyan, Hasmik, Jean-Rémi Duquet, and Jean-Marc Robert. Trust in new decision aid systems. In *Conference on l'Interaction Homme-Machine, Montreal, Canada, April 2006*, 115–122. New York: Association for Computing Machinery.
42. Li, Yung-Ming, and Yung-Shao Yeh. Increasing trust in mobile commerce through design aesthetics. *Computers in Human Behavior* 26, no. 4 (2010): 673–684. doi:10.1016/j.chb.2010.01.004.
43. Biros, David P., Mark Daly, and Gregg Gunsch. The influence of task load and automation trust on deception detection. *Group Decision and Negotiation* 13, no. 2 (2004): 173–189. doi:10.1023/B:GRUP.0000021840.85686.57.
44. Rouse, William. B., and Nancy M. Morris. On looking into the black box: Prospects and limits in the search for mental models. *Psychological Bulletin* 100, no. 3 (1986): 349–363.
45. Clark, David, Thomas Berson, and Marjory Blumenthal. *At the Nexus of Cybersecurity and Public Policy*. Washington, DC: National Academies Press, 2014.
46. Dzindolet, Mary T., Scott A. Peterson, Regina A. Pomranky, Linda G. Pierce, and Hall P. Beck. The role of trust in automation reliance. *International Journal of Human-Computer Studies* 58, no. 6 (2003): 697–718. doi:10.1016/S1071-5819(03)00038-7.
47. Jian, Jiun-Yin, Ann M. Bisantz, and Colin G. Drury. Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics* 4, no. 1 (2000): 53–71. doi:10.1207/S15327566IJCE0401_04.
48. Rotter, Julian. B. A new scale for the measurement of interpersonal trust. *Journal of Personality* 35 (1967): 651–665.
49. Singh, Indramani L., Robert Molloy, and Raja Parasuraman. Automation-induced "complacency": Development of the complacency-potential rating scale. *International Journal of Aviation Psychology* 3, no. 2 (1993): 111–122. doi:10.1207/s15327108ijap0302_2.
50. Beliveau-Dunn, Jeanne. Tackling the cybersecurity skills gap. *Cisco Blog.* Last modified January 21, 2014. https://blogs.cisco.com/education/tackling-the-cybersecurity-skills-gap.
51. Computer Network Defense. *Defense Technical Information Center.* Last modified June 10, 2015. http://www.dtic.mil/doctrine/dod_dictionary/data/c/10869.html.
52. National Institute of Standards and Technology. *NICE Cybersecurity Workforce Framework*. Last modified July 7, 2017. https://www.nist.gov/itl/applied-cybersecurity/nice/resources/nice-cybersecurity-workforce-framework.
53. Goodall, John R., and Mark Sowul. VIAssist: Visual analytics for cyber defense. In *IEEE Conference on Technologies for Homeland Security, Boston, MA, May, 11–12, 2009*, 143–150. Waltham, MA: Institute of Electrical and Electronics Engineers. doi:10.1109/THS.2009.5168026.
54. National Institute of Standards and Technology. *The National Cybersecurity Workforce Framework*. 2016. http://csrc.nist.gov/nice/framework/.
55. Craik, Kenneth. *The Nature of Explanation*. Cambridge, MA: Cambridge University Press, 1943.

56. Cooke, Nancy J., Rene'e Stout, and Eduardo Salas. (2001). A knowledge elicitation approach to the measurement of team situation awareness. In *New Trends in Cooperative Activities: System Dynamics in Complex Settings*, edited by Michael McNeese, Mica Endsley, and Eduardo Salas, 114–139. Santa Monica: Human Factors and Ergonomics Society, 2001.
57. Crandall, Beth, Gary Klein, Robert R. Hoffman. *Working Minds*. Cambridge, MA: MIT Press, 2006.
58. Hausawi, Yasser M. Current trend of end-users' behaviors towards security mechanisms. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*, edited by Theo Tryfonas and Ioannis Askoxylakis, 140–151. Cham: Springer, 2016. doi:10.1007/978-3-319-39381-0_13.
59. Davis, Gary. Navigating today's connected world. *McAfee Securing Tomorrow Blog.* Last modified February 27, 2017. https://securingtomorrow.mcafee.com/consumer/consumer-threat-notices/connected-life/.
60. Rodríguez-Priego, Nuria, and René van Bavel. *The Effect of Warning Messages on Secure Behaviour Online*. Luxembourg: Publications Office of the European Union.
61. O'Connell, Mary Ellen. Cyber security without cyber war. *Journal of Conflict and Security Law* 17, no. 2 (2012): 187–209. doi:10.1093/jcsl/krs017.
62. Donmez, Birsen, Linda N. Boyle, John D. Lee, and Daniel V. McGehee. Drivers' attitudes toward imperfect distraction mitigation strategies. *Transportation Research Part F: Traffic Psychology and Behaviour* 9, no. 6 (2006): 387–398. doi:10.1016/j.trf.2006.02.001.
63. Schaefer, Kristin. *The Perception and Measurement of Human–Robot Trust*. PhD diss., University of Central Florida, Orlando, FL, 2013.
64. Hancock, Peter A., G. M. Hancock, and J. S. Warm. Individuation: The $N = 1$ revolution. *Theoretical Issues in Ergonomics Science* 10, no. 5 (2009): 481–488. doi:10.1080/14639220903106387.
65. Beggiato, Matthias, and Josef F. Krems. 2013. The evolution of mental model, trust and acceptance of adaptive cruise control in relation to initial information. *Transportation Research Part F: Traffic Psychology and Behaviour* 18 (2013): 47–57. doi:10.1016/j.trf.2012.12.006.
66. Fogg, B. J., and Hsiang Tseng. The elements of computer credibility. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Pittsburgh, PN, May, 15–20, 1999*, 80–87. New York: Association for Computing Machinery. doi:10.1145/302979.303001.
67. Stanton, Neville A., Mark S. Young, and Guy H. Walker. The psychology of driving automation: A discussion with Professor Don Norman. *International Journal of Vehicle Design* 45, no. 3 (2007): 289–306. doi:10.1504/ijvd.2007.014906.
68. Pak, Richard, Nicole Fink, Margaux Price, Brock Bass, and Lindsay Sturre. Decision support aids with anthropomorphic characteristics influence trust and performance in younger and older adults. *Ergonomics* 55, no. 9 (2012): 1059–1072. doi:10.1080/00140139.2012.691554.
69. Stedmon, Alex W., Sarah Sharples, Robert Littlewood, Gemma Cox, Harshada Patel, and John R. Wilson. Datalink in air traffic management: Human factors issues in communications. *Applied Ergonomics* 38, no. 4 (2007): 473–480. doi:10.1016/j.apergo.2007.01.013.
70. Parasuraman, Raja, and Christopher A. Miller. Trust and etiquette in high-criticality automated systems. *Communications of the ACM* 47, no. 4 (2004): 51–55. doi:10.1145/975817.975844.
71. Looije, Rosemarijn, Mark A. Neerincx, and Fokie Cnossen. Persuasive robotic assistant for health self-management of older adults: Design and evaluation of social behaviors. *International Journal of Human-Computer Studies* 68, no. 6 (2010): 386–397. doi:10.1016/j.ijhcs.2009.08.007.
72. Byrne, Evan A., and Raja Parasuraman. Psychophysiology and adaptive automation. *Biological Psychology* 42, no. 3 (1996): 249–268. doi:10.1016/0301-0511(95)05161-9.
73. Moray, Neville, Toshiyuki Inagaki, and Makoto Itoh. Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied* 6, no. 1 (2000): 44–58. doi:10.1037/1076-898X.6.1.44.
74. Cai, Hua, and Yingzi Lin. Coordinating cognitive assistance with cognitive engagement control approaches in human–machine collaboration. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans* 42, no. 2 (2012): 286–294. doi:10.1109/TSMCA.2011.2169953.