

Common Nonsense about Password Security and the Expert–Layperson Knowledge Gap

Brett M. Frischmann and Alexandria Johnson

INTRODUCTION

Creating and managing passwords is familiar to most of us. We use passwords on computing devices every day. To access social media accounts, check the balance on online banking apps, or send an email, individuals must authenticate themselves, often by logging into a computer system with a username and password. Password creation and management practices require knowledge that may seem commonsensical. Yet common sense about passwords and security is often misleading nonsense. Many people hold incorrect beliefs about what makes a password secure, the nature and origin of security threats, and what practices compromise or conversely strengthen security. Substantial research has shown that user-chosen passwords are highly predictable and follow similar patterns; individuals create weak passwords based on memorability rather than on secureness, reuse passwords, and often use personal information in passwords that is easily discovered or guessed. Yet, less research has been done to understand the origins and persistence of common nonsense about password security. Several research questions arise: What beliefs do users have about creating secure passwords, security threats, and best practices? Where do these beliefs come from? Do users realize that they are selecting easily guessable passwords, and if so, is this intentional? Do users believe that the passwords they create are secure? Are users more focused on convenience and memorability than security? Are individuals aware of password security risks? Do individuals

Charles Widger Endowed University Professor in Law, Business and Economics, Villanova University, Charles Widger School of Law; Affiliated Faculty, The Vincent and Elinor Ostrom Workshop in Political Theory and Policy Analysis, Indiana University Bloomington; Affiliate Scholar, Center for Internet and Society, Stanford Law School; Trustee, Nexa Center for Internet and Society, Politecnico di Torino. J.D. Georgetown University Law Center; M.S., Columbia University; B.A., Columbia University.
J.D. Candidate Duke University School of Law; B.A. Political Science, Villanova University; B.A., Criminology, Villanova University.

GOVERNING EVERYDAY MISINFORMATION WORKSHOP

BRETT FRISCHMANN

COMMENTS BY JACK BRIGHTON



PASSWORD SECURITY and MISINFORMATION

PASsword123

- HOW WE GENERATE
- NATURE OF THREAT
- PRACTICES



COMMUNITIES OF PRACTICE

INTERACTIONS WITH STRENGTH METERS HELP PERPETUATE MISINFORMATION



UPDATING PASSWORDS CAN ACTUALLY HELP HACKERS



KNOWLEDGE CHANGES

WHAT GAPS NEED TO BE ADDRESSED?

DIFFERENT ROLES & INSTITUTIONAL BARRIERS



LACK OF COMMON LANGUAGE IS HINDRANCE



WHERE DO PEOPLE GET THEIR INFORMATION

STARTING FROM 1ST GRADE!

INCENTIVES for SHARING INFORMATION

LOVE/HATE RELATIONSHIP
EMPLOYEE   IT STAFF

PUBLIC SERVICE CREDIT



seeyourwords.com

FIGURE 5.1 Visual themes from common nonsense about password security and the expert-layperson knowledge gap.

not understand these risks or do individuals understand the risks and create weak passwords anyway?

While the gap between lay and expert knowledge about password security is well documented, less well understood are why the gap persists and what are the origins of incorrect beliefs and misguided security practices. In this paper, we examine the gap and identify knowledge governance dilemmas that help explain its persistence. We use the Governing Knowledge Commons (GKC) framework to structure our study and frame the sets of conceptual and empirical questions we investigate. The security knowledge ecosystem is complex. We focus on password security.

We initially framed our study in terms of two knowledge communities – lay persons and security experts, but we learned during interviews with experts that a better structure would differentiate and explore the interactions between three communities – lay persons, professionals/practitioners, and experts. Notably, it may be a stretch to describe lay persons and professionals/practitioners as communities. Actors in these classifications are members of various communities within which password security knowledge is shared and acted upon, but unlike the password security expert community, these various communities are not organized around or otherwise focused on password security knowledge. More important than community definitions, at least for purposes of this study, are the relationships, interactions, knowledge flows, and governance dilemmas among lay persons, professionals/practitioners, and experts.

We employed two primary methods: (i) a systematic literature review to gather information about password security, password behaviors, expert knowledge on passwords, and lay knowledge about passwords; and (ii) semi-structured interviews of password security experts to supplement the findings from the literature review and gain first-hand information and perspectives. We briefly describe each and provide more detail and data in an Appendix available via the GKC repository (an open data portal).¹

Our literature review involved five steps. First, we collected peer review publications, focusing on research articles and user studies. We searched major research databases, including Google Scholar, Science Direct, ACM, and Springer Link. Keywords included variations of password, user security perceptions, password creation behaviors, and password security awareness methods. Second, we screened research articles by assessing if the title, abstract, and full article matched the relevant subject matter. We also ensured that the research literature met five inclusion criteria.² Third, we identified additional publications by examining references of all the articles we had included and repeating the second step. Fourth, we

¹ <https://hdl.handle.net/2142/117212>.

² The article was (i) peer-reviewed, (ii) written in English, (iii) available full text, (iv) involved a user study, survey, questionnaire, and/or interviews, and (v) primarily focused on password security perceptions, behaviors, and practices.

identified additional articles through suggestions by experts during interviews and by drawing from the publication pages of experts we interviewed. The resulting sample size is ninety-seven articles. ($N = 97$). Fifth, we qualitatively analyzed the research literature using both the GCK framework and a basic set of research questions regarding password security beliefs and practices.³

After protocol review by the Villanova University IRB, we conducted eight semi-structured interviews with professionals directly involved in password security and/or usability.⁴ We used the systematic literature review and the GKC framework to provide structure and generate interview questions.⁵ We recorded, transcribed, analyzed, and summarized each of the interviews. Then, once all interviews were completed, we reviewed the transcripts again. We summarize the results of our qualitative analysis of the literature and interviews below.

BACKGROUND: PASSWORDS AND SECURITY

Basic Overview

Passwords have become a recurring aspect of everyday life in our digital networked world because of the many different devices and online accounts we use daily that require authentication. A 2020 study by Nordpass, a password manager website created by Nord Security, revealed that the average user has around 100 passwords, a 25 percent increase from 2019 (Bobba & Paruchuri 2022).

Passwords are also familiar to most people because of their offline use in other social contexts, for example, as a signal of membership in a club. Passwords take many forms and have a long history. People learn about passwords not only in practice but also in literature. One of our favorite examples, described by Martin Paul Eve (2016) in his book *Password*, is Daedalus's labyrinth, "designed as a spatial-control mechanism for determining the unique identity of a single individual based on knowledge of its topology. . . . For everyone but Daedalus, the labyrinth was supposed to be, quite literally, a death trap." Theseus escaped by circumventing the identification function and thus, according to Eve, "is one of the earliest species of geek that we now would call a hacker or cracker (3–4).

³ See Appendix <https://hdl.handle.net/2142/117212>.

⁴ We initially contacted interviewees with whom we have had professional contact in 2022, and then we added a few experts based on recommendations from the initial interviewees. We interviewed the following experts: Steven M. Bellovin, Zachary Bornstein, Lorrie Cranor, Comac Herley, Arvind Narayanan, Bruce Schneier, David Thaw, and Rick Wash. We provide more detail about the interviewees in the Appendix. After interviews were completed, we consulted Xavier de Carné de Camavalet, who provided written comments on our study.

⁵ We provide more detail, including interview questions and data in the Appendix.

Passwords are a knowledge-based means for authentication.⁶ Knowledge (of a password) serves as a proxy for identification in a system designed to control access to resources or systems. Alternative means of authentication that do not rely on knowledge may be based on “something you have, such as a token or a particular mobile phone [or] something you are, that is, some form of biometric” (Bellovin 2016, 107). Steve Bellovin explains that one must understand the “*systems* nature of authentication,” meaning: “The total environment – who will use it, how you deal with lost credentials, what the consequences are of lack of access or access by the wrong person, and more.” Ultimately, he concludes: “The most important question of all is how people will actually use the authentication technology in the real world” (107).

So how do people use passwords? What do people believe about the passwords they use? Do they understand password strength, security threats, and relevant consequences? These are some of the basic background questions we sought to answer in our review of the literature.

Literature Review and Interviews

Our literature review aimed to answer questions about password security beliefs, where these beliefs originate, how these beliefs are transmitted more broadly, and how the beliefs shape everyday actions and practices. Many articles we reviewed used either surveys alone or surveys with lab experiments as their methodology.⁷ We group our findings into two categories: lay knowledge and expert knowledge.

Lay Knowledge (Dilemmas) Concerning Password Security

We collected data on user beliefs about the nature and origins of security threats, what users believe makes a password secure, and what users believe compromise or strengthen security. Across the articles that focused on understanding users’ perceptions, beliefs, and behaviors, some common practices and beliefs among users were as follows:

Practices:

- choose passwords from a limited set of alpha-numeric characters;
- use names, slang words, dictionary words, or consecutive digits as passwords;⁸

⁶ Every security expert we interviewed responded to the preliminary questions – Why do we use passwords? What problem do passwords solve? – with authentication.

⁷ Many studies recruited subjects attending colleges and universities, a demographic that may not be representative of the larger population.

⁸ A 2021 Nordpass study, <https://nordpass.com/most-common-passwords-list/>, revealed the five most common passwords across fifty countries were:

1. 123456
2. 123456789
3. 12345
4. qwerty
5. password

- use personal information when creating passwords because it makes the passwords more memorable and easier to recall;
- reuse identical passwords across multiple sites.

Beliefs:

- adding numbers makes a password more secure than using only letters;
- adding an exclamation mark or other symbol at the end of a password makes it more secure;
- using a word that is difficult to spell as a password is more secure than an easy-to-spell word;
- a password is secure if the password is longer and contains uppercase letters, digits, and symbols;
- attacks on passwords are targeted at specific individuals;
- using personal information, such as a birthday, in a password is secure if such information is not on their social media accounts;
- password strength meters accurately measure the security of a password.

Misconceptions about password security persist despite growing public awareness of security threats. There are many reasons. We highlight those that surfaced most often and prominently in the literature and interviews.

First, lay people struggle to create and manage an ever-growing list of secure passwords.⁹ In a (somewhat dated) large-scale study of password habits, Dinei Florencio and Cormac Herley (2006) found that on average, people tend to have twenty-five accounts that require passwords, actively use around seven distinct passwords, and reuse passwords across sites/accounts. The researchers concluded that their large-scale study confirmed conventional wisdom about the large number of passwords that users maintain, the poor quality of those passwords, and the frequency of password reuse. In our interview, Herley confirmed that these basic observations from over fifteen years ago remain salient. The number of passwords people create and manage has risen significantly, possibly by an order of magnitude, and people continue to reuse passwords. Herley emphasized that people sometimes exercise common sense effectively when creating, managing, and even reusing passwords. There is a rough and often implicit cost-benefit analysis that people must engage in, given the growing number of password-protected accounts. Contrary to security advice not to reuse passwords often given by experts, Herley told us, “I reuse passwords all the time. I do it without fear, without shame.” Using his own experience as an example, he explained that for 200 or so password-

⁹ Most experts suggested that password managers are a useful solution to this problem. As Bruce Schneier explained: “Password Managers change everything. If I don’t have to remember passwords, and if I don’t have to repeatedly type them, I can choose strong random passwords. Even better, I can have my password manager choose them for me. [For example,] I just asked PasswordSafe for a password. It gave me ‘k%wo=-x{Y_XTpwWz@L.’”

protected accounts, 150 are low-value accounts for which weaker and even reused passwords might be justified. There are just too many accounts to create and keep track of strong, distinct passwords, and if a low-value account is hacked, there is less downside so long as it does not provide access to a high value account or enable another type of attack (e.g., phishing). Thus, one password security dilemma is knowing how to partition among low- and high-value (risk) accounts, and a second one is knowing how to respond when an account is compromised.

Second, lay people tend to have an incomplete understanding of the nature and variety of password security threats, which has changed substantially over the past few decades. The experts we interviewed all emphasized there are many different threats to consider when evaluating password security. Some threats involve “guessing a password,” and these threats are often what people think about when considering what is necessary for a secure password.¹⁰ Password guessing can be as simple as it sounds – a person trying to guess the password. That is a threat model most people can understand. But it can get much more complicated, and correspondingly more difficult for lay people to understand, when the person (guesser/attacker) uses different tools, ranging from surveillance tools (e.g., to collect information about likely passwords)¹¹ to computational tools (e.g., a password-guessing algorithm). Further, it can be difficult to appreciate the differences between attacks directed a specific target and undirected attacks, online versus offline attacks,¹² and other variations of guessing attacks.

To make matters more complicated, many threats to password-protected resources do not involve password guessing at all. For example, a person might look over your shoulder and observe you entering your password. Or a person might trick you into downloading malware onto your computer and that software may capture your keystrokes and thus your password(s). Or a person may hack the system for which a password is required and thereby obtain the password (along with others); if that password had been used for more than one password-protected account (password reuse), then those other accounts could be compromised. Notably, each of these threats occurs without any guessing. While it may be easy for someone to catch the person looking over one’s shoulder and respond to that threat, the other threats require more knowledge for lay people to be able to manage their security. For

¹⁰ According to Herley, “30 or 40 years ago, maybe the dominant threat against passwords was a password guessing attack . . . that hasn’t been the case for a long, long time, and the threats we have against password secure resources tend not to be guessing.”

¹¹ Surveillance tools that enable the person to observe you entering your password obviate the need for guessing and thus fit into the next category we discuss.

¹² In his interview, Arvind Narayanan explained: “Online guessing attack is where somebody tries putting in different passwords into the login screen. An offline guessing attack is where somebody breaks into the server and steals the password database; usually those passwords might be hashed, which is a type of encryption, but it can be reversed.” And so, one relevant but often overlooked password security consideration is “how well does your password resist password hashing?”

example, our literature review and interviews suggested that people tend to underappreciate the risk of password reuse because they fail to understand the external effects from a hacked system. (The hack of system 1 where password A is obtained compromises systems 2, 3, . . . N where A is also used as a password. Lay people who understand password security threats in terms of targeted attacks or guessing attacks may fail to appreciate fully the risk of password reuse.)

Misunderstanding the nature and variety of threats can strengthen incorrect beliefs and lead to poor security practices. For example, according to two studies (Ur et al. 2016a, 2016b) and (Ur et al. 2015) that directly examined user perceptions of password security threats, most users believed that attackers were strangers located far away from them; the users were concerned with attackers having access to and using their personal information. In these studies, users avoided using their own personal information such as birthdays and pet names, but some used the personal information of their family and friends instead to create passwords. The authors concluded that the users' misconceptions about password security directly related to their misunderstanding or lack of knowledge about automated password guessing attacks, which would be able to guess those passwords. Thus, another dilemma to consider is how user (mis)perceptions of security threats shape user beliefs about what constitutes a secure password or practice.

A third reason concerns how lay people learn about password security. Our literature review did not provide much insight on this issue. For the most part, the studies we reviewed aimed to uncover lay persons' beliefs and practices, but the studies did not examine learning processes or the origins of those beliefs. There are some studies that criticize security advice and suggest that lay people reject the advice because it is too difficult to follow given the ever-growing number of passwords to manage. (We discuss these studies below.) The experts we interviewed suggested that people learn through their everyday experiences with passwords (supporting, if not confirming, our hypothesis that password knowledge is often in the realm of everyday common sense knowledge). People engage with password creation regularly, encountering password composition rules and password security meters. In addition to their instrumental function (creating secure passwords), these tools serve an important, generally overlooked, pedagogical function, which is to teach users about password security (see Frischmann & Selinger 2018; Weizenbaum 1976). Unfortunately, these tools can easily mislead and teach the wrong lessons. Consider the following stylized example that we raised during interviews.¹³

Jeremy must create a new password. He encounters a typical password creation interface, a text box to enter characters, password composition rules displayed

¹³ During interviews, we discussed slight variations: **Brett**, **Bretti**, and **Bretti!** The difference in names makes a slight difference because of the additional length (one extra character). We switched to Jeremy for the write-up only because the example better illustrates the potential shift in a strength meter assessment from weak to strong.



Password strength test

This strength tester runs on your local machine and **does not** send your password over the network.

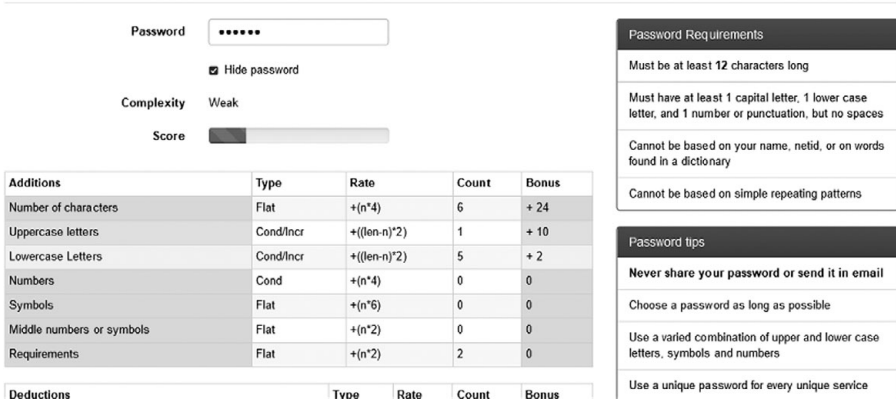


FIGURE 5.2 Password strength test.

prominently to one side, and a password strength meter that updates the color (red, yellow, green) with each new character displayed prominently on the other side.

- Step 1.** Jeremy begins by typing the following characters: **Jeremy**
The security meter remains *red*.
- Step 2.** Jeremy continues by adding the following character: **1**
The security meter turns *yellow*.
- Step 3.** Jeremy continues by adding the following character: **!**
The security meter turns *green*.¹⁴

Figures 5.2 and 5.3 demonstrate the first and third steps using the Password Strength Test, available at the UIC Academic Computing and Communication Center.¹⁵

What might Jeremy learn from this experience? What have these tools taught him about password security? During interviews, the experts noted how there are various ways in which the tools could teach the wrong lessons (i.e., misinformation about password security).

First, the security meter implies marginal security improvements that may not be genuine or meaningful improvements. Adding a single number and/or symbol at the

¹⁴ This is a stylized hypothetical. Most password security meters would not shift from red to yellow to green based on the changes we suggest. Our point is only to illustrate how password meters teach users about password security.

¹⁵ See www.uic.edu/apps/strong-password/.

Password strength test

This strength tester runs on your local machine and **does not** send your password over the network.

Password
 Hide password

Complexity **Strong**

Score

Additions	Type	Rate	Count	Bonus
Number of characters	Flat	$+(n^4)$	8	+ 32
Uppercase letters	Cond/Incr	$+\text{((len-n)}^2)$	1	+ 14
Lowercase Letters	Cond/Incr	$+\text{((len-n)}^2)$	5	+ 6
Numbers	Cond	$+(n^4)$	1	+ 4
Symbols	Flat	$+(n^6)$	1	+ 6
Middle numbers or symbols	Flat	$+(n^2)$	1	+ 2
Requirements	Flat	$+(n^2)$	5	+ 10

Deductions	Type	Rate	Count	Bonus

Password Requirements

- Must be at least 12 characters long
- Must have at least 1 capital letter, 1 lower case letter, and 1 number or punctuation, but no spaces
- Cannot be based on your name, netid, or on words found in a dictionary
- Cannot be based on simple repeating patterns

Password Tips

- Never share your password or send it in email**
- Choose a password as long as possible
- Use a varied combination of upper and lower case letters, symbols and numbers
- Use a unique password for every unique service

FIGURE 5.3 Password strength test.

end of a common name does not substantially improve the strength of the password. In the specific example, **Jeremy**, **Jeremy1**, and **Jeremy!** are all relatively weak passwords. The marginal difference between the first two passwords is negligible; the third is an improvement, but it is not substantial. The differences among the three do not impose much of a burden on automated guessing attacks.¹⁶

Second, in more absolute terms, the security meter suggests **Jeremy!** is a strong password when in fact, it is not. This is an incredibly important source of consequential misinformation. Jeremy, like most lay persons, is likely to trust the digital tool, rely on its “advice” and the weak password, and go no further to create an even stronger password or learn more about password security. How can this possibly occur? In practice, security meters are based on compliance with password composition rules and essentially that means strength depends on checking a series of boxes, such as including characters from different character classes (upper and lowercase letters, numbers, symbols), having enough characters (e.g., length of password is 8, 12, or possibly more characters), and not being on a blacklist of already known common passwords (Carnavalet & Mannan 2014). For most password strength meters, the first two examples in the hypothetical would remain in the red (weak) because they lacked a symbol and would be too short (password length under eight

¹⁶ We checked a few online tools for checking password strength and estimated times for the password to be cracked. While the estimates varied, they all suggested that an automated password guessing program would take seconds to crack the first and second passwords and minutes or hours to crack the third password.

characters); the third example would pass some, however, because it checks those two boxes.¹⁷ To be clear, the point we are illustrating with this hypothetical is not limited to the **Jeremy!** example or even password strength meter tools. Below, we discuss other examples related to password security, but it is important to highlight how this examination is relevant to other digital networked technologies and social dilemmas concerning mis- and dis-information.¹⁸

Third, above and beyond the specific attempt to create a password, are the lessons learned about password security that Jeremy may carry into his next password creation experience or that he may relate to other people. We know very little about these types of effects. None of the studies we reviewed considered them.

Security meters could be designed to be better pedagogical tools. For example, when Jeremy added the numeral 1 to his name, the tool could suggest more secure alternatives using the same addition, for example by placing the number in the middle rather than the end of the string of characters. Regardless of whether Jeremy chose that option, he could learn why it was more secure with just a little explanation. Conversational prompts during password creation, for example, could teach while also nudging users toward stronger passwords (Khern-am-nuai 2022). As Arvind Narayanan suggested after discussing the stylized hypothetical: “Well-designed password meters can both help users create stronger passwords and teach them what makes a stronger password.” Experts have found that password meters with a variety of visual appearances led users to create longer passwords, but significant increases in resistance to password-cracking algorithms were only achieved using meters that scored passwords stringently (Ur et al. 2012). These findings support the push for wide-scale deployment of more stringent password meters to train users to create stronger passwords more routinely. Yet, as we discuss below, this has not happened. Professionals and practitioners who design and deploy password creation tools typically balance security, usability, and market incentives to keep consumers using their services. Further, it is not clear that consumers want to learn about security when they are asked to create passwords, which is often when they are excited to try out a new service.¹⁹ (Again, we must emphasize that these observations are not

¹⁷ See, for example, the Password Strength Test, available at the UIC Academic Computing and Communication Center, www.uic.edu/apps/strong-password/ (weak rating for the first two, but strong rating for the third); Is My Password Vulnerable?, available at <https://nordpass.com/secure-password/> (weak rating for all three).

¹⁸ See generally Frischmann and Selinger (2018), *Re-Engineering Humanity*. See also Haenschen, Frischmann, and Ellenbogen (2020) (manipulating the Facebook birthday notification tool, causing it to distribute fake birthday notifications, to examine the pedagogical function of the tool and how it shapes human thinking).

¹⁹ As Bruce Schneier noted, “people are asked to create passwords at moments when they are not ready to do it. You know how it happens. You’re signing up for some new service, all excited about trying it out. And then you’re suddenly asked to create a password. What you want most is for that screen to go away. You don’t want to think about security. So, of course you do a bad job.”

necessarily limited to password creation tools; we strongly suspect that many other digital tools and interfaces could be designed to be better pedagogical tools but face similar dynamics.)

Another common experience lay people have with passwords is forgetting them. That is, people often forget a password and then engage with password recovery tools, for example, responding to a set of security questions (sometimes called “secret questions” or “challenge questions”) that substitute for the password as a knowledge-based means of authentication (Bonneau et al. 2015a). Password recovery tools also have instrumental and pedagogical functions. People need to access password-protected resources, and the recovery tools serve that need. The security dilemma is two-fold: First, the password recovery tools are not always secure. The types of security questions a person must answer may be easier to guess than a password (Bonneau et al. 2015a). There are more secure password recovery tools (e.g., SMS and email-based recovery; Bonneau et al. 2015a). But again, businesses face tradeoffs between security, usability, and convenience (Bonneau et al. 2015a). Second, since password recovery tools are often easier to use than remembering a strong password, some people (learn to) adopt a strategy of relying on them rather than passwords as a means for authentication and obtaining access to resources. We can see the pedagogical function of password security tools in terms of how experience using the tools shapes everyday security beliefs and practices.

When comparing the behavior of experts to that of nonexperts, Iulia Ion, Rob Reeder, and Sunny Consolvo found that experts report installing software updates, using two-factor authentication, using unique passwords, and using passwords managers to remain safe online, while nonexperts report using antivirus software, using strong passwords, only using known websites, and changing their passwords frequently (Ion et al. 2015). As a result of their findings, Ion et al. suggest that better messaging along with systems and usability work are necessary to get nonexperts to follow the same security practices that experts do.

Expert Knowledge (Dilemmas) Concerning Password Security

In this section, we explain what we learned from the literature review and interviews about expert knowledge, how it has changed over the past few decades, and what are the “best practices” currently recognized by security experts. Our abbreviated history of expert knowledge about password security focuses on major themes related to knowledge dilemmas within the community, rather than specific actors or their contributions. While the literature review provides support for the findings we highlight, most of the insights are drawn from the interviews.

In the early days of computing and networking, password security experts mostly were academic researchers with expertise in computer science, mathematics, and engineering. These experts thought about password security as a mathematical problem. Security experts mostly worried about automated guessing attacks – that

is, guessing attacks leveraging computational tools.²⁰ The offense and defense of password security was, at its core, a computational problem. As such, expert knowledge about password security was based heavily on the mathematical theory of computation. Expert knowledge thus led to certain prescriptions, namely that users should create sufficiently strong passwords to withstand an attack. Experts measured strength in terms of complexity (e.g., Shannon entropy), which usually involved a calculation based on the length of the password and the number of different character classes. For a reasonably accessible and thorough explanation, see *Appendix A: Estimating Entropy and Strength*, in the NIST Electronic Authentication Guideline (Burr et al. 2013b). Notably, as we discuss below, security experts now regard complexity/entropy as a *bad* proxy for guessability and instead prefer to evaluate password strength using an adversarial guessing approach (Lee et al. 2022).

Expert knowledge about password security shifted dramatically over the past few decades, and it continues to change. According to interviewees, during the 2000s, significant changes in expert knowledge emerged as knowledge grounded in theory failed to match reality (everyday life experience). Bonneau et al. discusses the evolution of passwords suggesting that estimates of password strength, models of user behavior, and policies related to password composition work in theory but can be unsupported in practice and even contradicted when observed empirically, possibly leading the research community to focus on the wrong threats (Bonneau et al. 2015b). We highlight two themes that surfaced repeatedly during interviews – first, changes in the technological environment (or reality on the ground), and second, corresponding changes in the expert community. We then turn to the role of professionals and practitioners in the password security ecosystem and explore how these actors mediate between the experts and lay people. We conclude with a brief discussion of the latest expert knowledge regarding best practices and knowledge dilemmas associated with translating such knowledge to actual practice by professionals, practitioners, and lay people.

CHANGES IN THE TECHNOLOGICAL ENVIRONMENT Over the past few decades, digital networked computing technologies evolved considerably and diffused widely. The number of users and the number of accounts each user must manage grew incredibly.²¹ Password-protected resources became part of everyday life experience

²⁰ Password theft, whether by someone looking over your shoulder or more sophisticated forms of espionage, existed as a threat model, but it was less prominent as a threat for academic security researchers to focus on because there were fewer targets (opportunities for such attacks). As noted in the text, this changed with the diffusion of digital networked technologies and proliferation of password-protected accounts.

²¹ “Password strength restrictions go back to about 1979 when a power user might have 3 logins on 3 different computers. And so very few computers, and . . . almost nothing is a service login as opposed to a system login. Users had no local storage and no local computing, and the rules made a lot more sense then. The situation today is completely different. [E]veryone’s got the local storage. My wristwatch has got more storage than the computers I was using in the late

for billions of people. This shift dramatically expanded the attack surface where the offense and defense of password security played out and, as a result, fundamentally altered the epistemic demands on password security experts. Not surprisingly, security threats multiplied and changed. Initially, security experts (mostly) agreed that strong passwords were an appropriate solution for the threat of guessing attacks. But as password security became an everyday life concern for lay people (at home, work, and everywhere else), this conventional expert wisdom began to fail for a few reasons.

First, despite reasonable consensus among experts on what constitutes a strong password (complexity/entropy, though that would change), people nonetheless often chose and still choose terribly weak passwords (see section “Expert Knowledge (Dilemmas) Concerning Password Security”). Leaked datasets of millions of actual passwords, for example the 2009 RockYou.com database leak, confirmed what many considered to be common sense: Lay people were unable to keep up with the computational arms race and ever-growing number of passwords,²² and they often defaulted to weak but memorable passwords and reused such passwords across different accounts. (As Herley and colleagues pointed out, lay people often were exercising common sense, performing rough cost–benefit analyses, and using stronger passwords for more important accounts.) This realization in the mid-2000s led to a shift in focus among security researchers toward the needs and capabilities of actual human beings involved in authentication processes. *Human factors and usability matter!* This realization also focused attention on the increasingly important role of professionals and practitioners who designed, deployed, and managed digital networked technologies, including password creation tools (e.g., human–computer interfaces, password composition rules, strength meters).

Second, automated guessing attacks became more sophisticated and varied. Attackers could use an array of widely available tools, including dictionaries, large datasets of stolen passwords, and even the various password creation tools themselves, to improve their guessing algorithms. Attackers easily could learn about how people tend to respond to password composition rules, for example by adding **!** at the end of a character string deemed weak by a strength meter or by replacing letters with similar symbols (**\$** for **s**). This knowledge advantaged attackers and confounded the Shannon entropy calculation used to determine password strength.

For some threats, such as online guessing attacks made at the login interface, experts determined effective countermeasures, such as limiting the rate at which guesses can be made by injecting exponentially increasing time delays between each

seventies and more compute power. I have hundreds of logins. They're mostly service logins to a website rather than a system logins that would give me access to system commands, and I'm not unique in that. And there's many more things about the, so the environment has changed completely.” (Steve Bellovin Interview)

²² Arguably, there is a tragedy of the commons style dilemma lurking here. However, it is not one we explored during our research.

wrong guess.²³ Despite the mathematical beauty of this solution, such friction-in-design measures are not, however, universally implemented. Again, math theory falls victim to pragmatic considerations of everyday life; human factors and usability matter. Such delays are inconvenient and that can be off-putting for consumers of digital services, and so, as companies engage in a (rough) cost–benefit tradeoff, they may forsake security in the name of consumer convenience (profits). Of course, as David Thaw pointed out during our interview, this countermeasure is effectively used to secure bank ATMs from PIN guessing attacks and ought to be used much more widely. Bruce Schneier made a similar point, noting that “systems that lock you out after only a few incorrect guesses work great; it’s why your ATM card password can be secure with only four digits.” But he warned of “a denial-of-service trade-off: I can deliberately lock you out of any system that has that feature by using your username and guessing wrong.”²⁴

For offline guessing attacks (e.g., following a database leak), countermeasures often are the burden of professionals operating web services. Xavier de Carné de Carnavalet (2022) observed:

Proper password hashing, for example, should be applied to protect plaintext passwords; otherwise, even seemingly hard-to-guess passwords would be readily available in the [leaked] database. Inappropriate hashing algorithms such as MD5 and SHA1 have been extensively used, and one can still find recent breaches with such algorithms used. Besides the lag between expert and professional knowledge explained later [in this chapter], this is also the result of outdated resources being promoted on the first page of Google search results [and] how young professionals learned about authentication at school.²⁵

Third, beyond guessing attacks, the threats multiplied, became more varied, and hybridized. Some attacks targeted vulnerable weak links. For example, attackers recognized that password hosts and login commands could be subverted. Phishing attacks, compromised servers, and compromised client hosts make it easier for attackers to steal passwords (Bellovin 2016). If an attacker has accomplished one of

²³ These measures would not counter offline guessing attacks, which for the sake of brevity, we do not discuss extensively.

²⁴ Xavier de Carné de Carnavalet explained: “To be effective, delays should be applied to accounts rather than IP addresses. Therefore, as delays could be the result of an attacker trying to guess an account password, they may be denying service to the rightful owner of the account.” Pre-publication comments on “Common nonsense about password security and the expert–layperson knowledge gap.” 2022. On file with authors.

²⁵ See Naiakshina et al. (2020, 8) (“Some [participants] further suggested that there is a lot of outdated information on the Internet with respect to security practices (PJ6, PS8), making Internet research rather challenging.”); Naiakshina et al. (2017) (exploring poor password storage practices among developers); Krombholz et al. (2017, 1347) (study of HTTPS configuration by professionals: “Our participants reported that they came across outdated or simply wrong information in online tutorials.”). Of course, the perpetuation of misinformation, including outdated information about security practices, on the internet is a more general problem.

these attacks, then a strong password does not provide any defense (Florencio et al. 2007); the attacker doesn't care about special characters, or any other suggested characteristics users are told will make their passwords stronger and more secure. Thus, security experts realized that even if necessary, strong passwords are not sufficient protection for password-protected accounts and resources. David Thaw told us that strong passwords were counterproductive in the sense that (i) they could contribute to "security theatre" by providing a false sense of security to users so long as they comply with password composition rules, and (ii) draw attention away from other more important security threats and countermeasures.²⁶ Lee et al. (2022, 572) similarly offer as a hypothesis to help explain "the disconnect between industry and the academic community" the idea that "Password policy is security theater: measures such as character-class PCPs, even if ineffective, may give users a false sense of security, and websites use them for this reason."

Initial thinking about security measures to counter the threat of an organizational attack also was grounded in the math theory mindset. Conventional wisdom among experts suggested that one way to minimize the downside risk of a data breach at a company where a password file (database) has been compromised is to require people change their passwords periodically; then in the event of a breach, there is a good chance that by the time someone tries to exploit the leaked data, it has already been replaced with a new password. Bellovin explains why this wisdom fails. Users dislike frequent requests to change their passwords and tend to use similar patterns when rotating passwords (e.g., adding a number or increasing an existing number by one), making the replacement passwords even easier to guess. Again, sophisticated attackers who pay attention to human factors and behavioral patterns easily learn to predict what people will do when forced to periodically change their passwords. Armed with such knowledge, attackers can more easily guess passwords using a hacked dataset as a baseline (input) for their guessing algorithm. The conventional wisdom thus backfires. We return to this example below as a mini-case study that highlights relevant knowledge dilemmas between experts, professionals/practitioners, and lay people.

The emergence of new threats and their successful exploitation forced the expert community to confront the notion that password security is not just a mathematical problem solved with strong passwords. Instead, it is multifaceted and necessarily embroiled with the many fallibilities of humans and organizations.

CHANGES IN THE PASSWORD SECURITY EXPERT COMMUNITY Interviewees mostly described the password security expert community as academic researchers who

²⁶ Thaw further noted: "There are some attacks on [password authentication systems] about which scientists reasonably disagree as to whether or not the system provides sufficient defense on its own . . . so I respect that there can be reasonable disagreement, [but] I'm on the side [that] the system can provide reasonable defense and password complexity doesn't meaningfully add to that, but there's fair reasonable disagreement."

published research in peer review journals and attended academic conferences. Some industry researchers, often at research labs, who participated in these same activities also were considered members of the expert community.²⁷ Here we highlight changes in community membership (types of experts included) and corresponding changes in the types of research and knowledge developed within the expert community. One of our interviewees described the transformative shift in the expert community that occurred roughly over the course of a decade (late 2000s to 2010s) as the triumph of empirical observation over math theory, a demonstration that human factors matter and that usability needed to be considered by security experts. Essentially, changes in the technological environment described above drove changes in the expert community as the community expanded to integrate different expertise. In the first few decades of computing and networking, the security expert community included mostly computer scientists and engineers who approached password security more or less as a security optimization problem. For example, where guessing attacks were the threat to password-based authentication systems, password strength needed to be optimized considering the attacker's resource constraints; thus, the research question was how strong a password must be to withstand an automated guessing attack, taking into account expected computational resources of attackers. A similar mathematical, risk-based framing occurred for the threat of data breaches and stolen password files, and this framing led to security prescriptions against password reuse and in favor of organizational policies requiring periodic password changes.

In the 2000s to early 2010s, research from experts outside of the traditional computer security community highlighted how other sets of real-world constraints needed to be considered. Specifically, human factors and usability constitute real-world constraints on the effectiveness of strong passwords as knowledge-based authentication mechanisms. Lorrie Cranor stated that until this new research from usability researchers surfaced and gained traction within the security community, "there was basically no knowledge about usability of passwords other than just intuition." She further explained:

Experts knew what was easier or difficult for them personally, but beyond that there was basically no knowledge. Most of the knowledge about security was based on math. We can calculate how many possible passwords exist if we limit ourselves to characters and keys on the keyboard. You can do the math and figure out the password space. So people were looking at that sort of thing, but what they weren't looking at is the fact that there are a lot of combinations of characters that no human would ever come up with on their own. The space of realistic passwords is

²⁷ Interviewees explained that communications within this community mostly occurs in traditional academic forums such as conferences, academic publications, and interpersonal communications (email, conversations, sharing drafts and research ideas). In addition, experts share knowledge using online media, such as blogs and social media, as well as traditional media, such as newspapers and television appearances.

much smaller than the space of all possible passwords unless you have a computer randomly generating them for you. Over the past 15 years, we've learned a lot about what kinds of passwords humans *actually* create when left to their own devices. And we've also learned about how to nudge humans towards creating stronger passwords.²⁸

The expert community gradually broadened beyond conventional computer security experts to include experts in usability, information technology, and other related scientific and engineering disciplines as well as experts from adjacent social sciences such as psychology, economics, and cybercrime. The community remained an academic one, in the sense that these experts participated in academic conferences and published research in peer review publications, regardless of whether they were affiliated formally with academic institutions. This led to the development of new cross-disciplinary conferences and specialties, such as usable security.

When asked how expert knowledge had changed over the past two decades, Arvind Narayanan replied:

Computer security [is] a field that has often fetishized mathematical knowledge over human subjects experiments, for various cultural reasons, historical accidents, etc., and so . . . a lot of the earlier experts' wisdom and decrees regarding password security often involved burdening users, and we never really tested to see how effective [those burdens] are, and so a lot of the things that experts know better today are the opposite of things they might have said 20 years ago. So what are some of those things?

Never write down your password. That's one piece of bad password wisdom. The problem with that is because our memories are finite, and not writing down passwords means that people are going to reuse passwords, and that's a much bigger problem today. Most experts (today) will say it's OK to write down passwords, but think about the physical safety of where you write them down.

A second one today is the importance of multifactor authentication; passwords are never going to be your only line of defense. It doesn't matter how clever your password is. It's probably going to get breached at some point.

A third one is that a lot of the password-related advice these days would be directed at companies as opposed to users.

. . .

I think the security expert community is just a little bit more humane now than twenty years ago. Rather than treating users as the problem, [experts] recogniz[e] that people are overburdened as they go about their everyday lives, and designers have the primary responsibility for security.²⁹

²⁸ Author interview with Arvind Narayanan, Professor of Computer Science, Princeton University. 2022.

²⁹ Author interview with Arvind Narayanan, Professor of Computer Science, Princeton University. 2022.

These observations resonated with some articles we reviewed from that transitional period that noted how experts sometimes misunderstood users. Herley (2009) noted that some experts may suggest that users are lazy, choose weak passwords, and ignore expert security advice; but he argued that these actions may be entirely justifiable. Users' rejection of security advice often is rational, Herley suggested. According to Herley, much of the public security advice for users is complex, outdated, and does little to address actual threats that users face, offering a poor cost–benefit tradeoff, and thus users reject the advice. Dinei Florencio, Cormac Herley, and Paul C. van Oorschot (2014) explored how users should manage large numbers of accounts and found that mandating that users only use strong passwords and not reuse passwords leaves users with an impossible task as the number of accounts they have has grown.

INTERMEDIATING ROLE OF SECURITY PROFESSIONALS AND PRACTITIONERS Changes in the technological environment were driven by and at the same time drove growth and expansion of an intermediate layer of computer and network security professionals and practitioners (hereinafter, “professionals” for brevity).³⁰ The expert security community came to recognize the importance of these actors and engaged with them more directly. Unfortunately, we were unable to fully explore this engagement during our research. It surfaced during interviews in a few different ways. Some interviewees noted how some industry security researchers actively participated in the academic security research community, whether by attending conferences or publishing in peer review journals. Others noted how IT, security, and other professionals often learned about password security while at university and then went on to careers where security was often one of many responsibilities. This intermediate layer of professionals played an increasingly important role in determining password security practices, by setting organizational policies and deploying technologies and services that required users to create and manage passwords.

While we were unable to fully map and explore the knowledge flows between these different communities (experts, professionals, lay people), we highlight observations drawn from interviews to help frame future research.

³⁰ The professional/practitioner community includes both security and IT professionals, such as Chief Information Security Officers (CISO), Chief Security Officers (CSO), security engineers, architects and analysts at companies, and even ethical hackers. These types of jobs entail implementing and testing security programs across an organization, overseeing the operations of the IT department, overseeing the physical and cybersecurity of a company, protecting company assets from threats, implementing quality control, conducting internal and external security audits, and various other related activities. Members of this community often join professional organizations such as Association for Computing Machinery (ACM), IEEE Computer Society, Information Systems Audit and Control Association (ISACA), Society for Information Management (SIM), and User Experience Professionals Association (UXPA). These types of professionals tend not to publish academic articles; instead, they may give talks and training sessions to other members of the community and lay people.

One theme that surfaced repeatedly in our interviews is the idea of a disconnect between expert knowledge and professional knowledge and practice, which, in turn, seems to trickle down and potentially generate lay person misunderstanding. This disconnect is, at least in part, the result of a lag in the knowledge flows between experts and professionals. The flow of knowledge from experts to professionals can be slow and delay the ability of professionals to implement new standards and translate that information into forms accessible to lay people. Reasons for this lag include the rapidly changing technological environment and different, and sometimes competing, incentives between the expert community and the professional community. According to David Thaw, it is easy to propagate information about security and best practices through the professional community because it is not held to evidence-based standards (at least, not the same as academics). Professionals have access to academic research, but that does not mean they keep up to date with it. They have mixed responsibilities and priorities and may view security primarily in terms of compliance.

Three notable concerns arise. First, professionals responsible for password security, business practices, and interface design may perceive security as one of many different competing considerations (business, technical, etc.). Second, a compliance mindset may lead professionals to ignore academic security research and instead rely on security standards set by insurance companies, auditors, other industry actors, or government. Finally, professionals may be overconfident and inaccurately see themselves as security experts or as being sufficiently up-to-date on security.

Expert consensus on what actions constitute best practices also struggles to keep up with the rapidly changing technological environment. As a result, what was accepted as best practices five to ten years ago can become outdated. Experts must correct their views first and then translate that knowledge to the broader public, including professionals and lay people. Knowledge transfer between experts as well as between experts and professionals may happen too slowly; technological changes may occur faster than expert knowledge can update and percolate. As a result, lay users may be the last community to become aware of new password security developments, and their security behaviors may be at odds with the behaviors experts characterize as best practices. Based on this lag-induced knowledge gap, lay users may unknowingly become accustomed to outdated security practices believing that they are increasing their security when they are in fact doing the opposite.

Experts and professionals also may have competing incentives, directly impacting how information is shared, which information is regarded as important, and how best practices and security advice get implemented. According to various interviewees, academic incentives prioritize publishing articles to increase prestige within the community. Arvind Narayanan emphasized his view that academic researchers have a responsibility to translate their findings to the public, including professionals and lay people, since most academics receive funding from taxpayers. He suggested that most fail to do this work and that universities should consider crediting such efforts in order to realign academic incentives with the public interest.

Interviewees suggested that the professional community has a compliance mindset. As a result, the professional community tends to evaluate security based on process rather than outcomes. One of the ways professionals evaluate security is through security checklists from organizations such as NIST or companies such as Deloitte. These checklists are sometimes forced upon professionals by their auditors and often are not updated to reflect the latest expert knowledge. Thus, a question arises as to whether professionals believe they are following the best security practices when they comply with these checklists. A related concern is that many professionals have been in the industry for years, but they are not getting retrained or refreshing their knowledge to reflect evolving expert knowledge. As a result, updated best practices and expert knowledge may not reach professionals and trickle down to lay people. While the lack of consistent training and refresher courses may be out of the control of professionals and more in the hands of the companies they work for, Rick Wash suggests that many professionals care more about checklists and not getting blamed for any problems than staying educated about security itself. Bruce Schneier similarly suggested that organizational inertia coupled with these compliance incentives do not allow for expert consensus to be integrated successfully in the professional community. Finally, Xavier de Carné de Carnavalet (2022) ties the observations of Wash and Schneier back to the concept of security theatre:

bothering users with ineffective and sometimes counterproductive measures gives them a sense that a system is “secure”, i.e., it tries to prevent attacks by requesting additional steps. . . . This could also explain why security practitioners are slow to adopt the latest NIST standard that deprecates PCPs. Those were at least some visible security steps, and it seems that improving security should be synonymous [with] adding constraints, not removing them.³¹

(Some) Best Practices

The conventional wisdom (consensus views) among security experts about best practices have changed over time as we have already described. Today, there are some identifiable best practices that security experts mostly agree on. We do not provide a comprehensive analysis of current best practices.³² Instead, we focus on the password creation action arena³³ and highlight a recent study by Lee et al. 2022, which sought to empirically study whether 120 of the most popular websites followed

³¹ Xavier de Carné de Carnavalet (2022) further noted: “Finally, recommending the opposite of what professionals have been praising for years may not be easy to do while keeping their face. This would be admitting they were wrong.”

³² A comprehensive study would require further empirical work.

³³ Refer to Figure 1.1: “An action arena ‘refers to the social space where participants with diverse preferences interact, exchange goods and services, solve problems, dominate one another, or fight (among the many things that individuals do in action arenas)’ (Ostrom2005, 14) – in other words, the place at which the exogenous variables combine in particular instances, leading over time to observed patterns of interactions and outcomes. A particular action arena involves

best practices. To do this, the authors identified the following “established” best practices based on prior research in security and usability:

1. **Blocklists:** (i) Check user passwords against blocklists that include leaked and easily guessed passwords. (ii) Block user passwords that appear on such lists. (iii) Prompt user to create a different password (Lee et al. 2022). (We might add: Inform user about the reason.)
2. **Strength meters and minimum strength requirement:** (i) Provide accurate, real-time strength estimates and feedback. (ii) Set minimum strength requirements based on an appropriate measure of password strength, namely guessability rather than complexity, Shannon entropy, or compliance with composition rules (Lee et al. 2022).
3. **Composition policies.** (i) Do not require specific character classes. (ii) Set minimum password length of eight characters (Lee et al. 2022).

Lee et al. (2022) explain how these practices are supported by recent research in security and usability. While precluding users from creating passwords found on blocklists is a low-cost security measure, Tan et al. (2020) tested different password security requirements in two experiments and found security–usability tradeoffs among different blocklist configurations. Also, in their study, strength meter and text feedback informed users about the reason for disallowing a password. Other research has shown that such feedback can nudge users to create stronger passwords.

Lee et al. (2022, 563) explain that minimum strength “requirements and strength meters are both effective and user-friendly.” However, the preferred means for evaluating strength has changed from complexity to guessability, which is the “number of guesses needed to crack a password” (563). Determining strength thus involves testing the password against an adversarial neural network rather than calculating complexity/entropy. These developments led to a corresponding shift away from character-class requirements (e.g., password composition rules that require characters drawn from specific character classes) and toward minimum strength requirements of “at least 10^6 [guesses] to prevent online guessing attacks” (564).

Remarkably, Lee et al. (2022) found that out of 120 popular websites, only 13 percent followed the established best practices noted above. More than half of the websites did not check user passwords and allowed users to use the most common, leaked, and easy-to-guess passwords (e.g., “12345678”). Almost a decade ago, Carnavalet and Mannan (2014) published a damning empirical study of password meters that showed how poorly password meters fared in terms of evaluating password strength and guiding users to create strong passwords, in part because many of the meters measured strength based on entropy rather than guessability.

specific action situations and specific actors, along with those actors’ identities and roles” (Frischmann, Madison, and Strandburg 2014, 14).

That study also showed inaccurate and inconsistent results across password meters, such that weak passwords were sometimes rated as strong. Yet Lee et al. (2022) found that little had changed. Surprisingly, most sites did not use password meters, despite plenty of research in the intervening period showing their value when properly deployed. The websites that did include password strength meters (i) misused the meters to nudge users towards including certain characters (i.e., to satisfy outdated character-class requirements) instead of serving the preferred best practice of encouraging users to freely construct strong passwords and (ii) did not include an evaluation of guessability (Lee et al. 2022). One positive note from Lee et al. (2022) was the observation that more than half of the websites have an eight-character minimum length requirement. The authors surmised, “Perhaps this is a result of updated guidance from NIST in 2017, which now recommends an 8-character minimum length for passwords, up from its previous recommendation of 6 characters” (Lee et al. 2022, 569).

Based on their findings, Lee et al. (2022) concluded that websites should review the best practices established by academic experts to adjust to their password policies. Moreover, the researchers concluded that because there seems to be a disconnect between academic password security experts and industry, future research should directly engage with system administrators to address the disconnect.

We conclude this section by noting another shift in best practice advice that surfaced during interviews.³⁴ One conventional wisdom dispelled by experts during interviews is the idea that people should not write down their passwords in a notebook or on a piece of paper; a few experts mentioned this example. They explained that experts now generally agree that writing down passwords is a good security practice so long as the physical password list is kept in a secure location because it helps people keep track of passwords and thus avoid (i) creating overly simple passwords to compensate for memorization difficulties and (ii) engaging with password recovery tools. The experts emphasized that the shift in advice reflects a better understanding of human factors and usability as well as prioritization of the security threats ordinary people face.

GOVERNANCE CHALLENGES

Three (meso-level) action arenas concerning everyday knowledge about password security surfaced in our study: The password creation action arena, the password recovery action arena, and the password expiration and replacement action arena. Each concerns a set of repeat interactions among stakeholders, where professional community stakeholders (e.g., employees of a company that design and operate the

³⁴ Other best practices noted by experts included using a password manager, two- or multifactor authentication, and even reusing passwords but only for low-value accounts.

password creation or password recovery tools or that implement a password expiration and replacement policy) determine the rules-in-use, informed to some degree by expert knowledge, and lay people take actions to create, recover, or replace passwords, possibly learning from their experiences. These action arenas present a series of knowledge dilemmas and governance challenges, which we discussed previously. We leave further analysis of them for future work and now look “upstream” from the password expiration and replacement action arena to examine the role of the National Institute of Standards and Technology (NIST) in shaping this action arena.

NIST on Password Expiration and Replacement: A Case Study

NIST, founded in 1901, serves as a nonregulatory agency funded by the United States Department of Commerce. The mission of NIST is to “promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life.” Most important for our chapter is NIST’s role as a standard-setting body. NIST standards generally establish best practices. Government agencies such as the FBI, NSA, and USDA must adhere to NIST guidelines, and in many sectors of industry, including cybersecurity, private companies tend to adopt NIST standards, although this can get complicated and be delayed as the specific case study below demonstrates.

To examine how NIST’s capacity to set regulations, standards, and best practices in the password security industry can lead to governance challenges, we explore NIST’s decision to change password expiration standards and the results in industry that followed. In 2003, NIST released “NIST Special Publication 800-63, Appendix A,” which advised users to protect and secure their accounts by creating passwords with random characters, capital letters, and numbers, and to change their passwords regularly (McMillan 2017). Federal agencies, large corporations, and universities followed the advice. However, Bill Burr, who authored the publication, has since stated that the advice was largely wrong. Burr wanted to use real password data as a foundation of his advice, but there wasn’t much of this data available to use as evidence for recommendations. Instead, security experts and professionals, including Burr, relied heavily on a paper written by Robert Morris and Ken Thompson (1979).

Recognizing the flaws and limitations of the 2003 publication, in June of 2017 NIST published “800-63-B: Digital Identity Guidelines: Authentication and Lifecycle Management.” In this publication, NIST established new guidelines for password security to replace widespread practices and policies that made authentication mechanisms weak. One of the recommendations given by NIST was for organizations to no longer require password expiration, stipulating that users should only change their passwords if there is evidence of compromise (Grassi et al. 2017). The exact wording of the recommendation is as follows “Verifiers SHOULD NOT

require memorized secrets to be changed arbitrarily (e.g., periodically). However, verifiers SHALL force a change if there is evidence of compromise of the authenticator.” (Grassi et al. 2017).

The change in the guidelines resulted from academic research and expert consensus that password expiration is counterproductive to good password security. As we described extensively above, decades ago, the field of password security focused heavily on mathematical knowledge and theory and not on human factors and usability. Since password security ideas were based on math, rarely were expert recommendations tested empirically with users to measure their effectiveness. The modern practice of empirically testing password recommendations is relatively new. Today, security experts collect data on how often common attack types actually happen, and this data combined with empirically testing password recommendations against human usability has enabled experts to identify the following problems associated with password expiration: Lay users, when forced to comply with password expiration policies, often change their passwords in a trivial way, such as adding a 1 at the end of an already established password, making the password easy to guess; this doesn't solve the problem (attack type) as this type of behavior is predictable from an attacker perspective. Academic researchers conducted research, collected data, and published results (Chiasson & van Oorschot 2015; Cranor 2016; Habib et al. 2018; Spafford 2006; Yinqian Zhang et al. 2010; Zhang-Kennedy et al. 2016), and it became obvious to them that periodic password changes needed to be eliminated as a supposed best practice. Highlighting the impactful nature of such academic expert research, when creating their updated password guidelines, NIST cited academic research that reported on the ineffectiveness of password expiration policies (Habib et al. 2017; Herley & van Oorschot, 2012; Komanduri et al. 2011).

While NIST took into account up-to-date research findings and changed its standards, it is not clear whether the professional community has followed suit. Periodic password expiration and replacement appears to remain a widespread practice. Compliance with NIST's new standard seems slow, at best. Frankly, we asked interviewees about it, and while most speculated that industry has not yet adopted the new standard, no one knew of an empirical study of the issue. When asked about the reasons for this lag, our interviewees stressed the observations we noted in the previous section concerning the compliance mindset, conflicting incentives, and insufficient training and keeping up-to-date with expert knowledge.

CONCLUSION

Common nonsense about password security is a type of misinformation and digital illiteracy problem. In some ways, it is distinguishable from fake news and propaganda types that depend on trolls, bots, or others who deliberately pollute knowledge networks. Yet there are similarities, particularly with respect to how people learn about and from the digital tools they regularly use as well as the perpetuation of

outdated information. This study examined different epistemic dynamics and dilemmas that may require different forms of governance. Unfortunately, we did not fully explore governance institutions within and between the different communities. That is an important topic for follow-on research. Our study shed light on the roles of experts and professionals, social demand for expert knowledge transfer via better and more widely accessible communication and education channels, and the societal risk of widening the lay–professional–expert knowledge gap. The GKC framework usefully structured our research methods and analysis, and it helped us identify different action arenas deserving of further study.

REFERENCES

- ‘About NIST’. NIST, 10 July 2009. www.nist.gov/about-nist.
- Bellovin, Steven M. 2016. *Thinking Security: Stopping Next Year’s Hackers*. New York: Addison-Wesley.
- Bobba, Sriram, and Vamsi Paruchuri. 2020. “Single Sign-on Using Contactless Smart Cards and Fingerprint Authentication.” *Lecture Notes in Networks and Systems* (October): 158–166. https://doi.org/10.1007/978-3-030-90072-4_16.
- Bonneau, Joseph, Elie Bursztein, Ilan Caron, Rob Jackson, and Mike Williamson. 2015a. “Secrets, Lies, and Account Recovery: Lessons from the Use of Personal Knowledge Questions at Google.” In *Proceedings of the 24th International Conference on World Wide Web (WWW ’15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 141–150. <https://doi.org/10.1145/2736277.2741691>
- Bonneau, Joseph, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2015b. “Passwords and the Evolution of Imperfect Authentication.” *Communications of ACM* (June). www.microsoft.com/en-us/research/publication/passwords-and-the-evolution-of-imperfect-authentication/.
- Burr, William E., Donna F. Dodson, and Ebad A. Nabbus et al. 2013a “Electronic Authentication Guideline.” National Institute of Standards and Technology (November). <https://doi.org/10.6028/NIST.SP.800-63-2>.
- Burr, William, Donna F. Dodson, and W. Timothy Polk (eds.) 2013b. “Electronic Authentication Guideline.” NIST Special Pub 800-63 Version 1.0, June 2004. (Later versions include Burr et al., NIST SP - 800-63-2, August.)
- Chiasson, Sonia, and Paul C. van Oorschot. 2015. “Quantifying the Security Advantage of Password Expiration Policies.” *Designs, Codes and Cryptography* 77 (2–3) (December): 401–408. <https://doi.org/10.1007/s10623-015-0071-9>.
- Cranor, Lorrie. 2016. “Time to Rethink Mandatory Password Changes.” Federal Trade Commission, 2 (March). www.ftc.gov/policy/advocacy-research/tech-at-ftc/2016/03/time-rethink-mandatory-password-changes.
- de Carné de Carnavalet, Xavier, and Mohammad Mannan. 2014. “From Very Weak to Very Strong: Analyzing Password-Strength Meters – NDSS Symposium.” (February). www.ndss-symposium.org/ndss2014/programme/very-weak-very-strong-analyzing-password-strength-meters/.
- Egelman, Serge, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov, and Cormac Herley. 2013. “Does My Password Go up to Eleven? The Impact of Password Meters on Password Selection.” (January). www.microsoft.com/en-us/research/publica

- tion/does-my-password-go-up-to-eleven-the-impact-of-password-meters-on-password-selection/.
- Florencio, Dinei, and Cormac Herley. 2006. "A Large Scale Study of Web Password Habits." 1 (November). www.microsoft.com/en-us/research/publication/a-large-scale-study-of-web-password-habits/.
- Florencio, Dinei, Cormac Herley, and Baris Coskun. 2007. "Do Strong Web Passwords Accomplish Anything?" USENIX (June). www.microsoft.com/en-us/research/publication/do-strong-web-passwords-accomplish-anything/.
- Florencio, Dinei, Cormac Herley, and Paul C. van Oorschot. 2014. "Password Portfolios and the Finite-Effort User: Sustainably Managing Large Numbers of Accounts." (August). www.microsoft.com/en-us/research/publication/password-portfolios-and-the-finite-effort-user-sustainably-managing-large-numbers-of-accounts/.
- Frischmann, Brett, Michael Madison, and Katherine Strandburg. 2014. *Governing Knowledge Commons*. Oxford University Press.
- Frischmann, Brett, and Evan Selinger. 2018. *Re-Engineering Humanity*. Cambridge: Cambridge University Press (April). <https://doi.org/10.1017/9781316544846>.
- Grassi, Paul A., James L. Fenton, and Justin P. Richer et al. 2017. *Digital Identity Guidelines: Authentication and Lifecycle Management*. Gaithersburg, MD: National Institute of Standards and Technology (June). <https://doi.org/10.6028/NIST.SP.800-63b>.
- Habib, Hana, Jessica Colnago, and William Melicher et al. 2017. "Password Creation in the Presence of Blacklists." www.ndssymposium.org/wpcontent/uploads/2017/09/USEC2017_01_3_Habib_paper.pdf.
- Habib, Hana, Pardis Emami Naeini, and Lorrie Faith Cranor et al. 2018. "User Behaviors and Attitudes under Password Expiration Policies." (August): 13–30. www.usenix.org/conference/soups2018/presentation/habib-password.
- Haenschen, Katherine, Brett Frischmann, and Paul Ellenbogen. 2021. "Manipulating Facebook's Notification System to Provide Evidence of Techno-Social Engineering." *Social Science Computer Review* 40 (6): 1–18.
- Herley, Cormac. 2009. "So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users." (April). www.microsoft.com/en-us/research/publication/so-long-and-no-thanks-for-the-externalities-the-rational-rejection-of-security-advice-by-users/.
- Herley, Cormac, and Paul van Oorschot. 2012. "A Research Agenda Acknowledging the Persistence of Passwords." *IEEE Security & Privacy Magazine* (January). <http://research.microsoft.com/apps/pubs/default.aspx?id=154077>.
- "How Secure Is My Password?" Accessed October 30, 2022. <https://nordpass.com/secure-password/>.
- Ion, Iulia, Rob Reeder, and Sunny Consolvo. 2015. "(No) One Can Hack My (Mind): Comparing Expert and (Non-Expert) Security Practices." (July): 327–346. www.usenix.org/conference/soups2015/proceedings/presentation/ion.
- Khem-am-nuai, Warut, Matthew J. Hashim, Alain Pinsonneault, Weining Yang, and Ninghui Li. 2022. "Augmenting Password Strength Meter Design Using the Elaboration Likelihood Model: Evidence from Randomized Experiments." *Information Systems Research* (March). <https://doi.org/10.1287/isre.2022.1125>.
- Komanduri, Saranga, Richard Shay, and Serge Egelman et al. 2011. "Of Passwords and People: Measuring the Effect of Password-Composition Policies." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*: 2595–2604. ACM. www.ece.cmu.edu/~lbauer/papers/2011/ch2011-passwords.pdf.
- Krombholz, Katharina, Wilfried Mayer, Martin Schmiedecker, and Edgar Weippl. 2017. "I Have No Idea What I'm Doing": On the Usability of Deploying HTTPS." In *Proceedings*

- of the 26th USENIX Conference on Security Symposium (SEC.17). USENIX Association, USA: 1339–1356.
- Lee, Kevin, Sten Sjöberg, and Arvind Narayanan. 2022. “Password Policies of Most Top Websites Fail to Follow Best Practices.” (August): 561–580. www.usenix.org/conference/soups2022/presentation/lee.
- Mazurek, Michelle L., Saranga Komanduri, and Blase Ur et al. 2013. “Measuring Password Guessability for an Entire University.” In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, (November 2013): 173–186. CCS '13. Association for Computing Machinery, New York. <https://doi.org/10.1145/2508859.2516726>.
- McMillan, Robert. 2017. “The Man Who Wrote Those Password Rules Has a New Tip: N3v\$R M1^d!” *Wall Street Journal*, 7 August, sec. Page One. www.wsj.com/articles/the-man-who-wrote-those-password-rules-has-a-new-tip-n3v-r-m1-d-1502124118.
- Morris, Robert, and Ken Thompson. 1979. “Password Security: A Case History.” *Communications of the ACM* 22 (11) (1 November): 594–597. <https://doi.org/10.1145/359168.359172>.
- Naiakshina, Alena, Anastasia Danilova, Eva Gerlitz, and Matthew Smith. 2020. “On Conducting Security Developer Studies with CS Students: Examining a Password-Storage Study with CS Students, Freelancers, and Company Developers.” In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, 1–13. <https://doi.org/10.1145/3313831.3376791>.
- Naiakshina, Alena, Anastasia Danilova, Christian Tiefenau, Marco Herzog, Sergej Dechand, and Matthew Smith. 2017. “Why Do Developers Get Password Storage Wrong? A Qualitative Usability Study.” In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17)*. Association for Computing Machinery, New York. 311–328. <https://doi.org/10.1145/3133956.3134082>.
- Ostrom, Elinor. (2005). *Understanding Institutional Diversity*. Princeton, NJ: Princeton University Press.
- “Password Meter – A Visual Assessment of Password Strengths and Weaknesses.” Accessed October 30, 2022. www.uic.edu/apps/strong-password/.
- Paul Eve, Martin. 2016. *Password*. 1st ed. Object Lessons. Bloomsbury Publishing. www.bloomsbury.com/us/password-9781501314889/.
- SearchSecurity. 2022. “What Is Cybersecurity? Everything You Need to Know | TechTarget.” Accessed October 30, 2022. www.techtarget.com/searchsecurity/definition/cybersecurity.
- Spafford, Gene. 2006. “Security Myths and Passwords.” April 19. www.cerias.purdue.edu/.
- Tan, Joshua, Lujio Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2020. “Practical Recommendations for Stronger, More Usable Passwords Combining Minimum-Strength, Minimum-Length, and Blocklist Requirements.” In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (October):1407–1426. Association for Computing Machinery, New York. <https://doi.org/10.1145/3372297.3417882>.
- “Top 200 Most Common Password List 2021.” Accessed October 30, 2022. <https://nordpass.com/most-common-passwords-list/>.
- Ur, Blase, Jonathan Bees, Sean M. Segreti, Fumiko Noma, and Jonathan Bees. 2016a. “Do Users’ Perceptions of Password Security Match Reality?” *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM: 3748–3760. doi.org/10.1145/2858036.2858546.
- Ur, Blase, Fumiko Noma, and Jonathan Bees. 2016b. “I Added ‘!’ at the End to Make It Secure: Observing Password Creation in the Lab.” *Proceedings of the Eleventh*

- USENIX Conference on Usable Privacy and Security*, USENIX Association. (May): 123–140.
- Ur, Blase, Patrick Gage Kelley, and Timothy Passaro et al. 2012. “How Does Your Password Measure Up? The Effect of Strength Meters on Password Creation.” 65–80. www.usenix.org/conference/usenixsecurity12/technical-sessions/presentation/ur.
- Weizenbaum, Joseph. 1976. *Computer Power and Human Reason: From Judgment to Calculation*. New York: W. H. Freeman.
- Zhang, Yinqian, Fabian Monrose, and Michael K. Reiter. 2010. “The Security of Modern Password Expiration: An Algorithmic Framework and Empirical Analysis.” In *Proceedings of the 17th ACM Conference on Computer and Communications Security* (October): 176–186. CCS '10. Association for Computing Machinery, New York. <https://doi.org/10.1145/1866307.1866328>.
- Zhang-Kennedy, Leah, Sonia Chiasson, and Paul van Oorschot. 2016. “Revisiting Password Rules: Facilitating Human Management of Passwords.” In *2016 APWG Symposium on Electronic Crime Research (ECrime)* (June): 1–10. <https://doi.org/10.1109/ECRIME.2016.7487945>.