Balancing Security and Usability in Encrypted Email

End-to-end encryption is the best way to protect digital messages. Historically, end-to-end encryption has been difficult for people to use, but recent tools have made it more broadly accessible, largely by employing key-directory services. These services sacrifice some security properties for convenience. A 52-person user study found that participants could learn to understand properties of different encryption models and make coherent assessments about when different tradeoffs might be appropriate. Participants recognized that a less-convenient key exchange model was more secure overall, but considered the key-directory approach to have security sufficient for most everyday purposes.

Wei Bai and Doowon Kim
University of Maryland

Moses Namara
Clemson University

Yichen Qian
University of Maryland

Patrick Gage Kelley
University of New Mexico

Michelle L. Mazurek
University of Maryland

Important communications now take place on digital channels, increasing privacy concerns. Users of email and chat services risk disclosure of their messages through attacks on communications services from outsiders, rogue employees, and even government subpoenas. The only way to truly assure confidentiality is to use end-to-end encryption, so that communications services have no access to the message content they deliver. Despite considerable evidence and mass media reporting about content breaches, end-to-end encryption has not yet been widely adopted for email and only relatively recently become common for chat services.

Researchers and security advocates have spent considerable effort investigating the reasons for this lack of adoption. More than 15 years of research have identified major usability problems with encryption tools, ranging from poorly designed user interfaces to the fundamental challenges of safe and scalable key distribution.[1-5]

Since 2014, however, Apple has applied seamless end-to-end encryption to its iMessage and FaceTime services. By centrally distributing public keys, Apple ensures the encryption is invisible to users. This design, however, leaves open the possibility that Apple itself could carry out a man-in-the-middle (MITM) attack to violate users' privacy, for example at the request of law enforcement authorities.[6] WhatsApp has implemented end-to-end encryption for text, voice, and video communications. WhatsApp also centrally distributes public keys; however, users can optionally verify each other's keys manually

1089-7801/17/$33.00 © 2017 IEEE

or via QR code. Google and Yahoo are reportedly developing similar approaches for email, with an added monitoring protocol that lets users and third parties audit the key directory for consistency.

Some privacy experts have suggested that because of the potential for an MITM attack, key-directory services that don't allow verification shouldn't be recommended to end users. One security researcher suggests that "iMessage remains perhaps the best usable covert communication channel available today if your adversary can't compromise Apple. ... If one desires confidentiality, I think the only role for iMessage is instructing someone how to use Signal."[6]

In a sense, the issue comes down to whether the benefit from many more people adopting encrypted communications is outweighed by the reduced security of the central key-distribution model. While security experts are best positioned to understand the technical differences between models, end users will ultimately choose which platforms to use. Researchers have considered the needs of some highly privacy-sensitive users, such as journalists and activists,[7,8] but our work is the first to ask average users for their opinions about these tradeoffs. This means that although security researchers might understand the risks and benefits of different tools, as a community we don't understand how an average user will weigh different factors in deciding whether to adopt or ignore various encrypted communication technologies.

Key Directories and Distributing Trust

Key management for email and other communication exists on a spectrum from completely decentralized to completely centralized. Pretty Good Privacy (PGP) uses a peer-to-peer key endorsement model; while appealing in theory, this decentralized approach requires significant user investment in key management and has not gained widespread adoption.

At the other end of the spectrum is Secure/Multipurpose Internet Mail Extensions (S/MIME). In this approach, all users have public-key certificates signed by a certification authority, which are distributed along with any signed emails sent by that user. S/MIME allows straight forward integration of encryption into email clients and is used by some corporate organizations, but it has not been widely adopted by consumers either.

Key directories are partially centralized approaches where users (or their devices) generate and manage their own private keys, but the communications service keeps a record of all public keys and distributes them to potential correspondents upon request. As these services increase in popularity, security researchers have investigated ways to reduce the likelihood of MITM attacks or other compromising behaviors. Most of these approaches rely on transparency — publicly available, cryptographically secure logs can be audited to detect when a key-directory service equivocates, or provides different key information to different clients.[9,10] (Equivocation presents a security threat because it suggests some users are being given incorrect keys, perhaps to allow decryption by third parties.) Google and Yahoo are exploring variations of this transparency approach for their end-to-end encryption extensions. Another approach is to split the key information between two different entities, requiring collusion to subvert the system.[11]

Email Encryption Usability Challenges

There have been nearly two decades of research into why encryption is difficult to use.[1-5] These studies have identified problems related to poor interface design, as well as deeper problems arising from the inherent difficulties of key management and from users' mistaken mental models of encryption. This research has led to important advances in the design of usable encryption interfaces. Our study examines a different problem: we explicitly minimize evaluation of the user interface design to focus on the encryption models' underlying properties.

Most encryption usability studies evaluate participants' success rates learning unfamiliar encryption and decryption tasks. They provide valuable insight into how effectively novices can learn a particular system, how specific user interface design choices impact users, and where the difficulties lie. However, users are rarely provided with multiple encryption systems to compare or are encouraged to consider their pros and cons. Studies that do compare usability among different interfaces or systems[3] don't explicitly consider the associated security tradeoffs.

Researchers also have studied social and cultural norms that discourage encryption. Often users believe they don't need encryption because they feel they have nothing to hide, or because they can't imagine anyone being interested in

Table 1. Partcipants' encryption-related tasks for the two main evaluated models.*

Task no.	Exchange model	Basic registration model
1	Generate public lock/private key pair	Generate and register public lock/private key pair
2	Exchange public lock with Alice	N/A
3	Send encrypted email to Alice	
4	Decrypt received email from Alice	
5	Exchange public locks with Bob and Carl	N/A
6	Send encrypted email to Bob and Carl	
7	Decrypt received email from Bob and Carl	
8	Imagine sending encrypted email to 10 people	
9a	Consider misconfiguration: lose Alice's public lock	N/A
9b	Consider misconfiguration: lose own private key	
9c	Consider misconfiguration: publicize own private key	

* Tasks differed slightly in the two models.

the messages they're sending.[7] Other people resist encryption because they're afraid it will draw attention or be perceived as paranoid.

A Study of User Preferences

We designed a within-subjects study to examine how participants would understand and value tradeoffs regarding the usability and security of end-to-end email encryption. Each participant was introduced to two general models for key management: exchange and registration. For both models, we described a public key as a *public lock*[12] for simplicity. In the exchange model, participants generate a lock/key pair and then must exchange locks with people with whom they want to communicate (via email, instant messaging, or posting the lock publicly). In the registration model, which simulates a key-directory service, participants actively register their public locks; all necessary correspondents' locks were preinstalled to the directory, so that participants could send and receive encrypted email immediately upon registering their own locks.

For each model, participants were asked to complete a series of simulated tasks (see Table 1) and also introduced to a brief, nontechnical review of security properties. For the exchange model, we described an MITM attack in which the attacker could intercept or replace public locks during the exchange process. For the registration model, we primarily described an MITM attack enabled by the directory service providing different keys to different users. We varied the order of activities across participants to offset ordering effects. Participants gave their opinions about each model immediately after completing the tasks and learning about security for that model, and also answered summative questions comparing the different models.

We adapted an existing Chrome extension, Mailvelope, to simulate both models. To ensure we tested the models rather than the interface, we deliberately minimized all interface differences except those required to exchange or register locks, and we tried to make the interface for lock exchange as simple as possible.

After participants worked with the registration model, we also described two variations on the basic registration approach. In confidentiality as a service (CaaS),[11] a third-party service acts as an intermediary to the email provider. Neither can read the encrypted email independently, but if they collude, both can read the email. In the auditing model (similar to that proposed by security researchers[9]), external auditors check keys distributed to different users to ensure they match (with some delay).

A total of 52 adults familiar with Gmail and Chrome participated in our study between August 2015 and February 2016. These participants were somewhat younger, more male, and more technical than the general US population, but most had little experience with computer security, measured using a scale developed by L. Jean Camp and colleagues.[13] We presented the auditing model, which was added during recruiting, to 24 participants.

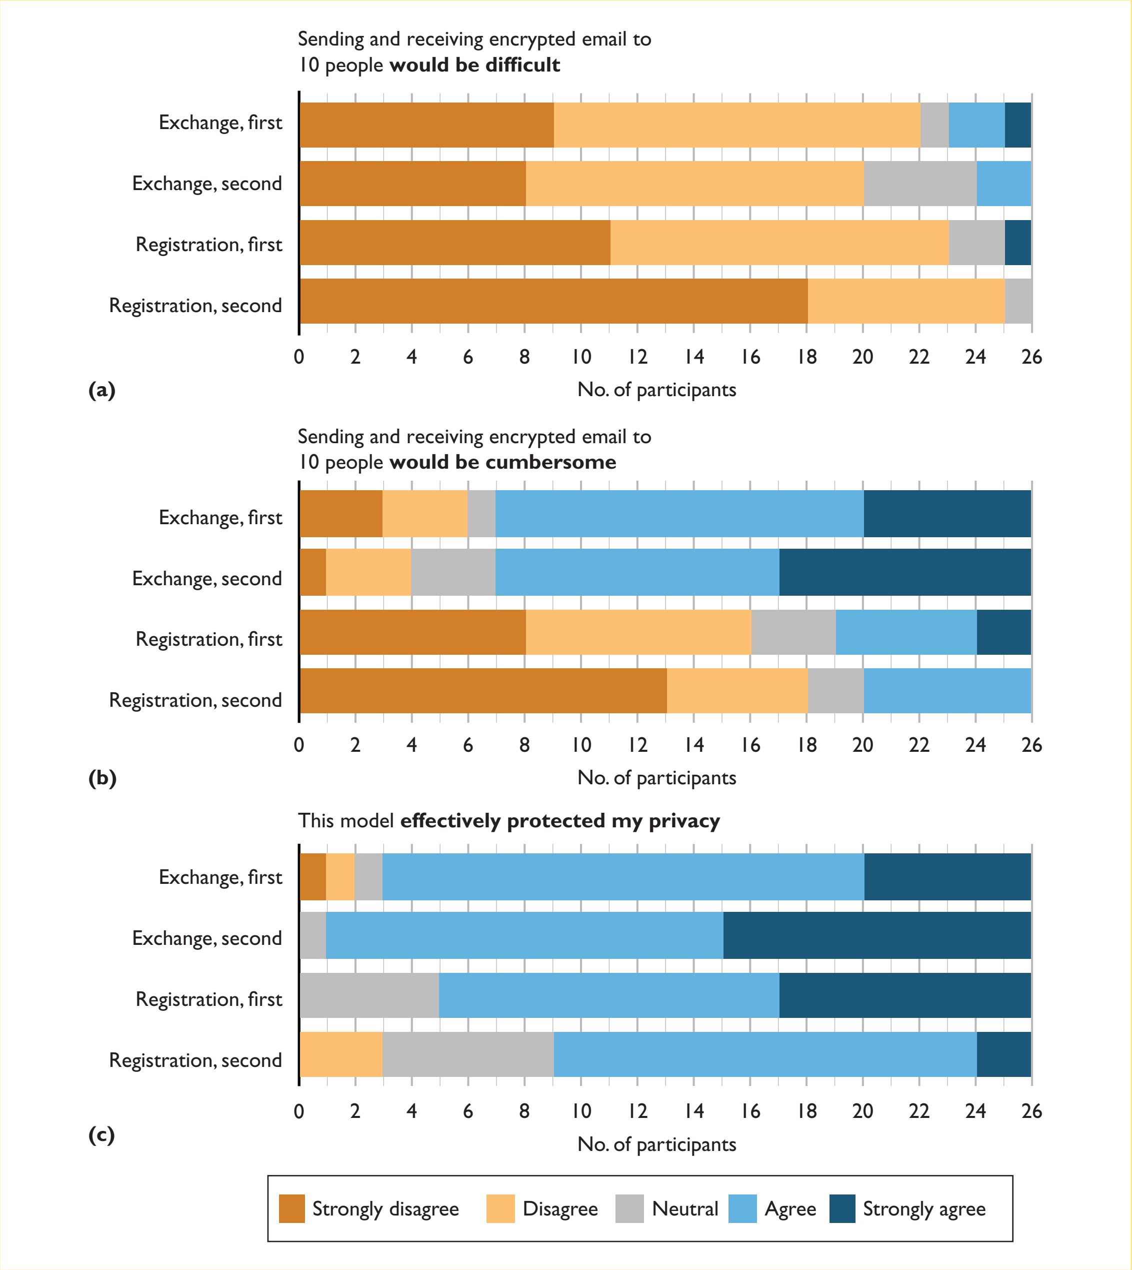

Figure 1. Participants rated the two main models for (a) difficulty, (b) cumbersomeness, and (c) privacy protection, on a five-point scale. The labels indicate the order in which participants evaluated each model: for example, "exchange, first" includes ratings for the exchange model among participants who used it before the registration model, while "exchange, second" indicates ratings for the exchange model among participants who used it after the registration model.

Registration Is More Convenient

Unsurprisingly, our participants found the registration system considerably more convenient than the exchange model. Using a cumulative-link (logit) mixed regression model (CLMM), we found that the exchange model was perceived to be both 20 times more cumbersome (tedious) and 5 times more difficult (challenging) than the registration model when sending an email to a group of 10 people. Figures 1a and 1b illustrate these differences.

As might be expected, the most tedious task was manually exchanging locks, and the most commonly mentioned issue was waiting for a correspondent's public lock. One participant was concerned that the exchange model was "time-consuming, especially sending urgent emails. I have no choice but to wait." But participants recognized that the main problem was the initial setup: "If their locks are already there, it would not be cumbersome."

While few participants considered any of the tasks terribly difficult, choosing a mechanism for exchanging locks (among provided options including email, chat, and posting publicly) was often considered the most difficult step. One had to "think about a safe way to exchange public locks," and another was concerned about making an exchange error while multitasking. Several participants mentioned that users with low digital literacy might need or prefer the registration model. For example, one participant recommended the registration model "especially to people that don't know very well how to use a computer," and another said the registration model is "easy to teach others to use."

The inconvenience of the exchange model could be mitigated somewhat by posting the key publicly or semipublicly, rather than sending it individually to different recipients. About one-third of our participants chose to post to a provided Facebook profile or lock server, but few mentioned the added convenience as a reason for their choice. Several participants chose the lock server because they thought it was more secure than other choices we provided.

Perceived Security Gap

Our participants understood and thought critically about the security properties we explained to them for each model. Perhaps surprisingly, they viewed the exchange model as only marginally more secure than the registration model (see Figure 1c).

Exchange Model: Can Lead to Vulnerability

Almost all participants (92 percent) agreed that the exchange model protected their privacy; however, many expressed concern that managing lock exchanges themselves would create vulnerabilities. More than half (52 percent) were concerned about the medium used to exchange locks, including that the lock server could be "manipulated or compromised," that an attacker could break into a Facebook account to post an incorrect lock, that WiFi in a coffee shop could be unsafe for transmitting locks, or that the Internet service provider could "sit between my recipient and me" and execute an MITM attack. Several participants noted that to "send the public locks and encrypted emails using the same email provider, it's not very secure."

Other participants worried about the ongoing responsibility of managing locks. One said, "Every time when I send or get a public lock ... there is a probability, even though not high, that my privacy is compromised. Then when I exchange public locks with many people, this probability will increase exponentially." Another noted that, "non-tech-experts may make mistakes." Several participants also mentioned that careless or compromised users could ruin the security of a whole group: "If one person is hacked, then the whole company is hacked."

Registration Model: Generally Trusted

As expected, many participants were concerned about the need to trust email providers in the registration model. Despite this, however, most participants (73 percent) agreed the system would protect their privacy. The order in which the models were introduced was significant. Participants who saw the registration model first were more comfortable with it than those who had already heard about the more-secure exchange model. Figure 1c shows the difference.

Many participants believed that even though email providers could compromise the security of the basic registration model, they would be unlikely to do this. Ten participants mentioned that they trust their own email provider. Others pointed out that some email providers are untrustworthy in well-known ways: "Some people ... will choose Yandex email from Russia [rather than Gmail], because they'd rather be intercepted by the Russian government, instead of the US government." Others would trust "certain big companies, not small companies," because big companies that must protect their reputations don't invade privacy "unless the government forces them to do so." Another was comfortable with "one large provider. I know whom to blame if a problem occurs."

CaaS and Auditing: Additional Perceived Security

More than 40 percent of participants thought the CaaS variation was more secure than the basic registration model, primarily because they believed the different companies would not collude. Examples of these comments included the following: "This separation makes me feel good," and "If one party is screwed up, you have another one to protect [your email]. You are still safe."

In contrast, about 20 percent of participants thought the basic registration model was more secure than the CaaS variation. Some believed adding a third party inherently made the system more complicated, and therefore less trustworthy. Others thought the potential for collusion invalidated the CaaS premise entirely. Two participants mentioned the importance of the potential gain: the two parties might not collude "for one or two persons' email, but for many, a group of people" or "if there are terrorists." Others were concerned about the trustworthiness of the third-party service; one said his opinion "depends on who the two entities are."

Most participants exposed to the auditing variation believed it improved security. One participant was happy that "somebody is supervising" lock distribution, and another said, "Obviously it's extra secure. Other parties are verifying it." The presence of many auditors reassured participants that collusion was unlikely; for example, one said, "It's less likely that all auditors [would] do something bad." Several participants, however, were concerned about the reliability of the auditors: "I want to know who these auditors are ... their reputations, and whether they are truly independent." Ten participants worried about the time lag for notification; for example, "Even an hour is too late ... Something bad has already happened." Others, however, were more pragmatic: "Immediate notification is ideal, but I don't expect immediateness in reality."

Which System Would You Use?

After exposing participants to different models, we asked whether they would use or recommend any. CLMM results show no significant difference in willingness to use the three models, but do show that basic registration was recommended less frequently than the exchange model. Participants who completed the encryption tasks before hearing about security properties were significantly less likely to say they would use or recommend any model than those who heard about security properties first. We hypothesize that participants who used the encryption extension before hearing about security anchored on the inconvenience of the tool rather than its privacy benefits. While this doesn't provide useful insight about comparing the different systems, it does underline the need for careful consideration about how new encryption tools are presented to the public.

We asked participants why they would or wouldn't use and recommend each system. Unsurprisingly, the perception of better security attracted participants to the exchange model, while poor usability drove them away. On the other hand, participants were split as to whether security or usability was the main reason to use a registration model.

Participants who said they would use the exchange model generally described using it for high-security information only, or only at a small scale. One participant said the exchange model is "the safest one. I want to use it in small scale, like one or two people ... like private and personal things. But I don't want to use it every day." In contrast, participants who said they would use the registration model mentioned "contacting a large number of customers" as well as "party and meeting announcements." Participants also mentioned using the exchange model only for "extremely sensitive information, such as SSNs," compared to using the registration model for information that was "overall private, but would not be a disaster if disclosed."

A few said they would not use the exchange model because it was too cumbersome, but would recommend it to others with stronger privacy requirements. Similarly, one said that "encryption is not necessary for me," but recommended the CaaS variation of the registration model because it is "easier to use [than the exchange model] and more secure than the vanilla [basic] registration system."

Most Choose the Auditing Model

We asked participants who heard about the auditing model whether they would use it, and 15 of the 24 said they would, and 10 preferred it to any other model discussed. One participant said, "It's best among all systems mentioned in the experiment, because somebody else is policing them, just like watchdogs. If someone is reading your email, they might be caught." Another preferred the auditing model to any other option because "instead of using [the attacker's] public lock blindly, I will get the update, 'Oh, that's the wrong public lock, you should not use this.'"

Four participants found the auditing model superior to the other registration models, but preferred the exchange model in at least some circumstances. One said the audit model is "slightly better than [the basic] registration model … But overall, the lock exchange system has extra steps, extra layers of security, so I like it best." Four participants said auditing was worst among all models discussed, either because they didn't trust the auditors or because the time lag was too great.

Evaluating Tradeoffs

In deciding which system(s) they preferred, participants explicitly and deliberately made tradeoffs among security and usability features. For example, one participant said the exchange model "makes it more private for me," despite the fact that "it takes time to exchange locks." Similar considerations motivated one participant to reject the registration model: "It's easy to send encrypted emails, especially to many people. But security concern is the reason I don't want to use it."

Other participants made the opposite judgment: one would use the basic registration model because it's "easy to use, and I think most of us trust our email provider," despite "some possible threats." Another said the exchange model "maybe gives me safer feelings, more protection," but would not use it because "the disadvantage is it is time consuming, cumbersome, tedious, more complicated." These comments and others demonstrate that participants understood pros and cons of the different models and thought carefully about how to balance them.

Encryption versus No Encryption

Many participants didn't believe most or any of their communication requires high security. One said "encryption is not necessary for me," and another said, "If I have some private information, I won't put it on the Internet."

On the other hand, five participants with concerns about the security of the registration model also mentioned that it does provide a valuable security improvement over unencrypted email. One participant said, "Doing encryption gives me a security sense that I lock my door," and another noted that in the registration model, "I have to trust the email provider, which is problematic, but … better than raw email."

Security Thinking

Our participants made several thoughtful points about encryption, security, and privacy that apply across models. One speculated that an extra benefit of encryption could be a reduction in targeted ads: the "email provider can collect data through my emails, and then present ads. … [Using this tool] the ads will not appear." While in reality unencrypted metadata might still enable ad targeting, this demonstrates depth of security thinking. In contrast, a different participant worried that encryption would bypass the email provider's virus-detection system. Others worried that sending encrypted mail might in itself catch the interest of attackers, or make the user appear paranoid. (Similar concerns were raised in other work.[7]) Perhaps because of these concerns, more than 20 percent of participants expressed interest in controlling who would be allowed to send them encrypted email.

Handling Misconfigurations

We asked participants to consider how they would handle various possible misconfigurations in each model (see Table 1). This both prompted them to consider usability issues related to longer-term key maintenance and offered a chance to evaluate their understanding of the different security models. Most participants (75 percent) responded to all five scenarios with a straightforwardly correct answer, such as asking a correspondent to resend a lost public lock or generating a new lock-key pair and redistributing the lock to all correspondents. Additional participants (13.5 percent) provided such answers to at least three of the scenarios. Other common answers included getting tech support from the company that developed the encryption extension and simply "I don't know."

Some Misconceptions Remain

Although most participants understood and reasoned effectively about the security properties we presented, some retained incorrect mental models that have implications for the ongoing design of encryption systems. One incorrectly believed that because he couldn't understand an encrypted message, no one else (including his email provider) would be able to, either. Several mistakenly believed it was important to keep the public lock secret to ensure emails couldn't be read by attackers: "The fewer people know my public lock, the safer." Several participants also had concerns and misconceptions about how keys are managed across multiple devices, regardless of model.

Encryption Systems for Average Users

Our participants evinced a strong interest in exploring, using, and understanding the encryption models we tested. We therefore conclude with design and education suggestions to help average users interested in encrypting their messages.

Explain the Chosen Encryption Model

As shown previously, users can — if they take the time to — understand the high-level concepts and risks of encryption, and even reason about (and sometimes correct) misconfigurations. Developers and system designers should explain in clear language the high-level risks of a given encryption approach, realizing that users might prioritize security properties differently, and trust users to make decisions accordingly. Alarmed denunciations of tools that don't offer perfect privacy might only scare users away from any encryption at all, given that many users already believe encryption is either too much work or unnecessary for their personal communications. Well-designed information about the encryption system can help users make more informed decisions and take steps to protect themselves. For example, showing users a representation of their lock allows them to verify it with others. This proactive step can remind users their content is encrypted and encourage trust via self-audit.

Explanations in Context

While it's always difficult to motivate people to consider security and privacy explicitly, leveraging visual cues and in-place reminders can increase users' familiarity with how and why encryption is applied to their messages. We found that participants who performed encryption tasks before learning about security benefits were less likely to say they would use or recommend models; as such, we believe it's critical to demonstrate these benefits up front, within the context of the messaging system. Because nearly all participants said they wanted encryption in some circumstances, and given recent media attention to encrypted chat apps, we believe explanations of security can be a selling point for a messaging or email system.

Increase Familiarity

One of the biggest limitations of our study, as well as most other encryption usability research, is participants' near-total lack of familiarity with encryption tasks. For example, one participant who saw the exchange model first continued to expect the annoyance of managing correspondents' locks even when dealing with the registration model later. This novice confusion might contribute to the effects found in our and other studies, making it difficult to model long-term usability. Additionally, if people become more accustomed to encryption overall, we can begin to address subtler misunderstandings that affected our participants' ability to make informed decisions (for example, the safety of widely publishing a public lock). We believe our findings demonstrate the benefits of continuing to try to educate users.

Explore Encryption Education Channels

Our participants were directly instructed to read our educational materials; however, real users often have neither the time nor the motivation to seek out this kind of information. This magnifies the role of journalists, security commentators, and other opinion makers whose recommendations users often rely on instead. When people do encounter encryption education (such as in media discussions, or in the context of everyday messaging tools), it's critical that this material be effective.

> **Well-designed information about the encryption system can help users make more informed decisions and take steps to protect themselves.**

We encourage future work that builds upon our findings by developing and testing instructional illustrations, videos, and demos; by refining metaphors and explanations to reduce misconceptions; and by understanding how hiding or exposing aspects of encryption systems affects users' decision making. The growing population of people who encounter and use newer features (such as lock verification) provides a strong opportunity for compelling future research.

References

1. S. Fahl et al., "Helping Johnny 2.0 to Encrypt His Facebook Conversations," *Proc. Symp. Usable Privacy and Security*, 2012; doi:10.1145/2335356.2335371.
2. S.L. Garfinkel and R.C. Miller, "Johnny 2: A User Test of Key Continuity Management with S/MIME and Outlook Express," *Proc. Symp. Usable Privacy and Security*, 2005, pp. 13–24.
3. S. Ruoti et al., "'We're on the Same Page': A Usability Study of Secure Email Using Pairs of Novice Users," *Proc. Conf. Human Factors in Computing Systems*, 2016, pp. 4298–4308.
4. S. Ruoti et al., "Confused Johnny: When Automatic Encryption Leads to Confusion and Mistakes," *Proc. Symp. Usable Privacy and Security*, 2013; doi:10.1145/2501604.2501609.
5. A. Whitten and J.D. Tygar, "Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0," *Proc. 8th Conf. Usenix Security Symp.*, vol. 8, 1999, pp. 13–24.
6. N. Weaver, "iPhones, the FBI, and Going Dark," *Lawfare*, blog, 4 Aug. 2015; https://www:lawfareblog:com/iphones-fbi-and-going-dark.
7. S. Gaw, E.W. Felten, and P. Fernandez-Kelly, "Secrecy, Flagging, and Paranoia: Adoption Criteria in Encrypted Email," *Proc. Conf. Human Factors in Computing Systems*, 2006; doi:10.1145/1124772.1124862.
8. S.E. McGregor et al., "Investigating the Computer Security Practices and Needs of Journalists," *Proc. 24th Conf. Usenix Security Symp.*, 2015, pp. 399–414.
9. M.S. Melara et al., "CONIKS: Bringing Key Transparency to End Users," *Proc. 24th Conf. Usenix Security Symp.*, 2015, pp. 383–398.
10. M.D. Ryan, "Enhanced Certificate Transparency and End-to-End Encrypted Mail," *Proc. Network and Distributed System Security Symp.*, 2014; https://eprint.iacr.org/2013/595.pdf.
11. S. Fahl et al., "Confidentiality as a Service — Usable Security for the Cloud," *Proc. IEEE 11th Int'l Conf. Trust, Security, and Privacy in Computing and Communications*, 2012 pp. 153–162.
12. W. Tong et al., "Why King George III Can Encrypt," *Freedom to Tinker*, blog, 2014; https://freedom-to-tinker.com/2014/06/06/why-king-george-iii-can-encrypt.
13. L.J. Camp, T. Kelley, and P. Rajivan, *Instrument for Measuring Computing and Security Expertise*, tech. report TR715, School of Informatics and Computing, Indiana Univ., 2015.

Wei Bai is a PhD student in the Department of Electrical and Computer Engineering at the University of Maryland. His research interests include network security and privacy with an emphasis on human factors. Bai has an MS in electrical and computer engineering from the University of Maryland. Contact him at wbai@umd.edu.

Doowon Kim is a PhD student in computer science at the University of Maryland. His research interests include security and privacy, as well as usable security. Kim has an MS in computer science from the University of Utah. Contact him at doowon@cs.umd.edu.

Moses Namara is a PhD student in human-centered computing at Clemson University. His research interests include usable privacy and security, and human factor issues related to the design of privacy-enhancing technologies. Namara has a BSc in computer science from the University of Maryland. Contact him at mosesn@clemson.edu.

Yichen Qian is an undergraduate student majoring in computer engineering and minoring in advanced cybersecurity experience for students at the University of Maryland. He expects to graduate in 2018 and plans to pursue an MBA with a concentration in information systems. Contact him at yqian1@terpmail.umd.edu.

Patrick Gage Kelley was an assistant professor in the Computer Science Department and the Organization, Information, & Learning Sciences Department at the University of New Mexico. His research interests include privacy, visualization, and technology ethics, through user-focused education and design. Kelley has a PhD in computation, organizations, and society from Carnegie Mellon University. Contact him at me@patrickgage.com.

Michelle L. Mazurek is an assistant professor in the Computer Science Department and the Institute for Advanced Computer Studies at the University of Maryland. Her research interests include understanding and designing tools and techniques to improve security- and privacy-relevant decision making. Mazurek has a PhD in electrical and computer engineering from Carnegie Mellon University. Contact her at mmazurek@umd.edu.