Effective Communication of Cyber Security Risks

Dr. Jason R. C. Nurse

jason.nurse@cs.ox.ac.uk
http://www.cs.ox.ac.uk/people/jason.nurse

Cyber Security Centre,
Department of Computer Science,
University of Oxford
Oxford, United Kingdom

Abstract

By now, it should be well-known that technology alone cannot solve the security problem. A central component to achieving security in organisations, businesses and home environments is supporting user awareness of security and designing security functionality that is highly usable. In this paper, we concentrate on this problem with special emphasis on the increasingly important issue of how to effectively communicate cybersecurity risks. This focus is motivated by the prevalence and success of online attacks, the large amount of people online nowadays (some aware of security risks, but a majority, complete novices), and the wide spectrum of activities that users engage in online, that is, everything from banking to social media (each with its own security risks). Specifically, this work reflects on our recent research on addressing this concern through focus on the fields of Information Trust, Risk Communication and Security Usability. The outcome of that reflection has been the definition of some key recommendations for trustworthy and effective communication of cybersecurity risks that can be applied across a variety of security contexts (e.g., online interfaces, security tools and security operation centres). We present a subset of these in this paper, with special note to some of the most important ones for system designers. The next step of our work is to critically evaluate these recommendations and refine them where possible, towards creating the most useful security communication practices. We believe that once adopted, these practices will have a significant positive affect on the decisions that user and individuals make regarding security risks online.

Keywords: cybersecurity risk, security usability, information trustworthiness, recommendations.

1 Introduction and Motivation

Although it has taken some time, industry and academia are gradually appreciating the importance of the human element as it pertains to achieving adequate security and trust in systems. Typical examples of this appreciation can be seen through the investments in security awareness / understanding training for employees within businesses [1], and the growing emphasis (via research projects and publications) on the topic in the academic domain [2, 3]. There are several driving forces behind this development, but undoubtedly, one of the most significant is the increase in attacks which exploit humans and again highlight why they are commonly regarded as the weakest link in the security chain. Attacks have not only increased in quantity but also in sophistication and complexity. Spear-phishing and targeted social engineering are good illustrations of prevalent attacks [4, 5] that have been found to persuade individuals to provide everything from user names and passwords to banking and financial details.

Another driving factor is the difficulty in using security in applications and systems, and security software. Simply put, a large amount of systems do not design for or prioritise usable security. Early work such as "Why Johnny can't encrypt" [6] discussed this aspect in detail, and recent work titled, "It's too complicated, so I turned it off!" [7], continues to highlight the issues. The consequence of poorly conceived and confusing security designs is that users then adopt measures to circumvent security controls, or simply ignore or badly configure them. This can be viewed in a range of areas, from the configuration of home routers (often too difficult for the less technical to setup therefore default configurations and passwords are commonly used), to the posting of complicated passwords (routinely required by today's systems and to be changed every two weeks) on post-it notes stuck to the bottom of office desktop screens.

At times users are well-aware of the consequences of careless security actions, however, in a significant amount of cases there is no proper communication of the security risks by systems and applications. This unfortunately leads to a lack of understanding and subsequently, poor decisions. Research articles in [7-9] provide evidence of the challenges as it relates to firewalls and several other security packages, but there are countless more examples that can be seen in the literature with security functionality in browsers, word processing software and encryption tools. Combining the issues discussed above, one can begin to see the depth and breadth of the security awareness, communication and understanding task. For businesses, this is even more crucial as breaches in security due to misconfigurations, spear-phishing or careless employee behaviour can result in loss of customer data, large government fines and irreparable damage to the organisation's reputation.

The remainder of this paper is structured as follows. In Section 2 we present our approach to solving the security communication problem, through reflection on the fields of Information Trust, Risk Communication and Security Usability. Section 3 pulls out key contributions from those fields and defines a number of recommendations for the communication of cybersecurity risks. Finally, Section 4 concludes the paper and presents our planned aims for future work.

2 Steps to Solving the Security Communication Problem

As a result of the concerns discussed in the previous section, there is an acute need to effectively communicate the security risk of actions and decisions to system / software users in the hope of facilitating understanding and reducing the successfulness of these attacks. In our research, we have considered this topic in detail and assessed three key areas towards defining an appropriate way forward to solving these security problems. The first area was that of Information Trust. Understanding how individuals perceive interfaces (or information) and decide whether or not to trust them is an essential initial component of risk communication. If employees do not trust the source of a security bulletin regarding the consequences of sharing passwords, or trust that a browser security warning is justified, they will ignore it. For this task therefore, we undertook a comprehensive analysis of the trust domain and the published articles with the special aim of identifying factors that influence people's decisions on information's trustworthiness; see [10] for detail. This resulted in the definition of several influential factors including, the presentation and format of the risk message (i.e., the information), how specific it is in its coverage of the topic, its relevance, how up-to-date it was, and its consistency and believability. From this, we were able to gain a much better understanding of the motivations for trust and how trust could be built into an interface or risk message going forward.

With this basis, we then progressed to the mature field of Risk Communication. Risk Communication is the interactive process of exchanging information regarding a risk (its nature, meaning, consequences, likelihood and response options) to people to enable them to make informed judgements [11]. Although this topic usually pertains to the medical and disaster contexts, several of the arguments and principles can be applied to the security field, as shown in [8]. One example is the importance of knowing who the security risk messages will be read by and ensuring that messages are tailored appropriately for that audience. In addition to tailoring the message, there is also a need to customise how it is presented [12]; the presentation factor linking once again to the trustworthiness of information. More directly, this raises the question of whether the risk information should be presented visually (graphs, images), verbally (textual content) or numerically (figures and percentages); each of these has a range of positives and negatives. Even after one of these techniques is selected, there are questions around which specific method is used, how to appropriately frame a risk message, and what are the most apt presentations that avoid bias and inaccurate perception? Fortunately, there have been numerous guidelines proposed and successfully used [4]. Most importantly for our purposes, these identify starting points which can be picked up on, applied and evaluated in the security-communication domain.

There has already been some research on Cybersecurity-risk Communications and decision making, but as seen in the reports from Section 1, it is clear that much more needs to be done. Some of the most noteworthy articles are: the analyses on factors that influence people's perceptions on security (e.g., person's knowledge, potential impact of risk, risk controllability) [13, 14]; contributions on developing effective security interfaces and alerts [9, 15]; and research on understanding the decision and thought processes via mental models for security [16]. All of these are towards gaining a better understanding and thereby further supporting users in making well-informed security decisions. Another interesting set of articles can be viewed in [17, 18], as they focus on the psychology of security and questions such as, why do users make bad security decisions? The reasons emerging are rather intriguing and range from the realities that, users often do not believe that they are at risk, losses are perceived disproportionately to gains, cognitive biases tend to negatively influence individuals, and at times, users are unmotivated and simply desire the quickest and 'good enough' solution (as opposed to the optimal one) [17, 18]. The existence of these aspects once again stresses the importance of more support for users in terms of effective designs of interfaces and risk messages, and better security awareness.

The last area is that of Security Usability. The aim of this field is simple, i.e., to make security features, functionality and interfaces usable [4]. Unfortunately, achieving that aim has been anything but straightforward. An often cited reason for this is that these two non-functional requirements natively conflict. To take an example, good usability practice would argue for making passwords easier to use and remember, conversely, good security measures purport strong (complicated) passwords that should be changed every few weeks and ideally, never written down. The challenges to usable security are well-documented [6, 19-21] and cover a wide spectrum of problems. These span confusing and clumsy interfaces, predominant use of technical security jargon with minimal documentation, limited human working memory leading to difficulties in remembering numerous logins and passwords, and the fact that security can inhibit functionality (e.g., blocking access to a potentially malicious Web site) thereby potentially annoying users. Similar to the other domains reviewed above, there have been a host of guidelines and best practices proposed to assist with usability; these mostly draw on the developments in the software Usability and Human-Computer Interaction (HCI) fields [22, 23]. As can be inferred from the on-going research in this area however, the problem is not at all solved and even if there are advances in security usability research, these have not been widely adopted in the end-user software on offer. Several publications over the years, both from academia and industry, illustrate this (e.g., [4, 6-9, 21, 24]).

3 Recommendations for Effective Communication of Cybersecurity Risks

Having briefly reflected on the three areas of Information Trust, Risk Communication and Security Usability, this section presents a number of clear, easy to understand and use recommendations for effectively communicating security risks to a system's users. We draw on best practices where available but especially look to identify potentially useful recommendations (e.g., those from the non-cybersecurity domains such as Risk Communication) that may be worth further consideration in security; our future work, as will be discussed later, is actually on assessing these recommendations in a variety of security contexts. The set of recommendations outlined references heavily on our previous contributions in [4, 8], and respective detail and additional guidance can be found in those articles. For the benefit of readers, we also narrow our scope to present what we believe are 10 of the most relevant recommendations for communicating cybersecurity risk. These are now presented below; where appropriate, we highlight the related original recommendation, especially where we adapt proposals from the Risk Communication field.

It is crucial to plan how cybersecurity risks will be communicated. System designers should be clear on: (i) the goal of the communication (e.g., is it to educate users or draw them away from a security decision that may be too risky); (ii) what type of security messages and communication strategies would be most useful (in [25] for example, the authors emphasise strategies reliant on visuals and mental models); and lastly (iii) the characteristics (e.g., level of knowledge and education, literacy and numeracy, attitudes/beliefs about the security issue) of individuals targeted by security-risk messages (e.g., knowledgeable Web users might desire more specifics than novice users regarding a security risk posed by a potentially malicious Web site). It is also important to explain possibly unfamiliar terms or complex security aspects – if users are not able to properly understand a risk, it is unlikely they will appropriately treat it. We note that current tools arguably do not allow for much personalisation (even if only in selecting a pre-set configuration level) and thus, generally operate on a one-size-fits-all basis.

Design with the understanding that humans possess a limited processing capacity. Designers should focus on reducing the cognitive effort required by individuals in processing security-risk information and/or security-related interfaces [26, 27]. This may be done by cutting back on the initial amount of security details, and as much as is possible, keeping communications simple [28]. This suggestion will need to be tempered by the current context, as certain users (e.g., experts or security analysts/administrators whose job it is to monitor all levels of system security) may prefer to be presented with detailed information initially. The presentation and format, relevance and specificity of information also become key factors in increasing a user's trust in a security-risk message displayed [10]. Methods that appreciate all of this recommendation's aspects may be deployed in practice but we can find only a few somewhat related studies on general performance and effectiveness.

The meaning of information presented in security-risk messages should be clear. To facilitate comprehension and build trust in a risk message (e.g., a firewall security warning), there is the need for the message information to be specific and unambiguous. Unclear messages are more likely to be disregarded and ignored, especially if they inhibit the user from their core task on the system. Another point to consider is appropriate message framing, i.e., how the information is expressed or stated. This includes assessing whether positive (e.g., there is a 96% chance a Web site is legitimate) framing is more suitable than negative (e.g., there is a 4% chance the site is malicious) framing and vice versa [27].

Users should be presented with clear and consistent directions for action, i.e., options to respond to a security risk faced [29]. Comprehending a risk message is crucial but providing users with options for response is also a key part of the puzzle. Designers should therefore aim at assisting users in understanding and visualising what the actual result of a security-risk decision may be like. This is particularly pertinent in situations where users may be faced with unfamiliar choices. Narratives (descriptions with a resulting outcome, such as, this actual will lead to an increased potential for the system to be compromised; or if the flagged file is indeed malware, installing it may result in disruption of normal system services and use, invasion of privacy, and so on) may also be helpful here in facilitating an appreciation of the risk and its severity. [27]

Limit use of technical and security-specific terms and jargon. To use security features, users have to be able to understand what they mean. With this in mind, designers should use technical and security-specific terms sparingly and where they are used, consider giving explanations. This is particularly useful for end-user systems and novice users. [21]. Of course however, the situation may be different in software to be used by security specialists, but even then, descriptions or quick to reference term glossaries may be very useful.

Be mindful when communicating cybersecurity risks numerically. When communicating cybersecurity risks numerically, there are a few points worthy of note. For example, users with high-numeracy levels are likely to pay more attention to risk figures, while low-numerate users may rely more on emotions, mood states and expert guidance [30]. Additionally, to avoid individuals dismissing small risks (e.g., 1% or less) or risks from familiar events (e.g., security information or warning messages from a particular source), an explicit and noticeably different message to this effect should be used [26]. Further recommendations on this format of communication can be found in [30].

Be mindful when communicating cybersecurity risks visually. For the visual communication of security risks, designers should note: (i) no single visual will work perfectly in all situations – icons, indicators, graphics and charts all have slightly more useful application contexts [31]; (ii) to promote educated judgements, displays should be representative of actual quantities/probabilities [26], this is particularly relevant if showing security-risk levels or virus infection statistics graphically (e.g., in pie charts for risk reports); (iii) if graphs are used (e.g., to show attack likelihood), these and any conclusions that might be drawn from the visuals should be explained clearly and not left up to an individual's sole interpretation [26], this would reduce the likelihood of misinterpretation and overly subjective judgements.

Be mindful when communicating cybersecurity risks verbally. When communicating cybersecurity risks verbally, it may be best to allow multiple formats to present security-risk information as various authors have expressed that verbal messages are not to be completely relied on [32]. This is especially relevant for security as it is common to see messages quoting that attacks are *likely* or *probable*. The core issue therefore is how to ensure that these terms mean the same to all users. The second aspect to be aware of is that of context and its potential influence on user perceptions [32]. Context might span who the intended system users are, where they are, what they are likely to be doing in the system, and the gravity of the security decision they currently face.

Provide help, advice and documentation for security. When necessary, users should be able to easily locate and view help manuals and system documentation for cybersecurity functions. If users cannot find or determine how to use these features, they are likely to be avoided [21, 34]. Hand-in-hand with this aspect is communicating informative security feedback to users when appropriate. Feedback should be clear, informative, sufficient, not too technical and where appropriate, give suggestions for going forward and responding to a current security risk [34]. A good example is the use of a function key (traditionally, F1) within a security screen to quickly load documentation for that security function.

Make security functionality visible and accessible. Similar to other application features, security should be visible and easily accessed. Hiding cybersecurity functionality within advanced or disparate parts of an interface are likely to make the user's task more difficult and ultimately hamper system usability. Another important aspect here is that users should be made aware of the current security state of the system. In many ways this is a form of passive feedback of cybersecurity. Some simple examples include, the word "Secured" on some encrypted or password-protected documents, active icons when security functions are being executed on a system, and padlocks within browsers to indicate browsing using Secure Sockets Layer (SSL)/Transport Layer Security (TLS); also potentially a prompt when browsing takes one away from a secured site. [21, 33, 34]

Having listed and briefly discussed the main recommendations, the next section concludes this paper and presents our directions for future work.

4 Conclusions and Future Work

In this paper, we re-iterated the importance of the human element as it pertains to achieving adequate security and trust in systems, and briefly reviewed relevant research developments in that domain. Our work assessed and reflected on the fields of Information Trust, Risk Communication and Security Usability for inspiration, before defining several recommendations. In this paper, we presented ten of the most crucial and relevant of these recommendations targeted at increasing the effectiveness of cybersecurity-risk communications.

The next stage of our research focuses on the critical evaluation of the full set of recommendations, particularly those adapted from the Risk Communication field. Validation of these recommendations is imperative if we are to properly support users in understanding security and risks, make informed decisions, and ultimately reduce the successfulness of online attacks. We envisage that this investigation will involve several progressive steps. These are: the identification of a set of case scenarios where various facets of cybersecurity-risk communication could be assessed, the development of prototype systems and/or add-on functionalities in line with scenarios to provide a practical basis for evaluation, and finally, in depth user studies to critically investigate the effectiveness and trustworthiness of cybersecurity-risk communications incorporating the proposed recommendations.

Topics such as the numeric and verbal communication of cybersecurity risks and personalisation for perceptual and individual factors are especially of interest, as these have not been addressed in great detail as far as it relates to this research field. We have already begun similar trials in the use of information trustworthiness advice in decision making and had promising results [35]; communicating the trustworthiness of information is arguably analogous to conveying the risk associated with users acting on that information. Furthermore, some of our other recent work has centred on understanding the risks to security and privacy in online social media [36], an extremely topical concern given the enormous amount of users on these sites. With this foundation, we aim to develop guidelines, practices and tools to properly communicate the associated risk information to users in the hope of modifying their 'bad' behaviour (which may relate to oversharing practices, leaking of sensitive business information online, social influence, and so on).

References

[1] PricewaterhouseCoopers LLP.: Cybercrime: protecting against the growing threat Global Economic Crime Survey, 2011.

[2] KPMG: The five most common cyber security mistakes: Management's perspective on cyber security. 2013.

[3] Blythe, J., Camp, J. and Garg, V.: Targeted risk communication for computer security," in 15th International Conference on Intelligent User Interfaces, 2011, pp. 295–298.

[4] Nurse, J. R. C., Creese, S., Goldsmith, M. and Lamberts, K.: Guidelines for Usable Cybersecurity: Past and Present, in Proceedings of the 3rd International Workshop on Cyberspace Safety and Security (CSS), 5th International Conference on Network and System Security (NSS), IEEE. pp. 21-26, 2011.

[5] Sophos Ltd. Security Threat Report 2013: New Platforms and Changing Threats, 2013.

[6] Whitten, A. and Tygar, J. D.: Why Johnny can't encrypt: A usability evaluation of PGP 5.0, in 8th USENIX Security Symposium, pp. 169–184, 1999.

[7] Raja, F., Hawkey, K. Jaferian, P., Beznosov, K. and Booth, K. S.: It's too complicated, so I turned it off!: expectations, perceptions, and misconceptions of personal firewalls, in 3rd ACM workshop

on Assurable and usable security configuration, pp. 53–62, 2010.

[8] Nurse, J. R. C., Creese, S., Goldsmith, M., and Lamberts, K.: Trustworthy and effective communication of cybersecurity risks: A review, in Proceedings of the 2011 Workshop on Socio-Technical Aspects in Security and Trust (STAST) at 5th International Conference on Network and System Security (NSS), IEEE. pp. 60-68, 2011.

[9] Ibrahim, T., Furnell, S., Papadaki, M. and Clarke, N.: Assessing the usability of end-user security software, in Trust, Privacy and Security in Digital Business, ser. Lecture Notes in Computer Science, S. Katsikas, J. Lopez, and M. Soriano, Eds. Springer, 2010, vol. 6264, pp. 177–189.

[10] Nurse, J. R. C., Rahman, S. S., Creese, S., Goldsmith, M, and Lamberts, K.: Information quality and trustworthiness: A topical state-of-the-art review, in International Conference on Computer Applications and Network Security (ICCANS). IEEE. pp. 492–500, 2011.

[11] National Research Council (NRC) USA, Improving Risk Communication. National Academy of Sciences, 1989.

[12] Lipkus, I.: Numeric, verbal, and visual formats of conveying health risks: suggested best practices and future recommendations, Medical Decision Making, vol. 27, no. 5, pp. 696–713, 2007.

[13] Huang, D.-L., Rau, P.-L. and Salvendy, G.: A survey of factors influencing people's perception of information security, in Human-Computer Interaction. HCI Applications and Services, ser. Lecture Notes in Computer Science, J. Jacko, Ed. Springer, vol. 4553, pp. 906–915, 2007.

[14] Gabriel, I. and Nyshadham, E.: A cognitive map of people's online risk perceptions and attitudes: An empirical study, in 41st Hawaii International Conference on System Sciences, pp. 274–283, 2008.

[15] Bravo-Lillo, C., Cranor, L., Downs, J. and Komanduri, S.: Bridging the gap in computer security warnings: A mental model approach, in IEEE Security & Privacy, vol. 9, no. 2, pp. 18–26, 2011.

[16] Camp, L. J.: Mental models of privacy and security, in IEEE Technology and Society Magazine, vol. 28, no. 3, pp. 37–46, 2009.

[17] West, R., Mayhorn, C., Hardee, J. and Mendel, J.: The weakest link: A psychological perspective on why users make poor security decisions, in Social and Human Elements of Information Security: Emerging Trends and Countermeasures. IGI Global, pp. 43–60, 2009.

[18] West, R. The Psychology of security, in Communications of the ACM, vol. 51, no. 4, pp. 34–40, 2008.

[19] Yee, K.-P.: Aligning security and usability, in Security & Privacy, vol. 2, no. 5, pp. 48–55, 2004.

[20] Furnell, S.: Why users cannot use security, in Computers & Security, vol. 24, no. 4, pp. 274–279, 2005.

[21] Furnell, S.: Security usability challenges for end-users, in Social and Human Elements of Information Security: Emerging Trends and Countermeasures. IGI Global, pp. 196–219, 2009.

[22] Dix, A., Finlay, J., Abowd, G. D. and Beale, R.: Human-Computer Interaction, 3rd ed. Prentice hall, 2004.

[23] Nielsen, J., and Hackos, J.T.: Usability engineering. Boston: Academic press, 1993.

[24] Enex TestLab: Usability of Endpoint Security. 2011.

[25] Blythe, J., Camp, J. and Garg, V.: Targeted risk communication for computer security, in 15th International Conference on Intelligent User Interfaces, pp. 295–298, 2011.

[26] Lipkus, I.: Numeric, verbal, and visual formats of conveying health risks: suggested best practices

and future recommendations, in Medical Decision Making, vol. 27, no. 5, pp. 696–713, 2007.

[27] Hibbard, J. and E. Peters, E.: Supporting informed consumer health care decisions: data presentation approaches that facilitate the use of information in choice, in Annual Review of Public Health, vol. 24, no. 1, pp. 413–433, 2003.

[28] Wiedemann, P., Schutz, H. and Clauberg, M.: Lessons learned: Avoiding pitfalls in risk communication, in International Conference and COST 281 Workshop on Emerging EMF Technologies, Potential Sensitive Groups and Health, 2006.

[29] Rudd, R., Comings, J. and Hyde, J.: Leave no one behind: improving health and risk communication through attention to literacy, in Journal of health communication, vol. 8, pp. 104–115, 2003.

[30] Peters, E.: Numeracy and the perception and communication of risk, in Annals of the NY Academy of Sciences, vol. 1128, no. 1, pp. 1–7, 2008.

[31] Pattinson, M. and Anderson, G.: How well are information risks being communicated to your computer end-users?" Information Management & Computer Security, vol. 15, no. 5, pp. 362–371, 2007.

[32] Visschers, V., Meertens, R., Passchier, W. and De Vries, N.: Probability information in risk communication: a review of the research literature, in Risk Analysis, vol. 29, no. 2, pp. 267–287, 2009.

[33] Johnston, J., Eloff, J. H. P. and Labuschagne, L.: Security and human computer interfaces, in Computers & Security, vol. 22, no. 8, pp. 675–684, 2003.

[34] Chiasson, S., Biddle, R. and Somayaji, A.: Even experts deserve usable security: Design guidelines for security management systems, in Symposium on Usable Security and Privacy (SOUPS) Workshop at Usable IT Security Management (USM), pp. 1–4, 2007.

[35] Nurse, J. R. C., Creese, S., Goldsmith, M., and Lamberts, K.: Using Information Trustworthiness Advice in Decision-Making, in 2012 Workshop on Socio-Technical Aspects in Security and Trust (STAST 2012) at The 25th IEEE Computer Security Foundations Symposium (CSF). IEEE. pp. 35-42, 2012.

[36] Creese, S., Goldsmith, M., Nurse, J. R. C., and Phillips, E.: A Data-Reachability Model for Elucidating Privacy and Security Risks Related to the Use of Online Social Networks, in 11th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (IEEE TrustCom-12). IEEE. pp. 1124-1131. 2012.