DIMACS Technical Report 2007-11
July 2007

Using Cartoons to Teach Internet Security

by

Sukamol Srikwan Markus Jakobsson[1]
School of Informatics
Indiana University
Bloomington, IN 47406

[1]Markus Jakobsson is a permanent member of DIMACS. This work was supported by a Microsoft Trusted Computing grant.

ABSTRACT

While good user education can hardly secure a system, we believe that poor user education can put it at serious risk. The current problem of online fraud is exasperated by the fact that most users make security decisions, such as whether to install a given piece of software or not, based on a very rudimentary understanding of risk. We describe the design principles behind `SecurityCartoon.com`, the first cartoon-based approach aimed at improving the understanding of risk among typical Internet users. We argue why an approach like ours is likely to produce better long-term effects than currently practiced educational efforts with the same general goals. This belief is based on the apparent difference between our approach and currently used alternatives. At the heart of these differences are the four guiding principles of our approach: (1) A *research driven content selection*, according to which we select educational messages based on user studies; (2) *accessibility* of the material, to reach and maintain a large readership; (3) user *immersion* in the material, based on repetitions on a theme; and (4) adaptability to a changing threat.

1 Introduction

Online fraud is threatening organizations and individuals alike, and many fear that it can turn into a weapon of electronic warfare within the not so distant future. There is a strong consensus that we, as a society, need to improve our resilience against this threat. This goal can be reached using at least three principal approaches: *Software-based security initiatives*, *legal and regulatory efforts*, and *educational approaches*. While the approaches are complementary, they are not entirely independent. For example, legal and regulatory efforts are limited by technological issues for detection and enforcement. Likewise, the impact of client-side software initiatives is affected by educational efforts relating how to use the technology, and how to maintain the integrity of deployed software. In turn, regulatory efforts fuel software development and deployment, and recent FFEIC guidance [19] encourages financial institutions to educate their clients.

While technical efforts to fight the problem proliferate, and legal and regulatory approaches are rapidly catching up, we argue that the development of educational efforts have been left behind. Consumers are faced with a bewildering array of advice of how to stay safe against identity thieves, but we are not certain that any of the efforts manage to communicate a basic understanding of what to do and why. Current advice comes in many forms, from the terse online resources of financial institutions to in-depth self-help books describing how to to obtain access to credit reports. Consumers are advised to buy and use paper shredders; look for icons indicating that sites are hacker safe, use encryption, and that are members of the Better Business Bureau. At the same time, the typical Internet user does not know how to identify a phishing email [48], but often [29] relies on checking spelling and identifying known deceit techniques. Many consumers do not realize how easy it is to clone an existing site (e.g., using a tool like WebWhacker [57]) but interpret convincing website layout as a sign of legitimacy. It is not surprising that the average consumer has a rudimentary understanding of the threat, both due to the fact that he or she does not understand the intricacies of the Internet, and due to the difficulties of communicating complex notions to users that would rather not be involved at all. To make it worse, phishing is both a matter of technology and psychology [30, 49], and there is ample evidence (see, e.g., [39]) supporting that most people *want to trust* what they see.

Phishing education is not easy. We explain pitfalls and difficulties, and describe how a detailed understanding of these—and of typical user behavior—can help guide efforts. While such improved efforts can be expressed using a variety of available media, we describe only one particular approach in detail: A cartoon format. This particular format was selected since it is accessible to a large portion of the population and allows the use of stories to illustrate complex processes.

Outline. We begin (section 2) by explaining the difficulty of educating users about security, supporting the beliefs that underlie the educational approach described in this paper. We then describe related efforts (section 3). This is followed by an in-depth description of some problems we perceive with alternative educational approaches (section 4) and a description of our goals (section 5). In section 6, we then turn to describing our approach. Here, we give examples that illustrate how we communicate various concepts and address the problem.

2 Why Security Education is Difficult

We believe that the educational aspect of security has not been given the attention it deserves, and that many skeptics have prematurely concluded that any involvement of the end user is doomed to fail. We hope that an improved understanding of the issues surrounding both online fraud and how people relate to potential threats may help develop educational approaches with better impact, and that at least *some* risky online behavior can be significantly curbed. However, the task of reaching this goal is far from trivial, due to the complexity of phishing, crimeware and associated threats. The problem is neither entirely technical nor entirely social in its nature, but a combination of the two, and there are numerous security vulnerabilities associated with this combination.

Current educational efforts aimed at encouraging safe online behavior have limited efficacy. We argue that many educational efforts expect too much of the audience. Namely, many educational efforts aim at a rather technically knowledgeable reader, and therefore do not reach the typical Internet user. For example, the Federal Trade Commission (FTC) advises people [20] to forward suspect emails with full headers, but without explaining what a full header is or how to view or forward it. Even though everybody with a networked computer has access to information that allow them to learn how to view full headers, many may be insufficiently motivated to find and read such information. As another example, APWG [2] advises users that unless an email is digitally signed, one cannot be sure it was not forged or spoofed. However, recent studies [25] have shown that typical users do not know how to benefit from the security guarantees offered by digitally signed email; moreover, most banks are reluctant to deploy technology to sign outgoing messages due to concerns that doing this may drastically increase the number of calls made by recipients who are not familiar with signed email. Handling these calls would incur a cost that is believed to dwarf the savings obtained from the increased security. Thus, signed emails are not as common as technologists believed they would be; this is also an effect of the failure to deploy a public key infrastructure at a large scale.

At the same time as many educational efforts may expect users to be savvier than they are, many other efforts over-simplify the message in an effort to make it digestable to a general audience. For example. financial institutions often warn users not to follow hyperlinks in email messages. As of mid-2007, phishers have started to adapt to users being wary of clicking on links and suggest to targeted individuals that they copy and paste URLs into the address bar. Given that no large-scale effort has warned against this particular attack, the new twist may very well be rather successful. The use of this trick by phishers also

illustrates the need to teach people not only simplistic rules but instead to explain *why* one should or should not do a particular thing. Understanding the threat will also help people adapt to changes. Thus, to create a lasting impact of the educational efforts, it is crucial not to teach people to recognize known phishing attempts, but instead, to recognize the patterns behind them. This allows people to generalize and truly understand, but takes an even greater effort—both on behalf of the educator and the user.

Some overly simplified educational messages may elevate the risk compliant users expose themselves to. As an example of this, consider the suggestion "you should only communicate information such as credit card numbers or account information via a secure website or the telephone" [2]. While this advice would reduce the vulnerability to many attacks that are common at the time of writing, it may also increase people's vulnerability to vishing attacks, and to phishing/vishing hybrid attacks in which the user is sent an email that advises her to call the impersonated service provider using a number supplied in the email. This type of attack may become common if takedown efforts become increasingly effective. This example shows that some advice may cause problems if users follow it without much critical thought, which in turn suggests that simple rules may not be suitable for education in which there is an attacker who may try to leverage his efforts on the approaches of the educational campaigns. This observation indicates that phishing education may be harder than other types of education, since it requires the user to act against a threat that in itself is molded by what users understand. While software can be constantly patched to harden it against an evolving threat, it is harder to re-educate users to counter such changes.

Coming back to the conflicts between too demanding education and overly simplified education, we see that at the heart of the dilemma is the fact that *simple* rules do not capture the problem well while *complex* rules do not captivate the audience. It is not easy to explain complex problems to users that often will not be highly motivated to learn until they have already suffered a bad experience. In addition, most people think they are being reasonably careful, and we all like to think that bad things happen to others only. As a result, current mainstream efforts do not attempt to cover advanced topics, recognizing that the typical user would be unlikely to spend enough time to understand anything but the basics. For example, while research [28, 29] has identified a common inability to distinguish legitimate domain names from those used in cousin-name attacks or sub-domain attacks, these are not topics that are addressed in educational efforts targeting typical Internet users. Examples of deceptive URLs of these types are www.citibank-secure-commection.com and www.chase.pin-reset.com, both of which have been shown [29] to inspire confidence among typical users. The same study showed that short URLs inspire more confidence than longer alternatives. There are plenty of examples of unnecessarily long URLs, and we do not even have to go beyond our bibliography to find such examples [12].

Another example of a neglected topic of high security relevance is malware. While almost all educational efforts indicate the importance of having up-to-date anti-virus software installed, typical service providers fail to explain the complexity of malware to their clients, or how to defend themselves against the threat. As an example, most people do not fully appreciate that software that *they* opted to install—maybe because a friend suggested to do

so—could in fact be malware. Still, this is a serious security concern: A recent study [52] demonstrated a yield of more than 50% for for one particular type of socially propagated malware attack.

A problem that exasperates the effort of educating users of security is that it is not sufficient to explain the problems to the target audience, but one must also change their behavior. It is often ignored that there is a tremendous discrepancy between what typical users *know* and what they *practice*. An example of this is illustrated by the recent studies involving eyeball tracking [58], in which it was concluded that most users rarely look for SSL indicators, much less choose to interact with these. Yet, most people understand that there is an association between security and the presence of an SSL lock. To make it worse, many phishing attacks use SSL [41], so teaching users to look for an SSL lock may not be the best approach. Security education has to highlight the need for behavioral changes. We believe this is best achieved by illustrating the causality between behavior and security outcome, a task for which the cartoon format is well suited.

Many educational efforts—and in particular those employed by financial institutions—offer highly concentrated advice, often on an itemized format. While the material is correct and almost always constructive, it explains *what* to do as opposed to *what and why*. As such, it may be a better resource for users who already understand the problem than the larger masses who do not. We think of this type of resource as a dictionary. Whereas nobody would attempt to learn a new language from a dictionary alone, dictionaries fill an important role for people with a reasonable background understanding.

Most existing techniques are best suited to deal with a static threat, which phishing arguably is not. There are several reasons for this. Education that is not often accessed by a given person (such as what is provided by financial institutions) fails to communicate changes simply due to the fact that the target audience only accesses the material relatively infrequently. (While financial institutions, for example, could provide their clients with updated advice as often as they wish, this could backfire by breeding fear, and thereby make people withdraw from online banking.) Game based approaches, on the other hand, suffer from a high structural cost associated with structural changes. Cartoons are lightweight enough to permit integration of new material at a low cost, and can be designed to avoid the intimation factor that lists of abstract do-and-don'ts typically seem to suffer.

While it is commonly agreed that typical Internet users do not well understand how to stay safe against phishing, it has also been shown [32] that available training material is surprisingly effective *when used*. A reasonable conclusion one can draw from this is that typical consumers do not make sufficient use of available material. To entice the typical banking client to spend hours instead of seconds learning about security, it is necessary to make the educational message both accessible and enjoyable. This paper describes a cartoon-based effort [51] aimed at reaching this goal, and describes our underlying design principles. These relate to the manner in which material is selected; how the material is made accessible to a typical reader; and how repetition on a theme is used to create an immersion approach.

3 Related Work

There is an increasing literature trying to educate typical users of threats. Consumers are offered a range of books (e.g., [6, 8, 13, 22]) explaining identity theft, how to stay safe, and how to recover from it. These books often focus on a taxonomy of the common scams, e.g., what happens in a 419 scam and why one must secure wireless networks. The exposition requires a fair amount of dedication from consumers for them to make substantial progress.

There is a variety of online resources aimed at improving the understanding of risk among the public. Banks (e.g., [12, 11]) and companies like eBay [16], Microsoft [37], and PayPal [45] also supply their own pages to warn users about the dangers of phishing and how to avoid falling prey. Pages like eBay's Spoof Email Tutorial [17] teach basics about email spoofing as well as technical details regarding URLs and domain names, but with a very clear emphasis on recognizing spoofs on eBay, and with a focus on commonly existing psychological twists only. Paypal [46] allows users to take a five-question test that is intended to identify risky behavior. However, it is likely that a large portion of users get full score. This is not only due to the simplicity of the questions, but is also the result of the subject-expectancy effect — this is a cognitive bias that occurs when a subject expects a given result and therefore unconsciously manipulates an experiment or reports the expected result. There is a risk that users taking the test come away feeling overly confident of their abilities, rather than slightly humbled and intrigued to learn more.

Most financial institutions maintain a list of recommendations, typically on bullet form and fitting on one page. The advice is useful, but seldomly communicates any understanding of *why* it is useful. It has been criticized [14] as not being a helpful educational tool. As such, it may often be ignored or forgotten, especially under duress, which is a commonly used approach of phishers. Apart from banks, many non-financial service providers also provide guidance. AT&T maintains sites [3, 5, 4] that stands out for having a larger degree of interactivity than most, which is believed to promote learning.

A recent academic effort [31] attempts to teach phishing awareness to a general audience using a computer-game based approach; their approach shares the third design value (accessibility) with our approach, but achieves this goal in a different way, and using true user interaction. We believe that the game-based approach may have many benefits, but that it may fall short in that it may reach a somewhat limited demographics at this point. In related efforts, [4, 42], AT&T and NetSmartzKids use games to teach Internet security to children, and [26] uses animation to reach the same goal.

Another approach that relies on user interaction to assess risks is the so-called *phishing IQ test* [35] that scores the user's ability to recognize phishing instances. This approach, which can be thought of as an indirect educational approach, however, has been criticized [9] for not measuring the ability of the test taker as much as his or her fears of phishing. Embedded education [34] teaches people about phishing during their normal use of email.

There are also several movie-based efforts to teach security awareness, see e.g., [18, 56]. These particular examples are highly technical, and are only accessible to expert audiences. Other movies, such as [33], are easier to access, but are also relatively brief, and therefore only

communicate a rather limited understanding of the threats and associated countermeasures. The prospective of an extensive effort using a movie format to communicate threats appears to offer significant benefits.

Some educational efforts, such as [50], cover slightly irrelevant educational messages, such as encryption over wireless networks. While encryption is helpful to avoid eavesdropping, it is not a form of access control, and does not secure the access points that use it. It is not clear that the educator always has a clear understanding of the threats. This is a pity, given the recent wealth of insights into the nature of the problem of phishing. In particular, there have been several recent advances towards understanding how typical users react to deceit [15, 23, 24, 29, 30, 59].

We use cartoons as the main media of communication. Graphical representations have advantages, given that words only have context in the culture of the speakers. That is certainly true when it comes to communication of technology-related messages. Scientists are lost in their own sub-culture based vocabulary. In contrary, the concept that "a picture speaks a thousand words" has been used in numerous settings to communicate educational messages of importance. Historical, political and social issues have been expressed via editorial cartoons in newspapers and magazines for more than 150 years [44]. Very few people, regardless of their culture and nationality, would not recognize a picture of a soldier and a tank and the associated concept of war. Such images brings alert and worry to most people, young or old, a fact that highlights that humans are visual by nature. Our minds cannot avoid making association between images and experiences. Cartoons can be drawn to enhance specific messages while suppressing others, which provides an excellent tool for communicating complex concepts. A simple cartoon is a form of amplification through simplification [36].

HIV/AIDS education using cartoons was done in the New York subway for a long period of time, in a way that was very clearly targeted to teenage readers. Copyright law is taught using a cartoon approach [7]. Teaching security, however, is slightly different from teaching other topics since—given the adversarial setting—one must assume that the adversary has access to the educational material and will do he best to find ways of taking advantage of the way material is (or is not) taught. This must be a consideration when designing the educational effort: The effort may, in itself, change the nature of the topic associated with the education.

Some material intended for the public to better understand the threats associated with identify theft focus on server-side issues, and does not make recommendations to the reader. An example of this is [1]. While there may be no direct modification of behavior as a result of this, it may help bring the problem to the attention of the general public, which may in turn facilitate legislative efforts. This is another important role of education.

4 Some Problems with Practiced Approaches

Following the principle that we learn from our mistakes, we review some common drawbacks associated with mainstream educational techniques. We classify this into general categories, and give examples to help make these categories concrete.

- **Advice that is hard to follow.** Many current educational efforts suffer from giving advice that is hard to follow, and which may, as a result, become entirely redundant at some point. An example of type is the common piece of advice *"do not click on links"* – almost all companies with an Internet presence send email with hyperlinks. Many send unsolicited emails, and a large number of them send such emails to users who have not signed up for their services. The advice not to click on links is helpful to security in some instances, but following it would hamper the web browsing experience to many users would find it hard to follow. Similarly, consumers are often advised to disable javascript. This, too, is advice that is hard to follow, given that web 2.0 applications rely on it being enabled. Consequently, more than 95% of all Internet users have it enabled. Another example of advice that is hard to follow, taken from [47], suggests to the user to *"Disable JavaScript or do not visit untrusted and trusted websites at the same time."* Whereas many users would be able to identify a clearly trusted site (say, their bank) as well as many clearly untrusted sites, few would be able to correctly identify randomly chosen sites. Most nobody can determine whether a site is trustworthy or not before having visited it.

- **Valid but not very important advice.** Some advice is valid, but not very central to the security of the user. If we believe that there is an abundance of important advice to be given, and that the typical user can only retain a given amount of information, then this type of advice would amount to a lost opportunity. An example of this kind of advice is [50]: *"Enable encryption on wireless routers immediately upon setting up a home network."* While this protects against theft of service, it has little impact on the greater problem of phishing, and is probably entirely overshadowed in terms of importance by a piece of advice that was not given: *"Select a good password for your router immediately upon setting up a home network. Never use the factory-set password, as strangers on the Internet can easily look this up and access the security settings of your router."* Typical Internet users are unaware of this threat, and do not know that if their router becomes corrupted, then it can steal passwords [40], mount pharming attacks [54], and block access to anti-virus updates [53]. Most users do not realize the ease with which a router can be infected [54, 55, 53].

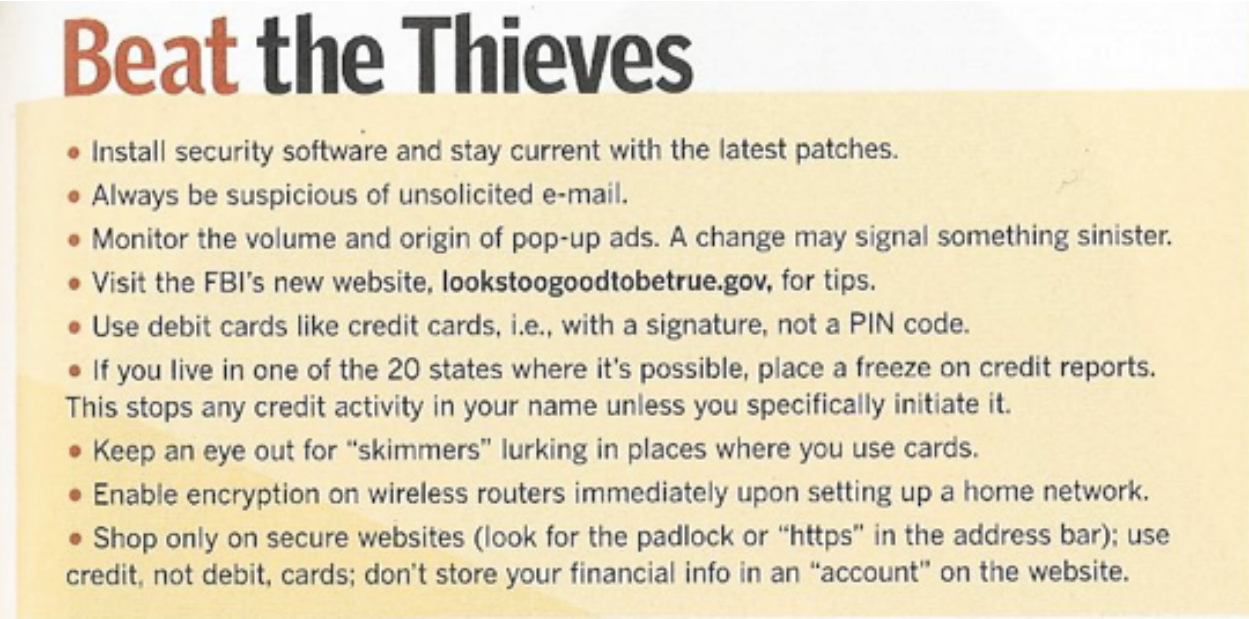

Beat the Thieves

- Install security software and stay current with the latest patches.
- Always be suspicious of unsolicited e-mail.
- Monitor the volume and origin of pop-up ads. A change may signal something sinister.
- Visit the FBI's new website, **lookstoogoodtobetrue.gov,** for tips.
- Use debit cards like credit cards, i.e., with a signature, not a PIN code.
- If you live in one of the 20 states where it's possible, place a freeze on credit reports. This stops any credit activity in your name unless you specifically initiate it.
- Keep an eye out for "skimmers" lurking in places where you use cards.
- Enable encryption on wireless routers immediately upon setting up a home network.
- Shop only on secure websites (look for the padlock or "https" in the address bar); use credit, not debit, cards; don't store your financial info in an "account" on the website.

Figure 1: A sidebar from a phishing education effort in the Reader's Digest [50]. A user who is made nervous about patches of security software, but who does not have a good understanding of what is a patch and what is not, may be deceived to install undesirable software that claims to be patches. In this case, it may be better to instead teach how to configure one's machine to perform automated patching, or to explain the more complex picture of malware. Similarly, while it is good to be wary of unsolicited email, the typical user receives the user contradictory advice in the form of large quantities of legitimate but unsolicited email—sometimes even from the same institution that asked the user to be suspicious of such email. The third piece of advice is well-meaning, but hard to follow for the typical user. What does it mean to monitor the origin of popup ads? What should be done if there is a drastic change? Further down, the user is encouraged to enable encryption on his or her router, but the much more important advice to set a strong password is left out.

- **Potentially dangerous education.** Some advice may be both valid and important, but still pose a potential risk to users. In figure 1, the reader gets the advice *"Install security software and stay current with the latest patches."* However, the reader is not told how to practically follow the advice, and may become more vulnerable to an email that tells the user to install a given security patch. If such an email is spoofed to appear to come from the system administrator of the user in question, then it may appear fully reasonable to many to quickly install the required "patch". To make sure that the advice does not pose a risk to the user, it may sometimes be important to qualify exactly how to follow the advice, and what not to do. This is particularly important given that phishers are likely to try to take advantage of habits that his intended victims have been trained to acquire. Another example of this type is to ask users to look for an SSL lock, implicitly equating the presence of that with security.

This is a problem given the ease for phishers also to use SSL [41].

If we see the creation of habits as a form of education, then another example of potentially dangerous education is a technique financial institutions commonly use to authenticate themselves to their clients. Namely, many financial institutions authenticate themselves to clients by including "E-mail intended for your account ending in: XXXX" or a similar text in email communication. However, and as was shown by [29], common Internet users do not find a statement such as "E-mail intended for your account beginning with: XXXX" less trustworthy than "E-mail intended for your account ending in: XXXX" in spite of the fact that the four first digits are possible to infer from the identity of the issuer (a fact that will not be lost on the phisher). Therefore, training users to accept "last four" as a form of authentication makes them vulnerable to an attack in which the "first four" is used—at least in the absence of a carefully crafted educational campaign aimed at addressing this potential problem.

- **Advice that is not absorbed.** The most significant problem of many current approaches, though, might be that they do not entice the intended readership to spend considerable time and effort understanding the messages. A recent study [] found that existing phishing education succeeds in communicating important security information, *if accessed*. The latter is the gist of the problem: Typical user education is *not* accessed by the typical consumer, or is, at the very least, not given hours of attention and afterthought. As an analogy, we argue that many currently used approaches for teaching Internet security to typical consumers are like an ski lesson in which there is no practice or feedback, but where the lesson consists of the single message *"Don't fall."*

5 Goals

Our mission can be formulated, quite simply, as "Help typical users improve their security against online threats". This is not quite as straightforward as it may seem, but can be broken down in the four following partial goals:

Goal 1: Know your audience. It is important to understand what the typical user's limitations are—both in terms of what he or she knows, and what he or she can reasonably be taught. As we have argued, it is a common mistake to either expect too much from the user, or to expect so little that the resulting educational message is reduced to soundbites.

It is tempting for security specialists to attempt to guess what problems typical users may have based on what problems they themselves perceive. This, however, is a fallacy, and one that led many to believe that the existence of SSL would lead to the failure of phishing as soon as awareness of the problem rose. User experiments, while often challenging to perform [21], provide a better understanding of the likely vulnerabilities.

Goal 2: Know your enemy. It is equally important to recognize what the threats *are*, and what they may *become*. Namely, it is not helpful to only secure users against current threats, given the proven malleability of the threat. To understand the threats and their likely development, it is crucial to recognize that there is both a technological component and a human factor. The educational effort, too, must take both of these aspects into consideration.

Knowing how typical users relate to both user interfaces and to deceit helps anticipate potential trends in phishing. This can also be done by making technological hypotheses, from which one may infer the likely resulting threats. See [27] for an example of this approach.

Goal 3: Know how to tell the story. A third goal is to enable communication of insights. To make this possible, the material to be taught needs to be *accessible*. Many researchers agree that security is not a primary goal of the typical user (see, e.g., [38]). The educational technique must take this into consideration by making the material accessible to the target audience. This is closely related to knowing one's audience.

Goal 4: Know how people learn. A fourth and final goal is to facilitate learning. It is not sufficient for the educational material to be appropriately selected and communicated, since learning is not an instantaneous process. To maximize the benefit of the material, repeated exposure is necessary, and the material must be designed with this in mind.

6 Our Approach

Having identified the goals to be achieved, we have also developed an approach to reach them. Our approach, which is examplified in figures 2 to 6, is based on four core principles: (1) Research driven content selection; (2) accessibility; (3) immersion; and (4) adaptability. We will review these below, and give examples describing how our cartoon approach embody these principles.

Research Driven Content Selection: The educational material is tailored around an understanding of actual user behavior, as observed in a collection of studies. For example, it was shown in [28] that average users know that an IP address (instead of a regular domain name) signals danger, but do not have any understanding of the risks of cousin-name or subdomain attacks. Similarly, another study [52] showed a success rate of more than 50% for a simulated attack in which victims received an email from a friend in which they were told to run a program with a self-signed certificate. This suggests that typical users do not distinguish well between *trusting their friend's taste* and *trusting their friend's ability to identify a threat*. Correspondingly, our approach describe both user-propagated threats and the risks of user-installed programs having malicious content.

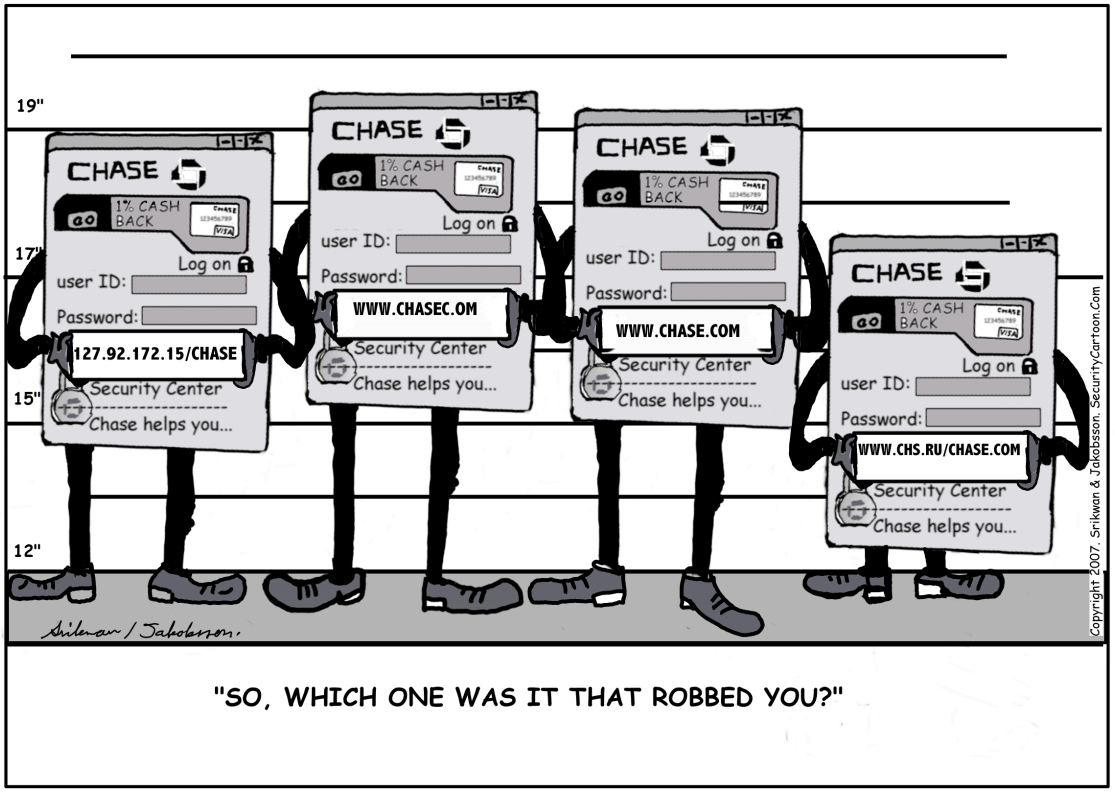

Figure 2: This cartoon strip intends to increase the understanding of URLs, which is a difficult topic to teach. Typical computer users do not understand the difference between domains and sub-domains, and are easily fooled by cousin-name attacks, as supported by [28]—this is an example of how we use research driven content selection. None of the large financial institutions attempt to teach their clients about URLs, in spite of being and important topic for people to master in order to stay safe. One reason might be that URLs are complex, and an exhaustive set of rules describing their nature would be daunting for the typical consumer. The strip shows four Chase webpages in a line-up. The reader is asked who robbed him, forcing him to start comparing them. The only real difference between them are the URLs. This is intended to reinforce the notion that copying logos are functionality is not beyond the reach for phishers, and to force the reader to start comparing the URLs. This is an example of how to achieve accessibility of a difficult topic.

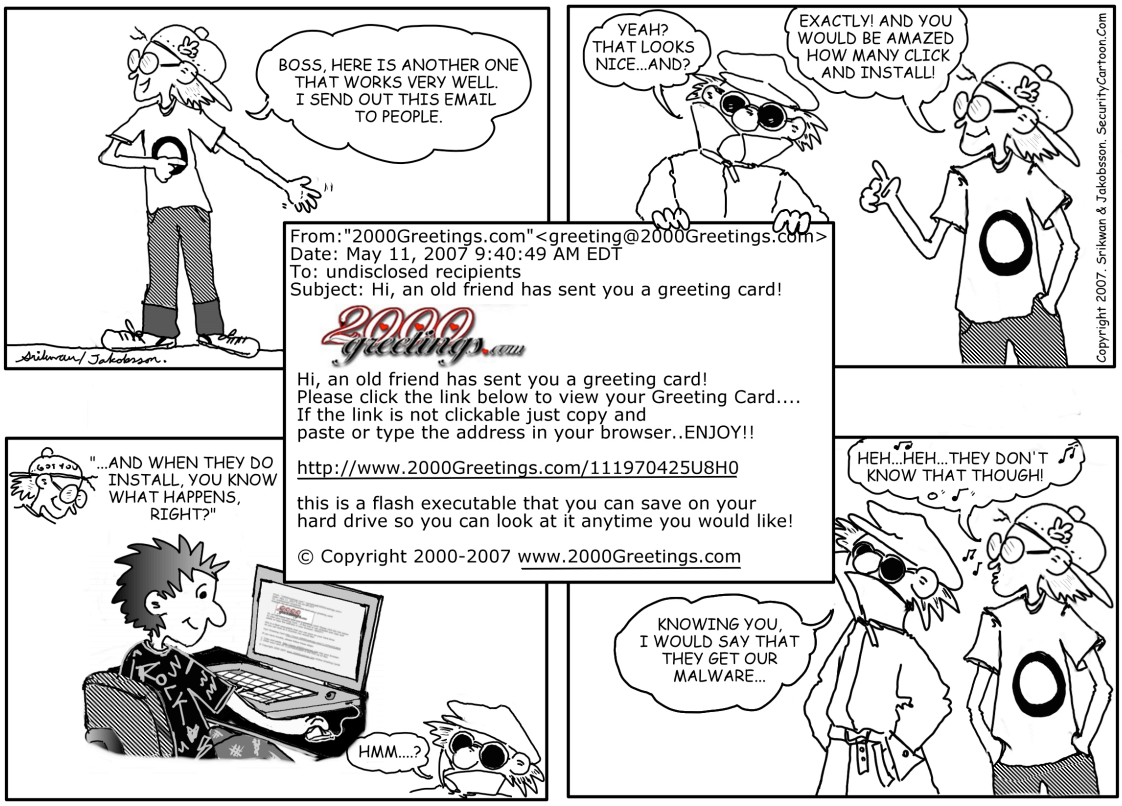

Figure 3: This cartoon strip shows how easily readers can be updated about recently occurring threats. When the message does not depend on the medium, one can achieve a high degree of adaptability of the educational message to a changing threat. This strip was developed in response to a wave of malware-attacks occurring in the late spring of 2007, where the recipient of the email was enticed to download content on the premises that it was a postcard from a friend. As described in [10], the attack vector involved both technical and human vulnerabilities in some instances, making this a rather sophisticated attack. It is worth noticing that many educational efforts do not warn users against this type of threat, as they often warn users about emails sent from strangers—an email like this is interpreted by many as being sent (or at least initiated) by a friend.

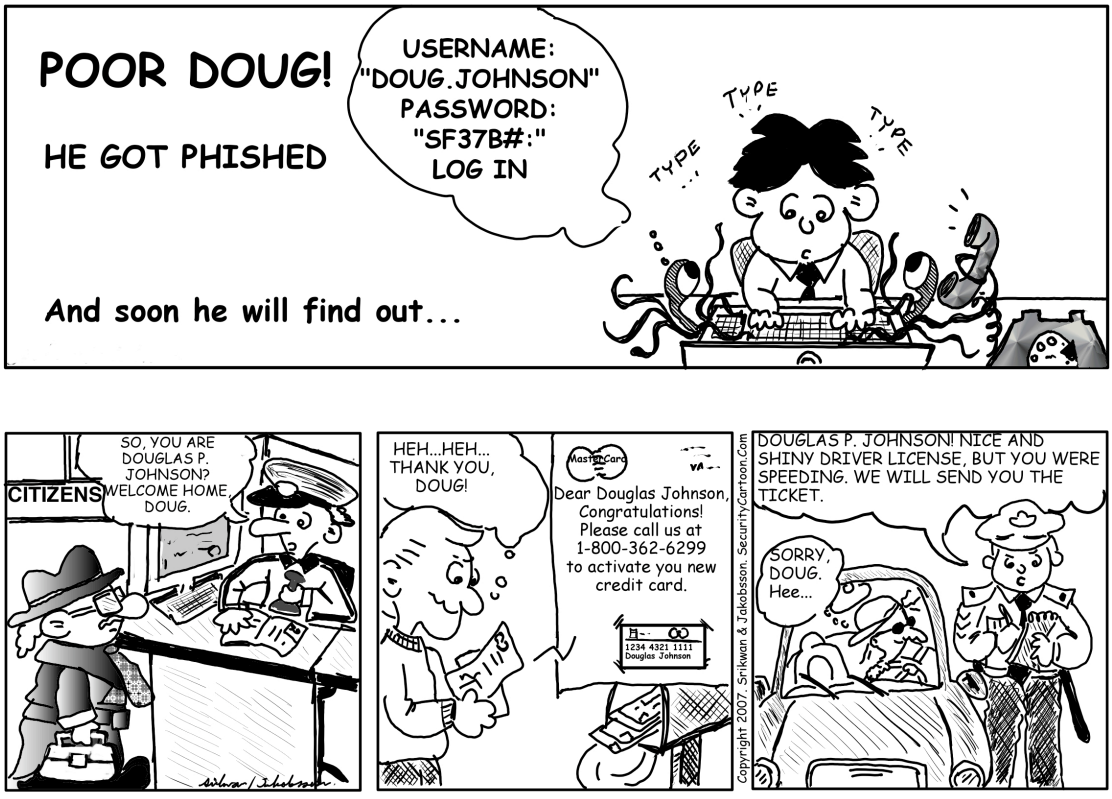

Figure 4: Most people are not willing to change their behavior unless they are given a good reason. Therefore, it is not enough to warn people about threats and say what they should and should not do. It is also important to explain what happens if they do not. This strip shows what happens to a victim of a keylogger (identified by the tentacled creature sticking out from his keyboard), and is a continuation on a sequence of strips in which keyloggers are introduced and the victim is infected by one.

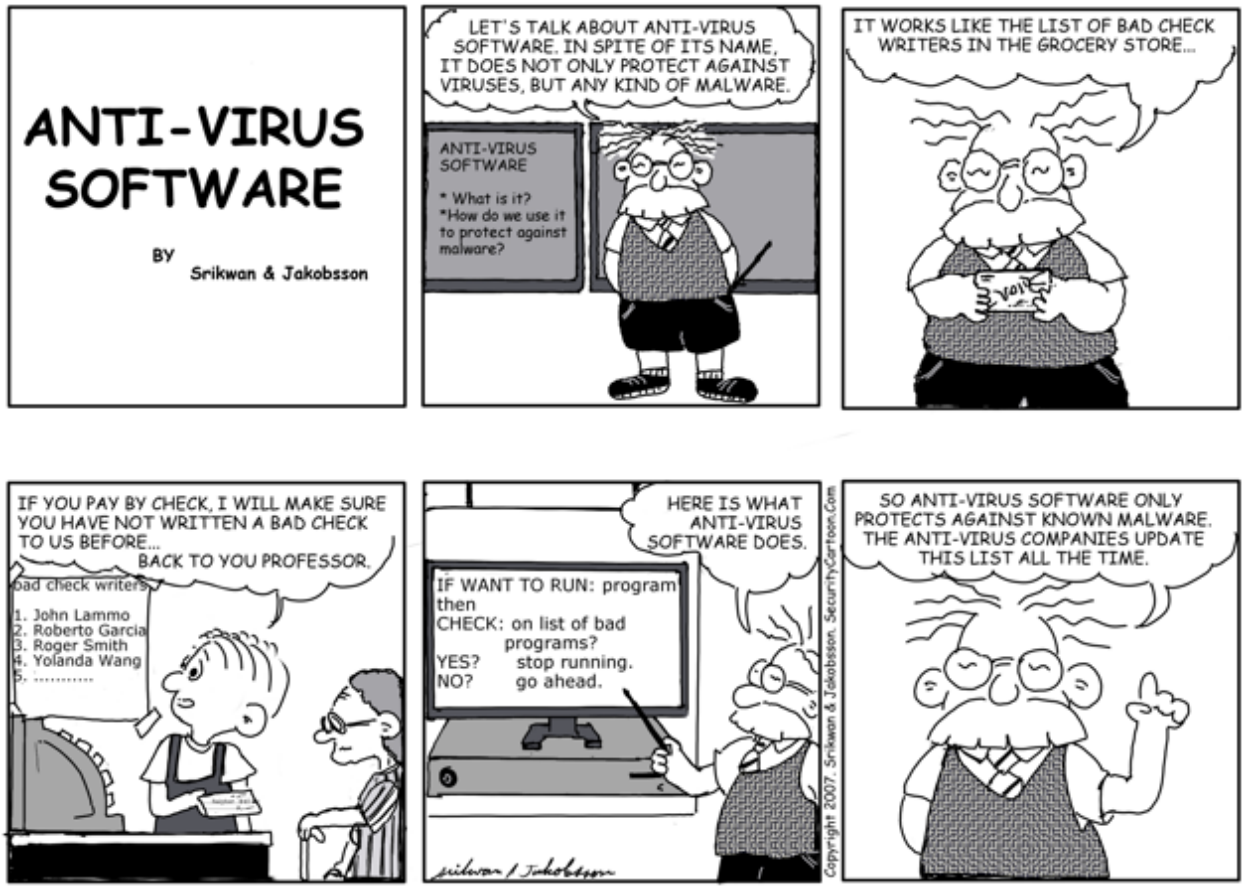

Figure 5: Many computer users put too much faith in anti-virus software, and many do not understand why it has to be constantly updated. This strip is designed to explain how AV software works, and identify when it may not detect an attack, both to attempt to set the reader's expectations at a reasonable level and to discourage readers from letting their AV subscriptions expire.

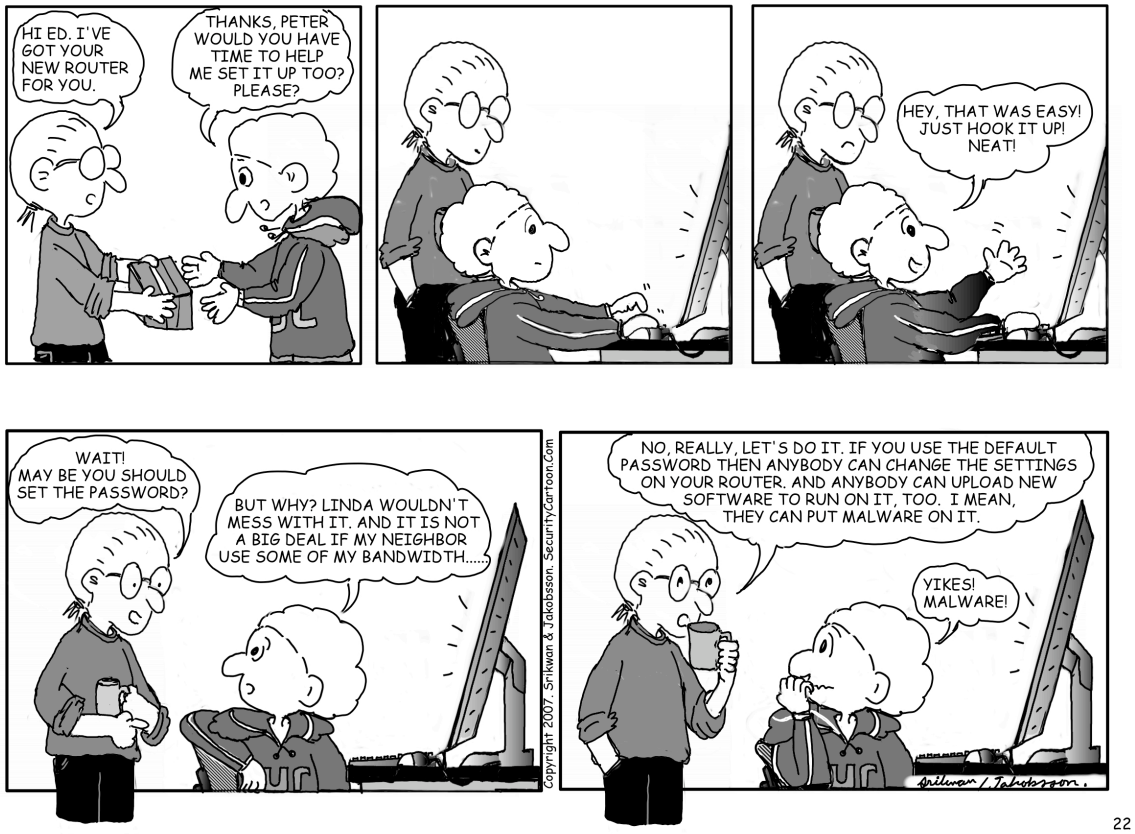

Figure 6: Many educational efforts focus on what to do, and what not to do, but neglect to describe why. As a result, people draw their own conclusions, which sometimes lead to a degraded security awareness. This strip attempts to highlight one of these issues in the context of the configuration of home routers. Research [55] shows that less than 50% of consumer routers have been correctly configured, corresponding to a notable practical security vulnerability in the context of potential attacks such as [53, 40].

Accessibility: It is important to design the educational message in a way that does not alienate or bore the intended readership. Combining warnings with explanations and a "role-playing" of consequences allows the reader to gain an intuitive understanding of why to follow given advice. We believe that this is important for a long-term adherence to the educational message. We let our educational message be framed as a cartoon, to avoid a compact "textbook" feeling, and to help the user identify with the characters in the story. We must remember that phishing education—and the exposure to it—is voluntary. Unless a given reader is highly motivated (which is not common) then the educational medium must make it easy to keep reading.

Immersion: To reach the desired results, we believe that the same message—framed in slightly different ways—needs to be repeatedly communicated to the reader. It is not enough to understand a concept once; we believe that the reader needs to practice acting in accordance with the message as well. This requires insight and understanding. This, of course, cannot be achieved simply by repeated use of one and the same message, but minor variations are necessary to retain the reader's interest. This principle is fully aligned with how other material is taught and learnt: Addition and subtraction, for example, are learnt by practice. Knowing the theory alone would not allow many to obtain the skills to do arithmetic.

Adaptability: To help defend against a changing threat, it is important to be able to reflect changes rapidly, and to be able to communicate recent trends in a manner that allows quick adoption of new practices. It is evident that cartoons permit this to be done, given the relatively low production costs. Furthermore, based on the accessibility and immersion principles, new advice will reach the target audience relatively quickly.

Acknowledgments

We want to thank Fred Cate for inspiring discussions leading to the creation of this material. We also wish to thank Lynn Goldstein and Sid Stamm for helpful feedback on the material underlying this paper, Liu Yang for latex support, and numerous colleagues and students of ours for their encouragement. Last but not least, we want to thank Microsoft. This work was partially funded by the Microsoft Trusted Computing grant (2005).

References

[1] M. Alexander, "Your Medical Records Stolen!", Reader's Digest, November 2006, pp. 86–93.

[2] Anti-Phishing Working Group, "Consumer Advice: How to Avoid Phishing Scams," http://www.antiphishing.org/consumer_recs.html, accessed July 2, 2007.

[3] A T & T, "Identifying and Protecting Against Phishing and Other Suspicious E-mails," http://att.centralcast.net/att_safety/Phishing/, accessed July 2, 2007.

[4] A T & T, "Internet Safety Game for Kids," http://www.att.com/gen/general?pid=1391, accessed July 2, 2007.

[5] A T & T, "Customer Education," http://www.att.com/gen/landing-pages?pid=6456, accessed July 2, 2007.

[6] F. W. Abagnale, "Stealing Your Life: The Ultimate Identity Theft Prevention Plan," Broadway Books, ISBN-13: 9780767925860

[7] K. Aoki, J. Boyle, J. Jenkins, "Bound By Law," www.law.duke.edu/cspd/comics/

[8] M. Arata, "Preventing Identity Theft for Dummies," John Wiley & Sons, ISBN-13: 9780764573361

[9] V. Anandpara, A. Dingman, M. Jakobsson, D. Liu, H. Roinestad. "Phishing IQ Tests Measure Fear, Not Ability," Extended abstract, USEC '07.

[10] Bryan Betts, "Unwanted e-card conceals a Storm," http://www.theregister.co.uk/2007/06/29/ecard_storm_trojan/, published June 29, 2007.

[11] Citibank, "E-mail Fraud & Security - Learn About Spoofs," http://www.citi.com/domain/spoof/learn.htm, accessed February 8, 2007.

[12] Chase, "Phishing," http://www.chase.com/ccp/index.jsp ?pg_name=ccpmapp/shared/assets/page/Phishing, Accessed February 8, 2007.

[13] J. M. Collins, "Investigating Identity Theft: A Guide for Businesses, Law Enforcement, and Victims," John Wiley & Sons, ISBN-13: 9780471757245

[14] M. Crawford, "Phishing education for banking customers useless," July 2, 2007. http://www.computerworld.com.au/index.php?id=1486962899&eid=-255

[15] R. Dhamija, J.D. Tygar, M. Hearst, "Why Phishing Works," Proceedings of the Conference on Human Factors in Computing Systems (CHI2006)(2006)

[16] eBay, "Reporting Spoof (Fake) Emails," http://pages.ebay.com/help/confidence/spoof-email.html, Accessed February 8, 2007.

[17] eBay, "Spoof Email Tutorial," http://pages.ebay.com/education/spooftutorial/, Accessed February 8, 2007.

[18] F-Secure Phishing Demo, http://www.youtube.com/watch?v=D54nTfLhRr4, Accessed July 3, 2007.

[19] Federal Financial Institutions Examination Council, "Authentication in an Internet Banking Environment," www.ffiec.gov/pdf/authentication_guidance.pdf, October 12, 2005.

[20] Federal Trade Commission's advice on how to report spam and scams, http://www.ftc.gov/bcp/conline/edcams/spam/report.html

[21] P. Finn and M. Jakobsson, "Designing and Conducting Phishing Experiments," In IEEE Technology and Society Magazine, Special Issue on Usability and Security, 2007.

[22] M. J. Frank, "From Victim to Victor: A Step-by-Step Guide for Ending the Nightmare of Identity Theft," Porpoise Press, ISBN-13: 9781892126047

[23] B.J. Fogg, C. Soohoo, D.R. Danielson, L. Marable, J. Stanford, E.R. Tauber, "How Do Users Evaluate the Credibility of Web Sites?: A Study with Over 2,500 Participants," Proc. DUX(2003)

[24] B.J. Fogg, J. Marshall, O. Laraki, A. Osipovich, C. Varma, N. Fang, J. Paul, A. Rangnekar, J. Shon, P. Swani, M. Treinen, "What Makes Web Sites Credible?: A Report on a Large Quantitative Study," Proc. CHI (2001) 61-68

[25] S. L. Garfinkel and R. C. Miller, "Johnny 2: a user test of key continuity management with S/MIME and Outlook Express," Proceedings of the 2005 Symposium on Usable Privacy and Security, 2005, pp. 13 – 24

[26] Hector's World, http://www.hectorsworld.com/, Accessed July 3, 3007.

[27] M. Jakobsson, "The Human Factor in Phishing," In Privacy & Security of Consumer Information '07, http://www.informatics.indiana.edu/markus/papers/aci.pdf

[28] M. Jakobsson and J. Ratkiewicz, "Designing Ethical Phishing Experiments: A study of (ROT13) rOnl auction query features," In Proceedings of the 15th annual World Wide Web Conference, 2006.

[29] M. Jakobsson, A. Tsow, A. Shah, E. Blevis, Y.-K. Lim, "What Instills Trust? A Qualitative Study of Phishing," Extended abstract, USEC '07.

[30] M. Jakobsson and S. A. Myers (Eds.), *Phishing and Countermeasures: Understanding the Increasing Problem of Electronic Identity Theft.* ISBN 0-471-78245-9, Hardcover, 739 pages, December 2006.

[31] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. Cranor, J. Hong, and E. Nunge, "Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish," In Proceedings of the 2007 Symposium On Usable Privacy and Security, Pittsburgh, PA, July 18-20, 2007.

[32] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, and J. Hong, "Teaching Johnny Not to Fall for Phish," CMU Technical Report, February 8, 2007

[33] Livesecurity, "US Bank phishing attack exposed," `http://www.youtube.com/watch?v=n2QKQkuSB4Q`, posted March 27, 2007.

[34] P. Kumaraguru, Y. W. Rhee, A. Acquisti, L. Cranor, J. Hong, E. Nunge, "Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System," Technical Report CMU-CyLab-06-017, November 9, 2006

[35] MailFrontier Phishing IQ Test II, `http://survey.mailfrontier.com/forms/msft_iq_test.html`

[36] S. McCloud, "Understanding Comics: The Invisible Art," HarperCollins Publishers, Inc. New York, 1993.

[37] Microsoft.com, "Recognize phishing scams and fraudulent e-mails," `http://www.microsoft.com/athome/security/email/phishing.mspx`, Accessed Feb 8, 2007.

[38] R. Miller, S. Garfinkel, F. Menczer, R. Kraut, "When User Studies Attack: Evaluating Security By Intentionally Attacking Users," Panel at SOUPS 2005, Slides available at `http://cups.cs.cmu.edu/soups/2005/program.html`

[39] K. D. Mitnick and W. L. Simon, "The Art of Deception: Controlling the Human Element of Security," John Wiley & Sons, ISBN-13: 9780764542800.

[40] S. Myers and S. Stamm, "Trawler Phishing," In "Crimeware," M. Jakobsson and Z. Ramzan, Eds., Symantec Press, 2007, ISBN 0321501950.

[41] Netcraft News, "More than 450 Phishing Attacks Used SSL in 2005," `http://news.netcraft.com/archives/2005/12/28/more_than_450_phishing_attacks_used_ssl_in_2005.html`, Accessed July 3, 3007.

[42] NetSmartzKids, "Teaching Kids What to Watch Out for Online," `http://www.netsmartzkids.org/indexFL.htm`, accessed July 2, 2007.

[43] NGSSoftware, "The Phishing Guide - Understanding & Preventing Phishing Attacks," http://www.ngssoftware.com/papers/NISR-WP-Phishing.pdf, Accessed February 8, 2007.

[44] The Ohio State University Libraries, "Thomas Nast," 2002. http://cartoons.osu.edu/nast/

[45] PayPal, "Protect Yourself from Fraudulent Emails," https://www.paypal.com/cgi-bin/webscr?cmd=_vdc-security-spoof-outside, Accessed July 14, 2007.

[46] Paypal, "Can you spot phishing?", https://www.paypal.com/fightphishing, Accessed July 14, 2007.

[47] Secunia Research, "Multiple Browsers Tabbed Browsing Vulnerabilities," http://secunia.com/secunia_research/2004-10/advisory/, accessed July 3, 2007.

[48] S. Schechter, R. Dhamija, A. Ozment and I. Fischer, "The Emperor's New Security Indicators," Proceedings of the IEEE Symposium on Security and Privacy, May 2007.

[49] B. Schneier, "Secrets and Lies: Digital Security in a Networked World," John Wiley & Sons, ISBN-13: 9780471453802

[50] E. Shanahan, "ID Thieves' New Tricks," Reader's Digest, June 2006, pp. 82-87

[51] S. Srikwan, M. Jakobsson, www.SecurityCartoon.com, Accessed May 16, 2007.

[52] S. Stamm, M. Jakobsson, M. Gandhi, "verybigad.com: A study in socially transmitted malware," http://www.indiana.edu/~phishing/verybigad/

[53] S. Stamm, Z. Ramzan, M. Jakobsson, "Drive-By Pharming," Technical Report TR641, Indiana University, Dec 2006

[54] A. Tsow, "Phishing With Consumer Electronics – Malicious Home Routers," In Models of Trust for the Web, a workshop at the 15th International World Wide Web Conference (WWW2006), May 22-26, Edinburgh, Scotland.

[55] A. Tsow, M. Jakobsson, L. Yang, S. Wetzel, "Warkitting: the Drive-by Subversion of Wireless Home Routers," Anti-Phishing and Online Fraud, Part II Journal of Digital Forensic Practice, Volume 1, Special Issue 3, November 2006

[56] vladnik, "Fly-by malware installation demo," http://www.youtube.com/watch?v=oU1gcprFEPU, posted July 20, 2006

[57] Webwhacker 5.0, http://www.bluesquirrel.com/products/webwhacker/

[58] T. Whalen and K.M. Inkpen, "Gathering evidence: use of visual security cues in web browsers," In Proceedings of the 2005 Conference on Graphics interface (Victoria, British Columbia, May 09 - 11, 2005). ACM International Conference Proceeding Series, vol. 112. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, pp. 137–144.

[59] M. Wu, R. Miller, S. Garfinkel, "Do Security Toolbars Actually Prevent Phishing Attacks?" Proc. CHI (2006)