# Security Education against Phishing:

## A Modest Proposal for a Major Rethink

**Iacovos Kirlappos and M. Angela Sasse** | University College London

**When tempted by a good deal online, users don't focus on security warnings; rather, they look for signs to confirm a site's trustworthiness. User education needs to focus on challenging and correcting the misconceptions that guide current behavior.**

Phishing—tricking computer users to disclose personal information, credit card details, user names, and passwords—has been a major problem for the past 15 years. The probability of an online shopper coming across a phishing website is alarmingly high, because many appear in popular Web search engines results. In a recent UK police operation, seven of the top 10 Google results for a popular brand of boots were found to be fraudulent websites. In addition, one in 12 online buyers of event tickets reported having been caught by a scam ticket website, with the average loss for each victim being £80. Most sites are taken down quickly once identified, but new ones spring up daily, making the process of identifying and closing all of them impossible for crime-prevention authorities.

Disclosing financial details to scam sites can lead not only to immediate monetary losses but also to identity theft and its consequences (damage to a person's credit rating or a person being linked to illegal activities). Even though some banks cover customers who have had their credit card details stolen, this solution isn't sustainable as a long-term solution. These problems can lead to an overall loss of trust in online shopping and deter consumers from engaging in online financial transactions.

Two main approaches have been used to protect users against phishing: antiphishing indicators and user education. Rachna Dhamija and her colleagues explained that the first approach is ineffective because a significant percentage of users ignore passive indicators.[1] Even when users notice the indicators, they often don't understand what they signify. In addition, inconsistent positioning on different Web browsers makes the task of identifying a phishing site difficult. Stuart E. Schechter and his colleagues reported that 53 percent of their study participants still attempted to log in to a site after their task was interrupted by a strong security warning.[2] In the same study, removing the HTTPS indicator had no effect on participants' willingness to enter their personal details in a site, and 97 percent of participants entered personal details even after site authentication images were removed. Both papers' findings lead us to conclude that effective education must complement any technical antiphishing measures to improve users' ability to detect phishing sites.

## User Education

Both government organizations and academic institutions have put significant effort into user education. To improve public understanding of security, the US

 Copublished by the IEEE Computer and Reliability Societies 

Computer Emergency Readiness Team offers "advice about common security issues for nontechnical computer users" on its site (www.us-cert.gov/cas/tips). Ponnurangam Kumaraguru and colleagues developed the PhishGuru training system to teach users how to identify phishing attacks.[3] The system sends out simulated phishing emails and delivers training messages when users click the included URLs. Its effectiveness was tested with 515 participants; 28 days after the first email, despite being given training more than once, 17.5 percent of participants still entered personal details into simulated phishing websites. This was a significant improvement from the 40.1 percent a control condition revealed before the study, but still leaves one in five users vulnerable. The same research group developed Anti-phishing Phil, an online game to teach users not to fall for phishing by explaining how to identify phishing URLs, where to look for cues in Web browsers, and how to use search engines to find legitimate sites.[4] They reported improved user ability to detect phishing websites after receiving training. The false-positive rate (phishing site identified as real) was reduced from 30 percent to 14 percent, and the false-negative rate (nonphishing site identified as spoof) was reduced from 34 percent to 17 percent. Despite those reductions, 31 percent of users were still unable to differentiate between good and bad sites.

Cormac Herley argued that teaching users to check URLs is the wrong strategy because even diligent application of what is being taught offers users only limited protection against phishing.[5] In his view, the effort/benefit ratio means users should ignore this advice, especially if the actual risk of financial loss is low.

Another reason current education and training efforts aren't effective is because they assume that users are keen to avoid risks, and thus likely to adopt behaviors that might protect them. But in reality, most online shoppers are looking for good deals. They start from a search engine and are presented with links to various websites that present (often very tempting) offers. The opportunity to save a significant amount of money on something they need, or acquire something they might normally not be able to afford, makes users vulnerable. Frank Stajano and Paul Wilson identified this as the *need and greed* principle, which scammers exploit successfully: once scammers know what users want, they can easily manipulate them.[6] To address this problem, the UK Office of Fair Trading (www.oft.gov.uk) launched campaigns aiming to increase consumer awareness of fake shopping websites. The slogan "if it sounds too good to be true, then it probably is" appears in the campaign materials as well as in communications by law enforcement officers—so far with little success.

In line with Herley, we argue that current security education on phishing offers little protection to users who assess a potentially malicious site in this frame of mind.[3,4] Security education needs to consider the drivers of user behavior in this situation—the cues users look for and how they interpret them. Successful security awareness, education, and training must do more than warn users of dangers—they must target the misconceptions that underlie user actions. Although we focus on phishing, a shift in perspective could help security researchers and practitioners develop more effective security awareness, education, and training in other areas of computer security.

## Trust Cues in Online Transactions

Online shoppers face a situation of risk and uncertainty: they must provide payment details and personal information to websites and can't be certain they'll receive the goods they expect in return. Many online shoppers will take risks to gain benefits, and they look for trust cues to reduce the degree of uncertainty about the outcome—a trustworthy transaction partner is more likely to deliver. Jens Riegelsberger and his colleagues developed a framework of trust signals that both transaction partners can emit, focusing on ways to incentivize trustworthy behavior and incorporating signals to assess trustworthiness, such as a site's "professionalism" (for example, the absence of technical failures, breadth of product palette, and usability) and social embeddedness (for example, a retailer's reputation among consumers' friends and relatives).[7] Combined with Dan Kim and his colleagues' findings that consumer trust directly and indirectly affects purchasing intentions,[8] we can assume that users' willingness to engage in a transaction increases when the perceived risk is low.

Marios Koufaris and William Hampton-Sosa conducted a study on the development of trust in online companies by first-time customers, identifying four factors affecting users' purchasing decisions:[9]

- perceived reputation of the company,
- perceived usefulness and ease of use of the website,
- perceived security control, and
- the selection of products available (if the company has a wide range of products, it's more trustworthy).

However, the use of closed questions in their surveys, specifically aiming to confirm those four factors, didn't allow the surveys to reveal any other website properties that affect user decisions. Kim and his colleagues also discussed the effectiveness of third-party seals as an assurance of trust, concluding that they decrease consumers' perceived risk, but that consumers

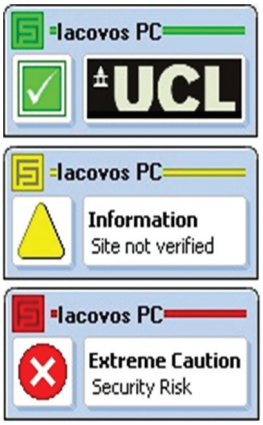


**Figure 1.** Solid displays the traffic-light convention color code. It appears as a small box outside the user's browser.

know very little about their purpose and what protection they offer.[8]

Scammers exploit trust development principles, both in the real world and online,[6] but the implications of users' trusting behavior haven't been considered in the phishing context.

## Study Description

We originally designed our study to evaluate the effectiveness of Solid (www.solidauthentication.com), a new active antiphishing tool that uses traffic-light security indicators to signal whether a website is genuine or fake (see Figure 1). A green indicator accompanied by the company's logo appears when a website's details match those expected. A yellow indicator appears when the webpage fails some part of the authentication test. When Solid identifies a webpage as malicious, a popup window appears before the webpage loads and the tool window turns red, indicating that a security risk exists and suggesting a redirection to the registered retailer's real website. Users have the option to close the current window and proceed to the risky website. If they choose to do that, the indicator remains red. If the website isn't registered with the tool, the indicator turns gray.

We recruited participants using the University College London (UCL) psychology subject pool, which is open to the public. Participants had to be over 18, regularly use online shopping, and be able to visit the lab for one hour of testing. The standard reward for their participation was a £15 Amazon voucher; participants received an additional reward if they chose safe websites in the experiment. We tested 36 participants in total:

- 17 (47 percent) were male and 19 (53 percent) female,
- their average age was 24 years,
- their average computer experience was 12 years,
- their average daily Internet browsing was 4.5 hours,
- they received an average of 14 emails per day,
- 35 (97 percent) of them had checked their account balance online at least once,
- 34 (94 percent) had transferred money to other people's accounts using online banking services at least once,
- all had bought goods online in the past,
- 19 (53 percent) had configured a firewall,
- 18 (50 percent) had designed a website,
- 8 (22 percent) had registered a domain name,
- 7 (19 percent) recalled using Secure Shell (SSH) connections in the past,
- 9 (25 percent) had been victims of phishing or knew someone who had been, and
- 12 (33 percent) had been victims of Internet scams or knew someone who had been.

We divided participants equally between two conditions: 18 used the active antiphishing tool, and 18 did not. They were asked to buy tickets for a music festival, presented with six websites, and asked to decide within five minutes which one to buy from. We used this time frame to replicate the *time principle* Stajano and Wilson identified as an attack tactic[6]—in this case, very plausible, because tickets for popular events tend to sell out quickly. To replicate the risk that ticket buyers face when buying from unknown retailers online, the reward given to participants varied, depending on which website they chose, based on the following scenario:

> You want to buy tickets for Friday, 27 August, for the LED electronic music festival at the moment they go on sale. You have £60 available. You know that festivals sell out very quickly, so you only have five minutes to buy them. You searched in Google for "LED festival tickets" and came upon six websites that claim to sell tickets. You now need to choose from which site to buy. Your additional reward from the experiment is the amount of money you initially have available (£60) minus the price of the tickets on the website you chose to buy from. If you buy from a fraudulent website, then you get no extra reward (only the 15 pounds that are paid for your participation in the experiment). You can browse in the websites with no limitations. Warnings will be given to you when two minutes and one minute remain.

All the websites in the experiment were local copies of legitimate retailers downloaded from the Internet.

**Table 1. The websites used in the experiment with the corresponding prices and colors.**

| Website | Ticket price | Tool color |
|---|---|---|
| Gigantic (www.gigantic.com) | £50 | Green |
| HMV Tickets (www.hmvtickets.com) | £50 | Green |
| See (www.seetickets.com) | £25 | Red |
| Skiddle (www.skiddle.com) | £20 | Gray |
| Sold-out Ticket Market (www.soldoutticketmarket.com) | £40 | Gray |
| View London (www.viewlondon.co.uk) | £20 | Yellow |

**Table 2. Distribution of participants' potential rewards on the basis of the color of the website they chose.**

| Potential payoff | Solid indicator | Number of participants | |
|---|---|---|---|
| | | Solid group | Control group |
| £10 | Green | 10 | 5 |
| £20–40 | Gray or yellow | 8 | 12 |
| £35 | Red | 0 | 1 |

We modified our DNS server so that the URL structure and website appeared to the participants in the same way as if they were browsing. We modified Solid to display the colors shown in Table 1.

Table 2 shows that most participants who used Solid chose the safe (green) options. In addition, none chose the website marked as red. A chi-squared test revealed a significant shift in participant decisions when Solid was used ($\mathcal{X}^2(1)$ p = 0.03324). Although this could be argued to be a success, a significant number still chose sites labeled as potentially risky (gray or yellow) over the ones clearly labeled as safe. Why did so many participants ignore the potential risks when a safe alternative existed?

## Identified Trust Factors and User Misconceptions

In the debrief interviews following the experiment, we asked each participant to explain what affected his or her website choice. No guiding questions were used—participants were free to report any factors that affected their final choice. During this discussion, the websites were left open so that participants could refer back to them. The interviews were audio-recorded, transcribed, and analyzed using grounded-theory coding techniques.[10] The results show that security indicators were only one of several different signals that our participants used to assess a website's legitimacy. We identified eight factors that affected the participants' choice of websites: previous experience, logos and certifications, advertisements, social networking references, inclusion of charity names, amount of information provided, website layout, and company information.

All eight Solid users who chose potentially unsafe yellow or gray sites said the potentially higher reward was an incentive to ignore the green site—confirming the need and greed principle.[6] Participants mentioned on average three additional factors that affected their decisions.

### Previous Experience

Previous experience with a website and familiarity with a brand induce users' willingness to trust it. Only one participant had shopped from any of the six websites, but 18 (50 percent) said they'd heard of the brand names, and this played a key role in their choices—suggesting a *trust halo effect*.[11] An example of this is the View London website, which five participants (14 percent) had used to read venue reviews, but never to buy event tickets. Brands such as View London and HMV are popular in the UK—the first because of its review pages and the second because of its many street retail outlets that sell music and gaming products. But none of our participants were familiar with their ticket-selling operations. Scammers can exploit this very broad concept of "being familiar" with a brand by creating fake websites, claiming to be online outlets of familiar brands.

### Logos and Certifications

Five websites displayed some form of trust logo, and 10 participants said those played a major role in their decisions. The VeriSign Secured logo turned out to be


**Figure 2.** The VeriSign Secured logo. Six study participants said they trusted this sign, but none could explain what it stands for.


**Figure 3.** The Internet Shopping Is Safe logo. This image affected four participants' decision to trust a website.


**Figure 4.** The Hitwise No. 1 Award Winners logo. Only one participant said this logo affected her decision to trust a website.

the most popular one (see Figure 2). Six participants (17 percent) said they trusted this sign because they'd seen it on other trusted websites, but none of them could explain what the logo stands for and why a website displaying it should be secure. Only two participants checked whether the logo was a clickable link for information about the merchant.

The Internet Shopping Is Safe logo was displayed on one website (see Figure 3), and four participants (11 percent) reported it affected their choice. Another logo, the Hitwise No. 1 Award Winners logo, was displayed on one website (see Figure 4), but only one participant mentioned it affected her choice. Three participants (8 percent) mentioned financial organizations' logos, such as the credit/debit cards accepted: "They accept Visa, MasterCard, and American Express, so they must be real." Only two participants (6 percent) checked whether the logos were clickable links or displayed a valid certificate or registration number with the relevant authority. In total, 13 participants claimed that logos affected their choices, but none could explain why this signaled trustworthiness.

### Advertisements

A website's affiliation with known entities is often interpreted as a trust signal. Five participants (14 percent) mentioned advertisements by well-known companies; they argued that a reputable company wouldn't pay scammers to advertise on their website. It hadn't occurred to them that scammers might include ads to make their site look legitimate and that the companies advertised might not be aware of this use of their material.

### Social Networking References

The growing popularity of social networking websites is starting to affect online commerce; scammers exploit this by suggesting their site is associated with these websites. Inclusion of links to Facebook and Twitter pages boosted seven participants' (19 percent) confidence in a site; they believed that links to those sites couldn't be fraudulent because any scam victims could post negative feedback to warn others. Social networking sites' iconic status is a key weakness if users don't understand how easily scammers can fake an association with the sites.

The presence of user feedback can also contribute to trust development. Four participants (11 percent) gave positive comments about their perceived correlation between user feedback and trust. This was particularly clear in the cases of a website that included pictures of users who left feedback and a website whose members "are planning to attend an event," confirming past findings that richer media representations can induce a positive trust bias.[12]

### Inclusion of Charity Names

The inclusion of the name of a charity (Oxfam) on one website (www.gigantic.com), accompanied by a claim that it donates 10 percent of its profits to the charity, led two participants (6 percent) to believe the site was genuine. Benevolence is an intrinsic trust property,[7] and real-world scams exploit this, using charities as a pretext—for example, collecting donations of money or clothes that the scammers actually keep. Online scammers can also exploit this, because users aren't aware of the potential misuse of charity names and don't attempt to verify the claims they see on websites.

### Amount of Information Provided

The amount of information the website included on the event of interest was reported as an important

factor by six participants (17 percent). All websites included information on the event (gate opening times, facilities, instructions how to get to the venue, and so forth), but those that displayed the information on the main event page attracted participants more. Again, addition of rich media, such as maps, made websites appear "more real" and trustworthy. In general, participants seemed to follow the maxim that the more effort is put into a website's development, the less likely it is to be the site of scammers, who want to make money fast.

### Website Layout

Seven participants (19 percent) mentioned that the website's design structure appeared familiar because it was similar to other legitimate websites. This similarity led them to assume the site ought to be genuine. Interacting with particular websites leads to *mental anchoring* of trustworthy sites' design and appearance, against which they assess trustworthiness of a new site on a first-time interaction. Participants were also reassured by indicators of routine business—in this case, availability of tickets for a variety of events. They simply assumed a scam site would try to target a particular event.

### Company Information

The level of detail the website provided on the company behind it also affected participant decisions. Five participants (14 percent) mentioned the presence of the company's registration number, tax reference numbers, direct telephone numbers, ticket delivery information, or claims that they are official ticket outlets increased their confidence in the website. But as with logos and privacy policies, none of the participants knew how to verify this information and didn't attempt to do so.

## Effective Antiphishing Education

Our analysis reveals a significant gap between the signals security experts would like users to consider when assessing a website's legitimacy and those they actually use when faced with a tempting offer. Our findings—which unite and confirm a set of observations from previous studies—suggest that advice given in current user education is largely ignored because it focuses on indicators that users don't understand or trust. We need to deploy new approaches targeting user misconceptions that lead to insecure behavior, attracting user attention on how they might be targeted and correcting those misconceptions.

### What Should We Teach Users?

To help users, we need to explain how and why trustworthiness indicators they use successfully in the real world fail them online. As Rick Wash put it, users form their own *folk models*—their own understanding of how things work in computer security—which aren't necessarily accurate and can potentially lead to erroneous decisions.[13] In our experiment, participants used those to justify their decisions to ignore expert advice: they ignored Secure Sockets Layer locks and URLs, and used their own heuristics to assess a site's legitimacy.

Reliance on indicators and models from the physical world leave users vulnerable in many ways:

- Participants were surprised when told after the experiment that fake versions of real websites can be uploaded by anyone online, or that someone can create a website claiming to be someone else.
- The fact that 13 participants used trust logos to guide their choices might seem encouraging, but only two checked whether those logos were clickable links, seeking more information on the certification and the merchant. None of our participants could explain what protection those logos might offer; they reacted to their mere presence as safety indicators.
- The blind trust users place in sites that suggest a link with social networking sites demonstrates their popularity, and a worrying potential for exploitation by scammers. Our participants didn't consider that anyone can create a page or profile on those sites or that spammers can add social networking logos to their fraudulent sites.
- The other design elements participants reported (amount of information provided, website layout, and company information) can also be easily mimicked. Our participants seemed unaware that although signals of high levels of investment are reliable indicators of real-world retailers, design elements can be copied in a matter of seconds.

Current security education approaches don't target the misconceptions we identified. Users don't understand how scammers operate, and they make assumptions about how the online environment operates on the basis of their real-world experiences. Effective security education needs to challenge users' assumptions about trust signals and their decision processes and replace them with trust signals and strategies for assessing risks in an online environment. Just as in the physical world, some users will willingly take risks online. So, security education should equip users to assess the potential risks and benefits correctly, rather than tell them to avoid going to any potentially risky site.

### How Should We Teach Users?

The first step toward effective user education is to

recognize that awareness, education, and training are three distinct steps of a process to improve user competence.[14] The role of security awareness is to attract users' attention and help them realize that there is a problem that might affect them. This is a necessary first step to render them receptive to education and training measures. Security awareness measures must capture users' attention using strong visual elements, surprise, or humor. In the case of phishing, existing perceptions need to be challenged, including users' perceptions of their ability to assess the risks involved in online transactions and what reliable trustworthiness indicators are. An example would be an advertisement, online or in print, that shows two very similar websites with the caption that reads, "One of these websites belongs to [a famous bank]; the other is run by a criminal gang in Elbonia waiting to steal your user name and password and empty your account at [famous bank]. Can you tell which is which?" Once users realize they can't tell the difference, or choose the wrong site, they're more likely to pay attention to a subsequent pointer to a site that offers education (to improve their knowledge) or training (to improve their skills).

An example of security education delivery in this context would be an online game in which users can collect or lose points by answering questions about the trust and assurance indicators on a professional-looking website. For instance, if they point to an ad on the site, they would be presented with the statement "The presence of an ad by [famous brand] indicates this is a legitimate site, because [famous brand] would not pay to advertise on a phishing website" and asked to rate it as true or false. Explanations of why an answer is true or false can help to correct misconceptions and reinforce correct statements. High scores or badges can promote secure behavior among individual users or groups in an organizational setting.

## How Can We Reach Users?

Another fundamental aspect of effective security education delivery is the choice of communication channels to disseminate awareness, education, and training information to users. To date, two different approaches have been used:

- general public awareness and education campaigns (both online and offline), and
- context-specific warnings and indicators (online).

In public-awareness campaigns, users are informed about the risk of scams and sometimes told about possible ways to protect themselves, but no training is delivered. The effectiveness of those campaigns is questionable. Approaches like the UK police campaigns don't provide any useful information or skills to consumers. Many legitimate online retailers sell goods at prices significantly lower than street prices, which is a major draw for online shopping. So, how can consumers tell when a good deal becomes too good to be true? Generic warnings like this might deter many who would benefit most from lower online prices—people with lower incomes—from shopping online altogether, because they can least afford to take the risk.[7]

A more promising approach is to provide awareness, education, and training in the context of the services users aim to access. Consumers are more motivated if warnings are specific to risks they know and care about, and these warnings are more likely to be accessible when explained by peers who have a similar perception of risks and pitfalls. An example worth following is eBay, which has created an online community in which users can post tutorials (often featured from eBay's homepage) on how to identify counterfeit goods or how to avoid scams. A UK bank uses another context-specific approach, asking its customers for partial PINs and passwords to access their online banking accounts (for example, digits 2, 1, and 4 of the PIN and digits 2, 6, and 9 of the password). Its login page explicitly mentions that users should never disclose their full PIN and password to a website, aiming to teach their customers to protect themselves from word capture attempts using phishing attacks.

Both of the above measures increase user awareness of how scammers might target them when using those specific websites as well as aim to educate them by explaining how to avoid falling for these types of attacks. But this isn't training, which needs to not only present correct behaviors to users but also test their understanding of the communicated information and correct any identified misconceptions.[14]

A potential user-training approach is to feature short tutorials on retailer and bank websites. For instance, after informing users about the potential of a criminal gang in Elbonia, we need to draw their attention to the differences between a legitimate and a scam website (for example, your bank would never ask you to disclose your full PIN and password). To ensure correct skill acquisition, after the tutorial, users should be asked to distinguish between a few examples of legitimate and scam sites on the basis of the principles they learned. To encourage participation, retailers could launch competitions with prize drawings.

## Lessons Learned from Misconceptions

Trust symbols such as logos and certifications are currently either misinterpreted or go unnoticed. Trust seals are effective only if users can recognize them, know what protection they offer, and check their

legitimacy.[7] Because this isn't the case, broader awareness campaigns using a range of information channels are necessary. First, attract people's attention to the presence of those seals, then explain what the problem is and what measures are in place to protect them (in this case, a browser add-on) and provide them with information on what needs to be done on their side. Active antiphishing tools, which interrupt the user's primary task only when a threat is identified, seem to be an effective measure against phishing attacks,[15] and Solid had a significant effect in deterring participants from known bad sites. But improving user defenses against future scams requires an additional step: whenever a tool detects unauthorized use of trust symbols, it should present users with information on what went wrong, increasing their awareness of the problem and the potential risks they face while shopping online. This needs to be done in the browser when users visit sites that carry those seals, so that users don't need to download and install additional software to be protected. In addition, whenever a tool identifies a risk, it should provide short tutorials with strong eye-catching visual artifacts, ensuring users understand the nature of the problem and what the messages delivered to them mean as well as correcting any potential user misconceptions. The information should be short and descriptive so that it doesn't appear as too much effort to users, as they might then ignore it.

> **Rather than flooding users with information, we need to consider how users make decisions, both in business and personal settings, and tailor new security solutions based on this.**

Users trust sites that appear familiar. This can be used to retailers' advantage—established brands can provide easy recognition and reassurance to customers. But customers will expect trusted institutions to guarantee a transaction and help them if something goes wrong. This can enable consumers to engage in transactions in which the perceived risk level is higher than what they would otherwise accept. An example of a well-trusted organization is PayPal—a payment method that provides users with the advantage of hiding their card details from the seller and guarantees to refund its customers if transactions go wrong. Including support for payment methods like this on a site could increase consumers' willingness to buy from it. However, the presence of those mechanisms alone isn't enough—again, users need to be made aware of the potential problems they might encounter when shopping online (for example, receiving counterfeit products, receiving nothing, or having credit card details compromised) and the extent to which they are protected, provided they comply with a manageable set of rules. We can achieve this by getting big retailers to use those mechanisms and provide visual elements to explain to users how they are protected. Statements such as "paying by EasyPay ensures your card details are not shared with anyone when buying online" can increase customer confidence in e-commerce. This could be accompanied by short tutorials labeled with phrases such as "how am I protected?" that explain to the users in more detail what can go wrong in an online transaction and how they are protected. Any approach attempting to do this should be consistent across online retailers/service providers to avoid flooding users with information, causing confusion instead of aiding their education and skill acquisition.

Again, users need to be made aware of how easy it is for attackers to mimic visual trust cues.[7] An engaging, though perhaps controversial, approach would be to create a YouTube video demonstrating "how to create your own phishing website in 10 minutes and five easy steps" and spreading the word through social networking sites.

Our findings suggest the need for a change of direction in security awareness, education, and training. Instead of flooding users with warnings and repeatedly telling them to behave as security experts would like them to, effective security awareness starts with the users' perspectives and decision-making processes, imperfect though they might be. Users form their own risk models and use a set of heuristics to assess the trustworthiness of the websites with which they interact. Having identified users' misconceptions, we need to connect with them through specific awareness, education, and training campaigns.

Campaigns should address retailers as well. We identified some examples of bad practice among legitimate retailers that don't provide reliable trust signals or allow scammers to exploit potential vulnerabilities in their website design. They need to be made aware of how their websites—and their customers—are attacked and how they can help customers distinguish between their legitimate website and counterfeits. This could help them protect their customer base and their reputation.

Our proposed approach to security education can

be generalized beyond antiphishing to the extended security community. Rather than flooding users with information, we need to consider how users make decisions, both in business and personal settings, and tailor new security solutions based on this. ■

## References


1. R. Dhamija, J.D. Tygar, and M. Hearst, "Why Phishing Works," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 06), ACM, 2006, pp. 581–590.
2. S.E. Schechter et al., "The Emperor's New Security Indicators," *IEEE Symp. Security and Privacy*, IEEE CS, 2007, pp. 51–65.
3. P. Kumaraguru et al., "School of Phish: A Real-World Evaluation of Anti-phishing Training," *Proc. 5th Symp. Usable Privacy and Security* (SOUPS 09), ACM, 2009, pp. 1–12.
4. S. Sheng et al., "Anti-phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish," *Proc. 3rd Symp. Usable Privacy and Security* (SOUPS 07), ACM, 2007, pp. 88–99.
5. C. Herley, "So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users," *Proc. New Security Paradigms Workshop*, ACM, 2009, pp. 133–144.
6. F. Stajano and P. Wilson, "Understanding Scam Victims: Seven Principles for Systems Security," *Comm. ACM*, vol. 54, no. 3, 2011, pp. 70–75.
7. J. Riegelsberger, M.A. Sasse, and J.D. McCarthy, "The Mechanics of Trust: A Framework for Research and Design," *Int'l J. Human-Computer Studies*, vol. 62, no. 3, 2005, pp. 381–422.
8. D. Kim, D. Ferrin, and H. Rao, "A Trust-Based Consumer Decision-Making Model in Electronic Commerce: The Role of Trust, Perceived Risk, and Their Antecedents," *Decision Support Systems*, vol. 44, no. 2, 2008, pp. 544–564.
9. M. Koufaris and W. Hampton-Sosa, "The Development of Initial Trust in an Online Company by New Customers," *Information & Management*, vol. 41, no. 3, 2004, pp. 377–397.
10. B.G. Glaser and A.L. Strauss, *The Discovery of Grounded Theory: Strategies for Qualitative Research*, Aldine, 1967.
11. L. Leuthesser, C.S. Kohli, and K.R. Harich, "Brand Equity: The Halo Effect Measure," *European J. Marketing*, vol. 29, no. 4, 1995, pp. 57–66.
12. J. Riegelsberger, M.A. Sasse, and J.D. McCarthy, "Rich Media, Poor Judgement? A Study of Media Effects on Users' Trust in Expertise," *Proc. British HCI Conf.*, 2005, pp. 267–284.
13. R. Wash, "Folk Models of Home Computer Security," *Proc. 6th Symp. Usable Privacy and Security* (SOUPS 10), ACM, 2010, pp. 1–16.
14. M.A. Sasse et al., "Human Vulnerabilities in Security Systems," white paper, Cyber Security Knowledge Transfer Network, 2007.
15. S. Egelman, L.F. Cranor, and J. Hong, "You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings," *Proc. 26th SIGCHI Conf. Human Factors in Computing Systems*, ACM, 2008, pp. 1065–1074.





**Iacovos Kirlappos** is a PhD student in the Security Science Doctoral Research Training Centre and the Department of Computer Science at University College London. His research interests include human behavior in security and information security management. Kirlappos has an MSc in human-computer interaction and master of research in security science from University College London. He's a member of the Association for Information Systems. Contact him at iacovos.kirlappos.09@ucl.ac.uk.

**M. Angela Sasse** is the professor of human-centered technology in the Department of Computer Science at University College London. A usability researcher by training, she started researching human behavior in security and privacy in the '90s, and now heads the Information Security Research Group at UCL. Sasse has a PhD in computer science from the University of Birmingham. She's a member of the British Computer Society and the British Psychological Society. Contact her at a.sasse@cs.ucl.ac.uk.