

So Long, And No Thanks for the Externalities: The Rational Rejection of Security Advice by Users

Cormac Herley
Microsoft Research
One Microsoft Way
Redmond, WA, USA
cormac@microsoft.com

ABSTRACT

It is often suggested that users are hopelessly lazy and unmotivated on security questions. They chose weak passwords, ignore security warnings, and are oblivious to certificate errors. We argue that users' rejection of the security advice they receive is entirely rational from an economic perspective. The advice offers to shield them from the direct costs of attacks, but burdens them with far greater indirect costs in the form of effort. Looking at various examples of security advice we find that the advice is complex and growing, but the benefit is largely speculative or moot. For example, much of the advice concerning passwords is outdated and does little to address actual threats, and fully 100% of certificate error warnings appear to be false positives. Further, if users spent even a minute a day reading URLs to avoid phishing, the cost (in terms of user time) would be two orders of magnitude greater than all phishing losses. Thus we find that most security advice simply offers a poor cost-benefit tradeoff to users and is rejected. Security advice is a daily burden, applied to the whole population, while an upper bound on the benefit is the harm suffered by the fraction that become victims annually. When that fraction is small, designing security advice that is beneficial is very hard. For example, it makes little sense to burden all users with a daily task to spare 0.01% of them a modest annual pain.

Categories and Subject Descriptors: K.6.0: Economics

General Terms: Economics

Keywords: Security advice, user education

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSPW'09, September 8–11, 2009, Oxford, United Kingdom
Copyright 2010 ACM 978-1-60558-845-2/09/09 ...\$5.00.

1. INTRODUCTION

The range of attacks directed against Internet users is vast and growing. Their computers are constantly targeted by viruses, worms, port scanning software, spyware, adware, malware, keyloggers, rootkits, and zombie and botnet applications. One study reports that an unpatched Windows PC will be compromised within 12 minutes of connecting to the Internet [1]. Things get yet worse: according to Schneier “Only amateurs attack machines; professionals target people.” Users are the famously weak link in any security chain. It is easier to get information or passwords by social engineering than direct assault or brute-force. The best way to get software onto any machine is to get the user to install it and human error is behind many of the most serious exploits [41, 43].

The main response of the security community to these threats against the human link has been user education. Users are given instructions, advice and mandates as to how to protect themselves and their machines. See, *e.g.* the US-Cyber Emergency Response Team (US-CERT) tips for end users [13]. Most large web-sites offer security tips to users, as do software vendors. Yet the relationship between users and user education has been a rocky one. Adams and Sasse [21] found that low motivation and poor understanding of the threats leads users to circumvent password security policies. This is certainly borne out by other data: a study of password habits in 2007 [26] found that users still choose the weakest they can get away with, much as they did three decades earlier [45]. This is a discouraging finding, since few issues have seen more sustained effort at user education. There is considerable evidence that the failure of user education in the password space is repeated in other areas [36].

There are several ways of viewing this. A traditional view is that users are hopelessly lazy: in the face of dire descriptions of the threat landscape and repeated warnings, they do the minimum possible. A second view, advanced by a growing body of usable security researchers suggests that security tasks must be made more usable

and less cumbersome. Some argue that user education is key. In this paper we argue for a third view, which is that users' rejection of the security advice they receive is entirely rational from an economic viewpoint. The advice offers to shield them from the direct costs of attacks, but burdens them with increased indirect costs, or externalities. Since the direct costs are generally small relative to the indirect ones they reject this bargain. Since victimization is rare, and imposes a one-time cost, while security advice applies to everyone and is an ongoing cost, the burden ends up being larger than that caused by the ill it addresses.

We examine three areas where much effort has been put (to little apparent effect) in user education: password rules, phishing site identification, and SSL certificates warnings. In each we find that the advice is complex and growing, but the benefit is largely speculative or moot. For example, it makes little sense to invest effort in password strength requirements if phishing and keylogging are the main threats. It does not pay to learn URL reading rules to recognize phishing sites when the direct losses borne by users average less than a dollar a year. It's hard to blame users for not being interested in SSL and certificates when (as far as we can determine) 100% of all certificate errors seen by users are false positives.

Thus users ignore new advice for several reasons. First, they are overwhelmed. Given the sheer volume of advice offered no user has any real prospect of keeping up (*e.g.* the US-CERT advice [13] contains 51 tips, each of which fans out to at least a full page of text). Their effort budget for security matters is, in effect, exhausted (*i.e.* at a certain point any advice for which $\Delta\text{Cost} > 0$ cannot be accepted unless some other advice is abandoned). Second, in some cases the benefit is entirely moot, or is perceived by the user to be moot. For example, the benefit of choosing a strong password is entirely moot (*i.e.* $\Delta\text{Benefit} \equiv 0$) if the user has a keylogger on his machine. Equally, if an unpatched Windows machine "is infected within 12 mins" [1], then a user may wonder what is the point of even basic precautions? Third, the claimed benefits are not based on evidence: we have a real scarcity of data on the frequency and severity of attacks. So the absolute reduction of risk for any attack is speculative (*i.e.* $\Delta\text{Benefit} = ??$). Finally, security advice helps users *reduce exposure to the direct costs* of an attack while *increasing their indirect costs*. However, for many Internet crimes the externalities are many times greater than the direct dollar losses. And users are liable for only a part, if any, of the direct losses. This means that a fractional decrease in the direct losses (dollars) in exchange for an increase in externalities (effort) is simply a poor tradeoff. To make this concrete, consider an exploit that affects 1% of users annually, and they waste 10 hours clearing up

when they become victims. Any security advice should place a daily burden of no more than $10/(365 \times 100)$ hours or 0.98 seconds per user in order to reduce rather than increase the amount of user time consumed. This generates the profound irony that much security advice, not only does more harm than good (and hence is rejected), but does more harm than the attacks it seeks to prevent, and fails to do so only because users ignore it. In the model we set forward it is not users who need to be better educated on the risks of various attacks (as Adams *et al.* [21] suggest), but the security community. Security advice simply offers a bad cost-benefit tradeoff to users.

Few terms are as overworked as Kuhn's "Paradigm Shift" [51] to describe the revolution that is necessary when one way of thinking is incapable of keeping pace with developments. However the conditions that preceded the Copernican revolution mirror the current state of affairs: an existing system can be kept going only with constant patching. But the rate at which new patches are needed causes ballooning complexity, which ultimately cannot be supported. According to Kuhn:

"Given a particular discrepancy, astronomers were invariably able to eliminate it by making some particular adjustment in Ptolemy's system of compounded circles. But... astronomy's complexity was increasing far more rapidly than its accuracy and that a discrepancy corrected in one place was likely to show up in another."

We face an analogous problem with security and user training. Each attack necessitates a change in the advice users are given. While the warning may be important, and the information good, it increases the complex model with which we cram users' brains.

The rest of the paper is as follows. We next examine three different cases where cumulative advice brings questionable benefit. In Section 3 we look at passwords, in Section 4 training users to parse URLs, and in Section 5 certificate errors. Section 6 covers related work. In Section 7 we examine what lessons can be drawn from this.

2. COSTS, BENEFITS AND EXTERNALITIES

Users perform an implicit cost/benefit calculation when deciding whether to follow security advice or not. The cost is the effort to follow the advice, while the benefit is avoidance of the harm that the attack might bring. The harm includes the monetary loss (if any) that victims endure, but also the time and effort they must spend resolving the situation with the bank. Like many economic activities, Internet crime generates negative externalities: indirect costs not borne by the criminal

[39]. In spam for example, the amount of money made by the spammer may be very small relative to the infrastructure and bandwidth costs, and the time wasted by recipients. For example, Kanich *et al.* [32] document a campaign of 350 million spam messages sent for \$2731 worth of sales made. If 1% of the spam made it into inboxes, and each message in an inbox absorbed 2 seconds of the recipient’s time this represents 1944 hours of user time wasted, or \$28188 at twice the US minimum wage of \$7.25 per hour. Thus, one portion of the externalities of the campaign is more than 10× the direct dollar gain to the spammer. Of course the spammer doesn’t care whether the externalities are ten, a thousand, or a million times his direct gains: those are someone else’s costs, whereas he cares only about his gains.

Similarly, for other Internet crimes: the direct dollar amount gained by the criminal is far from a complete accounting of the damage caused. Let’s consider attacks involving online banking. Table 1 summarizes the costs of the attackers, the banks and users. As far as direct costs are concerned it is a zero-sum game: the attackers gain as much as the banks and victims combined lose. At present it appears that banks make whole any losses suffered by users (*e.g.* after initially refusing, Bank of Ireland refunded victims of a 2006 attack in full [2]). See also Section 3.3 for a description of bank reimbursement guarantees. Again, the attacker doesn’t much care how the banks and users divide the losses among themselves. For the externalities the picture is very different: it is a negative sum game. The attacker gains nothing, but the attacks generate substantial costs for the banks and users, far in excess of the direct losses. The externalities for the bank can include increased support call volume, damage to reputation and increased costs that results from reluctance of users to bank online. According to the Paypal CISO [10]: “Phishing was not just impacting consumers, in terms of general loss, it was impacting their view of the safety of the Internet and it was indirectly damaging our brand.” Indirect costs for victim users include the time they spend resolving the fraud case with their bank. Indirect costs for non-victim users includes the effort that they make to follow security advice, and possibly greater costs if they become afraid of banking or shopping online.

The goal of security advice is to protect users from certain attacks. Password strength rules protect them from brute-force and guessing attacks. URL reading protects them from phishing attacks. Identifying certificate errors protects them from MITM or web-spoofing attacks. If the user follows the advice the hope is that he will reduce or eliminate the risk of being a victim. However this addresses only the direct costs of the attack. Thus, for security advice, the Δ Benefit (reduction of direct losses) comes at the expense of Δ Cost (increase of effort). It is hard to make this calculation for an in-

	Direct Costs	Indirect costs (<i>i.e.</i> externalities)
Attackers	Gain	Don’t Care
Banks	Loss	Reputation
Victim Users	Possible Loss	Clean-up Effort
Non-victim Users	None	User Education

Table 1: Costs of online financial fraud. The direct costs are zero-sum: the attackers gain as much as the banks and victims lose. The externalities are indirect costs imposed on banks and non-victim users as they seek to avoid and deal with the consequences of the attacks. For many forms of fraud the externalities are many times greater than the direct costs.

dividual user. However aggregate estimates across the whole population are easier to reason about. Thus, we can try to determine whether

$$\sum_{\text{All Users}} \Delta\text{Cost} < \sum_{\text{All Users}} \Delta\text{Benefit}.$$

Of course it can be difficult to trace or predict the portion of a reduction in losses that springs from a particular piece of security advice. However if the *increase* in externalities is greater than the *total* direct losses, *i.e.*

$$\sum_{\text{All Users}} \Delta\text{Cost} > \text{Total Direct Losses}$$

then a piece of advice certainly represents a poor cost benefit tradeoff for the user population. For example, a piece of security advice that requires an hour per year for the average user to follow should reduce direct costs *to the users* by at least $\$180e6 \times 2 \times 7.25 = \2.6 bn (again using twice the minimum hourly wage of \$7.25 and an online population of 180 million) to be worthwhile. We will find that this is almost never the case with the attacks that we examine. Instead we find the direct costs are small, or unquantifiable, or borne by the banks rather than users, or are theoretical, protecting users against potential rather than actual losses. Thus the advice offered to users creates a greater burden than the attacks that it purports to save them from and is completely counter-productive.

3. PASSWORD RULES

3.1 The Costs

Passwords, and the rules that govern their choice, use and maintenance, are one of the main points of interaction between ordinary users and the security community. The habit of users of choosing weak passwords has caused web-sites to set policies that force minimum strength rules. Strength rules generally constrain passwords with respect to:

1. Length
2. Composition (*e.g.* digits, special characters)
3. Dictionary membership (in any language).

Many sites offer password strength meters that allow users to gauge the quality of passwords. Web-sites with a very loose policy may merely insist on a minimum length. At the other extreme are the rules for truly strong passwords [17]. For example Paypal recommends that a new password “is at least 8 characters long, is not a word you can find in the dictionary, includes both capital and low case letters, and contains at least one special character.” In addition there are many rules for how a user should handle the password once chosen. Again there is variation between the instructions offered by different sites. Commonly these rules include the following:

4. Don’t write it down
5. Don’t share it with anyone
6. Change it often
7. Don’t re-use passwords across sites.

Rules 4-7 are merely the most common policies usually given to users. Additional rules often cover such matters as never caching a password at a third-party proxy, or re-using old passwords (*e.g.* cycling back to a previously-used password when a change is forced). For a more complete list see [17].

Different web-sites will have policies that are restrictive to different degrees. This may be deliberate: it can help ensure that users do not share passwords between sites (*i.e.* violate Rule # 7) if they have very different strength rules. In fact using a password that is unique to that account is a requirement of many banks [36]. However, this increases the burden on users further. Florêncio and Herley estimate that users have an average of 25 password accounts to manage [26], and re-use is common, the average password being used at 3.9 different sites.

Insisting that users choose a unique strong password for each, which they change often and never write down is clearly a large burden. Adams and Sasse [21] surveyed users about password memorability, and also conclude that choosing secure memorable passwords proves a difficult task for many users.

3.2 The Benefits (potential)

It is clear that password policies impose a significant burden on users. However there is far from unanimous agreement on the benefits of many of these requirements.

Rules 1-3 cover password strength. Florêncio *et al.* [27] suggest that strength rules for web passwords accomplish very little when a lockout rule can restrict access. In this case a simple 6-digit PIN can suffice. Only when there is an off-line attack on the password does strength become very important. Strength above this minimum accomplishes very little.

Rule 4 enjoins users to avoid writing their password down. However, many security experts question this advice [3, 4]. It’s clear that writing passwords in plain view is bad practice, however, keeping them written in a safe place, such as a wallet, only increases the risk from someone who has access to the wallet. If the threat is an anonymous attacker rather than a knowledgeable opponent then following Rule 4 carries no benefit.

Rule 6 will help only if the attacker waits weeks before exploiting the password. So this amplifies the burden for little gain. Only if it is changed between the time of the compromise and the time of the attempted exploit does Rule 6 help.

Let’s examine the incremental cost/benefit tradeoff for a user who wishes to comply with Rule 7, *i.e.* he will no longer re-use passwords across sites. As estimated by Florêncio and Herley a typical user has 25 accounts and 6.5 passwords, each used at 3.9 sites (implying a lack of compliance with Rule 7). Thus, to comply ΔCost becomes a $3.9\times$ magnification of the number of passwords he must choose and remember. What can we say of $\Delta\text{Benefit}$; *i.e.* what risk is eliminated? Without observing Rule 7, *if* the user shares the same password between sites A and B, *and* the password from A is compromised *and* the one from B is not, *and* the attacker knows his userID at site B, *then* the site B account is exposed. So this is the risk eliminated in observing Rule 7. This would appear to include only the cases where the user is phished (rather than keylogged) or a rogue employee steals the credentials from A. This appears a minor reduction of risk for a $3.9\times$ magnification of password management effort.

Finally, none of Rules 1-5 help at all against phishing and keylogging; *i.e.* the advice is moot and the entirety of the effort wasted against these threats. Rules 6 and 7 are of marginal benefit. Thus, even if a user strictly observes each of the rules indicated above they are by no means safe from exploits that involve password theft.

3.3 The Benefits (actual)

The main attacks against passwords would appear to be [27]: phishing, keylogging, a brute-force attack on the user’s account, a bulk-guessing attack on all accounts at the server, and special-access attacks (guessing, shoulder surfing and console access).

As far as we are aware, there is no data available on strength related attacks on passwords of web-sites that maintain lockout policies. The recent attack on Twitter

appears to have succeeded since they were one of the few large sites that did not have such a policy [16]. It is harder still to separate out actual the costs (and thus the possible Δ Benefit) for Rules 4-7.

However, the Paypal CISO [5] states that “Forty-one basis points is the total fraud number” on Paypal’s system. Thus 0.49% of Paypal’s Transaction Volume, or \$290 million, would appear to upper bound all of the password related attacks. Given that Paypal had 70 million active users in 2008, all of the annual security advice should consume no more than $\$290/70 = \4.14 or about seventeen minutes of twice minimum wage time per year. But even this loose bound assumes that users are liable for the loss and can address it by following security advice.

Finally, in estimating the Δ Benefit in direct losses to the user we must determine how the attackers gains are shared (as losses) between the banks and the users. On this score it appears that banks currently absorb almost all losses. For example, Wells Fargo, in their online security guarantee states [15] “We guarantee that you will be covered for 100% of funds removed from your Wells Fargo accounts in the unlikely event that someone you haven’t authorized removes those funds through our Online Services.” Similarly, Fidelity’s Customer Protection Guarantee reads [12] “We will reimburse your Fidelity account for any losses due to unauthorized activity.” Other major banks have similar guarantees. While this may not be true of all banks, in the US unauthorized transfers from financial accounts are governed by Regulation E of the Federal Reserve Board [11]. This covers all transfers except by check and credit card, and limits the user’s liability to \$50 if the loss is reported within two days of discovery. Interestingly, even in cases involving negligence the user’s liability is limited: “Negligence by the consumer cannot be used as the basis for imposing greater liability than is permissible under Regulation E. Thus, consumer behavior that may constitute negligence under state law, such as writing the PIN on a debit card or on a piece of paper kept with the card, does not affect the consumer’s liability for unauthorized transfers.” Thus, users are entirely rational to reject any increase in effort which offers to save them from direct losses which appear small, and in most cases, borne by the banks.

4. TEACHING USERS TO RECOGNIZE PHISHING SITES BY READING URLS

4.1 The Costs

With the advent of phishing and other spoofing attacks it has become clear that users are easily confused as to what domain their browser is connected to. For example, a phishing site can be indistinguishable from the Paypal login page. Early phishing attacks were of-

ten hosted at numeric IP addresses and advertised in emails littered with spelling and grammatical errors. It thus seemed sensible to point out that these were obvious indications that might save users from attack. Phishers quickly evolved. It became common to spoof the actual name of the institution under attack using spelling mistakes that were visually similar to the target URL. Thus phishing URLs such as `www.paypa1.com` and `www.bankofthevest.com` became common. This new form of the attack required new advice: be aware of address-bar typos and visually similar URLs. More recently, phishers have used wildcard DNS entries to satisfy users’ expectations of what the URL should look like [9]. For example `www.paypal.com.login.evil.com`. This requires yet more revision of the instructions to users. Users who have listened to previous instructions may expect that when going to Paypal the the URL should contain “www.paypal.com.” However, in this attack, the URL does contain “www.paypal.com,” but not in the right place. The instructions must be revised again. In fact, this path of user education, leads in the direction of teaching users the rules for parsing URLs. Table 2 shows the evolving complexity of the URLs that phishers have employed. Users must understand that what appears as a link in an email or document is not necessarily the advertised link. For example

```
<a href="www.evil.com">www.PayPal.com</a>
```

may look to the user as a link to Paypal, but will of course take them to `www.evil.com`. This requires an understanding that the path is different from the host. In reading the host it requires an understanding that numeric IP addresses have unknown owners. That the DNS system is hierarchic; that the dot has special status, that hosts are read right to left and that the top-level domain and second level domain are special. That the second level domain is generally the most important indicator of who controls the site.

What started as seemingly reasonable advice to protect users from harm is evolving into a requirement to teach them how to read URLs. However, the “L” in URL stands for “locator”, not “location.” A URL is computer program rather than a pointer, albeit a simple one in a constrained language. Further, even if we teach users to read URLs, it requires that users have an expectation of what should appear in the address bar. Cached queries from both the Google and Yahoo! search engines load from numeric IP addresses, users who configure their own router will often reach the interface via a private numeric IP address, *e.g.* `192.168.*.*`. Thus warning users even about numeric IP addresses requires context and caveats. Many sites redirect users from one domain to another. Thus Paypal loads content from `www.paypalobjects.com` as well as `www.paypal.com` and so on. Many banking sites have URLs that do not

Address	Message to users
192.34.23.1	Numeric IP addresses are suspect
www.paypal.com	Address-bar typos
www.paypal.so	Incorrect top-level domain
www.geocities.com/www.paypal.com	Institution should appear in host rather than path
www-paypal-com.evil.com	Punctuation matters: ‘-’ ≠ ‘.’
www.paypal.com.evil.com	Domains are read right to left

Table 2: Increasing sophistication of phishing URLs requires increasing complexity of the security advice to users.

resemble or contain the bank name (*e.g.* Bank of Ireland online banking is done at `www.365online.com`).

Other exceptions abound. URLs such as `www.boi.com.nyud.net` are encountered when using a distributed Content Delivery Network such as CoralCDN [34] or link-translating proxy such as URRSA [37]. Pages that look indistinguishable from the PayPal login page are to be found via any of thousands of anonymizing proxies. Some very large sites, such as Amazon and eBay do have versions of their sites that load from different top-level domains. Thus advice on reading URLs must make clear that `www.amazon.co.uk` is a legitimate site that is controlled by Amazon, while `www.bankofthewest.co.uk` is not controlled by BankOfTheWest.

4.2 The Benefits (potential)

The main difficulty in teaching users to read URLs is that in certain cases this allows users to know when something is bad, but it never gives a guarantee that something is good. Thus the advice cannot be exhaustive and is full of exceptions.

Recall this is all to help the user avoid a single type of attack (phishing). Let’s again examine the incremental cost/benefit situation for a user. A user who conscientiously follows the rules on URL parsing shoulders a considerable burden. The Δ Benefit is that he avoids some subset of phishing sites. As before, this benefit is moot if there is a keylogger on his machine. If that is the case then his Δ Benefit for taking the trouble to recognize phishing sites is wasted.

4.3 The Benefits (actual)

As before the actual Δ Benefit to the user depends on how the bank and the user split their loss (*i.e.* the attacker’s gain). As mentioned in Section 3.3 it appears that most banks shoulder the entire loss [15, 12]. In this case Δ Benefit to the user is zero: he has no direct loss, even if he is phished.

Suppose not however: suppose that instead the entire burden falls on users. In previous work [28] we estimated US annual phishing losses at \$60 million. This upper bounds the annual value Δ Benefit of any advice that helps a single user avoid phishing at $\$60e6/180e6 = 0.33$ or 33 cents (assuming online population of $180e6$

in the US). That is, the best case value of any advice that helps users avoid phishing is worth less in direct losses than the cost of a first class stamp. Even for minimum wage users any advice that consumes more than $0.33/7.25 = 0.045$ hours or 2.6 minutes annually is a poor tradeoff. Thus any piece of advice that requires more than 2.6 minutes *per year* to follow is unprofitable from a cost benefit point of view.

This observation produces the following ironies. First, banks have more to fear from the indirect losses such as support costs generated by their own customers than their direct losses to phishers. For example, Wells Fargo has 48 million customers [14]. An agent-assisted password reset (estimated at \$10 per reset) by 10% of their users would cost \$48 million, easily dwarfing Wells Fargo’s share of the overall \$60 million phishing losses. Second, users are burdened more by the security advice surrounding phishing than all of the direct losses. In both cases it is the externalities of phishing rather than the direct costs that represent the true burden.

5. CERTIFICATE ERRORS

5.1 The Costs

SSL was put in place to protect content from a Man-In-The-Middle attack as it flows over the network. Notwithstanding some recent attacks it appears to serve this function well. Thus when SSL connected to `https://www.paypal.com` the content is encrypted between the browser and Paypal’s server. However, to verify that the browser is connected to the correct site the user must be able to verify and check the certificate. There is a great deal of evidence that users do not understand or notice the `https://` indicator [24, 48]. Further, they do not appear to notice the lock icon, or understand that it must appear in a certain position. For example, Dhamija and Tygar [23] find that users believed that a lock icon that appeared in the content of the page indicated that it was secure.

Thus to be able to verify that the browser is SSL connected to a particular site the user needs to understand that certain parts of the browser are different from others; *i.e.* that the “chrome” at the top and the bottom is different from the content served by the site. Even here

there is an exception, since what appears in the address bar is controlled by the site. For example, any site can trivially employ a picture of a lock as their `favicon.ico` image, and thus have a lock icon always appear in the address bar.

In addition, since the vast majority of web content is non-SSL, the user needs to know when he should check for a certificate. Here there is no simple rule. Some sites use SSL for the POST event that sends the password to the server, but otherwise leave content unencrypted (*e.g.* gmail, hotmail, yahoo, facebook, myspace *etc.*). On these sites the user will never see a certificate (but will see a certificate error if it is mis-configured). On other sites, including many bank sites, the main page is non-SSL, but if the user clicks on the login page the browser switches to SSL (*e.g.* citibank). Still others encrypt anything related to a user's login session (*e.g.* Paypal). Thus, even if a user can be relied upon to check certificates, it is far from simple to tell a user when he should see one. Sobey *et al.* [30] suggest that Extended-Validation (EV) certificates may actually generate more confusion than existed before their introduction. Sunshine *et al.* [31] find that users are effectively trained to ignore certificate warnings by seeing them repeatedly when there is no real security threat.

5.2 The Benefits (potential)

The protection from a MITM attack is a powerful incentive to use SSL. However, to eliminate the possibility of a MITM attack the user must type the entire URL, including the method. For example, consider the following ways of navigating to PayPal:

1. Type `https://www.paypal.com`
2. Type `http://www.paypal.com` and get redirected
3. Type `paypal` Cntrl-Enter (browser adds `www.` and `.com`)
4. Search for "paypal" using google and click link
5. Click bookmarked site `https://www.paypal.com`
6. Click bookmarked site `http://www.paypal.com` and gets redirected

In 2, 3, 4, and 6 the user goes over the open network unencrypted and doesn't get the protection of SSL. Paypal redirects requests for `http://www.paypal.com` to `https://www.paypal.com` (*i.e.* directs the browser to use SSL), but by then it could be too late. For example, a bad router can take the user to a spoof site `www.paypal.com.bad.com` and provide a perfectly valid certificate. Thus, even to get protection from a MITM attack the user must either bookmark the SSL site, or type the full URL and method; *i.e.* use method 1 or 5. There is evidence that few users do this [6]. Instead typing into the search bar appears to be a main means by which users navigate to sites.

5.3 The Benefits (actual)

Browser vendors have invested considerable effort in making it harder to ignore certificate errors. In Firefox version 3, when encountering an expired, invalid or self-signed certificate the user sees an interrupt page explaining that the SSL connection failed. If he chooses to add an exception he sees another interrupt page with more warnings and a choice to add an exception or "get me out of here." If he elects (again) to add an exception he must click to get the certificate, view the certificate, and then add the exception. Internet Explorer 8 is somewhat less intrusive, but the procedure also seems designed to suggest that adding exceptions is very risky. Is it? Ironically, one place a user will almost certainly never see a certificate error is on a phishing or malware hosting site. That is, using certificates is almost unknown among the reported phishing sites in PhishTank [7]. The rare cases that employ certificates use valid ones. The same is true of sites that host malicious content. Attackers wisely calculate that it is far better to go without a certificate than risk the warning. In fact, as far as we can determine, there is no evidence of a single user being saved from harm by a certificate error, anywhere, ever.

Thus, to a good approximation, 100% of certificate errors are false positives. Most users will come across certificate errors occasionally. Almost without exception they are the result of legitimate sites that have name mismatches, expired or self-signed certificates. Thus the average user has seen certificate errors, but purely as an annoyance and never as something that saved him from harm. Of course, even if 100% of certificate errors are false positives it does not mean that we can dispense with certificates. However, it does mean that for users the idea that certificate errors are a useful tool in protecting them from harm is entirely abstract and not evidence-based. The effort we ask of them is real, while the harm we warn them of is theoretical.

6. RELATED WORK

There has been a great deal of work in the last few years on the failure of user education to achieve the desired goals. In fact, we cannot give more than a sampling of the work in this area. The New Security Paradigms Workshop (NSPW) has published numerous papers advocating for user-centric approaches to security. A panel discussion by Greenwald *et al.*[49] for example, examined the user boycott of security and questioned whether they have too much rather than too little security. Since 2005 the Symposium on Usable Privacy and Security (SOUPS) has provided a forum for Usable Security research.

Zurko and Richards [38] introduced the term user-centered security in 1996 to refer to systems that have usability as their primary goal. Adams and Sasse [21]

found that choosing strong memorable passwords was a serious challenge for most users, and that many were unmotivated owing to a poor understanding of the risks. Dhamija and Tygar [23] report that users are confused as to the distinction between the content of a web-page and the chrome of the browser. They also find that getting users to act on the absence of a security indicator is very difficult. Wu *et al.* [53] find that users ignore the warnings provided by anti-phishing toolbars. Anandpara *et al.* [52] find that many phishing IQ tests measure fear but not ability to tell good sites from bad. Whitten and Tygar [24] find that most users are unable to successfully navigate an encryption software package. Schechter *et al.* [48] find that users are largely oblivious to the presence or absence of the lock icon when logging in to a bank account, and are easily persuaded to ignore the absence of the SiteKey mutual authentication image. Egelman *et al.* [47] looked at the effectiveness of browser phishing warnings. Jakobsson *et al.* [35] detail the results of a user study to determine which cues induce trust and find that users' assessment of trustworthiness often relies on cues not designed as security features. Sunshine *et al.* [31] find in a survey of over 400 users that a majority ignored SSL warnings in a wide variety of conditions. Stewart and Martin [25] study the efficacy of warnings. They suggest, that to be effective, warnings must communicate the risks clearly, and give easily understood instructions for avoiding the harm. Jackson *et al.* [22] point out that EV certs do not necessarily improve users' understanding of the security context. Sobey *et al.* [30] suggest that EV certificates have actually generated more confusion. Mannan and van Oorschot [36] carry out a detailed examination of the usability of online banking. A study involving 123 users finds a large gap between banks security policy expectations and users' actions.

Florêncio *et al.* [27] first do a detailed examination of the costs and benefits associated with passwords. Numerous others have touched this issue over the years. Bellovin [46] also questions whether accepted wisdom such as password advice on security checklists are accomplishing their desired goals.

Anderson *et al.* [42] first proposed the comprehensive examination of security from an Economics perspective. He observed, for example, that economists have long studied how misaligned incentives often produce undesired outcomes, and many of these results carry lessons for security. Since 2001, the Workshop on the Economics of Information Security has explored these and other areas of overlap between economics and security. For example, there has been much interesting work on the establishment of a market for security vulnerabilities [20] and the Economics of Privacy [18]. The notion of security as an externality (*i.e.* the fact that the direct dollar losses are far from being a complete accounting

of the problem) has also been examined by Anderson [44]. Herley and Florêncio [28, 29] document that direct losses due to phishing and certain other forms of cybercrime are far lower than generally estimated.

The usable security community has produced many approaches that reduce the cost of teaching users. Security Cartoon [8] attempts to convey real world advice in a format that is easily digested by users. PhishGuru [40], addresses phishing education to users after they have responded to a fake phishing message. Similarly APWG/CMU Phishing Education Landing Page Program (which replaces discovered phishing pages with an educational site: <http://education.apwg.org/r/about.html>) is an excellent example that manages to target the at-risk population at no cost to the larger population.

The most closely related work is that of Beautelement *et al.* [19]. They find that bypassing security policies on how data is to be handled is a widely employed practice. For example, mandates that sensitive data on laptops and portable drives be encrypted are routinely ignored. They also introduce the idea of compliance budget, which formalizes the understanding that users, and organizations, do not have unlimited capacity to follow new instructions and advice. In the language of Beautelement *et al.* our work can be seen as a demonstration that users are effectively in an impossible compliance regime. Many authors have drawn attention to the large usability gap in security offerings. Zurko [33], for example, suggests that usability problems are a sign of failure to know the audience, a sign of fundamental trouble in any business. Adams and Sasse [21] point out that many security policies encourage an adversarial relationship with users. This paper can be seen as an extension of this line of enquiry.

7. ANALYSIS AND IMPLICATIONS

While we argue that it is rational for users to ignore security advice this does not mean that the advice is bad. In fact much, or even most of it is beneficial. It's better for users to have strong passwords than weak ones, to change them often, and to have a different one for each account. That there is benefit is not in question. However, there is also cost, in the form of user effort. In equilibrium, the benefit, to the user population, is balanced against the cost, to the user population. If observed user behavior forms the scales, then the decision has been unambiguous: users have decided that the cost is far too great for the benefit offered. If we want a different outcome we have to offer a better tradeoff. We examine next how we got things so wrong, and look at ways to make things better.

7.1 Users Understand Risks better than We do

In one view of the world users are ignorant of the

risks they face and must be educated to save them from themselves. “If only they understood the dangers,” the thinking goes, “they would behave differently.” However, this pre-supposes that we understand the risks better than users. Do we? Do we have evidence to demonstrate that users who follow advice fare better than those who ignore it? And that the difference is worth the extra effort? Remember, to get users to change we have to persuade them, not merely that there is benefit, but that, in deciding that the benefit is not worth the cost, they have seriously miscalculated.

Here we run into the well-known lack of data in security. What percent of users have their accounts compromised because of password strength? How many had compromises because they wrote the password down, did not change it often enough or re-used across sites? We quite simply don’t know. As with passwords (arguably the most visible aspect of user security) so with the other subjects of advice: in many cases we simply have no information on successful compromise rates for most attacks and thus can’t show that the cost-benefit calculation is favorable.

7.2 Worst-Case Harm and Actual Harm are not the Same

In the absence of actual compromise data the security community often speaks of worst-case risk. While worst-case analysis is a necessary tool in analyzing systems and protocols it turns out to be a poor way of motivating users. For example, the worst-case outcome of ignoring a certificate error is falling prey to a MITM attack, having money stolen and spending a great deal of time on cleanup. The average outcome is that nothing of the kind happens: the user endures a few annoying warnings and proceeds as before. In fact, as we’ve said, if there’s evidence that users are ever saved from harm by certificate errors we are unaware of it. Similarly, the worst-case outcome of sharing a weak password across several accounts is that one is brute-forced and all are compromised. If one of them is an email contact address for further accounts, the effects might ripple outward even more. However, the average outcome is very different: the vast majority of users ignore the strength and re-use advice [26] without apparent ill effect.

This difference between worst-case and actual outcomes causes a profound disconnect between what security advice offers and what users respond to. The worst-case is a factor $1/p$ greater than the actual harm (where p is the fraction of users that fall victim to some attack annually). When p is small (*i.e.*, annual victimization rate is low) this factor can be enormous. For certificate errors we’ve seen that $p \approx 0$. Even for an attack as visible as phishing $p = 0.0037$ [28, 50], which ensures that the worst-case harm is orders of magnitude more severe than what actually happens. Thus

it is users who show a better understanding of the actual risks than those who would school them. The wisdom of the crowd discerns that ignoring some threats brings little actual harm, even when they receive warnings about what might happen. Thus a main part of the problem with security advice is that we hugely exaggerate benefits. The advice is offered as protection against worst-case harms, while users care only about average or actual harm. Further the actual harms of some attacks, such as phishing appear greatly exaggerated [28], indicating that if we knew the actual harms it is likely that user behavior might still not change.

7.3 User Effort is not Free

In addition to overestimating benefits, advice almost always ignores the cost of user effort. The incremental cost of forcing users to choose an 8-character strong password, as opposed to allowing a 6-digit PIN, is hard to measure, but is certainly not zero. And ignoring it leads to a failure to understand the rational and predictable nature of user response.

There are about 180 million online adults in the US. At twice the US minimum wage one hour of user time is then worth $\$7.25 \times 2 \times 180e6 = \2.6 billion. A minute of user time per day is a $\$7.25 \times 2 \times 180e6 \times 365/60 = \15.9 billion per year proposition. This places things in an entirely new light. We suggest that the main reason security advice is ignored is that it makes an enormous miscalculation: it treats as free a resource that is actually worth \$2.6 billion an hour. It’s not uncommon to regard users as lazy or reluctant. A better understanding of the situation might ensue if we viewed the user as a professional who bills at \$2.6 billion per hour, and whose time is far too valuable to be wasted on unnecessary detail. Echoing Adams and Sasse [21] we might say: the user is your boss’s boss’s boss. This would help ensure that we ask for a minute of user time (the boss’s) only when absolutely necessary.

When we ignore the costs of security advice we treat the user’s attention and effort as an unlimited resource. Advice, policies and mandates thus proliferate. Each individual piece of advice may carry benefit, but the burden is cumulative. Just as villagers will overgraze a commonly held pasture, advice-givers and policy-mandaters demand far more effort than any user can give.

7.4 Designing Security Advice is not an Unconstrained Optimization

A consequence of ignoring cost is that advice is evaluated purely against its ability to reduce risk. We argue that a realistic evaluation must also include cost. For example, suppose a fraction p of users suffer harm H from particular exploit annually. Security advice which reduces the likelihood of being a victim by Δp

delivers benefit: $\Delta\text{Benefit} = \Delta p \cdot H$. Any advice for which $\Delta p \geq 0$ can then be argued to be beneficial. This appears to be the commonly applied criterion. We argue for a higher standard, and suggest that unless $\Delta\text{Cost} < \Delta p \cdot H$ the advice does more harm than good and is destined to be ignored. Thus, rather than the unconstrained optimization (where $\Delta\text{Benefit}$ is be maximized) we have a constrained one (where it is be maximized subject to having $\Delta\text{Benefit} > \Delta\text{Cost}$).

The difference between the constrained and unconstrained optimization is stark. For certificate errors we saw that $p \approx 0$. Thus, while almost any advice can satisfy $\Delta\text{Benefit} \geq 0$, almost none has $\Delta\text{Cost} < \Delta p \cdot H$. That is, most advice, when imposed on the whole population, ends up having larger cost than the existing total harm $p \cdot H$ which is very small. For password strength and policies, p is unknown for most attacks. Again, it is still easy to argue that $\Delta\text{Benefit} \geq 0$, but hard to meet the higher threshold. Finally, for phishing, where we have an estimate that $p = 0.0037$ [28, 50], advice that satisfies the lower criterion appears downright harmful under the second. That is, when $\Delta\text{Cost} > \Delta p \cdot H = \Delta\text{Benefit} \geq 0$, while the advice is in theory delivering benefit, it is, if followed, doing more harm than good.

7.5 The Economic Harm of Security Advice

We’ve seen that when the cost is greater than the benefit ($\Delta\text{Cost} > \Delta p \cdot H = \Delta\text{Benefit} \geq 0$) security advice does more harm than good. However, it can be even worse. If the cost is greater than the entire harm caused by the attack ($\Delta\text{Cost} > p \cdot H$) then the advice doesn’t merely do more harm than good, it does more harm than the attack it addresses. For example, suppose some security advice reduces the risk of becoming a phishing victim by 50%. If phishing victimizes 0.37% of users per year [28, 50] and each victim wastes 10 hours sorting it out, to be beneficial the daily effort of following the advice should be less than $0.0037 \times 0.5 \times 10/365$ hours or 0.18 seconds per day. Clearly, a user who makes the effort to read URLs to identify phishing sites will spend more time than this. Thus the advice is, in expectation, doing more harm than good. But worse, the advice is doing more harm than phishing itself unless it consumes less than 0.36 seconds per user per day. That is, suppose identifying phishing sites by reading URLs consumed a minute per user per day, or 365 minutes per year. An upper bound on the benefit is the entire elimination of the risk (*i.e.* reduce to zero the 0.37% chance that the user wastes 10 hours of cleanup time). In asking 365 minutes to reduce an expected loss of $0.0037 \times 10 \times 60 = 2.22$ minutes we are doing $365/2.22 = 164\times$ more harm than the attack itself. Hence the attack consumes 2.22 minutes per user per year on average, while the defence consumes 365. Mapping to dollars using the hourly

rate of users introduced in Section 7.3 we find that the cost of phishing to victims in terms of clean-up time is $\$2.6e9 \times 2.22/60 = \96 million per year, while the cost of the advice is \$15.9 billion. Throwing in the \$61 million in direct losses makes no material difference to the conclusion. Thus, this advice fails even the basic “first do no harm” principle.

7.6 So What Can We Do?

We do not wish to give the impression that all security advice is counter-productive. In fact, we believe our conclusions are encouraging rather than discouraging. We have argued that the cost-benefit tradeoff for most security advice is simply unfavorable: users are offered too little benefit for too much cost. Better advice might produce a different outcome. This is better than the alternative hypothesis that users are irrational. This suggests that security advice that has compelling cost-benefit tradeoff has real chance of user adoption. However, the costs and benefits have to be those the user cares about, not those we think the user ought to care about. We outline some general directions.

First, we need better understanding of the actual harms endured by users. There has been insufficient attention to the fact that it is mainly time, and not money, that users risk losing when attacked. It is also time that security advice asks of them. A main finding of this paper is that we need an estimate of the victimization rate for any exploit when designing appropriate security advice. Without this we end up doing worst-case risk analysis, and this can lull us into thinking that we are offering orders of magnitude more benefit than is actually the case.

Second, user education is a cost borne by the whole population, while offering benefit only to the fraction that fall victim. Thus the cost of any security advice should be in proportion to the victimization rate. This implies that it may be difficult, or impossible, to design advice for really rare exploits. This also suggests that user education technologies that can target the at-risk population (such as PhishGuru [40] and the APWG landing page) will show a far better cost-benefit ratio.

Third, retiring advice that is no longer compelling is necessary. Many of the instructions with which we burden users do little to address the current harms that they face. Retiring security advice can be similar to declassifying documents, with all cost and no benefit to the person making the decision. However, an ever-growing burden for users leads to rejection of all advice.

Fourth, we must prioritize advice. In trying to defend everything we end up defending nothing. In attempting to warn users of every attack, and patch every vulnerability with advice we have taken the user (our boss’s boss’s boss) off into the weeds. Since users cannot do everything, they must select which advice they will follow

and which ignore. In failing to prioritize we deny them any effective means to make this selection in a sensible way. In fact prioritizing advice may be the main way to influence the security decisions that users make. When we provide long lists of unordered advice we abdicate all opportunity to have influence and abandon users to fend for themselves.

Finally, we must respect users' time and effort. Viewing the user's time as worth \$2.6 billion an hour is a better starting point than valuing it at zero. We must understand that when budgets are exhausted, attention to any one piece of advice is achieved only by neglect of something else. Thus, when we exaggerate one danger we reduce the attention that can be paid to another. When we exaggerate all dangers we simply train users to ignore us.

8. CONCLUSION

"Given a choice between dancing pigs and security, users will pick dancing pigs every time." While amusing, this is unfair: users are never offered security, either on its own or as an alternative to anything else. They are offered long, complex and growing sets of advice, mandates, policy updates and tips. These sometimes carry vague and tentative suggestions of reduced risk, never security. We have shown that much of this advice does nothing to make users more secure, and some of it is harmful in its own right. Security is not something users are offered and turn down. What they are offered and do turn down is crushingly complex security advice that promises little and delivers less.

How can we help users avoid harm? This begins with a clear understanding of the actual harms they face, and a realistic understanding of their constraints. Without these we are proceeding blindly. Users, we have seen, are not irrational: exhaustive lists that seek to avoid all potential harms are not helpful to them and are ignored. If we want a different outcome we must present a better tradeoff. How did we manage to get things so wrong? In speaking of worst-case rather than average harm we have enormously exaggerated the value of advice. In evaluating advice solely on benefit we have implicitly valued user time and effort at zero.

Note: this paper is not to be read as an encouragement to end-users to ignore security policies or advice. The opinions expressed are those of the author.

Acknowledgement: the author wishes to especially thank Richard Ford and Mary Ellen Zurko for numerous detailed suggestions that helped enormously. He also thanks all of the attendees of NSPW 2009 for many questions and comments that improved the presentation, and Dinei Florêncio for discussions that have greatly influenced his thinking.

9. REFERENCES

- [1] <http://isc.sans.org/survivaltime.html>.
- [2] <http://www.vnunet.com/vnunet/news/2163714/bank-ireland-backtracks>.
- [3] http://www.theregister.co.uk/2005/07/19/password_schneier/.
- [4] http://www.schneier.com/blog/archives/2005/06/write_down_your.html.
- [5] http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1256995,00.html.
- [6] http://www.readwriteweb.com/archives/will_mainstream_users_ever_learn.php.
- [7] <http://www.phishtank.com>.
- [8] <http://www.securitycartoon.com>.
- [9] Making Waves in the Phishers Safest Harbor: Exposing the Dark Side of Subdomain Registries. http://www.antiphishing.org/reports/APWG_Advisory_on_Subdomain_Registries.pdf.
- [10] Phishers get more wily as cybercrime grows. <http://www.reuters.com/article/technologyNews/idUSTRE53G01620090417?feedType=RSS&feedName=technologyNews>.
- [11] Regulation E of the Federal Reserve Board. <http://ecfr.gpoaccess.gov/cgi/t/text/text-idx?c=ecfr&sid=0283a311c8b13f29f284816d4dc5aeb7&rgn=div9&view=text&node=12:2.0.1.1.6.0.3.19.14&idno=12>.
- [12] The Fidelity Customer Protection Guarantee. <http://personal.fidelity.com/accounts/services/findanswer/content/security.shtml.cvsr?refpr=custopq11>.
- [13] US-Cyber Emergency Response Readiness Team: CyberSecurity Tips. <http://www.us-cert.gov/cas/tips/>.
- [14] Wells Fargo News Release, Jan 1, 2009. https://www.wellsfargo.com/press/2009/20090101_Wachovia_Merger.
- [15] Wells Fargo: Online Security Guarantee. https://www.wellsfargo.com/privacy_security/online/guarantee.
- [16] Wired: Weak Password Brings 'Happiness' to Twitter Hacker. <http://blog.wired.com/27bstroke6/2009/01/professed-twitt.html>.
- [17] Department of Defense Password Management Guideline. Technical Report CSC-STD-002-85, U.S. Dept. of Defense, Computer Security Center, 1985.
- [18] A. Acquisti and J. Grossklags. Uncertainty, Ambiguity and Privacy. *WEIS*, 2005.
- [19] A. Beaument, M.A. Sasse and M. Wonham. The Compliance Budget: Managing Security Behaviour in Organisations. *NSPW*, 2008.

- [20] A. Ozment and S. Schechter. Milk or wine: does software security improve with age? *Usenix Security*, 2006.
- [21] A. Adams and M. A. Sasse. Users Are Not the Enemy. *Commun. ACM*, 42(12), 1999.
- [22] C. Jackson, D.R. Simon, D.S. Tan and A. Barth. An Evaluation of Extended Validation Certificates and Picture-in-Picture Phishing Attacks. *Proc. Usable Security*, 2007.
- [23] R. Dhamija and J. D. Tygar. The battle against phishing: Dynamic security skins. *Symp. on Usable Privacy and Security*, 2005.
- [24] R. Dhamija, J. D. Tygar, and M. Hearst. Why phishing works. *CHI*, 2006.
- [25] D.W. Stewart and I. M. Martin. Intended and Unintended Consequences of Warning Messages: A Review and Synthesis of Empirical Research. *J. of Public Policy and Marketing*, 1994.
- [26] D. Florêncio and C. Herley. A Large-Scale Study of Web Password Habits. *WWW 2007, Banff*.
- [27] D. Florêncio, C. Herley, and B. Coskun. Do Strong Web Passwords Accomplish Anything? *Proc. Usenix Hot Topics in Security*, 2007.
- [28] C. Herley and D. Florêncio. A Profitless Endeavor: Phishing as Tragedy of the Commons. *NSPW 2008, Lake Tahoe, CA*.
- [29] C. Herley and D. Florêncio. Nobody Sells Gold for the Price of Silver: Dishonesty, Uncertainty and the Underground Economy. *WEIS 2009, London*.
- [30] J. Sobey, R. Biddle, P.C. van Oorschot and A.S. Patrick. Exploring User Reactions to New Browser Cues for Extended Validation Certificates. *ESORICS*, 2008.
- [31] J. Sunshine, S. Egelman, H. Almuhiemedi, N. Atri and L. F. Cronor. Crying Wolf: An Empirical Study of SSL Warning Effectiveness. *Usenix Security*, 2009.
- [32] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM Conference on Computer and Communications Security*, pages 3–14, Alexandria, Virginia, USA, October 2008.
- [33] M. E. Zurko. User-Centered Security: Stepping Up to the Grand Challenge. *ACSAC*, 2004.
- [34] M. J. Freedman and E. Freudenthal and D. Mazières. Democratizing Content Publication with Coral. *NSDI*, 2004.
- [35] M. Jakobsson, A. Tsow, A. Shah, E. Blevis and Y-K Lim. What Instills Trust? A Qualitative Study of Phishing. *Proc. Usable Security*, 2007.
- [36] M. Mannan and P.C. van Oorschot. Security and Usability: The Gap in Real-World Online Banking. *NSPW*, 2007.
- [37] Z. Mao and C. Herley. A Robust Link-Translating Proxy Mirroring the Whole Web. *ACM SAC 2010*.
- [38] M.E. Zurko and R. T. Simon. User-Centered Security. *NSPW*, 1996.
- [39] N.G. Mankiw. Principles of Economics. *4-th ed.*, 2007.
- [40] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, J. Hong. Testing PhishGuru in the Real World. *SOUPS*, 2007.
- [41] R. Anderson. Why Cryptosystems Fail. In *Proc. CCS*, 1993.
- [42] R. Anderson. Why Information Security is Hard. In *Proc. ACSAC*, 2001.
- [43] R. Anderson. Security Engineering. In *Second ed.*, 2008.
- [44] R. Anderson and T. Moore. The Economics of Information Security. *Science Magazine*, 2006.
- [45] R. Morris and K. Thompson. Password Security: A Case History. *Comm. ACM*, 1979.
- [46] S. Bellovin. Security by Checklist. *IEEE Security & Privacy Mag.*, 2008.
- [47] S. Egelman, L. F. Cronor and J. Hong. You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. *CHI*, 2008.
- [48] S. Schechter, R. Dhamija, A. Ozment, I. Fischer. The Emperor's New Security Indicators: An evaluation of website authentication and the effect of role playing on usability studies. *IEEE Security & Privacy*, 2007.
- [49] S.J. Greenwald, K.G. Olthoff, V. Raskin and W. Ruch. The User Non-Acceptance Paradigm: INFOSEC's Dirty Little Secret. *NSPW*, 2004.
- [50] T. Moore and R. Clayton. Examining the Impact of Website Take-down on Phishing. *Proc. APWG eCrime Summit*, 2007.
- [51] T.S. Kuhn. The Structure of Scientific Revolutions. 1962.
- [52] V. Anandpara, A. Dingman, M. Jakobsson, D. Liu, and H. Roinestad. Phishing IQ Tests Measure Fear, Not Ability. *Proc. Financial Crypto*, 2007.
- [53] M. Wu, R. Miller, and S. L. Garfinkel. Do Security Toolbars Actually Prevent Phishing Attacks. *CHI*, 2006.