

Why Trust Seals Don't Work: A study of user perceptions and behavior

Iacovos Kirlappos, M. Angela Sasse, University College London, Department of Computer Science, London, United Kingdom,
email: {i.kirlappos,a.sasse}@cs.ucl.ac.uk

Nigel Harvey, University College London, Department of Psychology, London, United Kingdom,
email: n.harvey@ucl.ac.uk

Abstract. Trust seals, such as the *VeriSign* and *TRUSTe* logos, are widely used to indicate a website is reputable. But how much protection do they offer to online shoppers? We conducted a study in which 60 experienced online shoppers rated 6 websites – with and without trust seals - based on how trustworthy they perceived them to be. Eye tracking data reveals that 38% of participants failed to notice any of the trust seals present. When seals were noticed, the ratings assigned to each website were significantly higher than for the same website without a seal, but qualitative analysis of the interview data revealed significant misconceptions of their meaning (e.g. “*presence of seals automatically legitimizes any website*”). Participants tended to rely on self-developed – but inaccurate – heuristics for assessing trustworthiness (e.g. perceived investment in website development, or references to other recognizable entities). We conclude that current trust seals currently do not offer effective protection against scam websites; and suggest that other mechanisms – such as automatic verification of authenticity are required to support consumers’ trust decisions.

1 Introduction

Trust plays a vital role in the commercial world: people and organizations cooperate to achieve mutual benefits, and the success of business transactions depends on both parties behaving in a collaborative way. The wide success of e-commerce since the early 2000’s [1] posed a major challenge for consumers and merchants: how to reach a transaction-enabling level of trust between them, without the traditional trust development medium - face-to-face interaction [2]. Attackers soon exploited the opportunities this new setting created: they started setting up fake online stores, pretending to sell popular products at tempting prices, but actually stealing consumers’ money and credit card details [3],[1]. At best, consumers receive counterfeit goods. At worst, they receive no goods, lose money and suffer identity theft. Some financial institutions, like credit card issuers, have introduced buyer protection mechanisms that cover their customers for any monetary losses [4], but consumers still have to go through time-consuming processes to obtain new credit cards, and monitor their accounts and credit reports to prevent identity theft using their stolen credentials.

A number of different measures have been introduced to address this problem: anti-phishing tools, trust seals and user education. But the number of scam websites and the reported losses are still alarmingly high [5-8] – UK card fraud crime amounted to £365.4 million in 2011 [9], and in the US online merchants lost \$2.7 billion to fraud in 2010 only [10]. The persistence of criminals operating online suggests that it is worth their effort.

Trust seals were created to make it easier for consumers to identify trustworthy websites. Their effectiveness has been discussed by a number of research reports [11-18], but all used experimental designs that explicitly drew respondents’ attention to presence of the trust seals (e.g. surveys). This paper presents an experiment in which we observed participants’ reaction to the same websites with and without trust seals, without directing their attention to them. We also conducted a detailed debrief to elicit their “*folk perceptions*” [19] of trustworthiness indicators in a website, including trust seals. We identified a number of trust-development heuristics consumers use to verify a website’s authenticity, and identify those as targets for future security awareness approaches.

2 Background

2.1 E-commerce and Trust

Trust plays a significant role in online environments, as it enables transactions between parties that are separated in both space and time. Riegelsberger et al. [2] outline the basic trustor-trustee interaction in technology-mediated inter-

actions (Figure 1): a consumer (trustor) uses the signals (1) emitted by the merchant (trustee) to assess their trustworthiness before proceeding to the trusting action (2a) [2], which increases their exposure to a trustee’s potential misbehavior, but provides the potential for positive gains if the merchant fulfills (3a). In e-commerce, increased exposure comes as a result of sharing financial and personal details with a website, as consumers now rely on the merchant’s behavior to reduce the likelihood of a negative outcome (e.g. goods not arriving, selling of counterfeit products, credit card details compromised, identity theft) [20]. As a result, the higher the perceived trustworthiness of the website, the more likely a consumer is to proceed to initiating a transaction.

Other researchers have also stressed the importance of trust to enable successful commercial transactions over the Internet. Nielsen [21] defines e-commerce related trust as “A user’s willingness to risk time, money and personal data on a website”, and others have underlined its importance for the success of e-commerce [11-13],[22]: the lower the transaction-related uncertainty appears to be to a consumer, the more likely they are to act in a way that renders them vulnerable to the behavior of an online merchant.

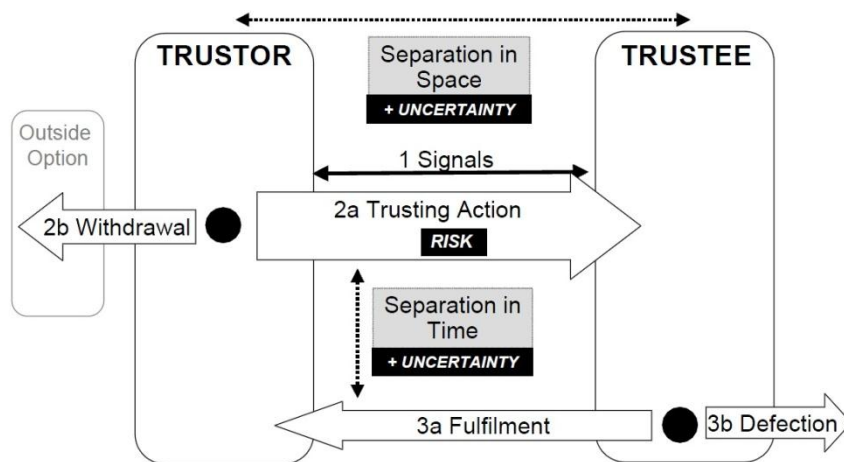


Fig. 1: The basic trustor-trustee interaction (Source: Riegelsberger et al. [2])

Another challenge to designing for trust in e-commerce environments is that, after a small number of successful transactions, consumers expect all further transactions to be successful, too – they now are in a state of *reliance*, rather than trust [2]. In this case, consumers extrapolate their past positive or negative experiences with some websites to new ones, resulting to a “*trust spillover*” effect: After a number of successful transactions, consumers are less likely to check for the *trust-warranting* [2] properties they searched for on the very first time they bought something online, or they may spend less time doing so. As a result trustworthy behavior of some merchants can lead to trust in the entire online market [23-25].

2.2 How Attackers Exploit Trust

Consumers shopping online are looking for ‘good deals’: trying to save money on regular purchases, or acquire something they would otherwise not be able to afford. In this situation, the classic economics problem of *information asymmetry* [26] applies: the lack of personal interaction results to consumers having less knowledge on the intentions of the merchant to fulfill in the transaction or the quality of the products offered [14],[27],[28]. The merchant, on the other hand possesses the advantage to wait until they have received full payment before shipping out any products [20]. This asymmetric nature of the interaction makes the exploitation of human needs easier for attackers, who are extremely skillful in exploiting human vulnerabilities [29]: Once they know what consumers want, they tempt them with “*too good to miss*” deals, but ship either counterfeit products or nothing at all.

2.3 Trust signaling

The trustor’s perceived uncertainty in the outcome of a transaction can be reduced by communicating information about the trustee’s ability and motivation to fulfill. In online markets, technology is not only the medium over which

an interaction happens, but also the medium over which both parties signal the trust-warranting properties required to allow the formation of positive expectations on the behavior of each other, which then provide the level of trust required by the trustor to initiate the interaction. Riegelsberger et al. [30] identify two main types of signals of trustworthiness trustees can emit:

- *Symbols*: They have arbitrarily assigned meaning, and were designed specifically to signal the presence of trust-warranting properties, so the trustor needs certain knowledge to be able to decode them. In e-commerce, trust symbols are all the mechanisms that aim to directly signal a merchant's trustworthiness to potential customers (e.g. trust seals issued by certification authorities - CAs).
- *Symptoms*: These are trust signals that were not specifically designed to signal trust-warranting properties; they are given off as by-products when honest actors go about their business, at little or no cost to honest actors. But if an attacker were to try and mimic these, it would be at great cost. In an e-commerce setting, symptoms are all the website properties and information consumers draw on to assess the trustworthiness of a merchant (e.g. well-known brand name and reputation amongst friends and relatives).

The opportunistic attempt of non-trustworthy actors to appear trustworthy by emitting manipulated signals (symptoms or symbols) is defined as *mimicry* [31]. Mimicry will occur only if emitting the signal required to appear trustworthy comes to a lower cost to untrustworthy actors than the potential benefit of doing so [2]. Symbols are much easier to mimic than symptoms, as attackers can simply place those in their websites at minimal costs, while symptoms require more effort, which can outweigh any potential benefits for an attacker.

2.4 Trust Symptoms - User Trust Assessment Heuristics

Consumers use a wide range of symptoms to assess the trustworthiness of a merchant. Those are mostly based on self-developed heuristics¹, a number of which have been reported by past research:

- Perceived professionalism and reputation of a company, e.g. well-known branding [13],[22],[28]
- Ability of the merchant to fulfill – usually revealed by positive user reviews [22],[32]
- Relationship with other known entities e.g. other well-known merchants [33]
- Willingness to customize products and services [13]
- Usefulness and ease of use of the website [21],[28],[32]
- Perceived security control e.g. providing reassurances in case of fraud [32]

The major drawback of all the aforementioned reports is that the proposed heuristics are either not supported by experimental verification ([21],[28],[32]), or were based on experimental designs that did not accurately capture the complete picture of the actual trust development process:

1) Experiments in [22] and [13] were based on pre-defined hypotheses by researchers (i.e. “*Would you trust this website if it had this property?*”), which results to reporting only the effect the properties they targeted have on a website's perceived legitimacy (e.g. [22] tested for the effect of size and reputation on trust development, but explicitly presented size and reputation information for each website to the participants). This setup also hinders the ecological validity of the results as it explicitly draws the participants' attention to website properties they may have failed to notice by themselves.

2) In [33] the focus of the experiment conducted was to test for the effectiveness of an anti-phishing indicator; trust development factors were identified afterwards, based on analysis of self-reports by participants in post-experiment interviews.

¹Merriam Webster heuristic definition: “involving or serving as an aid to learning, discovery, or problem-solving by experimental and especially trial-and-error methods”

2.5 Trust seals

Trust seals are extra-legal, symbol-based trust-signaling mechanisms, introduced to provide trustworthiness information on a merchant to potential customers from Trusted Third Parties (TTPs) – Certification Authorities (CA). They are logos added in websites to signal that a certified organization (TTP) has granted the right to use those, based on some rules of conduct (e.g. reliability as a merchant, correct private data handling or website security). They are used to facilitate trust-building in online commerce environments [14], decreasing the perceived transaction-related risk by consumers, thus increasing their willingness to engage in it [12].

The purpose of trust signaling mechanisms is similar to the way risk communication is used to advise the general public on issues of public concern: they act as advisors to consumers on the risk they accept by engaging in a transaction. Twyman et al. [34] classify trust in advisors in two major categories:

- *Trust in motives*: Consumers identify the similarity of values between them and the TTP: they both benefit from successful interactions and it is of interest to the TTP to provide them with reliable information.
- *Trust in competence*: Consumers have received reliable information from a specific TTP in the past so they are more likely to trust the information they receive from it. This explains past research reports that consumers are more likely to buy from an unknown website that bears a trust seal than one which does not [11].

Both forms of trust can be destroyed, if they are manipulated by untrustworthy actors. Attackers can easily add trust seals to their websites (*mimicry* – [31]) and negative experiences with trust seals will result to the corresponding signals losing significance [2]. This can undermine trust in the competence of the seals, and the trust certification approach in general.

Research opinion on the effectiveness of trust seals is divided: some researchers report that trust seals help to improve consumer purchasing decisions [11],[15-17], but others report that they fail to do so [13][18]. The problem with all the previous research is that they drew participants' attention to the presence of trust seals and explicitly asked if they influenced their decision to trust a website. Testing if consumers take trust seals into account if they notice their presence, is a valid question, but accounts only for a sub-set of the decision-making process. Previous experiments also did not test whether trust seals lead to *correct trust decisions*; if a trust seal is present, consumers are more likely to buy from it - but that might include buying from fraudulent sites that carry mimicked seals. We need trust signals that help consumers to make better decisions, not just manipulate their trust perceptions.

2.6 Public Awareness Campaigns

A number of awareness campaigns have been set up by governments and commercial organizations to inform consumers on the potential dangers they may face while shopping online [35-38]. They provide a range of advice on what consumers can use to protect themselves:

- Make sure they have antivirus, firewall and anti-spyware software installed and keep operating system and browser up to date.
- Check merchants out before first time purchases. Locate contact details and whether refunds are provided in case things go wrong.
- Verify the website's legitimacy using *https://* indicators and closed padlocks; also never make purchases through unsecured wireless networks.
- Only provide a website with the information required to complete the transaction.
- Check who the website is registered to and how long has it been registered.
- Check for website reviews on the Internet.

Whilst we would not argue that this advice is wrong, it ignores some key factors that drive consumers' behavior in these situations:

- When presented with a 'good deal', consumers may be tempted to accept a higher risk in order to reap the potential benefits. This risk propensity can be leveraged by attackers emphasizing the "limited offer duration" (*time principle* – [29]), putting time pressure on consumers to quickly seize the deal, otherwise lose it.

- Providing consumers with widely varying advice ([36] mentions 9 different website properties that require 3-4 verification checks each) causes confusion. Faced with too much information, consumers try to reduce it to a manageable level, but the process of selecting factors is haphazard [39]. Finally, consumers can follow the advice, and still fall for a scam, as many scam websites are well-designed and resemble legitimate ones, in their attempt to appear trustworthy. When this happens, consumers' trust in the competence of the advisors is undermined, making it less likely that they will pay attention to advice from that source next time [34].

2.7 Summary

The continuing high level of online scams suggests that existing security measures and advice are not working. Even though trust seals are widely used, there are conflicting reports on their effectiveness. Past research lacked ecological validity: to accurately capture the trust development process, consumers need to be presented with trust seals in the same way as they would in a home setting, without the experimenters drawing their attention to any specific trust development factors.

3 Experiment

3.1 Aim

The aim of our experiment was to evaluate the effectiveness of trust seals in a realistic setting (see 3.2) that would improve on the validity and applicability of our findings. We designed a study to test the following hypothesis:

H1: Website ratings will increase when participants notice the presence of trust seals

We used eye-tracking and screen recordings to guide a set of post-experiment interviews, where we questioned our participants on their eye-gaze fixations during the experiment. We analyzed interview data using qualitative methods, aiming to:

- Capture the participants' *perception of the meaning* of trust seals.
- Identify other elements they used to assess the trustworthiness of a website.

3.2 Method

We asked participants to browse through six websites that sell tickets for a music festival in London (*Wireless Festival in Hyde Park*) and asked them to rate each website based on how likely they were to buy from it. We chose online ticket sales for our study because they represent a large and constantly growing number of online scams: the UK National Fraud Authority reported a number of half a million ticket scam victims in the UK in 2010, each losing an average of £80 [40].

Apparatus and Materials. A pre-experiment questionnaire was designed to identify participant demographics, computer experience, online shopping habits and past experience with internet scams. During the experiment, screen and eye-gaze recordings were taken using a Tobii X50 eye tracker and Tobii Studio 2.0.4 software. The experiment took place in a usability laboratory on a computer running Windows XP and websites were displayed using Mozilla Firefox 3.5 web browser. Post-experiment interviews with participants were audio recorded.

Websites. Using a search engine, 6 websites selling tickets for the event were identified, and downloaded locally using the HTTPRack free website copier tool (<http://www.httrack.com>). Three of the websites had a trust seal positioned on the main page and other parts of the ticket selection process. To test for the effectiveness of trust seals two different conditions were created (Table 1):

- *Original:* All websites were used unmodified in the experiment.
- *Modified:* Trust seals were removed from the websites that originally carried those and a fake trust symbol was placed in the other three websites (that originally did not have them) as a plain image, without any links to verify its authenticity. The fake symbol was positioned in easy to spot positions in the websites.

Table 1: The conditions assigned to the websites used in the experiment

Website name	Original	Modified
<i>www.eventim.co.uk</i>	No trust seal	Trust seal
<i>www.getmein.com</i>	Trust seal	No trust seal
<i>www.gigantic.com</i>	No trust seal	Trust seal
<i>www.hmvtickets.com</i>	Trust seal	No trust seal
<i>www.seetickets.com</i>	No trust seal	Trust seal
<i>www.skiddle.com</i>	Trust seal	No trust seal

The local copies of the websites were setup on a university server and the DNS mapping was modified so that the participants could see the real website URL (e.g. *www.eventim.co.uk*) in the address bar.

Participants. Participants were recruited through the university’s psychology department subject pool. They had to be over 18, use online shopping regularly, and be available to visit the lab for a 1-hour session. They all received payment of £12 for their time. The university’s ethics procedures on experiments involving human participants were followed. (No application to the Research Ethics Committee was required since our participants were not identifiable, no personal information was kept after the experiment, and there was no deception involved).

62 participants took part in the study, but data from two had to be discarded due to lack of accurate eye-tracking recordings. Of the remaining 60:

- 36 (60%) were female and 24 (40%) male.
- Their average age was 24 years (Standard Deviation = 4.9).
- They had an average computer experience of 12 years (SD = 3.3).
- They browse the Internet daily for 4.9 hours (SD = 2.88).
- They receive 18 (SD = 14.7) emails per day.
- 51 (85%) of them have checked their account balance online.
- 50 (83%) had transferred money to other people’s accounts using online banking services.
- 11 (18%) had configured a firewall in the past.
- 21 (35%) had designed a website.
- 15 (25%) had registered a domain name.
- 17 (28%) had been victims of an online scam, or knew someone that has been.

Procedure. A between-subjects design was chosen to prevent habituation effects. Participants were equally divided between the two conditions. The websites were pre-opened in six browser tabs in randomized order, and participants had 5 minutes to browse through those. After signing the consent form and completing a screening questionnaire, they were presented with the experiment scenario: “*You want to attend the Wireless Festival 2011 in Hyde Park. You have used a search engine to find six websites that claim to sell tickets. Friends have warned you that ticket festivals sell out very quickly so you only have five minutes to look at the websites. You can browse through the websites with no limitations. Warnings will be given to you when two and one minutes are left. After the end of the 5-minute period you will be asked to indicate how much you trust each of the websites presented to you. To do so you need to assign a grade between -2 and +2 (-2,-1,0,1,2) to each website with -2 being the lowest and 2 the highest*”. After reading the scenario, they were asked to confirm they understood how the rating grades reflect their level of trust for each website.

During the 5-minute browsing period the experimenter was present, but participants were told they could not ask questions during this part of the session. They were allowed to distribute their browsing period in any way they wanted across websites, so they carry out all the checks as if they were shopping on their own computers (e.g. check delivery policy, FAQs etc.) and when they had enough information to make a decision proceed to the next website. Participants were not prohibited from using external sources (i.e. other websites) to check for a website’s reputation, but none attempted to do so during the experiment.

After participants rated the websites, there was a de-briefing session: the eye-tracking recording of their browsing period was replayed to them, and questions about their behavior asked, based on their eye-gaze fixations. When the recordings showed a fixation on any visual element of a website (e.g. reading through the text on a page), participants were asked to explain how each of those elements affected the trust rating they assigned to each website. Participants were then pointed to the trust seals in the sites and were asked to explain what they signal to them, and whether they knew how to verify their authenticity. This aimed to provide data that could be used to identify whether consumers perceive the meaning of trust seals correctly, which is important if trust seals affect their decisions, as incorrect understanding can result to misplaced trust. The questions on trust seals were asked at the final part of the interview, to avoid drawing the participants’ attention to their presence. The interviews were audio recorded and analyzed after the experiment, using a Grounded Theory analysis combining open, axial and selective coding procedures [41].

4 Results

4.1 Effectiveness of Trust Seals

The analysis of eye-tracking data revealed that only 12 (20%) participants noticed all three trust seals they encountered during their browsing session (Table 2), and more than a third did not notice any of them (23 – 38%). We tested our *H1* hypothesis by comparing the ratings participants assigned to a website when they noticed the presence of the trust seal in it against the ratings when the trust seal was not noticed or was not present. This revealed a significant tendency ($t(5) = 3.3786, p = 0.0099$) to rate websites higher when participants noticed a trust seal on a website (Table 3).

Table 2: Number of trust seals noticed by participants

No of seals noticed	No of participants
0	23
1	12
2	13
3	12

Table 3: The assigned ratings on websites when trust seals were present

Website name	Number of participants who noticed	Rating when noticed seals	Rating when not noticed seals or seals not present
<i>www.eventim.co.uk</i>	18	0.94	0.00
<i>www.getmein.com</i>	15	0.73	-0.04
<i>www.gigantic.com</i>	14	0.64	-0.11
<i>www.hmvtickets.com</i>	8	1.25	1.04
<i>www.seetickets.com</i>	11	0.27	0.37
<i>www.skiddle.com</i>	5	0.40	-0.40

After participants’ attention was drawn to the presence of trust seals, we asked what they signified, and received a variety of responses - all incorrect (see Table 4).

Table 4: The responses participants gave on the meaning of trust symbols

Comment	No of participants
Seals mean a website is safe for Credit Card details	18 (30%)
Completely ignore what trust seals are and what they mean	15 (25%)
They know that trust seals can be spoofed	11 (18%)
Payment method symbols mean the website is verified by the payment method company (e.g. VISA, PayPal, MasterCard)	11 (18%)
Seals provide confirmation that website is genuine (could not explain why)	10 (17%)
Authority exists that grants rights to use the seal to trustworthy merchants, punishing the misuse of those	9 (15%)
Seen some trust symbols in websites they use often, assumed that their presence in a website automatically signifies its legitimacy	7 (12%)
Seals are meaningless as they could be copied by anyone	6 (10%)
Seals mean a website has no viruses	1 (2%)

4.2 Factors Affecting Trust in Websites

The Grounded Theory analysis of the interview data revealed a number of factors other than trust seals that affected the participants' trust development decisions. These factors can be classified in two major categories: Those that affected the *perceived professionalism* of the company and those that affected the *perceived competence* of it as an online merchant.

1. **Trustee's professionalism.** Participants attempted to assess the professionalism of the company running a website, which they reported as a combination of many different factors (Table 5):

Table 5: The factors affecting the perceived professionalism of a website

Comment	No of participants
Perceived amount of effort invested in a website - indicated by factors like aesthetically pleasing design, well-formed layout.	45 (75%)
Presence of company information e.g. physical location, contact details etc	43 (72%)
Variety of products available	23 (38%)
Inclusion of Terms and Conditions/Privacy Policy	16 (27%)
Large amount of information on event of interest (opening times, venue information etc), good presentation of it with rich media (e.g. maps and pictures)	14 (23%)
Ease of use, self-explanatory labeling - to aid navigation around the website	8 (13%)
Well-formed URL - participants argued that scam websites have long, non-meaningful URLs	6 (10%)

2. **Trustee's competence.** Participants attempted to assess a website's competence by looking for a number of different website properties (Table 6):

Table 6: The factors affecting the perceived competence of a website

Comment	No of participants
Indicators of past trustworthy behavior - Name and reputation of a company were the major factors participants used to assess this: positive expectation about a merchant’s behavior was formed if participants recognized a company’s name or had previous experience with it (online or in the real-world).	45 (75%)
Trust transfer - Inclusion of other recognizable entities affected the decisions of participants e.g. claims by a website that they are subsidiary of <i>Ticketmaster</i> (UK’s biggest ticket merchant), advertisements of known companies or presence of a charity logo together with claims that part of the profits is donated to them.	30 (50%)
Social Networking links - Believed that they could find information on the merchant’s past behavior by following those links	28 (47%)
Assurances provided – The ticket purchases and financial details are safe and that tickets will be sent via secure postage	19 (32%)
User reviews – Present inside the website (did not check for off-site reviews)	16 (27%)

5 Discussion

5.1 Trust seals are not effective

Our findings suggest that trust seals do not improve on consumers’ ability to make accurate trust assessment of websites: despite a significantly increased rating amongst participants who noticed trust seals, only 20% noticed those on all 3 websites they encountered, and over a third of participants (38%) did not notice any of the 3. It is reasonable to suspect that the same is true for consumers - unless their attention is specifically drawn to the presence of a seal on a site. Our participants also had significant misconceptions about what the seals stood for (see Table 4). Those participants who noticed trust seals during the experiment interpreted the mere presence of those as proof of a website’s competence (hence the statistical significance in Table 3), and felt no need to check that they were genuine.

We also observed a trust “*spillover effect*”. Early research on trust seal effectiveness [42] pointed out that their presence does not legitimize a website, but consumers are still not aware of this: seven participants (12%) had previously seen some trust seals in websites they use, and incorrectly assumed that these mean the website is legitimate. This misconception makes consumers highly vulnerable to mimicry attempts.

Another problem with trust seals can be attributed to bad practice by merchants. In one of the websites we presented to participants, the trust seal present was a plain image, instead of linking to the verification pages provided by the seal issuers (e.g. whenever a VeriSign trust seal is present, it should be a clickable link, bringing up a verification page with the details of the company to which the website was registered [43]). If legitimate merchants implement trust seals incorrectly, the task of identifying mimicry attacks becomes almost impossible for consumers (even though our participants did not attempt to verify the seals).

The large number of different trust seals used is a further source of confusion for consumers, and thus undermines the effectiveness of trust seals. The *www.truste.com* [44] website lists 9 different certifications covering Privacy, Security (2 for SSL encryption, 2 for malware and vulnerability scans), Reputation and Reliability (which can be either review-based or granted by another authority). Such a complex system does not help consumers trying to detect online fraud - and how many consumers know what SSL encryption is, or what risks malware presents to them?

The creators of trust seals and website owners who use them expect consumers to search for trust seals, check their authenticity, and understand what protection their offer. Based on our results, we argue that these expectations are unrealistic. Usable security researchers have long argued that security is not the primary concern of people using

computer systems [45],[46]: the need to be careful about scams is a minor consideration in the context of the consumer's main activity – to find and buy something they want. Expecting consumers to interrupt this activity to find and check trust seals before every purchase is expecting too much. Like most security mechanisms, the effort involved is just too high for ordinary consumers [47] – so they either ignore them altogether, or associate those with a simple, but incorrect meaning. In both cases, consumers are left vulnerable.

5.2 Trust assessment heuristics and consumer awareness

The trust assessment heuristics we present in Section 4 partly confirm past research findings [13],[21],[22],[28],[32],[33], but the most worrying observation is that none of our participants attempted to verify the authenticity of the signals they used. This means that even simple mimicry attacks can succeed: attackers copy genuine websites, register well-formed web addresses and use search engine advertisements or phishing emails to direct consumers to them [7],[48]. None of the websites used in our experiment included a way to verify claimed affiliations (social network links, charity organizations or advertisements), which means that even if consumers were prepared to check for their authenticity, there would be little they could do. An example of how the identified heuristics can be manipulated by attackers are account takeovers, reported by eBay as a major source of threat for their customers [49], as consumers blindly trusting reputable retailers are left vulnerable against those attacks.

Our observations demonstrate that the advice given by awareness campaigns is not effective, either: No participant checked for *https://* indicators, padlocks in the address bar or who the website owner is. The advice “*Only provide a website with the information required to complete the transaction*” also seems ill-posed: what is more sensitive than the credit card information required to complete a transaction? Consumers transfer trust perceptions from physical world settings (e.g. the reputability of a brand name or claimed affiliations with other well-known organizations), unaware that these are easily and cheaply mimicked in the online world. This leaves them vulnerable to the techniques attackers use, like including well-known names in their website without any proof for their affiliations, and which awareness campaigns fail to address effectively: despite telling consumers to look for specific trust signals, they fail to equip them with the skills required to verify their authenticity, doing nothing to protect them from any potential mimicry attempts.

6 Conclusions

Our results demonstrate that trust seals do not effectively support consumers making decisions about websites. A significant part of consumers does not notice them, and most of those who do, do not understand what protection they offer and how to verify their authenticity. We thus argue that trust seals may currently do more harm than good, because they leave consumers vulnerable against even the simplest attacks (e.g. inclusion of fake trust seals in websites). To overcome these problems, a significant shift is required from the way trust signaling mechanisms are used today. Technology needs to be used to aid correct trust placement by automatically performing any verification required, alerting consumers when potential risks are identified, aiding their accurate assessment of the dangers they may face when they need to make trust-related decisions and reducing the potential of being victimized by online scams.

6.1 What needs to be done

Use automatic verification mechanisms. More radical measures are required to reduce the potential of successful mimicry attempts. Mechanisms that automatically verify a seal's authenticity need to be developed, which will alert consumers when seal misuse is detected. The backbone of this technology already exists: The SOLID authentication tool (developed by the UK firm First Cyber Security) “*gives the owner of a logo, trademark or certificate the ability to authenticate its use on other websites*”, using a *Secure On Line ID 3rd Party Validation* mechanism, which identifies unauthorized use of symbols registered by their original owners [50]. This, or other similar systems, can be used as the basis of a larger implementation, developed in collaboration with web browser creators, which will automatically alert crime prevention authorities and Internet Service Providers (ISPs) when scam websites are detected, who can then act to block traffic to those and take them offline. A widely-adopted automatic verification approach can also be used to provide shoppers with merchant information, like registration details of the company owning the domain name, contact details, where the product will be shipped from etc. – eliminating the need to find that information by looking around various websites on the Internet. The technology to implement this also exists - organizations like VeriSign already provide information on the owner of a website when consumers click on the VeriSign trust logo, but currently require consumers to notice the presence of that logo to do so.

The alerts an automatic verification mechanism presents to consumers should use meaningful messages, explaining what the identified problem is and how to protect themselves. Those messages should appear as active warnings, which are proven to be more effective than asking the consumers to stop and search for security indicators [51]. To avoid habituation issues, consumers' should only be interrupted when seal misuse is detected and presented with a short and clear warning that the website they are browsing is using unauthorized symbols. When no problem exists a passive window can be present in the consumer's browser providing information on the merchant. This will minimize the cognitive load imposed on them when they attempt to assess a merchant's trustworthiness and can result to more accurate trust-placement decisions. The success of any attempt to implement a mechanism like this requires the involvement of all interested stakeholders (merchants, certification authorities, ISPs, crime prevention authorities) and a good implementation can significantly improve on the public perception of e-commerce as a safe and trustworthy service.

Re-focus awareness campaigns. Research has already reported that current security awareness campaigns are not well-aligned with actual consumer behavior in online environments [33]; future campaigns should focus on widely-held misconceptions. Automatic website verification can significantly reduce the amount of information that needs to be communicated to consumers, and the effort they have to make to check the authenticity of a site. They only need to be made aware of the fact that they may be targeted by scams (e.g. a website you access may be fraudulent and you may receive nothing for the money you pay) and what they need to do to protect themselves (e.g. make sure your browser is up to date). This can result to a significant decrease in the confusion amongst them and aid safer decisions when shopping online.

6.2 Research limitations

The study aimed to create a scenario that would closely resemble the conditions under which consumers shop online: The need for them to accurately assess the trustworthiness of a merchant to avoid being victimized. A potential limitation that may have affected the ecological validity of our quantitative results is the fact that participants did not risk losing any money or having any personal details compromised, which would be the consequences of incorrect trust decisions while shopping online. This could have an effect on the ratings participants assigned to websites, but testing for this was not possible due to time and resource limitations. Despite that, the main issues we raise on the ineffectiveness of trust seals (failure to notice those and misunderstanding of their purpose) and the failure of awareness campaigns are well supported by the qualitative analysis of the interview data, where the identified misconceptions on trust seal meaning and trust development heuristics are unlikely to have been affected by this limitation.

7 References

1. The UK Cards Association, [http://www.financialfraudaction.org.uk/cms/assets/1/be%20card%20smart%20release%20final%20-%2024%20nov%2011%20\(nfa\).pdf](http://www.financialfraudaction.org.uk/cms/assets/1/be%20card%20smart%20release%20final%20-%2024%20nov%2011%20(nfa).pdf)
2. Riegelsberger, J., Sasse, M. A., and McCarthy, J. D. The mechanics of trust: a framework for research and design. In *International Journal of Human-Computer Studies*, 62(3), pages 381-422 (2005).
3. Financial Fraud Action UK, <http://www.financialfraudaction.org.uk/cms/assets/1/fraud%20figures%20release%2010%20mar%2010.pdf>
4. DirectGov UK, http://www.direct.gov.uk/en/Governmentcitizensandrights/Consumerrights/Howtocomplainaboutgoodsandservices/DG_196229
5. Publicservice.co.uk, http://www.publicservice.co.uk/news_story.asp?id=18293
6. Mail Online, <http://www.dailymail.co.uk/femail/article-2073344/Will-fall-Santa-frauds-Britain-flooded-designer-Christmas-gifts-actually-dangerous-fakes.html?ito=feeds-newsxml>
7. Retail Digital, http://www.retail-digital.com/consumer_trends/top-retail-scams
8. BBC News, <http://news.bbc.co.uk/2/hi/8392600.stm>
9. UK Cards Association, http://www.theukcardsassociation.org.uk/media_centre/press_releases_new/-/page/1323/
10. <http://www.internetretailer.com/2011/01/18/fraud-losses-fall>

11. Hu, X.R., Lin, Z.X., Zhang, H. Myth or reality: effect of trust promoting seals in electronic markets, *Proceeding of the Eleventh Annual Workshop on Information Technologies and Systems (WITS)*, New Orleans, Louisiana, 65–70 (2001).
12. Resnick, P., Zeckhauser, R., Friedman, E., Kuwabara, K. Reputation systems: facilitating trust in internet interactions, *Communications of the ACM* 43 (12) 45– 48 (2000).
13. Kim, D., Ferrin, D., and Rao, H. A trust-based consumer decision-making model in electronic commerce: The role of trust, perceived risk, and their antecedents. In *Decision Support Systems*, 44(2), pages 544-564 (2008).
14. Ba, S., Whinston, A.B, and Zhang, H. Building trust in online auction markets through an economic incentive mechanism. *Decis. Support Syst.* 35, 3 (June 2003), 273-286 (2003).
15. Kimery, K. M. and McCard, M., Third-party assurances: mapping the road to trust in e-retailing, *Journal of Information Technology Theory and Application*, Vol. 4 No. 2, pp. 63-82 (2002).
16. Rifon, N. J., LaRose, R. and Choi, S. M., Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures. *Journal of Consumer Affairs*, 39: 339–362 (2005).
17. Bakos, J. Y., C. Dellarocas. Cooperation without enforcement? A comparative analysis of litigation and online reputation as quality assurance mechanisms. *Proc. Internat. Conf. Inform. Systems*, Barcelona, Spain, 127–141 (2002).
18. Peterson, D., Meinert, D., Criswell, J. II, Crossland, M. Consumer trust: privacy policies and third-party seals, *Journal of Small Business and Enterprise Development*, Vol. 14 Iss: 4, pp.654 – 669 (2007).
19. Wash, R. Folk models of home computer security. In *SOUPS 2010: Proceedings of the 6th Symposium on Usable Privacy and Security*, SOUPS '10, pages 1-16, New York, NY, USA. ACM, (2010).
20. Tan, Y. and Thoen, W. Toward a Generic Model of Trust for Electronic Commerce. In *International Journal of Electronics Commerce*, 5, pages 61-74 (2000).
21. Nielsen, J., Molich, R., Snyder, S., and Farrell, C. *E-Commerce User Experience: Trust*. Fremont, CA: Nielsen Norman Group (2000).
22. Jarvenpaa, S., Tractinsky, N., Vitale, M. Consumer trust in an internet store, *Information Technology and Management* 1 (1–2) 45– 71 (2000).
23. Hoffman, D. L. Building consumer trust online. In *Communications of the ACM*, 42 (4), pages 80-85(1999).
24. Bolton, G.E., Katok, E., and Ockenfels, A. How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation. *Manage. Sci.* 50, 11, 1587-1602 (2004).
25. Ratnasingam, P. and Pavlou, P. A. Technology Trust in Internet-Based Interorganizational Electronic Commerce. *Journal of Electronic Commerce in Organizations* 1(1), 17-41 (2004).
26. Akerlof, G. The market for lemons: quality uncertainty and the market mechanism, *Quarterly Journal of Economics* 84 (3) 488–500(1970).
27. Handy, C., “Trust and the Virtual Organization,” *Harvard Business Review*, Vol. 73, No. 3, pp. 40-50 (1995).
28. Shneiderman, B. Designing trust into online experiences. *Communications of the ACM* 43(12), 57-59 (2000).
29. Stajano, F. and Wilson, P. Understanding scam victims: seven principles for systems security. In *Communications of the ACM*, 54 (3), pages 70-75, New York, NY, USA, (2011).
30. Riegelsberger, J., Sasse, M.A., and McCarthy, J.D. 2003. The researcher's dilemma: evaluating trust in computer-mediated communication. *Int. J. Hum.-Comput. Stud.* 58, 6, 759-781(2003).
31. Bacharach, M. and Gambetta, D. Trust as Type Detection. In: Castelfranchi, C. and Tan, Y. *Trust and Deception in Virtual Societies*. Kluwer: Dordrecht, 1-26 (2001).
32. Egger, F. N. Affective Design of E-Commerce User Interfaces: How to maximise perceived trustworthiness. In *Proceedings of International Conference on Affective Human Factors Design*, pages 317-324 (2001).
33. Kirlappos, I., and Sasse, M.A. Security education against phishing: A modest proposal for a major re-think. *IEEE Security and Privacy*, 99(Preliminary), (2011)
34. Twyman, M., Harvey, N. and Harries, C. Trust in motives, trust in competence: Separate factors determining the effectiveness of risk communication. In *Judgment and Decision Making*, 3, pages 111-120 (2008).

35. Google Good to Know, <http://www.google.co.uk/goodtoknow/online-safety/shopping/>
36. Stay Safe Online, <http://www.staysafeonline.org/>
37. DirectGov UK, http://www.direct.gov.uk/en/N11/Newsroom/DG_180506
38. DirectGov UK, http://www.direct.gov.uk/en/Governmentcitizensandrights/Consumerrights/Protectyourselffromscams/DG_195960
39. Harvey, N., Harries, C. & Fischer, I. Using advice and assessing its quality. *Organizational Behavior and Human Decision Processes*, 81, 252-273 (2000).
40. Action Fraud UK, <http://www.actionfraud.org.uk/festival-lovers-must-beware-of-ticketing-fraud-mar11>
41. Strauss, A., Corbin, J.: *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. SAGE publications, London (1998)
42. Edelman, B. Adverse selection in online "trust" certifications. In *Proceedings of the 11th International Conference on Electronic Commerce (ICEC '09)*. ACM, New York, NY, USA, 205-212 (2009).
43. Verisign: Report Seal Misuse, <https://www.verisign.com/support/contact/seal-abuse/index.html>
44. TRUSTe, <http://www.truste.com/consumer-privacy/comparing-web-privacy-seals>
45. Beauteument, A., Sasse, M. A., and Wonham, M. The compliance budget: managing security behaviour in organisations. In *NSPW '08: Proceedings of the 2008 workshop on New security paradigms*, pages 47-58 (2008).
46. Herley, C. So long, and no thanks for the externalities: The rational rejection of Security advice by users. In *Proceedings of the New Security Paradigms Workshop 2009*, pages 133-144 (2009).
47. Whitten, A., and Tygar, J.D. Why Johnny can't encrypt: a usability evaluation of PGP 5.0. In *Proceedings of the 8th conference on USENIX Security Symposium - Volume 8 (SSYM'99)*, Vol. 8. USENIX Association, Berkeley, CA, USA, 14-14 (1999).
48. GHD Repair, http://www.ghd-repair.co.uk/fake_ghds.html
49. Ebay, <http://pages.ebay.com/help/account/securing-account.html>
50. SOLID Authentication, <https://www.solidauthentication.com>
51. Wu, M., Miller, R.C., and Garfinkel, S.L. Do security toolbars actually prevent phishing attacks?. In *Proceedings of the SIGCHI conference on Human Factors in computing systems (CHI '06)*, Rebecca Grinter, Thomas Rodden, Paul Aoki, Ed Cutrell, Robin Jeffries, and Gary Olson (Eds.). ACM, New York, NY, USA, 601-610 (2006).