Computers in Human Behavior 41 (2014) 71–79

Contents lists available at ScienceDirect

Computers in Human Behavior

journal homepage: www.elsevier.com/locate/comphumbeh

Reading this may harm your computer: The psychology of malware warnings

David Modic *, Ross Anderson

University of Cambridge Computer Laboratory, JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom

ARTICLE INFO

Article history:
Available online 30 September 2014

Keywords:
Malware
Persuasion
Human computer interaction
Psychology

ABSTRACT

Internet users face large numbers of security warnings, which they mostly ignore. To improve risk communication, warnings must be fewer but better. We report an experiment on whether compliance can be increased by using some of the social-psychological techniques the scammers themselves use, namely appeal to authority, social compliance, concrete threats and vague threats. We also investigated whether users turned off browser malware warnings (or would have, had they known how).

1. Introduction

In life, as on the Internet, most of us are satisficers – we tend to favour actions and make decisions that are good enough, rather than optimal (Simon, 1956). As an energy-saving technique, this has benefits, but also drawbacks. When it comes to protecting oneself online, Akhawe and Felt (2013) and Herley (2010) have shown that Internet users work hard to ignore warnings and security notices. Existing theories and empirical work in criminology suggest this might be a problem. Situational crime prevention shows that offenders are more likely to take advantage of an environment that appears target-rich (cf. Felson & Clarke, 1998), while routine activity theory (RAT; Cohen & Felson, 1979) analyses crime incidence in terms of a motivated offender, a suitable target and an opportunity. However, there is comparatively little research on the causal link between ignoring warnings and being defrauded. One plausible explanation is that those who ignore the warnings might believe themselves to be less vulnerable because they might have less money to lose or are confident in their ability to resist scams. In reality, lack of funds does not pose a hurdle for determined scammers, who have been known to push prospective victims into taking loans (e.g. in investment scams; Stevenson, 2000) or entangle them in money laundering schemes (Zuckoff, 2005). Overconfidence in one's ability to resist fraud has also been shown to increase the likelihood of being scammed (Camerer & Lovallo, 1999; Fischer, Lea, & Evans, 2013).

While computer users are more likely to follow an inconvenient procedure if they are explicitly told it is for security purposes (Egelman et al., 2010), the daily exposure to an overwhelming amount of warnings remains an issue. This makes it hard for users to sort the real threats from the many trivial ones and the even greater number of false alarms (Bravo-Lillo et al., 2013). Users are willing to expend only a certain amount of effort and time on security concerns: that is, their *compliance budget* (Beautement, Sasse, & Wonham, 2008) is a limited resource. In brief, users would prefer to ignore warnings, but if that is hard enough they will comply with some of them, up to a point.

Thus there is a need for fewer but more effective of malware warnings, particularly in browsers. Earlier research tended to focus on the presentation of warnings; for example, passive warnings (that require no user action) tend to be almost universally ignored. Egelman, Cranor, and Hong (2008) found that active warnings helped deter 79% of their participants from visiting a potentially harmful website. Later research has moved towards the positioning of the dialogues, the amount of text, the length of the message and the amount of technical detail (Bauer, Bravo-Lillo, Cranor, & Fragkaki, 2013). Another recent approach has been to manipulate the content of security warnings (e.g. malware warnings; Egelman & Schechter, 2013; and SSL warnings; Sunshine, Egelman, Almuhimedi, Atri, & Cranor, 2009). The wording in warnings in such studies generally appears to be based on trial and error rather than on established psychological theories of communication or persuasion. In the present paper, we based our warnings on some of the social psychological factors that have been shown to be effective when used by scammers (Modic, 2013; Modic & Lea, 2013). Those factors which play a role in increasing potential victims' compliance with fraudulent requests, will also prove effective in warnings.

* Corresponding author. Tel.: +44 1223 767014; fax: +44 1223 334678.
E-mail addresses: david.modic@cl.cam.ac.uk (D. Modic), ross.anderson@cl.cam.ac.uk (R. Anderson).

http://dx.doi.org/10.1016/j.chb.2014.09.014
0747-5632/

1.1. Social psychological factors informing compliance with warnings

Previous research in the social psychology of persuasion (cf. Cialdini, 2001) has uncovered several factors that can influence our decision making abilities and increase compliance. Fischer, Lea, and Evans (2009) have shown in their report for UK Office of Fair Trade that social psychological mechanisms of persuasion such as influence of Authority and Social influence increase compliance with postal fraud. In addition Modic and Lea (2013) have developed a scale of Susceptibility to Persuasion that has been validated on victims of Internet fraud and has shown that same mechanisms we are using in this research, when used by scammers, are effective in reducing resistance of potential victims. The mechanisms effectively used by scammers on potential victims, would be likely just as effective when used by browser designers to increase scam resistance of potential victims.

1.1.1. Influence of Authority

Individuals are likely to respond to requests from authority figures across a range of cognate domains. For example, Titus and Dover (2001) show scammers using authority to elicit compliance with building inspector frauds and other scams. In postal fraud, Fischer et al. (2013) showed that potential victims were more likely to comply with requests from scammers with ostensible formal authority. Tyler and Degoey (1995) found that individuals are more willing to show self-restraint in social dilemmas when they perceive the requesters to be fair and honest. Trust in authority figures increases their influence. Murphy (2004) has shown that individuals are more likely to pay taxes when they trust the tax authorities. We thus hypothesize (H1) that warnings will be more effective when potential victims believe that they come from a trusted authority.

1.1.2. Social influence

Human susceptibility to group pressure or social influence is well supported empirically, from early line experiments by Asch (1956) to newer work: for example, Markus and Kitayama (1991) showed that individuals in different cultures construct their self-worth through comparison with other in-group members. Criminologists have found that individuals are more likely to comply with formal norms if they believe other members of their community also comply with them, while on the other hand visible disorder is a self-reinforcing cue for criminal activity (Kahan, 1997). Consumers susceptible to social influence may buy products a seller favours even if their preferences are different (Bearden, Netemeyer, & Teel, 1989). A malware warning exposing a potential threat to an individual's in-group might thus work across cultural contexts. We hypothesize that (H2) a warning constructed to solicit compliance with in-group norms would increase the likelihood of visiting a potentially harmful site even against the individuals' initial wishes.

1.1.3. Risk preferences

Individuals in general tend to act irrationally under risky conditions (Kahneman & Tversky, 1979; Munro, 2009; Rubinstein, 1997). They are willing to forgo privacy concerns to feel safer (Jagatic, Johnson, Jakobsson, & Menczer, 2007). And a study by Titus and Dover (2001) showed that repeated communication that varied the perceived risk of an unfavourable outcome increased compliance by potential victims. We hypothesize (H3) that straight talk would be effective in warnings; a concrete threat with clearly describes possible negative outcomes should increase compliance compared to a vague one.

1.2. The decision to keep malware warnings turned on

There is an underlying assumption in security-warnings research that most users will keep the warning mechanisms on their default setting (i.e. turned on). There is tangential empirical support for this claim. Spool (2011) has shown that nine out of ten individuals keep all the default settings in a popular text-processing package. The preference for things to stay the same (the status quo bias; see Ert & Erev, 2008; Kahneman, Knetsch, & Thaler, 1991) has strong support in other domains, from voters' propensity to keep existing political parties in power (Jost, Banaji, & Nosek, 2004) to consumer decision making (Anderson, 2003). Nevertheless, we wanted to test the status quo bias empirically for security warnings in order to determine how many individuals turn them off and why.

Although only a minority of users may turn off malware warnings, there are various possible reasons for departure from it. (a) Some individuals prefer to make their own decisions; Lee and See (2004) report that individuals are reluctant to trust automated systems, when they have insufficient information about their operation. (b) Others might want to turn the warnings off because they ignore them anyway and just click through them. (c) Some will turn warnings off because they impact their productivity and are a waste of time. As Herley (2009) shows, this sentiment is realistic to a point – only a small percentage of Internet users suffer a setback from ignoring security advice; and skilled users consider many warnings pointless (as a Google manager said to us: 'Surely I am invulnerable to phishing?'). (d) Some individuals might not understand the warnings and would thus prefer to not see them. (e) There might be too many false positives. Krol, Moroz, and Sasse (2012) show this to be an important issue in user decision-making. And finally. (f) Many non-Windows users might feel that malware warnings are only relevant to Microsoft Windows users.

We therefore want to understand the extent to which some individuals depart from the default option. We hypothesize (H4) that the main reason for turning off malware warnings is the wish to remain productive and avoid being derailed by security notices. Furthermore, we hypothesize that (H5) the reasons for turning the malware warnings off will differ between the participants who kept the warnings on and those who turned them off (or expressed a wish to do so).

1.3. Aims

The present study aims to show the following: (a) confirm the status quo bas when it comes to malware warnings and analyse the self-reported reasons of why it is violated; and (b) show which textual treatment of browser malware warnings is effective in eliciting compliance.

2. Method

2.1. Participants

Our respondents for this study were recruited via Amazon Mechanical Turk (mTurk). In total, 583 mTurkers responded, and were distributed evenly across five conditions. In several cases same respondents participated in more than one condition. These duplicates were omitted from further analysis. In addition we removed incomplete cases, leaving us with 496 valid responses. The respondents were paid $0.70 on average per finished survey. In the measured group, age was normally distributed with a peak between 26 and 30 years of age. Gender was self-reported as 207 (42%) female and 281 (57%) male (8 respondents refusing to answer). Most of the respondents were United States residents. More demographics are given in the Results Section.