Sci Eng Ethics
DOI 10.1007/s11948-017-9952-9

ORIGINAL PAPER

Ethics and Phishing Experiments

David B. Resnik[1] · **Peter R. Finn**[2]

Received: 26 April 2017 / Accepted: 27 July 2017

Abstract Phishing is a fraudulent form of email that solicits personal or financial information from the recipient, such as a password, username, or social security or bank account number. The scammer may use the illicitly obtained information to steal the victim's money or identity or sell the information to another party. The direct costs of phishing on consumers are exceptionally high and have risen substantially over the past 12 years. Phishing experiments that simulate real world conditions can provide cybersecurity experts with valuable knowledge they can use to develop effective countermeasures and prevent people from being duped by phishing emails. Although these experiments contravene widely accepted informed consent requirements and involve deception, we argue that they can be conducted ethically if risks are minimized, confidentiality and privacy are protected, potential participants have an opportunity to opt out of the research before it begins, and human subjects are debriefed after their participation ends.

Keywords Phishing · Cybersecurity · Human experimentation · Ethics · Informed consent · Deception · Debriefing

✉ David B. Resnik
resnikd@niehs.nih.gov

[1] National Institute of Environmental Health Sciences (NIEHS), National Institutes of Health (NIH), 111 Alexander Drive, Research Triangle Park, NC 27709, USA

[2] Department of Psychological and Brain Sciences, Indiana University at Bloomington, Bloomington, IN, USA

Published online: 15 August 2017

Introduction

Phishing is a fraudulent form of email that solicits personal or financial information from the recipient, such as a password, username, or social security or bank account number. The sender may pose as a trusted source, such as friend, colleague, employer, or financial institution to trick the recipient into disclosing the information (Federal Trade Commission 2017). The scammer may use the illicitly obtained information to steal the victim's money or identity or sell the information to another party.[1] The direct costs of phishing on individuals and organizations are exceptionally high and have risen substantially in the last decade. The Gartner Group (2007) estimated the overall cost of phishing in the U.S. in 2007 at $3.2 billion. According to the Ponemon Institute (2015), a U.S. business with 10,000 employees spends an average of $3.77 million per year to deal with phishing attacks. While most of the earliest phishing emails included obvious mistakes that could tip-off recipients, such as misspelled names, phishers have become increasingly sophisticated and adept at imitating inquiries from trusted sources. Some even personalize the email by including the recipient's name, title, or other information.

To protect businesses, government agencies, and private citizens from phishing attacks, it is important to understand the factors that affect susceptibility to phishing schemes so that cybersecurity experts and organizations can develop effective countermeasures (Buchanan et al. 2011; Finn and Jakobsson 2007). Some of these factors may include: the age, gender, educational level, or race/ethnicity of the recipient; the type of sender (e.g. friend, financial institution, etc.); the format of the email; and the content of the email (Finn and Jakobsson 2007).

There are several different approaches to studying susceptibility to phishing attacks.[2] The first is to survey or interview individuals concerning their experiences with such attacks. Investigators can analyze subjects' answers to determine what types of phishing emails they have received and how they have responded to them. A limitation of this approach is that research subjects may not provide accurate information, because they may be unaware of phishing attacks or have difficulty recalling them, or they are unwilling to admit that they have fallen for a phishing scheme. Another drawback is that it provides little useful information concerning causal factors that impact susceptibility to phishing because it does not involve experimental manipulation of behavior (Finn and Jakobsson 2007; El-Din 2012).

The second approach is to test participants' ability to distinguish between phishing emails (including associated websites) and legitimate inquiries under laboratory conditions. Although these studies can provide useful information about participants' "phishing IQ" they also do not help researchers understand susceptibility to phishing because they do not experimentally manipulate behavior. Furthermore, the data investigators collect may not reflect how people behave under

[1] There are several different types of phishing attacks (Federal Trade Commission 2017; Ponemon Institute 2015). In the article, we focus on what is often referred to as "spear phishing."

[2] Organizations may also use these approaches for training employees in how to avoid phishing attacks, but we will focus on research activities in this paper.

natural conditions, because the subjects know they are being studied and may modify their behavior accordingly (Finn and Jakobsson 2007; El-Din 2012).

A third approach is for investigators to work with organizations (such as employers or educational institutions) to conduct experiments that mimic actual phishing attacks. The organization would provide the investigator with email addresses to use in the experiment. To model real world conditions, recipients should not be able to easily ascertain whether they are receiving actual phishing emails or messages concocted by investigators. The email would ask recipients to submit personal or financial information to a website linked to the message. To protect privacy and confidentiality, the website would be secure and submitted information would not be stored, although the sender's email address would be stored and checked against the list provided by the organization. The investigators would only receive data concerning the recipient's response to the email, i.e. whether they have visited the website and attempted to enter information (Finn and Jakobsson 2007; Jagatic et al. 2006).[3]

The first two approaches do not raise any significant ethical issues because they use a low-risk study design that gives potential participants the opportunity to consent and does not involve deception. The third approach, however, raises significant ethical concerns, because it includes participants in research without consent and involves deception. In this article, we will examine the ethical issues related to phishing experiments under real world conditions and argue that they can be conducted ethically if participants are given an opportunity to opt out of them before they begin and are debriefed afterwards; risks are minimized; and confidentiality and privacy are protected.

Phishing Experiments and Consent

The most significant ethical problem with phishing experiments under real world conditions is that they violate informed consent, which is required by numerous laws, guidelines, and professional codes (Emanuel et al. 2000). For example, the Department of Health and Human Services (DHHS) regulations, also known as the Common Rule[4] because they have been adopted by 17 U.S. federal agencies, mandate that DHHS-funded investigators obtain consent from human subjects or their legally authorized representatives before enrolling them in research (Department of Health and Human Services 2009).[5] The World Medical Association's

[3] This approach is similar to the penetration testing methodologies described by Dimkov et al. (2010). A penetration test is an attempt to gain access to an organization's secure information. The purpose of the test is to obtain knowledge that will help the organization improve its security. A physical penetration test is an attempt to gain access to information by physical means, such as removing a laptop from the organization or using a USB drive to download information. To provide information that is accurate and reliable, penetration tests should model real world conditions and therefore may need to include some deception of employees.

[4] It is worth noting that the Common Rule does not apply to an organization that receives no U.S. federal funding for research involving human subjects. We would argue that organizations should voluntarily comply with the ethical principles underlying the Common Rule when conducting phishing experiments with human subjects, even if they are not required to do so by law.

[5] References are to the 2009 version of the Common Rule. On 19 January 2017, the Obama Administration published long-awaited revisions to the Common Rule; however, the Trump

(2013) *Helsinki Declaration*, the Council for International Organizations of Medical Sciences' (2016) guidelines, and laws or policies adopted by other countries include similar consent requirements (United Kingdom, Department of Health 2005; Canadian Institutes of Health Research et al. 2014).

Informed consent rules are based on respect for autonomy, a fundamental ethical principle embodied in Anglo-American and European law and supported by different moral theories and traditions (Brock 2008; Beauchamp and Childress 2008). Respect for autonomy implies that researchers should honor the choices of autonomous individuals concerning their research participation. An autonomous individual is someone who is capable of making reasonable choices based on information and options (Brock 2008; National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research 1979).

While respect for autonomy is an important ethical principle, it is not absolute. Even the most ardent defenders of individual freedom recognize that the government can limit autonomy for compelling reasons, such as to protect other people from harm (Feinberg 1987) or promote the common the good (Dworkin 1988). Laws that prohibit murder, rape, battery, theft, fraud, libel, and reckless driving restrict autonomy to protect people from harm, and laws that require children to attend school, mandate vaccinations for school children, restrict commercial development on private property, or allow public health authorities to quarantine individuals exposed to infectious diseases, restrict autonomy to promote the common good (Bayer et al. 2006; Gostin 2007; Selgelid 2005).

Waiving Consent Requirements

Since there are often compelling reasons to restrict autonomy to promote the common good, one could argue that sometimes it is ethically permissible to include human subjects in a study without consent if (1) the research addresses important questions of public concern, (2) the research cannot be conducted if the subjects must provide consent, and (3) involving subjects in the research without their permission does not significantly compromise their autonomy (Gelinas et al. 2016; Miller 2008). Indeed, the U.S. federal research regulations allow committees that oversee research, i.e. institutional review boards (IRBs),[6] to alter informed consent requirements or waive them entirely if they determine that:

1. The research involves no more than minimal risk to the subjects;

Footnote 5 continued
Administration may make additional changes to these regulations or delay their implementation. The changes to the Common Rule do not impact the discussion of phishing experiments in this paper because they do not affect waivers of informed consent requirements for social or behavioral research. Although the changes include a new category of social/behavioral research exempted from the regulations, i.e. research involving benign interventions, this exemption only applies if the subjects prospectively agree to the intervention, which would not occur in the phishing experiments discussed herein (Department of Homeland Security et al. 2017).

[6] In other countries these committees may be called research ethics boards or research ethics committees.

2. The waiver or alteration will not adversely affect the rights and welfare of the subjects;
3. The research could not practicably be carried out without the waiver or alteration; and
4. Whenever appropriate, the subjects will be provided with additional pertinent information after participation (Department of Health and Human Services 2009 at 45 CFR 46.116d).

For an example of unconsented research that many would regard as ethical, consider a quality improvement study conducted at 103 intensive care units (ICUs) at 67 Michigan hospitals (Miller and Emanuel 2008). The protocol randomly assigned ICUs to a group that would implement standard infection control measures to prevent infections related to the placement of intravenous catheters (the control group) or to another group that would implement standard infection control measures and a checklist to ensure compliance with the measures (the experimental group). The study could not be conducted if informed consent of all the patients would be obtained because the infection control measures must be implemented in an entire ICU within the hospital and not on a per patient basis (Miller and Emanuel 2008; Gelinas et al. 2016). One might argue that an IRB could waive informed consent requirements for this study because: (1) the research involves minimal risks (because the study would provide all patients with standard infection control measures), (2) the research does not significantly compromise the rights or welfare of subjects (since the study has minimal risks and hospitals routinely implement quality improvement procedures without consent), (3) the research could not be conducted if consent were required (noted above), and (4) the subjects will be told about the study (i.e. debriefed) when their participation is complete (Gelinas et al. 2016).

Although the phishing experiments discussed in this paper are different from hospital quality improvement studies because they involve social/behavioral manipulations instead of medical interventions, we propose that an IRB would be justified in waiving informed consent requirements for these experiments because they can yield important knowledge concerning issues of public concern and are likely to fulfill the conditions for waiving consent.

First, the phishing experiments cannot be carried out without a consent waiver or alteration of consent, because the participants' knowledge of the experiments could invalidate the results (discussed earlier).

Second, the subjects would be provided with additional information concerning their participation after the experiments are concluded. (Debriefing is discussed in more depth below.)

Third, the phishing experiments would probably involve no more than minimal risk. The DHHS regulations define minimal risk as "the probability and magnitude of harm or discomfort anticipated in the research are not greater in and of themselves than those ordinarily encountered in daily life or during the performance of routine physical or psychological examinations or tests (Department of Health and Human Services 2009 at 45 CFR 46.102i)." The phishing experiments would probably not entail risks greater than those ordinarily encountered in daily life

because the subjects' private information will be protected on a secure website and the interactions with subjects are not substantially different from other deceptive interchanges people often face in daily life, such as misleading advertisements, emails, and phone calls (Finn and Jakobsson 2007; Benham 2008). (Risks related to deception are discussed in more depth below.)

Opt-Out Enrollment

Fourth, the research would not adversely affect the rights or welfare of the subjects. Since the risks of the research are likely to be minimal, the research would not adversely impact subject welfare. The impact of the research on the rights of the subjects depends on the extent that the research infringes autonomy. We recognize that involving human subjects in phishing experiments without consent raises significant issues related to autonomy even if the risks of this research are minimal, since most people would like to be allowed to decide whether to participate in behavioral experiments.

One way of minimizing the infringement of autonomy would be to give potential participants the opportunity to opt out of these studies beforehand. To ensure that their knowledge of the research does not bias the data, some deception would be necessary. Potential participants could be given an opportunity to opt out of the research but would not be told that it would involve phishing experiments, so that participants who receive a phishing email would have no reason to suspect that it was a behavioral experiment. An organization (such as university or employer) that approves one of these experiments on its population of email users could announce that it is planning to study email patterns and trends at the organization to understand how to better protect users from cybersecurity threats. To address concerns about privacy and confidentiality, the announcement would inform users that the studies will not examine the contents of any emails but will only collect metadata on email activity and may involve an experimental manipulation. The announcement would inform email users how to opt out of the research. Users that opt out would not receive any phishing emails from investigators. The institution could make the announcement several times in different venues (e.g. via email, advertisements in newspapers and on the radio) to ensure that email users are aware that they can refuse to participate in the research. Only users who have reached the age of majority[7] would participate. To avoid involving users' friends in this experiment, the phishing email could appear to come from a trusted institution. While this proposal would not provide the degree of autonomous decision-making that occurs in fully informed consent, it would give email users some control over whether they would participate in these studies and thus minimize infringement on autonomy.

Some might argue that it would be more respectful of autonomy to ask subjects to opt into cybersecurity research rather than allowing them to opt out because opting in requires more mental effort and reflection than opting out, and some potential

[7] We assume the organization would have information about email users' age.

participants might not be aware that they can opt out (Annas 2000). Individuals who enroll by an opt-in procedure are therefore more likely to have made a fully autonomous choice than those who enroll by an opt-out procedure (Mackay 2015). We agree that opting in is more respectful of autonomy than opting out, but we are concerned that using an opt-in procedure may bias the study population (Cassell and Young 2002; Junghans et al. 2005), because individuals who opt in are more likely to be aware of and concerned about cybersecurity issues—and therefore less likely to respond to phishing emails—than those who do not opt in.[8]

Since the infringement on autonomy entailed by opt-out enrollment is minimal and generally accepted by the public (Vellinga et al. 2011), we do not think an opt-in procedure is necessary to adequately protect the rights of human subjects. However, we remain open to further discussion concerning this topic and an oversight committee might decide that an opt-in procedure is preferable.

Phishing Experiments and Deception

While our proposal for prior authorization for participation in cybersecurity research addresses some of the concerns related to informed consent, it still involves deception, which is ethically controversial (Finn 1995). The deception is necessary, we believe, to obtain data that accurately reflects how people respond to phishing emails under real world conditions.

Deceiving research subjects is controversial because it (1) interferes with informed consent by preventing subjects from receiving information pertinent to deciding whether to participate, (2) can undermine the trust that subjects have in the research enterprise when they find out that they have been deceived, and (3) can cause harm in some cases (Wendler and Miller 2008).

Evidence concerning the impact of deception on human subjects indicates that harms are usually minimal and temporary (Boynton et al. 2013; Epley and Huff 1998; Hertwig and Ortmann 2008; Pihl et al. 1981; Smith and Richardson 1983; Soliday and Stanton 1995). In a seminal study on the impact of deception, Smith and Richardson (1983) surveyed 495 undergraduate students who had participated in psychological research. They found that 28.7% of 195 students who had participated in experiments involving deception reported the perception of harm, such as anger, nervousness, discomfort, and humiliation, while only 17% of 269 students who had participated in studies not involving deception reported the perception of harm. They also found, however, that effective debriefing following the deceptive experiments (i.e. explaining to the subjects the true design and purpose of the research) eliminated participants' perception of harm (Smith and Richardson 1983).

[8] One way to help resolve this issue would be to conduct studies that compare opt-in and opt-out procedures to determine whether either method has a substantial enrollment bias. However, it may be difficult to obtain data for these studies because researchers will not be asking subjects for their consent and therefore will not have access to important information that might bias subject selection. The researchers in our proposed study would probably have access to some demographic information about the subjects, but they would not have data pertaining to other important variables, such as their awareness of or cybersecurity issues or their attitudes toward research participation.

While the evidence indicates that harms resulting from deceptive experiments are usually minimal and transient, Wendler and Miller (2008) argue that a small percentage of participants may experience significant and long-lasting harms resulting from participating in some studies that use deception. For example, several of the subjects who were deceived in Milgram's (1974) controversial obedience to authority experiments reported significant psychological distress, shame, and guilt after they realized that they were willing to act immorally by administering dangerous electric shocks to people (Sobel 1978).[9] However, the nature of the potential harm to subjects in the Milgram experiments (i.e., guilt or shame for doing something that one deems immoral) seems to be far greater than the impact of finding out one was duped in a phishing experiment. Recipients of the phishing emails are not being asked to disclose someone else's confidential information or break laws, for example. In fact, debriefing in a phishing experiment has the potential of having a long-lasting positive impact, if debriefing is structured in such a way that the subject learns how to avoid being a victim of future real phishing attacks Finn and Jakobsson 2007).

Debriefing is widely recognized as a method for minimizing the harms related to deception and holding researchers accountable to subjects (Miller et al. 2008; Oczak and Niedźwieńska 2007). Debriefing should inform the subjects about the nature and importance of the research and the rationale for the deception (Finn 1995; Miller et al. 2008). However, debriefing in some types of phishing experiments involving deception may result in psychological harm. Jagatic et al. (2006) conducted a phishing experiment on 1731 email users at Indiana University at Bloomington who were at least 18 years old. This was a "social phishing" experiment in that the email appeared to come from a friend.[10] After the experiment was over, they sent the subjects an email informing them about it. The email also gave the participants an opportunity to express their concerns by participating in a discussion forum on a website. Jagatic et al. (2006) found that many participants in their experiment expressed anger and frustration with the researchers and the IRB in the discussion forum.

There are a couple of reasons why this debriefing may not have been particularly helpful. First, the debriefing occurred online rather than in-person (Finn and Jakobsson 2007). Since individuals who learn that they have been deceived may feel harmed or violated, it is best for investigators to address their concerns in-person to earn their trust and reassure them that harms they may have experienced were not

[9] Milgram's experiments involved two types of human subjects, learners and teachers. The teachers presented the learners with lists of word-pairs they were supposed to memorize. The learners were hooked up to a machine that appeared to be capable of giving them an electric shock. The investigators instructed the teachers to administer a shock to the learners whenever they gave an incorrect answer. Most of the teachers continued administering shocks even when the learners cried out in pain and asked the experiment to stop. In reality, the learners never received a shock. The purpose of the experiment was to determine whether the teachers would obey instructions to give a shock to the learners. The teachers consented to participating in the study, but they were not told they were being deceived. Milgram debriefed the learners after their participation was complete and explained the true nature of the experiment to them. See Milgram (1974) for further discussion.

[10] "Social phishing" is more ethically problematic than the phishing experiments discussed in this paper since it involves two people who are involved in research without consent, i.e. the recipient and the friend.

deliberate. Because online debriefing can be impersonal and uncaring, it may exacerbate the subjects' negative emotional reactions. However, it was not practical for the investigators to have face-to-face meetings with 1731 participants. We propose that phishing experiments should use a smaller sample of participants (e.g. a few hundred)[11] to make it easier for investigators to debrief them in-person.

Second, since the individuals had no prior knowledge that they were enrolled in a research study, the email arrived unexpectedly, which could have added to their anger and frustration. One might argue that since debriefing subjects online may be more harmful than not debriefing them at all, debriefing may be inadvisable in phishing experiments that simulate real world conditions (Finn and Jakobsson 2007). However, we do not agree with this proposal because we think that more harm could occur if subjects are not debriefed. For example, if subjects were to learn that they participated in a phishing experiment without prior consent (for example, by reading a journal or newspaper article), they could become even more angry and distrustful than if they had been debriefed online because they might view the research as surreptitious and manipulative.

We think the best way to minimize potential negative impacts of online debriefing is to give potential subjects prior notice about cybersecurity research being conducted at the institution (by an opt-out process described above), which could help to soften the blow of debriefing. Since the subjects would have prior knowledge about their potential participation in cybersecurity research, the debriefing email would not be completely unexpected. In addition to explaining the nature and importance of the research and the rationale for deception, debriefing should include some education on the different types of phishing scams and how to avoid succumbing to them, and give participants the opportunity to ask additional questions or meet in person with the investigators. Thus, this study design could not only contribute to the advancement of knowledge concerning cybersecurity risks and countermeasures but also benefit the subjects by informing them about how to protect themselves from phishing attacks. The debriefing email could also invite subjects to participate in a follow-up survey concerning psychological, social, and economic characteristics that may indicate degree of vulnerability to an attack. Data from the follow-up survey could be especially useful in designing phishing countermeasures.

Wendler and Miller (2008) argue that one way of respecting autonomy and minimizing the potential harms related to deception is to ask potential participants to consent to some type of deception. Under this proposal, which Wendler and Miller (2008) call "authorized deception," investigators inform potential participants that the research may involve some form of deception, such as withholding information or providing misleading information. For example, the informed consent document for a clinical trial may notify participants that they will be randomly assigned to receive a placebo or an experimental treatment. The document may also inform participants that investigators will implement procedures, such as

[11] We recognize that smaller samples may not have enough statistical power to achieve significance results. To deal with this issue, investigators should carefully select a sample size that is adequately powered but also is not so large that in-person debriefing is impractical.

double-blinding, to prevent them from discovering whether they are receiving a placebo (Wendler and Miller 2008). In some cases, the procedures used to prevent subjects from learning that they are receiving a placebo may cause harm. For example, when a drug with known side-effects (such as nausea) is being tested against a placebo, it may be necessary to give participants a medication in the placebo pill that mimics the side-effects.

While authorized deception has considerable ethical appeal, it may compromise the validity of the results, because participants may try to determine whether or how they are being deceived and alter their behavior accordingly. In phishing experiments under real world conditions, if participants are informed that they are participating in research that may involve deception, they might decide to provide the information requested by the emails because they believe that these are only experimental manipulations; or they may make the opposite decision because they think the investigators do not want them to respond to the email. In either case, participants' knowledge of the deceptive nature of the experiment could impact the outcome. Since authorized deception may bias the results of these phishing experiments, we believe that the better approach for the phishing experiments is authorized participation (i.e. the opt-out procedure described above), not authorized deception.[12]

Conclusion

Phishing experiments under real world conditions can provide cybersecurity experts and organizations with valuable knowledge they can use to develop effective countermeasures and prevent individuals from being duped by phishing emails. Although these experiments contravene widely accepted informed consent requirements and involve deception, we argue that they can be conducted ethically if risks are minimized, confidentiality and privacy are protected, potential participants have an opportunity to opt out of the research before it begins, and subjects are appropriately debriefed after their participation ends. Since phishing research involving human subjects is relatively new, it is important for researchers, sponsors, organizations, and oversight committees to continue to discuss how to ensure that it meets scientific and ethical standards.

Acknowledgements This research was funded by the Intramural Program of the National Institute of Environmental Health Sciences (NIEHS), National Institutes of Health (NIH) (ZIAES-102646-08). It does not represent the views of the NIEHS, NIH, or U.S. government.

Compliance with Ethical Standards

Conflict of interest The authors have no conflicts of interest to disclose.

[12] Milgram's experiments involved authorized participation not authorized deception. See footnote 3.

References

Annas, G. J. (2000). Rules for research on human genetic variation—Lessons from Iceland. *New England Journal of Medicine, 342*(24), 1830–1833.

Bayer, R., Gostin, L. O., Jennings, B., & Steinbock, B. (Eds.). (2006). *Public health ethics: Theory, policy, and practice*. New York, NY: Oxford University Press.

Beauchamp, T., & Childress, J. (2008). *Principles of biomedical ethics* (8th ed.). New York, NY: Oxford University Press.

Benham, B. (2008). The ubiquity of deception and the ethics of deceptive research. *Bioethics, 22*(3), 147–156.

Boynton, M. H., Portnoy, D. B., & Johnson, B. T. (2013). Exploring the ethics and psychological impact of deception in psychological research. *IRB, 35*(2), 7–13.

Brock, D. E. (2008). Philosophical justifications of informed consent. In E. J. Emanuel, C. Grady, R. A. Crouch, R. K. Lie, F. G. Miller, & D. Wendler (Eds.), *The Oxford textbook of clinical research ethics* (pp. 606–612). New York, NY: Oxford University Press.

Buchanan, E., Aycock, J., Dexter, S., Dittrich, D., & Hvizdak, E. (2011). Computer science security research and human subjects: Emerging considerations for research ethics boards. *Journal of Empirical Research on Human Research Ethics, 6*(2), 71–83.

Canadian Institutes of Health Research, Natural Sciences and Engineering Research Council of Canada, and Social Sciences and Humanities Research Council of Canada. (2014). Tri-Council policy statement: Ethical conduct for research involving humans. http://www.pre.ethics.gc.ca/pdf/eng/tcps2-2014/TCPS_2_FINAL_Web.pdf. Accessed 14 July 2107.

Cassell, J., & Young, A. (2002). Why we should not seek individual informed consent for participation in health services research. *Journal of Medical Ethics, 28*(5), 313–317.

Council for International Organizations of Medical Sciences. (2016). International ethical guidelines for health-related research involving humans. https://cioms.ch/shop/product/international-ethical-guidelines-for-health-related-research-involving-humans/. Accessed 12 July 2017.

Department of Health and Human Services. (2009). Protection of Human Subjects, 45 Code of Federal Regulations 46.

Department of Homeland Security; Department of Agriculture; Department of Energy; National Aeronautics and Space Administration; Department of Commerce; Social Security Administration; Agency for International Development; Department of Housing and Urban Development; Department of Labor; Department of Defense; Department of Education; Department of Veterans Affairs; Environmental Protection Agency; Department of Health and Human Services; National Science Foundation; and Department of Transportation. (2017). Federal policy for the protection of human subjects. *Federal Register, 82*(12), 7149–7274.

Dimkov, T., Pieters, W., & Hartel, P. (2010). Two methodologies for physical penetration testing using social engineering. In *Proceedings of the annual computer security applications conference* (pp. 399–408). New York, NY: American Chemical Society.

Dworkin, G. (1988). *The theory and practice of autonomy*. Cambridge, UK: Cambridge University Press.

El-Din, R. S. (2012). To deceive or not to deceive! Ethical questions in phishing research. In *Proceedings of the British Computing Society, human–computer interaction 2012 Workshops*. http://ewic.bcs.org/upload/pdf/ewic_hci12_ec_paper2.pdf. Accessed 26 April 2017.

Emanuel, E. J., Wendler, D., & Grady, C. (2000). What makes clinical research ethical? *Journal of the American Medical Association, 283*(20), 2701–2711.

Epley, N., & Huff, C. (1998). Suspicion, affective response, and educational benefit as a result of deception in psychology research. *Personality and Social Psychology Bulletin, 24*(7), 759–768.

Federal Trade Commission. (2017). Phishing. https://www.consumer.ftc.gov/articles/0003-phishing. Accessed 26 April 2017.

Feinberg, J. (1987). *Harm to others*. New York, NY: Oxford University Press.

Finn, P. R. (1995). The ethics of deception in research. In R. L. Penslar (Ed.), *Research ethics: Cases and materials* (pp. 87–118). Bloomington, IN: Indiana University Press.

Finn, P. R., & Jakobsson, M. (2007). Designing ethical phishing experiments. *IEEE Technology and Society, 26,* 46–58.

Gartner Group. (2007). Phishing costs the U.S. economy $3.2 billion. Press release, 17 December 2007. https://www.finextra.com/news/fullstory.aspx?newsitemid=17871. Accessed 12 July 2017.

Gelinas, L., Wertheimer, A., & Miller, F. G. (2016). When and why is research without consent permissible? *Hastings Center Report, 46*(2), 35–43.
Gostin, L. O. (2007). General justifications for public health regulation. *Public Health, 121*(11), 829–834.
Hertwig, R., & Ortmann, A. (2008). Deception in social psychological experiments: Two misconceptions and a research agenda. *Social Psychology Quarterly, 71*(3), 222–227.
Jagatic, T. N., Johnson, N. A., Jakobsson, M., & Menczer, F. (2006). Social phishing. *Communications of the ACM, 50*(10), 94–100.
Junghans, C., Feder, G., Hemingway, H., Timmis, A., & Jones, M. (2005). Recruiting patients to medical research: Double blind randomised trial of "opt-in" versus "opt-out" strategies. *British Medical Journal, 331*(7522), 940.
MacKay, D. (2015). Opt-out and consent. *Journal of Medical Ethics, 41*(10), 832–835.
Milgram, S. (1974). *Obedience to authority*. New York, NY: Harper and Rowe.
Miller, F. G. (2008). Research on medical records without informed consent. *Journal of Law, Medicine & Ethics, 36*(3), 560–566.
Miller, F. G., & Emanuel, E. J. (2008). Quality-improvement research and informed consent. *New England Journal of Medicine, 358*(8), 765–767.
Miller, F. G., Gluck, J. P., Jr., & Wendler, D. (2008). Debriefing and accountability in deceptive research. *Kennedy Institute of Ethics Journal, 18*(3), 235–251.
National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. (1979). *The Belmont report: Ethical principles and guidelines for the protection of human subjects of research*. Washington, DC: Department of Health, Education, and Welfare.
Oczak, M., & Niedźwieńska, A. (2007). Debriefing in deceptive research: A proposed new procedure. *Journal of Empirical Research on Human Research Ethics, 2*(3), 49–59.
Pihl, R., Zacchia, C., & Zeichner, A. (1981). Follow-up analysis of the use of deception and aversive contingencies in psychological experiments. *Psychological Reports, 48*(3), 927–930.
Ponemon Institute. (2015). The costs of phishing & value of employee training. https://info.wombatsecurity.com/hubfs/Ponemon_Institute_Cost_of_Phishing.pdf?t=1499361243887. Accessed 6 July 2017.
Selgelid, M. J. (2005). Ethics and infectious disease. *Bioethics, 19*(3), 272–289.
Smith, S. S., & Richardson, D. (1983). Amelioration of deception and harm in psychological research: The important role of debriefing. *Journal of Personality and Social Psychology, 44*(5), 1075–1082.
Sobel, A. (1978). Deception in social science research: Is informed consent possible? *Hastings Center Report, 8*(5), 40–45.
Soliday, E., & Stanton, A. L. (1995). Deceived versus nondeceived participants' perceptions of scientific and applied psychology. *Ethics and Behavior, 5*(1), 87–104.
United Kingdom, Department of Health. (2005). Research governance framework for health and social care. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/139565/dh_4122427.pdf. Accessed 14 July 2017.
Vellinga, A., Cormican, M., Hanahoe, B., Bennett, K., & Murphy, A. W. (2011). Opt-out as an acceptable method of obtaining consent in medical research: A short report. *BMC Medical Research Methodology, 11,* 40.
Wendler, D., & Miller, F. G. (2008). Deception in research. In E. J. Emanuel, C. Grady, R. A. Crouch, R. K. Lie, F. G. Miller, & D. Wendler (Eds.), *The Oxford textbook of clinical research ethics* (pp. 315–324). New York, NY: Oxford University Press.
World Medical Association. (2013). Declaration of Helsinki: Ethical principles for medical research involving human subjects. https://www.wma.net/policies-post/wma-declaration-of-helsinki-ethical-principles-for-medical-research-involving-human-subjects/. Accessed 12 July 2017.