

IRBs and Security Research: Myths, Facts and Mission Creep

Simson L. Garfinkel

(slgarfin@nps.edu)

Naval Postgraduate School & Harvard University

April 7, 2008

Abstract

Having decided to focus attention on the “weak link” of human fallibility, a growing number of security researchers are discovering the US Government’s regulations that govern human subject research. This paper discusses those regulations, their application to research on security and usability, and presents strategies for negotiating the Institutional Review Board (IRB) approval process. It argues that a strict interpretation of regulations has the potential to stymie security research.

1 Introduction

As more security researchers turn their attention to usability and other human factors issues, many are surprised to discover that they must comply with regulations governing the use of human beings as experimental subjects.

These regulations, known collectively as “The Common Rule,” were created after a series of well-publicized abuses in the 1960s and 1970s. These regulations require those working with US Government funds to receive approval from their organization’s designated Institutional Review Board (IRB) before most research involving human subjects can commence.

There seems general understanding among researchers that hands-on laboratory usability experiments are covered under the IRB rules. But many other kinds of less-invasive research may still require IRB notification and approval. Furthermore, it appears that many researchers either do not understand their legal obligations, or else have simply chosen to ignore them.

While this paper concerns itself solely with US law, there are “approximately 900 laws, regulations, and guidelines that govern human subjects in 84 countries, as well as from a number of international and regulation organizations”[12]. A list can be found in the reference.

2 Legal Framework

2.1 45 CFR Part 46 Subpart A

The National Research Act (Pub. L. 93-348) created the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. This Commission identified a set of ethical principles, and later practical guidelines, for US Government-funded research involving human beings. Ultimately the Department of Health and Human Services wrote and adopted a set of regulations requiring recipients of USG funding to create their own institutional bureaucracies for overseeing re-

search that involves human subjects. These regulations, embodied in Title 45 Part 46 subpart A of the Code of Federal Regulations, exist for one purpose: *to safeguard the welfare of human research subjects.*

Although CFR 45 Part 46 applies only to work funded by HHS, regulations with the same language were adopted in 1991 by 14 other grant-giving USG agencies, including the Department of Defense (DoD), the National Aeronautics and Space Administration (NASA), and the National Science Foundation (NSF). For this reason 45 CFR part 46 subpart A is referred to as “the Common Rule for the protection of human subjects”[13].

2.2 The Institutional Review Board (IRB)

The Common Rule requires organizations performing federally supported research to designate an institutional review board (IRBs) to oversee human subject research. Many organizations fulfill this requirement by establishing their own IRBs. Each IRB must have “at least five members” with different backgrounds so that it can “adequately review [the] research activities commonly conducted by the institution”(§46.107 (a)). Members must include representatives of both genders, at least one scientist and one non-scientist, and one member who is not affiliated with the institution(§46.107(b,c,d)). Members are prohibited from voting on their own research(§46.107(e)).

2.3 IRB Coverage

The Common Rule is very clear: With the exception of six specific categories of research, the Common Rule “applies to all research involving human subjects conducted, supported or otherwise subject to regulation...”(§46.101(a))

Research is defined as “a systematic investigation, including research development, testing and evaluation, designed to develop or contribute to generalizable knowledge” (§46.102(d)). *Human Subject* is defined as “a living individual about whom an investigator (whether professional or student) conducting research obtains 1) Data through intervention or interaction with the individual, or 2) Identifiable private information”(§46.102(f)).

The IRB regulations allow organizations to augment the rules, adding requirements, and broadening them to include more areas of research(§46.112). At many schools, including MIT[8], Harvard[11] and UC Berkeley[3], IRB approval is required for *any* research involving human subjects, regardless of funding, and even if the research is exempt under the Common Rule.

2.4 IRB Exemptions

The Common Rule contains exemptions for some research involving human beings. For security research the relevant exemptions are:

- Research to be conducted on educational practices or with educational tests (§46.101(b)(1&2)).
- Research involving “existing data, documents, [and] records. . .” provided that the data set is either “publicly available” or that the subjects “cannot be identified, directly or through identifiers linked to the subjects” (§46.101(b)(4)).
- Research involving “survey procedures, interview procedures or observation of public behavior,” unless information is obtained that could identify the human subjects, and “any disclosure of the response outside the research could place the subjects at risk of criminal or civil liability or be damaging to the subjects’ financial standing, employability, or reputation” (§46.101(b)(2)(i&ii)).

Few Organizations allow investigators to decide if their own research is exempt under the guidelines—that would create a conflict-of-interest. For example, Harvard requires that “all research using human subjects” to be “reviewed or designed as exempt from review by a Harvard Institutional Review Board” [11]. In practice, all requests at Harvard to involve human subjects in research goes to a staff member at one of the Harvard IRBs; the staff then decides if the research is actually exempt or in need of review.

2.5 IRB Myths and Facts

Below are some popular myths about IRBs and the Common Rule that are apropos to computer security research.

Myth: Because the Common Rule exempts research involving subjects that cannot be identified, IRB approval is not required when using anonymized data Although this would certainly be convenient, most institutions only allow a determination of exemption to be made by the IRB itself.

Myth: “Pilot studies” do not require IRB approval. Although some schools have policies which define a kind of “pilot study” not requiring IRB approval, there is no support for this interpretation in the Common Rule, which makes no reference to “pilot” or “preliminary” studies.

Many universities (e.g. [15, 9, 4]) have specific language in their IRB guidelines stating that IRB approval is required for *all* research, even pilot studies that will not be published. Georgia State University’s policy [9] goes further, requiring consent forms to indicate if a study is a pilot study and requiring that the experimenter obtain additional IRB approval when the study progresses beyond preliminary stages.

But some organizations allow unapproved pilot studies: the School of Social Service Administration at University of Chicago allows small-scale pilot studies with less than

10 individuals to proceed without IRB approval assuming that “proper steps will be taken to protect human subjects,” sensitive data will not be collected, vulnerable populations will be excluded, and methods with no more than minimal risk will be used. However, SSA/UC requires IRB approval if the data collection in the pilot study will be used in any publication; if the data is to be used, IRB approval is required before data collection begins [5].

Myth: IRB approval is not required if you are working with data you already have. The Common Rule makes no such exception. If previously collected data will be analyzed using a methodology that is different than that which was described in the original IRB application, new approval may be required.

Fact: IRB approval is not required by the Common Rule when using publicly available data. The Common Rule states that research involving the “collection or study” of “existing data, documents [or] records” is exempt “if these sources are publicly available” §46.101(b)(4). But, as previously noted, most institutions require IRB approval for *all work* involving human subjects, even research exempt under the Common Rule.

3 Research Implications

The Common Rule places security researchers in a difficult position: Large quantities of data in our possession or which are easy to get appears to be off-limits for research without prior IRB approval.

3.1 Scenarios

In this section we present seven scenarios of usability & security research projects that potentially use human subject data and which could not have been anticipated when the Common Rule was drafted. We will briefly describe each one; in the following section we’ll discuss whether or not the scenario would require IRB approval.

Scenario 1: Security toolbar with anonymized summary statistics. Alice has developed an anti-phishing toolbar. To assist in development and research, the toolbar sends a small anonymized report to the experimenter once a day. Because each toolbar reports only once every 24 hours, it is easy for the experimenter to measure adoption and use of the toolbar.

Scenario 2: Web server logfile analysis. Bob’s research group operates a popular web-based discussion forum. Bob writes a program which analyzes the web-server’s log file to report the number of daily password resets. He also instruments the software to record the number of newly chosen passwords that do not pass the website’s password complexity rules. The research plan is to see how these numbers change as the rules become successively more restrictive over time. To assure that no personally identifiable information is collected, Bob configures the Apache server so that IP addresses are not logged.

Scenario 3: Popular security search terms. Christine is a graduate student who also writes articles for a major security-related website. She is working on a project that correlates search terms on the website with news stories. The security-related website prepares a report which shows, for each hour, the number of times each term is searched. The report is sent as a PGP-encrypted file to Christine's Gmail account.

Scenario 4: Building better spam filters. Don is creating a better spam filter and wishes to test it on his own inbox.

Scenario 5: Wi-Fi Security Survey. Elaine installs a copy of NetStumbler on a laptop and drives around the neighborhood with a GPS. She compares the names and locations of the Wi-Fi sites with a similar database available online. Her research shows that most of the older Wi-Fi access points that were open are now either closed or have been removed from service. Of the new Wi-Fi access points that have been deployed, most are closed.

Scenario 6: Hidden Data Survey. Guy writes a web crawler and downloads 100,000 Microsoft Word files from public websites. He analyzes the files and finds that approximately 15% contain significant amounts of hidden information. He randomly chooses 100 of the 15,000 files to confirm her software's findings.

Scenario 7: Online EXIFs. Felicity downloads 10,000 JPEGs from a social network website. By examining the camera serial numbers in the images he is able to determine which images were shot by the same camera. She then shows that he can reconstruct the "friends" networks—and shows which pseudonyms are sharing the same camera, indicating that they might be living together or might be the same person.

3.2 Analyzing the Scenarios

Scenarios 1, 2, 3, and 4 clearly require IRB approval under the Common Rule:

1. Alice needs IRB approval because she is recruiting, interacting with, and collecting information from her subjects. Furthermore, her users reveal their IP address when the toolbar reports its statistics; although IP addresses do not necessarily reveal personal information, they frequently do—especially in a university environment where an address may be assigned to a specific person.

2. Even though Bob is not collecting IP addresses, he still needs IRB approval because the information in the webserver logs was generated by human subjects and is not publicly available.

3. Likewise, Christine requires IRB approval because the data is generated by human beings and is not publicly available. Christine could avoid IRB involvement if the security website published the search terms on a public web page rather than encrypting them and sending them to her Gmail account. (Although it seems that this creative

way to bypass the Common Rule has exactly the opposite of the desired effect, presumably the website would do its own privacy audit before releasing such information.)

4. Don has thousands of legitimate and spam messages, but the Common Rule prohibits him from analyzing incoming email without IRB approval.

The other scenarios are more troublesome:

5. Elaine might not require IRB approval for his survey because she is not observing people: She is observing wireless access points. But these devices were configured (or not configured) by people. And in some rural areas the GPS coordinates might identify specific individuals.

6. Guy's research with Word documents literally involves "existing ... documents" that are "publicly available," so it should be exempt under the Common Rule. But the documents might have been inadvertently placed on the Internet and contain private information.

7. Felicity's research, like Guy's, probably does not fall under the Common Rule if she only collects documents (e.g., the photographs) that are publicly available. But what if the images were available to a large community but not the general public?

3.3 Researcher's Can't Say "Trust me"

Although researchers might be frustrated by the conclusions of this analysis, it is important to realize that the Common Rule is doing precisely what Congress intended when it passed the National Research Act: Congress wanted to put a stop to scientists saying "trust me." For decades, scientists had argued that good scientific practice and ordinary research ethics would protect the interests of their subjects. Experience proved otherwise.

With the National Research Act, Congress concluded that some scientists were not worthy of trust when it came to evaluating the impact of their own experiments on experimental subjects. And with good reason: some research involving human beings frequently requires deception, stress, or bodily risk. The Act recognizes that it is sometimes unreasonable to ask a scientist to be both an advocate for their research and their research subjects well-being.

For example, Alice could data mine her logfiles, seek out the personally identifiable IP addresses, perform Google searches to correlate IP addresses with email addresses, and then create a web page that identifies people who have "good" security practices (because they run her program) and "bad" practices (because they uninstalled her program.) Perhaps she might even send phishing attacks to the subjects to see how they respond. The IRB structure provides a place for someone who has had training to review her research protocol.

There are many ways to "anonymize" log files: sometimes the anonymization is incomplete and personally identifiable information can be recovered. One reason to require IRB review of research involving "anonymized"

logs is so that a neutral third-party can review protocol. Otherwise, we are just trusting the good judgement of the researcher—a person who has an inherent conflict-of-interest.

The analysis in Don’s case seems silly: after all, Don already has the email messages, and they were voluntarily sent to him by his friends and colleagues. But those people didn’t send Don the email for the purpose of being involved in an experiment. Part of the IRB process is to protect human beings who are involved in research without their knowledge or consent.

Since each of the scenarios above involve no more than minimal risk, we believe that an IRB would properly approve each protocol under the Common Rule’s “expedited review procedures”(§46.110). At many organizations the expedited review involves a form that is submitted by email and is administratively approved within days by either the IRB chair or a staff member. Although it appears to be a formality, expedited review has an important role: it forces the experimenter to create a written description of the research protocol. The mere act of writing down the protocol and discussing it with the IRB may help the experimenter to realize ways to further minimize the impact of the experiment on the human participants.

3.4 The Human Test

One useful test of the need to involve an IRB is to ask this question: *would the experiment be useful if the data were generated by a random process and not by a human?* For scenarios #2 and #3 above, the answer would clearly be “no.” The only reason that our hypothetical Bob is interested in reviewing password rule violations, and Christine is interested in search terms is precisely because this information is being generated by human subjects. If the passwords and search terms were randomly chosen, the research would not be worth doing.

If the experiment can be performed with randomly generated data, then use random data. This is an application of the “respect for persons”[2] ethical principle.

4 Working with the IRB

It is widely believed that different IRBs apply different standards when deciding whether or not to approve computer security research. This is likely a result of IRBs having different levels of experience with this kind of research.

We suggest several approaches that security researchers can take to improve the situation with their IRBs:

- Researchers should be intimately familiar with both the Common Rule and whatever local regulations their home institutions may have adopted.
- Researchers must learn how to make clear and cogent arguments that their research should be approved under the “expedited review procedures” on

the grounds that there is “minimal risk” to the experimental subjects.

- IRBs have the authority to waive informed consent requirements (§46.116(c,d)). Researchers should become familiar with this option and request it where appropriate.
- Researchers should be familiar with protocols that have been approved by other IRBs. The research community would also benefit from having an open repository of approved protocols.
- Security researchers should volunteer to serve on their organization’s IRBs. Bringing security expertise to the IRBs in this manner will help educate other IRB members and ease the way for other security research involving human subjects. (We have heard stories of IRBs that have blocked membership of computer scientists on the grounds that they were not biomedical researchers and the position on the IRB reserved for a non-scientist was already taken. Such positions are a misreading of the Common Rule, which specifies minimums but not maximums of IRB membership(§46.107)).

5 IRB “Mission Creep”

IRBs are centers of power: they have the ability to shut down research, and they operate with little institutional supervision. There are also complaints that IRBs have become more restrictive, in some cases acting as if their purpose is to safeguard their institutions from lawsuit, rather than to protect the welfare of experimental subjects. Perhaps as a result, some researchers have complained of a pervasive IRB “Mission Creep,” in which more and more research is being placed under the purview of IRBs.

In November 2005 the Center for Advanced Study at the University of Illinois College of Law held a conference on the topic of “IRB mission creep” and produced a 32 page paper, based on a two-year study, stating that IRBs were being stretched thin by being forced to pay excessive attention to research that poses little chance of risk. The report also called for “removing some kinds of activity from IRB review altogether,” especially journalism and ethnography[1]. The authors later published an editorial in *Science* making many of the same claims[10].

Journalism is particularly a problem, the Illinois group argued, because people who are the subject of journalistic research are frequently injured by the results. The example they give is that of President Richard Nixon, whose reputation was damaged and who lost his job as the result of the Watergate investigation. Although this research had great social value, it might not have been permitted by an IRB.

Katz argues that the Common Rule turns ethnographers into “IRB Outlaws” when they perform fieldwork by living with a host community to learn about it. The very

premise of field work is that results cannot be predicted, so it is impossible to get approval in advance from an IRB for what might happen[14]. Recently an entire issue of *American Ethnologist* was devoted to this topic[7].

A growing number of observers have criticized IRBs for being research censorship instruments. Cohen writes that an IRB that he sat on, which was primarily comprised of health care professionals, “received a qualitative, social science project to review. The IRB promptly disapproved it because it wasn’t science in the view of the IRB members”[6]. The protocol was eventually approved, but only after Cohen convinced the IRB to send the proposal out for external review.

6 Conclusion

Some might say that this call for IRB involvement in computer security research is part of IRB Mission Creep. But the Common Rule, as written, clearly applies to much work in the field of usability, psychology and security.

The problem here is that the Common Rule was written to cover biomedical and psychological research. It is clear that the authors never imagined a day that not just researchers, but most members of our society would have desktop computers containing personal information created by thousands of individuals with whom we have no direct relationship.

The penalties for performing research without approval include forced termination of research and loss of funding: we ignore the rules at our own peril. But in the long term, society would be better served with broader exemptions that could be automatically applied by researchers without going to an IRB.

Revisions to the Common Rule should also address a particularly wasteful practice: the intentional destruction of data which was collected without IRB approval: this practice certainly seems to violate the Belmont Report’s “respect for persons” principle.

The Common Rule was created because of abuses in medical and psychological research, but the Rule was very broadly written. If it cannot be revised or reinterpreted, the impact on computer security research may be severe.

7 Acknowledgements

This paper is partially funded by Naval Postgraduate School’s Research Initiation Program.

Brent Olde and Amelia Weinreb both provided useful insights in the preparation of this paper. Craig Martell, James Migletz, Beth Rosenberg and Adam Slagell all provided important feedback. Drafts of this paper were reviewed by Associate Professor Paul Ohm at University of Colorado Law School and by James M. Miller, an SAIC senior policy, program and legal analyst.

The author has worked with IRBs at MIT, Harvard and NPS; the author is currently a member of the NPS IRB.

References

- [1] Deb Aronson. *Improving the System for Protecting Human Subjects: Conteracting IRB “Mission Creep”*. Technical report, College of Law, College of Liberal Arts and Sciences, Office of the Vice Chancellor for Research, November 2005. <http://www.law.uiuc.edu/conferences/whitepaper/whitepaper.pdf>.
- [2] The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. *Ethical Principles and Guidelines for the protection of human subjects of research*, April 18 1979.
- [3] *Research Compliance: A Faculty Handbook*, 2008. Accessed January 2008, <http://rac.berkeley.edu/compliancebook/humansubjects.html>.
- [4] *Frequently Asked Questions*, January 2008. Accessed January 2008, http://www.capella.edu/portal/alumni/scontent/resources/research_center/FAQ.aspx.
- [5] *Does your study need IRB review?*, 2007. Accessed January 2008, <http://www.ssa.uchicago.edu/ssairb/doesyourstudy.shtml>.
- [6] Jeffrey Cohen. *IRB Expertise*, 2 October 2007. <http://hrpp.blogspot.com/2007/10/irb-expertise.html>.
- [7] *American Ethnologist*, Virginia R. Dominguez (Ed.). <http://www.aesonline.org/ae/334>.
- [8] MIT Committee on the Uses of Humans as Experimental Subjects. *COUHES Overview*, 2003. <http://web.mit.edu/committees/couhes/index.shtml>.
- [9] *Frequently Asked Questions*, January 2008. Georgia State University. Accessed January 2008, <http://www.gsu.edu/research/irb-faq.html>.
- [10] C. K. Gunsalus, Edward M. Bruner, Nicholas C. Burbules, Leon Dash *et al.* *Mission Creep in the IRB World*. *Science*, volume 312(5779):(2006) page 1441. <http://www.sciencemag.org/cgi/content/summary/312/5779/1441>.
- [11] *Approvals and Certifications*, 2008. http://vpf-web.harvard.edu/osr/proposal/prop_pr_certifications.shtml.
- [12] Office for Human Research Protections. *International Compiation of Human Research Protections*, 2008. <http://www.hhs.gov/ohrp/international/HSPCompilation.pdf>.
- [13] United States Department of Health and Human Services. *Federal Policy for the Protection of Human Subjects*, 2008. Accessed January 2008.
- [14] Jack Katz. *Ethical Escape Routes for Underground Ethnographers*, 2008. Unpublished draft, <http://www.sscnet.ucla.edu/soc/faculty/katz/pubs/UndergroundEthnographersDraft.pdf>.
- [15] *Student Research*, May 2007. UCSF School of Medicine. <http://medschool.ucsf.edu/studentresearch/irbpolicy.aspx>.