

A Refined Ethical Impact Assessment Tool and a Case Study of its Application

Michael Bailey¹, Erin Kenneally², and David Dittrich³

¹ Computer Science and Engineering, University of Michigan

² Cooperative Association for Internet Data Analysis, Univ. of California, San Diego

³ Applied Physics Laboratory, University of Washington

Abstract. Research of or involving Information and Communications Technology (ICT) presents a wide variety of ethical challenges and the relative immaturity of ethical decision making in the ICT research community has prompted calls for additional research and guidance. The Menlo report, a revisiting of the seminal Belmont report, seeks to bring clarity to this arena by articulating a basic set of ethical principles for ICT research. However the gap between such principles and actionable guidance for the ethical conduct of ICT research is large. In previous work we sought to bridge this gap through the construction of an ethical impact assessment (EIA) tool that provided a set of guiding questions to help researchers understand how to apply the Menlo principles. While a useful tool, experiences in the intervening years have caused us to rethink and expand the EIA. In this paper we: (i) discuss the various challenges encountered in applying the original EIA, (ii) present a new EIA framework that represents our evolved understanding, and (iii) retrospectively apply this EIA to an ethically challenging, original study in ICTR.

1 Introduction

Information communication technology research (ICTR) presents a wide variety of ethical challenges, touching on diverse research topics including botnets, spam, malware, phishing, etc. Examples of interesting ethical questions raised by such studies include: If someone has the ability to take control of a botnet, can they just clean up all the infected hosts? What risks do researchers face when they provide data to the community? How do theoretical exploits and concepts differ from existing vulnerabilities? What impact does the immediacy of an event (e.g., DDoS) have on our response to the event? [4] Unfortunately, the relative immaturity of ethical decision making and a lack of community standards has prompted calls for additional research and guidance [5].

1.1 ICTR

Before delving deeply into the above challenges, it is first instructive to briefly discuss ICTR, its goals and potential risks. Information cannot be separated

from the systems in which it is stored, processed, or through which it is transmitted. The umbrella term *Information and Communication Technology (ICT)* encompasses these systems, and implicitly the information (or data) that they store, transmit, and process. Research involving ICT often involves risks centered around the core properties of these systems information – confidentiality, integrity, and availability.

Harm that results from impacts on these properties can manifest in physical, psychological, legal, social, and economic damage. These non-informational risks are typically viewed in light of historical behavioral and biomedical research that involve physical procedures that can cause physical pain, bodily harm, or psychological traumas. Informational risks derive from inappropriate use or disclosure of information, which could be harmful to the study subjects or groups. Both categories of harm must be dealt with in ethical evaluation of research involving ICT, spread across all potentially affected stakeholder populations.

When research focuses primarily on ICT itself, indirect harm (either informational or non-informational) to humans can still occur. As ICT evolves and is more tightly integrated into our lives through process controls and cyber-physical systems such as automobile braking controls, smart energy meters, and embedded medical devices, the use and disclosure risks to ICT will increasingly put humans at risk. This necessitates shift from considering research in terms of human subjects involvement to that of human-harming potential [1].

1.2 The Menlo Report

The Menlo report [6], a revisiting of the seminal Belmont report [8], seeks to bring clarity to this arena by articulating a basic set of ethical principles for ICTR. The effort is the result of an interdisciplinary working group sponsored by DHS which commenced in mid-2009. The goal of this effort was to create an updated Belmont report for the field of ICTR. The report appeared for comment in the Federal Register at the end of 2011.

1.3 The EIA v1.0 and its Limitations

While the Menlo report describes fundamental principles, the gap between such principles and actionable guidance for the conduct of ICTR is large. In previous work we sought to bridge this gap through the construction of an ethical impact assessment tool (EIA) [10] we will refer to as EIA v1.0. The EIA v1.0 provided a set of guiding questions to help researchers understand how to apply the Menlo principles. While a useful tool, experiences in the intervening years have caused us to rethink and expand the EIA. Specifically, we believe the EIA v1.0 was successful in achieving its goal of *education*, highlighting the specific classes of ethical problems that need to be addressed. However, in spending the intervening years applying the EIA v1.0 ourselves to both our own work and numerous case studies of others work in the field, we feel two alternative goals now warrant attention. Specifically, those of *Consistency* and *Lowering Barriers to Use*.

An EIA that has a *Low Barrier to Use* will make it easier for researchers to use reasoning by analogy, to trend classes of ethical issues, to assure fairness, etc. It must be easy to use and map, in an understandable way, to existing processes and methodologies. In achieving *Consistency* in ethical analysis, researchers will be better suited to develop ethically defensible research protocols from the start, and others will have an easier time evaluating these protocols because of the clarity and consistency with which researchers describe which humans may be at risk, to what extent, and what protective measures researchers have implemented.

The EIA v2.0 we present here embodies the lessons we have learned to date and uses a least common denominator set of stakeholders that we believe makes it suitable for the majority of ICT research of minimal or low-to-medium risk. We wish to be clear that there are some research situations presenting higher risk, such as vulnerability research involving threat to life or real property, or large-scale computer crime situations, where even the EIA v2.0 may not be sufficiently fine grained or comprehensive to address all stakeholders listed in in Table 1, or all case studies documented in relation to the Menlo Report [7].

2 Ethical Impact Assessment (EIA)

In this section we present the EIA v2.0 framework, with special attention to places where it has been expanded or modified from v1.0 as our understanding has evolved.

2.1 Research Lifecycle

One common experience analyzing case studies using the EIA v1.0 framework was that we consistently repeated classes of risk in our analysis. In many cases these similarities were more an artifact of the phase of research, rather than the research methodology itself.

While we find that we are mostly concerned with experimental computer science, theoretical computer science can also pose risks to humans. In experimental computer science, “[t]he key ideas [are] an apparatus to be measured, a hypothesis to be tested, and systematic analysis of the data (to see whether it supports the hypothesis).” [2]. In such studies, we have robust models for thinking about the lifecycle of data (i.e., collection, use, dissemination) [12]. Explicitly examining the data lifecycle, it is evident that the ethical concerns differ by phase and that concerns repeat across studies in various classes.

In the EIA v2.0 framework, three activities are commonly called out: the *collection* of information (i.e, research data), the *use* of information or information systems in research (whether as vehicle for conducting research or as research subject), and the *disclosure* of research data or vulnerability information that could be used to cause harm. In this paper, we use these terms in a broad sense and emphasize that risks from information collection, use, and disclosure are transitive across stakeholder populations. Risk is present even when

the only data involved are facts and observations about the functioning of a cyber-physical device, and in cases when there is no information involved at all yet harm could arise from unintended consequences resulting from the manipulation of information systems that humans are dependent upon. This latter area is the hardest to evaluate with the EIA v2.0 framework as the focus on the data lifecycle does not cleanly accommodate all potential stakeholder populations, nor those risks that are not data related. We believe that continued evolution of the EIA into a richer and finer-grained framework will further enhance its consistency of evaluation and further lower the barriers to use.

2.2 Stakeholders Analysis

One of the major changes in the EIA framework since v1.0 [10] is the integration of a set of stakeholders as columns in the EIA spreadsheet.

Stakeholder Analysis identifies the key players in the situation in terms of their interests, involvement, and their relationship (i.e., producer or recipient) of outcomes such as benefit or harm. In previous case studies [3] we have adapted the definitions of stakeholders [11] used in other domains for ethical analysis. We also have found that some ICT research, such as studies of botnets and other ongoing computer crime activity, or vulnerability research where publication of research results could be used by malicious actors to cause grievous damage, require consideration of both *Positively Inclined* and *Negatively Inclined* stakeholders in order to fully understand the risk vs. benefit calculus over time [1]. These stakeholders are listed in Table 1.

The problems we seek to address through a comprehensive stakeholder analysis are *indirect harm* and *consideration of intermediaries*. Indirect harm may result from secondary effects, such as disrupting a service provider, which in turn affects the customers of that service provider and the customers of those customers (i.e. in a wholesale vs. retail sales relationship). Or it can be harm that occurs long after publication of vulnerability information as attackers make use of the information for criminal gain before system owners learn of patches and apply them to render services immune to attack. The complexity resulting from the involvement of ICT makes it hard to see what the impacts of ones actions may be. Enumerating the stakeholders helps elucidate the potential harms and benefits. We also find that there are a common set of re-occurring stakeholders, which is reflected in the EIA, however we acknowledge that the full range of Positively and Negatively Inclined stakeholders as depicted in Table 1 must be dealt with effectively in future iterations of the EIA framework.

2.3 Ethical Principles and Their Application

The EIA v1.0 framework was invented at a time when the Menlo report was still in its infancy and well before we had external feedback from reviewers of the document. In the interim, the Menlo Report has matured [6] and the EIA v2.0 framework has been modified to align with the current set of principles and their applications. These include:

Stakeholder Type	Positively Inclined	Negatively Inclined
Key [<i>Affect on producing outcome</i>]	Researchers Programmers Operations Staff Executives Law Enforcement	Criminals (Individuals/Gangs) Malware Programmers Botmasters Criminal Masterminds
Primary [<i>End users</i>]	Consumers (product/service) Enterprises (.edu, .com, .org) Manufacturers Government entities	Espionage Consumers Criminal Enterprises
Secondary [<i>Intermediaries in delivery</i>]	Service Providers Platform Providers Transit Providers Retailers	“Bullet Proof” Hosting Providers Malware Delivery Providers Malware Obfuscators Sellers of fake goods

Table 1. A complete breakdown of stakeholders for a Botnet research scenario. While both Positively and Negatively Inclined stakeholders are shown here, most ICT research involves neither criminal activity nor vulnerability disclosure and would thus not involve the Negatively Inclined Stakeholders.

- **Identification of Stakeholders** As research targeting or involving ICT can hide potentially harmed humans, a thorough analysis of stakeholders is a necessary pre-requisite to a comprehensive analysis of risks, benefits, identification of burdens, and mitigation of actualized harms.
- **Informed Consent** Researchers should obtain informed consent to collect, use or disclose data, or to interact with systems in ways that could have a negative impact on those systems.
- **Harms** Researchers should consider the full spectrum of harms to both persons and information systems (systems assurance, privacy, reputation, physical, psychological, economic)
- **Benefits** Researchers should identify benefits to all stakeholder populations, including (but not limited to) benefits to the broader society.
- **Balancing Risks and Benefits** Research should be designed and conducted not simply to maximize benefits and minimize harms, but to appropriately balance risk and benefits across all stakeholder populations.
- **Mitigation controls** Researchers should notify appropriate parties if research causes harm and have plans in place to efficiently and effectively resolve problems.
- **Fairness and Equity** The benefits and burdens of research should be apportioned fairly across all stakeholder populations.
- **Compliance** researchers should perform due diligence in regards to respecting laws, contracts, etc. in order to protect individuals and organizations.
- **Transparency and Accountability** Researchers should act in ways that garner trust with the general public by communicating intent, research methodology, risk-benefit analysis, and ethical reasoning.

2.4 Bringing it Together: The EIA

The EIA v2.0 framework (see Figure 1) assists researchers in formulating policies, processes, and methodologies that align with ethical principles throughout three research lifecycle phases. It illuminates all relevant ICT stakeholders, as well as both the benefits and human-harming risk potential of research in order to achieve ethically-defensible methodologies and results. A downloadable version is available at <http://www.eecs.umich.edu/~mibailey/EIA.xlsx>

3 Case Study

We illustrate the evaluative use of the EIA v2.0 framework by retrospectively applying it to a case study that provoked ethical debate within the research community. The Menlo Report and the EIA did not exist at the time, so use of the principles and assessment framework during the fundamental research design, implementation and publication was not possible. The researchers in this case study were advised by one of this paper’s authors, who was also substantially involved in the then-parallel Menlo effort. These deliberations influenced the EIA v1.0 framework and the subsequent evolution of both the Menlo Report and EIA framework. The post hoc analysis performed here exposes opportunities where researchers could have made more ethically-defensible decisions.

3.1 Background

Researchers at University of California San Diego (UCSD) undertook an

Research Lifecycle	Ethical Principles Considered	Application of Principles	Stakeholders					
			ICT Researchers	Data Subject / End User	Human Subjects Network / Platform / Service Provider	Malicious Actors	Society	Gov't / Law Enforcement
Research Collection	Respect for Persons	Uninformed Consent						
	Beneficence	Benefits						
	Mitigation of Realized Harms	Justice						
Research Use / Management	Respect for Law and Public Interest	Justice						
	Respect for Persons	Beneficence						
	Respect for Law and Public Interest	Justice						
Research Disclosure	Respect for Law and Public Interest	Justice						
	Respect for Persons	Beneficence						
	Respect for Law and Public Interest	Justice						

Fig. 1. The EIA worksheet

experiment to measure the conversion rate of unsolicited commercial e-mail as part of an empirical study to understand the quantitative value proposition of spam [9]. Lacking sufficient methods to indirectly measure spam conversion, the methodological challenges stemmed largely from the ethical implications of mimicking real spam campaigns. Specifically, key components of such operations involved building fake e-commerce sites, marketing them via spam, presenting sales transactions for the advertised goods, and distributing the various communications (e-mail marketing, processing recipient responses) via illicit botnets.

To address the obvious ethical and legal problems posed by spamming and botnet activities, researchers sought insight from *non-maleficence* theory⁴ and legal and ethical advisement. This guidance informed the research methodology which involved parasitically infiltrating the command and control infrastructure of an existing spamming botnet by accepting invitations to become proxy bots, or conduits between master servers and worker bots. Researchers then modified a subset of the spam the botnet was already distributing, so respondent users were directed to servers under researcher control, not those of the real spammer. Then researcher servers presented web sites that mimicked those actually hosted by the spammer, however they “de-fanged” them by removing functionality designed to compromise the user’s system or that would collect and disclose sensitive user information (e.g., name, address, credit card data).

3.2 Stakeholder Identification

ICT Researchers In addition to the obvious inclusion of the UCSD research team, it became clear that other researchers were also analyzing the same botnet. The “in vivo” nature of botnet studies warrants consideration of these other stakeholders who may be simultaneously undertaking various empirical studies.

Data Subject or End User Stakeholders here were the users of computers infected with the Storm bot (a.k.a., worker machines), and recipients of spam email sent through the botnet. This research impacted the collective rights and interests of not only the owners and users of computers that were infected with the Storm bot, but those being tricked by it.

Network, Platform, or Service Provider Parties to be considered here were network services providers for the botnet proxy hosts and command and control servers, Internet service providers (ISPs) of users with infected computers, webmail platform providers, registrars of mimicked illicit phishing sites, and the network community (the Overnet peer-to-peer platform) used by the botnet to communicate.

Society Beyond those directly affected by botnet infection, this research impacted the collective rights and interests of all users of computers that are affected by social engineering attacks involving spam and online fraud activities.

Government or Law Enforcement As the primary source of funding for the research, the National Science Foundation provided authoritative influence

⁴ Researchers should act in good faith and control risks, exposing end users to no more harm than they would face but for the research activities

and is thus a stakeholder. Similar to the rationale for considering other bot researchers, the research had the potential to impact law enforcement agencies (LEAs) in multiple countries who were investigating and attempting to enforce various laws against the parties responsible for the botnet's illegal activities.

3.3 Research Collection

Consent – Informed consent was obtained from the network provider for the proxy collector machines, the webmail platform providers, and the domain registrar for the researcher's mimicked phishing sites. Each had an interest in safeguarding the ICT resources it owned, controlled or managed, including the data associated with those resources. The researchers believed they could justify a waiver of informed consent from owners of worker machines and end user subjects of the research. Identifying and providing notice to the owners of thousands of compromised home computers was impracticable, given the scale and scope of the botnet. Informing both worker hosts and end user stakeholders about the research procedure, purpose, risk-benefit analysis, and withdrawal opportunities would negatively impact the scientific integrity of the research by altering the behavior that was attempted to be studied. A determination on waiver of informed consent due to impact on research integrity is often the responsibility of an IRB, not a researcher decision. A Menlo evaluation using the EIA framework raises questions about whether researchers should have debriefed end users who were deceived via the phished sites (fake pharmaceutical and e-card) via some form of pop-up alert.

Compliance – Legal due diligence analysis was performed to address a number of factors. Research activities respected federal and state laws concerning computer fraud (e.g., no unauthorized access to systems or networks, researcher proxy bots were invited to participate in botnet, researchers were authorized to log traffic to their own fake phish website, no exceeding access to webmail platform since Terms of Service were not violated, research action did not cause legally cognizable damage or harm), electronic communications privacy (e.g., no interception of traffic; proxy bots were a party to the communications, although there was possible violation if acquisition of bot communications would be deemed to require two-party consent), intellectual property (e.g., mimicked phished sites did not replicate the images that infringed copyright on the real phished sites, no circumvention of mediating devices), or contract laws (e.g., there were no agreements associated with nodes in the Overnet platform; researcher actions adhered to normal and expected functioning of Overnet protocols; use of webmail did not violate Terms of Service prohibiting sending of spam since those accounts were receiving users' responses to redirects). While researchers did engage ex ante ethical and legal risk analysis, federal regulation required that they should have consulted with their IRB prior to, rather than after, the completed research.

Harms – Researcher actions (i.e., botnet command rewriting, interposing Spam delivery, interposing user click-through) did not diminish the performance, availability or integrity of the networks or machines in the bot infrastructure.

There were no new machines compromised or worker bots created, nor did researchers cause corrective action to be undertaken by systems administrators. Privacy harms were avoided by not collecting, storing or transmitting any private personal information from either worker systems with whom the researcher proxy hosts communicated or from the mimicked sites. There was no reason for the researchers to believe that the study was interfering with LE investigation activities involving the botnet. Researchers minimized potential reputational harm to Webmail providers from spam-advertised product association by obtaining informed consent. With the fake e-card phished sites, researchers presented a benign executable that performed a simple *HTTP POST* to the researcher controlled backend server, and then exited. This could be interpreted as direct intervention with the environment of subjects who have not consented, however the potential for harm here was strictly minimized and there was no malicious intent.

Benefits Considered – This research aimed to enhance understanding of internet criminal activity and thus produces benefits to the broader society by improving user’s abilities to safely use ICT in their daily lives.

Mitigation – Researchers mitigated any harm to integrity or functionality of user’s systems from the botnet-directed spam by redirecting them to de-fanged fake phished site, only logging the user-agent string to determine if the exploit would have likely worked. The users were always asked to download the file, but where not actually provided with an executable (e.g., presented a 404 error).

3.4 Research Use or Management

Consent – The webmail and network provider’s consent to collect information for specific research activities extended to the ongoing use of those platforms for the limited duration of the experiment.

Compliance – Researchers designed their methodology to avoid running afoul of consumer protection laws (e.g., prohibiting the sending of commercial e-mail). Researchers acquiesced to being infected by the botnet and subsequently interposed as proxy bots within the existing bot infrastructure. This positioned researchers as a conduit, passively transmitting and observing the spam-related commands and data between the master servers which initiated and controlled the transmission of spam and the worker bots which carried out the directives. Actions that altered command messages (spam template, dictionary entries) to include researcher-controlled sites arguably did not alter the spam liability evaluation since the primary purpose of the deception employed by researchers was not related to advertising or promoting a commercial product or service, but rather, to study users’ susceptibility to engage these campaigns. Measurements associated with fake phishing sites respected intellectual property rights of legitimate brand owners by not replicating known trademarked or copyrighted material from the legitimate sites. In the event that the cloned phished sites (e-card and pharmacy sites) did include protected intellectual property unbeknownst to researchers, they were well-positioned to exercise a “fair use” defense. As with

collection, researchers should have obtained IRB approval prior to engaging in research.

Harms – Researcher’s actions did not expose end users to more harm than they would face but for the research activities, and steps were taken to reduce harm from the Storm bot. The probability and magnitude of any harm or discomfort anticipated in the research was not greater than that ordinarily encountered by users in normal use of the Internet. The only sensitive data retained was internet protocol addresses of worker bots as needed for research measurement, and they were discarded immediately after statistics were collated. The measurement infrastructure did not create new qualitative or quantitative harm to other protected computer systems – absent researcher involvement, the same users would have received the same spam e-mails from the same worker bots. Researcher proxies were passive actors that did not initiate the transmission spam e-mail, compromise hosts, or contact worker bots asynchronously. The modification of messages strictly reduced harm to users who followed the embedded links. Additional burden was not placed on hosting network resources. Research proxy nodes did not transmit or distribute any illicit information or program, send e-mail, mount or participate in denial of service attacks, crawl for or scrape e-mail addresses, compromise or otherwise introduce user accounts, or interfere with the ability of users systems to protect themselves or use the network. Researcher nodes acted in accordance with expected P2P infrastructure functions, including respecting communications protocols that maintained topological consistency with the rest of the infrastructure, and receiving and forwarding commands.

Foreseeable harms related to legitimate intellectual property rights holders were addressed in several ways. Researchers did not duplicate the phished sites that were copies of legitimate websites stolen by scraping (i.e., cloning or copying the text, logos, artwork or design templates). Rather, they replicated the general look and feel. Legitimate domain names were not spoofed, forged, or otherwise hijacked. To avert trademark likelihood of confusion harms, researchers did not obtain economic or commercial benefit, nor were not unjustly enriched by mocking the legitimate website design.

Benefits Considered – Research management and use of the measurement infrastructure provided empirical knowledge of end user susceptibility to spam marketing campaigns, botnet structure and function, and un-quantified behavior underlying the spam value proposition. Collateral individual user benefits included thwarting visits to malware-infected phishing sites and further communications with botnet command channel.

Mitigation controls – Researchers were sensitive to possible interruption of network services from retaliatory denial of service against the network hosting the proxy bots and were prepared to discontinue their utilization if that harm manifested.

3.5 Research Disclosure

Harms – Researchers did not disclose any sensitive individual or organizational information, including the internet addresses of infected worker machines or confidential network data. This was done to prevent foreseeable harms to privacy, reputation, and systems assurance associated with botnet victimization and vulnerability. Any relatively small burden borne by recipients of spam was balanced against the larger benefit to society from performing beneficial research. Researchers could have been more mindful of risks to themselves as a stakeholder class, specifically pertaining to probable reputation harms from not adequately disclosing their efforts related to ethics considerations in the design and execution of their research.

Benefits Considered – In addition to previously mentioned benefits, disclosure of research results could enhance understanding of the structure and function of digital criminal enterprises in the interests of law enforcement investigations, take-downs, and prosecutions.

Mitigation controls – While researchers did not have actual and specific knowledge of LE or other research involvement in the botnet study, there was no overt effort made to avoid collision.

Fairness & Equity – The selection and targeting of end user subjects and owners of worker machines was outside researcher control. Similarly, selection of network and application providers was likely a function of the Overnet network.

Transparency – Although the ethical controls were implicit in research design, researchers did not explicitly disclose details about the plethora of ethical considerations that informed their research. While researchers did offer a high level description of ethical undertakings, the EIA suggests that transparency and accountability could have been strengthened by more granular, a priori disclosure of the methodology and results in various publicly-available conference publications and presentations. However, unless conference committees make accommodations in paper length limitations, researchers will be de incentivized from elucidating ethical considerations in their published work.

4 Conclusion

We have described the second iteration of an ethical impact assessment framework that operationalizes the application of principles described in the Menlo Report. We are continuing to evolve this framework and other tools for the ethically-justifiable design and assessment of research involving ICT. It reflects the iterations and refined collaborative thoughts that occurred between the chosen case study and this paper. We are continuing to improve this tool so that it most effectively assists in ethical design and assessment of research involving ICT that carries a probable risk for human harming activities.

References

1. Katherine Carpenter and David Dittrich. Bridging the Distance: Removing the Technology Buffer and Seeking Consistent Ethical Analysis in Computer Security Research. In *1st International Digital Ethics Symposium*. Loyola University Chicago Center for Digital Ethics and Policy, 2011.
2. Peter J. Denning. ACM President's Letter: What is experimental computer science? *Commun. ACM*, 23:543–544, October 1980.
3. David Dittrich, Michael Bailey, and Sven Dietrich. Have we Crossed the Line? The Growing Ethical Debate in Modern Computer Security Research. In *(Poster at) Proceedings of the 16th ACM Conference on Computer and Communication Security (CCS '09)*, Chicago, Illinois, USA, November 2009.
4. David Dittrich, Michael Bailey, and Sven Dietrich. Towards Community Standards for Ethical Behavior in Computer Security Research. Technical Report 2009-01, Stevens Institute of Technology, Hoboken, NJ, USA, April 2009.
5. David Dittrich, Michael Bailey, and Sven Dietrich. Building an Active Computer Security Ethics Community. *IEEE Security and Privacy*, 9(4):32–40, 2011.
6. David Dittrich and Erin Kenneally (eds.). The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. <http://www.cyber.st.dhs.gov/wp-content/uploads/2011/12/MenloPrinciplesCORE-20110915-r560.pdf>.
7. David Dittrich and Erin Kenneally (eds.). Applying Ethical Principles to Information and Communication Technology Research: A Companion to the Department of Homeland Security Menlo Report, January 2012.
8. The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research National. The Belmont Report - Ethical Principles and Guidelines for the protection of human subjects of research, 1978. U.S. Government Printing Office. DHEW Publication No. (OS) 78-0008. Reprinted in Federal Register 44 (April 18, 1979): 23192.
9. Chris Kanich, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M. Voelker, Vern Paxson, and Stefan Savage. Spamalytics: an empirical analysis of spam marketing conversion. In *CCS '08: Proceedings of the 15th ACM conference on Computer and communications security*, pages 3–14, 2008.
10. Erin Kenneally, Michael Bailey, and Douglas Maughan. A Tool for Understanding and Applying Ethical Principles in Network and Security Research. In *Workshop on Ethics in Computer Security Research (WECSR '10)*, Tenerife, Canary Islands, Spain, January 2010.
11. Stella Mascarenhas-Keyes. Ethical Dilemmas in Professional Practice in Anthropology. <http://www.theasa.org/networks/apply/ethics/analysis/stakeholder.htm>, July 2008.
12. Mary Vardigan, Pascal Heus, and Wendy Thomas. Data Documentation Initiative: Toward a Standard for the Social Sciences. *International Journal of Digital Curation*, 3:107–113, 2008.