

The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research

September 15, 2011

Executive Summary

This report proposes a framework for ethical guidelines for computer and information security research, based on the principles set forth in the 1979 Belmont Report, a seminal guide for ethical research in the biomedical and behavioral sciences. Despite its age, the Belmont Report's insightful abstraction renders it a valuable cornerstone for other domains. We describe how the three principles in the Belmont report can be usefully applied in fields related to research about or involving *information and communication technology*. ICT research raises new challenges resulting from interactions between humans and communications technologies. In particular, today's ICT research contexts contend with ubiquitously connected network environments, overlaid with varied, often discordant legal regimes and social norms. We illustrate the application of these principles to information systems security research – a critical infrastructure priority with broad impact and demonstrated potential for widespread harm – although we expect the proposed framework to be relevant to other disciplines, including those targeted by the Belmont report but now now operating in more complex and interconnected contexts.

We first outline the scope and motivation for this document, including a historical summary of the conceptual framework for traditional human subjects research, and the landscape of ICT research stakeholders. We review four core ethical principles, the three from the Belmont Report (Respect for Persons, Beneficence, and Justice) and an additional principle *Respect for Law and Public Interest*. We propose standard methods to operationalize these principles in the domain of research involving information and communication technology: identification of stakeholders and informed consent; balancing risks and benefits; fairness and equity; and compliance, transparency and accountability, respectively. We also describe how these principles and applications can be supported through assistive external oversight by ethical review boards, and internal self-evaluation tools such as an Ethical Impact Assessment.

The intent of this report is to help clarify how the characteristics of ICT raise new potential for harm and to show how a reinterpretation of ethical principles and their application can lay the groundwork for ethically defensible research.

Working Group Participants

This report is the product of a series of workshops and meetings held over a period of sixteen months. The participants at these meetings are listed alphabetically below. In addition, the authors thank the dozen or so ICTR community members whose feedback was invaluable to assuring that this document reflects the ground truth sentiments of the professionals at the front lines of ICT research ethics.

- Michael Bailey, University of Michigan
- Aaron Burstein, University of California Berkeley
- KC Claffy, CAIDA, University of California San Diego
- Shari Clayman, DHS Science & Technology
- David Dittrich, Co-Lead Author, University of Washington
- John Heidemann, University of California, ISI
- Erin Kenneally, Co-Lead Author, CAIDA, University of California San Diego
- Douglas Maughan, DHS Science & Technology
- Jenny McNeill, SRI International
- Peter Neumann, SRI International
- Charlotte Scheper, RTI International
- Lee Tien, Electronic Frontier Foundation
- Christos Papadopoulos, Colorado State University
- Wendy Visscher, RTI International
- Jody Westby, Global Cyber Risk, LLC

Contents

A	Introduction – Focus and Motivations	5
A.1	Who is the Target Audience for this Report?	5
A.2	Historical Context	6
B	Restatement of Belmont Principles in the ICTR Context	7
C	Application of the Principles	7
C.1	Stakeholder Perspectives and Considerations	8
C.2	Respect for Persons	9
C.2.1	Informed Consent	9
C.3	Beneficence	10
C.3.1	Identification of Potential Benefits and Harms	10
C.3.2	Balancing Risks and Benefits	11
C.3.3	Mitigation of Realized Harms	12
C.4	Justice: Fairness and Equity	12
C.5	Respect for Law and Public Interest	13
C.5.1	Compliance	13
C.5.2	Transparency and Accountability	14
D	Implementing the Principles and Applications	14

A Introduction – Focus and Motivations

This report attempts to summarize a set of basic principles to guide the identification and resolution of ethical problems arising in research of or involving *information and communication technology* (ICT)¹. ICT is a general umbrella term that encompasses networks, hardware and software technologies that involve information communications pertaining to or impacting individuals and organizations. ICT has increasingly become integrated into our individual and collective daily lives, mediating our behaviors and communications and presenting new tensions that challenge the applications of these guiding principles.

ICT research (ICTR) involves the the collection, use and disclosure of information and/or interaction with this ubiquitously connected network context which is overlaid with varied, often discordant legal regimes and social norms. The challenge of evaluating the ethical issues in ICTR stems in large part from the attributes of ICT: scale, speed, tight coupling, decentralization and wide distribution, and opacity. This environment complicates achieving ethically defensible research for several reasons. It results in interactions with humans that are often indirect, stemming from an increase in either logical or physical “distance” between researcher and humans to be protected over research involving direct intervention. The relative ease in engaging multitudes of distributed human subjects (or data about them) through intermediating systems speeds the potential for harms to arise, and extends the range of stakeholders who may be impacted. Also, legal restrictions and requirements have expanded considerably since the 1980s, and ICTR is unquestionably subject to a variety of laws and regulations that address data collection and use. While it is true that these individual complications are shared by traditional biomedical and behavioral research, this report seeks to manage the tension resulting from the simultaneous confluence of these complicating factors that occur with regularity in ICTR.

There is a need to interpret and extend the traditional ethical framework to enable ICT researchers and oversight entities to appropriately and consistently assess and render ethically defensible research². Such a framework should also support current and potential institutional mechanisms that are well served to implement it, such as a research ethics board (REB). We build on the foundation set by the *Belmont Report*, which articulates three fundamental ethical principles and guiding applications of these principles for protecting human subjects of biomedical and behavioral research: respecting persons; balancing potential benefits and harms; and equitably apportioning benefits and burdens across research subjects and society³. The guidelines in this report are applicable to research that has the potential to harm humans, regardless of whether those humans are the direct research subjects or are indirectly at risk of harm from interactions with ICT. This report explains how the traditional framework fits within the context of the computer science sub-discipline of information security research. Specifically, this domain addresses ICT vulnerabilities, digital crime, and information assurance for critical infrastructure systems. These are areas where harms are not well understood yet are potentially significant in scope and impact. The framework proposed herein is germane to other disciplines that involve the use of ICT, including those targeted by the Belmont Report that now operate in ICT contexts.

A.1 Who is the Target Audience for this Report?

This report offers guidance primarily for ICT researchers (including academic, corporate, and independent researchers), professional societies, publication review committees, and funding agencies. Secondarily, this report aims to assist those who administer and apply these princi-

ples, such as oversight authorities (e.g., REBs), policy makers, attorneys, and others who shape and implement human subject protection policies and procedures.

This report does not recommend particular enforcement mechanisms, which are more appropriately relegated to institutional oversight processes. To the extent that enforcement of ethical practices is inconsistent across and within academic and non-academic ICTR, we intend this report to improve consistency in ethical analyses and self-regulation for both individuals and organizations striving toward ethically defensible research.

A.2 Historical Context

Despite a long history of well-publicized abuses, it took over a decade for the ethical standards prescribed in the Belmont Report to first be defined in the Code of Federal Regulations (CFR). Language from 45 CFR 46, which covers biomedical and behavioral research funded by the Department of Health and Human Services (HHS), was later adopted by all executive branch departments in what is known as the *Common Rule*⁴. It ushered in a government-wide requirement for REB oversight of research protocols to protect human research subjects. Prior to this point, there was no regulated oversight mechanism and biomedical and behavioral researchers relied on subjective, ad hoc, and inconsistent ethical compasses to guide their decision making.

In parallel during the 1970s, a U.S. Defense Advanced Research Projects Agency (DARPA) project was designing and implementing a communications architecture to support cooperative time-sharing of computational resources across large government-funded laboratories. Although this network architecture would eventually evolve into the global Internet, the community at the time was small, trusted, non-commercial, and research-oriented. The resulting testbed did not create or make available databases with millions of personally-identifying records, accessible globally to records globally available to anyone, nor did it support critical services or communications. Unsurprisingly, early ICT research evolved without significant concern for human subjects, leading to instances where ethical considerations were either absent or misapplied because researchers failed to understand their relevance, or lacked any standards for assessment, accountability, or oversight. Cases include interactive studies of malicious software and platforms, engagement in active counterattack measures, exploitation and disclosure of systems vulnerabilities, and collection and sharing of sensitive information. The demonstrated potential for harm in ICTR illustrates the need to re-conceptualize the traditional *human subject protection* paradigm that underpins ethical oversight in other fields.

The Common Rule was intended to protect persons who might be harmed from research, not simply those who are actively participating in research. Specifically, *human subject* means, “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” (45 CFR 46.102(f)). A key trigger for invoking ethics board review under the Common Rule is whether there is “intervention” or “identifiable private information”. How these measures are construed in ICTR is significant since millions of humans could be engaged indirectly in a given research project via their use of ICT. Even a simple link traffic characterization study could involve millions of computers used by humans who are not themselves the direct subjects of research.

An evolved paradigm, therefore, considers what constitutes *human-harming* rather than simply *human subjects* research when applying ethical principles to protect humans who may be impacted by the research. Examples of potentially human-harming ICT artifacts that researchers may interact with include avatars in online virtual worlds, malware controlling compromised machines, embedded medical devices controlling biological functions, or process

controllers for critical infrastructure. The significant changes brought about by ICT since the commencement of formal regulated research necessitates a reconceptualization of the application of ethical principles for research involving ICT.

B Restatement of Belmont Principles in the ICTR Context

The first three rows of Table 1 summarize the three core principles and their application as outlined in the Belmont Report⁵. We offer an additional principle to guide ethical considerations in ICTR research, listed in the fourth line of Table 1. We call this principle *Respect for Law and Public Interest* because it addresses the expansive and evolving, yet often varied and discordant, legal controls relevant to communication privacy and information assurance (i.e., the confidentiality, availability, and integrity of information and information systems). While respect for the law and public interest is implicit in Belmont’s application of Beneficence, several challenging factors suggest these issues merit explicit consideration in the ICTR context: the myriad laws that may be germane to any given ICTR; conflicts and ambiguities among laws in different geo-political jurisdictions; the difficulty in identifying stakeholders, a necessary prerequisite to enforcing legal obligations; and possible incongruence between law and public interest.

Principle	Application
Respect for Persons	Participation as a research subject is voluntary, and follows from informed consent; Treat individuals as autonomous agents and respect their right to determine their own best interests; Respect individuals who are not targets of research yet are impacted; Individuals with diminished autonomy, who are incapable of deciding for themselves, are entitled to protection.
Beneficence	Do not harm; Maximize probable benefits and minimize probable harms; Systematically assess both risk of harm and benefit.
Justice	Each person deserves equal consideration in how to be treated, and the benefits of research should be fairly distributed according to individual need, effort, societal contribution, and merit; Selection of subjects should be fair, and burdens should be allocated equitably across impacted subjects.
<i>Respect for Law and Public Interest</i>	<i>Engage in legal due diligence; Be transparent in methods and results; Be accountable for actions.</i>

Table 1: Proposed guidelines for ethical assessment of ICT Research.

C Application of the Principles

The challenges of ICTR risk assessment derive from three factors: the researcher-subject relationships, which tend to be disconnected, dispersed, and intermediated by technology; the proliferation of data sources and analytics, which can heighten risk incalculably; and the inherent overlap between research and operations. In order to properly apply any of the principles

listed above in the complex setting of ICT research, it is first necessary to perform a systematic and comprehensive stakeholder analysis.

C.1 Stakeholder Perspectives and Considerations

Stakeholder identification includes consideration of several factors: the degree to which information involved in the research identifies individuals (including their digital identities), groups and organizations and what behaviors, communications, or relationships are associated with such identification. Harms related to exposing the identity of research subjects engaging in sensitive behaviors, communications, or relationships, which they assume to be private, can extend beyond the direct research subject to family, friends or other community relations. While this is also true of some research where the subject is the primary party at risk, in ICTR these harms may often be broader because ICT can amplify both the disclosure as well as the number of stakeholders impacted.

Further, ICTR often involves stakeholders that are non-research entities who rely on information and systems that are involved in the research and who may be harmed by its unavailability or corruption. Groups or organizations (e.g., companies or networks) may warrant different consideration from that of individuals, especially when applying the principles of Beneficence and Justice. Impinging on the privacy of an organization may be ethically justified if it yields substantial social benefit through an increase in cybersecurity, while it may be far less justifiable to violate an individual's privacy interests for the same objective.

ICT Researchers In commercial, academic, and government sectors, ICT researchers have a vested interest in pursuing, sharing, and applying empirically grounded scientific knowledge. Research in economics, network science, security, and social behavior may inform operations, policies, and business models.

Human Subjects, Non-Subjects, and ICT Users Traditional biomedical and behavioral research requires protection of natural persons and certain data that identifies them. In ICTR, the target of research may be an information system or associated data, which complicates the assessment of potential harm to users of that system or data. Primary considerations include the ability to interact with ICT without suffering harms such as disruption of access, loss of privacy, or unreasonable constraints on protected speech or activities. Victims of computer crimes are potential human non-subjects of research.

Malicious Actors A subset of ICTR involves criminal activity or potential exploitation of vulnerabilities in the design or implementation of ICT. Malicious actors can use published research results for nefarious purposes that increases harm more than the intended societal benefits of the research. It is not the role of researchers to judge guilt or innocence, however the existence of malicious stakeholders is not hypothetical and requires that researchers consider how to control disclosure of security knowledge that could lead to harm in the wrong hands.

Network/Platform Owners and Providers Network owners or providers are typically commercial entities who are vested in safeguarding their physical and intellectual property, pursuing innovation and wealth, and building business and customer relationships. They are concerned about obligations associated with such representation. As intermediaries between a research and end users, they may be in a position of authority to serve as proxies for consent on behalf of their customers when it is otherwise impracticable for the researcher to individually obtain informed consent from end users.

Government: Law Enforcement Public law enforcement is mandated to advance criminal justice by protecting individuals and fostering public safety. Law enforcement also has an interest in research that improves its strategic, tactical, or operational efficacy in preventing, investigating, and responding to illegal activities. Examples include countering new and complex criminal ecosystems and instruments of crime such as botnets.

Government: Non-Law Enforcement Local, state, and federal government agencies are responsible for providing public services, protecting the rights of their citizens, and establishing law and policy governing social conduct. Research is an important vehicle through which the government can promote social good and innovation. For example, cybercrime research may enhance understanding of infrastructure risks, online social networks, or economic markets of criminal enterprises; influence the deployment of commercial countermeasure technologies; and inform the interpretation or reform of relevant laws and policies.

Society ICTR implicates the collective rights and interests of owners and users of networks and data to know, influence, and choose how and when to engage with information communications networks and systems. Society benefits from knowledge that improves policies, laws and the administration of justice, and the well-being of the lives of its citizens. Society may likewise be harmed through actions that negatively impact information systems infrastructures, or through the collection, use, or disclosure of information that may assist criminals as much if not more than ICT system developers and operators.

C.2 Respect for Persons

In the Belmont Report, the principle of Respect for Persons reflects two tenets: individuals should be treated as autonomous agents, and persons with diminished autonomy are entitled to protection. This principle has been applied by involving as research subjects only those with sufficient understanding or awareness to provide *informed consent*, or by obtaining informed consent from legally authorized representatives (e.g., parents of minors, relatives of unconscious patients, or guardians of those incapable of deciding for themselves). In the ICTR context, the principle of Respect for Persons includes consideration of the computer systems and data that directly interface, integrate with, or otherwise impact persons who are typically not research subjects themselves.

C.2.1 Informed Consent

Informed consent is a process during which the researcher accurately describes the project and its risks to subjects and they accept the risks and agree to participate or decline. Subjects must be free to withdraw from research participation without negative consequences. Where feasible, researchers should obtain informed consent to collect, use, or disclose sensitive identifying data, or to interact with information systems in ways that could negatively affect those systems or their users. This consent should be obtained from individually identifiable persons with whom a researcher interacts or obtains identifying information and secondary stakeholders who could be impacted by the ICTR (e.g., customers using the systems being studied.)

Researchers should inform subjects that they may not benefit from the research, although society may benefit in the future. Researchers should be mindful that leveraging intended benefits to coerce or entice consent from subjects fails the voluntary participation element of informed consent. Examples include suggesting that research participants will receive improved or enhanced services, or that services will be degraded or withheld if a subject declines

participation in or withdraws from a study. Informed consent for one research purpose or use should not be considered valid for other research purposes. When an individual is identified with a group or organization, individual consent does not imply consent from other members of the group. Finally, informed consent for one research purpose or use should not be considered valid for different research purposes.

There are justifiable reasons why it may be impracticable to obtain informed consent. In ICTR the frequency of this occurring may be greater than in traditional human-centered research. Of the three components in the informed consent process – notice, comprehension, and voluntariness – providing notice may be particularly challenging given the scale and scope of many operational ICTR environments. It may be impracticable, it may not be technically feasible to identify subjects, or it may interfere with scientific integrity of the results.

It may be infeasible to identify, or obtain consent from millions of users whose everyday communication generates traffic across a heavily aggregated backbone link in a traffic modeling study. Or it can be onerous to attempt to inform the owners of hundreds of thousands of compromised home computers that are being used as a single instrument of criminal activity (i.e., a botnet) under study. In cases where it is impracticable to identify and obtain consent from users, an intermediate party with appropriate authority may provide proxy consent. An Internet service provider (ISP) may hold such authority on behalf of its subscribers via explicit agreements or contracted terms of service.

Sometimes informing stakeholders about the research procedure, purpose, risk-benefit analysis, and withdrawal opportunities impacts the scientific integrity of research results. Informing research subjects that some web sites are fake during a research experiment on phishing vulnerabilities could negatively impact the research validity by altering the subject's behavior. Appropriate Respect for Persons in such deception research can typically be achieved by debriefing the subjects after the research is completed. Debriefing is typically required when deception is used in order to mitigate harm resulting from loss of trust in researchers by those subjects who were deceived.

There may be a conflict between satisfying ethical review requirements and separate legal protections. Even though a researcher obtains a waiver of informed consent due to impracticability reasons, this may not remove legal liabilities under communications privacy statutes that prohibit acquiring or disclosing communications without the consent of the communicating parties.

When a researcher believes waiver of informed consent is warranted, he should clearly describe the justification for departing from the principle of consent.

C.3 Beneficence

In the Belmont Report, the Beneficence principle reflects the concept of appropriately balancing probable harm and likelihood of enhanced welfare resulting from the research. Translating this principle to ICTR demands a framework for systematic identification of risks and benefits for a range of stakeholders, diligent analysis of how harms are minimized and benefits are maximized, preemptive planning to mitigate any realized harms, and implementing these evaluations into the research methodology.

C.3.1 Identification of Potential Benefits and Harms

As with traditional human-centered research, ICT researchers should identify benefits and potential harms from the research for all relevant stakeholders, including society as a whole,

based on objective, generally accepted facts or studies. Since communication technologies intermediate so much of our lives, designing, conducting and evaluating ICTR may demand attention to potential societal benefits and harms related to: systems assurance (confidentiality, availability, integrity); individual and organizational privacy; reputation, emotional well-being, or financial sensitivities; and infringement of legal rights (derived from constitution, contract, regulation, or common law). Challenges identifying harms and benefits in ICTR environments include: the scale and rapidity at which risk can manifest; the difficulty of attributing research risks to specific individuals and/or organizations; and our limited understanding of the causal dynamics between the physical and virtual worlds.

Compared to our institutionalized or socially internalized understanding of harm related to physical interactions with human subjects, as a society we are relatively inexperienced regarding qualitative and quantitative assessment of damages and harms in the digital realm. One helpful approach to identifying harms is to review the laws and regulations that apply to an ICTR activity, and analyze the underlying individual and public interests that the research might negatively impact. The inevitable gap between the trajectories of technological progress, legislative process and judicial interpretation limits a purely legalistic approach to identifying harm. Regard for the fundamental purposes of applicable laws can help mitigate these gaps. Because laws may be unclear or open to interpretation, a narrow focus that only considers acts impacting the integrity or availability of information and information systems might overlook a broader range of harms that may not be explicitly protected by law.

C.3.2 Balancing Risks and Benefits

A simplistic interpretation of Beneficence is the maximization of benefits and minimization of harms. This is not to say that all harm must be eliminated completely and every possible benefit must be identified and fully realized. Rather, the researcher should systematically assess risks and benefits across all stakeholders. Researchers should be mindful that risks to subjects are being weighed against the benefit to society, not to either the research subjects or the researchers themselves. Researcher actions should be measured using a standard of a *reasonable researcher*, who exercises the knowledge, skills, attention, and judgment that the community requires of its members to protect their interests and the interests of others.

When ICT is involved, burdens and risks can extend beyond “the human subject,” making the quantification of potential harm more difficult than with direct intervention. It can be difficult to balance risks and benefits with novel research whose value may be speculative or delayed, or whose realized harm may be perceived differently across stakeholders. If there are plausible risks, researchers bear the burden of showing specific, evidence-based consideration that they can manage those risks.

In a direct intervention research scenario, balancing is partially addressed through the informed consent process. Waivers of informed consent and the use of deception in research complicates the balancing process. When informed consent is waived, the researcher typically must show that the study involves minimal risk, that there are valid scientific reasons for the consent waiver, and the subject must be de-briefed at the conclusion. Additionally, seeking to waive post-deception debriefing is a further complication.

Risks must also be considered in light of everyday events that users encounter. This includes events such as programs crashing, malicious software accessing and infecting networked computers, or the exposure of electronic communications. Adhering to a zero-risk standard will severely limit the ability to perform needed research, possibly precluding entirely the study of criminal behavior that can be observed only in a live production environ-

ment. Researchers, therefore, must sometimes take calculated risks. Borrowing from the theory of *nonmaleficence*, researchers should act in good faith and control risks so as to expose stakeholders to no more harm than they would face outside the research setting. A researcher studying live malicious software may need to run the software on his own platform and observe its interactions with criminals controlling it. Even with reasonable measures to detect and reduce potential harm, the malicious software being studied could still accidentally infect other computers. However, regardless of researcher actions, those computers would have been infected when the malicious software propagates at the attacker's direction. Ethically defensible Beneficence lays on a spectrum between unequivocal adherence to averting risk and mirroring the risks posed by malicious actors.

C.3.3 Mitigation of Realized Harms

Despite appropriate precautions and attempts to balance risks and benefits in ICTR, research may cause unintended side effects that harm stakeholders. In anticipation, researchers should consider preempting the escalation of realized harms by notifying affected parties or otherwise engaging mitigation actions. To that end, researchers should develop mitigation procedures and checklists, such as a contact list of parties to notify, if such unintended consequences ensue. Other potential harms that are reasonably foreseeable may have a low probability of occurring, but have a high impact. Researchers should anticipate such worst-case scenarios and make appropriate preparations to respond in a manner and scope that shows due diligence on the part of the researcher. It may be necessary or prudent to involve the researchers' own institutional risk management and oversight authorities and media relations.

ICTR may involve records containing sensitive data about individuals, or potentially cause disruption to millions of computers around the world. ICT researchers must be aware of these harms as not only primary risks, but also secondary, collateral risks (e.g., to customers of primary data subjects or computer owners) and be prepared to responsibly inform affected stakeholders. In many cases, it is impracticable to notify all affected individuals, but it may be feasible to notify service providers or other entities who have the authority and capability – derived from their relationship with the affected stakeholders – to mitigate harm. A mitigation strategy should admit the variance in capacity and/or willingness of the notified entity to understand and act on the notification.

C.4 Justice: Fairness and Equity

In the Belmont Report, the principle of Justice is applied through fairness in the selection of research subjects, and equitable distribution of the burdens and benefits of research according to individual need, effort, societal contribution, and merit. Fairness should guide the initial selection of the subjects, as well as the apportionment of burdens to those who will most likely benefit from the research. Research design and implementation should consider all stakeholders' interests, although conflicting interests may render equal treatment impracticable. In the ICTR context, this principle implies that research should not arbitrarily target persons or groups based on attributes including (but not limited to): religion, political affiliation, sexual orientation, health, age, technical competency, national origin, race, or socioeconomic status. Neither should ICTR target specific populations for the sake of convenience or expediency.

It is important to distinguish between purposefully *excluding* groups based on prejudice or bias versus purposefully *including* entities who are willing to cooperate and consent, or who are better able to understand the technical issues raised by the researcher. The former raises

Justice concerns, while the latter demonstrates efforts to apply the principles of Respect for Persons and Beneficence and still conduct meaningful research.

Challenges to obtaining informed consent from users might motivate a researcher to work with a service provider who has direct contractual relationships with its network's users, as well as associated authority to proxy notice and consent on their behalf. Such decisions to engage entities that can offer research consent by proxy may raise fairness and equity concerns. Each provider with whom a researcher may interact will have varying levels of understanding and ability (or willingness) to act. If a researcher is required to get unanimous and uniform responses from all autonomous entities, it may be impossible to perform beneficial research. On the other hand, moving forward with risky research without the involvement, or at least awareness, of autonomous entities is undesirable as it may increase the potential for greater harm.

From an equity standpoint, open public disclosure of system vulnerabilities demands that researchers consider how the burdens and benefits of publicizing newly discovered vulnerability balance out. The burdens might be borne by the developers, yet actually might benefit malicious actors more in the short-term than developers or users of those systems. The calculation of benefits is actually a function of time, where malicious actors may act faster at exploiting vulnerability information than benevolent actors can act in mitigating the vulnerabilities.

C.5 Respect for Law and Public Interest

Respect for Law and Public Interest is implicit in the Belmont Report's application of Beneficence. In the context of ICTR, we include it as a separate principle with two applications – *Compliance* and *Transparency and Accountability*. The second application refers to transparency of methodologies and results, and accountability for actions. Transparency and accountability serve vital roles in many ICTR contexts where it is challenging or impossible to identify stakeholders (e.g., attribution of sources and intermediaries of information), to understand interactions between highly dynamic and globally distributed systems and technologies, and consequently to balance associated harms and benefits. A lack of transparency and accountability risks undermining the credibility of, trust and confidence in, and ultimately support for, ICT research.

C.5.1 Compliance

Researchers should engage in due diligence to identify laws, regulations, contracts, and other private agreements that are applicable to their research, and should design and implement ICTR that respects these restrictions. While legal controls that call for compliance can be numerous and wide-ranging, those that should inform ethical assessments cluster categorically around computer crime and information security, privacy and anonymity, intellectual property, computer system assurance, and civil rights and liberties. More specifically, ICT research may implicate rights and obligations related to: identity theft; unsolicited bulk electronic mail; privacy in electronic and wire communications; notification of security breaches; copyright and other intellectual property infringement; data security and destruction; child pornography; spyware and phishing; fraudulent deception; financial privacy; economic espionage; constitutional privacy; health information security and privacy; industry standards and best practices; and contractual privacy and acceptable use policies.

Respect for public interest can often be addressed by obeying relevant laws. If applicable laws conflict with each other or with the public interest, and a decision is made to not comply with legal obligations that are viewed as unethical, researchers should have ethically defensible justification and be prepared to accept responsibility for their actions and consequences.

C.5.2 Transparency and Accountability

Transparency is a mechanism to assess and implement accountability, which itself is necessary to ensure that researchers behave responsibly. These applications interact to ultimately generate trust in ICTR by the public. Transparency-based accountability helps researchers, oversight entities, and other stakeholders avoid guesswork and incorrect inferences about whether, where, and how ethical principles are addressed. Transparency entails clearly communicating the purposes of research – why data collection and/or direct interaction with ICT is required to fulfill those purposes – and how research results will be used. It also involves clear communication of risk assessment and harm minimization related to research activities.

Accountability demands that research methodology, ethical evaluations, data collected, and results generated should be documented and made available responsibly in accordance with balancing risks and benefits. Data should be available for legitimate research, policy-making, or public knowledge, subject to appropriate collection, use, and disclosure controls informed by the Beneficence principle. The appropriate format, scope and modality of the data exposure will vary with the circumstances, as informed by Beneficence determinations.

D Implementing the Principles and Applications

This document describes foundational ethical principles and their applications at a level intended to span a broad range of current and future research that will undoubtedly be affected by changes in ICT. For federally funded biomedical and behavioral research, the responsibility for evaluating whether a research project comports with these principles lies with REBs, which in the United States are known as Institutional Review Boards (IRBs). This report contends that ICTR will benefit from similar oversight, and the proposed guidelines will assist ICT researchers and oversight authorities identify, preempt and manage ethical risks. Current ICTR that does not fall under the purview of REBs would also benefit from community-derived self-regulation guided by this report. Proactively and transparently engaging in ethical assessment of ICT research will help move the research community mindset in the direction of embedding ethics into ICTR design as productively and safely as possible, and more practically influence policy and governance at these crossroads.

Notes

¹The term *information and communication technology* was coined by Denis Stevenson in a 1997 report to the United Kingdom government, *Information and Communication Technologies in the UK Schools: An Independent Inquiry* <http://rubble.heppell.net/stevenson/ICT.pdf>

²This report offers pragmatic guidance in the application of these fundamental principles to ICTR, and avoids taking a position in the philosophical debate about the uniqueness of computer ethics. For an overview of the philosophical debate, see, Bynum, Terrell, “Computer and Information Ethics”, *The Stanford Encyclopedia of Philosophy* (Winter 2008 Edition), Edward N. Zalta (ed.). <http://plato.stanford.edu/archives/win2008/entries/ethics-computer/>

³*The Belmont Report*, the touchstone document guiding human subjects research in the biomedical and behavioral research fields, was named after the conference center where it was drafted in 1976.

(See <http://ohsr.od.nih.gov/guidelines/belmont.html>.) This document similarly takes its name from the city where a substantial portion of the working group meetings that resulted in this document took place in 2009-2010.

⁴Fifteen government departments and agencies performing research involving human subjects adopted 45 CFR 46 Subpart A in what is known as the *Common Rule*. Each has its own guidance on the interpretation of their section of the CFR. Refer to guidance appropriate to the funding source.

⁵See <http://ohsr.od.nih.gov/guidelines/belmont.html>