CHAPTER TWENTY-FOUR

Informed Consent by Design

BATYA FRIEDMAN, PEYINA LIN, AND JESSICA K. MILLER

CONSUMER PRIVACY ATTITUDES HAVE SHIFTED SIGNIFICANTLY OVER THE PAST DECADE, from being a concern of a minority in the 1980s, to being a concern of the majority. As of 2001, more than three-quarters of Americans (77%) were "very concerned" about potential misuse of their personal information.[1] Consumers not only want to be informed about how their personal information will be used, but also want the opportunity to choose: 88% of Internet users would like "opt-in" privacy policies that require Internet companies to ask consumers for permission to share their personal information with others "always" (based on a four-tier scale from "never" to "always"). In addition, if consumers could choose not to have their personal information collected, 56% would "always" opt out, and 34% would "sometimes" opt out.[2]

While informed consent may not eliminate all privacy concerns for all groups, it can help to gain users' trust by accurately articulating business practices for personal information and by allowing users autonomous choice.

[1] Alan Westin, "Opinion Surveys: What Consumers Have to Say About Information Privacy" (2001) [cited Dec. 9, 2004]; *http://energycommerce.house.gov/107/hearings/05082001Hearing209/Westin309.htm*.

[2] The McGraw-Hill Companies Inc., "Business Week/Harris Poll: A Growing Threat" (2000) [cited Dec. 9, 2004]; *http://www.businessweek.com/2000/00_12/b3673010.htm*.

How, then, might the information systems community design for informed consent? We take up that challenge here. Toward answering this question, we ground our work in the interactional theory and tripartite methodology of Value Sensitive Design.[3, 4, 5] Our approach integrates conceptual, technical, and empirical investigations. We consider the impact of information systems on both direct and indirect stakeholders.

Introduction

Changes in consumer attitudes are occurring against the backdrop of two major trends: (1) the evolution of the concept of privacy; and (2) the erosion of historical protections for privacy.

The conception of privacy and, correspondingly, that of informed consent, has evolved through changes in political, legal, economic, social, and technological spheres. In earlier times, privacy (and the security it afforded) existed primarily in relation to physical property. Consequently, protections were given against trespasses of physical property and against battery to a person's body. Liberty meant "freedom from actual restraint."[6] In the late 1800s, advances in photographic technology made it possible to take pictures surreptitiously. Thus, the implicit consent given when "sitting" for a portrait became an inadequate safeguard against the improper capturing of one's portrait for circulation and profit. In reaction to these changes, in a landmark 1890 *Harvard Law Review* article, "The Right to Privacy,"[7] Samuel Warren and Louis Brandeis urged the courts to recognize an individual's "right to be let alone," to be free from unwarranted intrusions into personal affairs. While Warren and Brandeis were referring primarily to publishable records (e.g., photography and popular media articles), inherent to their plea was the change in meaning of "the right to life," from merely protecting physical property to "the right to enjoy life" spiritually, emotionally, and intellectually. Important for purposes here, the article established a clear relationship between the "right to privacy" and informed consent: that "the right to privacy ceases upon the publication of the facts by the individual, or upon his consent." In addition, having the ability to consent—to prevent or allow publication—afforded peace of mind, relief, and freedom from fear of injury.[8]

[3] Batya Friedman (ed.), *Human Values and the Design of Computer Technology* (Cambridge: CSLI Publications Center for the Study of Language and Information, 1997).

[4] Batya Friedman, "Value Sensitive Design," in William Sims Bainbridge (ed.), *Berkshire Encyclopedia of Human-Computer Interaction* (Great Barrington, MA: Berkshire Publishing Group, 2004), 769–777.

[5] Batya Friedman, Peter Kahn, and Alan Borning, "Value Sensitive Design and Information Systems," in D. Galleta and Ping Zhang (eds.), *Human-Computer Interaction in Management Information Systems* (Armonk, NY: M.E. Sharpe, in press).

[6] Samuel D. Warren and Louis D. Brandeis, "The Right to Privacy," *Harvard Law Review* IV:5 (1890).

[7] *Ibid*.

[8] *Ibid*.

Paralleling this evolution in conceptions of privacy, protections for privacy have eroded, in large part brought about by technological advances. For example, in earlier times, the mere format in which personal information was collected (paper) afforded reasonable privacy protection. Inconvenient access to information, limited reproduction capabilities, and the inability to easily link, sort, and process information acted as "natural" shields protecting personal data. Nowadays, networking and digitization have stripped the public of that "natural shielding," and continuously beg for the reconceptualization of privacy and its protection. How is this relevant for today's businesses and for technology design?

Accounting for privacy in technology design does not merely address a moral concern, it safeguards the design host company from financial risks and public relations backlashes. For example, the built-in unique processor-identifier of Intel's Pentium III processor, also known as the Processor Serial Number (PSN) technology, was introduced to help corporations manage inventories.[9] Even though the PSN is a number attached to a processor, not a person, and even though Intel took precautionary actions by creating a control utility that allowed users to turn off the identifier, the introduction of security at the potential expense of privacy generated consumer distrust. Intel was faced with consumer boycotts and legal action attacks from privacy advocacy groups. According to a 1999 report on CNET news,"The serial number has sparked a definite negative emotional core with segments of the population."[10] Other technologies, such as commonplace workstations that contain microphones without hardware on/off switches[11] and the more exotic Active Badge Location system of the Palo Alto Research Center (PARC),[12] are examples of designs that similarly compromised privacy for the sake of other functionalities and that were met with resistance from some groups.

In this chapter, we first present a conceptual model of informed consent for information systems. Next we present three cases in which the conceptual model has been applied: cookie handling in web browsers; secure connections for web-based interactions; and Google's Gmail web-based email service. Each case discussed here involves a widely deployed technology in use at the time our investigations were conducted; each invokes privacy or security concerns for the public at large; and each highlights a unique set of challenges and design possibilities for informed consent. In reporting on these cases, our goal is not only to articulate how these three specific systems might be improved, but also more generally to illustrate how the model can be used proactively to design information systems that support the user experience of informed consent. Finally, we propose design principles and business practices to enhance privacy and security through informed consent.

[9] Intel Corporation, "Intel Pentium III Processor" (2003) [cited Dec. 10, 2004]; *http://support.intel.com*.

[10] Stephanie Miles, "How Serious Is Pentium III's Privacy Risk?" (1999) [cited Dec. 1, 2004]; *http://news.com.com/How+serious+is+Pentium+IIIs+privacy+risk/2100-1040_3-222905.html*.

[11] John C. Tang, "Eliminating a Hardware Switch: Weighing Economics and Values in a Design Decision," in Batya Friedman (ed.), *Human Values and the Design of Computer Technology* (Cambridge: CSLI Publications Center for the Study of Language and Information, 1997), 259–269.

[12] Roy Want, Andy Hopper, Veronica Falcao, and Jon Gibbons, "The Active Badge Location System," *ACM Transactions on Information Systems* 10 (1992), 91–102.

VALUE SENSITIVE DESIGN

Value Sensitive Design[a] emerged in the 1990s as an approach to the design of information and computer systems that accounts for human values in a principled and comprehensive manner throughout the design process. While emphasizing the moral perspective (e.g., privacy, security, trust, human dignity, physical and psychological well-being, informed consent, intellectual property), Value Sensitive Design also accounts for usability (e.g., ease of use), conventions (e.g., standardization of technical protocols), and personal predilections (e.g., color preferences within a graphical interface).

Key features of Value Sensitive Design involve its interactional perspective, tripartite methodology, and emphasis on direct and indirect stakeholders:

Interactional theory

Value Sensitive Design is an interactional theory: values are viewed neither as inscribed into technology nor as simply transmitted by social forces. Rather, people and social systems affect technological development, and new technologies shape (but do not rigidly determine) individual behavior and social systems.

Tripartite methodology: conceptual, empirical, and technical

Value Sensitive Design systematically integrates and iterates on three types of investigations. *Conceptual investigations* comprise philosophically informed analyses of the central constructs and issues under investigation. For example: what values have standing? How should we engage in tradeoffs among competing values (e.g., access versus privacy, or security versus trust)? *Empirical investigations* focus on the human response to the technical artifact and on the larger social context in which the technology is situated. The entire range of quantitative and qualitative social science research methods may be applicable (e.g., observations, interviews, surveys, focus groups, measurements of user behavior and human physiology, contextual inquiry, and interaction logs). *Technical investigations* focus on the design and performance of the technology itself, involving both retrospective analyses of existing technologies and the design of new technical mechanisms and systems. The conceptual, empirical, and technical investigations are employed iteratively such that the results of one type are integrated with those of the others, which, in turn, influence yet additional investigations of the earlier types.

Direct and indirect stakeholders

Direct stakeholders are parties who interact directly with the computer system or its output. *Indirect stakeholders* are all other parties who are otherwise affected by the use of the system. For example, online court record systems impact not only the direct stakeholders, such as lawyers, judges, and journalists who access the court records, but also an especially important group of indirect stakeholders: the people documented in the court records.

[a] Adapted from Batya Friedman, “Value Sensitive Design” in William Sims Bainbridge (ed.), *Berkshire Encyclopedia of Human-Computer Interaction* (Great Barrington, MA: Berkshire Publishing Group, 2004), 769–777.

A Model of Informed Consent for Information Systems

The model of informed consent for information systems we present here was first developed in 2000 by Friedman, Felten, and Millett[13] in the context of online interactions.[14] This model is based on six components:

- Disclosure
- Comprehension
- Voluntariness
- Competence
- Agreement
- Minimal distraction

The word *informed* encompasses the first two components: disclosure and comprehension. The word *consent* encompasses the following three components: voluntariness, competence, and agreement. In addition, the activities of being informed and giving consent should happen with minimal distraction, without diverting users from their primary task or overwhelming them with intolerable nuisance.

Disclosure

Disclosure refers to providing accurate information about the benefits and harms that might reasonably be expected from the action under consideration. What is disclosed should address the important values, needs, and interests of the individual, explicitly state the purpose or reason for undertaking the action, and avoid unnecessary technical detail. The information should also disabuse the individual of any commonly held false beliefs. Moreover, if the action involves collecting information about an individual, then the following should also be made explicit:

- What information will be collected?
- Who will have access to the information?
- How long will the information be archived?
- What will the information be used for?
- How will the identity of the individual be protected?

Comprehension

Comprehension refers to the individual's accurate interpretation of *what* is being disclosed. This component raises the question: how do we know when something has been adequately comprehended? While there is no easy answer here, at least two methods seem

[13] Batya Friedman, Edward Felten, and Lynette I. Millett, "Informed Consent Online: A Conceptual Model and Design Principles," *CSE Technical Report* (Seattle: University of Washington, 2000).

[14] See also Ruth R. Faden and Tom L. Beauchamp, *A History and Theory of Informed Consent* (New York: Oxford University Press, 1986).

viable: (1) being able to restate in different words what has been disclosed, and (2) being able to apply what has been disclosed to a set of hypothetical events. Take, for example, a web-based recommendation system, such as an e-commerce site recommending products based on the customer's prior purchases or on purchases of other customers with similar profiles. Based on information disclosed to the customer—about what data is being collected and how it will be used—can the customer answer reasonable questions about the data's use, such as:

- Will information about the customer's last three purchases be included in the recommendation system?
- Will some other user of the recommendation system be able to determine what the customer has purchased in the past?
- Will information about the customer's past purchases be a part of the recommendation system two years from now?

In face-to-face interactions, two-way dialog, facial expressions, and other physical cues help ensure the adequate interpretation of any information disclosed. Technologically mediated interactions, however, lack many of the cues and opportunities to ascertain and ensure comprehension. Typical web-based interactions present the disclosure of information in a web page or dialog box, and users are expected to agree to or decline participation by clicking on a button. Rarely are email or chat facilities provided during the process of disclosure. As more and more interactions move online or become mediated by technology, ensuring comprehension becomes more challenging. Nevertheless, comprehension is a crucial component of informing, and it should not be dismissed.

Voluntariness

A voluntary action is one in which an individual could reasonably resist participation should she wish to. Voluntariness, then, refers to ensuring that the action is not coerced or overly manipulated.

Coercion is an extreme form of influence that controls by compulsion, threat, or prevention. A canonical example of coercion occurs as follows: Person A holds a gun to Person B's head and says, "Fly me to Havana or I'll shoot." Often, coercion can occur without notice when there is only one reasonable way for individuals to receive certain needed services or information (or, if other ways do exist, they are too costly in terms of finance, time, expertise, or other costs to be viable options.) This form of coercion is a serious concern for online interactions and other technology mediated interactions. As more and more critical services move online in their entirety, such as applying for medical insurance or to universities for higher education, individuals who want to obtain these services or information will have to engage in web interactions. Given the lack of substantive choice among web sites and web browsers, users can be, in effect, coerced into web-based interactions that compel them to give up personal information or engage in other activities.

Manipulation of certain forms can also undermine voluntariness. Manipulation can roughly be defined as "any intentional and successful influence by a person by noncoercively altering the actual choices available to the person or by nonpersuasively altering the person's perceptions of those choices."[15] The key here is that manipulation alters the individuals' choices or perception of choices by some means other than reason. This sort of manipulation can be achieved in at least three ways:

- *Manipulation of options.* The first way entails manipulation of the options presented to the individual such that the presentation encourages certain choices or behaviors. For example, consider an e-business that asks the user for more information than is necessary to complete a purchase but does not indicate to the user that completing some fields is optional.
- *Manipulation of information.* The second way entails manipulation of information. This manipulation uses information intentionally to overwhelm the individual or to provoke or take advantage of an individual's fear or anxiety. For example, some web sites have packaged information into multiple cookies—information that could have been packaged more concisely—so that the user who elects to accept cookies on a case-by-case basis would be bombarded with numerous requests to set cookies from a single site. As a result, the user may turn on the "agree to all cookies" option to avoid the overwhelming requests for information, or fail to notice an undesirable cookie.
- *Psychological manipulation.* The third way is psychological. This form of manipulation includes any intentional act that influences a person by causing changes in the individual's mental processes by any means other than reason. Flattery, guilt induction, and subliminal suggestions are a few relevant influences. Recent work by Reeves and Nass and their colleagues[16, 17] indicates that individuals are vulnerable to psychological manipulation in online interactions, particularly with respect to psychological manipulations from the technology's interface. For example, in their research, Reeves and Nass have shown that users respond to flattery from a computer, judge computers that criticize rather than praise others to provide more accurate information, and apply gender stereotypes to the technology simply on the basis of subtle gender cues, to name but a few of their results. Web sites that use psychological manipulation—for instance, to flatter the user into divulging information or into attributing greater accuracy to online recommendations—may violate the criterion of voluntariness.

[15] *Ibid.*, 354.

[16] B. Reeves and C. Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places* (New York: Cambridge University Press, 1996).

[17] C. I. Nass, Y. Moon, J. Morkes, E. Kim, and B. J. Fogg, "Computers Are Social Actors: A Review of Current Research," in Batya Friedman (ed.), *Human Values and the Design of Computer Technology* (New York: Cambridge University Press, 1997), 137–162.

Competence

Competence refers to possessing the mental, emotional, and physical capabilities needed to give informed consent. For example, a person with Alzheimer's may lack the mental capability to make her own medical decisions. Or, in the online environment, a 15-year-old may have the technical competence, but lack the mental and emotional capability to make reasoned judgments about when to provide personal information to e-businesses and in online chat rooms.

For example, at roughly the same time as the Year 2000 Census was being conducted in the United States, the Barbi web site presented Barbi as a census taker who asked web site visitors—mostly young girls under the age of 12—to help Barbi with her census work by completing a form that requested personal information. Troublesome from the perspective of informed consent, these young girls often lacked the mental and emotional competence necessary to judge the appropriateness of the information they volunteered to the web site.

Designers of web sites and other technologies targeted to children and adolescents will need to be especially cognizant of the component of competence. For example, the United States Children's Online Privacy Protection Act (COPPA) requires written parental consent when web sites collect information from children aged 12 and under. Another model is that used by many Institutional Review Boards; it suggests obtaining informed consent from both the adolescent and the adolescent's guardian for adolescents between the ages of 13 and 17 (inclusive). However, the case of adolescents is not straightforward: a tension exists between ensuring that adolescents can adequately assess the impacts of the type of information being collected from them and maintaining adolescents' privacy (as this is typically the time in a person's life that privacy vis a vis one's parents begins to become an important value). This tension should be considered when determining whether it is necessary to obtain consent from the adolescent's guardian as well as from the adolescent.

Agreement

Agreement refers to a reasonably clear opportunity to accept or decline to participate. Aspects to consider include:

- Are opportunities to accept or decline visible, readily accessible, and ongoing?
- Is agreement by the participant ongoing?

In traditional human subjects research, the component of agreement is ongoing. Participants may withdraw their agreement to participate at any time and for any reason (participants do not need to provide a reason for discontinuing participation). While the arena of research differs in important ways from interactions with information systems, and while considerable complexity exists concerning how to apply the guidelines from one to the other, still the aspect of ongoing agreement may have relevance for information systems. For example, in the case of recommendation systems, users could be provided with the opportunity to withdraw or prevent further use of their data from the recommendation system at any time. Many of today's recommendation systems lack such an ability.

A related issue for ongoing agreement arises in the context of discussion groups and online chat rooms where dialog that may often feel like "ethereal" online conversation is archived and, in reality, is more permanent and accessible than most other forms of communication. In these online forums, participants in the flurry of heated conversation may forget that they have agreed to have their online conversation recorded and archived and, if given the opportunity, might even suspend that agreement. Mechanisms that periodically remind participants that online dialog may be archived (and perhaps allow participants to remove dialog from the archive) could help preserve informed consent in these interactions.

Not all forms of agreement need to be explicit. As a society, we have a good deal of experience with implicit consent where, by virtue of entering into a situation, the individual has, in effect, agreed to the activities that are broadly known to occur in that context. For example, when a player steps out onto the football field in football garb and enters the game, the individual has implicitly agreed to participate in the normal activities of the game—namely, to being bumped, bashed, and smashed by other players who have entered into identical agreements. Implicit consent holds in this case because the other components have also been met: disclosure and comprehension (via reasonable expectation), competence (if we assume that the individual is of a reasonable age and of sound mind and body), and voluntariness (if we assume that the individual was not coerced or manipulated to dress in football garb and to go out onto the field). For implicit consent to hold for information systems, similar criteria need to be met.

Minimal Distraction

This criterion for informed consent arose from empirical investigations in which users, overwhelmed by the activities of "being informed" and "giving consent," became numbed to the informed consent process and disengaged from the process in its entirety. Specifically, minimal distraction refers to meeting the preceding criteria of disclosure, comprehension, competence, voluntariness, and agreement without "unduly diverting the individual from the task at hand."[18] This criterion is challenging to implement, because "the very process of informing and obtaining consent necessarily diverts users from their primary task,"[19] yet it is crucial if informed consent is to be realized in practice.

With a model of informed consent for information systems in hand, we turn now to examine how this model can be used to analyze, assess, and improve the design of existing information systems.

[18] Batya Friedman, Daniel C. Howe, and Edward Felten, "Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design." Thirty-Fifth Hawaii International Conference on System Sciences (Hawaii, 2002).

[19] *Ibid*.

Possibilities and Limitations for Informed Consent: Redesigning Cookie Handling in a Web Browser

In this section, we demonstrate that the proposed model of informed consent can be used to:

- Assess how effectively a particular system design supports informed consent
- Guide successive design
- Identify how the underlying technology may constrain the range of possible solutions to support users' informed consent

Specifically, we examine the role the web browser plays in obtaining informed consent for cookies.

What Are Cookies and How Are They Used?

A cookie is a text file stored on the user's machine that can be used to maintain information about the user, such as an identifier or a log of the user's navigation on the web site. One accepted way companies use cookies is to customize their web site to the user (e.g., Amazon uses cookies to remember what users like to buy and what users put in their shopping baskets). However, cookies can also be abused by surreptitiously collecting information about the user.

When a user wants to retrieve a particular web page from the Internet, the user opens a web browser and enters the web page's address. The browser sends the request for the web page to the appropriate web server. The web server then retrieves the requested web page and sends it back to the user's browser where it is displayed. This process involves a couple of additional steps if the web server wants to set a cookie. When sending the requested web page back to the browser, the server sends the browser a request to store the cookie. Depending on how it has been programmed, and on any cookie-related preferences that have been set, the browser may or may not store the cookie as requested. If the browser stores the cookie on the user's computer, the browser will volunteer the cookie each time the user revisits that web page and possibly any other web pages in that domain, depending on the scope of the cookie.

Web Browser as Gatekeeper to Informed Consent

In the previous description, the browser acts as a gatekeeper by determining which web server requests to fulfill. In addition, with respect to informed consent, the web browser plays at least two other critical gatekeeping roles:

- The web browser controls whether the user is notified about a server request and, to a large extent, controls the content of that notification. Thus, the components of disclosure and comprehension largely reside in the web browser.
- The web browser controls whether the user has an opportunity to agree to or decline the web server's request (e.g., prompting for user input each time a server requests to place a cookie as opposed to the browser handling the request without user input). Thus, the component of agreement also resides in the browser.

Admittedly, a proactive web site could supplement the functionality provided by the web browser by explicitly addressing disclosure and agreement (e.g., privacy policies). However, relying on all web sites to do this individually would result in ad hoc methods and require users to become familiar with each web site's policies and practices.

Web Browser Development and Progress for Informed Consent: 1995–1999

With a conceptualization for informed consent online in hand, Millett, Friedman, and Felten[20] conducted a retrospective analysis of how cookie handling in Netscape Navigator and Internet Explorer evolved with respect to informed consent over a five-year period, beginning in 1995. Specifically, they used the criteria of disclosure, comprehension, voluntariness, competence, and agreement to evaluate how well each browser in each stage of its development supported the users' experience of informed consent for the use of cookies. (At this early stage in their work, the criterion of minimal distraction had not yet been identified.)

Through this analysis, they found that cookie technology had improved over time regarding informed consent. For example, there had been increased visibility of cookies, options for accepting or declining cookies, and access to information about cookie content. However, as of 1999, some startling problems remained:

- While browsers disclosed to users some information about cookies, the right kind of information—that is, information about the potential harms and benefits from setting a particular cookie—was still not disclosed.
- In Internet Explorer, if a user wanted to decline all third-party cookies, the burden fell on the user to do so one cookie at a time.
- Users' out-of-the-box cookie experience (i.e., the default setting) was no different in 1999 from what it had been in 1995: to accept all cookies. That is, the novice user installed a browser that accepted all cookies and disclosed nothing about that activity to the user.
- Neither Internet Explorer nor Netscape Navigator alerted a user when a cookie was sent back to a site, as opposed to when a cookie was stored.

Redesigning the Browser

After completing the retrospective analysis, Friedman, Howe, and Felten[21] considered how to redesign the browser to better support informed consent. First, they identified four overarching design goals:

- Enhance users' understanding of specific cookie events
- Enhance users' global understanding of the common uses of cookie technology, including what a cookie is and its potential benefits and risks

[20] Lynette I. Millett, Batya Friedman, and Edward Felten, "Cookies and Web Browser Design: Toward Realizing Informed Consent Online," CHI (Seattle, WA, March 2001).

[21] Friedman, Howe, and Felten.

- Enhance users' ability to manage cookies
- Achieve these goals while minimizing distraction for the user

By iterating through three design prototypes, each followed by small-scale usability studies (see the earlier sidebar, "Value Sensitive Design"), Friedman *et al*. redesigned the Cookie Manager tool of the Mozilla browser (the open source version of Netscape Navigator).[22]

In consideration of the design goals and design strategies, Friedman *et al*.[23] implemented a *peripheral awareness* mechanism by implementing a small application they named the "Cookie Watcher" docked in Mozilla's existing sidebar. In their final design (see Figure 24-1), users were notified in real time not only about the occurrence of cookie events, but also about the domain and type of cookie being set. Visual cues such as background color and font style were used to represent domain and duration information, respectively, as follows: third-party cookies were displayed in red; cookies from the same domain were displayed in green; italicized fonts were used for session cookies; and bold fonts were used for cookies with durations of more than a year.

Two just-in-time interventions were implemented:

- Users can click on any installed cookie (once displayed in the sidebar) to bring up a cookie manager tool. With that tool, the user can learn more information about the specific cookie, delete the cookie, and ban that site from resetting cookies.
- At the bottom of the sidebar, users can click on a "Learn About Cookies" button to prompt a Cookie-Information Dialog Box with information about the potential benefits and harms of cookies, and label information on what the colors and font sizes represent.

Participants of the usability studies commented favorably on the just-in-time management tools. Their comments suggested that the Cookie Watcher helped to enhance understanding about cookies as well as eased cookie management. Evidence showed that direct access to information on individual cookies supported the design goals.

Technical Limitations to Redesigning for Informed Consent

Although the Mozilla prototypes made good progress toward achieving the design goals, there remained changes that Friedman *et al*. were unable to make as a result of the underlying technology. For example, many users would like to know not only when a web site wants to set a cookie, but also when a web site wants to *use* a cookie. However, because the web browser automatically volunteers cookies whenever the user revisits the domain, it is currently not possible to provide users with that information.

In order to fully disclose when and why a particular cookie is being set and then used—and its potential harms and benefits—changes would need to be made to the network

[22] Mozilla version 0.8 was used in the prototype designs.

[23] Friedman, Howe, and Felten.

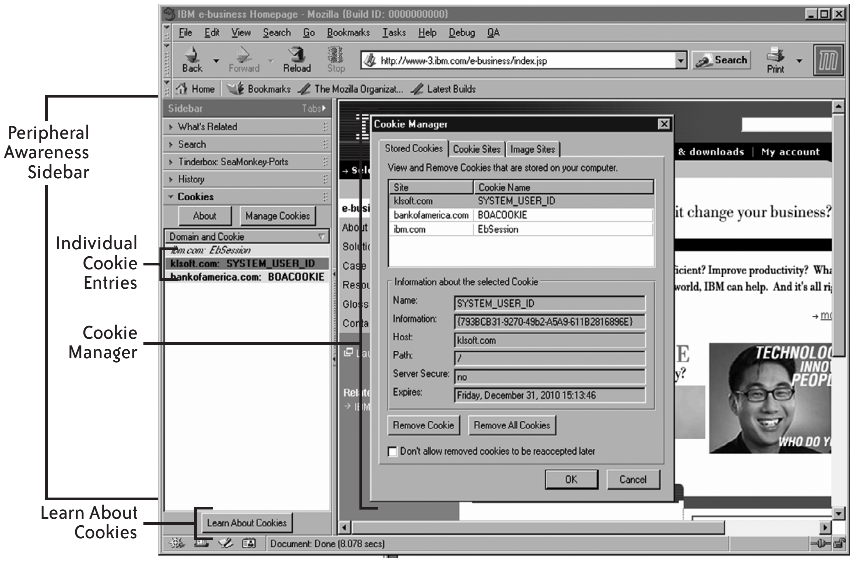
Peripheral Awareness Sidebar
Individual Cookie Entries
Cookie Manager
Learn About Cookies
IBM e-business Homepage - Mozilla {Build ID: 0000000000}
File Edit View Search Go Bookmarks Tasks Help Debug QA
Back Forward Reload Stop
Search
Print
Home Bookmarks The Mozilla Organizat... Latest Builds
Sidebar Tabs
What's Related
Search
Tinderbox: SeaMonkey-Ports
History
Cookies
About
Manage Cookies
Learn About Cookies
Cookie Manager
Stored Cookies
Cookie Sites
Image Sites
Remove Cookie
Remove All Cookies
OK
Cancel

http://www-3.ibm.com/e-business/index.jsp

Domain and Cookie

ibm.com: EbSession
klsoft.com: SYSTEM_USER_ID
bankofamerica.com: BOACOOKIE

View and Remove Cookies that are stored on your computer.

Site	Cookie Name
klsoft.com	SYSTEM_USER_ID
bankofamerica.com	BOACOOKIE
ibm.com	EbSession

Information about the selected Cookie

Name:	SYSTEM_USER_ID
Information:	{793BCB31-9270-49b2-A5A9-611B2816896E}
Host:	klsoft.com
Path:	/
Server Secure:	no
Expires:	Friday, December 31, 2010 15:13:46

☐ Don't allow removed cookies to be reaccepted later

Document: Done (8.078 secs)

FIGURE 24-1. Taking advantage of Mozilla's sidebar structure, a new peripheral awareness sidebar (the Cookie Watcher) was developed to dynamically notify users whenever a cookie is set (as shown above); in addition, two just-in-time mechanisms were implemented and supported in the Cookie Watcher: (1) by clicking on a cookie, the user can bring up a cookie manager tool (as shown above); (2) by clicking on the Learn About Cookies button, the user can bring up a dialog box with information about the potential harms and benefits of cookies, and about what the colors and font styles in the Cookie Watcher represent (not shown)

protocol that, in turn, would necessitate changes to the web browser and the remote web site, as follows:

- The underlying network protocol (in this case, HTTP) would need to support the description of the harms and benefits of the cookie as well as provide a means for the web site to request to *use* (not just set) a cookie.
- The browser then would need to be programmed to display these harms and benefits in a clear and accessible way both when the cookie is being set and when the cookie is requested.
- Finally, the web site would need to provide accurate information by filling in the appropriate HTTP fields.

While each entity has a critical role to play, the network protocol constrains possibilities for the other two.

Reflections

This section has highlighted how the model of informed consent can be used to evaluate and design information systems. We note that some of the design ideas presented here, such as the use of peripheral awareness mechanisms, can now be seen in the current version of both Mozilla (version 1.73) and Internet Explorer (version 6). For example, both browsers now use a peripheral awareness mechanism in the form of a small "eye" icon displayed in the bottom righthand corner of the browser window to indicate when a web site attempts to set a cookie that is restricted or blocked by the users' privacy preferences. Thus, the design methods presented in this section have pragmatic value.

We have also explicated the interaction between the underlying technical infrastructure (in this case, the HTTP protocol) and what solutions can be designed and implemented to support users' informed consent. Along these lines, the Platform for Privacy Preferences (P3P) (discussed in Chapter 22 in this volume) represents one recent effort that works around the protocol limitations to provide a mechanism that evaluates web site privacy practices against user-specified privacy preferences.

Finally, we point to the interaction between technical implementation and business practice: even if there were no technical limitations on redesigning the web browser for informed consent, in order for the web browser to provide complete disclosure to the user, web sites would need to provide an accurate and clear description of what information their cookies collect and how the collected information will be used. In this and other ways, business practice must work in consort with technical implementations.

Informing Through Interaction Design: What Users Understand About Secure Connections Through Their Web Browsing

The process of informing the user can happen as the user interacts with the system instead of through simple, explicit text disclosure. That is, in addition to the user's existing knowledge about how the system functions, the visual cues during interaction and the text displayed on the interface (web pages, browser, etc.) may lead the user to develop an idea or mental model of how the system functions. An issue of concern arises when there is a mismatch between the disclosed text and the interaction cues—in particular, when the latter heavily influences the user's perception of how the system works. As a result, in a best-case scenario, the user could end up confused but not jeopardize any personal data; in a worst-case scenario, the user could construct inaccurate mental models about the security of the system and make poor decisions on what actions to take or not to take to protect personal information.

With the design strategy of informing through interaction in mind, in this section we describe a study by Friedman, Hurley, Howe, Felten, and Nissenbaum[24] on how users across diverse communities conceptualize web security.

Participants

Seventy-two individuals, 24 each from a rural community in Maine, a suburban professional community in New Jersey, and a high-technology community in California, partici-

[24] Batya Friedman, David Hurley, Daniel C. Howe, Edward Felten, and Helen Nissenbaum, "Users' Conceptions of Web Security: A Comparative Study," in *Extended Abstracts of CHI* (2002), 746–747.

pated in an extensive (two-hour) semistructured interview concerning users' conceptions, views, and values about web security. Equal numbers of men and women participated from each community. We report here on one section of the interview that focused on users' mental models of web security. Both verbal and nonverbal techniques were used to assess users' understandings.

Users' Conceptions of Secure Connections

Participants were asked to define and portray secure connections in various ways, as we describe in the following subsections.

Definition of a secure connection

Participants were first asked to define a secure connection. Participants' definitions of a secure connection encompassed one of the following concepts:

- *Transit.* Protecting the confidentiality of information while it moves between machines on the Web
- *Encryption.* The specific mechanism of encoding and decoding information
- *Remote site.* Protecting information once it has arrived at its destination on the Web

High-technology participants (83%) provided correct definitions of a secure connection more frequently than rural participants (52%) ($p < .05$) did. Statistically, there was no difference in responses between the high-technology (83%) and suburban (68%) participants.

Recognition of a connection as secure or not secure

Next, participants were shown four screenshots of a browser connecting to a web site and were asked to recognize a secure connection. For each screenshot, participants were asked to state whether the web connection was secure or not secure, as well as to provide the rationale for their evaluation.

Table 24-1 shows the types of evidence participants used to evaluate a connection. As shown, participants depended primarily upon six types of evidence.

1. HTTPS protocol—for example, "usually, it says http for nonsecure or standard and https for secure, the s meaning secure".
2. Icon—for example, "[the site is secure] just because the key is there".
3. Point in transaction—for example, "it looks like one of the main pages on the site and usually main pages are nonsecured connections".
4. Type of information—for example, "that at least has the indication of a secure connection; I mean, it's obviously asking for a Social Security number and a password".

5. Type of web site—for example, "I can't imagine a bank would be online and not have security measures in there".
6. General distrust—for example, "I'm wary of the computer itself…I basically don't think any of the sites are secure".

TABLE 24-1. Percentage of types of evidence participants used to evaluate a connection as secure or not secure

Type of evidence[a]	Correct evaluation		Incorrect evaluation	
	Not secure	Secure	Not secure	Secure
1. HTTPS protocol	16	20	0	9
2. Icon (lock or key)	45	53	45	18
3. Point in transaction	11	2	0	9
4. Type of information	2	18	27	27
5. Type of web site	2	0	27	0
6. General distrust	5	0	0	18
7. Blue line	3	4	0	0
8. Amount/presence of Information	1	0	0	0
9. Accessibility of site	2	0	0	9
10. Text from web site	6	0	0	9
11. Alerts on screen	2	2	0	0
12. Security conventions	1	0	0	0
13. Transaction completed	1	2	0	0
14. Unspecified	3	0	0	0
15. Uncodeable	2	0	0	0

[a] Some participants provided multiple types of evidence. All types of evidence were coded for each participant.

Secure connections were recognized by roughly half the participants evenly across the three communities. In contrast, nonsecure connections were correctly recognized more frequently by high-technology participants (92%) than by either rural (59%) or suburban (50%) participants ($p < .05$).

Visual portrayal of a secure connection

Finally, to elicit participants' models about web security, participants were asked to revise a drawing of the Web that they had made earlier in the interview to reflect a secure connection. Participants sketched primarily five different representations:

- Screenshot (12%), a symbol on the screen such as the key icon
- Direct connection (12%), a direct line between the user's computer and the target web site

- Secure boundary (14%), a barrier, such as a firewall, that surrounds or protects the user's computer, a server, or the target web site
- Encryption (40%), scrambling the information while it is in transit, including both message encoding and decoding in more sophisticated drawings
- No difference (11%), drawings that remained unchanged from the participant's initial drawing

Participants' drawings were then analyzed for their representation of a secure connection as something that applies to information while it is in transit from one machine to another (a correct understanding) (see Figure 24-2), or as something that applies to a specific "place" on the Web (an incorrect understanding) (see Figure 24-3).

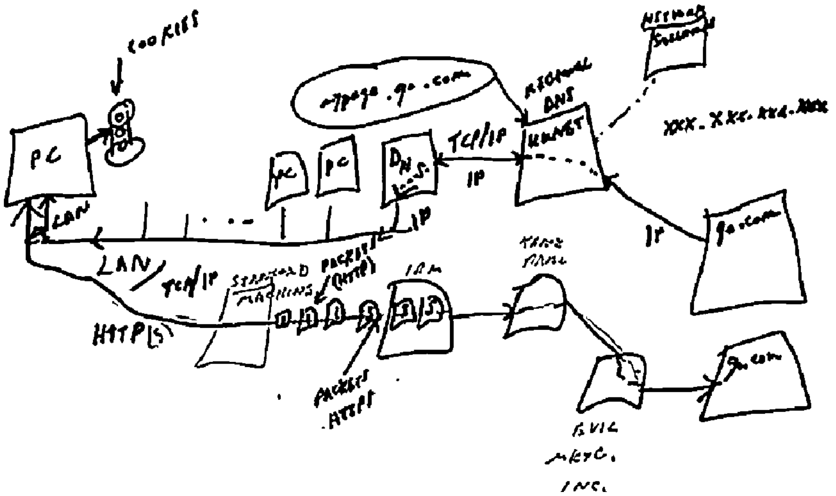

FIGURE 24-2. Participant drawing showing security as transit; the drawing shows a secure connection in terms of encryption while the information is in "transit"; the darker solid lines represent the secure connection

High-technology participants (74%) provided transit (i.e., correct) representations more frequently than did either rural (33%) or suburban (46%) participants ($p < .05$).

Reflections

Based on empirically derived typologies, results (Table 24-1) suggest that many users across diverse communities inaccurately evaluated a connection as secure when it was not, and vice versa. In addition, users who correctly recognized connections as secure or not secure sometimes did so for incorrect reasons. Furthermore, the high-technology participants did not always have more accurate or sophisticated conceptions of web security than did their rural and suburban counterparts.

Through this study, we highlighted two main points: that informing can happen through interaction, and that users develop mental models that shape their understanding about the system and about its security.

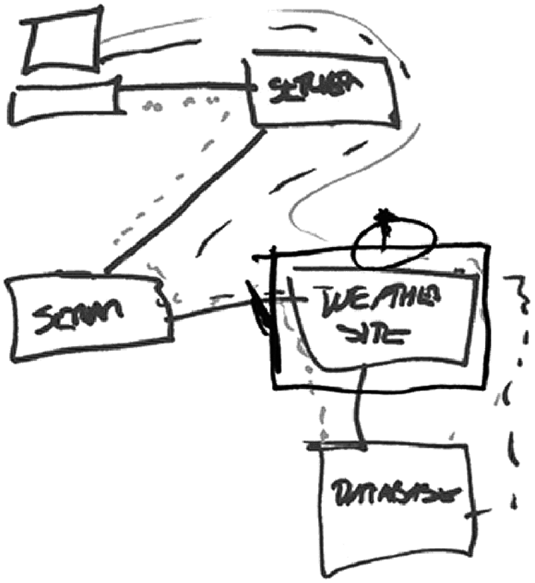

FIGURE 24-3. Participant drawing showing security as a place; the drawing shows a secure connection in terms of a secure boundary around a specific "place" on the Web; the darker solid lines represent the secure connection

We also mentioned that poor, inadequate, or misguided mental models about security may lead to poor privacy behavior—believing personal information is adequately protected when it is not, or not taking the appropriate actions to protect personal information. These negative consequences could be reduced through web browser design that helps users construct more accurate understandings of a secure connection. Such design work can profit from this study's typologies. For example, the most frequently used icons to represent the security status of a connection—the key or padlock—convey the idea of a "place" that can be made secure. Such a conception runs counter to the more accurate meaning of a secure connection (referring to the security of the information in transit). More generally, well-designed interactions—conceptualized to match the underlying security model and validated empirically with users—can lead users to construct reasonable models for system security. In so doing, the interaction design can go a good distance toward tacitly informing users of potential privacy and security risks.

The Scope of Informed Consent: Questions Motivated by Gmail

In the first two cases, we provided "proof-of-concept" projects for ways in which the information systems community can design for informed consent. In our third case—Google's Gmail web mail (web-based email) system—we examine the scope of informed consent—namely, are there issues concerning privacy and security that informed consent cannot reasonably address? And, if so, how do these issues affect informed consent?

What Is Gmail?

Gmail (*http://www.gmail.com*) is Google's web mail system, currently in beta testing. Similar to other free email services, such as those provided by Yahoo! Mail or Hotmail, Gmail provides three key additional features:

- A larger amount of storage space for one's email than is typically provided (as of December 2004, 1 GB of storage as compared to Yahoo! Mail's 250 MB or Hotmail's 250 MB)
- The ability to use Google search technology to search one's email messages
- Grouping an email and the replies to it as a conversation

As with most free web mail services, Gmail subscribers see advertising alongside their email. However, unlike other free web mail providers, Gmail determines which advertisements to display based on the *content* of the subscriber's email message. For example, if a Gmail subscriber receives an email message from a friend asking to borrow the subscriber's bicycle, the subscriber would likely see advertisements related to online bicycle vendors alongside the email message (see Figure 24-4).

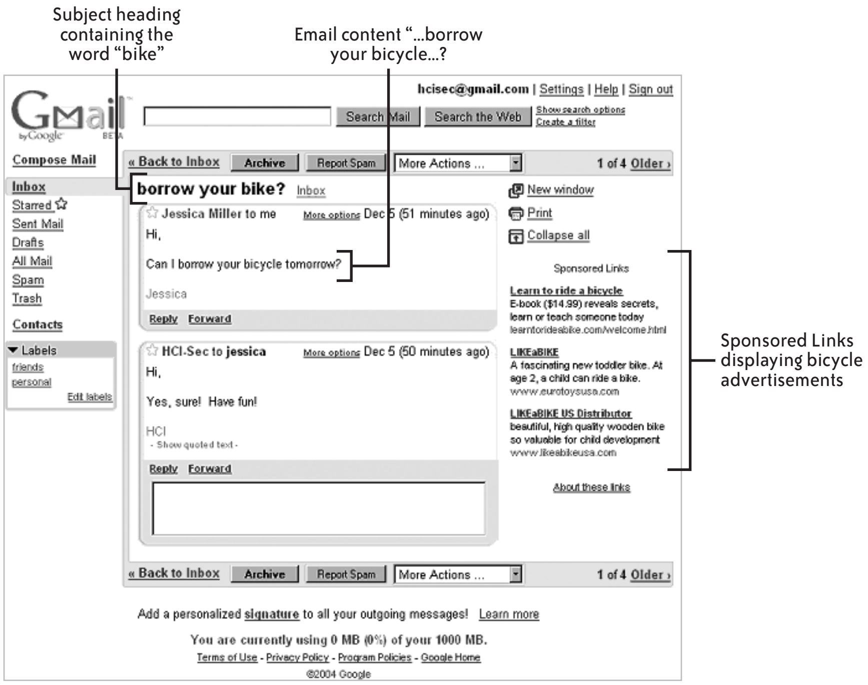

Subject heading containing the word "bike"

Email content "...borrow your bicycle...?

Sponsored Links displaying bicycle advertisements

hcisec@gmail.com | Settings | Help | Sign out

borrow your bike? Inbox

Jessica Miller to me More options Dec 5 (51 minutes ago)

Hi,

Can I borrow your bicycle tomorrow?

Jessica

HCI-Sec to jessica More options Dec 5 (50 minutes ago)

Hi,

Yes, sure! Have fun!

HCI
- Show quoted text -

Sponsored Links

Learn to ride a bicycle
E-book ($14.99) reveals secrets, learn or teach someone today
learntoridesbike.com/welcome.html

LIKEaBIKE
A fascinating new toddler bike. At age 2, a child can ride a bike.
www.eurotoysusa.com

LIKEaBIKE US Distributor
beautiful, high quality wooden bike so valuable for child development
www.likeabikeusa.com

About these links

Add a personalized signature to all your outgoing messages! Learn more

You are currently using 0 MB (0%) of your 1000 MB.

Terms of Use - Privacy Policy - Program Policies - Google Home

©2004 Google

FIGURE 24-4. Viewing an email message in Gmail; advertisements targeted to the content of the body of the message—in this case, bicycles—appear to the right of the email message, under the label Sponsored Links

How Gmail Advertisements Work

Let's look briefly at both the Gmail business model and its technology:

- *The business model.* In a nutshell, Google charges advertisers only when a user clicks on an advertisement. When advertisers submit their advertisements to Google, they negotiate a "rate-per-user-click" and designate keywords they believe to be most relevant to their advertisement.[25]
- *The technology.* The following automated process occurs dynamically each time a Gmail subscriber clicks on an email entry. The Gmail system retrieves the message and scans the text (attachments are not scanned) for keywords (provided earlier by advertisers) in the body of the email message.[26] Based on the results of scanning the message, as well as on how well the advertisement keywords match the email content, the amount advertisers pay per user-click, and the prior click-through rate (i.e., the number of clicks divided by the number of times the advertisement has been displayed), Google computers select and determine the order in which to display advertisements. The selected advertisement is displayed near the subscriber's message; no link is established between the advertisement and the email message. The only information Google relinquishes to its advertisers is the number of times their advertisement was chosen for display and the click-through rate. No personal information about subscribers or the content of email messages is released to advertisers.

Gmail and the Six Components of Informed Consent

In the context of this business model and technical implementation, how well does Gmail meet the criteria for informed consent? To make this assessment, we analyzed Gmail's Terms of Use,[27] Program Policy,[28] and Privacy Policy[29] in relation to Gmail's functionality as described above. Between these documents, Gmail's registration interface, and the user interface, Google addresses each of the different components of informed consent in a reasonably thorough manner.

Disclosure

As of December 2004, the Gmail privacy policy explicitly states:

What information will be collected about subscribers and their activities

For example, "personal information including your first and last name, a user name…and a password," "log information…browser type you used, your Internet

[25] Google, *Google Adwords* (2004) [cited Dec. 2004]; *https://adwords.google.com/select/*.

[26] According to Ana R. Yang (product marketing manager, Gmail), personal email (Dec. 4, 2004).

[27] Google, *Gmail Terms of Use* (2004) [cited Dec. 2, 2004]; *http://gmail.google.com/gmail/help/terms_of_use.html*.

[28] Google, *Gmail Program Policies* (2004) [cited Dec. 1, 2004]; *http://gmail.google.com/gmail/help/program_policies.html*.

[29] Google, *Gmail Privacy Policy* (2004) [cited Nov. 15, 2004]; *http://gmail.google.com/gmail/help/privacy.html*.

Protocol address, and the date and time of day…the unique ID provided by our cookie and the URL of the last site you visited."

Who will have access to the information

For example, "No human reads your email…without your consent," "Google will never sell, rent or share your personal information…with any third parties for marketing purposes without your express permission."

What the information will be used for

For example, subscribers' accounts will be used "internally to deliver the best possible service to you, such as improving the Gmail user interface, preventing fraud within our advertising system, and better targeting related information."

How the identity of the individual will be protected

For example, employees who have access to user information are monitored and educated to protect the user's privacy.[30]

Overall, Google discloses an impressive amount of the right type of information to enable informed consent. We note one anomaly: Google provides only vague details about how long information is kept.

Comprehension

The Privacy Policy uses fairly clear, jargon-free English; this goes a good distance toward helping ensure comprehension. To further support comprehension, Google could provide an opportunity for dialog and clarification about Gmail policies and practices by publishing a way to contact the appropriate Google personnel (e.g., email address, phone number, or online chat).

Voluntariness

Gmail subscribers' consent can be considered voluntary for two reasons:

- Gmail's Terms of Use and Privacy Policy contain reasonably neutral language that does not attempt to coerce or manipulate users to use Gmail.
- Other free web mail services are available, albeit with substantially less storage space and perhaps not as high-quality search technology, that do not scan users' email for the purpose of displaying content-targeted advertisements.

Competence

Competence is difficult to assess in online interactions, and how to do so remains an open problem for the information systems community. Within the bounds of current knowledge and practice, Gmail's Terms of Use addresses the competence of minors by stating:

[30] *Ibid*.

"Due to the Children's Online Privacy Protection Act of 1998 (which is available at *http://www.ftc.gov/ogc/coppa1.htm*), you must be at least thirteen (13) years of age to use this Service."[31]

Agreement

Google requires explicit consent with Gmail's Terms of Use before the user receives a Gmail account. In practice, the user must click a large button at the end of the registration process signaling agreement to the Terms of Use. However, the Terms of Use does not provide information about whether subscribers can withdraw from the agreement and, if so, whether it is possible to delete their Gmail accounts. Indeed, it is possible to do both, but subscribers must be able to navigate through several different layers of menus to do so.

Minimal distraction

Setting aside users' tendency not to read agreement policies during online registration processes, we address the criterion of minimal distraction under the confinements of online interactions. Links to Gmail's Terms of Use and Privacy Policy are always at the bottom of a user's email, thus making it easy for the user to see what he has agreed to while not being distracted from the task at hand: reading and sending email.

Two Questions Related to Informed Consent

Despite Google's reasonable handling of informed consent, privacy advocates have made claims that scanning personal email to place content-targeted advertisements is still a privacy invasion. Their claims revolve around two primary questions:[32]

- Does using machines to read personal email constitute a privacy violation?
- Should the consent of indirect stakeholders (i.e., email senders) also be obtained?

We turn now to discuss these two questions and their relation to informed consent.

The question of machines reading personal content

This question is not so much one of informed consent as it is one of privacy. Namely, what do we mean by privacy and, in turn, who or what can violate a person's privacy? Consider the following two definitions of privacy:

- *Parent's definition.* W. A. Parent defines privacy as the "condition of not having undocumented personal information about oneself known by others."[33] This definition

[31] Google, *Gmail Terms of Use.*

[32] Chris Jay Hoofnagle, Beth Givens, and Pam Dixon, *Violation of California Civil Code § 631 by Google Gmail Service*, Electronic Privacy Information Center (EPIC, 2004) [cited Nov. 15, 2004]; *http://www.epic.org/privacy/gmail/agltr5.3.04.html*.

[33] W. A. Parent, "Recent Work on the Conception of Privacy," *American Philosophical Journal Quarterly* 20 (1983), 341.

requires an “other” who can “know” something. Putting aside artificial intelligence claims for machine agency, this definition implies that the information must be imparted to another human. Google assumes such a definition when making the claim that Gmail’s content-targeted advertising does not violate privacy because no *human being* ever reads subscribers’ email.

- *Bloustein’s definition.* In contrast, consider an alternative definition by Edward Bloustein: privacy preserves one’s inviolate personality, independence, and dignity.[34] For Bloustein, a privacy violation concerns the act of intrusion upon the self, independent of the state of mind (or knowledge) of the intruder. Privacy advocates troubled by Google’s practices likely share Bloustein’s view of privacy when they claim that by extracting meaning from a subscriber’s email message—whether with a machine reader or a human reader—Google is violating privacy. Moreover, some privacy advocates are more concerned about machine than human readers because of computer efficiency, memory capacity, and power.

The question of indirect stakeholders

Many privacy advocates also assert that content-targeted advertisements should not be allowed without the consent of *all* parties involved in an email exchange. While Gmail does obtain the consent of the subscriber, it does not obtain the consent of the email sender. Yet personal content from an email sender is scanned to generate content-targeted advertisements to the email receiver.

Google argues that Gmail’s practices are not a violation of the sender’s privacy “since no one other than the recipient is allowed to read their email messages, and no one but the recipient sees targeted ads and related information.”[35] A further argument against obtaining sender consent could be made by adopting the United States’ model of ownership for physical mail. As soon as physical mail is delivered to receivers, it is considered to be the receivers’ property; at that point, receivers may do whatever they please with the mail. Thus, the question becomes, “Is the model of ownership for physical mail an adequate or desirable model for ownership of email?”

Reflections

In this case, we have argued that privacy advocates’ claims against Gmail are not claims about informed consent per se, but rather, claims that concern differing views of privacy and ownership. Google has done a reasonable job obtaining informed consent under its notions of privacy and property, but it is possible that these notions are inappropriate for email and need to be reconsidered. Overall, the case of Gmail illustrates that at times,

[34] Edward J. Bloustein, “Privacy As an Aspect of Human Dignity,” in Ferdinand David Schoeman (ed.), *Philosophical Dimensions of Privacy: An Anthology* (Cambridge University Press, 1964), 156–202.

[35] Google, *More on Gmail and Privacy* (2005) [cited July 30, 2005]; *http://mail.google.com/mail/help/more.html*; <*https://iexchange.ischool.washington.edu/exchweb/bin/redir.asp?URL=http://mail.google.com/mail/help/more.html*>.

considerations beyond the scope of informed consent (e.g., privacy, ownership) need to be understood in order to ascertain when informed consent is relevant.

Design Principles for Informed Consent for Information Systems

In our view, design refers not only to the design process and resulting technology, but also to policy and business practice. Given this broad context for design, and based on the model of informed consent for information systems and the three cases, we now distill 10 design principles:

1. Decide what technical capabilities are exempt from informed consent.

   Obtaining users' informed consent comes with a high cost to users. Information must first be disclosed and comprehended by users, who must then have an opportunity to agree or decline. And while all of this is being done, the user has been diverted from the task at hand—the thing that the user really wanted to do. Let us refer to these costs for obtaining the user's informed consent as "the nuisance factor." If all interactions with information systems required explicit informed consent, the nuisance factor would be unmanageable. Users would simply crumble under the burden.

   Fortunately, a fair number of interactions with information systems may be considered exempt from the need to obtain users' informed consent. But how are designers to determine which interactions are exempt? While there are no hard and fast rules, the Belmont Report[36] on human subjects research offers some useful guidelines. According to the report, an individual's participation is considered exempt from the need to obtain informed consent when the following three conditions are met:

- Participation can in no way put the individual in physical, legal, psychological, or social jeopardy. To this list of harms, for interactions with information systems we also include that participation does not place the individual's privacy, data, or hardware in jeopardy.
- The purpose and sponsorship of the activity is known (or clearly stated to the individual).
- No coercion is involved.

   Designers can invoke these three criteria to scrutinize information system–based interactions for exemption from informed consent. Granted, it may be difficult to make these judgments in advance; however, defensible assessments should be made, and if the initial assessments are in error, then remedies should be implemented.

[36] The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, "The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research" (1978).

2. Invoke the sanction of implicit consent with care.

   On the surface, implicit consent seems a reasonable umbrella to cover most online interactions. After all, users do not interact online—as in the canonical example—with a gun held to their heads. However, unless users have comparable (with respect to costs such as time, effort, knowledge, and expense) alternative access to comparable services, products, information, and so forth, then information system use may not be regarded as wholly noncoercive. Given the rapidity and widespread movement with which access to goods and services has moved online and the corresponding movement to discontinue or minimize traditional means of access, the viability of alternative comparable access to goods and services is at times slim and getting slimmer. In this climate, we advocate presuming that implicit consent is not a viable option, and only in special circumstances and after careful consideration invoking the sanction of implicit consent.

   Further challenges for implicit consent arise from the criterion of disclosure. The disclosure issue can be understood as follows: although the user may be told what mechanisms are enabled, the user may not be aware of the full implications of those mechanisms. For example, while many users were aware of and enabled cookies, few users understood the implications of cookies for individual privacy (until an extensive public discussion took place).

3. Understand the scope of informed consent and how informed consent interacts with other values.

   Informed consent concerns the important but limited activities of “informing” those affected by using, or the use of, an information system and obtaining their “consent” for that participation. Thus, informed consent closely interacts, but is not synonymous, with other values such as privacy, trust, and security. Nonetheless, how those other values are defined (e.g., can a machine invade a person’s privacy?) has implications for what activities may require consent—for example, obtaining user consent for machine reading of personal content in an email message.

4. Consider both direct and indirect stakeholders.

   Designers and usability engineers too often focus only on those users who interact directly with the information system without considering the system’s impact on others, whom we refer to as indirect stakeholders. At times, it may be important to obtain informed consent from indirect as well as direct stakeholders (see the earlier sidebar, “Value Sensitive Design”).

5. Put users in control of the “nuisance factor.”

   Different users place differing degrees of importance on different types of harm. Correspondingly, how much of a nuisance factor a user is willing to tolerate to obtain informed consent will depend on the particular user. Rather than mandating a single mechanism for obtaining informed consent for all users in all situations, designers need to provide users with a range of mechanisms and levels of control so that users are positioned to manage the nuisance factor in accordance with their concerns.

Successful designs will likely contain a reasonable balance among overarching controls (e.g., “never accept cookies” or “always accept cookies”), micromanaged controls (e.g., “ask about each cookie”), and intermediate controls that mix well-chosen overarching controls with micromanaged subsets (e.g., “decline all third-party cookies and ask me about all other cookies”).

6. Defaults matter.

   It is well established that most users do not change preference settings. Thus, default settings should err on the side of preserving informed consent. Notably, for many years, the default setting for cookies on current browsers was “accept all cookies,” which neither informs nor obtains consent from users. Default settings will also need to take into account the nuisance factor to obtain users’ informed consent (see Design Principle 5).

7. Avoid technical jargon.

   Follow the well-established interface principle to avoid technical jargon in favor of clear language that directly addresses the user’s values, needs, and interests.

8. Inform through the interaction model.

   In addition to using words to explicitly inform users, take advantage of the mental models that users construct through their interaction with a system to reinforce reasonably accurate conceptions of how information is secured and flows through the system. Avoid interaction models that may lead to misleading or ambiguous user conceptions of information flow (e.g., in web browsers, the use of locks that suggest a secure “place” to indicate a secure connection for information while it is in “transit” over the Internet).

9. Field test to help ensure adequate comprehension and opportunities for agreement.

   Because information systems will likely rely on automated means to realize informed consent, designers face significant challenges in ensuring adequate disclosure, comprehension, and opportunities for agreement. Thoughtful interface design will need to be coupled with equally thoughtful field tests to validate and refine the initial designs. Moreover, because informed consent carries a moral imperative, the components of disclosure, comprehension, and agreement need to work reasonably well for all users. Thus, it becomes a requirement (and not simply better practice) to include a reasonable range of both representative and atypical users in the field tests.

10. Design proactively for informed consent.

    More frequently than not, information systems are conceived of and implemented without consideration of informed consent. Once introduced, practices evolve around these new interactions, and these, too, develop without consideration of informed consent. When issues of informed consent at last come to the fore, designers face a nearly insurmountable task: to retrofit the information system capability and interaction. The solution, in part, is to design proactively for informed consent.

Acknowledgments

This work was funded, in part, by NSF Award IIS-0325035.

About the Authors

Batya Friedman is a professor in the Information School and an adjunct professor in the Department of Computer Science and Engineering at the University of Washington, where she co-directs the Value Sensitive Design Research Laboratory. Dr. Friedman's research interests include human-computer interaction, especially human values in technology and design methodology. Her recent work has focused on informed consent, privacy, trust, and human dignity engaging such technologies as web browsers, urban simulation, robotics, open source, and location-enhanced computing.

http://www.ischool.washington.edu/vsd/

Peyina Lin is a Ph.D. student in the Information School at the University of Washington, where she works in the Value Sensitive Design Research Laboratory. Ms. Lin holds an M.Sc. in information management (Syracuse University) and an M.A. in telecommunication digital media arts (Michigan State University). Her research investigates the impact of information technologies on community, sense of place, and value tradeoffs, such as privacy-security, and privacy-social awareness.

http://students.washington.edu/pl3/

Jessica K. Miller is a Ph.D. student in the Department of Computer Science and Engineering at the University of Washington, where she works with Batya Friedman and Alan Borning in the Value Sensitive Design Research Laboratory. Ms. Miller is principally interested in human-computer interaction, and specifically how user-centered design methodologies can be used to develop technologies that impact social relationships such as the Internet, wearable computing devices, and pervasive computing.

http://www.cs.washington.edu/jessica/