

An Empirical Study of Consumer Perceptions and Comprehension of Web Site Privacy Policies

Matthew W. Vail, Julia B. Earp, *Member, IEEE*, and Annie I. Antón, *Senior Member, IEEE*

Abstract—The U.S. legislation at both the federal and state levels mandates certain organizations to inform customers about information uses and disclosures. Such disclosures are typically accomplished through privacy policies, both online and offline. Unfortunately, the policies are not easy to comprehend, and, as a result, online consumers frequently do not read the policies provided at healthcare Web sites. Because these policies are often required by law, they should be clear so that consumers are likely to read them and to ensure that consumers can comprehend these policies. This, in turn, may increase consumer trust and encourage consumers to feel more comfortable when interacting with online organizations. In this paper, we present results of an empirical study, involving 993 Internet users, which compared various ways to present privacy policy information to online consumers. Our findings suggest that users perceive typical, paragraph-form policies to be more secure than other forms of policy representation, yet user comprehension of such paragraph-form policies is poor as compared to other policy representations. The results of this study can help managers create more trustworthy policies, aid compliance officers in detecting deceptive organizations, and serve legislative bodies by providing tangible evidence as to the ineffectiveness of current privacy policies.

Index Terms—E-commerce, end user computing management, healthcare, privacy management, privacy policy, trust, user computer interaction.

I. INTRODUCTION

THE INTERNET has created a new avenue for organizations to leverage technology to create new revenue streams, lower the cost of doing business, improve customer satisfaction, and attract new customers [10]. Objectives such as these often compel managers to collect large amounts of customer information to aid in their strategic decision making processes [30]. Modern marketing, for example, employs the Internet to aid in the collection and aggregation of information to personalize online experiences and deliver customized services [10]. Personalized services may be critical to the success of customer relationship management [10]; therefore, the more information

gathered about customers, the better a company is able to meet their customers' needs [36].

Increasingly, consumers are concerned that their information will be disclosed to third parties [45] or used for purposes other than those for which it was collected [48]. Consumer concerns have only heightened as media coverage of consumer privacy issues has increased over the past decade [44]. The public has been sensitized to privacy breaches such as JetBlue [7], which violated its privacy policy by disclosing the travel records of five million customers to a government contractor, and RealJukeBox, which collected personally identifying information, matched it with file information from users' PCs, and then, sold the combined data [48]. Consequently, consumers are hesitant to share their personal information with online companies because they fear that their privacy will be violated [10], [42], [47], [48].

Studies have repeatedly shown that lack of consumer trust in online companies has a direct effect on their purchase behaviors and is a major obstacle to the success of e-commerce business models [9], [13], [21], [33], [41]. Thus, a company must garner trust from the consumer to minimize the perceived risk of privacy violations. One way that companies seek to increase trust is by posting a privacy policy notice on their Web site. The simple existence of a privacy policy, regardless of content, has been shown to allay some consumer privacy concerns [4]. However, simply posting a privacy policy is not sufficient to establish consumer trust.

A 2001 Forrester study showed that almost half of online consumers read privacy policies prior to making an online purchase [18], and consumers have been found to have more trust in privacy policies that they perceive as comprehensible [34]. Therefore, to increase consumer trust, it is essential that companies post privacy policies that are both concise and comprehensible. However, research from the perspective of both consumers [33] and researchers [6], [8], [27], [47] has shown that current Internet privacy policies are often too long and confusing.

In this paper, we address the need for online organizations to create more concise and comprehensible privacy policies by reporting on an empirical investigation in which we examined consumer perception and comprehension of typical online privacy policies (henceforth, we refer to these as typical online privacy policies (TOPPs) currently displayed on major Web sites. We also examined three types of alternative privacy policy representations not commonly used [henceforth, we refer to these as atypical privacy policies (APPs)] and compared consumer perception and comprehension of TOPPs versus APPs. Four key research questions drove our empirical study.

1. How well do consumers comprehend TOPPs?
2. Do consumers comprehend APPs better than TOPPs?

Manuscript received May 1, 2006; revised October 1, 2006. Review of this manuscript was arranged by Department Editor A. Chakrabarti. This work was supported in part by the U.S. Department of Commerce under Grant BS123456 and in part by the National Science Foundation (NSF) under Information Technology Research (ITR) Grant 0325269.

M. W. Vail is with the National Institute of Standards and Technology, Boulder, CO 80305 USA.

J. B. Earp was with Rice University, Houston, TX 77005 USA. She is now with the Department of Physics, Colorado State University, Fort Collins, CO 80523 USA.

A. I. Antón is with the Electrical Engineering Department, University of Colorado, Boulder, CO 80309 USA, on leave from the National Research Institute for Metals, Tsukuba, Japan.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TEM.2008.922634

3. Are consumers' perceptions of APPs more positive than those of TOPPs?
4. Does consumer comprehension of a policy align with their perception of the policy?

The results of this study are expected to aid companies in bolstering consumer trust in their organization in order to encourage increased product purchasing and revenue. Furthermore, the results can help managers determine which policy representation is most appropriate for their organization and can aid the Federal Trade Commission (FTC) in refining and enforcing the fair information practices (FIPs) [20], the principles that serve as the basis for privacy best practices in the U.S.

The remainder of this paper is organized as follows. First, we discuss relevant research on privacy policies, readability, organizational trust, and healthcare legislation. Next, we cover the methodologies employed to conduct our empirical investigation, as well as the results of the experiment. Finally, we discuss the results of the experiments and the implications of this research for managers and software designers.

II. BACKGROUND AND HYPOTHESES DEVELOPMENT

This section provides an overview of the relevant related work, including Internet privacy policies, the U.S. Health Insurance Portability and Accountability Act, trust and its implications for vendors, and previous examinations of Internet privacy policy readability. In this section, we also motivate and present six hypotheses that guided our empirical assessment of consumer comprehension and perception of different privacy policy representations. The hypotheses are tested and the results interpreted in Sections III and IV, respectively.

A. Privacy Policies

Privacy policies are notices that are posted on a Web site, accessible to the public, and describe an organization's information practices—how they collect, use, and disclose information. Though not always true in practice, privacy policies are supposed to reflect a Web site's actual privacy practices and serve as a contract between the Web site and the consumer. These policies are of great value to consumers because they are often the only means of gaining insight into an organization's data collection and disclosure practices.

In 1998, an FTC survey revealed that only 14% of the Web sites surveyed disclosed their privacy practices [16]. A similar survey conducted by the FTC in 2000 revealed that 88% of the Web sites surveyed disclosed their privacy practices [15]. In the absence of a legislative requirement, this increase may indicate that companies recognized consumers' privacy concerns and attempted to dispel those fears by posting privacy policies. Since 2000, federal and state regulations have been enacted that require Web sites to post a conspicuous, plainly written privacy policy on their Web site. Although federal regulations requiring privacy policies are currently restricted to the domains of healthcare, finance, and child protection, other organizations are obligated to conform to similar state-level policy requirements.

B. Trust: Implications for E-Vendors

Trust has been defined as "the belief that the trustee will act cooperatively to fulfill the trustor's expectations without exploiting its vulnerabilities" [40]. There is extensive literature regarding the nature of trusting beliefs, but the majority can be viewed as a three-dimensional (3-D) set of specific beliefs, which are generally classified as *ability*, *benevolence*, and *integrity* [21]. *Ability* beliefs reflect confidence that a trustee has the capability of performing the job; *benevolence* beliefs reflect confidence that a trustee will not act opportunistically, even when given the chance; and *integrity* beliefs reflect confidence that a trustee will act honestly and keep its promises [40].

Ability, benevolence, and integrity trusting beliefs can be influenced by a number of factors, including Web site usability [46], signaling [46], and reputation [4]. Furthermore, there can exist a correlation between ability, benevolence, and integrity trusting beliefs. For example, if an organization is deceptive and does not keep promises made to a consumer, the consumer's integrity beliefs are affected, as well as the benevolence beliefs. Benevolence and integrity beliefs have also been found to increase when a Web site posts a strong privacy/security statement [46]. This finding, in conjunction with another study [34], indicates that a comprehensive, concise, and comprehensible privacy policy is a key component in building consumer trust in an online organization. This finding may be of particular interest to smaller e-vendors who lack an established reputation that may have otherwise afforded them inherent ability, benevolence, and integrity trusting beliefs.

The role of trust in establishing a relationship with the consumer is both challenging and pivotal to e-commerce business models [9], [21], [24], [33], [41]. Trust has been shown to be crucial in transactional, buyer-seller relationships that include an element of risk, such as those in e-commerce [43]. To consumers, the risk involved with e-vendors are typically high because the nature of the communication is generally devoid of social interaction [21], making it difficult, if not impossible, for consumers to verify that they are communicating with a trusted party. Consumers are circumspect about engaging in transactions with e-vendors with whom they have not previously established a trusting relationship because there are no guarantees that the e-vendor will not engage in harmful or opportunistic activities. In fact, researchers have suggested that consumers avoid e-vendors that they do not trust [43]. Therefore, trust is a significant antecedent of online commerce participation [21], and thus, a crucial aspect of e-commerce success.

Trust has a considerable effect on the purchasing behavior of consumers and the e-vendor's subsequent revenue; but, e-commerce success begins by sending positive signals (e.g., privacy seals, privacy policy, etc.) to the consumer to foster positive trusting beliefs. Consumers must perceive an organization to be trustworthy enough to overcome the potential risk of privacy violations before they feel comfortable enough to share their personal information. As previously mentioned, these fears can be reduced by posting a clear and concise privacy policy on one's Web site. Based on these findings, we propose the following three hypotheses:

- H1) Individuals feel more secure sharing personal information with organizations that display an APP than with organizations that display a TOPP.
- H2) Individuals feel that organizations that display an APP are more likely to protect their information than organizations that display a TOPP.
- H3) Individuals feel more confident in their understanding of an APP than a TOPP.

In our study, two of the three APPs we constructed were significantly more concise than TOPPs. The APPs we have developed contain the same content as the TOPPs, but they are presented to participants using an alternative format. Based on previous work that has shown users associate policy length with quality [4], we propose the following additional hypothesis.

- H4) Individuals feel that TOPPs are explained more thoroughly than the concise APPs.

C. Interactional Justice Perspectives on Consumer Privacy

Culnan and Bies [12] discuss the role of justice perspectives as a useful theoretical framework for analyzing consumer privacy concerns during an exchange that occurs for an economic benefit. They discuss three types of justice relevant to consumer privacy—distributive, procedural, and interactional. *Distributive justice* refers to the perceived fairness of the overall outcome of a transaction from the consumer's perspective; *procedural justice* refers to the fairness of the procedures and how those procedures are enacted; and *interactional justice* refers to the fairness of interpersonal treatment that an individual receives from another over the course of a transaction or relationship [12], [32]. In this paper, we are primarily concerned with interactional justice; we seek to demonstrate new approaches that support consumers' expectations of fairness, rather than the overall outcome or procedures involved in a transaction.

In the context of consumer privacy in the United States, the FIPs address justice concerns by providing guidelines for privacy protection and seeking to ensure that consumers are treated fairly. The FIPs consist of five principles that serve as the basis for best practices in privacy: notice/awareness, choice/consent, access/participation, security/integrity, and enforcement/redress [20]. In this study, we categorize the content of privacy policies according to the privacy taxonomy developed by Antón *et al.* [5]. This privacy taxonomy uses the FIPs as a foundation and expounds upon them to provide a more exhaustive privacy taxonomy.

An organization adhering to these principles signals to consumers that the organization can be trusted. The most fundamental FIP principle is notice/awareness, and it states that consumers should be given notice of an entity's privacy practices prior to collecting information from them. Such notices usually manifest in the form of privacy notices (e.g., privacy policy). Privacy notices have been cited as influencing whether consumers are likely to initiate a relationship with one organization over its competitors [10]. Therefore, posting a clear and comprehensible privacy policy may positively influence consumers' sense of interactional justice, thereby increasing trust in the organization.

D. Readability

Recent investigations into privacy policy composition and complexity found that in both the financial and healthcare domains, no Web site privacy policy was comprehensible by most English-speaking people in the U.S. [6], [8], [22], [27], [47]. On average, financial and healthcare related Internet privacy policies require consumers to have a reading equivalency of at least 2 years of college education to fully comprehend them [6], [8], [22], [27], [47]. However, studies have shown that only 52.1% of the general population has obtained this level of education [37]. Furthermore, literacy research has shown that many people read three–five grades lower than their highest level of schooling [23]. These findings suggest that the majority of the U.S. population is unable to comprehend the content of most Web site privacy policies. Consequently, we believe that users are unable to make informed decisions regarding the Web sites with which they can safely share their information.

Attempts to overcome problems associated with privacy policies, and thus, reduce the burden on users, have resulted in machine-readable policy specification languages, such as the Platform for Privacy Preferences (P3P) [49] and A P3P Preference Exchange Language (APPEL) [31]. The P3P is a specification language that allows organizations to specify their privacy practices in a machine-readable Extensible Markup Language (XML) format, and APPEL is an exchange language that allows users to specify their privacy preferences. P3P-enabled policies can be read by automated user agents (e.g., privacy critics [2] and Privacybird [11]), but they only alert users of policies that are likely to cause user concern. In order to take advantage of P3P's capabilities, the Web site and the user must both be willing and able to use the appropriate tools. If this is done properly, it allows the user to exercise preferences about the Web site's privacy practices. The idea is that by filtering out the noise and focusing users' attention on the privacy elements that are in contrast to the user's previously stated preferences, users are more likely to be engaged. However, researchers have identified numerous challenges that have hindered the widespread acceptance of P3P and related tools [3], [25], [28], [50]. Despite its shortcomings, P3P will provide some important lessons for the design of future solutions [1]. Until these solutions are user-friendly and widely deployed by organizations, Web sites need to consider restructuring existing, paragraph-form privacy policies to benefit users as they interact with the Web site. Based upon this prior research, we formulated the following hypothesis:

- H5) Individuals comprehend APPS better than TOPPs.

Given that hypotheses H1–H4 expect users to feel more positively toward APPs than TOPPs, and hypothesis (H5) expects users to comprehend the APPs better than TOPPs, it follows that we also expect user perception and comprehension to be in alignment. Based on this, we formulated the following hypothesis:

- H6) Individuals feel more positively toward more comprehensible privacy policies.

III. RESEARCH METHODOLOGY

The objectives of our empirical study were: 1) to gauge user perception of various alternatives to natural language privacy policies; 2) to measure user comprehension of the alternatives; and 3) to compare user perception with user comprehension to determine whether perception and comprehension are aligned.

A. Designing APPs Using Content Analysis Approach

We developed our various APPs using a goal-mining process. Goal mining is traditionally employed to extract goals from data sources by applying goal-based requirements analysis methods [8]. Goal-mining heuristics have also been applied to privacy policies to extract and classify privacy goals in order to analyze organizational privacy practices [5], [6], [8]. The goal-mining process begins by analyzing a policy document to gain an indepth understanding. The first step involves extracting goals by analyzing each statement within the policy document and asking, “What privacy goal(s) does this statement exemplify?” and “What privacy goal(s) does this statement obstruct or thwart?” Any statement that contains an *action word* (verb) is a goal candidate. Consider the following excerpt from the Drugstore.com¹ privacy policy:

We may *enter* into an agreement with other companies or include individuals to perform functions on our behalf. These functions may include sending promotional e-mails on our behalf to such companies customers²; serving advertisements on our behalf; providing marketing assistance; processing credit card payments; and have access to information necessary to perform their functions.

When analyzing this statement, we identify action words. The first action word in this example is “enter.” We then identify *actors* (person performing the action), *targets* (who the action is being performed on), *instruments* (how the action will be performed), and *purpose* (why the action is being performed). By asking the goal identification questions, we extracted the following goals from the previous example.

- G₁: USE customer email address for marketing and promotional purposes.
- G₂: SHARE customer information (CI) with subsidiaries to recommend services to customer.
- G₃: SHARE CI with third parties to perform marketing services on our behalf.
- G₄: SHARE CI with third parties to provide valuable financial services we do not offer (e.g., credit card).

After identifying the goal components, we classify each goal statement according to *keywords* and express it in structured natural language statements [8]. Each action word maps to a keyword in the privacy taxonomy in [5]. In the aforementioned example, the action word “sending” maps to the keyword “SHARE.”

Identified goals are classified as either protection goals or vulnerabilities. *Privacy protection goals* express ways in which sensitive information is protected. *Privacy vulnerabilities* reflect

ways in which sensitive information may be susceptible to privacy invasions. Goals not relevant to privacy or privacy-related functionality are not classified because they are outside of the scope of privacy and security.

Once goals have been classified, they are further categorized by type [5]. There are five types of privacy protection goals: notice and awareness, choice and consent, access and participation, integrity and security, and enforcement and redress. There are seven types of vulnerabilities: information collection, monitoring, personalization, storage, transfer, aggregation, and contact. A comprehensive description of each category can be found in [5].

In this study, we employ goal mining to extract privacy protection goals and vulnerabilities from healthcare privacy notices. We then use these goals to reconstruct the implicit privacy requirements met by the privacy policies, while eliminating unnecessary text that can either mask the true meaning of a policy or cause the policy to be too complex to understand. We utilize the conciseness of the goals expressed in structured natural language, as well as the classification benefits of the goal taxonomy, to create our APP alternatives to TOPPs.

B. Experimental Design

The experiment employed a 3×4 factorial design and presented each participant with a privacy policy to read. The policies originated from three different Web sites in the healthcare domain. We chose to examine healthcare privacy policies for three key reasons: 1) healthcare Web sites are required by the Health Insurance Portability and Accountability Act² (HIPAA) to provide privacy policies that are accessible to the public, “plainly written,” and disclose several types of privacy practices, which provides for detailed policies that cover most of the privacy categories in the Antón *et al.* privacy taxonomy; 2) people are most concerned about the collection of their financial and healthcare information, so they are more likely to read policies related to the disclosure of such types of information [47]; and 3) extensive research previously conducted into the readability of existing healthcare privacy policies [6], [27], [47] substantiates the need for alternative policy representations that are more consumer centric and comprehensible.

These three healthcare policies were each presented using four different representations, we refer to these four representations as *variants* as follows.

- 1) *Policy*: This variant is the original privacy policy, in paragraph-form, obtained from a given Web site. This is the most common approach to privacy policy representation, serving as the TOPP for our investigations. Fig. 1 is an excerpt from a policy in this form; this is the actual policy text displayed on an active Web site, but with the company name replaced with “BrandX.”
- 2) *Goal/vulnerability statements*: In this representation, the policy is expressed as a list of privacy goals and vulnerability statements. To create the list of privacy goals and vulnerability

¹Drugstore.com privacy policy. Downloaded from <http://www.drugstore.com> on September 15, 2003.

²Health Insurance Portability and Accountability Act of 1996, 42 U.S.C.A. 1302 d to d-8 (West Supp. 1998).

When you place an order, we will ask you to set up "your account," which includes your name, e-mail address, mailing address, credit card number and expiration date, as well as certain other information when you order prescriptions. Using your account information, we will send you communications that we believe are relevant to you, including eMedalert, prescription refill and renewal reminders, newsletters or emails. If you prefer not to receive optional email or other communication from us, you may adjust your account to prevent such communications. If we receive updated account information from our shippers or other third parties, we may revise your account for you so that we can efficiently process your orders, deliver your packages or otherwise communicate with you. If you would like to review or revise the information we have in your account, you may access such information by clicking on the "your account" tab on any screen.

Fig. 1. Original text excerpt from a TOPP.

- COLLECT PII when placing an order
- USE PII to offer products/services
- OPT-OUT from receiving emails from our company
- UPDATE PII automatically using information received from 3rd parties
- ALLOW customer to modify/remove their PII

Fig. 2. Goal list excerpt from goal/vulnerability APP.

statements that accurately correspond to the statements in the given Web site's actual privacy policy, we employed the goal-mining approach (see Section III-A). Fig. 2 presents an example of this type of representation; the goals in this example were mined from the natural, paragraph-form example in Fig. 1.

3) *Categorical*: In this representation, we express the policy document as a list of privacy goals and vulnerability statements that have been categorized. The goals were extracted from the given Web site's actual privacy policy, using goal mining, and then, organized into categories according to the Antón *et al.* privacy taxonomy [5]. In this alternative, participants were first presented with a list of the taxonomy categories. Each category had at least one goal associated with it from the original Web site privacy policy. Participants could then click on a category heading hyperlink to view a list of goals/vulnerability statements, presented in bulleted form that are relevant to the category of interest. Fig. 3 illustrates how the categories appeared to participants, while Fig. 4 illustrates what participants would have seen when they clicked on one of the category headings.

4) *Goals/vulnerabilities in policy*: In this representation, participants were presented with the given Web site's original natural language privacy policy, but the format differs from the *policy* variant [variant (1)]. Within the policy, statements that contain privacy goals or vulnerability statements rele-

Access/Participation
This category contains policies relevant to denying access to pages or services if customers do not provide their PII.

Choice/Consent
This category outlines ways users have control over how what information is collected from them and whether the information can be transferred to others.

Enforcement/Redress
This category outlines the mechanisms in place to enforce privacy, and prescribes general guidelines that companies and their employees should follow.

Fig. 3. Example categories list for the categorical APP.

Choice/Consent

Definitions:

PHI - PHI stands for Personal Health Information. This includes any information that is related to one's medical history such as prescriptions, family illnesses, past treatments, current treatments, etc.

BrandX's Choice/Consent policies:

- We will disclose PHI at request of patient
- Allow consumers to opt-out from receiving emails from our company
- Allow customers to opt-out from sharing website usage information with 3rd parties
- Allow customers to opt-out of sharing information with 3rd parties

Back to the Categories

Fig. 4. Example list of goals from a choice/consent category for the categorical APP.

vant to consumer privacy were bolded and highlighted. When a participant hovered their mouse over a statement, the statement turned to blue and a popup box appeared containing the protection goal or vulnerability statement extracted from the original text. In this way, participants were presented with both the original text and its corresponding privacy goal or vulnerability. Fig. 5 illustrates what users saw if assigned to this variant. Notice the goal/vulnerability statements bolded and italicized within the policy, as well as the blue popup window containing the associated goal/vulnerability that appeared when the user hovered their mouse over a statement.

The other design factor referred to the healthcare Web site that provided the content for the four policy representations listed earlier. These policies were obtained from the following Web sites: Novartis.com, Drugstore.com, and HealthCentral.com. Each Web site policy contained a different ratio of

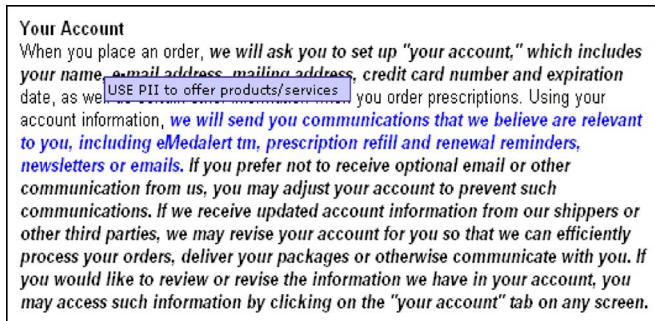


Fig. 5. Example of goals/vulnerabilities in policy APP.

vulnerability statements to privacy protection goals. The number of vulnerability statements and number of protection goals provide a straightforward way to measure how protective or how potentially dangerous a Web site might be to the consumer's sensitive information. The Novartis.com policy contained more protection goals (23) than vulnerability statements (9), and the Drugstore.com policy contained more vulnerability statements (36) than protection goals (19). The HealthCentral.com policy served as our *control* factor because it contained the same number of vulnerability statements (12) as protection goals (12).

Table I shows the 12 possible policy/representation combinations that were evaluated in this study. Each column represents a different policy format (natural language policy, goals and vulnerability statements, categorical, goals/vulnerabilities in policy), whereas each row represents different policy content (more vulnerable policy, more protective policy, control policy). The number in each cell corresponds to the number of participants who received and submitted a questionnaire for that policy/representation combination.

To prevent name-brand recognition bias from influencing the participants' trusting beliefs, and thereby the results of the experiment, we removed all references to the names of the original policy authors and replaced the company name with "BrandX." To further prevent bias, we randomly assigned participants to one of the 12 policy representations.

Participants were first asked to imagine themselves in a scenario where they were in need of an Internet supported healthcare service (e.g., purchasing from an online pharmacy, retrieving information from a physician online). The scenario continued with the participant imagining a visit to "BrandX's" Web site to obtain the needed service. Similar to many healthcare Web sites, BrandX would ask the participant to provide some personally identifiable information (PII). The experiment provided the participant with a privacy policy to read as a means of learning how the PII would be collected and used.

Once participants finished reading the policy, they were presented with a set of questions based upon the content of the given policy. There were three question categories: perception, comprehension, and demographic. The *perception* questions gauged how users felt about what they read and were presented using a five-point Likert-scale anchored by "strongly disagree" (1) and "strongly agree" (5). The *comprehension* questions measured how well users comprehended the content of the policy and

were presented in the form of multiple-choice quiz questions. Participants were not notified ahead of time that they would be quizzed on the content because we wanted them to focus on the given scenario in order to simulate realistic purchasing behaviors. However, we did warn users that they would not be able to return to the policy once they proceeded. The *demographic* questions measured the demographic makeup of the sample and were presented in the form of drop-down, multiple-choice questions. Respondents were also presented with the question: "Why didn't you read the entire privacy policies of the Web site?" Respondents could choose from one of five broad responses, including: "I read the entire set of privacy policies of the Web site."

The same perception and demographic questions were asked of all participants, regardless of which of the 12 policy representations they received. The comprehension questions differed depending on which policy they received, since each organization has their own set of privacy practices. Participants received three comprehension questions that corresponded to the content of the specific policy.

C. Pilot Testing

In March of 2005, we conducted a focus group, consisting of five privacy experts, to address the need for alternatives to TOPPs. The outcome of the focus group culminated in a pilot study that was conducted a month later with 95 participants. The participants were both graduate and undergraduate students in management and computer science. The initial experiment had four policy formats using content from three different Web sites. The *policy*, *goals/vulnerabilities*, and *goals/vulnerabilities in policy* variants were used in the pilot study, but the *categorical* variant was not used. The *goals/vulnerabilities* variant was presented as a list that separated privacy goals from vulnerabilities. Participants were also notified in the pop-ups of the *goals/vulnerabilities in policy* variant whether a policy statement was a privacy goal or vulnerability. We eliminated this separation to eliminate any terminology or classification bias.

Participants were presented with one of two possible survey formats: 1) participants were presented with the *policy* or *goals/vulnerabilities* variant, followed by a four-item survey to assess user perceptions or 2) participants were presented the *policy* or *goals/vulnerabilities* variant, followed by a mid-experiment survey to assess user perceptions, followed by the *goals/vulnerabilities in policy* variant, and then, a final four-item survey. This design aimed to gauge and compare user perceptions of the various policy formats.

The participants were also encouraged to provide feedback regarding the experiment, the procedure and the survey. Based on some preliminary statistical analysis and comments from the participants, we deleted some of the survey items and reworded others. Following an initial set of survey and experiment revisions, we held two focus groups consisting of privacy experts, privacy students, and survey experts. The focus groups had 15 and 8 participants, respectively. The purpose of the second focus group was to analyze the survey comments provided by the pilot experiment participants and comments from the first focus

TABLE I
EXPERIMENTAL DESIGN AND SAMPLE SIZE

		Variant				Total
		Original Policy	Goals/Vulnerabilities	Categorical	Goals/Vulnerabilities in Policy	
Policy	<i>Drugstore.com</i> (more vulnerable)	76	80	77	93	326
	<i>HealthCentra.com</i> (control)	86	87	81	83	337
	<i>Novartis.com</i> (more protective)	77	81	83	89	330
	Total	239	248	241	265	993

group. Based on the discussion, we revised the format of the experiment. We reduced the number of policies and the number of questionnaires for each participant. This led to our current experimental design where each participant saw one policy format and answered one set of questions. The resulting questionnaire had seven-scale items, 11 demographic items and between three and five multiple-choice questions, depending on the variant being applied.

Comprehension questions, which were not asked of pilot participants, were incorporated into the final questionnaire. Based on pilot study feedback, as well as focus group inputs, we suspected that user perception and confidence in their understanding might not accurately reflect their actual comprehension. Even though people may think that they comprehend a topic, there is no way to be certain unless you test their understanding. Thus, comprehension questions in the form of multiple-choice quiz questions were incorporated.

D. Survey Distribution and Data Collection

The final survey was linked from a National Science Foundation (NSF) sponsored Web site, and was advertised to a variety of Internet users worldwide. Participants were solicited through a variety of outlets, including links to the survey from various university Web pages, general news sites, alumni mailing lists, professional mailing lists, e-mail, and word of mouth. To increase the variability in the data and generality of the survey, we launched a marketing campaign designed to target all demographic audiences. Participants of the survey were offered an entry into a prize drawing, to take place at the conclusion of the survey. The survey was available October 25, 2005 to December 10, 2005 via the Web at an NSF-sponsored project site.

We received 1215 total responses, but used some built-in mechanisms to distinguish between valid and invalid responses. Based on the pilot study, it was relatively certain that participants could not read the directions, as well as the policy, in less than 30 s. Therefore, we eliminated participants who spent a combined total of 30 s or less reading the instructions and policy. This eliminated 212 responses. We eliminated 10 additional responses by analyzing the responses to some carefully formed validation items included in the survey. The final dataset resulted in 993 usable responses.

IV. FINDINGS

A. User Perception

The user perception questions were presented to the participants in a five-point Likert-scale form. Each answer was assigned a numeric value: strongly disagree = 1, disagree = 2, unsure = 3, agree = 4, and strongly agree = 5. Table II presents a summary of the perception responses for all participants. Analysis of variance (ANOVA) methods and Tukey's least significant difference test were used to analyze the participant responses and test our hypotheses.

When asked whether they feel secure sharing personal information with BrandX, after viewing their privacy policy, participants tend to feel more secure with the *policy* (mean = 2.72) and *goals/vulnerabilities in policy* (mean = 2.85) variants than with the *goals* (mean = 2.44) and *categorical* (mean = 2.58) variants ($F_{3972} = 7.01$; $p = 0.0001$). However, there is no statistically significant difference between the *policy* and *goals/vulnerabilities in policy* variants or between the *goals* and *categorical* variants. These results do not support hypothesis (H1), which states that participants feel more secure with APPs than with TOPPs. In contrast, they feel more secure with the policies expressed using a paragraph format.

Another question asked the participants whether they believe that BrandX would protect their personal information more than other companies. Our survey did not provide users with a reference set of other Web sites. Rather, this question requires participants to draw upon previous experiences to compare to this Web site. More importantly, we are asking this question to compare variants and not using it to gauge how secure participants actually feel as compared to other Web sites. The *goals/vulnerabilities* (mean = 2.6) and *categories* (mean = 2.6) participant responses significantly differed from those of the *goals/vulnerabilities in policy* (mean = 2.8) participants ($F_{3972} = 3.04$; $p = 0.0283$). This result does not fully support hypothesis (H2), which states that participants feel more protected by APPs than with TOPPs.

When asked whether they believe BrandX's privacy practices were explained thoroughly in the policy they read, participants tend to believe that policies expressed in natural, paragraph form (e.g., *policy* and *goals/vulnerabilities in policy*) are the most thorough ($F_{3972} = 17.95$; $p < 0.0001$). This result supports hypothesis (H4), which says that participants believe that TOPPs are more thorough than concise APPs (i.e., *goals* and *categorical*

TABLE II
DESCRIPTIVE STATISTICS (LIKERT SCORES)

	Natural Language Policy		Goals/Vulnerabilities		Categorical		Goals/Vulnerabilities in Policy	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
<i>I feel secure sharing my personal information with BrandX after viewing their privacy practices</i>	2.72	1.03	2.44	1.02	2.58	1.06	2.85	1.01
<i>I believe BrandX will protect my personal information more than other companies</i>	2.71	0.97	2.60	0.93	2.60	0.90	2.80	0.98
<i>I feel that BrandX's privacy practices are explained thoroughly in the policy I read</i>	3.24	0.91	2.77	1.07	2.88	1.06	3.33	0.96
<i>I feel confident in my understanding of what I read of BrandX's privacy policy</i>	3.27	1.03	2.95	1.05	3.22	1.02	3.31	1.01

policies). This result may be explained by the fact that policies expressed in natural, paragraph form are generally more verbose than the *categorical* or *goal/vulnerabilities* variants that present concise policy statements in list form. A longer policy may provide the illusion of being more thorough even though it may not actually be more informative.

When asked whether they feel confident in their understanding of what they read in BrandX's privacy policy, participants felt equally confident with all policies except the *goals* policy ($F_{3972} = 5.76$; $p = 0.0007$). This result does not support hypothesis (H3) that users are likely to feel more confident in their understanding of APPs than TOPPs. Indeed, participants felt as confident in their understanding of two of the three APPs as the TOPP, but not more so. Participants felt statistically less confident in the *goals* APP. We suspect that this is due to users feeling intimidated by an unfamiliar privacy policy representation that is less user-friendly than natural language phrases.

B. User Comprehension

The comprehension questions were presented in the form of a multiple-choice quiz. For each participant, a quiz score was calculated as the percentage of the number of questions answered correctly. The formula for the score was as follows: score = (questions correct/total questions asked) \times 100. This yielded a percentage value between 0 and 100 and represented the participant's true comprehension, not their perception of comprehending. We chose to represent the scores as percentages, rather than on a three-point scale, to better represent the percentage of questions that participants typically answered correctly.

The results (see Table III) clearly illustrate that the TOPP was the most difficult representation to comprehend ($F_{3969} = 11.55$; $p = 0.0001$). In fact, participants only answered one-third of the privacy related questions correctly when given the *policy*

TABLE III
AVERAGE COMPREHENSION SCORE FOR EACH POLICY REPRESENTATION (ALL PARTICIPANTS)

Variant	N	Average	SD
<i>Natural Language Policy</i>	239	35.70	28.15
<i>Goals/Vulnerabilities</i>	248	43.82	31.99
<i>Categorical</i>	241	52.14	35.11
<i>Goals/Vulnerabilities in Policy</i>	265	43.27	31.21
<i>Average Overall</i>		43.74	

variant. The standard deviation of each variant score ranges from 28.15 to 35.11. The survey presented participants with between three and five comprehension questions. Therefore, the observed standard deviation is consistent with participants being within \pm one question of the mean comprehension score.

The comprehension of the TOPP increased when privacy statements were highlighted within the policy and accompanied by goal statements, as in the *goals/vulnerabilities in policy* variant. The comprehension scores increased further when participants were presented with the policy expressed as a list of goal statements, as in the *goals/vulnerabilities* variant. The comprehension scores were the highest for the *categorical* variant.

Recall that respondents were asked: "Why didn't you read the entire privacy policies of the Web site?" One of the possible answers was: "I read the entire set of privacy policies of the Web site." By eliminating all users who did not choose the latter response, we were able to omit all participants who did not read the entire privacy policy, and thereby, further analyze the readability of the policy variants. The average comprehension score for each variant for participants who read the entire policy is presented in Table IV. First, it did not matter whether participants read the entire policy or not, the *categorical* variant was always the easiest to comprehend (based on the quiz score) and the

TABLE IV
AVERAGE COMPREHENSION SCORE FOR EACH POLICY REPRESENTATION
(ONLY THE PARTICIPANTS WHO READ THE ENTIRE POLICY)

Variant	N	Average	SD
<i>Natural Language Policy</i>	135	40.00	27.86
<i>Goals/Vulnerabilities</i>	113	55.46	31.69
<i>Categorical</i>	150	64.22	30.68
<i>Goals/Vulnerabilities in Policy</i>	135	49.38	32.54
<i>Average Overall</i>		52.88	

policy variant was always the most difficult. Second, even among participants who read the entire privacy policy, only about half of the comprehension questions were answered correctly. These results overwhelmingly support hypothesis (H5), which stated that participants would comprehend APPs better than TOPPs.

C. User Perception and User Comprehension Alignment

Although participants feel most secure and protected by the goals/vulnerabilities in the policy variant, their comprehension of the content is not as good as with other variants. This finding does not support hypothesis (H6), which states that participants would feel more positively toward policies they comprehend better. Even among participants who read the entire policy, those who read the *policy* variant, which they perceive as being the most secure and thorough variant, answer only a little more than one-third of the questions correctly. In contrast, despite not feeling as secure or protected by the categorical policies, participant comprehension scores were much greater when given these policies. This misalignment of user perception with user comprehension is disconcerting because consumers may be more inclined to trust a company with a policy that lacks clarity and readability. This leads to less-informed decisions that could result in the increase of unanticipated and unwanted uses and disclosures.

D. Demographic Considerations

Approximately 67% of the participants were male, 49% were between the ages of 18 and 35, and 75% had earned a college degree or higher. These demographics are comparable to profiles reported in other Internet user studies [14], [29]. Additionally, 32% of the participants have made online healthcare product purchases and 42% engage in online healthcare research at least once a month. A listing of demographic data can be found in Table V.

After analyzing the data, we conclude that demographics had no statistically significant ($p < 0.05$) impact on user perception between the four variants or between the three policies. Furthermore, with the exception of a single age group (age 57 and older), demographics had no statistically significant affect on the comprehension scores between the four variants or between the three policies. In this case, individuals age 57 and older scored an average of 37.75 on the comprehension questions as compared to the overall mean = 43.74 ($F_{9962} = 2.5, p = 0.008$).

TABLE V
PARTICIPANT DEMOGRAPHICS

	# of Participants	% of Participants
Gender		
<i>Female</i>	297	29.91
<i>Male</i>	663	66.77
<i>Rather Not Say</i>	26	2.62
<i>No Response</i>	7	0.70
Age		
<i>Less Than 18</i>	8	0.81
<i>18-21</i>	122	12.29
<i>22-28</i>	207	20.85
<i>29-35</i>	160	16.11
<i>36-42</i>	144	14.50
<i>43-49</i>	120	12.08
<i>50-57</i>	116	11.68
<i>57+</i>	98	9.87
<i>Rather Not Say</i>	14	1.41
<i>No Response</i>	4	0.40
Education		
<i>Some High School</i>	14	1.41
<i>High School Graduate</i>	19	1.91
<i>Some College</i>	194	19.54
<i>College Graduate</i>	210	21.15
<i>Some Graduate School</i>	110	11.08
<i>Master's Degree (or equivalent)</i>	241	24.27
<i>Doctorate</i>	148	14.9
<i>Professional Degree (MD/JD)</i>	38	3.83
<i>Other</i>	6	0.6
<i>Rather Not Say</i>	9	0.91
<i>No Response</i>	4	0.40

E. Additional Observations

As previously mentioned, we asked participants, "Why didn't you read the entire privacy policy of the Web site?" In each variant, the majority of participants read the entire policy (see Table VI). The most common reason why participants did not read the entire privacy policy, however, was that the policy was too long. One important observation is that participants most often read the entire privacy policy when given the *categorical* variant. As a result, a lesser percentage of participants felt that the *categorical* policies were too long.

Participants also felt that, compared to the other variants, the *goals* variant was not organized very well. This may be due to the succinct and blunt wording of the goal statements.

V. DISCUSSION

Our empirical results validated prior work [6], [8], [27], [47] by discovering that a TOPP was the most difficult policy representation to comprehend. In fact, users only answered one-third of the comprehension questions correctly when given the *policy* variant. The TOPP representation does not adequately convey an organization's privacy practices, leaving consumers uninformed and vulnerable to privacy violations by deceptive organizations.

TABLE VI
VARIANT PERCENTAGES FOR EACH RESPONSE TO “WHY DIDN’T YOU READ THE ENTIRE PRIVACY POLICY OF THE WEB SITE?”

	Variant				
	Policy	Goals	Categories	Goals/Policy	Average
<i>The policy was too hard to understand</i>	1.69	5.69	4.64	2.34	3.59
<i>The policy was too long</i>	35.86	34.15	26.16	42.97	34.23
<i>The policy was not organized well</i>	4.66	13.01	5.49	3.13	6.39
<i>I have no interest in the privacy policies of institutions I share my personal information with</i>	0.42	2.85	2.11	1.95	1.8
<i>I read the entire privacy policy of the website</i>	56.36	44.31	61.6	49.61	53.49

The comprehension of the TOPPs increased when privacy statements were highlighted within the policy and accompanied by goal/vulnerability statements, as in the *goals/vulnerabilities in policy* variant. This illustrates that consumers can greatly benefit from a privacy policy representation that simply supplements the policy with highlighted goal/vulnerability statements. The inherent benefit with this variant is that the policy is presented to the consumer in a familiar format, yet provides additional information to the interested consumer to help them better understand the organization’s privacy practices.

The comprehension scores increased when users were presented with the policy expressed as a list of goal/vulnerability statements, as in the *goals* and *categorical* variants. We expected this result since goal/vulnerability statements are uniform and eliminate extraneous and unnecessary information. For example, in the Drugstore.com policy, the *goals* variant contained one-fourth of the number of words contained in the *policy* variant. With less information to read and retain, consumers do not experience information overload and can retain essential information regarding the uses and disclosures of their personal information.

The comprehension scores were highest for the *categories* variant. This policy variant organizes the goal/vulnerability statements found in the *goals* variant into the privacy taxonomy category to which it belongs. This variant presents users with a list of the different categories and allows them to choose which policy category they wish to further explore. Because this variant is comprised of the same goals/vulnerabilities that are found in the *goals* variant, we expected to see similar comprehension results. Instead, we found that the comprehension of the policies actually increased when they were categorized. This may suggest that users will retain information better when it is presented to them in an organized and more abstract, succinct manner. The category names and descriptions may also serve as a mnemonic device that enables consumers to mentally classify the information in the policy for easier retention and recall.

One of the more compelling observations is that, despite which policy content users were given (i.e., Healthcare.com, Drugstore.com, or Novartis.com), users tend to feel more secure and protected by policies that are expressed in natural, paragraph form (i.e., *policy* and *goals/vulnerabilities in policy* variants) than with policies expressed as a list of goal/vulnerability statements (i.e., *goals* and *categorical* variants). We attribute this to

users feeling more comfortable with the natural, paragraph-form representation that strives to present a warm and caring impression. Most of the typical, paragraph-form policies begin with reassuring language that encourages the user to feel positively toward the Web site. For example, the following text is found at the beginning of an actual policy posted by HealthCentral.com³ during our study:

Dear Friends,

First and foremost, HealthCentral is deeply committed to preserving your privacy. We have established stringent rules of privacy and responsibility in order to protect the rights of HealthCentral users.

This perception may also be due, at least in part, to the level of familiarity associated with this type of policy representation. It is very likely that the majority, if not all, of the participants were familiar with policies expressed in natural, paragraph form; yet it is unlikely that any of the participants were familiar with goal-based policies. Participants may have been wary of a policy format with which they had no prior experience; this is reflected in the results of our study. With repeated exposure to the APPs, we expect the level of familiarity to increase the consumers’ sense of security. Furthermore, making users aware that APPs are statistically easier to comprehend than TPPs may also increase the consumers’ sense of security toward APP representations.

Despite feeling more secure and protected by the natural, paragraph-form policy representations, participant comprehension scores were much higher with goal-based policies. This misalignment of end user perception with comprehension is disconcerting because users may be more inclined to trust a company whose policy lacks clarity and readability. This may result in consumers making poorly informed decisions that could result in the increase of unanticipated or unwanted uses and disclosures. Deceitful organizations could even exploit these results to garner user confidence, and then, proceed to violate their privacy, which would further exacerbate long-term consumer privacy concerns.

A. Discussion: Limitations

Our participants have an average education level of more than a college degree—they tend to have taken at least one class in

³“HealthCentral.com Privacy Policy.” <http://www.healthcentral.com>. Downloaded on September 19, 2003.

a graduate program. Although this does not parallel the average education level of Internet users over age 25 (14.4 years or two years of college), we can still make important inferences from this study. Table III shows the exceptionally low “quiz” scores of our sample participants. It is likely that a more average sample would score even lower on this portion of the survey. Additional work is needed to better understand how organizations should communicate privacy management practices to their consumers.

As with any survey, there are always concerns about whether participants are completely honest in their responses [26]. Several measures were taken to avoid incorporating dishonest users’ responses into the participant dataset, including preventing users from revisiting the policy to look up answers to comprehension questions, requiring that the questionnaires be completed before submission, ensuring the anonymity of participants would be preserved, and removing responses from the dataset that were deemed invalid (e.g., by eliminating responses from participants who spent a combined total of 30 s or less reading the instructions and policy). Given that users can revisit a policy at any time in the real world, we plan to replicate this study and allow users to revisit the policy.

As a result of this study, we discovered that a disparity exists between user perceptions of privacy policy representations and how well users comprehend the policies. Even though users felt more secure with, protected by, and comfortable with the *goals/vulnerabilities in policy* variant, they did not comprehend them as well as the *categorical* privacy policies. Recognizing that users more readily comprehend categorical policies than natural, paragraph-form policies, researchers need to exploit this finding so that organizations can communicate privacy management practices more effectively and efficiently to consumers. This is especially important in healthcare where many consumers are searching for solutions to sensitive types of medical problems.

B. Implications and Conclusion

This paper contributes to the emerging body of research on privacy policy readability, techniques to address consumer privacy concerns, and overcoming trust obstacles involved in online transactions. In this paper, we discuss the development of a validated empirical survey instrument that enabled us to evaluate user perception and comprehension of typical and alternative privacy policy representations.

Our findings have many practical implications for managers, compliance officials, and legislative bodies. Consumers’ concerns continue to grow with the increase in highly publicized privacy breaches [44], [45], [48]. Convincing consumers to invest trust in an online organization, especially one with which a consumer has no prior relationship, has become an impediment to the revenue streams of several online commerce organizations [9], [21], [24], [33], [41]. Privacy policies are an effective way of establishing trust with consumers, but policies that are incomprehensible may deter potential consumers.

The results presented in this paper illustrate that current privacy policy representations are not sufficient for conveying an organization’s privacy practices. On average, users were only

able to answer one-third of the privacy-related comprehension questions correctly after reading a privacy policy. This number is staggeringly low and consistent with predictions made by researchers in previous work [6], [8], [27], [47]. In this study, we found that consumers comprehend all three APPs better than the TOPP. Based on previous research [34], we know that consumers are more willing to engage in transactions with organizations where they perceive the policy as being more comprehensible. The implication is that managers would benefit from exploring alternatives to current privacy policy representations. Many online organizations, including eBay⁴, Barnes & Noble⁵, and Google⁶ have already adopted policies that exhibit some form of categorization of policy content. Such policies aid consumers in organizing and prioritizing information; therefore, consumers are more apt to read the policy, which may help foster a more trusting consumer–organization relationship.

As legislation and consumer concerns continue to drive the evolution of privacy policies, managers will be forced to explore and justify more reasonable policy alternatives. Posting concise, comprehensible policies may serve to increase a consumer’s sense of interactional justice with an organization, while complying with applicable legislation. Our study will aid managers in choosing a policy representation that is appropriate in meeting their organizational needs, while addressing the consumer and legislative demand for more comprehensible privacy policies.

Our study highlights a paradox between trust and comprehension. Participants perceived the TOPP to be more trustworthy than the APPs, despite the TOPP being the least comprehensible. Though consumer perception of APPs would likely increase with increased exposure and education as to the benefits of APP representations, many organizations may be reticent to adopt APP representations under a free market model. In this case, it would be the responsibility of legislative bodies and compliance organizations, such as the FTC, to mandate the adoption of an APP format that raises consumer comprehension to an acceptable degree.

Legislative bodies at the federal and state levels may use the results of this study to substantiate the need for more stringent guidelines that govern the privacy policy content and presentation. We have clearly shown that existing privacy policies are not sufficient for conveying the privacy practices of an organization to consumers. This study may provide legislators insight into the shortcomings of existing policies, while providing reasonable alternatives that improve consumer awareness.

Compliance organizations, such as the FTC, may employ the results of this study to ensure that organizations do not exploit the complexity of their policies to the detriment of consumers. For example, our study illustrates that consumers feel the length of the policy is associated with thoroughness when, in fact, length can often be used to mask and obfuscate practices that

⁴eBay.com privacy policy. Accessed on May 6, 2006 from <http://www.ebay.com>

⁵Barnes & Noble privacy policy. Accessed on May 6, 2006 from <http://www.barnesandnoble.com>

⁶Google.com privacy policy. Accessed on May 6, 2006 from <http://www.google.com/intl/en/privac.html>

are not consumer-friendly. A deceptive organization may exploit this finding by providing a lengthy policy that concentrates the most relevant, and most potentially disturbing, privacy related information, for example, toward the bottom of their policy. Our findings may aid the FTC in detecting such organizations and establishing more appropriate policy guidelines that are less likely to allow deceptive practices.

ACKNOWLEDGMENT

The authors would like to thank IBM for providing door prizes for the survey participants.

REFERENCES

- [1] M. Ackerman, "Privacy in pervasive environment: Next generation labeling protocols," *Pers. Ubiquitous Comput.*, vol. 8, no. 6, pp. 430–439, 2004.
- [2] M. Ackerman and L. Cranor, "Privacy critics: UI components to safeguard users' privacy," in *Extended Abstr. ACM Conf. Human Factors Comput. Syst. (CHI 1999)*, vol. 2, pp. 258–259.
- [3] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, "An XPath-based preference language for P3P," in *Proc. 12th Int. World Wide Web Conf. (WWW 2003)*, pp. 629–639.
- [4] E. B. Andrade, V. Kaltcheva, and B. Weitz, "Self-disclosure on the Web: The impact of privacy policy, reward, and company reputation," *Adv. Consum. Res.*, vol. 29, pp. 350–353, 2002.
- [5] A. I. Antón and J. B. Earp, "A requirements taxonomy to reduce Web site privacy vulnerabilities," *Requir. Eng. J.*, vol. 9, no. 3, pp. 169–185, 2004.
- [6] A. I. Antón, J. B. Earp, M. W. Vail, N. Jain, C. Gheen, and J. M. Frink, "HIPAA's effect on Web site privacy policies," *IEEE Security Privacy*, vol. 5, no. 1, pp. 45–52, Jan./Feb. 2007.
- [7] A. I. Antón, Q. He, and D. Baumer, "The complexity underlying JetBlue's privacy policy violations," *IEEE Security Privacy*, vol. 2, no. 6, pp. 12–18, Jan./Feb. 2004.
- [8] A. I. Antón, J. B. Earp, D. Bolchini, Q. He, C. Jensen, and W. Stufflebeam, "The lack of clarity in financial privacy policies and the need for standardization," *IEEE Security Privacy*, vol. 2, no. 2, pp. 36–45, Mar./Apr. 2004.
- [9] F. Belanger, J. Hiller, and W. J. Smith, "Trustworthiness in electronic commerce: The role of privacy, security, and site attributes," *J. Strateg. Inf. Syst.*, vol. 11, no. 3/4, pp. 245–270, 2002.
- [10] M. Brown and R. Muchira, "Investigating the relationship between Internet privacy concerns and online purchase behavior," *J. Electron. Com. Res.*, vol. 5, no. 1, pp. 62–70, 2004.
- [11] L. Cranor, M. Arjula, and P. Guduru, "Use of a P3P user agent by early adopters," in *Proc. ACM Workshop Privacy Electron. Soc.*, 2002, pp. 1–10.
- [12] M. J. Culnan and R. Bies, "Consumer privacy: Balancing economic and justice considerations," *J. Soc. Issues*, vol. 59, no. 2, pp. 323–342, 2003.
- [13] G. Dhillon and T. Moores, "Internet privacy: Interpreting key issues," *Inf. Resour. Manage. J.*, vol. 14, no. 4, pp. 33–37, 2001.
- [14] J. B. Earp, A. I. Anton, L. Aiman-Smith, and W. H. Stufflebeam, "Examining Internet privacy policies within the context of user privacy values," *IEEE Trans. Eng. Manage.*, vol. 52, no. 2, pp. 227–237, May 2005.
- [15] Federal Trade Commission, Privacy online: Fair information practices in the electronic marketplace (a report to Congress). [Online]. Available: <http://www.ftc.gov/reports/privacy2000/privacy2000.pdf>
- [16] Federal Trade Commission, Privacy online: A report to Congress. [Online]. Available: <http://www.ftc.gov/reports/privacy3/priv-23a.pdf>
- [17] First Administrators, Inc., Penalties for noncompliance. [Online]. Available: <http://www.firstadministrators.com/hipaa/penalties.asp>
- [18] Forrester Research, "J Bennet, its my life," *Wall Street J.*, Oct. 2001.
- [19] S. Fox. (2003, Jul.), "Internet health resources," Pew Internet and American Lift Project [Online]. Available: <http://www.pewinternet.org/PPF/r/95/report.display.asp>
- [20] The Code of Fair Information Practices, U.S. Department of Health, Education and Welfare, Secretary's Advisory Committee on Automated Personal Data Systems, Records, Computers, and the Rights of Citizens, VIII. (1973) [Online]. Available: http://www.epic.org/privacy/consumer/code_fair_info.html
- [21] D. Gefen, E. Karahanna, and D. W. Straub, "Trust and TAM in online shopping: An integrated model," *MIS Quart.*, vol. 27, no. 1, pp. 51–90, 2003.
- [22] M. A. Graber, D. D'Alessandro, and J. Johnson-West, "Reading level of privacy policies on Internet health Web sites," *J. Fam. Pract.*, vol. 31, no. 7, pp. 642–645, 2002.
- [23] M. Hochhauser, "Why patients won't understand their HIPAA privacy notices [Online]. Available: <http://www.privacyrights.org/ar/HIPAA-Readability.htm>
- [24] D. L. Hoffman, T. P. Novak, and M. Peralta, "Building consumer trust online: How merchants can win back lost consumer trust in the interests of e-commerce sales," *Commun. ACM*, vol. 42, no. 4, pp. 80–85, 1999.
- [25] G. Hogben and Giles (2002), "A technical analysis of problems with P3P v1.0 and possible solutions," in *W3C Workshop Future P3P*, Position Paper [Online]. Available: <http://www.w3.org/2002/p3p-ws/pp/jrc.html>
- [26] C. Jensen and C. Potts, "Privacy practices of Internet users: Self-report versus observed behavior," *Int. J. Human Comput. Stud.*, vol. 63, pp. 203–227, Jul. 2005.
- [27] C. Jensen and C. Potts, "Privacy policies as decision-making tools: An evaluation of online privacy notices," in *Proc. ACM Conf. Human Factors Comput. Syst.*, Apr. 2004, pp. 471–478.
- [28] G. Karjoth, M. Schunter, E. V. Herreweghen, and M. Waidner, "Amending P3P for clearer privacy promises," in *Proc. 14th Int. Workshop Database Expert Syst. Appl.*, 2003, pp. 445–449.
- [29] C. Kehoe, J. Pitkow, and K. Morton, Eighth WWW user survey 1997. [Online]. Available: <http://www.cc.gatech.edu/gvu/user-surveys/survey-1997-10/>
- [30] O. M. Khalil and T. D. Harcar, "Relationship marketing and data quality management," *SAM Adv. Manage. J.*, vol. 64, no. 2, pp. 26–33, 1999.
- [31] M. Langheinrich, "A P3P preference exchange language 1.0 (APPEL1.0)," *W3C Working Draft*, Apr. 2002.
- [32] V. Martinez-Tur, J. M. Peiro, J. Ramos, and C. Moliner, "Justice perceptions as predictors of customer satisfaction: The impact of distributive, procedural, and interactional justice," *J. Appl. Psychol.*, vol. 36, no. 1, pp. 100–119, 2006.
- [33] D. H. McKnight and N. L. Chervany, "What trust means in e-commerce customer relationships: An interdisciplinary conceptual typology," *Int. J. Electron. Com.*, vol. 6, no. 2, pp. 35–60, 2002.
- [34] G. R. Milne, A. J. Rohm, and S. Bahl, "Consumers' protection of online privacy and identity," *J. Consum. Aff.*, vol. 38, no. 2, pp. 217–232, 2004.
- [35] A. D. Miyazaki and S. Krishnamurthy, "Internet seals of approval: Effects of online privacy policies and consumer perceptions," *J. Public Policy Mark.*, vol. 19, no. 1, pp. 54–61, 2000.
- [36] S. Nicovich and T. B. Cornwell, "An Internet culture?: Implications for marketing," *J. Interact. Mark.*, vol. 12, no. 4, pp. 22–33, 1998.
- [37] National Telecommunications and Information Administration, A nation online: How Americans are expanding their use of the Internet [Online]. Available: <http://www.ntia.doc.gov/ntiahome/dn/>
- [38] A. Nöteberg, E. Christiaanse, and P. Wallage, "Consumer trust in electronic channels: The impact of electronic commerce assurance on consumers' purchasing likelihood and risk perceptions," *e-Serv. J.*, vol. 2, no. 2, pp. 46–67, 2003.
- [39] R. Park, Online health information [Online]. Available: <http://www.imediadconnection.com/content/5375.asp>
- [40] P. Pavlou and M. Fygenson, "Understanding and predicting electronic commerce adoption: An extension of the theory of planned behavior," *MIS Quart.*, vol. 30, no. 1, pp. 115–143, 2006.
- [41] P. A. Pavlou, "Consumer acceptance of electronic commerce: Integrating trust and risk with the technology acceptance model," *Int. J. Electron. Com.*, vol. 7, no. 3, pp. 69–103, 2003.
- [42] Princeton Survey Research Associates International, Leap of faith: Using the Internet despite the dangers [Online]. Available: <http://www.consumerwebwatch.org/pdfs/princeton.pdf>
- [43] F. F. Reichheld and P. Schefter, "E-loyalty: Your secret weapon on the Web," *Harv. Bus. Rev.*, vol. 78, no. A, pp. 105–113, 2000.
- [44] J. A. Roznowski, "A content analysis of mass media. Stories surrounding the consumer privacy issue 1990–2001," *J. Interact. Mark.*, vol. 17, no. 2, pp. 52–69, 2003.
- [45] L. M. Sama and V. Shoaf, "Ethics on the Web: Applying moral decision-making to the new media," *J. Bus. Ethics*, vol. 36, pp. 93–103, 2002.
- [46] A. E. Schlosser, T. B. White, and S. M. Lloyd, "Converting Web site visitors into buyers: How Web site investment increases consumer trusting beliefs and online purchase intentions," *J. Mark.*, vol. 70, pp. 133–148, 2006.
- [47] K. Sheehan, "In poor health: An assessment of privacy policies at direct-to-consumer Web sites," *Amer. Mark. Assoc.*, vol. 24, no. 2, pp. 273–283, 2005.

- [48] E. C. Turner and S. Dasgupta, "Privacy on the Web: An examination of user concerns, technology, and implications for business organizations and individuals," *Inf. Syst. Manage.*, vol. 20, no. 1, pp. 8–18, 2003.
- [49] W3C, Platform for privacy preferences (P3P) project [Online]. Available: <http://www.w3.org/P3P>
- [50] T. Yu, N. Li, and A. I. Antón, "A formal semantics for P3P," presented at the ACM Workshop Secure Web Serv. (SWS), Oct. 2004.



Matthew W. Vail received the B.S. and M.S. degrees in computer science from North Carolina State University, Raleigh, in 2004 and 2006, respectively. He is currently pursuing the Ph.D. degree in computer science at the University of Massachusetts, Amherst.

His current research interests include computer security and privacy, networking, and software engineering.

Mr. Vail is a member of the Association for Computing Machinery (ACM), ThePrivacyPlace.org, and the Laboratory for Advanced Software Engineering Research (LASER).



Julia B. Earp (M'99) received the B.S. degree in mathematics and the second B.S. degree in statistics in 1991 both from North Carolina State University, Raleigh, the M.S. degree in statistics in 1992, and the Ph.D. degree in information technology in 1997 both from Virginia Polytechnic Institute and State University, Blacksburg.

She is currently an Associate Professor in the College of Management, North Carolina State University. She is an Active Senior Research Collaborator with theprivacyplace.org. Her current research interests include Internet security and privacy issues from several different perspectives,

including data management, consumer values, systems development, and policy. She is the author or coauthor of several papers on information privacy and security in journals such as the *Communications of the ACM*, *IEEE SECURITY AND PRIVACY*, *IEEE TRANSACTIONS ON ENGINEERING MANAGEMENT*, *Requirements Engineering Journal*, *Computers and Security*, and *International Journal of Organizational Computing and Electronic Commerce*.

Dr. Earp is a member of the Association for Computing Machinery (ACM) and Association for Information Systems (AIS).



Annie I. Antón (SM'03) received the Ph.D. degree in computer science from Georgia Institute of Technology, Atlanta, in 1997.

Since 1998, she has been with the Faculty of North Carolina State University, Raleigh, where she is currently an Associate Professor. Her current research interests include specification of complete, correct behavior of software systems that pose significant risks as a consequence of failures and misuse. She is the Co-Founder and the Co-Director of the NCSU E-Commerce Studio as well as the Founder and the

Director of ThePrivacyPlace.org.

Dr. Antón is an Associate Editor of *IEEE TRANSACTIONS ON SOFTWARE ENGINEERING* and the Cognitive Issues Area Editor of the *Requirements Engineering Journal*. She is a U.S. National Science Foundation CAREER Award Winner, a Computing Research Association (CRA) Digital Government Fellow, CSO Magazine's 2005 Woman of Influence in the Public Sector, and was a member of the Institute of Defense Analysis (IDA)/Defense Advanced Research Projects Agency (DARPA) Defense Science Study Group during 2004–2005. She is a member of the Association for Computing Machinery (ACM).