Available online at www.sciencedirect.com

SCIENCE DIRECT®

Int. J. Human-Computer Studies 63 (2005) 203–227

International Journal of Human-Computer Studies

www.elsevier.com/locate/ijhcs

Privacy practices of Internet users: Self-reports versus observed behavior

Carlos Jensen[a,*], Colin Potts[a], Christian Jensen[b]

[a] *Graphics, Visualization and Usability Center, Georgia Institute of Technology, Atlanta, GA 30332, USA*
[b] *Department of Economics, Southern Methodist University, Dallas, TX 75275, USA*

Abstract

Several recent surveys conclude that people are concerned about privacy and consider it to be an important factor in their online decision making. This paper reports on a study in which (1) user concerns were analysed more deeply and (2) what users said was contrasted with what they did in an experimental e-commerce scenario. Eleven independent variables were shown to affect the online behavior of at least some groups of users. Most significant were trust marks present on web pages and the existence of a privacy policy, though users seldom consulted the policy when one existed. We also find that many users have inaccurate perceptions of their own knowledge about privacy technology and vulnerabilities, and that important user groups, like those similar to the Westin "privacy fundamentalists", do not appear to form a cohesive group for privacy-related decision making.

In this study we adopt an experimental economic research paradigm, a method for examining user behavior which challenges the current emphasis on survey data. We discuss these issues and the implications of our results on user interpretation of trust marks and interaction design. Although broad policy implications are beyond the scope of this paper, we conclude by questioning the application of the ethical/legal doctrine of informed consent to online transactions in the light of the evidence that users frequently do not consult privacy policies.

Keywords: Privacy; Design; Survey; Economic models; E-commerce; Decision-making; Policy

[*] Corresponding author. Tel.: +1 404 385 1102; fax: +1 617 373 5121.

E-mail addresses: carlosj@cc.gatech.edu (C. Jensen), potts@cc.gatech.edu (C. Potts), christia@mail.smu.edu (C. Jensen).

1071-5819/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ijhcs.2005.04.019

1. Introduction

Several recent surveys conclude that people are concerned about privacy and consider it to be an important factor in their online decision making (Cranor et al., 1999; Culnan, 1999; Earp and Meyer, 2000; Culnan and Milne, 2001; Jupiter, 2002). According to one study, privacy concerns are the most frequently cited reason for not engaging in e-commerce (Jupiter, 2002). Indeed, the increasing prevalence of data collection, sharing and storage mean that this may be an increasingly prudent position for consumers to adopt (FTC, 2000; Adkinson et al., 2002).

Most studies of user concerns about privacy have been done using a survey methodology. These studies report surprisingly high rates among respondents of such behaviors as reading a privacy policy when visiting a site or taking concrete steps to protect their privacy. Informal analysis of log-file data, however, suggests that the true rates are much lower (Jensen and Potts, 2004).

This paper presents the results of an empirical study comparing users' self-reported with their observed behavior in a simulated e-commerce scenario. In particular, we examined which visible indicators of privacy invasions or privacy guarantees were effective in swaying consumers' purchase decisions. We also examined what effects gender, level of experience, and other demographic variables have on reported and observed behavior. Finally we investigated the salience of categorization schemes for users privacy concerns based on survey responses. One such scheme used in Internet-based market research is the Westin privacy segmentation (Harris et al., 1998), in which people are classified into one of three groups; "privacy Fundamentalists", "privacy pragmatists", and "privacy unconcerned." Such schemes imply that users can be classified systematically and that a user's category helps predict the user's online behavior. Only by comparing self-reports with online behavior can such assumptions be verified.

2. Method

This study was conducted online, with subjects recruited through email announcements to mailing lists, and advertisements on academic websites. Over 175 volunteer subjects, predominantly from the United States, participated in the study. Subjects came from diverse backgrounds, though approximately two thirds were currently involved in education (students, faculty and researchers). Subjects were asked a series of multiple choice demographic questions. Subjects were anonymous; they did not need to, and were not given an opportunity to give any personally identifying information. Subjects did not receive compensation for their participation, there was no deception in this study; subjects knew the purpose of the study when they decided to participate.

The study was divided into four separate but interrelated sections: (1) A basic demographic survey. (2) A survey of privacy values and attitudes. (3) A set of questions challenging users' knowledge of specific technologies and how they affect privacy. (4) An experiment presenting subjects with a series of pair-wise comparison tasks to determine the effect privacy indicators have on actual behavior. Subjects typically completed all four sections in one sitting, though they had the option to interrupt the study and return to it later. In all, subjects spent between 45 and 60 min on this study.

2.1. Demographic survey

In addition to collecting information on age, gender, and geographic location, we asked subjects about their educational and computer experience. Subjects, on average, were more highly educated (16.2 years of education) than the general Internet population (14.4 years of education) (NTIA, 2002; Jensen and Potts, 2004). Because of the self-selected nature of survey participation, we expect this population to be somewhat more concerned and knowledgeable of online privacy issues than the norm.

While 8.6% of the survey participants claimed never to buy things online, the majority purchasing things online at least once per month. On average, users reported their maximum online purchase to have been around $1000.00. These statistics indicate a survey population comfortable and familiar with e-commerce.

Our sample contained a larger group of men (74%) than women (26%), and subjects' ages averaged 30. Computer experience was high; the average respondent reported 7 years of online experience. Over 90% of participants reported having access to the Internet both from home and work, spending 25 h online per week.

The only statistically significant gender difference in the demographics was that women reported lower levels of computer expertise. The population average was 4.2 on a 5-point Likert scale, women averaged 3.6, while men averaged 4.4 ($p = .01$). Women consistently reported higher levels of concern with privacy, and online privacy in particular, though none of these differences proved statistically significant.

In terms of exposure to fraud and identity theft, our sample (see Table 1) matched data reported for the general population by the Federal Trade Commission (Synovate, 2003). Consistent with the findings of a recent study (Javelin, 2005), the majority of reported cases of identity theft and credit-card fraud originate offline rather than online.

Table 1
Victimhood and self-protection

	All (%)	Women (%)	Men (%)
Victims of identity theft	6.8	10.5	6.3
Victims of online identity theft	2.3	10.5	0.0
Victims of credit card fraud	14.3	26.3	11.9
Victims of online credit card fraud	4.6	5.6	4.6
Have installed software to protect online privacy	37.9	43.8	39.4
Have taken other steps to protect online privacy	42.7	58.8	40.3

Percentage of survey participants who claimed to have been the victims of identity theft or fraud, or taken steps to protect themselves.

One interesting finding from this section was the surprisingly large number of subjects claiming to have taken steps such as installing some form of privacy or security software. In the results section of this paper we explore the relationship between the installation of such software and the Westin privacy segmentation of users.

2.2. Privacy values

The second part of this survey consisted of a number of 5-point Likert-scale questions relating to attitudes and expectations on privacy, both online and offline. This section also focused on subjects' use of privacy policies, asking them to rate their likelihood of reading a sites' privacy policy based on the type of site, the activities they were engaged in, and their familiarity with the site. The exact questions and subjects' responses are reported in the results section.

This section of the survey was used to map our subjects to the Westin privacy segmentation (Harris, 2003). This index divides the population into three groups based on their level of concern with regards to privacy. This segmentation has been widely adopted, and is widely used to direct marketing and research efforts. Subjects are categorized based on their answers to three questions:

- Consumers have lost all control over how personal information is collected and used by companies.
- Most businesses handle the personal information they collect about consumers in a proper and confidential way.
- Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today.

Subjects giving privacy-oriented answers to all three questions are classified as "Privacy fundamentalists," those giving no privacy-oriented answers are classified as "Privacy unconcerned", while those in-between are classified as "Privacy pragmatists".

We chose not to use the three Westin classification questions because we wished to ask questions more directly related to online privacy. We placed our subjects into one of the three Westin categories based on the general pattern of their responses. We picked questions which corresponded closely to those used in the Westin surveys. Because we did not ask the same questions as Westin, the mapping is imprecise, but corresponds sufficiently for the purposes of our analysis. To avoid confusion we refer to our mapping as the Westin equivalence.

We chose to use five questions to map to the Westin groups rather than three, resulting in a more robust definition of the three categories. These questions were chosen to match what we considered to be essential properties of the three categories of users. These questions were as follows:

- I am concerned about online identity theft.
- I am concerned about my privacy online.

- I am concerned about my privacy in everyday life.
- I am likely to read the privacy policy of an ecommerce site before buying anything.
- Privacy policies accurately reflect what companies do.

We classified a participant as a "Fundamentalist" if he or she gave a privacy-oriented response to four of these five questions (and no negative answers). A participant was classified as "Unconcerned" if he or she gave no privacy-oriented responses (and at most one neutral response) to these five questions. The remaining participants were classified as Pragmatists.

Ninety-three participants completed the attitude survey and were classified as shown in Table 2. The rightmost column gives the proportions of respondents in a recent poll who were classified according to the corresponding Westin categories (Harris, 2003). The results from our classification are in line with the results of the Harris-Westin privacy polls conducted in recent years. The only slight difference is that our classification led to a more evenly divided population, with greater percentages falling in the Fundamentalist and Unconcerned categories and fewer in the Pragmatist category. The values we observed for each category was within the range of what has been reported in surveys in the past three years. It is not clear how much of this effect can be attributed to our defining questions as opposed to the way we selected participants for this study.

2.3. Knowledge challenge

One of the consistent problems with privacy surveys is the tendency subjects have of over-reporting their understanding of privacy-related issues and their willingness to act in order to protect their privacy. In order to test users and determine how big this perception gap is, we included a set of knowledge challenges in our survey. These challenges were focused on three commonly used and discussed technologies which may impact user privacy: Cookies, Web-bugs and P3P privacy policies. We chose these technologies in particular because they are parts of the vocabulary users are frequently assumed to be familiar with when setting privacy preferences (for instance in Microsoft's Internet Explorer 6.0).

Table 2
Population privacy classification

	Harris-Westin Polls				Survey—2004 (Count)
	1999 (%)	2000 (%)	2001 (%)	2003 (%)	
Fundamentalist	25	25	34	26	34% (32)
Pragmatist	54	63	58	64	43% (40)
Unconcerned	22	12	8	10	23% (21)

Percentage of the population as classified by the Westin Privacy Segmentation, and our Westin equivalence test.

When participants claimed to know what these technologies were, they were asked to rate their level of concern as well as select a reason why this technology may impact their privacy. Subjects were given a list of five possible reasons, two of which were correct, and three of which were incorrect. Users could select any number of these reasons, and we counted any answer which contained at least one correct option as a correct answer. We used these responses to gauge what percentage of subjects was truly familiar with a technology.

2.4. E-commerce experiment

To further test reported behavior against actual behavior, we included an e-commerce experiment to complement the survey sections. Each participant was presented with eight pairs of simulated e-commerce web-pages, one pair at a time, and asked to select which site they would prefer to buy from. Fig. 1 shows an example testing the difference between the use of the TRUSTe symbol and credit card icons. Subjects knew these were not real e-commerce sites, and that no money was being exchanged. The contents and design of the pages were in all cases similar

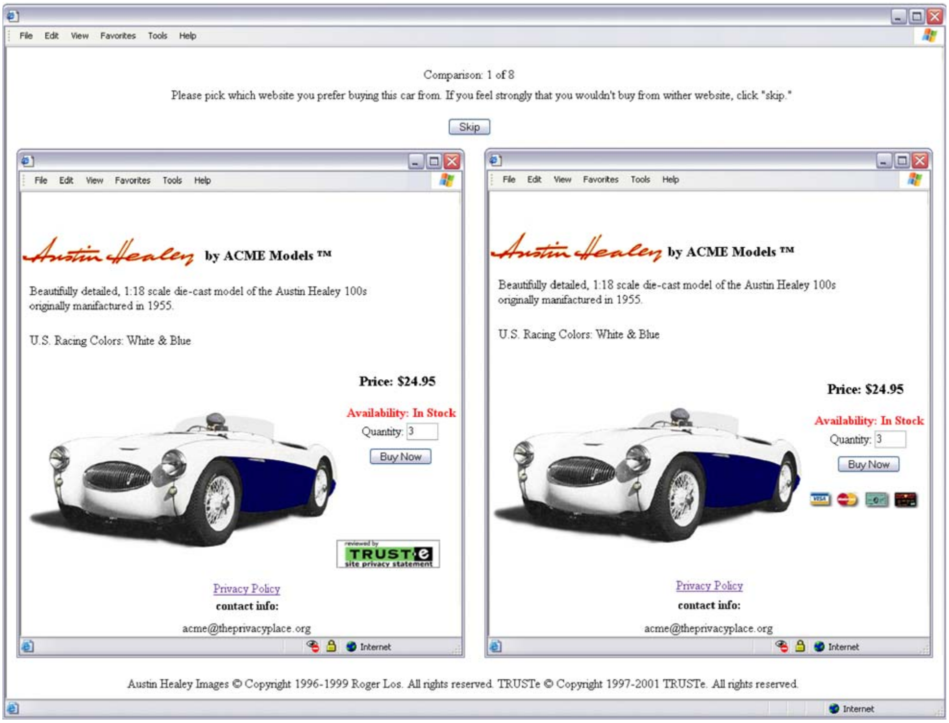

Comparison: 1 of 8

Please pick which website you prefer buying this car from. If you feel strongly that you wouldn't buy from wither website, click "skip."

Austin Healey by ACME Models ™

Beautifully detailed, 1:18 scale die-cast model of the Austin Healey 100s originally manifactured in 1955.

U.S. Racing Colors: White & Blue

Price: $24.95

Availability: In Stock

Quantity: 3

Buy Now

reviewed by TRUSTe site privacy statement

Privacy Policy

contact info:

acme@theprivacyplace.org

Austin Healey by ACME Models ™

Beautifully detailed, 1:18 scale die-cast model of the Austin Healey 100s originally manifactured in 1955.

U.S. Racing Colors: White & Blue

Price: $24.95

Availability: In Stock

Quantity: 3

Buy Now

Privacy Policy

contact info:

acme@theprivacyplace.org

Austin Healey Images © Copyright 1996-1999 Roger Los. All rights reserved. TRUSTe © Copyright 1997-2001 TRUSTe. All rights reserved.

Fig. 1. Screenshot from e-commerce experiment: In each scenario subjects were asked to identify which site they preferred buying from. The websites were identical except for two factors. In this case, one site uses a TRUSTe logo and the second uses credit card icons.

and involved a controlled variation of twelve factors commonly cited as affecting e-commerce decision-making.

The independent variables were: (1) Price of item; (2) visible indication of Secure Socket-Layer (SSL) encryption; (3) use of third-party cookies and P3P; (4) providing an e-mail address, (5) a telephone number, or (6) a postal address for the company; (7) the presence of a privacy seal (TRUSTe), (8) the presence of credit-card symbols (Visa, Mastercard, America Express, and Discover), and (9–12) four distinct types of privacy policies.

Given these 12 factors, there were 66 possible experimental conditions. Subjects were asked to indicate their preference in eight randomly selected scenarios. This resulted in a total of 1521 responses, or an average of 23 observations for each cell.

In addition to tracking selections, this part of the experiment also tracked subjects' use of policies in their decision-making. When a policy was present in one or both pages, we tracked whether the subject opened the policy page. We cannot, however, ascertain how much of the policy was read or how carefully.

2.4.1. Description of e-commerce factors

For the manipulation of the price we chose to use a fixed amount in order to reduce variability. Subjects in this condition were offered a 20% discount ($5.00) on their purchase.

To indicate the use of third-party cookies and P3P we used the Internet Explorer icon for blocked third-party cookies, which is placed in the lower right-hand corner of the browser (see Fig. 2a). While this is a feature unique to Internet Explorer, more than 90% of survey participants used this browser, and therefore were assumed to be familiar with the icon. This is of course a negative factor, the presence of this icon means the website attempted to do something which many users would be opposed to, so the dependent variable was defined as the absence of this icon.

Secure communication in the form of SSL was simulated in a similar fashion, the use of the familiar and ubiquitous SSL "lock" (see Fig. 2b). This icon also resides in the lower right-hand corner of the browser, but is common to all browsers, and has been in use for a number of years.

Finally, the four policies used were defined along two axes, whether they addressed issues considered important by users or by companies, and whether they were policies which would have a positive or negative impact on the users' privacy. In this way we derive the following four policies:

- User centered—Good.
- User centered—Bad.

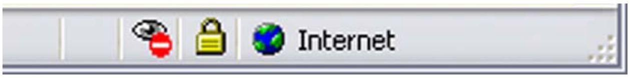

Fig. 2. (a& b) Cookie blocked (P3P) and SSL encryption icons used in the Microsoft Internet Explorer browser, respectively. The overwhelming majority of study participants used Internet explorer and should be familiar with these icons.

- Company Centered—Good.
- Company Centered—Bad.

The definition of user and company centered policies comes from a survey of user concerns and privacy policy inventory (Earp and Meyer, 2000). In that study the authors found that users were most concerned with notices disclosing information transfer/sharing practices, notification practices, and information storage practices, in order of importance. Corporations, as evident in examinations of privacy policies, are most concerned with disclosing information on security practices (assurances), collection mechanisms, and consent/assent policies, in order of importance. To simplify matters, our policies only addressed the relevant three issues (notification, data collection, and data sharing, or data collection, security, and data use). Furthermore, the "good" policies gave assurances to users on all three issues, and the "bad" policies admitted adverse practices on all three issues.

3. Results

The results of the demographic survey are reported above in the description of the study participants. In the following subsections, we present the results of the attitude survey, the self-assessment of knowledge of, and attitudes toward, technology and the simulated e-commerce scenario.

3.1. Attitudes toward privacy

Attitudes toward privacy were assessed by means of the 5-point Likert-scale attitude survey. "Agree/Strongly agree" and "Disagree/Strongly disagree" responses were pooled in Table 3, so there are three major columns, not five. These are broken down into three figures: the total responses in that category, the female responses, and the male responses. Since some participants did not indicate their gender on the demographic survey, the total may not match the average of the two sub-groups.

Both an independence test ($p < .10$) and an analysis of a logistic regression model ($p < .10$) agree that females are trending towards under-representation in the unconcerned group, but that this does not reach statistically significant levels. This means that women tend to report higher levels of concern than the men in the five key questions used to define the different user categories.

Independence tests suggest that males and females rate their concerns about privacy (general), threats posed by cookies, and their predisposition to rechecking policies differently ($p < .05$). The logistic regression models suggest that females tend to score these higher, but this trend is not statistically significant.

We also note that the level of concern for online identity theft and credit-card fraud is marginally higher than concerns for offline identity theft and credit-card fraud. This goes against expectations, as more subjects reported having experienced problems offline than online. This was consistent with recent findings indicating that most identity theft occurs offline (Synovate, 2003).

Table 3
Participant privacy attitudes and concerns

	Agree			Neutral			Disagree		
	All (%)	F (%)	M (%)	All (%)	**F (%)**	M (%)	All (%)	F (%)	M (%)
I am concerned about online identity theft*	61.3	79.0	55.9	20.4	10.5	22.1	18.3	10.5	22.1
I am concerned about online credit card fraud	66.7	84.2	60.3	16.1	5.3	19.1	17.2	10.5	20.6
I am concerned about my privacy online*	72.0	89.5	69.1	15.1	0.0	17.7	12.9	10.5	13.2
I am concerned about my privacy in everyday life*	59.1	73.7	52.9	23.7	15.8	26.5	17.2	10.5	20.6
I am likely to read the privacy policy of a site I visit for the first time	23.7	47.4	17.7	15.1	21.1	14.7	61.3	31.6	67.7
I am likely to read the privacy policy of a site which does not ask me for information	7.5	15.8	2.9	6.5	5.3	7.4	86.0	79.0	89.7
I am likely to read the privacy policy of an ecommerce site before buying anything*	43.0	79.0	35.3	25.8	21.1	25.0	31.2	0.0	39.7
I am likely to re-check the privacy policies of sites I frequently visit	7.5	10.5	4.4	9.7	10.5	8.8	82.8	79.0	86.8
What privacy policies say frequently influences my decision whether to visit or use a websites	19.4	26.3	16.2	37.6	31.6	36.8	43.0	42.1	47.1
Privacy policies accurately reflect what companies do*	16.1	15.8	14.7	50.5	52.6	50.0	33.3	31.6	35.3
Privacy policies are easy to find	36.6	21.1	38.2	35.5	52.6	32.4	28.0	26.3	29.4
It is important to me that websites publish privacy policies	68.8	63.2	69.1	19.4	31.6	17.7	11.8	5.3	13.2

Response rates to privacy attitudes survey items. Questions used to map participants to the three Westin categories are marked with a "*".

Looking at the three Westin-classes of users, we do find a high internal consistency in the answers across this section. Pragmatists rate their concern about online credit-card fraud, online identity theft, and privacy in everyday life significantly lower than Fundamentalists ($p<.05$ for all). These users also differed in their ratings of how much policies influence their decisions, how trustworthy policies are, whether policies are read on the first visit to a site, whether to check the policy when buying something online, and the likelihood that policies will be re-checked (again significantly lower than Fundamentalists ($p<.05$ for all)). Unconcerned users rate the same questions significantly lower than Pragmatists ($p<.03$ for all).

3.2. Knowledge of, and attitudes toward, privacy-relevant technology

A number of differences emerged when participants were asked if they knew about certain privacy-relevant technologies and then asked a follow-up question to probe their knowledge. The results are summarized in Table 4 as ‘‘claim’’ and ‘‘demonstrate’’ knowledge for the three technologies in question: P3P, cookies, and web-bugs. False report shows the percentage of subjects who claimed knowledge but failed to demonstrate it. The demonstrate row shows what proportion of the total population actually proved knowledgeable about these technologies. Thus, of the 21.5% who claimed to know P3P; only 25.0% could answer the probe question correctly, or 5.4% of all participants. It is important to remember that our subjects were more highly educated about computers and privacy than the average user, and that we set our threshold for knowledge pretty low. These numbers are therefore likely upper-bounds.

According to this survey, claiming knowledge about a technology does not mean much. Across the board, less than a quarter of participants who claimed to know a technology could answer simple questions about it. For P3P and Web-bugs, we find that on the whole, only 5–6% of subjects actually understand these technologies.

Only in the case of cookies do we see the majority of subjects (over 90%) claiming knowledge. Though significantly more people know about cookies than the other two technologies, the disparity between claimed knowledge and proven knowledge is

Table 4
Key technology familiarity

	P3P			Cookies			Web-bugs		
	All (%)	F (%)	M (%)	All (%)	F (%)	M (%)	All (%)	F (%)	M (%)
Claim knowledge	21.5	21.7	23.4	90.3	95.7	89.1	34.8	34.8	36.5
False report (of those who claim)	75.0	80.0	73.3	84.5	90.9	80.7	82.8	75.0	84.0
Demonstrate knowledge (overall)	5.4	4.3	6.3	14.0	8.7	17.2	5.4	8.7	5.9

Percentage of survey population to claim to understand technologies, miss-judge their understanding, and percentage of knowledgeable participants over-all.

actually larger than in the other cases, especially for women, who perform poorly on these questions. In this light, it may be argued that the lower computer experience scores reported by women might reflect accurate self-assessments, and not differences in confidence. The interesting exception to this is that more women seem to know what web-bugs are then men do.

Of the technologies examined here, cookies have, by far, received the most publicity, something evident by the very high recognition rate. These results show that, the vast majority of users do not have any real knowledge why or how cookies pose a risk to them. Despite this lack of knowledge, participants registered moderate to high levels of support for the adoption of these technologies, or in the case of cookies, the ability to control their use.

In terms of risk or benefit evaluations (Table 5), P3P and cookies were viewed as moderate risks, whereas web-bugs were viewed almost unanimously as a high-risk to personal privacy. Throughout this section we find that there are no statistically significant gender differences (the number of women who demonstrated knowledge of cookies is so small that statistical tests are inconclusive).

The only difference found in terms of the Westin equivalence was that the Unconcerned rated their concern about web-bugs significantly lower than Fundamentalists, ($p < .005$). This means that while privacy Fundamentalists are no more knowledgeable than Pragmatists or the Unconcerned, they do worry more about the risks. It is therefore possible that this segmentation is not so much based on the subjects' knowledge of risks, but on other risk estimates and sensitivities.

Table 5
Technology and risk perception

		Yes/agree			Don't know/ neutral			No/disagree		
		All (%)	F (%)	M (%)	All (%)	F (%)	M (%)	All (%)	F (%)	M (%)
P3P	It is important to me that sites adopt p3p policies	66.7	75.0	73.3	23.8	25.0	13.3	9.5	0.0	13.3
	p3p can help protect my privacy	47.6	50.0	46.7	47.6	50.0	46.7	4.8	0.0	6.7
Cookies	It is important to me to know about and control the use of cookies	72.6	73.3	73.4	17.9	26.7	14.1	9.5	0.0	12.5
	Cookies are a threat to my privacy	45.2	60.0	42.2	35.7	20.0	42.2	19.1	20.0	15.6
Web-bugs	It is important for me to know about and control the use of web-bugs	71.9	100.0	68.0	18.8	0.0	20.0	9.4	0.0	12.0
	Web-bugs present a threat to my privacy	71.9	75.0	72.0	18.8	25.0	16.0	9.4	0.0	12.0

Rate of survey participants expressing interest in and concern about key privacy technologies.

The Westin index is often used in marketing and deployment decisions with regards to privacy tools. It is assumed that the Fundamentalists are the drivers of this market, while the other two categories of users are largely uninterested. We found that the only statistically significant relationship between those claiming to have downloaded and installed countermeasures and the Westin index was that Unconcerned users reported to have done so to a less extend ($p<.05$). There were no significant differences between Fundamentalists and Pragmatists.

3.3. A further look at demographics and experience

There were some significant relationships between survey items. There was a significant correlation between self-reported frequency of online purchases and the maximum purchase amount (those reporting more frequent purchases also reported spending more money ($p<.0001$). This makes sense, as both numbers say something about their comfort and confidence with e-commerce.

We also found a correlation between the maximum online purchase amount and the reported frequency of victimization in credit card fraud or identity theft, online or offline. Victims tend to have spent more money online than non-victims (hundreds of dollars versus tens) ($p<.05$). This also makes sense, those with the highest online purchase amounts were also those with the most frequent transactions, and therefore the highest level of exposure. In both cases, differences emerge when comparing those who are moderately frequent buyers with those who rarely buy, and between those who have made moderately large purchases with those who have only made small purchases. Very frequent purchases and the purchase of very big-ticket items do not appear to be associated with victimization, possibly because we have few subjects in these categories.

Having stratified the participants according to our approximations to the Westin categories, it was possible to investigate whether their attitudes toward privacy as indicated by Westin category was associated with online experience or expertise. This was not the case, as shown by Table 6. Note that, as in most of the tables in this paper, not all subjects were classified as Fundamentalists, Pragmatists or

Table 6
Westin equivalence and e-commerce

	Frequency of purchase					Maximum purchase amount			
	Never (%)	Less than month (%)	Monthly (%)	More than monthly (%)	Weekly (%)	N/A (%)	Tens (%)	Hundreds (%)	Thousands (%)
Fundamentalist	6.3	43.8	21.9	18.8	9.4	6.3	18.8	40.6	34.4
Pragmatist	0.0	42.9	33.3	14.3	9.5	0.0	28.6	38.1	33.3
Unconcerned	2.7	54.1	21.6	13.5	8.1	0.0	13.5	59.5	27.0
All	8.6	43.4	28.0	13.1	6.3	6.9	18.3	48.0	26.3

Mapping of e-commerce activity to Westin equivalence.

Unconcerned, therefore the statistics for the total sample in the study may not match the sums of the subgroups.

There is a rather puzzling lack of association between Westin category and subjects reports of victimization (fraud or identity theft). One would expect victims to be more concerned about privacy and therefore be classified as Fundamentalists or Pragmatists. Again, this lack of correlation may be caused by the low number of victims in this survey, or it may be the case that victims accept this as the risk of doing business online. What we do find is that women tend to be underrepresented in the Unconcerned group, though this not statistically significant ($p<.10$).

3.4. Inspection of privacy policies

In the simulated e-commerce scenarios, 97% of the trials made at least one privacy policy accessible. Thus in almost all trials, a participant could check a privacy policy if he or she wanted to. In fact, Table 7 reveals that policies were only consulted in 25.9% of cases where a policy was available. This number is similar to the rates at which subjects self-reported they would (23.7%) in the general case, but much lower than the 43.0% reported for e-commerce scenarios. Over half of the participants (58.2%) consulted at least one policy. The mean number of policy look-ups for these participants was 3.18 in eight trials.

The likelihood that a participant would consult at least one policy was unrelated to the participant's Westin category and gender. In other words, women and Fundamentalists are no more likely to read policies than men or the privacy Unconcerned, respectively.

We believe that the policy consultation numbers are inflated because subjects knew they were being observed, and what the purpose of the experiment was. They therefore likely took more care and were more thorough in their decision-making process than they normally would.

3.5. Factors influencing simulated purchases

The data from the e-commerce experiment were analysed using a binary logistic regression technique in which a best-fitting regression model was constructed for a subset of independent variables measuring the differences between the two

Table 7
Policy consultation rate

	Policy consulted (where available) (%)
Fundamentalist	29.7
Pragmatist	26.2
Unconcerned	26.8
All	25.9

Percentage of scenarios or trials where the subject consulted a privacy policy, by Westin equivalence.

e-commerce web-pages compared. In this section we present several models for different sub-groups we define.

A binary logistic model assigns coefficients to the independent variables so as to match the behavior of the dependent variable as closely as possible across all observations. The coefficients in this model can be used to determine the estimated relative importance of each factor in the decision-making process. The higher the coefficient assigned to a factor, the more it influenced the decision-making.

One of the twelve independent variables, the use of 3rd party cookies and P3P, was found not to be significant in any model, and is therefore dropped from further consideration. There are several potential explanations for why this variable proved to be insignificant across the board. It could be that subjects are not concerned about the use of 3rd party cookies and P3P in websites. In the challenge section of this survey, this was the technology users expressed least concern with. It is also possible that subjects were not familiar with the indictor used, that it brought about the wrong associations, or that it was simply not sufficiently visible.

The best-fit model including all independent variables as factors (except 3rd party cookies and P3P) leads to a fit of 8.4% (Table 8). This means that this model accounts for 8.4% of the variance in the sample. While this is not a great fit for such a model, it is not unexpected given the large number of factors which we do not control for in this experiment, and the natural variance in peoples decision-making strategies and sensitivities with regard to privacy.

The tables in this section are normalized so the coefficient of the most important variable in each model is set to 100% and the others drop according to their relative coefficient. For instance, in this case, the inclusion of a contact email contributes 56.7% less to the users' decision than the inclusion of a TRUSTe seal. The Logistic

Table 8
Best fit model with all factors

Variable	Contribution (%)	Rank	Probability
TRUSTe	100.0	1	$p<0.0001$
Policy-User-Good	93.5	2	$p<0.0001$
Policy-Corp-Good	86.2	3	$p<0.0005$
Policy-Corp-Bad	74.7	4	$p<0.0001$
Policy-User-Bad	55.4	8	$p<0.0001$
Contact Phone	74.6	5	$p<0.0001$
Contact Address	69.5	6	$p<0.0001$
Price Cut	62.3	7	$p<0.0001$
Credit Card	50.9	9	$p<0.0001$
SSL	48.8	10	$p<0.0005$
Contact Email	43.3	11	$p<0.001$
McFadden R^2	8.4		

Ranking of experimental factors by order of contribution towards explaining user actions. Contribution measured as percentage of most influential factor. Note that "Policy-User-Bad" is presented out of sequence to illustrate relative importance of policies.

regression model produces actual numbers for the coefficients, but one cannot use these to compare between models as with the percentages.

Note that Policy-User-Bad is placed out of order from the other factors (to group it with the other policy options). The order in terms of significance of the policies was as expected, users preferred policies which addressed their concerns, or in the negative case, did not confirm their fears (Policy-User-Bad ranks significantly lower than Policy-Corp-Bad). It is interesting to note what strong effect policies have despite being inspected in only a quarter of the cases. This indicates that in many cases it is that the presence of a policy has a positive effect on users, not its content.

To get a better sense of what is going on, and to study different decision-making priorities and processes, we divided the sample into relevant sub-groups. The sub-groups examined will be men versus women, the three Westin categories of users, and the way subjects use privacy policies.

The contribution of each factor to the model in each sub-sample is given in Tables 9–11. Empty cells represent cases where the factor did not prove to be statistically significant in the decision-making model. Note that values cannot be directly compared across columns, only their relative values and ranks.

When looking at the differences between men and women's decision-making (Table 9), we were not surprised to see that men followed the pattern of the overall population (with some local swapping of factors). Men constituted 74% of the survey sample; it was therefore natural that they greatly influence the global model.

The model for women proved interesting because it offered a far better fit than the global or male logistic regression model (16.1% versus 8.4% globally). Women also eliminated two factors from their model, "Contact Email" and "Credit Card". Furthermore, women exhibited clustering behavior in terms of the ranking of factors. TRUSTe was a very influential factor, with the next most influential factor, "Policy-Corp-Good", contributing 27.3% less to the decision. At the bottom of this

Table 9
Decision-making by gender

Variable	All (%)	Rank	Men (%)	Rank	Women (%)	Rank
TRUSTe	100.0	1	91.0	2	100.0	1
Policy-User-Good	93.5	2	100.0	1	63.3	3
Policy-Corp-Good	86.2	3	83.8	3	72.7	2
Policy-Corp-Bad	74.7	4	80.2	4	41.5	5
Policy-User-Bad	55.4	8	49.2	10	36.8	8
Contact Phone	74.6	5	76.2	5	40.7	6
Contact Address	69.5	6	73.2	6	29.4	9
Price Cut	62.3	7	61.4	8	55.8	4
Credit Card	50.9	9	64.9	7		
SSL	48.8	10	46.6	11	37.5	7
Contact Email	43.3	11	54.9	9		
McFadden R^2	8.4		7.8		16.1	

Decision-making and relative importance of factors by gender.

Table 10
Decision-making by Westin equivalence

Variable	All (%)	Rank	Pragmatists (%)	Rank	Unconcerned (%)	Rank
TRUSTe	100.0	1	100.0	1		
Policy-User-Good	93.5	2	83.6	2	100.0	1
Policy-Corp-Good	86.2	3			47.5	2
Policy-Corp-Bad	74.7	4	69.1	5		
Policy-User-Bad	55.4	8				
Contact Phone	74.6	5	71.4	4		
Contact Address	69.5	6	77.7	3		
Price Cut	62.3	7	49.6	9		
Credit Card	50.9	9	63.1	7		
SSL	48.8	10	69.0	6		
Contact Email	43.3	11	51.8	8		
McFadden R^2	8.4		13.9		12.7	

Decision-making and relative importance of factors by Westin Equivalence. Note that no model emerged for the Fundamentalists.

scale we see a tight cluster of factors, all within 5 percentage points of each other (rank 8–5).

The next way to divide up and examine the population was according to the Westin privacy classifications (Table 10). We present two models, one for the Pragmatists and one for the Unconcerned. None of the variables were statistically significant for the Fundamentalists, and no model could be found.

The simpler of the two models was that for the Unconcerned, consisting of only two variables, yet accounting for 12.7% of the variability of the sample. This is significantly better than what we accomplish for the general population. What we find is that the presence or absence of the "Policy-User-Good", and to a lesser extent the "Policy-Corp-Good", determined the users' choice. For the Pragmatists, the model was much more complex, including 9 of the 11 factors, and apart from a marked preference for TRUSTe indicators, there were no dramatic jumps in the weight of one factor to its nearest neighbor.

The final way of dividing the sample which we examined in this paper was according to the users' relationship with, and use of, the privacy policies (Table 11). We have seen clear indications that privacy policies greatly influence users' choice in most models, yet we know that almost half of the users never looked at a policy, and that policies were only consulted in a quarter of the trials. We therefore examined the sample two different ways. First we compare the trials in which the user checked a policy with the ones where no policy-check was conducted (*Per-Trial Policy Behavior*). Then we divided the sample based on user behavior, the group of users who checked a policy at least once were compared to the group of users who never checked policy (*Per-User Policy Behavior*). The first resulted in a roughly 25–75% split of the sample, while the second resulted in a roughly 50–50% division.

Here we see how strongly policies influence the decision-making process. In the model of trials where users consulted the policy, it accounted for 22.6% of the

Table 11
Decision-making by policy-related behavior

Variable	General case		Per-trial policy behavior				Per-user policy behavior			
	All (%)	Rank	Checked (%)	Rank	Unchecked (%)	Rank	Some checks (%)	Rank	Never checks (%)	Rank
TRUSTe	100.0	1	54.9	3	96.2	2	91.3	3	100.0	1
Policy-User-Good	93.5	2	100.0	1	73.6	6	100.0	1	48.0	8
Policy-Corp-Good	86.2	3	83.8	2	71.8	7	91.5	2	49.6	7
Policy-Corp-Bad	74.7	4	43.3	4	82.4	3	66.9	4	76.2	2
Policy-User-Bad	55.4	8	38.0	5	55.3	11	59.8	6		
Contact Phone	74.6	5			100.0	1	64.2	5	73.8	3
Contact Address	69.5	6			78.3	4	56.8	8	73.2	4
Price Cut	62.3	7			57.2	9	59.8	6	45.9	9
Credit Card	50.9	9			78.1	5	35.1	10	69.6	5
SSL	48.8	10			58.2	8	51.2	9		
Contact Email	43.3	11			57.2	9	30.7	11	49.6	6
McFadden R^2	8.4		22.6		6.9		12.4		7.2	

Decision-making and relative importance of factors by per-user and per-experiment policy-reading behavior.

variance in the sample. Furthermore, we saw clear evidence of discrimination between the good and the bad policies, with users showing a strong preference for the good. Interestingly, the only non-policy factor to remain in this model was the TRUSTe seal, coming in between the good and the bad policies. This shows the strength of positive associations users have with these types of trust-marks.

When we looked at the unchecked policy model, we saw that policies are still important in the decision-making, though the ordering of the good and bad policies is arbitrary since they were not actually read. Interestingly enough, this is the only model which rated the presence of "Contact Phone" as the most significant factor. What we saw was clear evidence of how people used these factors to determine "trustworthiness", not based on fact but rather on appearance and first impression. Policies are important, not just because of what they say, but because they are there. As we saw, this model offered a far worse fit for the data, demonstrating the finality that policy checking brings to decision-making.

When we look at the second category we see similar behavior. For those who never check policies we find that TRUSTe dominates over the other factors, and that policies, though never read, have a fairly powerful effect. Policy-checkers naturally closely match the behavior of the policy checked group, as they on average checked policies in half the trials.

4. Discussion of results

In this section, we discuss the results and possible threats to their validity. Broader implications are discussed in the next section.

4.1. Privacy classification

The Westin privacy segmentation is a way of dividing and thinking about privacy sensitivities which has been widely adopted and embraced by industry and academia. In our survey we did not use the same questions as Westin, but we were able to find very similar, very cohesive groups, especially in the privacy questionnaire. It is therefore likely that we identified the same groups as the Westin surveys have identified in the past.

We think that it is interesting and important to demonstrate that these groups are significant and identifiable outside the context of the three traditionally posed questions, and that they employ very different decision-making strategies. It is also interesting to see where these groups started to lose their significance, especially in the analysis of experiment. One of the most surprising findings was that we were unable to find a logistic regression model for the Fundamentalist group, where none of the twelve independent variables were significant.

There are several potential explanations for the lack of a model for the Fundamentalists. This group may itself be a collection of very different sub-groups, all highly concerned about personal privacy, but with very different decision-making strategies. Or perhaps the Fundamentalists are not really influenced by any of the factors we included in this study. This would mean that this group, though as likely as any other to read privacy policies, is not influenced by them, or does not trust what policies say. It is also possible that Fundamentalists are looking for different types of information in policies than the rest of the population.

Another interesting finding with regards to these three user groups is that Fundamentalists are no more likely to install or use software to protect themselves than the Pragmatist group, though they are different from the Unconcerned. This means that the Pragmatists should not be ignored as consumers or early adopters of privacy enhancing technologies.

4.2. Gender

There are interesting indications throughout the study of gender differences. Women were underrepresented in the study and therefore some gender-specific questions could not be answered.

The Westin classification may confound gender, since a smaller proportion of women were classified as Unconcerned than were men, a result that approached statistical significance. We do not know whether this means that women tend to be more risk-averse, more pessimistic or skeptical about the motives and honesty of online vendors, less knowledgeable about the technology in question, or more knowledgeable about the risks of online transactions (e.g. having suffered more from the consequences of identity theft or fraud or being familiar with the fate of friends and associates who have been). Our study design and the number of people falling in some of these categories (particularly victims of fraud or identity theft) are such that we cannot investigate the reasons further.

The stratified model for women was by far the best fitting model for the experimental scenario. Although the presence of the TRUSTe seal was one of the most significant variables in all models, it is noteworthy that for women, the presence or absence of the seal was a much more significant factor than any other.

4.3. Indicators

Third-party cookies and P3P was the only indicator not to prove significant in any model. There are several potential explanations for this. It could be that subjects are not concerned about the use of third-party cookies, or are not familiar enough with P3P to make decisions based on this indicator. In the challenge section of this survey, cookies were the technology users expressed least concern with. However, is also possible that subjects were not familiar with the indictor used on the web page to signal the presence of third-party cookies. We did not test subjects' knowledge of indicators. However it is consistent with this interpretation that cookies were the technology for which subjects' self-reported knowledge diverged most from their ability to answer the challenge question.

Perhaps users were unable to process information about cookies coherently. This is an unlikely explanation. A more likely one, given the low effect of the SSL encryption indicator is that these icons are too inconspicuous, and that many users do not notice or pay attention to these indicators.

4.4. Policies

The order of significance of the policy variables in the regression analysis was as expected: users preferred policies which addressed their concerns rather than the company's; in negative cases, they preferred policies that did not confirm their fears (that is Policy-User-Bad ranks significantly lower than Policy-Corp-Bad).

It is interesting to note that the presence of a visible link to a privacy policy has a major effect on purchasing behavior, even though only a quarter of the policies were consulted. In most cases, users had more confidence in a site simply because it had a policy.

It was the Unconcerned users who were most influenced by the content of the policies. The picture that emerges here is of users who take a more casual approach to the evaluation of privacy risks, yet are strongly swayed by the assurances made in policies. Since one of the questions used to categorize users referred to the trustworthiness of policies, it is not surprising that the Unconcerned were more affected by the policies alone than other, potentially more skeptical users.

5. Implications

We conclude with a discussion of the implications of our results in the following four areas:

- The design of user interface indicators so that users understand and may act on privacy-relevant information.
- The wisdom of classifying users into categories along a single dimension of privacy concern.
- Implications for research methodology of the contrast between the results obtained from self-reports and those obtained through experimental economic scenarios.
- Public policy implications, such as the regulation and legislation of how and when users must be notified of privacy practices and policies, together with limits to the notion of informed consent.

5.1. Design

While this work provides important guidance for business, policymakers, and management, it also provides important insights for interface designers. In our experiment there is a set of variables which we can call "trust-marks", factors which may not say anything about the site's privacy practices, but which are interpreted as such by users.

One such factor, the TRUSTe marker, very important in most models, actually says nothing about the practices of the site. This marker serves as a guarantee that the site discloses a minimum set of information in its policy, and that it follows the practices it claims rather than what the policy says about these practices. The TRUSTe marker should have been a powerful indicator, but only in conjunction with a privacy policy.

Privacy policies themselves serve as "trust-marks", evident we think from the impact they have on users who never consult them. What we are seeing is that the presence of a policy has a significant effect on decision-making regardless of whether the policy was read or not. The impact a policy has is of course more powerful when it is read, but it is not negligible when it is not.

Other factors which can be classified as "trust-marks" are the credit card icons and the contact information variables. The credit card icons are interesting because they do not in fact imply any promise of fraud prevention or privacy protection. Just about every e-commerce site accepts some form of credit card payment (some operate on electronic payment systems such as Paypal), and it is therefore not clear why consumers should find these icons reassuring.

As to the inclusion of contact information, there was a strong preference for phone information over mailing or email information. This means that users were looking for ways of holding companies accountable, or indicators for a company's willingness to dealing directly with them should they have any problems as a result of this transaction. One interesting question is whether users would actually test to see if the phone number was valid before buying from a site. In these tests, the phone number was plainly invalid.

The impact these "trust-marks" have on decision-making, across all user groups, shows a clear need for designers to develop privacy-enhancing technologies which give users simple and clearly visible trust indicators. If these markers are not clearly visible they may be ignored, as we saw with the SSL encryption icon and Cookie-blocking and P3P. These indicators may be too inconspicuous for users to notice. While it is possible that users do not place a lot of value in these factors, we believe that the reason that relatively meaningless indicators, such as the credit-card icons, are preferred is because they are more clearly visible.

What we found in this study, like others, was that only a minority of subjects read policies with any frequency. The information contained in these policies is considered highly significant, and highly influential in users' decision-making, but is rarely sought out. In this experiment, where the rate of policy consultation was likely inflated by subjects knowing the purpose of the experiment, we found that subjects only consulted policies in a quarter of trials. Other studies have shown this rate to be much lower in real life, by as much as a factor of ten (Jensen and Potts, 2004).

These findings all argue for the development of policy simplifications, standardization, or machine readable policies. Based on this data we can also make a strong case for the need to develop and implement standardized, simple visual indicators for the practices and technologies websites use, and the risks users are exposed to.

5.2. User classification

When we ran the logistic regression analysis on the Fundamentalist sample we were surprised to find that none of the variables were statistically significant, and that no model could be found. This is especially surprising because the Fundamentalists were the second largest user group with 34% of the sample population, surpassed by the Pragmatists (43%) and outranking the Unconcerned (23%), both of whom provided models accounting for over 12% of the variance. This lack of a model for the Fundamentalists is interesting, because the Westin categories are used in marketing particularly to isolate those consumers who are most likely to embrace privacy-protecting products and services. Our results, however, indicate that while Fundamentalists have a consistent concern with privacy, they do not form a cohesive group with respect to decision-making. Only the Unconcerned and Pragmatists are internally cohesive groups.

Colloquially speaking, it seems that while it may be more difficult to push the other groups' buttons, they do at least have some; the Fundamentalists, in contrast, don't seem to have a single set of buttons to push. Perhaps market researchers should turn their attention to how the concerns of the less concerned groups can be mobilized rather than concentrating on the more diverse concerns of the Fundamentalists.

The main alternative explanation for the lack of a model for the Fundamentalists, as discussed in the previous section, is that Fundamentalists were not influenced by the factors we used in this experiment. This seems an ad hoc explanation, however, given that the other groups behaved as expected and that there were twelve plausibly variables under investigation. Also worth noting is the finding from our survey that Fundamentalists were no more likely than others to install privacy-protecting technology. The most parsimonious explanation appears to be that Fundamentalists are not really that "fundamentalist" about privacy at all.

5.3. Research methodology

In general, the study demonstrates that users do not do what they say, and they do not know what they claim to know. Although the subjects of this survey consulted online privacy policies more often than previous log-based studies indicate users "in the wild" do, their behavior did not match their survey statements. Subjects were also generally not able to answer questions about privacy-related technology that they claimed to know, a trend particularly noticeable in the case of cookies, where they reported the highest knowledge.

Such results call for a reevaluation of the role of surveys in the study of Internet behavior. Surveys appear to be best suited to the evaluation of attitudes and opinions rather than behaviors or experience. Where issues arise, such as the role of perceived competence in decisions-making, the use of self-reports is invaluable as a baseline against which actual behavior can be compared. The self-reported data should not be taken as evidence of behavior, however. Indeed, the shakiness of our subjects' self-reports and judgment of knowledge leads us to wonder whether their experience reports should be taken seriously. How many years experience a user has and how intense the usage during that period may have been may be very different from what the user reports in a survey. In addition to demand characteristics of the survey situation (albeit in a situation of personal anonymity) simple forgetting and salience effects are likely to skew the user's recall and categorization of his or her experience.

In place of a full reliance on survey data, we see the need for more concerted attempts in the area of experimental economics in which users are put in realistic situations and required to indicate preferences or make economic decisions such as bidding in auctions or purchasing items with simulated funds. To capture the effect on decision-making of the context of previous decisions and of current affairs (e.g. news about technology vulnerabilities or protections) such studies should ideally be longitudinal rather than one-shot experiments. Such a shift in methodology would likely increase the ecological validity of research instruments and settings but would require a wholesale shift in how we plan studies, recruit subjects, and standardize instruments. As an example of standardization, consider the role played in the current study of the various online policies. Research into the effect of policy content on online behavior clearly requires that the policies used in different studies be systematically comparable if not identical.

The experiment reported here was not intended to be a full-fledged study in experimental economics. We did not investigate systematically the trade-off between privacy indicators and a range of price points or product attributes for the items for sale. Instead, we presented a 20% difference in price of an identical item available from two vendors. Different price differentials, or tradeoffs between price and product quality could interact with privacy awareness in complex ways. Nor did we vary the items for sale in a way that would assess the sensitivity of the models to the nature of the product. (It is unlikely, for example, that privacy indicators would have the same impact relative to price when the consequences of disclosure are more sensitive than model cars—such as pharmaceuticals or erotica. Although when users would behave more cautiously and when more recklessly and how these behaviors might interact with demographic and personality variables is hard to predict.)

In connection with the distinction between professed and actual knowledge, it would be interesting to know whether knowledge really is power and whether a little knowledge is a dangerous thing. Are users who have more knowledge about the privacy implications of Internet technology better able to make more effective discriminations among web sites and services? How does user knowledge relate to user confidence in making online purchasing decisions, and does a little knowledge (for example, of cookies in the case of our study) lead to overconfidence and/or reduced effectiveness? Do general demographic factors, such as amount of online experience, education, age and gender play a role in modulating the answers to these two questions? Unfortunately, our study design does not permit these analyses. It does suggest that they should be fascinating questions to answer in future research.

5.4. Public policy

Public trust in technology rests on the policies that regulate technology and on the doctrine of informed consent. This is not the forum to discuss the broader issues of policymaking and regulation in any depth. However, the notion of agreement and informed consent obviously relies heavily on user's understanding of both technology and the consequences of online behavior. It appears from the results presented above that many users have an incorrect understanding of their own knowledge of technology, their online behavior, and potential consequences. Not only do users frequently fail to consult online privacy policies, when they do, the policy may not help them make informed decisions. Recent studies have found that online privacy policies are difficult to find and demand levels of reading skill to understand that are not typical of the Internet population (Jensen and Potts, 2004; Antón et al., 2004) In view of recent court rulings in the US insurance industry stating that policies must be worded in plain language in order to be enforceable as contracts, the obscurity of privacy policies may call their validity into question.

To consult policies regularly would be very inefficient and dysfunctional unless the likely consequences of not doing so were punitive. Users therefore seem to adopt a strategy of sporadic checking, possibly triggered by the presence of suspicious indicators, in conjunction with the heavy use of proxies or surrogates.

The most significant of these proxies are trust markers. For these to serve as surrogates for detailed inspection of policies, trust marks need to be quality indicators, and not merely presence indicators. In the case of the TRUSTe mark, users appear to take its presence as evidence of the quality of the privacy policy, not merely that the vendor has a privacy policy and follows it.

Trust marks that are presence indicators but not quality indicators do not encourage deception by vendors, but they do make it possible. Unscrupulous vendors could use such marks as camouflage for policies and practices that users would not willingly agree to.

The results of this study show that even self-selected volunteers in a survey on online privacy, who are therefore likely predisposed to think about privacy issues, and who know that their online behavior is being monitored, still show remarkable ignorance and inappropriately placed trust in their actions. To avoid exploitation and consequent reduction of that trust, greater public awareness of privacy issues, the capabilities and limitations of privacy-enhancing technologies and the significance of policies and trust indicators are all necessary.

Acknowledgments

This research was supported by NSF ITR Grant 0325269. The authors thank TRUSTe for permission to use their mark in these experiments, and Roger Los of http://www.los.com/ and http://austinhealey.com/ for permission to use his artwork. We also thank Ji Han Bae and William Stufflebeam for their help in designing and running this survey.

References

Adkinson, W.F., Eisenach, J.A., Lenard, T.M., 2002. Privacy online: a report on the information practices and policies of commercial web sites. Progress and Freedom Foundation, Washington DC. Online: http://www.pff.org/issues-pubs/books/020301privacyonlinereport.pdf.

Antón, A.I., Earp, J.B., Bolchini, D., He, Q., Jensen, C., Stufflebeam, W., 2004. The lack of clarity in financial privacy policies and the need for standardization. IEEE Security & Privacy 2 (2), 36–45.

Cranor, L.F., Reagle, J., Ackerman, M.S., 1999. Beyond concern: understanding net users' attitudes about online privacy, AT&T Labs-Research Technical Report TR 99.4.3, http://www.research.att.com/library/trs/TRs/99/99.4/99.43/report.htm.

Culnan, M.J., 1999. Georgetown Internet Privacy Policy Survey: Report to the Federal Trade Commission. Georgetown University, McDonough School of Business, Washington, DC Online: http://www.msb.georgetown.edu/faculty/culnanm/GIPPS/gipps1.pdf.

Culnan, M.J., Milne, G.R., 2001. The Culnan-Milne survey on consumers & online privacy notices: summary of responses. Proceedings of Get Noticed: Effective Financial Privacy Notices. A Federal Trade Commission Workshop. Washington, DC. Online: http://www.ftc.gov/bcp/workshops/glb/supporting/culnan-milne.pdf.

Earp, J.B., Meyer, G., 2000. Internet consumer behavior: privacy and its impact on internet policy. Proceedings of the TPRC 28th Research Conference on Communication, Information and Internet Policy, Alexandria, VA.

Federal Trade Commission (FTC), 2000. Privacy online: fair information practices in the electronic marketplace: a report to congress. Online: http://www.ftc.gov/reports/privacy2000/privacy2000.pdf.

Harris et al., 1998. E-commerce & privacy: what net users want. Privacy & American business and price water house coopers LLP. Online: http://www.pandab.org/ecommercesurvey.html.

Harris Interactive, 2003. The Harris Poll® ♯17: most people are 'privacy pragmatists' who, while concerned about privacy, will sometimes trade it off for other benefits. Online: http://www.harrisinteractive.com/harris_poll/index.asp?PID=365.

Javelin Strategy & Research, 2005. 2005 Identity Fraud Survey Report. Online: http://www.javelinstrategy.com/reports/2005IdentityFraudSurveyReport.html.

Jensen, C., Potts, C., 2004. Privacy policies as decision-making tools: a usability evaluation of online privacy notices. Proceedings of CHI'04. Vienna, Austria, pp. 471–478.

Jupiter Research, 2002. Security and privacy data. Presentation to the Federal Trade Commission Consumer Information Security Workshop. Online: http://www.ftc.gov/bcp/workshops/security/020520leathern.pdf.

National Telecommunications and Information Administration (NTIA), 2002. A nation online: how Americans are expanding their use of the Internet. Washington, DC. Online: http://www.ntia.doc.gov/ntiahome/dn/.

Synovate, 2003. Identity theft survey report, prepared for the Federal Trade Commission (FTC). Online: http://www.ftc.gov/os/2003/09/synovatereport.pdf.