

Power Strips, Prophylactics, and Privacy, Oh My!

Julia Gideon
Carnegie Mellon University
jgideon@cmu.edu

Lorrie Cranor
Carnegie Mellon University
lorrie@cs.cmu.edu

Serge Egelman
Carnegie Mellon University
egelman@cs.cmu.edu

Alessandro Acquisti
Carnegie Mellon University
acquisti@cmu.edu

ABSTRACT

While Internet users claim to be concerned about online privacy, their behavior rarely reflects those concerns. In this paper we investigate whether the availability of comparison information about the privacy practices of online merchants affects users' behavior. We conducted our study using Privacy Finder, a "privacy-enhanced search engine" that displays search results annotated with the privacy policy information of each site. The privacy information is garnered from computer-readable privacy policies found at the respective sites. We asked users to purchase one non-privacy-sensitive item and then one privacy-sensitive item using Privacy Finder, and observed whether the privacy information provided by our search engine impacted users' purchasing decisions (participants' costs were reimbursed, in order to separate the effect of privacy policies from that of price). A control group was asked to make the same purchases using a search engine that produced the same results as Privacy Finder, but did not display privacy information. We found that while Privacy Finder had some influence on non-privacy-sensitive purchase decisions, it had a more significant impact on privacy-sensitive purchases. The results suggest that when privacy policy comparison information is readily available, individuals may be willing to seek out more privacy friendly web sites and perhaps even pay a premium for privacy depending on the nature of the items to be purchased.

Categories and Subject Descriptors

K.4.1 [Computers and Society]: Public Policy Issues—*Privacy*; K.4.4 [Computers and Society]: Electronic Commerce; H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Evaluation/Methodology*

Keywords

P3P, Privacy Policies, Search Engines, E-Commerce, User Studies

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee. Symposium On Usable Privacy and Security (SOUPS) 2006, July 12-14, 2006, Pittsburgh, PA, USA.

1. INTRODUCTION

Surveys and empirical studies have shown that while individuals claim to have a high level of concern about Internet privacy, they rarely take steps to actively protect their privacy online. A number of possible explanations for this behavior have been proposed, including the fact that it is generally difficult for individuals to obtain information about privacy-friendly alternatives, especially when making online purchasing decisions. We conducted a study to determine whether users would take privacy information into account when making online purchasing decisions if this information were made available alongside search engine results.

In a 1991 survey commissioned by Equifax, researchers found that individuals generally fall into three groups with regard to privacy concerns. The "privacy fundamentalists" are extremely concerned with how their personal information is used and therefore are generally unwilling to share it with anyone; the "privacy pragmatists" share some of these concerns but prefer to make decisions on a case by case basis; and the "privacy unconcerned" are generally willing to give away their personal information without much thought whenever it is requested of them [26]. While most individuals fall into the pragmatists category, the proportion who are unconcerned has diminished from 22% to 10% as of 2003 [24, 22]. One of the biggest privacy concerns among the concerned groups is how a company will use or share stored personal information. A 2000 survey conducted by the Pew Internet & American Life Project showed that 86% of respondents were concerned that companies they had done business with in the past may reuse their stored personal information without first seeking permission [9].

While most users claim to be concerned with Internet privacy, it is not clear that their behavior reflects their concerns. According to a 2003 survey, only 16% of Americans have purchased a privacy-enhancing product. Such products include credit reports, anonymous web browsing tools, and other tools to help prevent identity theft [11]. Additionally, some users' behaviors go contrary to what they say regarding their privacy concerns. In a 2001 experiment, 24% of self-described privacy fundamentalists disclosed personal information that was not required to complete a transaction [20]. Because of the propensity for users to describe their behaviors inaccurately, further user studies are needed to determine actual user behavior with regard to privacy.

Although many web sites post their privacy policies in an attempt to address consumer privacy concerns, very few

consumers bother to actually read them [7]. It is also unclear whether those who read privacy policies understand much of what they read as most privacy policies are rich in legal jargon, have no standard format, and use language that requires college-level reading comprehension skills [13].

The Platform for Privacy Preferences (P3P) was created in an attempt to solve some of the problems with privacy policies. P3P, which defines a standard XML format for privacy policies, was created by the World Wide Web Consortium (W3C) [4]. P3P user agents can check for machine-readable P3P privacy policies at every web site a user visits and compare these policies with the user's pre-defined privacy preferences. If the web site's privacy policy does not conform to the user's preference, the user agent may take actions such as warning the user, blocking cookies from the site, or blocking all access to the site, depending on the user's stated preferences and the features of the P3P user agent.

One shortcoming of most currently-available P3P user agents is that they provide privacy notifications to a user only after the user has accessed a particular web site. Thus, they expose some information to web sites before users have had an opportunity to receive a privacy notification, and they require users who are seeking out web sites with acceptable privacy policies to visit multiple web sites one at a time until they find one that matches their privacy preferences.

The Privacy Finder P3P-enabled search engine service was designed to make privacy policy information more accessible and to make it easier to compare privacy policies across multiple sites. When users select their privacy preferences and enter search terms, Privacy Finder examines every search result that is to be displayed in an attempt to compare the destination sites' P3P privacy policies with the user's preferences. The search results can thus be annotated with privacy information, and users can observe which sites comply with their preferences without having to first visit each site. Since Privacy Finder is visiting these sites, the user does not need to leave any identifying information at sites that he or she does not personally visit (when locating a P3P that is not in its cache, Privacy Finder will visit the site and only leave its IP address). Privacy Finder maintains a cache of all P3P policies found. According to the standard, every policy has an expiration date. Thus, Privacy Finder uses this information to determine when it needs to retrieve a current P3P policy from the destination web site.

Now that we have the ability to provide privacy information with search results, we are interested in finding out whether users make use of that information when selecting e-commerce web sites from which to make purchases. In this paper we discuss the user studies that we conducted to address this question. In Section 2 we examine previous studies on privacy and online trust decisions and provide some additional background on Privacy Finder. In Section 3 we describe our user study methodology, and in Section 4 we present our results and analysis. We discuss the limitations of our study and plans for future work in Section 5. Finally, we present concluding remarks in Section 6.

2. BACKGROUND

Privacy is an important issue for a majority of Internet users. In this section we will discuss previous work related to the valuation of privacy, trust of web sites, and e-commerce

privacy concerns. We will also introduce Privacy Bird and the Privacy Finder service.

2.1 Valuation of Privacy

Empirical studies on how consumers value privacy have highlighted a dichotomy between professed attitudes and actual behavior, raising questions about individuals' true valuation of privacy that researchers have tried to answer through experimental approaches.

On one hand, many individuals claim to value privacy so highly that they are willing to accept inconveniences in exchange for increased privacy. A 1998 Business Week/Harris Poll survey found that among the 77% of Internet users who had never purchased products on the Internet, 86% were holding back due to concerns about the use of their personal and financial information [25]. In 2000, a Price-WaterhouseCoopers study claimed that nearly two thirds of the consumers surveyed "would shop more online if they knew retail sites would not do anything with their personal information" [3]. In February of 2002, a Harris Interactive study found that the three biggest consumer concerns in the area of online privacy were companies sharing personal data without permission, the consequences of insecure transactions, and theft of personal data [10]. In the same year, Jupiter Research calculated that by 2006, \$24.5 billion in online sales would be lost due to privacy concerns [14].

On the other hand, numerous studies have shown that consumers often are willing to provide personal information in exchange for very small rewards. Another 2002 Jupiter Research study found that 82% of online shoppers would give personal data to new shopping sites in exchange for the chance to win \$100 [23]. Presenting the results of the 2003 Harris privacy poll, Taylor [22] notes that most people are concerned about privacy, but will "sometimes trade it off for other benefits."

These surveys paint a nuanced picture: in large numbers, Internet users claim to highly value their privacy; still, they are willing to trade off personal information for small rewards, or are unwilling to change their behavior when privacy threats arise. Several possible explanations for this dichotomy have been discussed in the literature [2, 18, 21, 1, 16].

In recent years, efforts have been directed towards empirical studies of consumers' valuation for privacy under different conditions. Researchers at Berlin Humboldt University simulated an online shopping environment in which an anthropomorphic 3-D shopping bot posed a variety of personal questions to shoppers. Many of these questions requested information unnecessary to the shopping task. In order to receive discounts on the purchase of certain goods, subjects answered a majority of the personal questions asked by the bot, even if they had previously claimed to have high privacy concerns. The authors also found that the content of the privacy statements associated with the bot had no effect on the amount of information disclosed by the subjects [20].

Another study used a second-price auction experimental setup to study the monetary value of private information to individuals. Using weight and age as two types of information that the subjects may be sensitive about (and therefore value), the authors found that subjects demanded higher prices to reveal information they viewed as having a higher deviation from group norms (for example those who were older or heavier than the other group members

on average demanded higher prices for revealing this information) [12]. In another experiment, the researcher used a contingent valuation survey approach to estimate the economic value subjects place on a change in the data protection laws that would give the subjects enforceable property rights over their personal information. The author found that while most survey participants expressed high sensitivity to privacy, their willingness to pay for such strong property rights was low — only 47.5% of those surveyed would pay for it an average of NZD 55.40 (USD 28.25) [17]. These and other studies suggest that consumers are willing to provide personal information in exchange for small rewards.

2.2 Trust of Web Sites

Consumer privacy concerns along with pressure from the US Federal Trade Commission (FTC) resulted in an increase in the posting of online privacy policies in recent years. However, there are many users who do not trust that online privacy policies truly reflect a company’s practices. Only 29% of those surveyed in 2001 agreed that they can strongly trust privacy policies, while 52% said that they were unsure. When asked if privacy policies should be distrusted 52% neither agreed nor disagreed, while 34% disagreed. Over half of those surveyed said they sometimes read privacy policies of web sites upon visiting them for the first time. This implies that users are interested in knowing the privacy practices of companies with whom they are unfamiliar. A slightly smaller percentage (about 45%) of users will look back over the privacy policy of a site if they believe that the policy has changed [7]. There is a disconnect here in that web site visitors want to know certain information from companies, but are not fully trusting of it once it is presented to them.

In the United States, only certain industries such as banking, insurance and healthcare are required to post privacy policies, and enforcement is far from uniform. The FTC has the authority to take action against companies that deviate from the practices expressed in their posted policies (even if those companies posted their policy voluntarily), however the FTC has limited resources and cannot pursue every company that posts a fraudulent privacy policy [7].

Privacy seals can help consumers determine whether web sites meet minimum standards when posting privacy policies. However, a 2005 study found that while most users understand that seals have something to do with privacy, most could not identify any of the most common seal programs, most did not know how a site earns a privacy seal, and few thought they were important in choosing a web site [15]. As many users do not read the web site privacy policies, privacy seals have many problems from a policy standpoint as well [7]. The specific practices of companies that display privacy seals can differ greatly, unbeknownst to most users. Thus, privacy seals tend to give a false sense of security. This is similar to the belief that the existence of a privacy policy is indicative of favorable privacy practices. For instance, 57% of Internet users incorrectly believe that web sites with privacy policies will not share personal information with third parties [24]. This indicates a general lack of understanding when it comes to the nature of web site privacy policies.

Research has also shown that many Internet users base trust decisions about web sites largely on the overall look and feel of the site. Adequately addressing privacy concerns is not a major factor in these trust decisions. When de-




Icon	Site..
	...matches privacy preferences.
	...conflicts with privacy preferences.
	...has an error in its P3P policy.

Table 1: Privacy Finder icons.

termining whether or not a web site is credible, most users largely use such factors as whether the web site is a known company outside of the Internet and whether or not the company has physical locations [8].

2.3 Privacy Finder

Privacy Finder is based on the Privacy Bird¹ P3P user agent, which displays colored bird icons and plays bird sounds to indicate whether or not a web site’s P3P policy matches a user’s privacy preferences. Privacy Bird users can click on the bird icons to bring up an English translation of a web site’s P3P policy in a standard format. Privacy Finder takes a similar approach, using a set of colored bird icons to annotate search results with information about whether each result matches or conflicts with a user’s privacy preferences. A green bird indicates that the site’s privacy policy complies with user preferences, while a red bird indicates a conflict. A yellow bird is displayed when the site has critical errors in its privacy policy such that Privacy Finder is unable to parse it.² Sites not posting P3P policies are not annotated with any icons. The specific icons used can be seen in Table 1.

Privacy Finder makes use of the Privacy Bird preference setting interface. It uses three standard preference settings — high, medium, and low privacy — as well as 12 warning conditions that users may individually select in order to customize their settings. The standard settings map to the 12 warning conditions, as shown in Table 2.

In addition to immediately providing users with privacy information for each search result, Privacy Finder also makes understanding privacy policies easier. When moused over, the bird indicator either explains to users that a site is in compliance with their preferences, or it will enumerate all the reasons why it conflicts with their preferences. Each bird indicator also serves as a link that takes the user to a page with a “privacy report” created by translating the computer readable P3P policy into English. The format of the privacy report emphasizes key sections of privacy policies that are likely to be of most interest to users; for example, information about a company’s data sharing practices and information about how to opt-out of data sharing and marketing solicitations. An example privacy report is shown in Figure 1.

3. METHODOLOGY

The goal of our study is to determine the effect of privacy information presented by a P3P-enabled search engine on

¹<http://www.privacybird.com/>

²Note that Privacy Bird displays a yellow bird to indicate a site with no P3P policy.

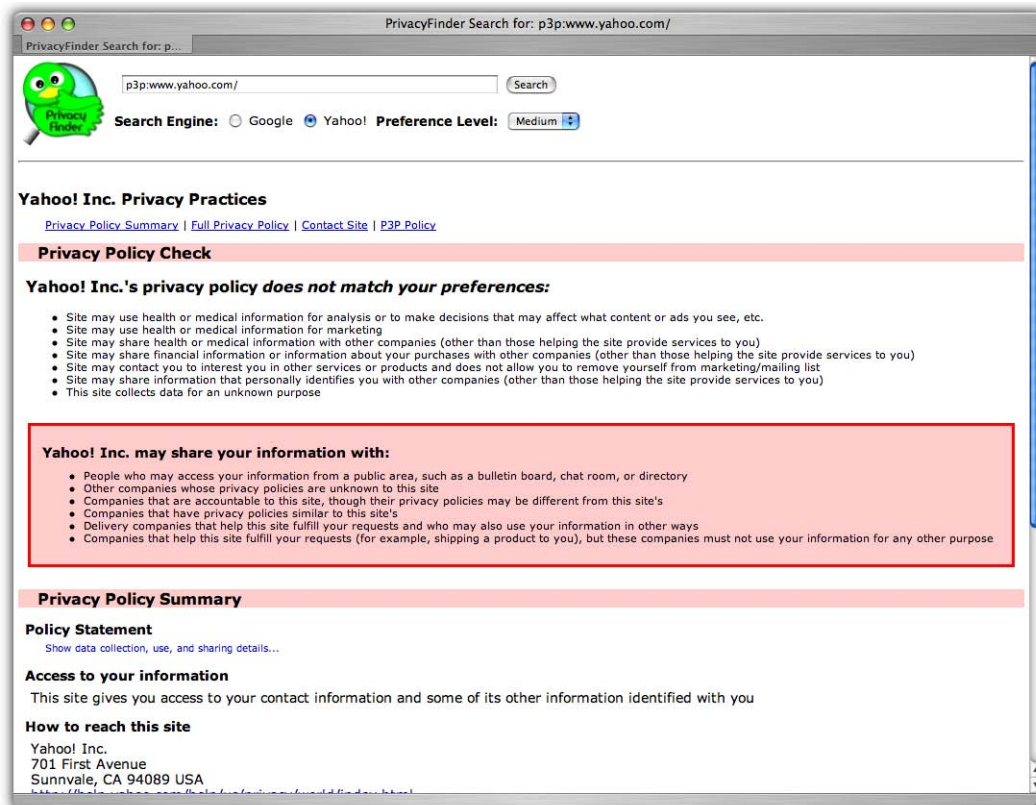


Figure 1: Privacy Finder’s privacy report screen.

online purchasing decisions. The study consisted of three stages: a screening survey, a laboratory experiment, and an exit survey.

Study participants were students at Carnegie Mellon University. Advertisements were posted throughout the campus and on a high-traffic online student bulletin board. Prospective participants interested in the study were instructed to respond via email. These methods of solicitation helped to ensure that participants were at the very least familiar with email communication, and were able to use computers for basic tasks. We had four prerequisites for this study. Participants had to be at least 18 years of age, had to have a personal credit card, had to have had at least one previous online shopping experience, and had to express at least a minimal level of privacy concern in response to the privacy-related questions on our screening survey.

We selected 24 participants and randomly divided them into two separate groups, a control and an experimental group. Both groups were told that they were taking part in an online shopping study. The control group searched for products with a version of the Privacy Finder that did not actually report any privacy policy information. The experimental group participants used a modified version of the full Privacy Finder service. Participants were told that they would be making purchases with their own credit cards and that they would be reimbursed for their purchases and paid an additional \$10 for their participation in the study.

3.1 Screening Survey

A screening survey containing twenty-two questions was

administered by email to those who responded to the advertisements that were placed online and around campus. The screening survey was used primarily to make sure participants met the four pre-requisites for the study. It was also used to gain a better understanding of participants’ privacy concerns and so that we could verify that the information presented by Privacy Finder addressed these concerns. Respondents who were deemed eligible to participate were later contacted to set up an appointment to complete the shopping experiment.

The self-reported privacy preferences of the twenty-four participants selected for the study can be seen in Table 3.

3.2 Laboratory Experiment

The laboratory experiment involved participants using a search engine to select web sites from which to purchase two specified products. In the subsections that follow, we explain our choice of products, our experimental setup, and the experimental protocol.

3.2.1 Product Selection

We decided to select two products for participants to purchase. We looked for one product that would be typical of a business or household purchase and would not raise any particular privacy concerns in and of itself (thus the privacy concerns associated with the purchase would largely be related to concerns about the use and disclosure of payment and contact information). We looked for a second “privacy sensitive” product that would be likely to raise additional privacy concerns because participants might feel uncomfort-

	4	3	2	1	0	Average
Site shares your financial information with other companies	21	3	0	0	0	3.88
Site does not allow you to be removed from marketing/ mailing lists	22	1	1	0	0	3.88
Site shares your health information with other companies	19	3	2	0	0	3.75
Site does not allow you to find out all the information it keeps on you	18	5	1	0	0	3.71
Site contacts you about other services or products via telephone	17	5	2	0	0	3.63
Site shares information that identifies you with other companies	16	4	3	0	1	3.42
Site uses your financial information for deciding web site content or ads	12	9	2	1	0	3.33
Site uses your health information for deciding web site content for ads	12	2	8	2	0	3.00
Site contacts you about other services or products via email or postal mail	6	10	3	5	0	2.71
Site uses information that identifies you to determine habits, interests, or other characteristics	7	2	11	2	2	2.42
Site shares information that does not personally identify you with other companies	3	5	10	3	3	2.08
Site uses information that does not personally identify you to determine habits, interests, or other characteristics	1	4	11	4	4	1.75

Table 3: User privacy preferences as captured by the screening survey. Questions were answered on a 5-point Likert scale, with a ‘0’ meaning that the individual “likes that practice a lot,” a ‘2’ indicating indifference, and a ‘4’ indicating that he or she “doesn’t like that practice at all.”

Warn when...	Low	Med	High
...site collects health or medical info for analysis or marketing.	X	X	X
...site shares health or medical info with others.	X	X	X
...site collects financial info for analysis or marketing.			X
...site shares financial info with others.		X	X
...site may contact me by telephone.			X
...site may contact me via other means.			X
...site does not allow me to opt-out from marketing lists.	X	X	X
...site uses personally identifiable info to analyze me.			X
...site shares personally identifiable info with others.		X	X
...site does not allow me to see the info collected on me.		X	X
...site uses non-personally identifiable info to analyze me.			X
...site shares non-personally identifiable info with others.			X

Table 2: Table of privacy preference levels.

able having other people know that they had purchased that product. Due to budgetary considerations, we looked for products that were typically available for around \$10. In addition, we needed to find products that were available from multiple web sites offering a range of P3P policies.

A person’s familiarity with a particular site can persuade them to buy from the site regardless of the site’s privacy practices [19]. The fact that a company is trusted can entice consumers to disclose information without considering other factors [5]. Similarly, consumers are apt to not read

the privacy policies of companies that are well-known or companies with whom they have done off-line business [6]. This leads to the need to select a product that is not associated with any well-known company. For instance, it would not be reasonable to ask participants to choose a site from which to buy a CD. Based on an informal poll of CMU graduate students familiar with the IT field, we ascertained that most are accustomed to buying from a select number of sites that sell music, such as Amazon.com, Sam Goody, or Tower Records. This forced us to focus on items that most people do not purchase regularly online, but are still readily available from multiple online vendors. Likewise, the specific product to be purchased must be similar across all of the search results. If participants are given two different sites with products that are of varying qualities, they may focus on the choice of products rather than the choice of merchants.

We selected a surge protector as a product typical of a business or household purchase and a box of condoms as a product likely to raise privacy concerns. A specific type of surge protector and a specific brand and type of condoms were specified. These items were selected after verifying that they met all of the above product selection criteria.

3.2.2 Experimental Setup

The experiment was conducted on laptop computers in our usability laboratory. Each computer had the Firefox web browser loaded and displayed the front page of a modified version of the Privacy Finder search engine. In order to reduce the effects of priming we removed the Privacy Finder name and logo and referred to the search engine as “Shopping Finder.” We also removed the privacy preference setting and configured the search engine to always use the “medium” privacy setting. Privacy Finder is able to use either the Yahoo! or the Google APIs for conducting searches. For the purpose of this experiment, we configured the search engine to always use the Yahoo! API. A screenshot of the Shopping Finder results page can be seen in Figure 2.

Search engine results change frequently and can vary depending on whether users capitalize search terms or make

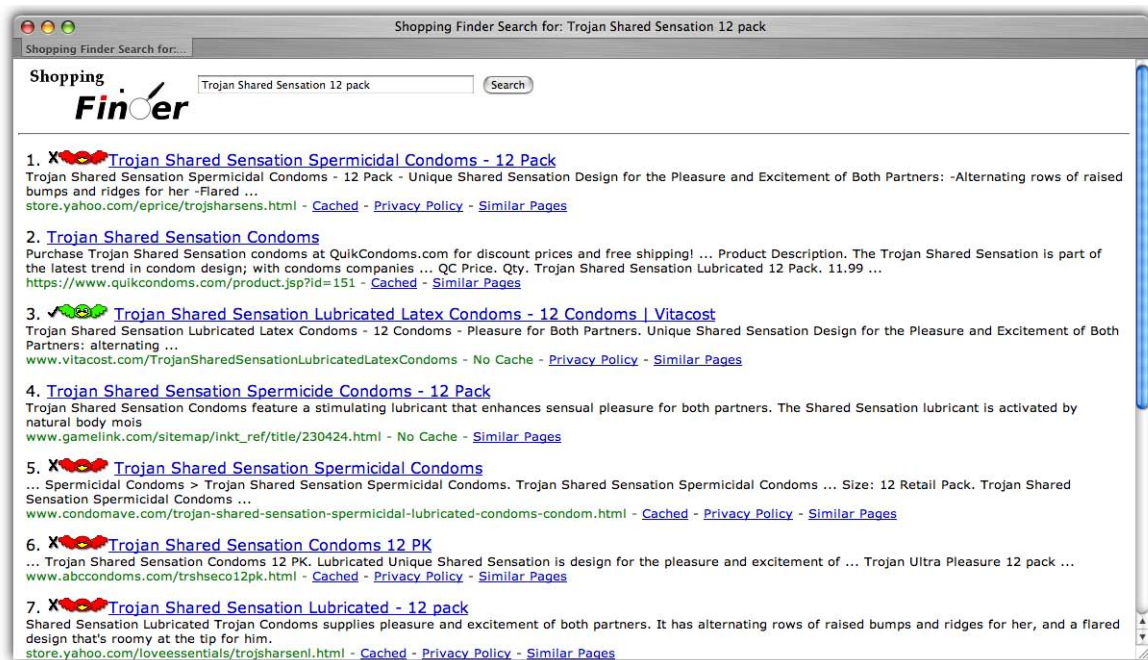


Figure 2: Shopping Finder’s search results screen.

minor typos when entering search queries. Therefore, in order to ensure that all participants viewed the same set of search results, we hard coded a set of 10 results for our two product purchase queries and displayed those results any time a user entered a search string that was the same as or similar to one of the queries we specified in our instructions to participants.

Two versions of “Shopping Finder” were prepared for our experiment. The version used by the experimental group displayed bird icons and privacy reports. The version used by the control group did not display bird icons or any privacy-related information.

3.2.3 Experimental Protocol

We conducted our experiment with one or two participants at a time. Each participant was seated in front of a laptop computer in our usability laboratory and monitored as they went through the online shopping scenario. After reviewing and signing an informed consent form, each participant was given a brief information sheet on shopping online. This was done to distract from the focus on privacy. The issues of product price, shipping prices, web site privacy policy, web site presentation, and product quality were all addressed. The experimental group was given vague information on how the Privacy Finder search engine decides what bird graphic to assign to web site results. They were told that a green bird will be associated with a web site that has a ‘good’ privacy policy, while a red bird will be associated with a web site that has a ‘bad’ privacy policy. The contents of the information sheets can be seen in Appendix A.

Each participant was then given written instructions to search for and purchase a “Universal surge protector six outlet” using their own credit card. Participants were told to compare three web sites before selecting one from which to

make their purchase. Participants were instructed to let the experimenter know when they had completed their purchase so that she could print a receipt for verification and reimbursement. In addition, participants were asked to write down the price of the chosen product and the URL of the store from which they purchased it.

After completing the surge protector purchase, participants were given similar written instructions to search for and purchase a “Trojan Shared Sensation 12 pack” using their own credit card.

The experiment was designed so that participant behavior would also be monitored through logs in order to not solely rely on self-reported information collected in the exit survey. We intended to record click stream data to verify information given by users regarding the number of web sites visited and the reported behaviors (such as reading privacy policies and privacy reports). Unfortunately, this information was not captured due to a technical glitch.

3.3 Exit survey

An exit survey was administered to participants after they completed both purchases. Participants were asked questions to determine whether factors such as previous shopping experiences, price, and web site privacy policies were taken into consideration when shopping online. The exit survey allowed participants to explain their rationale for choosing a particular web site. Those in the experimental group were also asked how Privacy Finder aided them in making their purchases.

Two different exit surveys were given, depending on which group the participant was a member of. The survey for the control group asked 21 questions. The first twelve questions dealt with the web sites from which the participant chose to make his or her purchases. This helped determine how often they shop online, whether they have purchased these sorts

Privacy Concern	Control: Surge Protector	Control: Condoms	Experimental: Surge Protector	Experimental: Condoms	Total
Confidentiality of financial information	0	1	6	5	12
Sharing of personal information with other companies	3	2	3	4	12
Unsolicited marketing	2	2	2	4	10
Confidentiality of packaging and delivery	0	3	0	0	3
Purchase history confidentiality	0	2	0	0	2
User tracking via cookies	0	0	1	1	2
Security of stored personal information	0	0	1	1	2
Confidentiality of medical information	0	0	0	1	1
Would prefer a physical store	0	0	0	1	1

Table 4: Privacy concerns mentioned in the exit survey.

of items online before (and if so, whether that influenced the decision this time), how many web sites they browsed before making each purchase, and the reason for choosing a particular web site to complete each purchase. The next four questions were with regard to privacy— how many web site privacy policies were read (and if any, why), the specific privacy concerns for each product, and whether the participant has more privacy concerns over one product than the other. The remaining five questions were designed to gather demographic information from each participant.

The survey given to the experimental group had all of the questions that were given to the control group, with the addition of eight questions regarding the Privacy Finder service (29 questions in total). These questions were designed to gather such information as whether each participant noticed the additional privacy information such as the birds or privacy reports, whether he or she understood what these features did, whether they addressed the participants’ privacy concerns, and if and how they were used when making purchasing decisions.

4. RESULTS

In this section we will examine the results of the screening survey, the exit survey, and the experiment itself. Both our screening and exit surveys indicate that information sharing and unsolicited marketing are major privacy concerns for our participants. However, we also found some more nuanced concerns that were specific to the items being purchased. We observed that those in the experimental group were willing to pay significantly more for the condoms than those in the control group. This indicates that when privacy information is made readily available, individuals may be willing to pay a premium for increased privacy protections, at least when spending someone else’s money. Similar results were seen with the surge protectors, though the average price difference was not significant.

4.1 E-Commerce Privacy Concerns

Overall, there is clear evidence that data sharing, marketing (especially telemarketing), and the ability to opt-out are the top consumer privacy concerns. We saw evidence of this in our screening survey, exit survey, and the experiment itself. Our observations confirm that the privacy information provided by Privacy Finder is relevant to users’ actual privacy concerns.

The screening survey asked participants to rate 12 web

site data practices on a 5-point Likert scale. As shown in Table 3, nine of these practices were disliked by the majority of participants. The most disliked practices were the sharing of health, financial, and identifiable information with other companies; not allowing individuals to be removed from marketing/ mailing lists; not allowing individuals to find out what information is kept on them; and telemarketing.

The exit survey asked participants to list the privacy concerns they had when making each of their purchases. As shown in Table 4, the most frequently mentioned privacy concerns across all participants were confidentiality of financial information, sharing of information with other companies, and unsolicited marketing. Data sharing and unsolicited marketing were of concern to participants in both groups when purchasing both products.

We noticed some differences in the concerns articulated by the control and experimental groups that were likely influenced by the privacy information provided by Privacy Finder. Six of the twelve participants in the experimental group mentioned that the security of their credit card information was their primary privacy concern, while this concern was absent in the control group.³ Additionally, the privacy concerns expressed by those in the experimental group were more likely to be addressed by a web site privacy policy (e.g. sharing of personal information with third parties), whereas the privacy concerns of those in the control group were less likely to be addressed by a web site privacy policy (e.g. shipping items in discreet packaging).

We also observed some differences in the types of privacy concerns associated with the two products. Concerns associated with buying condoms that were not mentioned when buy surge protectors included concerns about what would appear on their credit card statements, whether or not the company kept an order history, and whether or not the condoms would arrive in discreet packaging. A Wilcoxon Signed Rank Test across both groups showed that participants were significantly more likely to have a greater number of privacy concerns when purchasing the condoms ($p < 0.008$) than when purchasing surge protectors.

4.2 Impact of Privacy Indicators

³Privacy Finder does not actually provide information about credit card security, but half of our participants said they thought the green bird icon indicated that a site used encryption to secure credit card information.




Figure 3: Privacy preference compliance results for purchases made.

While the privacy reports help users understand a site’s full privacy policy, the first thing that a user sees when conducting a search is the colored bird indicating whether or not the site’s privacy policy complies with their preferences. It does not come as a surprise that most participants seem to have made their decisions about a site based on this feature alone. There is strong evidence that the presence of a bird (indicating a P3P policy) had an effect on purchasing decisions. Over 90% of the participants in the experimental group claimed that the presence of the bird influenced them, though it should also be noted that two of these participants bought products from sites with red birds. The presence of the bird also had a greater effect when condoms were purchased. This implies that while participants reported that price and the trustworthiness of the site were the primary decision making factors, privacy policies were taken into account when making more privacy-sensitive purchases (condoms). On the other hand, since the search for condoms yielded three times as many sites with red birds than the search for surge protectors, it is possible that this may have primed the participants.

Seven P3P-enabled sites were displayed in the condom search results. One of these sites featured a green bird, and was the third site in the list of ten search results. The remaining six P3P-enabled sites featured red birds (three sites were not P3P-enabled and thus did not feature any bird). Four P3P-enabled sites were displayed in the surge protector search results. Two of these, the third and fifth in the list of ten search results, featured green birds. Tables 5 and 6 show the product costs and privacy information for each site that appeared in the search results. The total number of purchases made at each type of site can be seen in Figure 3.

Participants who were presented with privacy policy information within the search results were more likely to make their condom purchases at a site with a “good” privacy policy than those who did not receive this information within

Site	Bird	Base	Total	Experimental	Control
1	Red	\$8.49	\$16.48	0	2
2	None	\$9.99	\$9.99	1	3
3	Green	\$9.89	\$14.88	8	2
4	None	\$16.95	\$21.90	0	0
5	Red	\$9.49	\$13.49	1	1
6	Red	\$6.40	\$11.35	0	4
7	Red	\$14.95	\$20.90	0	0
8	None	\$9.99	\$9.99	2	0
9	Red	\$8.99	\$12.99	0	0
10	Red	\$6.40	\$11.35	0	0

Table 5: List of available merchants for condom purchases. The first column lists the order that the site appeared in the search results. The next column lists the bird color (if any) for the site. The “Base” and “Total” columns list the base price for the item as well as the price including shipping, respectively. The last two columns show how many individuals from each group made purchases at the site.

Site	Bird	Base	Total	Experimental	Control
1	Red	\$9.99	\$15.99	1	1
2	None	\$9.35	\$15.85	1	0
3	Green	\$9.99	\$15.99	2	0
4	None	\$7.99	\$15.78	3	4
5	Green	\$12.50	\$14.50	2	1
6	Red	\$7.99	\$14.94	1	1
7	None	\$6.65	\$14.27	2	2
8	None	\$9.09	\$17.51	0	0
9	None	\$6.65	\$14.07	0	1
10	None	\$9.99	\$17.78	0	2

Table 6: List of available merchants for surge protector purchases. The first column lists the order that the site appeared in the search results. The next column lists the bird color (if any) for the site. The “Base” and “Total” columns list the base price for the item as well as the price including shipping, respectively. The last two columns show how many individuals from each group made purchases at the site.

the search results. Participants in the control group who expressed privacy concerns did not express those concerns through their actions. In the experimental group, eight participants purchased the condoms from the single green bird site. Three participants made their purchases from less expensive sites that did not feature a colored bird, and one participant purchased from one of the less expensive red bird sites. This stood in contrast with the control group (where no birds were actually displayed) as only two individuals purchased the condoms from the green bird site. Two subjects in the control group made their purchases at the first site listed, a red bird site that was more expensive than the green bird site. The remaining eight made their condom purchases at cheaper sites with red birds or no birds. A chi-square test indicated that these differences were highly significant ($p < 0.025$).

Privacy policy indicators also had an impact on surge protector purchases, but to a lesser extent than they did for condom purchases. More participants in the experimental group than in the control group purchased surge protectors from green bird sites, but a chi-square test did not yield significant results. In the experimental group, four participants made their surge protector purchases from one of the two green bird sites, while two purchased from the cheaper red bird sites. Six of the participants in the experimental group purchased from sites that did not display a bird (and therefore did not have a P3P policy). In the control group, only one participant purchased a surge protector from a green bird site. Nine participants purchased from cheaper sites, while two made purchases from more expensive sites (one participant did not do any comparison shopping and the other one reported that the site design was more professional looking and therefore “less risky”). This indicates that while some participants took the bird indicators into account, they did not have the effect that they did when making the condom purchases. Thus, privacy was not as much of a concern for the surge protector purchases.

Additional evidence that participants took privacy into account more when making condom purchases can be found by looking at the behavior of individual participants across their two purchases. In the experimental group, there were five cases where the participant purchased the surge protector at a site with either a red bird or no bird and then made their condom purchase at a site that had a green bird. However not all participants behaved this way: there were two participants who switched from a site with a green bird for the surge protector purchase to a site with either no bird or a site with a red bird for the condom purchase.

On the exit survey, participants in the experimental group were twice as likely as those in the control group to report that privacy policies influenced their purchasing decisions. Privacy was a deciding factor for eight members of the experimental group and three members of the control group when purchasing condoms, and for seven members of the experimental group and three members of the control group when purchasing a surge protector. However, price was still the primary deciding factor across both groups. In the control group, 11 participants said that price was one of the deciding factors when purchasing both the condoms and the surge protectors. In the experimental group, 10 participants said price was a deciding factor for the surge protector, and nine said it was a deciding factor for the condoms.

While price was the primary decision making factor, pri-

vacuity played an important role. The average purchase price for condoms in the experimental group was \$9.88, without shipping. The average purchase price for the condoms in the control group was \$8.49. This would imply that the participants were willing to pay slightly more for increased privacy protections. Since the prices of the items were not normally distributed, we performed a Wilcoxon Mann-Whitney Test and found the mean price differences to be marginally significant as $p = 0.088$. When factoring in shipping, the prices were \$13.96 and \$12.63, respectively. This difference in means, though, was only significant at $p = 0.248$. However, since participants were being reimbursed, these statistics only show that participants were willing to pay a premium for privacy when it was someone else’s money. This same effect can be seen with the surge protector purchases as those in the experimental group paid \$17.04 on average, while those in the control group paid \$16.47 on average, though this difference was not statistically significant.

4.3 Communicating About Privacy

As already discussed, the privacy information presented by Privacy Finder appears to have influenced participants’ purchasing decisions as well as the types of privacy concerns they articulated in our exit survey. However, participants’ exit survey responses suggest that Privacy Finder did not always communicate the intended messages clearly.

When asked what the green bird represents, six of the participants said that it means the site keeps financial information secure through the use of encryption. Had participants read the privacy reports we believe it would have been apparent to them that this is not what the green bird indicates; however, only four of the twelve participants read the privacy reports. Four participants said they did not know where to find the privacy reports, three said they were not interested enough to read them, and one did not specify a reason for not reading them. In any case, further studies are needed to determine the extent to which Privacy Finder is providing users with useful privacy policy information as well ways of making the information more easily accessible.

While it was clear that participants had privacy concerns, it is not clear that they were making any extra efforts to learn about web site privacy policies. Only a third of our participants claimed to have read web site privacy policies while making purchase decisions during our experiment. Two of the experimental participants mentioned that they read the privacy reports but not the web sites’ full privacy policies because they trusted the information provided by Privacy Finder and did not see a need to read further. For these participants, Privacy Finder may be doing exactly what it is meant to do.

When asked how the bird indicators helped them make a purchasing decision, five participants said that they avoided sites with the red birds entirely. This implies a higher level of trust for web sites that did not choose to disclose a P3P policy. If this is truly the case, then users are making a poor assumption. When no bird is displayed, it simply means that no privacy information is available for the site – this is not indicative of a favorable privacy policy. In fact, when a red bird is displayed, information about the site’s privacy policy is conveyed to the user, whereas when there is no bird, the worst case scenario about the site’s privacy policies should be assumed (unless the user goes through the effort of reading the web site’s human readable privacy policy).

5. LIMITATIONS AND FUTURE WORK

This study was a useful first effort to assess the effects of displaying privacy indicators in search results. However, there is much more work to do in this area. In this section we discuss some of the limitations of this study and our plans for future work that will address these limitations and explore some new directions.

Control over search results. Performing this study using real web sites and slightly-modified search engine results made for a more realistic experiment than we could have conducted using only simulated web sites and search results, and allowed our participants to make real purchases in which they actually faced a potential privacy risk. We did modify the search results somewhat because we needed to make sure all participants would see the same results and to assure a good distribution of P3P policies among the top 10 results. Future studies should more carefully control the search results presented so that there are fewer variables that might impact purchase decisions. For example, while the condom search presented a choice between web sites offering identical products, the surge protector search presented multiple brands of surge protectors with varying features. It would have been better to present a set of search results featuring an identical set of products. Other variables that might be better controlled for in the future include the perceived trustworthiness of the web sites in the search results (which is influenced largely by how well known each site is and how professional it looks), the range of prices offered (with and without shipping fees), the order in which sites appear in the search results, and the number of sites with “good” and “bad” privacy policies in the search results. This would make it easier to isolate the effects of privacy information from other variables and allow for a better comparison between privacy-sensitive and privacy-insensitive purchases.

Information participants looked at when making purchase decisions. Although we had planned to log which links our participants clicked on in the search results, including when they clicked on the privacy report links, this information did not get logged due to a bug introduced as a result of a last minute change to our experimental system. Thus, we have only self-reported data on how many sites each participant visited before making a purchase. Future studies should not only log this information, but also direct all web traffic through a proxy and record all of the participants’ clicks at the web sites they visit. This will enable us to determine whether or not participants checked a site’s shipping costs, reviewed a site’s privacy policy, or looked at other information that might have influenced their purchase decisions.

Misleading privacy indicators. Some of our participants appeared to assume that a site with a red bird icon was worse than a site with no privacy icon (indicating an unknown privacy policy). In addition, those who did not mouse over the red bird or read the privacy report may have considered all sites receiving a red bird as equally bad. In the latest version of Privacy Finder we have attempted to address these problems by eliminating the red and green bird icons. Instead, we have adopted a scoring system and use a set of four filled or empty squares to indicate a “privacy level.” No squares are displayed next to a site that does not have a P3P policy. A site that fully complies with the user’s stated preferences will have all four squares filled in. Sites that conflict with the user’s preferences have a proportion-

ate number of squares filled in based on the degree of the conflict. When calculating the degree of the conflict we use a scoring system that weights some conflicts more than others based on our research into which privacy issues tend to raise the most concerns with Internet users. In addition, to address the problem that one third of our participants were unaware that they could click on the bird icon to retrieve the privacy report, our new design includes an explicit “privacy report” link beneath each set of squares. More work is needed to test whether this new scoring system and associated icons is less misleading and more meaningful to users. It is believed that this system would also be more accessible to users who are colorblind as they no longer would have to make a distinction between red and green birds. Unfortunately, the presence of colorblind participants was not examined (though when given instructions about examining red and green icons, one would expect that someone who could not tell the difference between the two would have said something). Furthermore, it would be useful to test whether the mere presence of positive indicators influences purchasing decisions, even if participants are not told what the indicators mean and the indicators are not accompanied by privacy reports or other privacy-related information.

Priming. Priming might have been an issue in this experiment. The section of the instructions that discussed privacy policies (Appendix A) was longer and more in-depth for the experimental group than for the control group. This had the potential to inadvertently increase participant awareness to privacy such that they took privacy considerations into account more than they normally would when shopping online in their natural environment. In future studies we also plan to randomize for each participant the order in which they are requested to purchase the non-privacy sensitive and the privacy sensitive goods. Along these same lines, future studies might ask participants to perform other searches to increase familiarity with the Privacy Finder service before conducting the purchasing tasks. This would help participants focus more on the idea that they are testing a new search engine and less on privacy or online purchasing. A separate study involving the use of Privacy Finder by participants on their own computers over an extended period of time would provide complementary data that would offer insights into the use of privacy information in a natural environment.

Participants. Because some of the effects may be small and nuanced, a larger number of participants is needed to produce more significant results. Furthermore, in order to reach more generalizable conclusions future studies should not limit participants to college students.

Price sensitivity and privacy/price tradeoffs. The present study fully reimbursed participants for their purchases and thus did not provide an opportunity to test the extent to which participants were willing to trade off higher prices for greater privacy when using their own money. Furthermore, because privacy information was displayed in the search results but price information was not, a different level of effort was required for gathering price and privacy information. To address these issues we are developing a version of Privacy Finder that searches the Yahoo! Shopping Network and annotates results with both privacy information and price information. Future studies might use this version of Privacy Finder and pay participants a fixed participation fee rather than reimbursing them for the products purchased. This would provide an incentive for participants

to purchase lower priced items so that they could keep more of the money.

6. CONCLUSION

Privacy is a major concern for Internet users, but it is difficult for individuals to obtain enough information about web site privacy policies to take privacy into consideration when making purchasing decisions. Reading and understanding web site privacy policies is difficult and time consuming, and identifying web sites with acceptable policies can be extremely difficult. To make this process easier, we have developed a search engine that annotates search results with privacy information and presents privacy reports for each site in a standard format.

We conducted a first set of experiments aimed at determining the extent to which privacy information provided by a search engine influences online purchase decisions. Our results suggest that when privacy information is made readily available, many users will take it into account when making purchase decisions that require them to expose their credit card information and other personal information but do not require them to spend their own money. Furthermore, our results indicate that the type of product being purchased may also impact users' concerns about privacy and their interest in using privacy information when choosing a vendor. Future work is needed to find ways of presenting privacy information more clearly and additional studies are needed to understand the tradeoffs people make between privacy and price in purchase decisions.

7. ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under grant IGERT 9972762 in CASOS, and by the Pennsylvania Cyber Security Commercialization Initiative (PaCSCI). Additional support was provided by the Center for Computational Analysis of Social and Organizational Systems (CASOS), the Institute for Software Research International (ISRI), and CyLab at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation or the U.S. government.

The authors would also like to acknowledge AT&T for the development and release of the Privacy Bird source code, on which the code used for this project is based. The previous prototype for the Privacy Finder service was written by Simon Byers, David Kormann, and Patrick McDaniel while at AT&T Labs-Research.

Finally, we would like to acknowledge Janice Tsai for her input throughout the design of this study as well as the writing of this paper.

8. REFERENCES

- [1] A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the ACM Electronic Commerce Conference (EC 04)*, pages 21–29, New York, NY, 2004. ACM Press. <http://www.heinz.cmu.edu/acquisti/papers/privacy-gratification.pdf>.
- [2] A. Acquisti and J. Grossklags. Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. In *Proceedings of The 2nd Annual Workshop on Economics and Information Security (WEIS '03)*, 2003.
- [3] D. Allen. The great online privacy debate, 2000. http://www.ebusinessforum.com/index.asp?doc_id=1785&layout=rich_story.
- [4] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, and J. Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, April 2002. <http://www.w3.org/TR/P3P/>.
- [5] L. F. Cranor, J. Reagle, and M. S. Ackerman. Beyond concern: Understanding net users' attitudes about online privacy. *AT&T Labs-Research Technical Report TR 99.4.3*, 14 April 1999. <http://www.research.att.com/resources/trs/TRs/99/99.4/99.4.3/report.htm>.
- [6] M. J. Culnan. How privacy notices promote informed consumer choice, 2002. <http://www.cdt.org/privacy/ccp/notice1.pdf>.
- [7] M. J. Culnan and G. R. Milne. The culnan-milne survey on consumers and online privacy notices, 2001. http://intra.som.umass.edu/georgemilne/pdf_files/culnan-milne.pdf.
- [8] B. Fogg, J. Marshall, O. Laraki, A. Osipovich, C. Varma, N. Fang, J. Paul, A. Rangekar, J. Shon, P. Swani, and M. Treinen. What Makes Web Sites Credible? A Report on a Large Quantitative Study. In *Proceedings of the ACM Computer-Human Interaction Conference*, Seattle, WA, March 31 - April 4, 2001. ACM.
- [9] S. Fox, L. Rainie, J. Horrigan, A. Lenhart, T. Spooner, and C. Carter. Trust and privacy online: Why Americans want to rewrite the rules. August 20, 2000. http://www.pewinternet.org/pdfs/PIP_Trust_Privacy_Report.pdf.
- [10] Harris Interactive. First major post-9/11 privacy survey finds consumers demanding companies do more to protect privacy: public wants company privacy policies to be independently verified, 2002. <http://www.harrisinteractive.com/news/allnewsbydate.asp?NewsID=429>.
- [11] Harris Interactive. Identity Theft: New Survey & Trend Report, August 2003.
- [12] B. Huberman, E. Adar, and L. Fine. Valuating privacy. In *Proceedings of The Workshop on The Economics of Information Security*, Boston, MA, June 1-3, 2005.
- [13] C. Jensen and C. Potts. Privacy policies as decision-making tools: An evaluation of online privacy notices. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 471–478, Vienna, Austria, 2004.
- [14] Jupiter Research. Seventy percent of us consumers worry about online privacy, but few take protective action, 2002. http://www.jmm.com/xp/jmm/press/2002/pr_060302.xml.
- [15] T. Moores. Do consumers understand the role of privacy seals in e-commerce? *Communications of the ACM*, 48(3):86–91, 2005.
- [16] R. F. Murphy. Social distance and veil. *American Anthropologist*, 66(6):1257–1274, 1964.

- [17] E. Rose. Data users versus data subjects: Are consumers willing to pay for property rights to personal information? In *Proceedings of the 38th Hawaii International Conference on System Sciences*, 2005.
- [18] A. Shostack. Paying for privacy: Consumers and infrastructures. In *Proceedings of The 2nd Annual Workshop on Economics and Information Security (WEIS '03)*, 2003.
- [19] M. D. Smith and E. Brynjolfsson. Consumer Decision-Making at an Internet Shopbot. Technical Report 4206-01, MIT Sloan School of Management, October 2001. <http://ssrn.com/abstract=290334>.
- [20] S. Spiekermann, J. Grossklags, and B. Berendt. E-Privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. In *Proceedings of EC'01: Third ACM Conference on Electronic Commerce*, pages 38–47, Tampa, Florida, 2001. http://www.sims.berkeley.edu/jensg/research/eprivacy_acm.html.
- [21] P. Syverson. The paradoxical value of privacy. In *Proceedings of The 2nd Annual Workshop on Economics and Information Security (WEIS '03)*, 2003.
- [22] H. Taylor. Most People are “Privacy Pragmatists” Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits. 17, 2003. http://www.harrisinteractive.com/harris_poll/index.asp?PID=365.
- [23] B. Tedeschi. Everybody talks about online privacy, but few do anything about it. *The New York Times*, page C6, June 3, 2002.
- [24] J. Turow. Americans and online privacy: The system is broken, 2003. <http://www.asc.upenn.edu/usr/jturow/internet-privacy-report/36-page-turow-version-9.pdf>.
- [25] C. Varney. You Call This Self-Regulation? *Wired News*, June 9, 1998.
- [26] A. F. Westin. Harris-equifax consumer privacy survey (1991). Technical report, Equifax, Inc., Atlanta, GA, 1991.

APPENDIX

A. INFORMATION SHEET

Contents of the information sheet given to participants:

Points to Consider When Shopping Online

- **Price**
The same or similar products are often available at different web sites for different prices.
- **Privacy Policy (Control Group)**
Many web sites have privacy policies that describe the types of personal information the site collects and how they will use it.
- **Privacy Policy (Experimental Group)**
Many web sites have privacy policies that describe the types of personal information the site collects and how they will use it. The Shopping Finder search engine displays color coded pictures of birds in the search results to indicate the quality of a web sites privacy policy. A red bird signifies that the web site has a poor privacy policy, while a green bird indicates that the web site has a good privacy policy. If the search engine is unable to interpret a sites privacy policy it does not display any bird for that site. Users can click on a bird for more information about a sites privacy policy.
- **Product Quality**
Product descriptions, user reviews and brand names are information that can be used to assess the product quality.
- **Shipping Fees**
Shipping fees can increase the price of a product. The base price of a product can be deceiving when shipping fees are high.
- **Site Appearance or Presentation**
The appearance of a web site can be an indicator of a companys business practices.