Timing Is Everything? The Effects of Timing and Placement of Online Privacy Indicators

Serge Egelman, Janice Tsai, Lorrie Faith Cranor, Alessandro Acquisti
Carnegie Mellon University
Pittsburgh, PA
egelman@cs.cmu.edu, jytsai@andrew.cmu.edu, lorrie@cs.cmu.edu, acquisti@andrew.cmu.edu

ABSTRACT

Many commerce websites post privacy policies to address Internet shoppers' privacy concerns. However, few users read or understand them. Iconic privacy indicators may make privacy policies more accessible and easier for users to understand: in this paper, we examine whether the timing and placement of online privacy indicators impact Internet users' browsing and purchasing decisions. We conducted a laboratory study where we controlled the placement of privacy information, the timing of its appearance, the privacy level of each website, and the price and items being purchased. We found that the timing of privacy information had a significant impact on how much of a premium users were willing to pay for privacy. We also found that timing had less impact when users were willing to examine multiple websites. Finally, we found that users paid more attention to privacy indicators when purchasing privacy-sensitive items than when purchasing items that raised minimal privacy concerns.

Author Keywords

Privacy, privacy policies, website indicators, mental models, usable privacy and security

ACM Classification Keywords

H.1.2 User/Machine Systems, H.5.2 User Interfaces, D.4.6 Security and Protection

INTRODUCTION

Privacy is often cited as a primary concern among Internet users [1]. In response, many corporations have posted privacy policies [15]. However, these policies rarely help consumers because they often go unread [17], or do not address the most common consumer concerns [8, 19]. Even worse, privacy policies are difficult to understand. Anton et al. examined forty bank privacy policies and found that on average, a college education was needed to comprehend them [3]. A 2008 survey found that several years of graduate school are required to understand the privacy policies of the top Internet companies [20].

CHI 2009, April 3 - 9, 2009, Boston, MA, USA.

Indicators can be used to distill privacy policy information into intuitive icons. However, studies have shown that privacy and security indicators can fail users when they go unnoticed, when they force users to take extra steps to complete a task, or when other environmental stimuli outweigh the strength of the indicators [25, 12, 11]. Previous studies have shown that users may be willing to pay a premium to know when they are visiting a high privacy website [23]. But there is still an open question of *how* to effectively convey website privacy information.

We performed a laboratory study where we tightly controlled the price of two items offered by several online vendors such that participants would have to pay more money to purchase the items from vendors with better privacy policies—a privacy premium. We selected the privacy premiums based on the results of an online survey designed to determine the maximum amount online shoppers would be willing to pay for increased privacy. A total of 89 participants came to our laboratory and purchased two items using their own credit cards and providing their personal billing information: one item was selected to elicit heightened privacy concerns, and one was not. We created four conditions that used different privacy icons to annotate the websites such that we varied both *when* and *how* the icons were displayed.

Our results demonstrate, first, that many online shoppers will go to extra efforts to purchase from high privacy websites when privacy indicators are available. Second, we show that online shoppers who are less privacy-motivated will pay significantly more for privacy when privacy indicators are presented to them before visiting websites, rather than after they arrive at a website. Third, we demonstrate that online shoppers are more likely to take privacy indicators into account when purchasing privacy-sensitive items.

In the next section, we present related work on usability problems with privacy and security indicators, and related studies on privacy and purchasing behaviors. Next, we discuss our privacy premium survey and the methodology behind our laboratory study. We then present our results in terms of how the timing of the indicators impacted user behaviors, and how privacy decisions were related to the type of item being purchased. Finally, we conclude with limitations and future work.

BACKGROUND

Privacy indicators attempt to turn privacy policy information into intuitive icons. Unfortunately, current indicator designs are not very effective for a variety of reasons. To date, there have been few studies on optimal indicator designs. In this section we highlight previous studies of privacy and security indicators, we introduce the Platform for Privacy Preferences (P3P), we give an overview of our privacy-enhanced search engine, Privacy Finder, and we describe some related studies we have performed using Privacy Finder.

Online Privacy and Security Indicators

Many companies post "privacy seals" on their websites in an attempt to improve consumer confidence. Adkinson et al. estimated privacy seal adoption at 11% in 2001 [2], while Jensen et al. estimated privacy seal adoption at around 2% in 2006 [16]. For FY2007, TRUSTe claimed 2,241 participating websites worldwide, including 31 Fortune 500 participants [22]. Assuming privacy seals are pervasive enough to be recognized by consumers, do consumers properly understand what they represent?

Many Internet users erroneously believe that websites with seals have adopted consumer-friendly privacy practices. However, the presence of a privacy seal says little about the content of a company's privacy policy [18]. In fact, Edelman conducted a study of websites brandishing the TRUSTe privacy seal in 2006 and concluded that "sites that seek and obtain trust certifications are actually significantly less trustworthy than those that forego certification" [9].

If trustworthy privacy seals do exist, it is unlikely that users recognize them. In a study conducted in 2005, 15% of participants claimed to recognize an authentic-looking privacy seal created solely for the purpose of the study. At the same time, the legitimate privacy seals were recognized by only 26% of the participants, on average [18].

Online security and privacy indicators also fail when users do not notice them. In a usability study of web browser security indicators, Wu et al. found that 85% of participants evaluated the content of a website when making a trust decision, often incorrectly trusting the content more than the indicator [25]. This corroborates Fogg et al.'s finding that the "look and feel" of a website is often the strongest factor behind users' trust decisions [12]. Wu et al. also observed that 25% of participants failed to notice the security indicators at all. Studies conducted on previous SSL indicators and new Extended Validation (EV) SSL indicators have made similar discoveries: when not primed for security, users do not look for security indicators in the browser chrome [24, 21], perhaps because the user's locus of attention is on the website content. Thus, placing privacy and security indicators near a user's locus of attention will likely increase efficacy.

The Platform for Privacy Preferences (P3P)

The W3C's Platform for Privacy Preferences (P3P) was created to help users understand website privacy policies. P3P specifies a standard set of XML elements that can be used to construct machine-readable privacy policies. These policies can be posted on websites and then analyzed by user agents on behalf of Internet users. If a user agent encounters a privacy policy that does not conform to a user's stated privacy preferences, the user agent can take actions on behalf of the user by displaying a warning, rejecting cookies, or blocking the website entirely [5]. Byers et al. found that by 2003, P3P had already been adopted by over 30% of the most popular websites and 10% of their entire sample [4]. In 2005, Egelman et al. reexamined this sample and found that P3P adoption had increased by over 30%. They also found that on average, 32% of all Google queries yield at least one P3P-enabled search result [10]. In 2006, Jensen et al. compiled a sample of over 26,000 websites from around the world and used it to estimate P3P adoption at 25% [16]. The increasing rate of P3P adoption is beneficial to consumers because it facilitates the automatic dissemination of website privacy information; tools can be developed to distill privacy policies into simple indicators automatically.

In 1999, AT&T began developing their Privacy Bird P3P user agent for Internet Explorer. Privacy Bird displays a colored bird icon in the corner of the web browser to indicate whether a policy matches the user's stated privacy preferences. A red bird indicates a conflict with the user's preferences, while a green bird indicates a compliant policy. Cranor et al. conducted a survey of 309 Privacy Bird users and found that a common complaint was that privacy information was not displayed on many websites. They concluded Privacy Bird was still useful since 88% of the respondents said that being aware of website privacy policies caused them to alter their behaviors. Many claimed that they stopped visiting certain websites, sought opt-out information, and compared websites based on privacy policies [7]. However, a shortcoming of Privacy Bird is that to view a website's privacy information, users must first transmit certain clickstream data to visit that website. This also means that to compare the privacy policies of n different websites, a user must visit all n websites before making a decision. It is unclear whether or not a user will go through this process until he or she finds a satisfactory privacy policy.

Privacy Finder

Cranor et al. developed a prototype P3P-enabled search engine that allowed users to enter a set of search terms and retrieve a list of results annotated with red or green birds indicating whether or not each result complies with the user's stated privacy preferences [6]. Egelman et al. improved this search engine, named it Privacy Finder, and made it publicly available. One of the improvements was the addition of "privacy reports." Users of Privacy Finder can click on the privacy indicators to generate a summarized version of a website's privacy policy highlighting any conflicts it may have with the user's privacy preferences [10].[1]

Privacy Indicator Purchasing Studies

Gideon et al. conducted a user study of Privacy Finder in 2006. Participants were instructed to purchase a privacy-sensitive item—condoms—and a common household item—

[1] http://www.privacyfinder.org/

power strips. The search results for each product were preselected so that at least one green bird icon appeared along with several red bird icons. When purchasing the privacy-sensitive item, participants were more likely to purchase from websites with a positive indicator [13].

We performed a follow-up to Gideon et al.'s study in 2007. To determine whether participants cared about privacy or were just attracted to the indicators, we added a second control condition where the same indicators were labeled as representing irrelevant information (handicap accessibility) rather than privacy. We also changed the privacy indicators from red and green birds to a set of four boxes: the number of boxes colored green was inversely proportional to the number of conflicts with the user's privacy preferences; four green boxes indicated a privacy policy completely matched a user's privacy preferences. We removed the indicator from the website with the lowest price to test the effect of encountering an unknown privacy rating. We conducted an online survey to identify products that would raise participants' privacy sensitivities, but would unlikely result in participants dropping out of the study if asked to purchase them. We chose a vibrating sex toy as the privacy-sensitive item and a pack of AA batteries as an item that by itself would be unlikely to raise privacy concerns. Participants used their own personal information for the purchases and therefore may have had privacy concerns related to their information, regardless of the type of items they were purchasing [23].

We observed that participants paid a premium to buy both products from a website with a privacy indicator. However, we did not control the exact amount of the premium or keep it constant between the two products (i.e. we were unable to test the interaction between price and privacy sensitivities). Thus, it is unclear whether participants would have paid the same premium for the two products. We did not test whether participants would pay a privacy premium when the cheapest website had the worst privacy policy (rather than no privacy indicator). Finally, we never examined how alternate methods of displaying privacy indicators impacted purchasing decisions. Several of the other studies we have cited show how (not) to display indicators in browser chrome [25, 11, 24, 21], but few studies have offered methods for displaying privacy indicators alongside website content. One study examined the role of timing when displaying software license agreements [14], but we are unaware of previous studies that have examined the role of timing on privacy indicators. Thus, this paper focuses on the timing and placement of privacy indicators.

PRIVACY PREMIUM SURVEY

Before our experiment, we conducted an online survey to estimate the maximum premium people would be willing to pay to purchase from a website with a high privacy level. We recruited 676 Internet users through Craigslist and sweepstakes websites in June 2008. The survey contained five pages of Privacy Finder screenshots (Figure 1). Each screenshot depicted four search results for identical products with identical descriptions. The search results only differed based on the privacy indicator placed to their left and the price information placed to their right. Both the price and privacy level increased with each subsequent search result. Thus, the websites with the highest privacy ratings also had the highest prices. We assigned half the respondents to a between-group condition in which the cheapest website had no privacy indicator and the other half to a condition in which the cheapest website had the lowest privacy level. The product displayed in the search results alternated between the sex toy and pack of batteries that laboratory participants would be purchasing, with the order randomly selected. Respondents were given the following instructions:

> "Given only the information displayed in the search results, from which web site would you be most likely to make this purchase? Please choose one site, even if this is not a product you would be likely to ever purchase."

Respondents were exposed to two of three possible premiums for the highest privacy—denoted by four green boxes: \$0.25, \$0.75, and \$1.50. The premiums and associated privacy indicators are shown in Table 1. The privacy premiums were randomly assigned so that respondents saw the same premium for the first two pages (i.e. respondents saw the same premium for both products). The third page of the survey contained a control where one of the two products was randomly displayed with identical prices for each of the four search results. The privacy indicators varied so that we could examine whether participants would select the website with the highest privacy level in the absence of a premium.

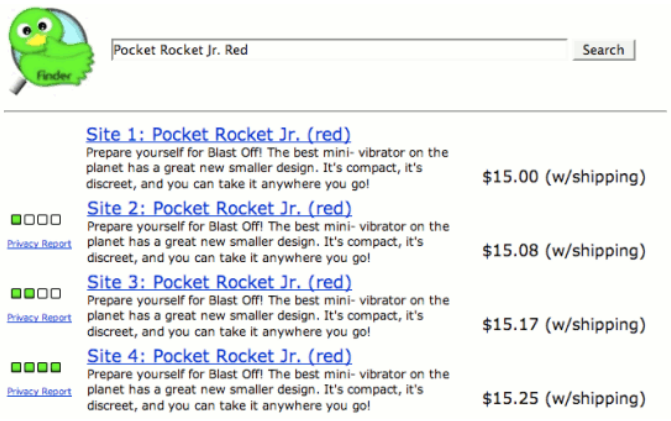

Figure 1. Example screenshot used in the privacy premium survey.

Indicator	Premium 1	Premium 2	Premium 3
☐☐☐☐	\$15.00	\$15.00	\$15.00
■☐☐☐	\$15.08	\$15.25	\$15.50
■■☐☐	\$15.17	\$15.50	\$16.00
■■■■	\$15.25	\$15.75	\$16.50

Table 1. The privacy premiums and associated privacy indicators used in the survey. The privacy indicator for the cheapest website was only displayed to half of the respondents.

The fourth and fifth pages followed the same protocol as the first and second pages, but participants were randomly assigned one of the two privacy premiums they had not already seen. However, we decided not to include these results in the analysis since we found evidence that participants' willingness to pay the subsequent premiums was highly dependent on the first premium to which they were exposed.

We combined the two between-group conditions for the analysis when we discovered that the only difference occurred when respondents encountered the highest privacy premium: those selecting the batteries were significantly more likely to select the first website—the cheapest one—when the indicator was absent ($t_{239} = 2.175$, $p < 0.031$).[2]

The ideal privacy premium for our laboratory study is the highest one that survey respondents would be willing to pay for both products; the survey responses likely provided an upper bound because the respondents reported how much they *would* pay without actually having to pay that amount. Using ANOVA to compare the three privacy premiums for each of the two products we found no significant differences between the three premiums when respondents considered the sex toy: most respondents indicated they were willing to pay any privacy premium presented to them. However, when the privacy premium was $1.50, respondents were more likely to purchase the batteries from cheaper vendors, and therefore unwilling to pay a premium for privacy ($F_{2,673} = 6.251$, $p < 0.002$). At the same time, respondents indicated they were still willing to spend $0.25 and $0.75 for increased privacy when purchasing the batteries. We concluded a privacy premium of $1.50 may be too high for our laboratory experiment.

A pairwise t-test confirmed that a $0.75 privacy premium would still allow us to observe differences between the two products. Respondents indicated they were willing to spend significantly more money for the sex toy—in exchange for greater privacy—than for the batteries ($t_{214} = 5.226$, $p < 0.0005$). We concluded that a $0.75 privacy premium would be low enough that laboratory participants would consider paying it for both products, while still allowing us to observe differences in behavior between the two product purchases.

METHODOLOGY

Our primary goal for this study was to examine whether the placement and timing of privacy indicators impacts purchasing decisions. In order to quantify differences in purchasing behaviors, we created a controlled privacy premium: participants who wanted a higher degree of privacy would have to pay a fixed amount for it. We also wanted to determine whether participants' behaviors would differ when purchasing a product that did not raise additional privacy concerns compared to a product that did. We designed the laboratory experiment to test the following hypotheses:

1. Participants will pay for increased privacy when they see privacy indicators.
2. Participants who see privacy indicators will pay more for the privacy-sensitive item than the item that does not raise additional privacy concerns.
3. Participants will be more likely to pay for increased privacy when they see privacy indicators alongside search results before visiting a website than when they see privacy indicators after clicking on search result links.
4. Participants will be more likely to pay for increased privacy when they see privacy indicators before they see the content of a website than when they see privacy indicators alongside the content of a website.
5. Participants who see privacy indicators after clicking on search result links will visit more websites than those who see privacy indicators alongside search results.

Study Design

We conducted a laboratory experiment during the summer of 2008 using participants from the Pittsburgh area. We recruited 89 participants using Craigslist and flyers on bus stops, telephone poles, and community bulletin boards. We used a screening survey to gather basic demographic data and to assess privacy concerns related to using the Internet and online shopping. Because the privacy indicators we tested were designed for use by individuals who have privacy concerns when shopping online, we used the same screening survey and screening methodology used in our previous study to screen out those who perceived little or no privacy risk when shopping online [23]. Based on this requirement, we screened out 16.39% (50 of 305) responses.

We chose a specific vibrating sex toy, the "Pocket Rocket Jr.," as the privacy-sensitive item. We instructed participants to purchase the red version so that our results would not be confounded by the availability of differing colors from different vendors. We chose an 8-pack of Duracell AA batteries as the item unlikely to raise additional privacy concerns beyond the act of providing personal information to an online vendor. We tightly controlled the price of each item by collaborating with four office supply vendors and four sex toy vendors who had varying privacy policies.[3] We asked the vendors to set specific prices based on their privacy policies and the results of our privacy premium survey.[4] Privacy Finder returned static results pages when specific search strings (or variants thereof) were submitted: "Pocket Rocket Jr. Red" and "Duracell AA 8-pack." Each of these two pages of search results contained five hits with varying prices and privacy ratings, as seen in Table 2. In both sets of search results we also included a fifth search result that did

[2]For a privacy premium of $1.50, users may purchase from a website with an unknown privacy policy (i.e. the cheapest website) if the item being purchased does not raise privacy concerns.

[3]We contacted over twenty vendors for each product until four vendors for each product agreed to participate. For the vendors who lowered their prices, we compensated them for the difference. We only contacted vendors who participants were likely unfamiliar with; a full list of the vendors appears in the Acknowledgements.

[4]We used a privacy premium of $0.75 based on the results of the survey. Due to vendor constraints we had to set the base price at $15.50 rather than the $15.00 we used in the premium survey.

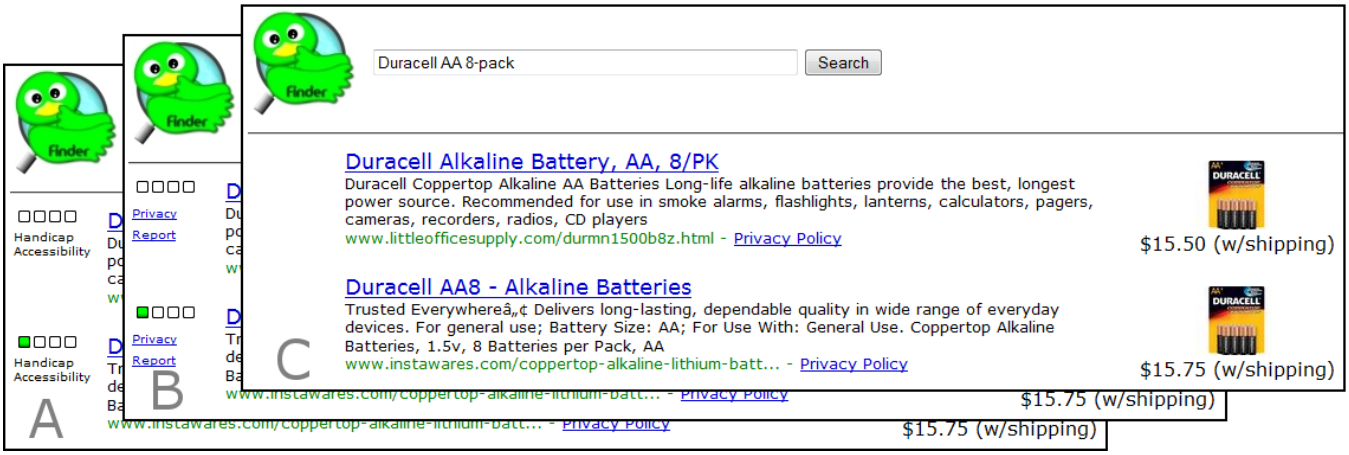

Figure 2. Screenshot of the search results for the four study conditions: (A) participants in the *handicap* condition saw the handicap accessibility indicators; (B) participants in the *privacy* condition saw the privacy indicators; and (C) participants in the *frame* and *interstitial* conditions did not have annotated search results.

Hit #	Indicator	Price
1	☐☐☐☐	$15.50
2	■☐☐☐	$15.75
3	■■☐☐	$16.00
4	■■■■	$16.25
5		$16.75+

Table 2. The prices and privacy ratings for both sets of search results, the batteries and the sex toy. Participants who wanted the highest level of privacy had to pay an additional $0.75 for each product.

not have a privacy rating. This website had the highest price of the five and was included because we were curious if any participants would pay more than the $0.75 privacy premium to buy from a website with an unknown privacy policy, and whether they would understand that the lack of any indicator corresponds to an unknown privacy policy.[5]

We randomly assigned participants to one of three experimental conditions or the control condition, balancing the gender of participants in each condition:

- Handicap Accessibility (control): Participants were shown annotated search results (Figure 2A). However, we labeled the privacy indicators as "handicap accessibility" so that the indicators were not associated with privacy. The links to the privacy reports (i.e. the machine-generated privacy policy summaries) were removed.[6] We used this condition to examine whether participants in the other conditions were genuinely thinking about privacy or whether they were choosing websites simply based on the presence of irrelevant green indicators.
- Privacy (experimental): Participants were shown annotated search results with privacy indicators (Figure 2B).
- Frame (experimental): Participants were shown search results that were not annotated (Figure 2C). Once a participant visited a website from the search results, a frame appeared at the top of the website that displayed the privacy indicator and a link to the privacy report (Figure 3). We created this condition to simulate the Privacy Bird experience: users who wanted to comparison shop based on privacy indicators would have to visit a website in order to see its privacy rating. We hypothesized that users would find this tedious and therefore make poor privacy choices, especially when purchasing the batteries since they would likely be less motivated to protect their privacy.
- Interstitial (experimental): Participants were shown search results that were not annotated (Figure 2C). Once a participant visited a website from the search results, they saw an interstitial—a full screen message—with the privacy indicator (Figure 4). We created the interstitial condition to examine whether the content of a website detracted from the privacy indicator. We wanted to control for users being able to view website content alongside the privacy indicator in the *frame* condition. We hypothesized that users would choose higher privacy in this condition because they would be making the decision solely based on the privacy indicator.

We found no significant differences between the average ages ($\mu = 30.24, \sigma = 12.253$) of the groups. Differences paid for each product by gender were not significant ($t_{87} = 1.73$, $p < 0.087$ for the sex toy; $t_{87} = 0.96$, $p < 0.34$ for the batteries). We therefore believe the groups consisted of comparable populations.

Our flyers solicited participants for a study on the usability of an online search engine so that we would not prime participants to privacy. The flyers informed participants that we would be paying them to shop online and that they would "Keep the Change!" When participants arrived for the ex-

[5]No subject purchased either product from this website, and we therefore do not mention it in the analysis.

[6]Privacy reports are not discussed anywhere else in this paper since too few participants clicked them for us to draw any conclusions.

Figure 3. Screenshot of a website in the *frame* condition.

Figure 4. Screenshot of a website in the *interstitial* condition.

periment, we handed them instruction sheets that labeled the various features of Privacy Finder: the search box, the list of results, the annotated price information, the product pictures, and the privacy indicators. All references to "Privacy Finder" were changed to "Finder" in order to reduce priming effects. Likewise, we scheduled all participants at least 72 hours after taking our privacy concerns screening survey.

We gave participants packets that instructed them to complete several information retrieval tasks in addition to the two purchasing tasks in order to familiarize them with the interface and to conceal the purpose of the study. The tasks included searches for boot prices, prices and average lifetimes of light bulbs, and the prices and available sizes of tote bags. After two information retrieval tasks, participants used Privacy Finder to find websites offering either the sex toy or the batteries and purchased these products. The order in which participants purchased these two items was assigned randomly. The instructions specified the search strings to use to find these products. Unbeknownst to participants, these search strings returned our static search results.

Participants conducted additional information retrieval tasks between the first and second purchases. If they had purchased the batteries first, they purchased the sex toy second, and vice versa. After the second purchase, participants completed an online exit survey that asked questions about their purchases and overall reactions. They were required to use their own credit card and billing information for both purchases so that they would treat the purchases as "real" purchases. However, we allowed them to ship unwanted items to our laboratory. To prevent gaming of the study, we gave participants $10 in cash for completing the laboratory experiment and then another $40 by mail once we had confirmation that their orders had been shipped.[7]

ANALYSIS

Our most significant finding was that the timing of privacy indicator display had a highly significant impact on the behavior of participants who chose to make a purchase on the first website they visited. Those participants paid for increased privacy only when their search results were annotated with privacy indicators; participants who saw the indicators at a later time were significantly more likely to ignore them. Participants who chose to comparison shop by visiting several websites before making a purchase were influenced by the privacy indicators regardless of when they were displayed. Likewise, participants' reliance on the privacy indicators also depended on whether or not they were purchasing the privacy-sensitive item, as well as the strength of the privacy indicator to which they were exposed.

In this section we describe how purchasing behaviors changed when participants were exposed to privacy indicators. Next, we examine how privacy concerns and purchasing behaviors varied based on the type of product being purchased. Finally, we detail how the timing of the privacy indicators resulted in very nuanced behaviors regarding the prices participants paid for the items, how website content had less of a role than we expected, and how timing had an impact on the number of websites participants visited.

General Effects of Privacy Indicators

Hypothesis 1: Participants will pay for increased privacy when they see privacy indicators.

We compared the average price paid by participants in the control (*handicap*) condition with the average price paid by participants in the three experimental conditions to determine whether participants would pay more to shop at sites with privacy indicators than they would to shop at sites with irrelevant green indicators. We performed an ANOVA to compare the prices paid for each product between each of the experimental groups and found that when purchasing the sex toy, participants in the three experimental groups paid significantly more than participants in the *handicap* condition ($F_{3,85} = 7.938$, $p < .0005$). However, while participants in the experimental groups also paid more for batteries than those in the *handicap* condition, we did not observe any significant differences in price paid for batteries between the conditions. We concluded that participants were influenced by privacy indicators rather than by irrelevant indicators. Table 3 shows the average premium that participants paid for each product across all four conditions.

[7]We asked participants to mail us invoices or email us tracking numbers for their purchases so that they would not plan to cancel their orders after they left our laboratory (which would make item prices less of a factor since they would not actually pay for them).

Condition	Battery Premium	Sex Toy Premium
Handicap	\$0.15	\$0.11
Privacy	\$0.34	\$0.52
Frame	\$0.26	\$0.41
Interstitial	\$0.39	\$0.49

Table 3. The average privacy premiums paid for both products across all four study conditions. This is the amount paid above the \$15.50 base price for increased privacy.

Our observed data corroborated the exit survey data: participants who did not see privacy indicators were less likely to consider privacy when making their purchases. We provided participants a text box on the exit survey to enter the biggest factor that they considered when making each purchase. In the *handicap* condition, 82% of participants indicated price was the primary factor during the battery purchase, and 86% indicated price for the sex toy purchase. At the same time, 9% said the website rating was the primary factor during the battery purchase, and 14% mentioned it for the sex toy purchase. In the other conditions, participants claimed price had a less important role, and the website rating was more important. In the *privacy* condition, 64% mentioned price for the batteries (36% cited the privacy rating), but only 36% mentioned price for the sex toy (55% cited the privacy rating); in the *frame* condition, 64% mentioned price for the batteries (18% cited the privacy rating), but only 46% mentioned price for the sex toy (36% cited the privacy rating); in the *interstitial* condition, 52% mentioned price for the batteries (35% cited the privacy rating), while 44% mentioned price for the sex toy (48% cited the privacy rating). As expected, when price played less of a role, the privacy ratings played more of a role in participants' purchasing decisions.

We tried to control the study by only selecting vendors that we believed would be unfamiliar to participants. During the exit survey three participants (3.4% of 89) disclosed that they had done business with our vendors in the past (two sex toy vendors and one battery vendor). However, when we asked them if previous experiences with a particular company were factors (using a 7-point Likert scale) for either purchase, we found no correlation between self-reported familiarity and where participants made purchases during the study.

Product-Specific Privacy

Hypothesis 2: Participants who see privacy indicators will pay more for the privacy-sensitive item than the item that does not raise additional privacy concerns.

We performed a pairwise t-test across both purchases to compare the prices paid for the sex toy with the prices paid for the batteries in each condition (Table 3), and found that participants paid significantly more—for higher privacy levels—for the sex toy than for the batteries in both the *privacy* ($t_{21} = 2.935$, $p < 0.008$) and *frame* ($t_{21} = 2.346$, $p < 0.029$) conditions.

Information	$\mu_{sex\ toy}$	$\mu_{battery}$	t_{88}	p-value
Credit card	4.92	4.55	2.938	.004
Email address	4.87	3.96	5.002	.0005
Physical address	4.29	3.45	4.738	.0005
Phone number	4.62	3.94	4.008	.0005
Purchase history	3.87	2.92	5.499	.0005

Table 4. Participants used a 7-point Likert scale to specify how concerned they were during each purchase when providing various types of personal information.

What we found most interesting was that participants in the *interstitial* condition did not pay significantly more for one product versus the other. Instead, they paid a privacy premium for both products. In this case, the effect of the privacy indicators being displayed as an interstitial diluted the role of product-specific concerns when the participants made their purchases. Thus, they were motivated to find the high privacy websites for both products.

We compared our observed data to the self-reported data that participants provided on our exit survey. In the exit survey we asked participants to rate their privacy concerns for both products on a 7-point Likert scale (six represented "extremely concerned," while zero represented "not concerned at all"). Participants reported an average concern level of 5.56 for the sex toy ($\sigma = 2.291$) and 3.56 for the batteries ($\sigma = 1.864$). We performed a paired t-test and determined that participants had significantly higher levels of concern when purchasing the sex toy ($t_{88} = 7.884$, $p < .0005$). Participants used another 7-point Likert scale to specify how concerned they were during each purchase when providing specific types of information: credit card numbers, email addresses, physical addresses, phone numbers, and purchase histories. For each piece of information, participants were significantly more concerned about what would happen to that information when they provided it for the sex toy purchase than for the batteries purchase, as shown in Table 4.

Participants who saw privacy indicators were able to address many of their privacy concerns by purchasing the sex toy from websites with better privacy policies. However, this was not the case for those in the *handicap* condition, who did not see the privacy indicators.

The Effect of Timing on Prices

Hypothesis 3: Participants will be more likely to pay for increased privacy when they see privacy indicators alongside search results before visiting a website than when they see privacy indicators after clicking on search result links.

Hypothesis 4: Participants will be more likely to pay for increased privacy when they see privacy indicators before they see the content of a website than when they see privacy indicators alongside the content of a website.

The results of our study indicate that the impact of timing was nuanced: Hypothesis 3 was correct for participants who clicked only one search result, but false for participants who visited multiple websites before deciding where to pur-

Condition	Websites	Batteries	(n)	Sex toy	(n)
Handicap	1	\$0.16	(13)	\$0.10	(16)
	>1	\$0.14	(9)	\$0.17	(6)
Privacy	1	\$0.41	(14)	\$0.46	(13)
	>1	\$0.22	(8)	\$0.61	(9)
Frame	1	\$0.03	(8)	\$0.06	(8)
	>1	\$0.39	(14)	\$0.61	(14)
Interstitial	1	\$0.03	(8)	\$0.19	(8)
	>1	\$0.58	(15)	\$0.65	(15)

Table 5. Average privacy premiums paid—above the base price of \$15.50—for each product by participants in the four study conditions. The study conditions are broken down based on whether participants visited multiple websites before making a purchase. The numbers in parentheses reflect the size of the groups.

chase. Table 5 shows the average prices paid for each product across the four study conditions, broken down based on whether participants visited more than one website.

One-click purchases

We performed an ANOVA to compare the amounts participants paid between the different conditions when they visited only one website before purchasing the batteries ($F_{3,39} = 4.772$, $p < 0.006$). We discovered that participants in the *privacy* condition paid significantly more than those in the *frame* ($p < 0.019$) or *interstitial* ($p < 0.019$) conditions.[8] This indicates that participants used the search result annotations to choose websites with increased privacy levels. However, when the privacy indicators were displayed after participants had selected websites from the search results, the participants ignored those indicators, perhaps because they were unwilling to return to the search results. Instead, they were focused on the purchasing task. For these participants the increase in privacy for the batteries was not worth the hassle of selecting new websites from the search results.

We observed slightly different behaviors when participants purchased the sex toys. Again, we observed significant differences between the study conditions ($F_{3,31} = 4.402$, $p < 0.009$), but now the differences were between the *privacy* condition and the *handicap* ($p < 0.012$) and *frame* ($p < 0.027$) conditions. Again, participants in the *privacy* group paid more for privacy when visiting only one website because they saw the privacy indicators before choosing a website to visit. The lack of a significant difference between the *privacy* and *interstitial* conditions is likely a random phenomenon that may disappear with a larger sample size.

Multiple-click purchases

Of the participants who visited multiple websites before purchasing an item, we found that the timing of the privacy indicators did not significantly impact the selection of the website from which they made their purchases. An ANOVA yielded significantly different prices paid for the batteries between the study conditions ($F_{3,42} = 5.424$, $p < 0.003$). Using post-hoc analysis we discovered that participants in the *interstitial* condition paid significantly more than participants in both the *handicap* ($p < 0.004$) and *privacy* ($p < 0.030$) conditions. However, there were no significant differences in battery prices when comparing the *frame* condition with the *handicap* and *privacy* conditions. This can likely be attributed to the role of website content—those who viewed content alongside the privacy indicator relied on the privacy indicator less. It is also likely that because the interstitial interrupted their immediate task and required their attention to dismiss it, the strength of this privacy indicator was greater than that of the other two.

[8]All post-hoc analysis throughout this paper was done using Tukey's HSD test.

Condition		Batteries	Sex Toy
Handicap	(22)	1.86 ($\sigma = 1.17$)	1.41 ($\sigma = 0.91$)
Privacy	(22)	1.86 ($\sigma = 1.36$)	1.73 ($\sigma = 1.12$)
Frame	(22)	3.05 ($\sigma = 1.79$)	3.09 ($\sigma = 1.77$)
Interstitial	(23)	3.09 ($\sigma = 1.78$)	3.04 ($\sigma = 1.69$)
Interstitial*	(23)	2.09 ($\sigma = 1.38$)	1.74 ($\sigma = 1.10$)

Table 6. The total number of search results visited (out of a maximum of five) before participants purchased each product. The last row shows the number of sites visited by members of the *interstitial* condition when they chose to proceed to the website in light of the privacy indicator.

The significantly stronger effect of the *interstitial* condition was only observed during the battery purchase: we observed significant differences between the conditions when examining prices paid by participants who visited multiple websites when purchasing the sex toy ($F_{3,40} = 8.860$, $p < 0.0005$), but this was because everyone exposed to privacy indicators—regardless of timing and placement—paid significantly more than those in the *handicap* condition ($p < 0.001$ for *handicap* vs. *privacy*, and $p < 0.0005$ for both *frame* and *interstitial* vs. *handicap*). This is interesting because it means that those who saw privacy indicators after choosing websites from the search results still ended up purchasing the sex toy from the higher privacy websites—it just took them longer to find them.

The Effect of Timing on Website Visits

Hypothesis 5: Participants who see privacy indicators after clicking on search result links will visit more websites than those who see privacy indicators alongside search results.

We further explored the role of timing by examining the number of search results visited by participants in the *frame* and *interstitial* conditions. Recall that these participants only saw privacy indicators after selecting search results. Table 6 shows the number of websites participants in all conditions visited on average before making a purchase. We performed an ANOVA and found significant differences between the conditions for both the battery ($F_{3,85} = 4.475$, $p < 0.006$) and the sex toy ($F_{3,85} = 8.394$, $p < 0.0005$) purchases.

Because we were primarily interested in how long it took participants to find the websites with the highest privacy levels, we performed another ANOVA, though this time only examining participants who purchased from the websites with four green boxes. When purchasing the batteries, participants in the *privacy* condition clicked significantly fewer

search results to find the website with the four green boxes ($F_{3,22} = 23.126$, $p < 0.0005$). Participants in the *interstitial* and *frame* conditions clicked 203% more search results on average than those in the *privacy* condition to purchase from this same website and obtain the same level of privacy ($p < 0.0005$ for both comparisons). Thus it took participants in the *interstitial* and *frame* conditions significantly longer to find the same high-privacy website that those in the *privacy* condition were able to locate with a single click.

Recall that in the *interstitial* condition, participants must acknowledge the privacy indicator before viewing the destination website. If instead of examining the number of search results clicked, we examine the number of websites viewed by those in the *interstitial* condition, we no longer see a significant difference between the *interstitial* condition and the *privacy* and *handicap* conditions. That is, when participants encountered the interstitial privacy indicator on a website with a low privacy level, they were more likely to return to the search results without viewing that website.

This distinction was also apparent when we examined the number of search results clicked prior to purchasing the sex toy from the website with the highest privacy level ($F_{3,33} = 21.039$, $p < 0.0005$): participants in the *interstitial* and *frame* conditions clicked an average of 168% more websites ($p < 0.0005$ for both comparisons) than those in the *privacy* condition. Again, participants in these three conditions did not differ on the level of privacy they achieved, it merely took them longer to achieve that same level of privacy when the indicators were displayed after search results were selected. Therefore, displaying privacy indicators alongside search results creates more efficient shopping experiences for most users, while also helping users who click fewer search results to achieve greater levels of privacy.

CONCLUSION

In this paper we showed that the timing of privacy information display impacts purchasing decisions: participants who decided to visit only one website to make their purchases paid significantly more money for a higher level of privacy when privacy indicators were presented alongside their search results; similar participants who did not see privacy indicators until after they had already selected a website were unwilling to spend time finding websites with higher privacy levels and instead made purchases from cheaper websites. Likewise, participants who did comparison shopping were just as willing to use interstitial and frame privacy indicators to find websites with higher privacy levels, even though this meant visiting significantly more search results.

Finally, we observed that privacy decisions depended on privacy concerns surrounding the items being purchased: participants had greater privacy concerns when making the sex toy purchases and therefore went out of their way to use the privacy indicators to find websites that offered higher levels of privacy, even if this meant paying a premium. Likewise, many participants were not willing to pay a privacy premium for the batteries because the product did not trigger the same level of privacy concern as the sex toy.

Limitations & Future Work

While we demonstrated that the timing of a privacy indicator's appearance has an impact on whether users visit websites with better privacy policies, there are still many unanswered questions. We did not compare the effect of privacy indicators with other relevant indicators such as customer ratings, nor did we explore the extent to which participants might view privacy indicators as a proxy for other indicators of trustworthiness unrelated to privacy. Two additional areas that we plan to focus on in future studies are how consumers make decisions about privacy premiums and how website content competes with indicators for a user's attention.

Privacy Premiums

We observed that participants were willing to pay premiums to receive higher levels of privacy. In this particular study we used a privacy premium of $0.75. However, we do not know if participants view privacy premiums as a percentage of a purchase price or as a flat rate. That is, would participants have paid this same premium on an item that cost half as much? Would participants pay a $1.50 privacy premium on an item that cost twice as much?

Website Content

Fogg et al.'s work on website credibility indicates that the "look and feel" of a website is the main factor when users make trust decisions [12]. However, we were surprised to discover that this was not always the case: many times participants placed more weight on the privacy indicators than the websites. That being said, it is unclear how exactly participants assessed the quality of the websites they visited. Future studies might examine how participants assess the look and feel of websites while also examining their reactions to privacy indicators.

ACKNOWLEDGMENTS

Thanks to Daniel Rhim for his assistance carrying out this study. We are also grateful to the companies who participated: EdenFantasys.com, Instawares, Little Office Supply, NiteTimeToys.com, Office Quarters, On Time Supplies, SheVibe, and The Dirty Bunny. This work was supported in part by the National Science Foundation under grant CCF-0524189 and by U.S. Army Research Office contract no. DAAD19-02-1-0389 (Perpetually Available and Secure Information Systems) to Carnegie Mellon Universitys CyLab. 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.

REFERENCES

1. ACKERMAN, M. S., CRANOR, L. F., AND REAGLE, J. Privacy in e-commerce: examining user scenarios and privacy preferences. In *EC '99: Proceedings of the 1st ACM conference on Electronic commerce* (New York, NY, USA, 1999), ACM, pp. 1–8.

2. ADKINSON, W., EISENBACH, J., AND LENARD, T. Privacy online: A report on the information practices

and policies of commercial web sites. Tech. rep., Progress & Freedom Foundation, 2002.

3. ANTON, A., EARP, J., HE, Q., STUFFLEBEAM, W., BOLCHINI, D., AND JENSEN, C. Financial privacy policies and the need for standardization. *IEEE Security & Privacy 2*, 2 (Mar-Apr 2004), 36–45.

4. BYERS, S., CRANOR, L. F., AND KORMANN, D. Automated Analysis of P3P-Enabled Web Sites. In *Proceedings of the Fifth International Conference on Electronic Commerce (ICEC2003)* (October 1-3, 2003), pp. 197–207. http://lorrie.cranor.org/pubs/icec03.html.

5. CRANOR, L. F. *Web Privacy with P3P*. O'Reilly and Associates, Sebastopol, CA, 2002.

6. CRANOR, L. F., BYERS, S., KORMANN, D., AND MCDANIEL, P. Searching for Privacy: Design and Implementation of a P3P-Enabled Search Engine. In *Proceedings of the 2004 Workshop on Privacy Enhancing Technologies (PET2004)* (May 26-26, 2004), pp. 314–328.

7. CRANOR, L. F., GUDURU, P., AND ARJULA, M. User Interfaces for Privacy Agents. *ACM Transactions on Computer-Human Interaction 13*, 2 (June, 2006), 135–178.

8. EARP, J., ANTON, A., AIMAN-SMITH, L., AND STUFFLEBEAM, W. Examining internet privacy policies within the context of user privacy values. *Engineering Management, IEEE Transactions on 52*, 2 (May 2005), 227–237.

9. EDELMAN, B. Adverse selection in online 'trust' certifications. In *Proceedings of the 2006 Workshop on the Economics of Information Security (WEIS'06)* (Cambridge, UK, 2006).

10. EGELMAN, S., CRANOR, L. F., AND CHOWDHURY, A. An analysis of p3p-enabled web sites among top-20 search results. In *Proceedings of the Eighth International Conference on Electronic Commerce* (August 14-16, 2006). http://lorrie.cranor.org/pubs/icec06.html.

11. EGELMAN, S., CRANOR, L. F., AND HONG, J. You've been warned: An empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the ACM Computer-Human Interaction Conference* (New York, NY, USA, April 2008), ACM Press.

12. FOGG, B., MARSHALL, J., LARAKI, O., OSIPOVICH, A., VARMA, C., FANG, N., PAUL, J., RANGEKAR, A., SHON, J., SWANI, P., AND TREINEN, M. 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.

13. GIDEON, J., EGELMAN, S., CRANOR, L., AND ACQUISTI, A. Power Strips, Prophylactics, and Privacy, Oh My! In *Proceedings of the 2006 Symposium on Usable Privacy and Security* (12-14, July 2006), pp. 133–144.

14. GOOD, N. S., GROSSKLAGS, J., MULLIGAN, D. K., AND KONSTAN, J. A. Noticing notice: a large-scale experiment on the timing of software license agreements. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2007), ACM, pp. 607–616.

15. HOCHHEISER, H. The platform for privacy preference as a social protocol: An examination within the U.S. policy context. *ACM Transactions on Internet Technology (TOIT) 2*, 4 (2002), 276–306.

16. JENSEN, C., SARKAR, C., JENSEN, C., AND POTTS, C. Tracking Website Data-Collection and Privacy Practices with the iWatch Web Crawler. In *Proceedings of the 2007 Symposium On Usable Privacy and Security (SOUPS)* (Pittsburgh, PA, 2007), ACM Press, pp. 29–40.

17. MILNE, G. R., AND CULNAN, M. J. Strategies for reducing online privacy risks: Why consumers read (or don't read) online privacy notices. *Journal of Interactive Marketing 18*, 3 (Summer 2004), 54–61.

18. MOORES, T. Do consumers understand the role of privacy seals in e-commerce? *Communications of the ACM 48*, 3 (2005), 86–91.

19. POLLACH, I. What's wrong with online privacy policies? *Commun. ACM 50*, 9 (2007), 103–108.

20. SHERMAN, E. Privacy policies are great—for phds, September 4, 2008. http://industry.bnet.com/technology/1000391/privacy-policies-are-great-for-phds/.

21. SOBEY, J., BIDDLE, R., VAN OORSCHOT, P., AND PATRICK, A. Exploring user reactions to browser cues for extended validation certificates. In *Proceedings of the European Symposium on Research in Computer Security (ESORICS'08)* (October 2008).

22. TRUSTE. TRUSTe Fact Sheet, 2008. http://www.truste.org/about/fact_sheet.php.

23. TSAI, J., EGELMAN, S., CRANOR, L., AND ACQUISTI, A. The effect of online privacy information on purchasing behavior: An experimental study. In *Proceedings of the 2007 Workshop on the Economics of Information Security (WEIS'07)* (Pittsburgh, PA, USA, 2007).

24. WHALEN, T., AND INKPEN, K. M. Gathering Evidence: Use of Visual Security Cues in Web Browsers. In *Proceedings of the 2005 Conference on Graphics Interface* (Victoria, British Columbia, 2005), pp. 137–144.

25. WU, M., MILLER, R. C., AND GARFINKEL, S. L. Do Security Toolbars Actually Prevent Phishing Attacks? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Held in Montreal* (2006), ACM Press, pp. 601–610.