

THE EFFECT OF ONLINE PRIVACY INFORMATION ON PURCHASING BEHAVIOR: AN EXPERIMENTAL STUDY

Janice Y. Tsai

Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
jytsai@andrew.cmu.edu

Serge Egelman

Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
egelman@cs.cmu.edu

Lorrie Cranor

Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
lorrie@cs.cmu.edu

Alessandro Acquisti

Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
acquisti@andrew.cmu.edu

Pre-publication version

Forthcoming in ISR, 2010

Abstract

Although online retailers detail their privacy practices in online privacy policies, this information often remains invisible to consumers, who seldom make the effort to read and understand those policies. This paper reports on research undertaken to determine whether a more prominent display of privacy information will cause consumers to incorporate privacy considerations into their online purchasing decisions. We designed an experiment in which a shopping search engine interface clearly and compactly displays privacy policy information. When such information is made available, consumers tend to purchase from online retailers who better protect their privacy. In fact, our study indicates that when privacy information is made more salient and accessible, some consumers are willing to pay a premium to purchase from privacy protective websites. This result suggests that businesses may be able to leverage privacy protection as a selling point.

Keywords: Privacy, Information Systems, Economics, Experimental Economics, E-Commerce

*Publish Version Can Be Found:
Information Systems Research
22(2), 254-268 (2011)
<http://doi.org/10.1287/isre.1090.0260>*

1. Introduction

Most Americans believe that their right to privacy is “under serious threat,” (CBS News, 2005) and express concern with businesses that collect their personal data (Harris Interactive, 2001; CBS News, 2005; P&AB, 2005; Turow *et al.*, 2005; Lebo, 2008; Consumer Union, 2008; Burst Media, 2009). According to surveys, such concerns affect consumers’ willingness to purchase online or register on websites (P&AB, 2005). Businesses address these privacy concerns by posting privacy policies (Culnan, 2000) or displaying privacy seals (Benassi, 1999) to convey their information practices. However, 70% of people surveyed disagreed with the statement “privacy policies are easy to understand” (Turow *et al.*, 2005), and few people make the effort to read them (Privacy Leadership Initiative, 2001; TRUSTe, 2006). Similarly, empirical evidence suggests that consumers do not fully understand the meaning of privacy seals (Moore, 2005). Various studies have also indicated that most people are willing to put aside privacy concerns, providing personal information for even small rewards (Acquisti and Grossklags, 2005a). In such cases, people readily accept trade-offs between privacy and monetary benefits (Hann *et al.*, 2007) or personalization (Chellapa and Sin, 2005).

In this paper we empirically investigate whether prominently displayed privacy information will cause consumers to incorporate privacy considerations into their online purchasing decisions. Answering that question may not only reveal a great deal about privacy-related consumer behavior, but also contribute to a long-standing debate: whether or not businesses can use privacy strategically, leveraging the protection of private information for competitive advantage (Gellman, 2002; Rubin and Lenard, 2002).

We present the results of an online concerns survey and an online shopping experiment conducted in a laboratory. We used the online concerns survey to identify the most pressing

types of online privacy concerns and to determine which types of products are most likely to elicit such concerns in a purchasing scenario. We then invited a different set of participants to test a new search engine whose search results were annotated with icons. These participants were asked to search for and purchase products online using the search engine shopping interface. In a between-subjects design, participants across different experimental conditions received different explanations of what the icons meant. In the “privacy information” condition, participants were told that the icons indicated a rating based on an analysis of the site's privacy policy. In two control conditions, the icons either indicated ostensibly irrelevant information (the site's handicap accessibility rating for sight-impaired users) or were absent. In all conditions, natural language privacy policies were available via the merchants' existing “Privacy Policy” links.

The icons presented privacy information in a prominent manner. We found that participants in the privacy information condition were more likely than those in other conditions to make purchases from websites offering medium or high levels of privacy, even when the price was higher than the price on other sites. Those in the control conditions generally made purchases from the lowest priced vendor. Furthermore, individuals presented with irrelevant indicators were less likely than those in the privacy information condition to take these indicators into consideration when making purchases.

Our results suggest that individuals are willing to pay a premium for privacy when privacy information is made prominent and intuitive. While many suggest that even privacy conscious consumers are unlikely to pay for online privacy (Shostack, 2003) or give up rewards to protect their data (Spiekermann *et al.*, 2001), our results suggest that businesses may be able to use information technology tools (such as those built upon computer-readable privacy policies) to present their privacy practices in a prominent and accessible way. Such a practice would allow

businesses to strategically manage privacy and leverage privacy protection for a competitive advantage.

The rest of this paper is organized as follows: in Section 2, we discuss related literature on privacy valuations, privacy policies and seals, and the privacy search engine used in our experimental study. In Section 3, we present the theoretical background underlying our study. In Section 4, we describe the methodology of our empirical study and the experimental hypotheses on which it was based. In Section 5, we present its results. We discuss limitations and implications in Sections 6 and 7, respectively.

2. Related Literature

2.1 Privacy valuations

Privacy is notoriously difficult to define. Smith *et al.* (1996) outline four dimensions of consumer privacy concerns: *collection* of personal information, *unauthorized secondary use* of personal information, *errors* in personal information, and *improper access* to personal information (see also Stewart and Segars, 2002). In online marketing, these dimensions of concern have been interpreted to refer to the *collection* of personal information, *control* over the use of personal data, and *awareness* of privacy practices and how personal information is used (Malhotra *et al.*, 2004). Other consumers' concerns (as defined by Brown and Muchira, 2004) focus on unauthorized secondary use and errors in personal information. When those concerns are elicited by the merchant's behavior, the individual may lose trust in the merchant (Camp, 2003). Milne and Gordon (1993) refer to the proper treatment of consumer information as an "implied social contract" with the customer. When a breach of confidentiality between the organization and the individual occurs, the violation of trust may entitle the victim to compensation (Solove, 2006). On the other hand, the guarantee of fair information practices can

counterbalance consumers' concerns about information sharing (Culnan and Artmstrong, 1999; Dinev and Hart, 2006).

Over time, surveys have consistently indicated that people are concerned with the ways businesses use their personal information. Ostensibly, those concerns prevent some consumers from making online purchases. A 2005 survey conducted by Privacy & American Business, for instance, found that concerns about the use of personal information kept 64% of respondents from purchasing from a company, while 67% of respondents declined to register at a website or shop online because they found the privacy policy to be too complicated or unclear (P&AB, 2005). On the other hand, consumers have also been found to provide personal information in exchange for small discounts or rewards. A 2002 Jupiter Research study found that 82% of online shoppers were willing to give personal data to new shopping sites in exchange for the chance to win \$100; 36% said they would allow companies to track their World Wide Web surfing habits in exchange for \$5 discounts (Tedeschi, 2002). In an experimental investigation, Spiekermann *et al.* (2001) found evidence that even individuals concerned with privacy are willing to trade privacy for convenience and discounts. As the authors noted, "most [study participants] stated that privacy was important to them, with concern centering on the disclosure of different aspects of personal information. However, regardless of their specific privacy concerns, most participants did not live up to their self-reported privacy preferences." Similar discrepancies have been found in other privacy scenarios involving consumer grocery cards (Acquisti and Grossklags, 2005a) and online social networks (Acquisti and Gross, 2006). The fact that privacy-related businesses have had such difficulties finding a market for their products (Brunk, 2002) further suggests that many consumers are reluctant to pay for privacy.

Several researchers, working to determine what drives consumer privacy valuations, have investigated how individuals trade privacy for monetary or intangible benefits. Hann *et al.* (2007) tried to quantify the value individuals ascribe to website privacy protection, finding that “among U.S. subjects, protection against errors, improper access, and secondary use of personal information is worth US\$30.49-44.62.” However, the conjoint analysis approach and the hypothetical nature of the study make it difficult to determine conclusively whether individuals will, in actuality, pay to protect their privacy. Chellappa and Sin (2005) found evidence of a tradeoff between consumers’ desire for personalization and their concern for privacy. Huberman *et al.* (2005) used a second-price auction experimental setup to study what price individuals put on specific pieces of private information (such as their weight).. They found that individuals wanted more money to reveal information that was “abnormal” or “undesirable”. In a contingent valuation survey of the value assigned to enforceable property rights to personal information, Rose (2005) found that survey participants expressed a high sensitivity to privacy, but that only 47% of them would be willing to pay for those property rights (an average of NZD 55.40 or USD 28.25). Hui *et al.* (2006) used a field experiment in Singapore to study the value of various privacy assurance measures. They also found that privacy statements and monetary incentives could induce individuals to disclose information.

A debate has therefore emerged in the literature, one centered on the seeming contradiction between people’s expressed privacy concerns *and* their willingness to trade-off privacy for even small benefits. Some believe this is evidence of inconsistent behavior, while others point to rational decision making processes and between-subject variance in privacy sensitivities¹

¹ Westin (1990) clustered individuals around three archetypal privacy sensitivities: the unconcerned, the pragmatist, and the fundamentalist.

(Acquisti and Grossklags, 2003; Shostack, 2003; Syverson, 2003; Acquisti, 2004; Acquisti and Grossklags, 2005a; Wathieu and Friedman, 2005). The literature has highlighted several factors that may affect individual privacy attitudes differently than they affect actual behavior; these factors include variability in individual privacy sensitivities, bounded rationality, behavioral or cognitive biases, such as immediate gratification or optimism bias (Acquisti, 2004), and information asymmetry (Akerlof, 1970). Information asymmetry in particular plays a double role in privacy valuations and decision-making. To use an example from the context of electronic shopping, before a consumer completes her first purchase with an online merchant, the merchant may have limited information about the consumer's taste, reservation price, identity, and so on (see Taylor, 2004; Acquisti and Varian, 2005). However (and more pointedly), *after* the purchase, the consumer may not know how the merchant will use the personal information she revealed as part of the transaction (Acquisti and Grossklags, 2005b). This lack of information arguably affects individual behavior in different ways. For one, consumers may perceive greater risk and uncertainty when dealing with merchants whose privacy policies are unknown; as a result, they may be less willing to complete transactions with those merchants. However, if the lack of information is so profound that consumers are not even aware that their personal information could be exchanged or misused, it may make them *more* likely to engage in such risky (from a privacy perspective) transactions.

2.2 From Asymmetric Information to Privacy Policies, Seals, and Privacy Finder

To avoid potential losses stemming from consumers' lack of information about privacy practices or their mistrust of online shopping, online industry has developed a number of solutions designed to assuage consumers' privacy concerns. *Privacy policies*, which have been widely adopted by online businesses, are one attempt to reduce information asymmetry (Milne and

Culnan, 2002). In principle, privacy policies fill the information gap between the consumer and the vendor by providing a complete picture of the vendor's information practices. In practice, however, perusing privacy policies has its share of transaction costs (McDonald and Cranor, 2009): for instance, the policies themselves may be difficult to understand (Hochhauser, 2003; Jensen and Potts, 2004) and may be time consuming to read. As a result, people rarely read them (Privacy Leadership Initiative, 2001; Jensen *et al.*, 2005; TRUSTe, 2006). When they do, they often make mistaken assumptions about their meaning: one study found that a majority of Americans who report having seen privacy policies on popular websites believe the presence of a link to a privacy policy means that their data is protected (Turow *et al.*, 2005). In short, individuals who know that a company or organization has a privacy policy may still lack enough information to make informed decisions.

Another self-regulatory solution (which has been adopted in a limited fashion) relies on third-party certification of a merchant's adherence to its own privacy policy through *privacy seal programs* (Benassi, 1999). Privacy seals may help reduce information asymmetry by reducing the cost a consumer incurs when accessing and assessing information about a merchant's data practices (Zhang, 2004). Privacy seals may also improve consumers' perceptions of the vendor (Miyakazi and Krishnamurthy, 2002). However, empirical evidence about the effect of privacy seals is mixed (Moores and Dhillon, 2003). Belanger *et al.* (2002) found no evidence that seals impact individuals' intention to purchase, while Moores (2005) found that consumers seem to misunderstand privacy seals. On the other hand, Rifon, LaRose, and Choi (2005) found that privacy seals enhanced users' trust in the Web site they were visiting, and Mai, Menon, and

Sarkar (2006) showed that firms bearing privacy seals tend to list higher prices than their competitors.²

Both privacy policies and privacy seals do not seem to consistently impact consumer decision-making – either because the information they provide remains invisible to consumers, or because it is ignored or misinterpreted.³ However, the question remains: how is consumer decision-making impacted when information about a merchant’s privacy practices is made more prominent (in a position where the consumer is unlikely to ignore it) and more accessible (represented in an intuitive manner)? Such changes could reduce the transaction costs associated with learning a merchant’s information practices and thus, arguably, also reduce the size of the information asymmetry gap between consumer and merchant.

In our experimental design, we made use of a tool called Privacy Finder (Byers *et al.*, 2004) to answer that question. Privacy Finder is a search engine that annotates a user’s Google or Yahoo! search results with “privacy meter” icons produced through an automated analysis of the P3P policies of the retrieved sites.⁴ These icons graphically represent how well a website’s privacy policy matches preferences specified by the user. Privacy Finder also generates “privelacy reports” for P3P-enabled websites. These reports present privacy information that is “of greatest concern to users” in a simplified format (Cranor *et al.*, 2006). As compared to the *status quo*, which we tested as the control condition (the merchant’s original privacy policy), the display of intuitive icons during the search stage of a consumer’s shopping experience offers a tool to test

² In related work, Tang, Hu, and Smith (2007) present a theoretical model contrasting privacy seals, privacy policies, and privacy regulation. Edelman (2006) studies adverse selection in online trust certifications.

³ Larose and Rifon (2007) find that explicit privacy warnings increase perceptions of information risks in individuals, but not in the presence of privacy seals.

⁴ The Platform for Privacy Preferences (P3P), a machine-readable format for privacy policies, was developed in 2002 to facilitate user access to privacy information. People use software tools to define their privacy preferences and determine if websites’ P3P privacy policies match those preferences (Cranor, 2002). The search engine used in our study translated these computer-readable privacy policies and displayed a “privacy icon” for each site with a P3P policy.

whether more prominent and accessible privacy information affects consumers' purchasing behavior. In an earlier study, we found preliminary evidence that online shoppers seek more privacy-friendly websites when privacy policy information is made available in search engines (Gideon *et al.*, 2006); however, we did not investigate whether consumers were willing to trade money for privacy. In Section 4, we explain how we modified Privacy Finder to examine that issue.

3. Theoretical Framework and Research Objectives

If privacy were a feature consumers truly value when making online transactions, privacy friendly merchants would gain a competitive advantage over their counterparts. The competitive advantage would potentially allow these merchants to command price premiums over the competition (Shapiro, 1983; Mai, Menon, and Sarkar, 2006). While trust building technologies have been shown to impact price premiums in online auction markets (Ba and Pavlou, 2002), the evidence for the privacy case, as highlighted in the previous section, is mixed at best. One of the factors introduced in the previous section to explain why privacy protection may increase a consumer's expected utility and yet fail to influence her behavior is asymmetric information. It is expensive for consumers to gain information about a company's data practices by looking at its privacy policy; as a result, consumers may not be consistently aware of — or do not focus upon — possible privacy concerns when transacting online. Furthermore, the prospective cognitive cost of reducing the information asymmetry about how a merchant handles consumers' information may be too large. Subsequently, privacy considerations may carry significantly less weight in a consumer's utility function than other factors, such as the vendor's price.⁵ If this is the case, providing clearer information about a merchant's privacy policy may reduce

⁵ Vila *et al.*, 2004 describe this process as a “lemons market” dynamics for privacy policies.

information asymmetry, decreasing the transaction costs associated with learning a merchant's information practices, and thereby increasing the weight of privacy considerations in the consumer's utility function and decision-making process.

We can represent this scenario within a simple microeconomic framework. Let us define the consumer's utility maximization problem when purchasing a good from an online merchant i as:

$$U(v, p_i, c_i, d_i), \quad [1]$$

where U represents the utility the consumer wants to maximize, v represents the consumer's valuation of the good (identical across the merchants selling the homogeneous good), p_i represents the price charged by merchant i , c_i is a proxy for the privacy concerns the individual associates with the purchase of the good from merchant i , and d_i represents other residual factors that may influence the consumer's utility when purchasing that good from that particular merchant. Naturally p , and c are expected to enter the utility function with negative signs, v with a positive sign, and d with an undetermined sign. For our explanatory purposes, it is not necessary to specify the functional relation between the various factors: we assume that consumers can purchase the same good from different merchants, and that merchants may have different prices, reputations, individual characteristics, and privacy policies that may elicit different privacy concerns. The consumer needs to choose the merchant i from which she will make a purchase.

If consumers acted as fully-informed, rational agents, c_i would accurately reflect the subjective weight of a consumer's privacy concerns (the expected, subjective privacy costs) when purchasing from merchant i . Incomplete information may reduce the weight of privacy concerns in the purchase decision. Conversely, other things being equal, prominent and accessible data about different merchants' privacy policies may increase the weight of c_i in the consumer

maximization decision; it would do so by alerting the consumer about privacy considerations and reducing the cost of comparing the privacy policies of different merchants. In a sense, by making privacy information more prominent, part of the consumer's attention gets shifted towards privacy, reducing the consumer's relative focus on price considerations.⁶ Such a change would be inferred by observing the consumer's choice of merchant i (with different perceived privacy costs c_i but also different prices p_i) for purchases. Using a revealed preferences argument, we expect consumers' purchase decisions to reveal the utility they expect to gain from the transaction, making it possible to estimate the weight they grant the various factors in Equation [1].

By using an experimental approach to control for merchants' privacy policies and prices, and by manipulating the level and type of privacy-relevant information provided to participants in the study, it becomes possible to test the hypothesis that the availability and accessibility of relevant privacy information will affect consumers' purchase selections. Given large enough control and treatment groups, we can assume that the unobservable factors embodied in d_i (such as respondents' heterogeneous preferences for certain merchants or perceived trustworthiness for specific sites) will be similarly distributed within different experimental groups. These factors will therefore not significantly interfere with the comparison of the relative effect of additional privacy information between the groups.

In Section 4, we present a study based on the above framework, one that tests whether privacy information can affect consumer purchasing behavior. Specifically, the objectives of the study were: 1) to determine whether the prominent display of privacy information causes privacy-

⁶ Under a limited capacity model of attention (see Kahnemann, 1973; McLeod and Jones, 1986), tasks and interrupts compete for individuals' limited attention resources and cognitive capacity (see also Yerkes & Dodson, 1908).

concerned users to take privacy into account when making online purchasing decisions; and 2) to determine whether privacy-concerned users are willing to pay a premium to make their purchases from more privacy-friendly merchants.

4. The Study

We used the Privacy Finder search engine to test the impact of prominent privacy information on purchasing behavior. Our study consisted of three parts: 1) an online concerns survey to determine what types of privacy concerns and products to include in the experimental part of the study (Section 4.1); 2) an online shopping experiment to investigate how the prominent display of privacy information affects the purchase behavior of privacy-minded users (Section 4.2); and 3) a post-experiment interview (Section 4.3). While the shopping experiment took place in a laboratory, the privacy and monetary incentives associated with the experiment were real, as detailed below.

In our experiment, we compared the way users currently obtain privacy policy information (a link to a privacy policy on the merchant's site) to a method in which privacy information was made more prominent and accessible, with search engine results presenting privacy icons. In all conditions, participants could still access privacy policies as they normally would – by clicking on the privacy policy link on a particular merchant's site.

4.1 Online Concerns Survey

We developed an initial online concerns survey with two high-level questions in mind. First, we wanted to examine the types of privacy concerns individuals have when they shop online (and the risk individuals associate with each of these concerns); this allowed us to design an experiment in which those concerns were addressed by the prominent privacy information

provided. Second, we wanted to determine the types of products that may or may not elicit privacy responses in a purchasing scenario.

The design details and demographic characteristics of participants in the online concerns survey are discussed in the online Appendix.⁷ Through the survey, we found that the scenarios participants rated with the highest likelihood of occurring were the same as those addressed by the Privacy Finder Search engine. We also identified two products for participants to purchase in our online shopping experiment. We wanted to find one product that would raise few significant privacy concerns and one that would be more privacy-sensitive, raising significant concerns for most participants. We posed the following question to our survey participants:

We will be conducting studies for an online shopping and privacy research project in which we will pay participants to make online purchases with their own credit cards. Each participant will receive enough money to cover the cost of the purchase plus \$10. If you were asked to participate, would you be willing to purchase the items below with your own credit card, and how concerned would you be about doing so?

Most participants showed little resistance to purchasing common products like office supplies online. We detected increasing hesitance as we moved to items that involved personal values and mental states, such as items related to sex and books on depression. When items were indicative of violent behavior, like bullets or a book on bomb-making, we found significant reservations. We used these insights to guide our selection of products for the experiment (See A1.3 in the online Appendix).

4.2 Online Shopping Experiment

We conducted the online shopping experiment in the Carnegie Mellon Usable Privacy and Security (CUPS) laboratory in Pittsburgh, PA. The experiment was designed so that participants

⁷ The Appendix is available at <http://andrew.cmu.edu/~jytsai/privacystudy.pdf>.

faced actual privacy concerns and monetary incentives. The participants were solicited to “test a new search engine interface.” The tasks participants were asked to complete included searching for trivia-like information and purchasing products online using the new search engine shopping interface. In a between-subjects design, participants across different experimental conditions were given different explanations of what the icons that accompanied their search results meant (see A2.4 in the online Appendix). In the rest of this section, we describe participant recruitment, the screening survey, the experimental protocol, the experimental design, and our hypotheses.

4.2.1 Participants Recruitment

Participants were recruited from the general Pittsburgh population; there was no overlap between the participants in the online shopping experiment and the respondents to our online concerns survey. Participants were sought for an “Online searching and shopping study,” with flyers posted around town, online in the *Volunteers* section of Craigslist, and via the Center of Behavioral and Decision Research at Carnegie Mellon. Participants had to be at least 18 years old, have a personal credit card to use during the study, and have experience shopping online. The flyer also advertised that participants would be paid to shop online using our money and would get to “Keep the change.”

4.2.2 Screening Survey

Interested participants were directed to a preliminary survey online. We received 272 complete responses. Our study was designed to target individuals concerned with privacy rather than the population at large: we assumed that our search interface would be helpful to people with some online privacy concerns. We calculated a “risk score” for each participant and used it to screen out those who perceived online shopping to involve little or no privacy risk. Based on this requirement, we screened out 12.5% of the total respondents. Participants who met our

requirements were contacted via email several weeks later to schedule a laboratory session. Due to the delay between the survey and the laboratory sessions, we believe there is little chance that the screening questions primed participants to think about privacy during the laboratory sessions.

We also used the screening survey to ask participants to rate the importance of various factors they might consider when choosing a website for a purchase. These factors and their mean ratings are detailed in the Appendix, Section A2.1. Participants reported that they primarily base purchasing decisions on price, followed by return policy. Shipping speed, customer service, privacy policy, website design, and customer reviews were rated as equally important. We used participant ratings of these purchasing factors to determine which have minimal impact on purchasing decisions – an insight that we used to design the experimental conditions. The factor “accessibility for sight-impaired users” was found to have almost no impact on purchase intentions.

4.2.3 Experiment Protocol

Participants were given an informed consent form when they arrived at our laboratory.⁸ After reading and signing the form, participants were given a “Search Engine Key.” This key served as instructional material (similar to Figure 1), explaining the meaning of the icons and other user interface features. Participants in the three experimental conditions had nearly identical information, but the explanations of the icons differed (see Section 4.2.5). To reduce any framing and priming effects, Privacy Finder was renamed *Finder*, and participants did not see or have access to the privacy preference settings. Instead, based on the results of the online concerns survey, Finder was configured to use the “medium” privacy setting. The “medium” setting

⁸ A chart representing the complete experiment protocol is provided in the online Appendix, Section A2.2.

calculates a warning based on the sharing of personal financial information, purchase information, or personally identifying information; a website's refusal to allow a user to remove their personal information from marketing lists; and the inability of users to view their own information on the site.




Figure 1: Search Engine Key presented to participants in the Privacy Information condition.

To familiarize participants with the interface and draw focus away from the purchasing tasks, participants across all conditions were asked to complete the same six search tasks; instructions for these tasks were provided one task at a time. Only the fourth and sixth tasks required participants to search for vendors selling a specified item (a pack of batteries and a sex toy – the order was randomized across participants) and *use their credit card* to actually purchase the product from the site of their choice. Participants were also asked to write down the website from which they had made their purchase along with the total price they paid. The web browsers were

configured so that all traffic passed through a proxy server to create logs noting the number of websites browsed, visits to the privacy reports, and visits to the privacy policies of the perused websites.

As noted above, we based our selection of the items participants had to purchase during the experiment on the online concerns survey. We selected products that had an average cost of \$15 per item, including shipping. These products also had to be available from a variety of real websites with diverse privacy policies. One item was an office supply product: an 8-pack of Duracell AA batteries; the other item was a vibrating sex toy, the “Pocket Rocket Jr.” Participants used their own credit cards to pay for the products, which meant that their personal information was exposed to real merchants during the study. The websites were actual, *real* merchant sites, and they were chosen due to the very small likelihood that they would be familiar to the participants (to avoid confounding biases from brand effects). However, though the participants did not know it, we had preselected which merchant websites would appear during the users’ searches for the online purchasing tasks. Purchasing either item (the batteries or the sex toy) forced individuals to reveal personal information (their credit card number) to unknown merchants; this arguably may have raised privacy concerns. However, one item (the sex toy) could be considered more personal and sensitive than the other, and may have therefore elicited greater concerns.

4.2.5 Experimental Design and Hypotheses

The Privacy Finder annotates search results with icons that represent a five-point privacy “meter” (see Table 1). The meter is composed of a set of four boxes that are shown as green (filled) or white (empty) based on an algorithm that accounts for the number of privacy preference mismatches between the site’s privacy policy and the user’s privacy preferences.

Thus, a site that violates most of the user’s preferences will have zero or one box filled, while a site with only a few mismatches might have two or three filled boxes. Sites without P3P policies are not annotated with a privacy icon. Privacy Finder also provides a link to the privacy report for each P3P-enabled website.

Icon	Site
■ ■ ■ ■	Matches privacy preferences
■ ■ ■ □	
■ ■ □ □	Does not match privacy preferences
■ □ □ □	
□ □ □ □	

Table 1: Privacy Finder’s privacy indicators

We modified Privacy Finder for online shopping, submitting search queries via the Yahoo! shopping interface and returning search results annotated with product photographs and price information, as well as the privacy information described above.

We randomly assigned participants to one of three experimental conditions.⁹ Across all conditions, participants viewed the same set of search results in the same order. Sites were selected based on their privacy policies and the price of the product. Therefore, a site with “4 green boxes” or “high privacy indicator” offered a high level of privacy protections regardless of whether or not participants were presented with privacy indicators in their set of search results. We compared participants’ purchasing decisions in the following between-subjects design to gauge the impact of providing privacy information:

- **Condition 1 (control condition)**, No privacy indicator: This group viewed search results

⁹ To determine the sample size for the study, we performed a power analysis for two proportions, evaluating whether 50% of the participants in the privacy condition would purchase from “high privacy” sites as compared to 10% in the other conditions ($\alpha = 0.05$, $\beta = 0.2$). To yield a power of 80%, 16 participants were required for each condition, for a total of 48 participants. In each condition, the participants were divided equally by gender.

without any annotations (as is the case with actual merchants in the *status quo*). Participants were given a version of the Search Engine Key that highlighted the type of data the search engine made visible: merchant names, product prices, photos, and so on. Search results during the experiment did not include any Finder icons. However, the natural language privacy policies were still accessible from the merchants' sites.

- **Condition 2 (control condition)**, Irrelevant information: This group viewed search results annotated with icons representing irrelevant information. Participants were given a Search Engine Key that highlighted the presence of green box icons indicating a high or low "rating calculated based on our analysis of the site's computer readable accessibility information for vision-impaired users." (Natural language privacy policies also remained accessible from the merchants' sites.)
- **Condition 3 (treatment condition)**, Privacy information: Privacy icons and links to privacy reports were presented to this group. Participants in this condition were given a Search Engine Key that highlighted the presence of green box icons indicating a high or low privacy "rating calculated based on our analysis of the site's computer readable privacy policy." During the experiment, the search results visible to participants in this condition included such icons.

We selected an irrelevant information condition (in addition to the baseline control condition of status quo information) to rule out the possibility that the presence of an icon by itself would have as much influence on purchase decisions as the presence of privacy information. In previous studies, other content-free symbols (including credit card logos) have increased participants' willingness to trust certain sites (Jensen *et al.*, 2005).

The between-subjects design allowed us to test the following hypotheses, derived from the theoretical framework described in Section 3:

Hypothesis 1: *Participants in the privacy information condition will be more likely than those in the no privacy indicator condition to purchase from websites annotated with icons.*

Hypothesis 2: *Participants in the privacy information condition will be more likely than those in the no privacy indicator condition to purchase from websites annotated with the four-green-boxes icon (the sites offering the best privacy policy).*

Hypotheses 1 and 2 follow from the theoretical background presented in Section 3. When individuals are uncertain or ignorant of a merchant's privacy practices and the resulting potential for privacy issues, privacy concerns have little influence over the decision to make a purchase (Acquisti, 2004). When merchants provide accessible privacy information, the consumer's utility function will give more salience and weight to privacy considerations; as a result, consumers in the privacy information condition should be more likely to purchase from merchants with better privacy policies.

In Hypothesis 2, we theorize that participants will be compelled to purchase from the site that offers the best privacy policy (four-green-boxes). This is not only because the privacy policy is available, but also because it is easy for the consumer to compare sites that offer high levels of privacy to those offering low and medium levels of privacy.

Hypothesis 3a: *Participants presented with prominent privacy information (those in the privacy information condition) will be more likely than those in the no privacy indicator condition to pay a premium to purchase from sites that have better privacy policies.*

Once salient information about privacy is provided and privacy considerations have a more significant role in the consumer's utility function, one would expect some consumers to trade money for privacy. The decision to make this trade depends on the relative strength of their

privacy and price sensitivities (see also Acquisti and Varian, 2005; and Taylor, 2004, for privacy models with price discrimination): the interplay of p_i and c_i in Equation [1].

***Hypothesis 3b:** In the absence of prominent privacy information, people will purchase where price is lowest.*

This hypothesis follows directly from basic microeconomic theory and is used purely as a control for Hypothesis 3a.

***Hypothesis 4:** Icons in the privacy information condition will affect purchase decisions more than icons in the irrelevant information condition.*

This hypothesis is inspired by the literature on “institutional-based trust” that studies structures and situations that affect trust-based individual decision-making (McKnight and Chervany, 2002). For instance, consumers often consider trust seals to be a proxy for merchant quality (Riegelsberger *et al.*, 2005). Hence, in the “irrelevant information” condition, the green icons visible through the interface may be interpreted as proxies of merchant quality regardless of their actual meaning (see also Jensen *et al.*, 2005). We wish to differentiate between the actual impact of privacy information and the impact of institutional-based trust; that is, we wish to rule out the possibility that consumers make decisions based solely on the presence of icons, regardless of their meaning. If Hypothesis 4 is supported, we will be able to conclude that our participants’ purchasing decisions were affected more by privacy considerations than by the search engine interface itself.

4.2.6 Incentives and Reimbursements

We paid participants a two-part “lump sum” payment of \$45 for their participation in the study. The participants kept the products and any money left over after the purchases were made. This

design created a price incentive, encouraging participants to purchase from merchants with lower prices. To best capture the “premium” that participants paid for privacy, we ordered search results based on both privacy level and price across all conditions. The first item was the least expensive and was sold by a web site without a P3P policy (thus no privacy information was readily available). With each subsequent result, both the privacy level and the price increased, as shown in Figure 2. Based on previous pilot studies, we found that participants were unlikely to browse beyond the first four search results. Thus, we did not focus on the specific order of privacy levels beyond the first four sites.




Figure 2: Search engine results interface for the Privacy Information condition.

User study payments were made in two installments to prevent gaming the study (for instance, canceling the purchase after the study). At the end of the session, participants were given \$10 in cash. Once the products shipped *and* the study participants sent us tracking numbers or product packing slips, they were mailed the remaining \$35 payment.

Due to product availability and the fluctuation of product and shipping prices, we used

marginally different sets of search results during the study¹⁰ (see A2.3 in the online Appendix) while keeping both the price and privacy policy distributions fairly constant. The premium for “high privacy” batteries ranged from 3-5% of the product cost, while the premium for the sex toy ranged from 7-10%. Due to retailer problems that occurred during the purchasing tasks, as well as some participants’ refusal to make some of the purchases, we continued to recruit participants until we had collected 48 complete responses for the study.¹¹

As stated above, participants paid for the products using their own credit cards and were later reimbursed a *fixed* amount. This means that both the privacy concerns (revealing personal information to a merchant site) and price incentives were real.

4.3 Exit Survey

Upon completion of the study tasks, participants completed an exit survey. We asked whether the privacy icon (if seen) played a role in their purchasing decisions, whether they understood what the icon represented, whether they read any of the privacy policies, and whether those privacy policies influenced their purchasing decisions. This set of self-reported data was compared with and complemented the quantitative results of our experiment.

5. Results

We found that participants in the privacy information condition were more likely to make purchases from websites offering medium or high levels of privacy (even when those sites charged higher prices), while those in the control conditions generally made purchases from the

¹⁰ The first (and cheapest) result for the batteries search was out of stock when 18 participants completed the experiment. Because, as a result, we could not use these participants’ battery purchase data, we recruited 18 additional participants. We retained the sex toy purchase data for those participants.

¹¹ Due to the nature of the privacy-sensitive product, two participants opted to cease their participation in the study, six opted out of the privacy-sensitive product purchase but completed the remainder of the study, and one decided not to purchase either item but completed the exit survey.

lowest priced vendor. This indicates that individuals are likely to pay a premium for privacy when privacy information is made more accessible. Furthermore, individuals presented with the same indicators as those used for the privacy group – but ostensibly attached to irrelevant merchant features – were less likely to take those indicators into consideration when making purchases. This demonstrates that the observed behavior cannot simply be attributed to an interest in purchasing from web sites labeled with attractive indicators.

5.1 Meaningful Privacy Information

Hypothesis 1: Participants in the privacy information condition will be more likely than those in the no privacy indicator condition to purchase from websites annotated with icons. – **Supported.**

One of the goals of this study was to determine whether participants presented with salient privacy information would be more likely to purchase from sites with privacy indicators than participants who did not see that information. As shown in Table 2, we found that to be the case.

	Conditions		Fisher's Exact <i>p</i>
	Condition 1: No Privacy Indicator	Condition 3: Privacy Information	
% of battery purchases made from sites with icons	11.1% n=2/18	77.8% n=14/18	<.0001
% of sex toy purchases made from sites with icons	16.0% n=4/25	66.7% n=14/21	<.005

Table 2: A between-conditions comparison of the proportion of purchases made from sites corresponding to those annotated with icons in the privacy information condition. To test for significance between these proportions we used the Fisher’s Exact test.

For both products, participants in the privacy information condition made a greater proportion of purchases from sites that displayed privacy icons. Participants in the no privacy indicator condition were significant less likely to purchase from the corresponding sites. These results

indicate that people choose sites with better privacy policies when they are provided with privacy information in a more salient format.

***Hypothesis 2:** Participants in the privacy information condition will be more likely than those in the no privacy indicator condition to purchase from websites annotated with the four-green-boxes icon (the sites offering the best privacy policy). – Supported*

When shopping for batteries, participants in the privacy information condition made significantly more purchases from the four-green-box “high privacy” site (47.4%) than participants in the no privacy indicator condition (5.6%), $\chi^2 = 10.6$, $df = 2$, $N = 53$, $p = 0.005$. For the sex toy purchases, participants in the privacy information condition also made significantly more purchases from the high privacy site (33.3%) than participants in the no privacy indicator condition (0%), $\chi^2 = 16.1$, $df = 2$, $N = 64$, $p = 0.0003$.

5.2 Privacy Premium

***Hypothesis 3a:** Participants presented with prominent privacy information (those in the privacy information condition) will be more likely than those in the no privacy indicator condition to pay a premium to purchase from sites that have better privacy policies. – Supported*

As stated previously, this experiment was also designed to determine whether individuals would be willing to pay a premium for enhanced privacy protection (though it is important to note that the goal of the study was not to quantify a specific premium for the selected products). When comparing the no privacy indicator condition to the privacy information condition, we found statistically significant privacy premiums of roughly 60 cents for both products, as detailed in Table 3. Note that, to achieve a realistic design, we relied on actual merchants’ prices. In the course of the study, due to product constraints and fluctuating prices, the first result for the batteries was replaced with a slightly cheaper result, while the first result for the sex toy was

replaced with a slightly more expensive result. All of these changes were on the order of a few cents; we found no evidence that these changes impacted purchase decisions. Based on t-tests, we found that individuals shown privacy information were significantly more likely ($p < 0.001$ in both cases) to pay a premium to purchase from sites with better privacy policies. This effect was present for purchases of the privacy-sensitive item as well as the non-privacy sensitive item.

	Condition 1: No Privacy Indicator	Condition 3: Privacy Information	Premium	p Value
Mean Price: Batteries	\$14.64	\$15.23	\$0.59	0.0007
Mean Price: Sex Toy	\$15.26	\$15.88	\$0.62	0.00005

Table 3: t-test comparisons of mean prices paid in the no privacy indicator condition and the privacy information condition.

Hypothesis 3b: In the absence of prominent privacy information, people will purchase where price is lowest. – Supported

Examining the number of purchases made at the websites offering the lowest prices, we see that participants in the control conditions tended to purchase both items from the least expensive website, as denoted in Table 4.

	Purchases from lowest priced site – Batteries	Purchases from lowest priced site - Sex Toy
Condition 1: No privacy indicator	83.3%	80.0%
Condition 2: Irrelevant Information	75.0%	66.7%
Condition 3: Privacy Information	21.1%	28.6%
Chi² Value	17.3	13.1
p Value	0.0002	0.002

Table 4: Chi² test comparing the proportions of purchases made at the sites offering the lowest price for the batteries and the sex toy.

5.3 The Impact of Icons

Hypothesis 4: *Icons in the privacy information condition will affect purchase decisions more than icons in the irrelevant information condition. – Supported*

	Conditions		Fisher's Exact <i>p</i>
	Condition 2: Irrelevant Information	Condition 3: Privacy Information	
% of battery purchases made from sites with icons	25.0% n=4/16	77.8% n=14/18	<.002
% of sex toy purchases made from sites with icons	27.8% n=5/18	66.7% n=14/21	<.02

Table 5: A between-conditions comparison of the proportion of purchases made from sites annotated with icons. To test for significance between these proportions we used the Fisher's Exact test.

When comparing the proportions of purchases made from sites with icons, we found statistically significant differences in purchase patterns between participants who were presented with privacy indicators and those who were presented with indicators representing irrelevant information (Table 5). Unlike the former, participants who saw icons associated with irrelevant information were not likely to purchase from sites annotated with green box icons. This implies that our results can be attributed primarily to the actual privacy signals carried by the icons.

Additionally, as detailed in Table 6, we detected no statistically significant differences between the two control conditions' purchasing patterns. This table indicates that there was no significant difference between the no privacy indicator and irrelevant information conditions in terms of purchases made at sites with icons.

	Conditions		Fisher's Exact <i>p</i>
	Condition 1: No Privacy Indicator	Condition 2: Irrelevant Information	
% of battery purchases made from sites with icons	11.1% n=2/18	25.0% n=4/16	0.39
% of sex toy purchases made from sites with icons	16.0% n=4/25	27.8% n=5/18	0.46

Table 6: A between-conditions comparison of the proportion of purchases made at sites with icons in the irrelevant information condition and the corresponding sites in the no privacy indicator condition.

Similarly, when using a t-test to compare the average purchase prices of the no privacy indicator group with the purchase prices of the irrelevant information group, we did not find significant differences in the prices paid for each product, as shown in Table 7.

	Condition 1: No Privacy Indicator	Condition 2: Irrelevant Information	Premium	<i>p</i> Value
Mean Price: Batteries	\$14.64	\$14.69	\$0.05	0.64
Mean Price: Sex Toy	\$15.26	\$15.30	\$0.04	0.65

Table 7: Comparison of mean price paid for each product in the control conditions. Based on a t-test, there was no significant difference between the control conditions.

Figure 3 also clearly depicts the different purchase patterns between conditions. For both items, a greater percentage of purchases were made at four-green-box sites in the privacy information condition than in the no privacy indicator and irrelevant information conditions. The proportion of purchases made at sites with irrelevant icons is somewhat larger than the proportion made at sites with no privacy indicator – however, as noted above, this difference is not significant. More importantly, while we may have found that irrelevant icons motivate some participants to purchase from certain sites, we also found that the impact of such icons is far less than the impact of clearly annotated privacy information.




Figure 3: The percentage of purchases made for each product, by level of privacy, for each condition.

5.4 Other Results from the Exit Survey

In the exit survey, we asked whether the privacy icon (if seen) influenced participants' purchasing decisions, whether participants understood what the icon represented, whether they read any of the privacy policies, and whether those privacy policies influenced their purchasing decisions. Overall, the privacy icons served as an effective means for communicating privacy information. In the "privacy information" condition, 92% noticed the icons (95% CI = 74% - 99%), and 32% of participants read the privacy reports (95% CI = 15% - 53.5%). In the exit survey, 60% of the participants in the privacy condition reported that privacy information influenced the sites they *visited* and the sites from which they *purchased* (95% CI = 38.7% - 78.9%).

Providing visible privacy information heightened privacy awareness for the batteries, an

innocuous item. When asked in the exit survey about their battery purchase decision, participants in the privacy information were more likely to write in “privacy” or “privacy policy” when identifying the factor that most influence their decision than participants in the no indicator condition (32% vs. 0%; Fisher’s Exact $p = .001$).

These results indicate that once people were provided with salient privacy information, they chose sites they considered privacy protective; furthermore, they perceived differences in the level of privacy offered by sites annotated with the high, medium and low privacy icons.¹²

6. Limitations

Our study was not designed to establish whether the premium consumers were willing to pay for privacy should be interpreted in absolute terms (roughly 60 cents) or relative ones (roughly 4% of the price of the goods in question). However, the literature in the areas of marketing and behavioral economics suggests a number of plausible inferences, which further experiments could help us validate. These fields of research indicate that consumers' valuations are highly dependent on framing (Kahnman and Tversky, 1984), relative changes in price, and relative comparisons (Kahneman and Tversky, 1979; Chen *et al.*, 1998). As exemplified by Equation [1], participants in our experiment could assess the price charged by privacy protective merchants (for instance, \$15.14 for a set of batteries) against two other reference points: 1) the value of protecting their privacy; and 2) the price charged by other (less protective) merchants. Since the benefits of privacy protection are often uncertain and intangible (Acquisti and Grossklags, 2005b), we can expect that consumers are likely resort to relative comparisons when they try to determine the value of protecting their privacy, and therefore will assess privacy premiums in

¹² Additional results from the exit survey are discussed in the online Appendix, Section A2.5.

relative (percentage) terms. However, evidence also suggests that the willingness to pay for privacy is, ultimately, bounded (Grossklags and Acquisti, 2007). With regard to the prices charged by other merchants, the literature suggests that, for low-price products, consumers pay more attention to price premiums expressed in percentage terms. For high-price products, however, consumers are more likely to be affected by price premiums expressed in absolute dollar amounts (see Chen *et al.*, 1998). In the case of our relatively inexpensive user study products (batteries and sex toys), consumers may have perceived a 4% premium – around 60 cents – to be an acceptable amount to pay for privacy; however, if the price of the items increased, a percentage of 4% would become a larger and larger amount in absolute dollar terms - an amount capable of dissuading more consumers from paying for privacy. Combining these two lines of reasoning, we can expect the privacy premium to be a percentage of the absolute price of a good that decreases as that absolute price rises; furthermore, this premium is likely bounded in absolute dollar terms: a consumer purchasing a \$20,000 luxury item may be willing to allocate \$20 to make her transaction more confidential (this amount would represent *more* than the 60 cent premium in our scenario), but arguably not as much as \$800 (the equivalent to our 4% premium). Future research will be necessary to pinpoint the exact trade-offs between price and privacy sensitivity.

Lastly, while our participants made purchases using their own credit cards, the purchases were made in a laboratory setting following a specific experimental protocol. This setting is not necessarily reflective of ordinary search activity. To better determine the impact of prominent privacy information in a more natural setting, we plan to conduct a field study in which participants are asked to use Privacy Finder over a period of months. This may allow us to measure the impact of privacy information on people's everyday searches.

7. Implications and conclusions

The goal of this study was to determine whether the availability and accessibility of privacy information affects individuals' purchasing decisions. In turn, investigating that question allowed us to discuss whether businesses can leverage privacy protection as a selling point. Our study focused on what occurs when a search engine prominently displays privacy ratings for web sites. We used a modified version of Privacy Finder to display the privacy policies of certain online shopping sites in a fashion that, arguably, reduces the information asymmetry that separates merchants and customers *vis a vis* the usage of the customer's data. Our experimental approach was designed to investigate the impact of more prominent and accessible privacy information on consumer purchasing behavior in a realistic setting; this approach differs from the current method of making privacy practices information available via privacy policies.

Our results offer new insight into consumers' valuations of personal data and provide evidence that privacy information affects online shopping decision-making. We found that participants provided with salient privacy information took that information into consideration, making purchases from websites offering medium or high levels of privacy. Our results indicate that, contrary to the common view that consumers are unlikely to pay for privacy, consumers may be willing to pay a premium for privacy.

The results of this study suggest that future research needs to estimate the relationship between privacy and price sensitivity; in addition, researchers must work to achieve a more granular understanding of the behavioral and cognitive factors that influence a consumer's decision when privacy information is made more accessible. Our results also indicate that businesses may use technological means to showcase their privacy-friendly privacy policies and thereby gain a

competitive advantage. In other words, businesses may direct their policies and their information systems to strategically manage their privacy strategies in ways that not only fulfill government best practices and self-regulatory recommendations, but also maximize profits. Specifically, if the adoption of P3P increases, businesses protective of customer privacy may be able to attract consumers by posting their P3P policies and signaling “good” privacy practices. Survey data indicates that online consumers greatly value insight into what will be done with their personal information and how they can control those processes (Malhotra *et al.*, 2004). While consumers are often unable to control the practices of those who collect their information, they can control who they share their information with and the type of information they provide.

8. Acknowledgements

We wish to thank Christina Fong for her input and assistance in developing this research. We would also like to thank Rammaya Krishnan, Michael Smith, and Hal Varian for helpful comments. The work was supported by the Army Research Office grant number DAAD19-02-1-0389, entitled “Perpetually Available and Secure Information Systems,” and by the IBM Open Collaborative Research Initiative.

References

- Akerlof, G. 1970. The market for lemons: quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3): 488 – 500.
- Acquisti, A. 2004. Privacy in Electronic Commerce and the Economics of Immediate Gratification. *Proceedings of ACM Electronic Commerce Conference (EC '04)*. New York, NY: ACM Press, 21-29.
- Acquisti, A. and Gross, R. 2006. Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. *Proceedings of Privacy Enhancing Technologies Workshop (PET '06)*.
- Acquisti, A. and Grossklags, J. 2003. Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. *Proceedings of Workshop on Economics and Information Security (WEIS '03)*.

- Acquisti, A. and Grossklags, J. 2005a. J. Privacy and Rationality in Decision Making. *IEEE Security and Privacy*, 3(1): 26-33.
- Acquisti, A. and Grossklags, J. 2005b. Uncertainty, Ambiguity, and Privacy, *Proceedings of Workshop on the Economics of Information Security (WEIS '05)*.
- Acquisti, A. and Varian H. 2005. Conditioning Prices on Purchase History. *Marketing Science*, 24(3): 1-15.
- Ba, S. and Pavlou, P. A. 2002. Evidence of the effect of trust building technology in electronic markets: price premiums and buyer behavior. *MIS Quarterly*, 26(3): 243 – 268.
- Belanger, F., Hiller, J. S., and Smith, W. 2002. Trustworthiness in electronic commerce: the role of privacy, security, and site attributes. *Journal of Strategic Information Systems*, 11: 245 – 270.
- Benassi, P. 1999. TRUSTe: an online privacy seal program, *Communication of the ACM*, 42(2):56 – 59.
- Brown, M. and Muchira, R. 2004. Investigating the Relationship between Internet Privacy Concerns and Online Purchase Behavior. *Journal of Electronic Commerce Research*, 5(1): 62-70.
- Brunk, B. D. 2002. Understanding the privacy space. *First Monday*, 7(10). http://firstmonday.org/issues/issue7_10/brunk/index.html.
- Burst Media. 2009. Online Privacy Still A Consumer Concern. Conducted by Burst Media, February 2009. http://www.burstmedia.com/assets/newsletter/items/2009_02_01.pdf.
- Byers, S., Cranor L., Kormann, D., and McDaniel, P. Searching for Privacy: Design and Implementation of a P3P-Enabled Search Engine. *Proceedings of the Workshop on Privacy Enhancing Technologies (PET 04)*, 26-28.
- Camp, L. J. 2003. Design for trust. In *Trust, Reputation and Security: Theories and Practice*, Rino Falcone, editor. Springer-Verlang, 2003.
- CBS News. 2005. Poll: Privacy Rights Under Attack. <http://www.cbsnews.com/stories/2005/09/30/opinion/polls/main894733.shtml>.
- Chellapa, R. and Sin, R. G. 2005. Personalization Versus Privacy: An Empirical Examination of the Online Consumers' Dilemma. *Information Technology and Management*, 6(2-3): 181-202.
- Chen, S-F. S., Monroe K. B., and Lou, Y.C. 1998. The Effects of Framing Price Promotion Messages on Consumers' Perceptions and Purchase Intentions. *Journal of Retailing*, 74(3): 353-372.
- Consumer Union. 2008. Consumer Reports Poll: Americans Extremely Concerned About Internet Privacy. http://www.consumersunion.org/pub/core_telecom_and_utilities/006189.html.
- Cranor, L. 2002. *Web Privacy with P3P*. O'Reilly & Associates, Sebastopol, CA.
- Cranor, L., Guduru, P., and Arjula, M. 2006. User Interfaces for Privacy Agents. *ACM Transactions on Computer-Human Interaction*, 13(2): 135 – 178.
- Culnan, M. J. and Armstrong, P.K. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust. *Organization Science*, 10(1): 104-115.

- Culnan, M. J. 2000. Protecting Privacy Online: Is Self-Regulation Working? *Journal of Public Policy and Marketing*, 19(1): 20-26.
- Dinev, T. and Hart, P. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. *Information Systems Research*, 17(1): 61-80.
- Edelman, B. 2006. Adverse Selection in Online “Trust” Certifications. *Proceedings of the Workshop on the Economics of Information Security (WEIS '06)*.
- Gellman, R. 2002. Privacy, consumers, and costs - how the lack of privacy costs consumers and why business studies of privacy costs are biased and incomplete. <http://epic.org/reports/dmfprivacy.html>
- Gideon J., Egelman, S., Cranor, L., and Acquisti, A. 2006. Power Strips, Prophylactics, and Privacy, Oh My! *Proceedings of the 2006 Symposium on Usable Privacy and Security (SOUPS '06)*. http://cups.cs.cmu.edu/soups/2006/proceedings/p133_gideon.pdf.
- Grossklags, J. and Acquisti, A. 2007. When 25 Cents is too much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. *Proceedings of the Workshop on the Economics of Information Security (WEIS '07)*.
- Hann, I. H., Hui, K. L., Lee, T., and Png, I. 2007. Overcoming Information Privacy Concerns: An Information Processing Theory Approach. *Journal of Management Information Systems*, 24(2): 13-42.
- Harris Interactive. 2001. Privacy On & Off the Internet: What Consumers Want. Tech report. http://www.aicpa.org/download/webtrust/priv_rpt_21mar02.pdf.
- Hochhauser, M. 2003. Why Patients Won't Understand Their HIPAA Notices. Privacy Rights Clearinghouse. <http://www.privacyrights.org/ar/HIPAA-Readability.htm>.
- Hui, K.-L., Teo, H.-H., Lee, S.-Y. T. 2007. The Value of Privacy Assurance: An Exploratory Field Experiment. *MIS Quarterly*, 31(1):19-33.
- Huberman, B., Adar, E., and Fine, L. 2005. Valuating Privacy. *Proceedings of the Workshop on the Economics of Information Security (WEIS '06)*.
- Jensen, C. and Potts, C. 2004. Privacy policies as decision-making tools: an evaluation of online privacy notices. *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, 471-478.
- Jensen, C., Potts, C., and Jensen, C. 2005. Privacy Practices of Internet Users: Self-reports versus Observed Behavior. *International Journal of Human-Computer Studies*, 63: 203 – 227.
- Kahneman, D. 1973. *Attention and Effort*. New Jersey: Prentice-Hall.
- Kahneman, D. and Tversky, A. 1979. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47, 263-291.
- Kahneman, D. and Tversky, A. 1984. Choices, Values, and Frames. *American Psychologist*, 39, 341-350.
- LaRose, R. and Rifon, N. J. 2007. Promoting i-Safety: Effects of Privacy Warnings and Privacy Seals on Risk Assessment and Online Privacy Behavior. *Journal of Consumer Affairs*, 41(1): 127–149.
- Lebo, H. 2008. The 2008 Digital Trends Report: Surveying the Digital Future. USC Annenberg School Center for the Digital Future. http://www.digitalcenter.org/pages/current_report.asp?intGlobalId=19.

- Mai, B., Menon, N., and Sarkar, S. 2006. Online Privacy at a Premium. *Proceedings of the 39th Hawaii International Conference on System Sciences (HICSS '06)*.
- Malhotra, N., Kim, S. S., and Agarwal, J. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, 15(4): 336-355.
- McDonald, A. and Cranor, L. 2009. The Cost of Reading Privacy Policies. Forthcoming in *I/S: A Journal of Law and Policy for the Information Society*.
- McKnight, H. D. and Chervany, N. L. 2002. What Trust Means in E-Commerce Customer Relationships: An Interdisciplinary Conceptual Typology. *International Journal of Electronic Commerce*, 6(2): 35 – 59.
- McLeod, R. J. and Jones, J. W. 1986. Making executive information systems more effective. *Business Horizons*, 29(5): 29-37.
- Milne, G. R. and Culnan, M. J. 2002. Using the Content of Online Privacy Notices to Inform Public Policy: A Longitudinal Analysis of the 1998-2002 U.S. Web Surveys. *The Information Society*, 18(5): 345-359.
- Milne, G. R. and Gordon M. E. 1993. Direct mail privacy-efficiency trade-offs within an implied social contract framework. *Journal of Public Policy Marketing*, 12(2): 206–15.
- Miyazaki, A. and Krishnamurthy, S. 2002. Internet Seals of Approval: Effects on Online Privacy Policies and Consumer Perceptions. *Journal of Consumer Affairs*, 36(1): 28-49.
- Moor, J. H. 1997. Toward a Theory of Privacy in the Information Age. *Computers and Society*, 27-32.
- Moore, T. 2005. Do consumers understand the role of privacy seals in e-commerce? *Communication of the ACM*, 48(3): 86 – 91.
- Moore, T. T. and Dhillon, G. 2003. Do privacy seals in ecommerce really work? *Communication of the ACM*, 46(12): 265 – 271.
- Privacy & American Business (P&AB). 2005. New Survey Reports an Increase in ID Theft and Decrease in Consumer Confidence. Conducted by Harris Interactive, May 2005. <http://www.pandab.org/deloitteidsurveypr.html>.
- Privacy Leadership Initiative. 2001. Privacy Notices Research Final Results. Conducted by Harris Interactive, December 2001. <http://www.ftc.gov/bcp/workshops/glb/supporting/harris%20results.pdf>.
- Riegelsberger, J., Sasse, M. A., and McCarthy, J. 2005. The Mechanics of Trust: A Framework for Research and Design. *International Journal of Human-Computer Studies*, 62(3): 381-422.
- Rifon, N. J., LaRose, R., and Choi, S. M. 2005. Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures. *Journal of Consumer Affairs*, 39(2): 339–362.
- Rose, E. 2005. Data Users versus Data Subjects: Are Consumers Willing to Pay for Property Rights to Personal Information? *Proceedings of the 38th Hawaii International Conference on System Sciences (HICSS '05)*.
- Rubin, P. H. and Lenard, T. M. 2002. *Privacy and the Commercial Use of Personal Information*. The Progress & Freedom Foundation, Washington, DC, USA.

- Shapiro, C. 1983. "Premiums for high quality products as returns to reputations," *Quarterly Journal of Economics*, 98: 659 – 679.
- Shostack, A. 2003. Paying for privacy: Consumers and infrastructures. *Proceedings of the Second Annual Workshop on Economics and Information Security (WEIS '03)*.
- Smith, H. J., Milberg, S., and Burke, S. 1996. Information privacy: Measuring individuals' concerns about organizational practices. *MIS Quarterly*, 20(2): 167-196.
- Solove, D. J. A Taxonomy of Privacy. *University of Pennsylvania Law Review*, 154(3): 477.
- Spiekermann, A., Grossklags, J., and Berendt, B. 2001. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. *Proceedings of the ACM Conference on Electronic Commerce*, 38–47.
- Syverson, P. 2003. The paradoxical value of privacy. *Proceedings of the Second Annual Workshop on Economics and Information Security (WEIS '03)*.
- Stewart, K. A. and Segars, A. H. 2002. An empirical examination of the concern for information privacy instrument. *Information Systems Research*, 13(1): 36-49.
- Tang, Z., Hu, Y. J., and Smith, M. D. 2007. Gaining Trust Through Online Privacy Protection: Self-Regulation, Mandatory Standards, or Caveat Emptor. Available at SSRN: <http://ssrn.com/abstract=555878>.
- Taylor, C. R. 2004. Consumer Privacy and the Market for Customer Information. *Rand Journal of Economics*,. 35(4): 631-51.
- Tedeschi, B. 2002. Everybody talks about online privacy, but few do anything about it. *New York Times*, June 3, 2002, Section C, Page 6, Column 1.
- TRUSTe. 2006. Consumers have a false sense of security about online privacy – Actions inconsistent with Attitudes. Conducted by TNS, December 2006. http://www.truste.org/about/press_release/12_06_06.php.
- Turow, J., Feldman, L., and Meltzer, K. 2005. Open to Exploitation: American Shoppers Online and Offline. A Report from the Annenberg Public Policy Center of the University of Pennsylvania.
- Vila, T., Greenstadt, R., and Molnar, D. 2004. Why We Can't be Bothered to Read Privacy Policies: Models of Privacy Economics as a Lemons Market. In *The Economics of Information Security*, L.J. Camp and S. Lewis, eds., Kluwer, 143–154.
- Wathieu, L. and Friedman, A, 2005. An Empirical Approach to Understanding Privacy Valuation. *Proceedings of the Workshop on the Economics of Information Security (WEIS '05)*.
- Westin, A. and Harris Louis & Associates. 1990. Findings from the Survey - Consumers in the Information Age for Equifax Inc. 2,254 adults of the national public. <http://www.privacyexchange.org/iss/surveys/eqfx.execsum.1990.html>.
- Yerkes, R., and Dodson, J. 1908. The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology of Psychology*, 18: 459-482.
- Zhang, H. 2004. Trust-promoting seals in electronic markets: impact on online shopping decisions. *Journal of Information Technology Theory and Application*, 6(4): 29 – 40.