IoTMarketplace: Informing Purchase Decisions with Risk Communication

Shakthidhar Reddy Gopavaram
SICE
Indiana Univeristy
Bloomington, IN
sgopavar@iu.edu

Jayati Dev
SICE
Indiana University)
Bloomington, IN
jdev@iu.edu

Sanchari Das
SICE
Indiana University)
Bloomington, IN
sancdas@iu.edu

Jean Camp
SICE
Indiana University)
Bloomington, IN
ljcamp@indiana.edu

Abstract—**It is common for people to declare that online privacy is important, to indicate that it is valuable, and to simultaneously behave in a manner inconsistent with these expressed preferences. This discrepancy between users' concerns and their behavior has been explained by three factors: information asymmetry, bounded rationality, and psychological biases. The conflict between expressions of concern and purchase decisions is amplified as the Internet of Things brings the potential for real-time multimedia surveillance, even at home. But there is not empirical evidence that privacy or security influence purchase decisions about IoT devices. In this work, we design an interface for an Internet of Things (IoT) Marketplace that enables participants to make more privacy aware purchases by addressing three of the factors that are associated with the privacy paradox. We then conduct a between subjects experiment to test the effect of the interaction on product selection. The results from this experiment show that when participants are presented with an interface that addresses all three factors, they purchase IoT devices that are more privacy preserving even if they have to pay a premium to do so. The results also show that when participants are presented only with privacy indicators they do not consistently make privacy-preserving device choices if the psychological biases affecting users' decision making are not also addressed by interaction design.**

Index Terms—**Privacy, IoT, Bounded Rationality, Psychological Biases, Marketplace, Purchase Decision, Risk Communication**

1. Introduction

The past decade has seen massive growth in Internet of Things (IoT) devices, from fitness trackers to household items. There are already eight to ten billion devices; and that number is expected to double. It is predicted that there will be approximately 18 billion IoT devices by 2022 [27]. The current low standards of security and privacy correspond to high risk of information loss. While it is well-known that IoT devices collect information about people and their environment the extent of that compilation is not well understood [19], [31]. Information is collected to meet basic functional requirements, provide tailored services, perform analytics, or for targeted advertising. For example, Nest Secure is a user installed home security systems that collects information from multiple sensors in order to detect and alert its users about intruders [33]. Nest Secure also has a microphone whose purpose and operation is not well understood, and which were not well documented [73]. Similarly, activity tracking devices like Strava collect user activity information (step count, heart rate etc.) and relay it back in a structured fashion to help Strava users measure their activity levels [71]. Strava also provided detailed heat maps that identify users' home addresses and maps classified military operations [14], [37], [53].

Unlike general purpose computers, IoT devices do not require users to actively interact with them to provide services. Some IoT devices cannot reasonably implement a user interaction without connecting to an additional device with a screen. Most IoT information is passively collected, processed, and disseminated invisibly to the user without the option of ad-blockers or blocking trackers. Thus most people are unlikely to be aware about what information is being collected about them and their surroundings; and beyond that, how it is being used.

One place where a user can find information about an IoT device's data collection and usage practices is its privacy policy. These policies provide valuable information that could, in theory, address the problem of information asymmetry [58]. However, these privacy policies are notoriously unusable. Past research has shown that privacy policies are often too complex and hard to comprehend [77]. Over the past few years, researchers have proposed several solutions to shorten and simplify privacy policies in order to make them easier to read and comprehend [13], [44], [45]. Yet, even if the privacy policies were made comprehensible, the time to read them would be prohibitive. Finally, even if people read and understood these privacy policies, it would be hard for them to compare IoT devices based on the privacy they offer and select a device that meets their privacy preferences. This is because solving the information asymmetry problem does not address the other two sources of the privacy paradox: bounded rationality and psychological biases [7], [49].

Bounded rationality is the realistic contrast to the *homo economicus* model of rational human decision-making which assumes the ability to implement a strictly rational calculus to determine the economically optimal choice (Information asymmetry addresses the reality that accurate information is not always available for these calculations). Bounded rationality recognizes the cognitive limitations of humans and the costs of obtaining accurate information. Limits on attention, time,memory and information availability prevent people from accurately computing the costs and benefits associated with certain actions. In the case of IoT devices, making a privacy-aware choice would require reading all the privacy policies, determining how these interact with the set of devices in their home, and then evaluating the implications of the data collection and usage policies associated with the different IoT devices. With this analysis complete, the person must be able to accurately select a device that offers the best cost-benefit trade-off. The theory of bounded rationality illustrates why this theoretically economic rational model of privacy policies is not, in fact, feasible. In complex situations like these, where the cognitive costs of risk assessment are high, users take mental short-cuts to make speedy decisions [78]. This leads to the issue of psychological biases.

Even if people had the cognitive capacity to process all the information presented to them, it would still be difficult to make rational decisions. This is due to the psychological biases that affect human decision making [4]. For instance, past research has shown that the human tendency to attribute significantly more weight to immediate benefits and discount future risks applies to security and privacy risks [4], [30], [38]. So people purchasing an device might attribute a higher value to the price and convenience offered by the device when compared to the potential costs associated with insufficient privacy. Past research in behavioral psychology also shows that users' decision making is influenced by several other factors like framing, defaults, context, and mode of elicitation [5], [10], [17], [18], [32], [34], [41], [42], [47], [52], [64], [81].

We built a marketplace for IoT devices designed to address information asymmetry, bounded rationality, and psychological biases. We implemented this as an interaction layered on the Amazon Marketplace and tested under which conditions we observed privacy-aware behavior.

Our design addresses information asymmetry by providing users with cognitively simple indicators for privacy for each device. The marketplace was designed to address bounded rationality by providing a simple comparison indicating the privacy and the price of a device relative to other options. This enables buyers to compare between devices with minimal mental effort. The design also addressed a subset of the psychological biases that impinge privacy decision-making by the leveraging framing and defaults to manipulate the status-quo bias in order to mitigate discounting of risk. The IoTMarketplace is discussed in more detail in Section 3.

We then conducted a between subjects experiment to understand the impact of our design. The results from the experiment indicate that presenting users with risk indicators for privacy does not consistently nor significantly result in changes in decision-making. Conversely, When the risk indicators are present and the default for the purchase is set to high privacy, our participants consistently and significantly made privacy preserving purchasing decisions, even at a higher cost. Thus by addressing psychological biases though the use of defaults, marketplaces could nudge buyers into more privacy-aware decisions. Given that both privacy and security are associated with negative externalities, there is an argument for providing such a framing in the marketplace.

In Section 2, we identify possible sources of our results by beginning with previous research on the effect of framing, defaults, and other psychological biases affecting users' decision making. In Section 3, we provide the rational and design of all three versions of the IoTMarketplace. In Section 4, we detail the experiment design; please note that the experiment was subject to IRB review. In Section 5 we provide the results from the experiment, and then return to the related work for explanations. Finally, we conclude with a discussion on the possible implications of our findings in Section 6

2. Related Work

It is common for people to express high levels of concern about online security and privacy without this concern being reflected in their online behaviors; including purchasing behavior [49], [60]. This discrepancy between peoples' online behavior and their expressed concerns is called the *privacy paradox* [16]. There are three common explanations for this paradox: information asymmetry, bounded rationality, and psychological biases in decision-making.

2.1. Information Asymmetry & Bounded Rationality

Information asymmetry occurs when a seller has information about the type or quality of a good that is unknown to the buyer [9]. In the case of privacy paradox, information asymmetry results from consumers' uncertainty about privacy risks associated with devices. For instance, sellers have information about data collection and usage practices associated with the services that they provide. All or some of this information is made available to the users in the form of privacy policies or user agreements. However, studies have shown that privacy policies are ineffective at communicating risk to the user [39], [77]. This is because very few people read these policies as they often span multiple pages and require a significant investment of time [57], [62]. In fact, a study conducted by McDonald et al. estimated that if all internet users in the United States were to read the privacy policy of every new website that they visited then they would spend about 54 billion hours reading privacy policies [57].

In addition, people rationally know that privacy policies are often subject to change without notification so reading them once would be inadequate [79]. Even if people were to read these privacy policies it would be hard for them to understand the contents as most of them contain complex or nuanced legal terminology [12], [40]. Jensen et al. evaluated the readability of privacy policies for a set of popular websites to find that most privacy policies were beyond the comprehension of people who had less than or equal to a high school education [40]. In some cases, people required the equivalent of a postgraduate education to understand the privacy policies [40]. Therefore, most people not having read or understood the contents of the privacy policies, usually make decisions with incomplete information. This leads to inaccurate estimates of costs and benefits which eventually results in choices that are not in their best interest (thereby creating the privacy paradox).

The bounded rationality of users also adversely affects their decision making [23]. For instance, even if people were to the read the privacy policies and understand its contents, it would be hard for them to accurately estimate the risks associated with different data collection and usage practices. This is because people have a limited ability to process information [30], [48]. Furthermore, to compensate for their bounded rationality people will take cognitive short-cuts to make quick decisions [78]. Such decisions will not accurately represent their preferences. Therefore, reducing information asymmetry and addressing bounded rationality requires clear communication of risk so that buyers can easily determine which devices are associated with more or less privacy risks. In the presence of privacy risk, such communication can be informed by risk communication practices.

2.2. Risk Communication

There is a significant body of research seeking to support privacy aware decision making by communicating the existence of privacy threats. For example, a study conducted by Tsai et al. showed that provision of privacy risk communication can decrease the discrepancy between users' expressed privacy concerns and their behavior [75]. In that previous work, Tsai and her coauthors asked participants who expressed high levels of concern about online privacy to select vendors of two products using a search engine. When the search engine only provided participants with links to the vendors' webpages and the price of the product, participants did not pay any attention to the vendors' privacy policies. They systematically chose the vendor that offered the lowest price. However, when the search engine provided participants with a simple privacy rating along with a link to the summarized privacy policy, participants paid a premium to buy products from vendors that offered higher levels of privacy. Studies conducted on mobile platforms have shown similar results [8], [64], [67], [80].

In the IoT, cost benefit trade-offs for disclosing personal information to obtaining personalized services can be quite complex [2], [50], [72]. At the same time, users indicate that it is difficult or even impossible to obtain privacy and security information about IoT devices [25]. Recent work has considered different approaches to resolving the specific issue in the consumer marketplace. Loi et al. proposed a systemic method to identify security and privacy shortcoming of an IoT device [55]. In their work, they conduct a series of tests to evaluate the devices along four dimensions: confidentiality, integrity, access control and reflection attacks. The results from the tests were then used to rate devices as secure, moderately secure or insecure. Blythe and Johnson proposed a protocol to develop a consumer security index for IoT devices [15]. There protocol consisted of identifying (1) security features that IoT devices must provide and (2) consumer preferences for the disclosure of security and privacy features that IoT devices provide. Both approaches can be aligned with the results of this work; particularly as sources for ratings or as simple information for the first link from the ratings.

In order to address information asymmetry in our study, we calculated a risk rating based on the privacy policies associated with the device and the corresponding manufacturer's app. The generation of aggregate privacy ratings based on the devices' privacy policies is not the focus of our work. However, since the success of our work is dependent on our ability to generate such ratings we include the following related work. Before the ubiquity of machine learning to process natural languages, the World Wide Web Consortium developed the Platform for Privacy Preference (P3P) to make privacy policies more usable. P3P is a machine readable language that enabled companies to state data collection and usage policies. The idea was to build P3P user agents that enabled people to set their privacy preferences and identify websites that met these preferences [20]. The two popular P3P user agents were Privacy Bird and Privacy Finder. While the former displayed a red, yellow, or green bird to indicate if the a website's privacy policies met the user's stated preferences the latter was a P3P enabled search engine that generated privacy ratings on a 5-point scale for each of its search results. User studies showed that both privacy bird and privacy finder improved users' privacy practices [75], [80]. P3P was not adopted by companies on a larger scale, which led to the eventual its eventual demise.

Wilson et. al. created and shared a corpus of 115 privacy policies with manual annotations for fine grained data collection and usage practices [82]. The availability of this corpus led to the development of Machine Learning models to automate the analysis of privacy policies [35], [54], [59], [83]. One of the machine learning approaches that directly informed our work was the framework proposed by Harkous et al. [35]. Their proposed framework uses a hierarchy of neural network classifiers to identify high-level and fine-grained data collection and usage policy details. They demonstrate the utility of their framework by building an application that given a link to the privacy policy could automatically generate disconnect icons [1] based on preset rules. Using the same framework a similar application can be built to generate aggregate privacy ratings for privacy polices based on preset rules or user preferences.

In this work, we manually generated privacy ratings for each device based on its privacy policy. Specifically, we evaluated the privacy policies based on five factors: data collection, data usage, control, unauthorized use, and improper access. For each factor, we assigned a score between one to five. The overall privacy score/rating was derived by calculating the average score across the five factors. Here we use positive framing so a higher score implies more privacy and a lower score implies less privacy. We provide more details on our choice of framing later in this section.

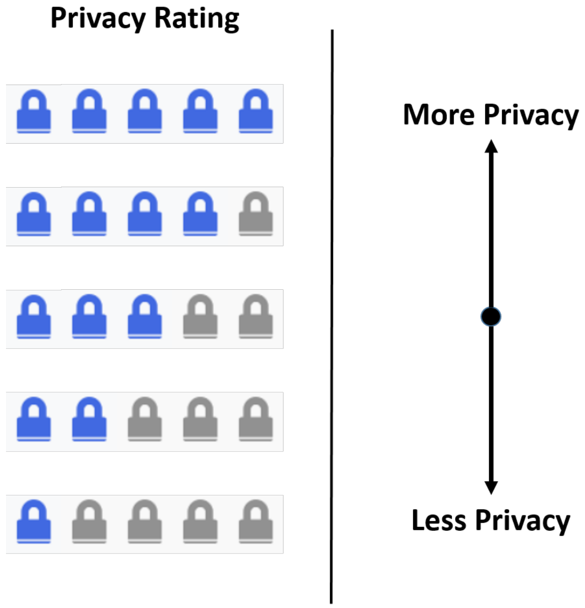

Figure 1. Privacy rating communicated using the lock icon. More locks imply more privacy.

The padlock icon was used to communicate the privacy rating. The choice of icon was primarily informed by the findings of the study conducted by Rajivan et al. [64]. In their study, Rajivan et al. investigated the effectiveness of multiple framing mechanisms and icons and found that positively framed risk indicators presented using the padlock icon were most effective. An illustration of privacy indicator used in this experiment is shown in Figure 1.

In summary, the experiment addressed information asymmetry and bounded rationality by building on previous work and used the standard lock icon to communicate this using positive framing with a tested icon.

2.3. Psychological Biases

Information asymmetry and bounded rationality are two of the many factors that contribute to the privacy paradox. There are other factors that impinge on peoples' preferences. This is because people don't have static predetermined preferences about their privacy. The lack of fixed preferences occurs even in intimate and high-value decisions such as organ donations and support for immunizations thus it is unsurprising that privacy is no exception [42]. Preference reversal occurs when decision-makers preferences are constructed at the time of elicitation rather than being predetermined [29], [43], [61], [70]. Preference reversal refers to an aggregate effect when changes in presentation of information causes the preferred option to change in the aggregate. The classic case of preference reversal is (applicably) in risk communication. The option of killing 10,000 people to save 90,000 is systematically rejected; while the option of an intervention to save 90,000 out of 100,000 people is systematically found to be acceptable despite the fact that these describe exactly the same intervention. In this work, we seek to create an interaction that results in more attention and valuation being paid to risk on the basis that "preference reversals in decision making under risk are accompanied by changes in attention to different attributes" [46].

Fluidity in preferences, both individual and aggregate, can be impinged by framing, defaults, context, trade-offs and the mode of elicitation. Manipulation of preferences to influence people's behavior at the time of decision is often referred to as nudging, to indicate that the changes are significant in the aggregate but not deterministic at the individual level. In this section, we will discuss some of the factors that influence people's decisions and then previous work on nudging in the security and privacy space.

This phenomenon where people attribute a higher value to the items they possess when compared to the items that they don't possess is called the endowment effect. In the design of our interaction we use defaults to communicate that the participant has privacy as a default, or that the lowest price is a default. Use of defaults is aligned with the original work on the endowment effect. Specifically in an early influential work by Knetsch, participants were either randomly provided with a mug or a chocolate bar and were later offered the option to exchange one for the other. The results showed that participants preferred to keep the item they were initially given despite there being no exogenous reason for this preference [47]. In other words, participants who were initially given a mug did not want to exchange it for a chocolate bar and participants who were initially provided with a chocolate bar did not want to exchange it for a mug. As a consequence of the endowment effect, peoples' maximum Willingness to Pay (WTP) for a good they don't possess is different from their minimum Willingness to Accept (WTA) for same good if they possess it. So, the value an individual attributes to a good changes based on if they perceive that they already have it; in general a good we are selling is perceived to have a higher value than when we are purchasing the same good [6], [34], [47].

[1] Disconnect icons are privacy icons that were developed a Mozilla led working group. The aim of these icon was to make it easy for people to understand the terms of the privacy policy and to communicate data collection and usage practices [24].

In our design, as discussed below, we hypothesized that people presented with lesser privacy and a lower price as a default will keep that preference, while those with high privacy as a default will similarly systematically choose privacy even when the price is higher. In previous experimental research in privacy WTP and WTA was defined as an individual's willingness to pay to protect against information disclosure versus an individual's willingness to accept compensation for disclosing information [34]. Influencing our design, in 2013 Acquisti et al. conducted an experiment in which the participants were either endowed with a $10 gift card or a $12 gift card [6]. The $10 gift card was anonymized i.e. purchases made though the gift card would not be linked back to the participants. On the other hand, the $12 gift card was an identified one i.e. purchases made though the gift card would be linked back to the participants. The participants who were endowed with the $10 gift card could agree to disclose their purchase information and receive an additional $2 (exchange the $10 gift card for the $12 one). Similarly, participants who were endowed with the $12 gift card could protect their purchase information from disclosure by paying $2 (exchange the $12 gift card for the $10 one). The results showed that more participants rejected the $2 offer to disclose their information when compared to the number of participants that were willing to pay the additional $2 to protect their information. These results clearly show a gap between WTA and WTP for privacy.

Another factor that influences decision-making in the presence of risk is the framing of risk mitigating action as a cost that must be paid to avoid risk or a benefit that enables risk avoidance. In the case of security and privacy behaviors, the options are the increased risk of information disclosure or the benefit of information protection. The communication of privacy as a benefit is described as positive framing and the communication of privacy as a risk is defined as negative framing. Positive framing is generally supported by work in the psychology of security to help people make more risk averse decisions [5], [64]. For example, Chen et al. conducted multiple experiments to determine the effect of framing on user choices. In these experiments they repeatedly presented participants with a list of 6 apps along with their privacy score that was either positively framed or negatively framed and they asked the participants to select 2 apps out of the 6. The results showed that participants paid more attention to the privacy score, had a better understanding of the score and made more risk averse decisions under positive framing when compared to negative framing [17].

In addition to information asymmetry, endowment effect and framing, status-quo bias also influences user behavior [65], [66], [68]. Status-quo bias consists of two primary components: (1) strong preference for the current state of affairs and (2) a strong preference for not taking any action also known as omission bias [65]. This strong preference for the current state of affairs is due to loss aversion [65], [74] . A change from status-quo implies that people would lose somethings while gaining other things. Since people are loss averse, they tend to attribute a higher weight to losses when compared to gains which explains their strong preference for the status-quo [65]. It must also be noted that status-quo bias is only present when people have to take an action. When there is no action involved people don't exhibit a status-quo bias [65]. Furthermore, people react more adversely to negative outcomes caused by taking an action as supposed to taking no action even if the negative outcomes are equivalent in both cases [65]. This fear of potential regret from taking an action compels people to not do anything when there is a potential negative outcome associated with an action. This fear of regret is one of the primary reasons for omission bias. Omission bias has been observed in situations where websites or other services ask users for their consent to collect personal information or for sending notifications and they check the consent box by default. In such cases more people are likely to take no action and stick to the default [41], [52].

Finally, it is important to note that people who have strongly held believes or attitudes towards issues are less susceptible to the biases mentioned above. For example, a survey conducted by Wilson et al. found that framing had no effect on peoples' decisions concerning termination of a child at risk of having heamophilia [56]. More specifically in the case of privacy, an experiment conducted by Lai et al. found that people who expressed high concerns for privacy were less susceptible to being influenced by status-quo when compared to people who were less concerned about privacy [52].

3. Marketplace Design

We designed an experiment to see if it is possible to create an interface for an IoTMarketplace that enabled users to easily compare devices based on their privacy while mitigating inherent biases or psychological biases, resulting in risk-averse decision making. In all versions of the marketplace, devices that offered a higher level of privacy were priced higher than those that offered a lower level of privacy. This was done to see if people would purchase devices that offered better privacy even if they charged a premium. The three versions of the IoT marketplace are as follows: (1) *Control Version*, (2) *High Privacy Default Version*, and (3) *Low Privacy Default Version*.

The *control version* of the marketplace contains no information pertaining to the level of privacy offered by a device. Given that all the devices in this experiment have similar features we expected the participants using this version of the marketplace to primarily make their decisions based on the price of the device.

The *high privacy default version* of the marketplace addresses all three factors associates with the privacy paradox (information asymmetry, bounded rationality, and psychological biases). Therefore, we expect participants using this versions of the marketplace to make more privacy preserving decision when compared to the participants using the control version.

H1 : More participants using the *high privacy default version* of the marketplace will purchase devices with a higher privacy rating when compared to the participants using the *control version*.

The *low privacy default version* of the marketplace addresses two out of the three factors associated with the privacy paradox. Specifically, this version of the marketplace does not address users' psychological biases. Therefore, we expect participants using this version of the marketplace to make fewer privacy preserving decision when compared to the participants using the *high privacy default version*.

H2 : More participants using the *high privacy default version* of the marketplace will purchase devices with a higher privacy rating when compared to the participants using the *low privacy default version*.

The rest if this section will discuss the design choices for each of the three versions in detail.

3.1. Control Version

The *control version* of the marketplace is basically used as a baseline to compare the effects of the interventions introduced in other versions of the marketplace. The control version of the marketplace provides users with the images, price and description for each device that is on sale. In this version of the marketplace, we do not communicate privacy risk to the user. The devices are presented as a single list and the order of devices is randomized for each and every participant to alleviate any biases caused by order. A screenshot of the control version of the marketplace is show in Figure 2. The design for this version of the marketplace was influenced popular e-commerce websites where products are often presented as a list with images, price, and a short description.

3.2. High Privacy Default Version

As noted in Section 2 the lack of effective risk communication could lead to people not considering privacy while making their purchase choices [75]. For the risk indicators to be effective, they need to be simple, concise, and easy to understand. As flooding users with a lot of information is also ineffective since users have a limited capacity for processing information [69]. Therefore, in this version of the marketplace we provide users with cognitively simple positively framed aggregate risk information using the padlock icon. The choice of framing and icon was primarily informed by the findings of a study conducted by Rajivan et al. [64]. In that study, after investigating the efficacy of multiple framing mechanisms and icons, they found that positively framed risk information presented on a 5-point scale using the padlock icon to be most effective. In addition to the aggregate risk rating, we also provided users with a short description of the rating. The illustrations of the aggregate risk score and the description can be found in Figure 3.

The focus of this study is on the behavioral aspects of purchase decisions for IoT devices. The generation of aggregate risk ratings is not the focus of this study. However, since our work is impinged by the ability to generate such ratings we reference the following previous works here. Harkous et al. proposed a framework for automating the analysis of privacy policies [35]. This framework could be used to generate aggregate privacy ratings based on the privacy policies. Prior to this applications like privacy bird generated privacy ratings on a 3 point scale based on user preferences [21], [22]. In this approach the application retrieved machine readable privacy policies and compared them to users' stated preferences. In our work the privacy ratings were manually generated by researchers.

Our goal was to select indicators to communicate privacy risk and in doing so reduce information asymmetry, enabling participants to make more privacy preserving purchases when compared to people with no indicators for privacy risk. However, multiple privacy studies have found that people tend to attribute significantly more weight to immediate short term benefits when compared to potential future losses [3], [4], [30], [38]. So they would likely be willing to give up privacy to save some money now. In order to alleviate such prejudice against privacy during decision making, we employ status-quo bias in our design to promote risk-averse behavior. There are two primary components to status-quo bias.

1) Loss Aversion: When performing an action that leads to change in state, people tend to attribute a higher weight to the losses caused by a change in state when compared to the gains.
2) Omission Bias: if there are potential negative outcomes associated with the change in state then people are more likely to stick to the status-quo in order to avoid the regret caused by the potential negative outcome.

We incorporated both *loss aversion* and *omission bias* in our design by presenting a default state which favors privacy over monetary gain.

In order to generate *loss aversion* and *omission bias* among users we do the following: (1) We categorized our devices based on their privacy rating, (2) ordered the categories in decreasing order of their privacy, and (3) we set the highest privacy category as the default and provided users with the ability to switch between categories by clicking on the respective tabs. A screenshot of our design is shown in Figure 4 (a). All tabs contained information about the privacy rating and the starting price for that category. This was done to focus users' attention on what they would gain and lose when switching between categories. So when a user who starts of with the highest privacy category as a default switches to a lower privacy category he/she would lose privacy but gain money (will save money as devices in a higher privacy category are priced higher). Based on the theory of *loss aversion* users would weigh their loss

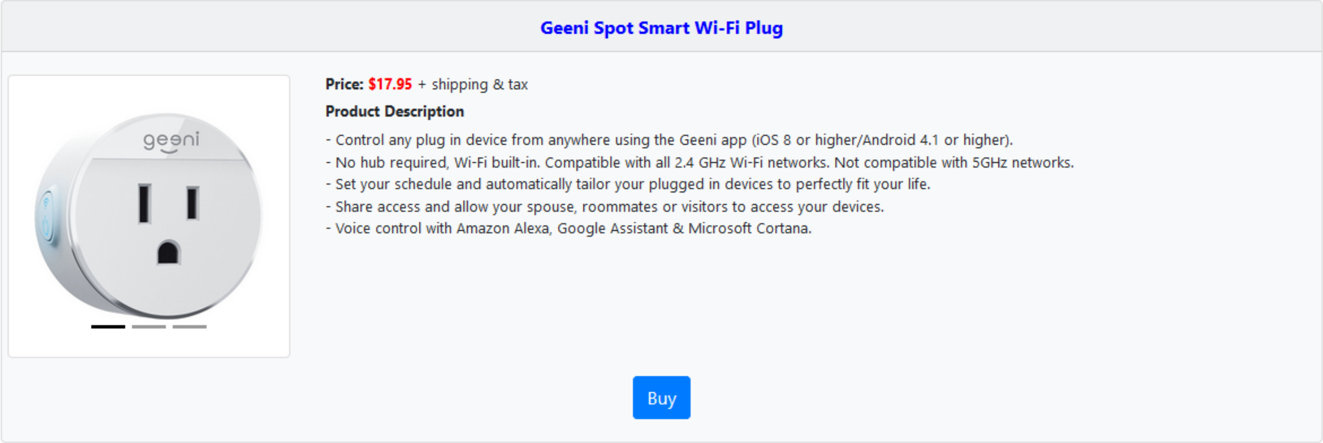

Geeni Spot Smart Wi-Fi Plug

Price: $17.95 + shipping & tax

Product Description

- Control any plug in device from anywhere using the Geeni app (iOS 8 or higher/Android 4.1 or higher).
- No hub required, Wi-Fi built-in. Compatible with all 2.4 GHz Wi-Fi networks. Not compatible with 5GHz networks.
- Set your schedule and automatically tailor your plugged in devices to perfectly fit your life.
- Share access and allow your spouse, roommates or visitors to access your devices.
- Voice control with Amazon Alexa, Google Assistant & Microsoft Cortana.

Figure 2. The control version of the IoTMarketplace does not present users with any privacy information. It just provides people with price and product description.

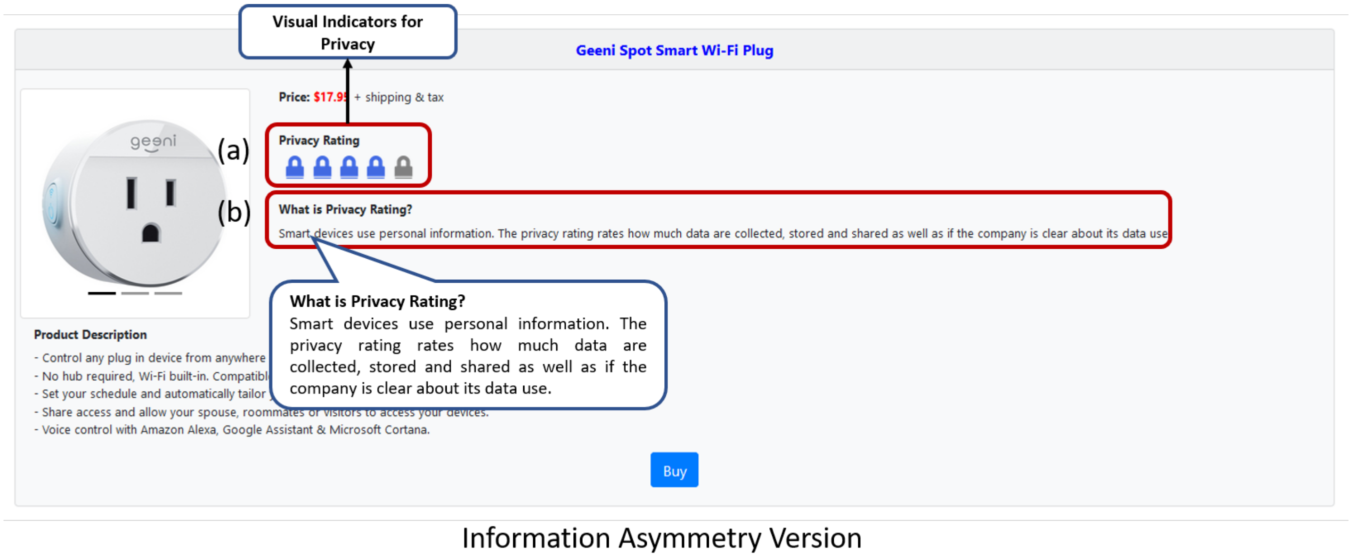

Visual Indicators for Privacy

Geeni Spot Smart Wi-Fi Plug

Price: $17.95 + shipping & tax

(a) **Privacy Rating**

(b) **What is Privacy Rating?**

Smart devices use personal information. The privacy rating rates how much data are collected, stored and shared as well as if the company is clear about its data use

What is Privacy Rating?
Smart devices use personal information. The privacy rating rates how much data are collected, stored and shared as well as if the company is clear about its data use.

Product Description

- Control any plug in device from anywhere
- No hub required, Wi-Fi built-in. Compatibl
- Set your schedule and automatically tailor
- Share access and allow your spouse, roommates or visitors to access your devices.
- Voice control with Amazon Alexa, Google Assistant & Microsoft Cortana.

Figure 3. (a)Positively framed privacy rating illustrated using the padlock icon. (b) A short description explaining the privacy rating.

in privacy higher than their gain in monetary saving. Furthermore, loss in privacy could have many adverse effects like financial loss, social embarrassment etc. Since people feel more regret from negative outcomes caused by an action when compared the same outcome occurring due to inaction [65]. We hypothesize that people will exhibit *omission bias* and would avoid purchasing a device from a lower privacy category to avoid potential regret.

3.3. Low Privacy Default Version

Similar to the *high privacy default version*, in this version of the marketplace all devices are categorized by their privacy rating. Users could also switch between categories by the clicking on the appropriate tab. However, in this version of the marketplace the categories were ordered in the increasing order of their privacy and the lowest privacy category was set as the default (as show in Figure 4 (b)). By changing the default and the order of categories we change effects of *loss aversion* and *omission bias*. Now when a user switches from the default category to a category with a higher privacy rating he/she loses money and gains privacy. So according to the theory of *loss aversion* users will attribute a higher weight to their monetary losses when compared to the gains in privacy. Furthermore, potential regret from not being able to have additional cash to spend on other purchases could lead to *omission bias*. Therefore, more people using this version of marketplace are likely to buy a device with a lower privacy rating in order to save money.

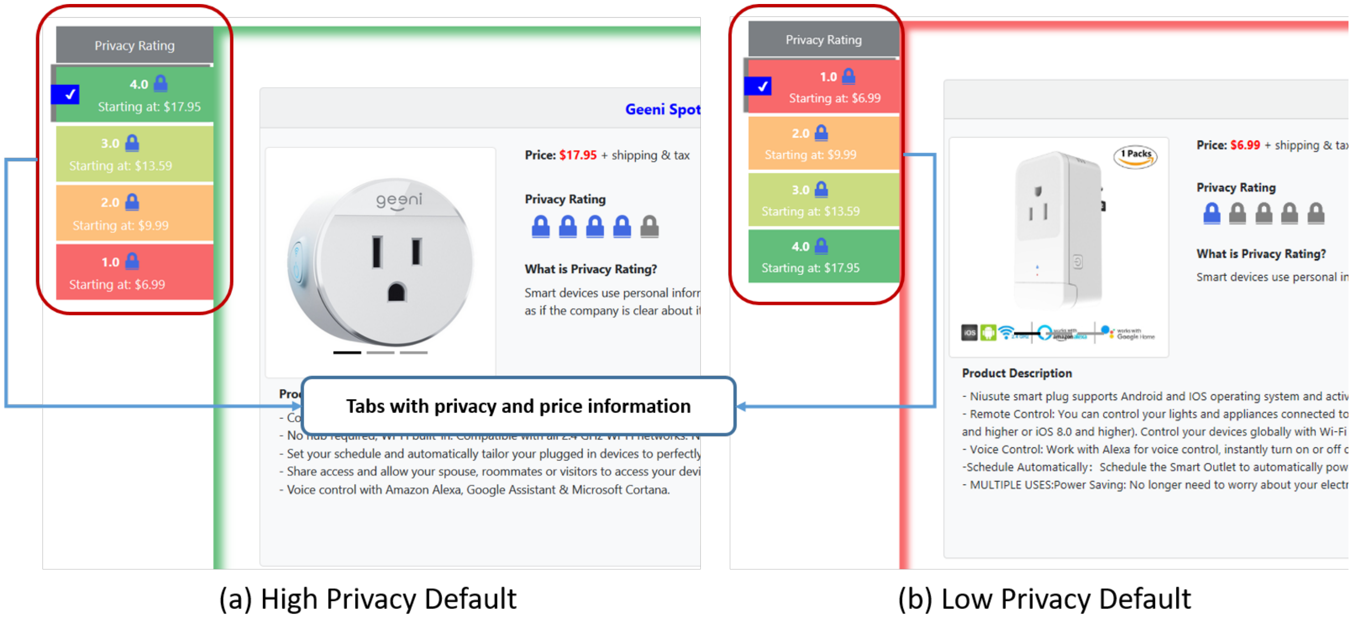

Smart devices use personal inforr
as if the company is clear about i

- Co
- No hub required, Wi-Fi built-in. Compatible with all 2.4 GHz Wi-Fi networks. N
- Set your schedule and automatically tailor your plugged in devices to perfectly
- Share access and allow your spouse, roommates or visitors to access your devi
- Voice control with Amazon Alexa, Google Assistant & Microsoft Cortana.

Smart devices use personal in

- Niusute smart plug supports Android and IOS operating system and activ
- Remote Control: You can control your lights and appliances connected to
and higher or iOS 8.0 and higher). Control your devices globally with Wi-Fi
- Voice Control: Work with Alexa for voice control, instantly turn on or off c
-Schedule Automatically: Schedule the Smart Outlet to automatically pow
- MULTIPLE USES:Power Saving: No longer need to worry about your electr

(a) High Privacy Default

(b) Low Privacy Default

Figure 4. (a) *High privacy default* has the category with privacy rating 4 as the default and the tabs are ordered in the descending order of privacy (or ascending order of price) (b) *Low privacy default* has the category with privacy rating 2 as the default and the tabs are ordered in the ascending order of privacy (or descending order of price)

4. Experiment Methodology

Our primary goal was to investigate if people using different versions of the IoT marketplace made different purchase choices. So we conducted a between subjects experiment with one control group and two experimental groups. While the participants in the control group used the *control version* of the marketplace, participants in the two experimental groups used either the *high privacy default version* or the *low privacy default version*. Here we name the groups after the version of the marketplace they use. The groups are: Control Group, High Default Group and Low Default Group.

Everyone that agreed to participate in the study, irrespective of the group they were assigned to, were initially presented with a set of instructions that told them how to purchase the device they selected. These instructions were purely mechanical and did not contain any information that would prime them for privacy. After reading the instructions people were allowed to move on to the next stage of the experiment where they were presented with three categories of products: Home Security cameras, Fitness Trackers, and Smart plugs. These categories of products were presented to the participants in a sequence i.e. they were first presented with a list of home security cameras and after selecting a device from that list they were presented with items from the next category. The first two categories were used to help participants familiarize themselves with the interface of the marketplace. Participants only bought products from the third category (Smart Plugs).

For the smart plug category, once a participant selected a device they wanted to purchase we redirected them to the product listing on Amazon where they made their purchase using the $25 amazon gift card that was provided to them. They were allowed to keep the device they purchased and any cash that was left on the Amazon gift card after making the purchase as compensation for participating in the study. After completing the purchase, the participants were asked to complete a short survey which included questionnaires about demographics, purchase decisions, expertise and privacy concerns.

For each category, we presented participants with a list of 8 devices. These were all real products that had a listing on Amazon. Specifically for the smart plug category, we browsed through a list of smart plug devices that were priced under $25 and manually analyzed their privacy policies to generate an aggregate risk score or privacy rating. We then selected a list of 8 devices such that (1) all the devices provided the same features, (2) all the devices were compatible with Alexa, Google Home, iPhone and Android devices, and (3) all devices that had a higher privacy score were priced higher. If all the products were priced the same and had the same features then people would obviously choose to purchase products with a higher privacy rating. Here we wanted to see if people would pay a higher price for privacy.

Participants for this study were recruited though ads on classifieds, flyers and email blasts. We recruited 20 participants per group. So we had a total of 60 participants. The design was approved by the Institutional Review Board(IRB).

5. Results

We found that participants in the high privacy default condition were more likely to purchase products with a higher privacy rating when compared to participants in the control and the low privacy default conditions. This indicates that people are more likely to pay a premium for privacy when privacy information is made salient and the higher privacy category is set as the default. Furthermore, despite having visual indicators for privacy, people in the low privacy default condition made purchases that were similar to the participants in the control group i.e. a significant number of participants in both groups purchased products with the lowest privacy rating. This demonstrates that making privacy information more accessible is not sufficient to make participants purchase more privacy preserving products. The design of the interface should also address the psychological biases that are an inherent component of human decision-making.

5.1. Demographics

All participants were over 18, and the sample skewed younger. Fifty-six percent of the participants were between 18-25 years old; 31.68% were between 25-35 years old; and 11.67% of the participants were older than 35. Out of the 60 participants 53.34% were men and 46.66% were women (32:M, 28F).

5.2. Time to Decision

We recorded the amount of time each participant required to select a smart plug device to purchase. Figure 5 compares the distribution of decision times of participants in different experimental groups. On average participants in the control and *low privacy default* conditions took 3.086 and 3.56 minutes, respectively, to select a device. While the average decision time of participants in the *low privacy default* group is higher than that of participants in the control group, the median for the *low privacy default* condition is less than that of the control group (median *low privacy default*: 2.80 mins *control*: 3.21 mins). Participants in the *high privacy default* condition on average took about 4.68 minutes to select a device to purchase. This was 1-1.5 minutes more than the participants in the control and the *low privacy default* conditions.

We conducted a one sided t-test to see if the results were statistically significant. The results from the test show that the decision time between the *high privacy default* and *low privacy default* is not statistically significant (t = 1.466, df = 32.564, p-value = 0.076). The decision time between the *high privacy default* and *control group* is statistically significant (t = 2.2013, df = 28.589, p-value = 0.018). Finally, the results between the *low privacy default* and *control group* are not statistically significant (t = 0.88393, df = 36.311, p-value = 0.1913).

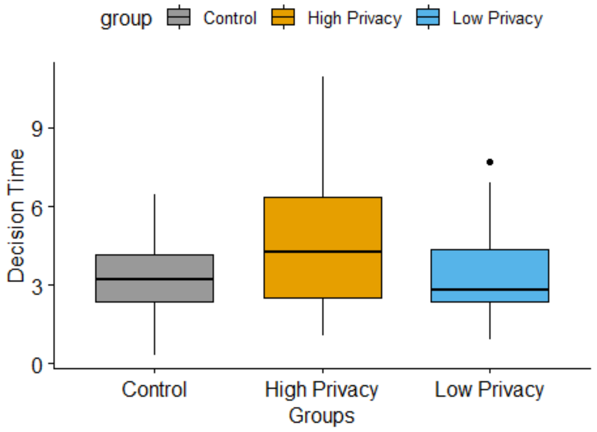

Figure 5. Box plot comparing the distribution of decision times of participants using the *control*, *low privacy default* and the *high privacy default* versions of the marketplace.

While the control version of the marketplace presented the devices as a single list, the other two versions of the marketplace categorized them into groups. So the participants using the *high privacy default* and the *low privacy default* versions of the marketplace had to switch between tabs to look at devices within privacy categories other than their default. Recall that the defaults would be immediate visible without switching tabs, and be the first listed tab. As the categorization of devices was the only difference in interaction between the control and the two experimental groups, it is reasonable to assert that this variation contributed to the difference in time to decision between the control and two experimental groups. Additionally, more participants in the *high privacy default* group viewed devices within all categories when compared to that of the *low privacy default* group. (This is discussed in more detail in Section 5.3). This could have contributed to the difference in decision time between the two groups. Fewer participants in the *low privacy default* group viewing devices within all privacy categories could also explain the lower median.

The difference in decision time between the control and experimental groups shows that people simply did not pick the first listing. They carefully considered the options presented to them before making their decision. However, the factors that most influenced participants decisions varies between groups.

5.3. Tabs Viewed

For both the *high privacy default* and the *low privacy default* conditions, the devices were divided into categories based on their privacy ratings. Each category consisted of two devices and participants could switch between the categories by clicking on the respective tabs. Here we report the categories that participants viewed before selecting a smart plug device to purchase.

Fifty-five percent of the participants in the *low privacy default* condition viewed the devices in all 4 categories before making a decision. The remaining 45% of the participants viewed 3 or fewer categories before making their decision. Participants that viewed 3 or fewer categories only explored the devices within the lower privacy categories i.e. they did not view the devices in the highest privacy category. In some cases, participants explored categories with a higher privacy rating for fitness devices and security cameras but only viewed the default category (lowest privacy category) before selecting a smart plug device. Note that both the price range and the privacy rating is visible on the tabs. This could indicate an overall disinterest in privacy, or an unwillingness to spend more on devices regardless of the privacy rating when the default presented offered a lower price.

Seventy-five percent of the participants in the *high privacy default* condition viewed the devices in all 4 categories before making a decision. The remaining 25% of the participants made decisions after viewing only the highest privacy category. Some of the participants in the 25% viewed devices in the lower privacy categories for the fitness trackers and security cameras but only viewed the highest privacy category for smart plug devices. The results show that participants in the *high privacy default* category were less likely to purchase devices with a lower privacy rating.

A significant portion of participants within the *high privacy default* and the *low privacy default* conditions made purchase decisions without viewing all devices. This implies that the decision made by these participants was to a great extent based on the privacy rating and the price of the device. By not viewing the devices within other categories these participants prevented themselves from being influenced by attributes (like appearance) associated with products within other categories. Therefore, by categorizing devices based on their privacy rating and setting a high privacy default, we can make participants attribute a higher value to the privacy offered by the device.

5.4. Descriptive Statistics

The bar chart in Figure 6 compares the distribution of privacy ratings for the products purchased by participants in the three experimental conditions. The chart shows that a lot more participants in the *high privacy default* condition purchased products with the highest privacy rating when compared to the participants in the *control* and *low privacy default* conditions. Alternatively, more people in the *control* and the *low privacy default* conditions purchased products with the lowest privacy rating when compared to *high privacy default condition.*

The bar chat also shows that the distribution of purchases made by people in the *control* and *low privacy default* conditions is close to identical. The same number of participants within the two conditions purchased products with the highest and lowest privacy rating.

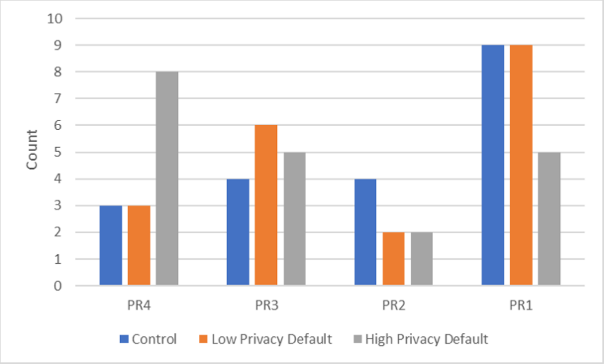

Figure 6. Bar chart comparing the distribution of purchases made by participants using the *control*, *low privacy default* and the *high privacy default* versions of the marketplace.

5.5. High Privacy Default vs Control

One of the goals of this study was to determine whether participants presented with salient privacy information and high privacy defaults would be more likely to purchase devices with a higher privacy rating when compared to participants with no privacy information and defaults.

H1 : More participants using the *high privacy default version* of the marketplace will purchase devices with a higher privacy rating when compared to the participants using the *control version*.

In order to determine this we performed single tailed Wilcoxon rank sum test. The results from our test shows that H1 is true (w=132.5, p-value = 0.029). These results indicate that privacy information along with high privacy defaults influence people to purchase devices with a higher privacy rating.

5.6. High Privacy Default vs Low Privacy Default

The explicit goal for including a different defaults was to compare a design with simple indicators to a design with simple indicators presented in a manner intended to address biases in decision-making.

H2 : More participants using the *high privacy default version* of the marketplace will purchase devices with a higher privacy rating when compared to the participants using the *low privacy default version*.

We designed the study to determine if participants in two different groups with accessible privacy information would make different purchases when the default were different. Specifically, we wanted to see if more participants in the *high privacy default* condition purchased products with a higher privacy rating when compared to the *low privacy default* condition. Once again we performed the single tailed Wilcoxon rank sum test. The results from our tests show that H2 is true (w = 260.5, p-value = 0.044). People in the *high privacy default* condition are more likely to purchase devices with a higher privacy rating compared to people in the *low privacy default* condition.

Despite having the same indicators for privacy, participants in the *high privacy default* group and the *low privacy default* group made significantly different purchase decisions. This indicates that participants' privacy preferences were constructed at the time of decision making rather than being predetermined. Due to this lack of fixed preference participants' decisions were influenced by the order of the categories and the default. Participants who were endowed with the highest privacy (started of with the highest privacy default) were less likely to give it up in exchange for saving money. In other words, they were less likely to accept payment for giving up their privacy. At the same time, participants who were not endowed with privacy (started of with the lowest privacy default) were less likely to spend a few more dollars to protect their privacy (they were not willing to pay to protect their privacy).

5.7. Control vs Low Privacy Default

Finally, we wanted to determine if people in the *low privacy default* condition were more likely to purchase devices with a higher privacy rating when compared to the *control* group. Specifically, we wanted to see if accessible privacy ratings alone were sufficient to make people purchase devices with a higher privacy rating. The inclusion of the low privacy default enabled a comparison of a marketplace with unusable privacy information to one where the information is not available. The results from the single tailed Wilcoxon rank sum test indicate that there is no significant difference between the *control* and *low privacy default* conditions (w = 192, p-value = 0.41). This indicates that presenting users with accessible privacy rating alone is not sufficient to make them purchase devices with a higher privacy rating. The design should also address the psychological biases associated with decision making.

People have a tendency to discount future risks for immediate short-term benefits [3], [4], [30], [38]. The *high privacy default* condition mitigates this issue by employing status-quo bias to favor risk-averse behaviour. However, the *low privacy default* condition does not address this issue. Therefore, a significant number of participants in the *low privacy default* group chose to save money now rather than protect themselves against future risks associated with loss in privacy.

5.8. Price

All devices that had a higher privacy rating were more expensive than those with a lower privacy rating. Therefore, any participant that decided to purchase a device from a higher privacy category had to pay a premium for it. More participants in the *high privacy default* category purchased devices with a higher privacy rating when compared to the participants in other experimental groups. It follows that more participants in the *high privacy default* condition paid a premium for privacy. Here we provide the results from our analysis on differences in prices of products purchased by participants in different experimental conditions.

Each privacy category consisted of two smart plug devices. The prices of both devices for a given category were higher than those from a lower privacy category; and their prices were different from each other. We computed the price premium by calculating the difference in price between the lowest priced product in the lowest privacy category and the lowest priced product in the category from which the participant purchased the device. In other words, this is the difference in starting prices between the two categories. An illustration of this can be found in Figure 7. We believe that this provides us a conservative estimate for the premium paid by the participants. Additionally, when switching between categories, participants are more likely to compute the difference in starting prices between categories as this information is prominently displayed on the tabs.

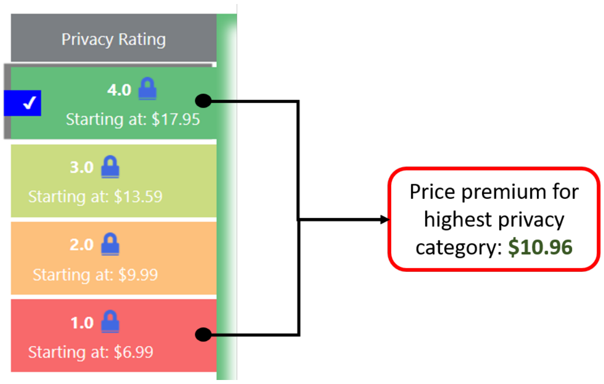

Figure 7. The price premium when a participant purchases a device from the highest privacy category is $10.96. This is calculated by computing the difference in starting prices between the highest and lowest privacy categories.

On average the participants in the *high privacy default* condition paid a premium of $6.18 with $6 being the median. The participants in the *low privacy default* condition on average paid a premium of $3.74 with $3 being the median. Participants in the control condition on average paid a premium of $3.44 with a median of $3. A comparison of means shows that the participants in the *high privacy default* condition on average paid $2.44 more than the participants in the control and *low privacy default* conditions.

We conducted a single tailed Wilcoxon's rank sum test to evaluate the statistical significance between groups. The results from the test shows that the price differences between the *high privacy default* condition and the *low privacy default* condition were statistically significant (W = 260.5, p-value = 0.044). The price differences between *high privacy default* and *control* conditions were also significant (W = 267.5, p-value = 0.029). Finally, the price differences between the *low privacy default* and the *control* conditions were not statistically significant (W = 208, p-value = 0.4147).

6. Discussion

The results from our experiment show that more participants in the *high privacy default* condition purchased devices from the highest privacy category when compared to the other two experimental groups. Conversely, those participants in the *low privacy default* condition made more purchases from the lowest privacy category. In fact, the purchases among the participants in the *low privacy default* were statistically indistinguishable from those in the control group, who had no indicators for privacy. Recall that both privacy rating and price range were presented in each tab. Thus one possible conclusion was that participants in the *low privacy default* group attributed a higher weight to the short term benefits of saving money when compared to potential future risks associated with loss of privacy.

Note that the control group had no privacy information, only price information. The design goal of the *low privacy default* was to provide information clearly differentiating between the privacy provided by different devices (i.e. addressing information asymmetry) and to offer cognitively easy comparisons of this information (i.e., addressing bounded rationality). With only the *low privacy default* we would be forced to conclude that there was no evidence that participants valued privacy.

The *high privacy default* provided exactly the same information as the *low privacy default* and also changed the ordering of the offerings. Defaults are powerful in user security and privacy practices [76]. For example, a comparison of nine-digit zip codes comparing homeowners' decisions to enable encryption to protect home wireless networks found that router defaults dominated education, income, and density in predicting user behavior [36].

The order of the offerings can itself have an effect, even in high-stakes decision-making [26]. Ordering effects are difficult to distinguish from the underlying ratings, familiarity, or popularity that often also impinge decision-making [1]. Yet ordering effects in isolation are known to be powerful. This is true even in offline information interactions; as these are also subject to biases in decision-making. For example, when search was dominated by listings in physical white pages directories professional service agencies often sought names that ensured their listed was first, e.g., AAA Plumbing or 1A Plumbers. In any domain evaluating the influences on decision-making in field research is difficult. However, these results in a controlled experiment provide strong indications of the potential of a comprehensive approach to inform decision-making by mitigating biases

The possibility that participants in all groups simply made the first choice is mitigated by two factors. First, the time to decision was different between groups. If all three participants simply made the decision to select the first listing, then the decision times should be similar. Recall that the average decision time for the *high privacy default* was 52% higher than that of the control group, and 31.4% higher than the *low privacy default*. In fact, the t-test showed that the decision time was statistically significant between the *high privacy default* group and the control group (p-value = 0.018). The t-test showed that the decision time was marginally significant between the *high privacy default* and the *low privacy default* conditions.

Second, the participants in the *high privacy default* had to choose to pay more money. Were all the devices the same prices, we could not distinguish these purchases from a simple status-quo decision. The difference in price indicates a difference in perceived benefit. Together, spending more time and spending more money argues that the differences are substantive as well as significant.

Addressing information asymmetry and bounded rationality proved insufficient in this interaction design to facilitate privacy-aware decisions. The design of the interface had to also address the psychological biases associated with decision making for there to be a significant difference. Previous research has focused on a range of interactions to provide privacy information, and to address the issue of information asymmetry in privacy and security. Experiments have used lock icons, eye icons, interactive eye icons, warning sounds, pop-up warnings, labels, and a little bird for the web [13], [44], [51], [64], [67], [80]. Yet the finding in these experiments have not been consistent when replicated. For example when Benton et al. repeated the experiment by Schlegel et al., using eyes as indicators, the results were inconclusive([13], [67] respectively). Rajivan et al had an identical control group as Feltman, but the percentage of participants who viewed permissions decreased precipitously([28], [64] respectively). Reproducibility in security and privacy is notoriously difficult, and this work may illuminate one of the reasons for this in behavioral analysis.

Our results indicate that some of the variance in previous work may be explained by having experiments that address a subset of the three factors we consider here. Similarly, experiments that include only nudging could be re-evaluated by expanding them to address information provision and risk communication.

The results show two very different behaviors in one experiment. Explicitly, we found both that people will trade privacy for small amounts of money (*low privacy default*) and people will pay for privacy (*low privacy default*). The difference between these groups is the presentation of the information. We have empirical results from a theoretically-grounded argument that the difference is that when people are endowed with privacy their willingness to sell that privacy for small amounts of money significantly decreases as opposed to the case when they perceive themselves to be paying for privacy.

The reality that there is very little privacy with Internet-connected devices and that data surveillance is ubiquitous creates a situation where people experience privacy as a willingness to buy situation. On the positive side this implies that fairly small changes can make a difference; for example, if Mozilla made privacy-enhancing browsing as a default or if Amazon listed devices ordered by privacy information. On the negative side, if the unwillingness to buy privacy and security is a result of the lack of security and privacy then this can be seen as a negative feedback situation or a collective action problem both of which are notoriously difficult to resolve.

The groups of participants were not distinguishable; however, we cannot empirically reject the possibility that there was some unobservable endogenous difference in our participants or in our subtle interactions with participants. While this is true of all evaluations of human behavior and supports repeated investigations into the phenomena referred to as the privacy paradox, it cannot be rejected out of hand. Only reproduction of the experiment can address this possibility. To mitigate this we have provided information (including visualizations) to enable reproduction of the experiment and would provide the code used in our experiment upon request.

Finally, we close the discussion by highlighting the potential of powerful marketplaces, none more than Amazon, for improving the level of privacy in the IoT ecosystem. The majority of the IoT devices passively collect information about their users. As noted above, interviews with people evaluating IoT devices found that information about privacy and security was difficult for them to locate, with some participants saying it was "impossible" [25]. We have illustrated that with positive framing and risk information there is a real demand for privacy that could profit the marketplace and create incentives for privacy-aware IoT.

7. Conclusions

People often express high degrees of concern for their online privacy but their behavior does not consistently align with their expressed concerns [11], [49], [60]. The gap between users' expressed concerns and their behavior has been found to be impinged by three factors: Information Asymmetry, Bounded Rationality and Psychological Biases [4], [63], [75]. Our goal in this work was to develop a marketplace that addressed all three of these phenomena in its interface design. We reduced information asymmetry by classifying the relative level of privacy risks in each device. The privacy policies of the device and the associated app were integrated into this privacy rating. The issue of bounded rationality was addressed by providing this rating information in a tab during the purchase. We also summarizing and provided an additional visualization with a simple 1-5 scale using the lock icons proven to be effective in previous research with mobile devices [64]. The goal was to present indicators to enable comparisons between devices at the decision time with minimal cognitive effort. Finally, we mitigated the effects of psychological biases that adversely affected privacy risk valuations by incorporating design aspects that encouraged risk-averse behavior. Specifically the inclusion of visual ratings not only addressed ease of comparison but also should make privacy more salient. The choice of a high privacy default in one condition was designed to create a feeling of endowment, so that the participant would have to choose to give up privacy (willingness to sell) rather than pay more to purchase privacy (willingness to buy). We then tested the effectiveness of the design by conducting a between subjects experiment.

In the experiment, participants in both the *low privacy default* and *high privacy default* condition were presented with the same indicators for privacy risk, but the framing was different. The resulting purchase decisions made by participants in these two groups were significantly different, with only the *high privacy default* participants making privacy preserving decisions significantly different that the control group. The difference between *low privacy default* and *control* were not statistically significant. We argue that the difference in behavior is due to the different in privacy perspectives generated by the two designs. As participants in the *high privacy default* condition have the highest privacy category as the default, switching to a lower privacy category implies monetary gain and a loss in privacy. Since people attribute a higher weight to losses caused by a change in state they are more likely to attribute a higher weight to the loss in privacy. The status-quo would predict that these participants would keep the highest privacy category and purchase a product within that category. Participants in the *low privacy default* condition start of with the lowest privacy category as their default. So choosing a higher privacy category would result in a loss in monetary savings. Again, these participants stayed with their default category and purchased devices that were priced low: losing privacy to save money.

The results from the experiment further strengthen the argument that user preferences are constructed at decision time rather than being predetermined [29], [43], [61], [70]. Therefore, any design that seeks to facilitate privacy preserving decisions should address the psychological biases associated with user decision making.

The results from the experiment also show that when information asymmetry, bounded rationality, and psychological biases were all addressed people would pay a premium for privacy. In our case, people in the *high privacy default* condition on average paid a $6 premium for privacy. This was $2.44 more than what participants in the control and *low privacy default* conditions paid. While design aspects of the interface can influence people to pay more for privacy this also depended on the difference in prices between categories. There may also be a threshold for the premium that people would be willing to pay for privacy. The threshold maybe higher for devices that are more privacy sensitive like baby monitor or home surveillance cameras and lower for less privacy sensitive for devices like smart plugs. More research is needed to evaluate the pricing of privacy in the presence of the endowment effect and very strong privacy as a default.

Acknowledgment

Removed for review

References

[1] Georgios Abakoumkin. Forming Choice Preferences the Easy Way: Order and Familiarity Effects in Elections. *Journal of Applied Social Psychology*, 41(11):2689–2707, 2011.

[2] Mark S. Ackerman and Lorrie Cranor. Privacy critics: UI Components to Safeguard Users' Privacy. In *CHI '99 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '99, pages 258–259, New York, NY, USA, 1999. ACM.

[3] A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security Privacy*, 3(1):26–33, Jan 2005.

[4] Alessandro Acquisti. Privacy in Electronic Commerce and the Economics of Immediate Gratification. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC '04, pages 21–29, New York, NY, USA, 2004. ACM.

[5] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.

[6] Alessandro Acquisti, Leslie K. John, and George Loewenstein. What Is Privacy Worth? *The Journal of Legal Studies*, 42(2):249–274, 2013.

[7] Idris Adjerid, Eyal Peer, and Alessandro Acquisti. The paradox of wanting privacy but behaving as if it didn't matter. *LSE Business Review*, 2018.

[8] Yuvraj Agarwal and Malcolm Hall. ProtectMyPrivacy: Detecting and Mitigating Privacy Leaks on iOS Devices Using Crowdsourcing. In *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '13, pages 97–110, New York, NY, USA, 2013. ACM.

[9] George A Akerlof. The market for lemons: Quality uncertainty and the market mechanism. In Peter Diamond and Michael Rothschild, editors, *Uncertainty in Economics*, pages 235 – 251. Academic Press, 1978.

[10] Ross Anderson and Tyler Moore. Information security: where computer science, economics and psychology meet. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1898):2717–2727, 2009.

[11] Susanne Barth and Menno D.T. de Jong. The privacy paradox - Investigating discrepancies between expressed privacy concerns and actual online behavior - A systematic literature review. *Telematics and Informatics*, 34(7):1038 – 1058, 2017.

[12] Masooda Bashir, Carol Hayes, April D. Lambert, and Jay P. Kesan. Online Privacy and Informed Consent: The Dilemma of Information Asymmetry. In *Proceedings of the 78th ASIS&T Annual Meeting: Information Science with Impact: Research in and for the Community*, ASIST '15, pages 43:1–43:10, Silver Springs, MD, USA, 2015. American Society for Information Science.

[13] K. Benton, L. J. Camp, and V. Garg. Studying the effectiveness of android application permissions requests. In *2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, pages 291–296, March 2013.

[14] Violet Blue. Stravas fitness heatmaps are a 'potential catastrophe'. *Engadget*, Feb 2018.

[15] JM Blythe and SD Johnson. The Consumer Security Index for IoT: A protocol for developing an index to improve consumer decision making and to incentivize greater security provision in IoT devices. *IET Conference Proceedings*, pages 4 (7 pp.)–4 (7 pp.)(1), January 2018.

[16] Barry Brown. Studying the internet experience. *HP LABORATORIES TECHNICAL REPORT HPL*, 49, 2001.

[17] Jing Chen, Christopher S. Gates, Ninghui Li, and Robert W. Proctor. Influence of Risk/Safety Information Framing on Android App–Installation Decisions. *Journal of Cognitive Engineering and Decision Making*, 9(2):149–168, 2015.

[18] Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. Nudging People Away from Privacy–Invasive Mobile Apps through Visual Framing. In Paula Kotzé, Gary Marsden, Gitte Lindgaard, Janet Wesson, and Marco Winckler, editors, *Human-Computer Interaction – INTERACT 2013*, pages 74–91, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[19] Richard Chow, Serge Egelman, Raghudeep Kannavara, Hosub Lee, Suyash Misra, and Edward Wang. HCI in Business: A Collaboration with Academia in IoT Privacy. In Fiona Fui-Hoon Nah and Chuan-Hoo Tan, editors, *HCI in Business*, pages 679–687, Cham, 2015. Springer International Publishing.

[20] Lorrie Cranor. *Web Privacy with P3P*. "O'Reilly Media, Inc.", 2002.

[21] Lorrie Faith Cranor, Manjula Arjula, and Praveen Guduru. Use of a P3P User Agent by Early Adopters. In *Proceedings of the 2002 ACM Workshop on Privacy in the Electronic Society*, WPES '02, pages 1–10, New York, NY, USA, 2002. ACM.

[22] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. User Interfaces for Privacy Agents. *ACM Trans. Comput.-Hum. Interact.*, 13(2):135–178, June 2006.

[23] André Deuker. Addressing the Privacy Paradox by Expanded Privacy Awareness – The Example of Context–Aware Services. In Michele Bezzi, Penny Duquenoy, Simone Fischer-Hübner, Marit Hansen, and Ge Zhang, editors, *Privacy and Identity Management for Life*, pages 275–283, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[24] Disconnect. Disconnect privacy icons. https://web.archive.org/web/20170709022651/disconnect.me/icons. [Online; accessed 28-June-2019].

[25] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pages 534:1–534:12, New York, NY, USA, 2019. ACM.

[26] Robert Epstein and Ronald E. Robertson. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences*, 112(33):E4512–E4521, 2015.

[27] Ericsson. Internet of Things forecast - Ericsson Mobility Report. https://www.ericsson.com/en/mobility-report/internet-of-things-forecast, Nov 2018.

[28] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android Permissions: User Attention, Comprehension, and Behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, pages 3:1–3:14, New York, NY, USA, 2012. ACM.

[29] Baruch Fischhoff. Value elicitation: Is there anything in there? *American psychologist*, 46(8):835–847, 1991.

[30] Christian Flender and Günter Müller. Type Indeterminacy in Privacy Decisions: The Privacy Paradox Revisited. In Jerome R. Busemeyer, François Dubois, Ariane Lambert-Mogiliansky, and Massimo Melucci, editors, *Quantum Interaction*, pages 148–159, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[31] Geoffrey A Fowler. Its the middle of the night. do you know who your iphone is talking to? *The Washington Post*, May 2019.

[32] V. Garg and J. Camp. Heuristics and Biases: Implications for Security Design. *IEEE Technology and Society Magazine*, 32(1):73–79, Spring 2013.

[33] Google. Nest Secure - Home Security & Alarm System. https://store.google.com/us/product/nest_secure_alarm_system?hl=en-US. [Online; accessed 30-June-2019].

[34] Jens Grossklags and Alessandro Acquisti. When 25 Cents is Too Much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. In *6th Annual Workshop on the Economics of Information Security, WEIS 2007, The Heinz School and CyLab at Carnegie Mellon University, Pittsburgh, PA, USA, June 7-8, 2007*, 2007.

[35] Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G. Shin, and Karl Aberer. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 531–548, Baltimore, MD, 2018. USENIX Association.

[36] Matthew Hottell, Drew Carter, and Matthew Deniszczuk. Predictors of home-based wireless security. In *In The Fifth Workshop on the Economics of Information Security, Cambridge, UK*, 6 2006.

[37] Jeremy Hsu. The Strava Heat Map and the End of Secrets. *WIRED*, Jan 2018.

[38] Thomas Hughes-Roberts. Privacy and Social Networks: Is Concern a Valid Indicator of Intention and Behaviour? In *International Conference on Social Computing, SocialCom 2013, Washington, DC, USA, 8-14 September, 2013*, pages 909–912, 2013.

[39] Privacy Leadership Initiative et al. Privacy notices research final results. *Conducted by Harris Interactive, December*, 2001.

[40] Carlos Jensen and Colin Potts. Privacy Policies As Decision-making Tools: An Evaluation of Online Privacy Notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pages 471–478, New York, NY, USA, 2004. ACM.

[41] Eric J. Johnson, Steven Bellman, and Gerald L. Lohse. Defaults, Framing and Privacy: Why Opting In-Opting Out. *Marketing Letters*, 13(1):5–15, Feb 2002.

[42] Eric J Johnson and Daniel G Goldstein. Defaults and Donation Decisions. *Transplantation*, 78(12):1713–1716, 2004.

[43] Daniel Kahneman and Amos Tversky. *Choices, Values, and Frames*, chapter Chapter 16, pages 269–278. World Scientific, 2013.

[44] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. A "Nutrition Label" for Privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, SOUPS '09, pages 4:1–4:12, New York, NY, USA, 2009. ACM.

[45] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1573–1582, New York, NY, USA, 2010. ACM.

[46] Betty Kim, Darryl Seligman, and Joseph Kable. Preference Reversals in Decision Making Under Risk are Accompanied by Changes in Attention to Different Attributes. *Frontiers in Neuroscience*, 6:109, 2012.

[47] Jack L. Knetsch. The Endowment Effect and Evidence of Nonreversible Indifference Curves. *The American Economic Review*, 79(5):1277–1284, 1989.

[48] Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies*, 71(12):1144 – 1162, 2013.

[49] Spyros Kokolakis. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers Security*, 64:122 – 134, 2017.

[50] Swapna Kolimi, Feng Zhu, and Sandra Carpenter. Contexts and Sharing/Not Sharing Private Information. In *Proceedings of the 50th Annual Southeast Regional Conference*, ACM-SE '12, pages 292–297, New York, NY, USA, 2012. ACM.

[51] Kat Krol and Sören Preibusch. Control Versus Effort in Privacy Warnings for Webforms. In *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society*, WPES '16, pages 13–23, New York, NY, USA, 2016. ACM.

[52] Yee-Lin Lai and Kai-Lung Hui. Internet Opt-in and Opt-out: Investigating the Roles of Frames, Defaults and Privacy Concerns. In *Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research: Forty Four Years of Computer Personnel Research: Achievements, Challenges &Amp; the Future*, SIGMIS CPR '06, pages 253–263, New York, NY, USA, 2006. ACM.

[53] Colin Lecher. Its not easy to opt out of the Strava heat map thats revealing secret locations. *The Verge*, Jan 2018.

[54] Frederick Liu, Shomir Wilson, Peter Story, Sebastian Zimmeck, and Norman Sadeh. Towards Automatic Classification of Privacy Policy Text. 12 2017.

[55] Franco Loi, Arunan Sivanathan, Hassan Habibi Gharakheili, Adam Radford, and Vijay Sivaraman. Systematically Evaluating Security and Privacy for Consumer IoT Devices. In *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy*, IoTS&#38;P '17, pages 1–6, New York, NY, USA, 2017. ACM.

[56] Theresa M. Marteau. Framing of information: Its influence upon decisions of doctors and patients. *British Journal of Social Psychology*, 28(1):89–94, 1989.

[57] Aleecia M McDonald and Lorrie Faith Cranor. The Cost of Reading Privacy Policies. *ISJLP*, 4:543, 2008.

[58] George R. Milne and Mary J. Culnan. Strategies for reducing online privacy risks: Why consumers read (or dont read) online privacy notices. *Journal of Interactive Marketing*, 18(3):15 – 29, 2004.

[59] Abhijith Athreya Mysore Gopinath, Shomir Wilson, and Norman Sadeh. Supervised and Unsupervised Methods for Robust Separation of Section Titles and Prose Text in Web Documents. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 850–855, Brussels, Belgium, October-November 2018. Association for Computational Linguistics.

[60] Patricia A Norberg, Daniel R Horne, and David A Horne. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007.

[61] John W. Payne, James R. Bettman, and Eric J. Johnson. Behavioral Decision Research: A Constructive Processing Perspective. *Annual Review of Psychology*, 43(1):87–131, 1992.

[62] Victoria C Plaut and Robert P Bartlett III. Blind consent? A social psychological investigation of non-readership of click-through agreements. *Law and human behavior*, 36(4):293–311, 2012.

[63] Stefanie Pötzsch. Privacy awareness: A means to solve the privacy paradox? In *IFIP Summer School on the Future of Identity in the Information Society*, pages 226–236. Springer, 2008.

[64] Prashanth Rajivan and Jean Camp. Influence of Privacy Attitude and Privacy Cue Framing on Android App Choices. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, Denver, CO, 2016. USENIX Association.

[65] Ilana Ritov and Jonathan Baron. Status-quo and omission biases. *Journal of Risk and Uncertainty*, 5(1):49–61, Feb 1992.

[66] William Samuelson and Richard Zeckhauser. Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1(1):7–59, Mar 1988.

[67] Roman Schlegel, Apu Kapadia, and Adam J. Lee. Eyeing Your Exposure: Quantifying and Controlling Information Sharing for Improved Privacy. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, SOUPS '11, pages 14:1–14:14, New York, NY, USA, 2011. ACM.

[68] Maurice Schweitzer. Disentangling Status Quo and Omission Effects: An Experimental Analysis. *Organizational Behavior and Human Decision Processes*, 58(3):457 – 476, 1994.

[69] Herbert A Simon. Rational choice and the structure of the environment. *Psychological review*, 63(2):129–138, 1956.

[70] Paul Slovic. The construction of preference. *American psychologist*, 50(5):364, 1995.

[71] Strava. Run and Cycling Tracking on the Social Network for Athletes. https://www.strava.com/features. [Online; accessed 30-June-2019].

[72] Jeff Sweat. Privacy paradox: Customers want control–and coupons. *InformationWeek*, (781):52, Apr 10 2000. Name - General Motors Corp; American Airlines Inc; 1-800-Flowers.com; Copyright - Copyright CMP Media Inc. Apr 10, 2000; Last updated - 2017-11-10; CODEN - INFWE4; SubjectsTermNotLitGenreText - United States; US.

[73] Taylor Telford. Google failed to notify customers it put microphones in nest security systems. *The Washington Post*, Feb 2019.

[74] Richard Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior Organization*, 1(1):39 – 60, 1980.

[75] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research*, 22(2):254–268, 2011.

[76] Markus Tschersich and Reinhardt Adriaan Botha. Understanding the impact of default privacy settings on self-disclosure in social networking services: Building a conceptual model and measurement instrument. In *19th Americas Conference on Information Systems*. Association for Information Systems, 2013.

[77] Joseph Turow, Lauren Feldman, and Kimberly Meltzer. Open to Exploitation: America's Shoppers Online and Offline. *A Report from the Annenberg Public Policy Center of the University of Pennsylvania*, page 35, 2005.

[78] Amos Tversky and Daniel Kahneman. Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157):1124–1131, 1974.

[79] Tony Vila, Rachel Greenstadt, and David Molnar. Why We Can't Be Bothered to Read Privacy Policies Models of Privacy Economics As a Lemons Market. In *Proceedings of the 5th International Conference on Electronic Commerce*, ICEC '03, pages 403–407, New York, NY, USA, 2003. ACM.

[80] Kim-Phuong L. Vu, Vanessa Chambers, Beth Creekmur, Dongbin Cho, and Robert W. Proctor. Influence of the Privacy Bird® user agent on user trust of different web sites. *Computers in Industry*, 61(4):311 – 317, 2010. Human-Centered Computing Systems in Industry - A Special Issue in Honor of Professor G. Salvendy.

[81] Ryan West. The Psychology of Security. *Commun. ACM*, 51(4):34–40, April 2008.

[82] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N Cameron Russell, et al. The Creation and Analysis of a Website Privacy Policy Corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1330–1340, 2016.

[83] Sebastian Zimmeck and Steven M. Bellovin. Privee: An Architecture for Automatically Analyzing Web Privacy Policies. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 1–16, San Diego, CA, 2014. USENIX Association.