IoT Marketplace: Willingness-To-Pay vs. Willingness-To-Accept

Shakthidhar Reddy Gopavaram[1], Jayati Dev[1], Sanchari Das[2], and L. Jean Camp[1]

[1] Indiana University, Bloomington, IN, USA {sgopavar,jdev,ljcamp}@iu.edu
[2] University of Denver, Denver, CO, USA {Sanchari.Das}@du.edu

Abstract. Willingness-To-Pay (WTP) is the most a person is willing to pay for a good or service. Conversely, Willingness-To-Accept (WTA) is the minimum amount a person is willing to accept for giving up a good or service. People often attribute a higher value for privacy in the WTA condition when compared to the WTP condition. In behavioral economics of privacy and security, this discrepancy between WTP and WTA has been explained by the endowment effect and the status-quo bias. In this study, we aim to emulate the effects of WTP and WTA through interface design. Specifically, we employed the principles of status-quo bias to build two versions of the IoT Marketplace. While one design of the marketplace emulated the WTA condition, the other emulated the WTP condition. In both versions of the marketplace, we communicated the aggregate privacy rating associated with the IoT device using the same indicator. We evaluated the effect the two interfaces had on product selection by conducting a study where participants were asked to purchase an IoT device using either a WTA interface or a WTP interface. Our results show that participants using the interface that emulated the WTA condition were more likely to pay a premium to purchase devices with a higher privacy rating when compared to the participants using the interface that emulated the WTP condition. We also append results showing that when asked to select items without making the purchases the same effect was more pronounced.

Keywords: Willingness-to-pay · Willingness-to-protect · Willingness-to-accept, Experimental economics · Privacy · Information security · Marketplace · IoT · Usable Privacy and Security · Endowment Effect

1 Introduction

The past decade has seen a 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 predicted to almost double to approximately 18 billion IoT devices by 2022 [12]. While these devices offer great convenience, they also collect a lot of sensitive information about the user and pose a significant threat to their privacy and security. For example, Strava, a fitness tracking platform, collected detailed activity information about it's users to generate a global heat map. The heat map exposed users' home addresses and mapped classified military operations [19, 21, 6].

Consumers are becoming increasingly aware of the privacy and security risks associated with IoT devices through media reports, opinions shared by friends, and by observing unexpected device behavior [11]. However, consumer awareness alone is not sufficient to protect individual privacy, "more fundamental behavioral responses must also be addressed if privacy ought to be protected" [1]. In this work, we explore the role of privacy in decision-making specifically by isolating the discrepancy between Willingness-To-Pay (WTP) and Willingness-To-Accept (WTA) as experimentally observable manifestations of the endowment effect and the status-quo bias.

Researchers have previously studied the differences in value attributed to privacy in WTP and WTA conditions by explicitly asking participants to either pay for privacy or accept payment in exchange for giving up their privacy [3, 15]. In those studies, the participants in the WTA condition often attributed a higher value for their privacy [3, 15]. Beyond privacy, WTP/WTA disparities have been found to be consistent across wide varieties of goods such as mugs, movie tickets, nuclear waste repositories and many others [18]. Past research also suggests that students tend to have lower WTA/WTP ratios than the general public. So conducting the experiment with the general population would increase the WTA/WTP ratio [18]. Across decades of empirical studies a general finding in WTA/WTP is that "a gain may be moderately valuable but a loss could be irreplaceable. and the difference between WTA and WTP would then be large" [17]. We began with a hypothesis that this general result would be applicable in the case of home-based IoT where there is the 'barn door' nature of information loss [32] and sensitive contexts which may result in severe loss of privacy [28]. Alternatively, living spaces can have the most sensitive privacy contexts so that there concerns of privacy of intimate sensitive data define the baseline value for privacy independent of framing. Additionally, we try to see if it is possible to emulate

the WTA and WTP scenarios without explicitly asking the participants to either give up their privacy for a monetary benefit or pay to keep their information private.

We recruited participants to be consumers in an experimental IoT marketplace, then randomly assigned them into WTP, and WTA groups. Participants in the WTP and WTA groups were provided with the same indicators for aggregate privacy risk, products, and prices. Each participant was given a $25 gift card. Without making a purchase they selected two devices, a fitness trackers and a security camera, as desirable. Then they purchased a smart plug, keeping the remaining balance on the card. In order to align with marketplace preferences, we selected participants who expressed a desire to make such a purchase. The marketplace presented smart plug devices which varied in price and privacy. All devices that had a higher privacy rating were priced higher than those with a lower privacy rating, exogenously to the experiment design. The results from this show that participants using the WTA framing were more likely to purchase privacy-preserving devices by paying a premium for it. Participants in the WTP group did not consistently make privacy-preserving choices despite having the same privacy information and purchasing devices for use in the same context.

In Section 2, we detail past work in decision making and risk communication that influenced the design of our IoT Marketplace interfaces. In Section 3, we provide the rationale for design of the WTA and WTP versions of the IoT Marketplace. In Section 4, we detail the experiment design; please note that the experiment was subject to Institutional Review Board (IRB) review. In Section 5, we provide the results from our experiment and then return to the literature for explanations. Finally, we conclude with a discussion on the possible implications of our findings in Section 6.

2 Related Work

In this section we provide an introduction to the three components of the experiment. First we address previous experiments in privacy valuation that similarly address the biases we consider. We then broaden the discussion to include examples of privacy valuation which use labels and indicators. We close by identifying the related work on privacy ratings that informed our calculation of these ratings for the selected IoT devices.

2.1 Endowment Effect and Status-quo Bias

It is common for people to attribute a higher value to items they possess when compared to the items they don't possess [20, 3, 15, 2]. This discrepancy in valuations between Willingness-To-Pay and Willingness-To-Accept payment has also been observed in the case of privacy. Specifically, past research has shown that people assign a higher monetary value to privacy in the WTA condition when compared to the WTP condition [3, 15]. In an early work applying this to privacy valuation Acquisti et al. conducted an experiment in which the participants were either endowed with a $10 gift card or a $12 gift card [3]. The $10 gift card was anonymized i.e. purchases made through the gift card could 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 could 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 being disclosed 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 that people are more likely to attribute a higher value to privacy in the WTA condition.

The discrepancy between WTA and WTP can be explained by the endowment effect and the status-quo bias. The endowment effect states that people are more attached to the items they possess, so they demand a higher price than what they are willing to pay for it [29]. In the example above, people who were initially endowed with the anonymized card were more attached to their privacy, so they were less likely to give it up for an additional $2.

Status-quo bias consists of two primary components: (1) strong preference for the current state of affairs and (2) strong preference for not taking any action, also known as omission bias [26]. This strong preference for the current state of affairs is due to loss aversion [26, 29]. A change from status-quo implies that people would lose some things while gaining other things. Since people are loss averse, they tend to attribute greater weight to losses when compared to gains which explains their strong preference for the status-quo [26]. It must also be noted that the 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 [26]. 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 [26]. The fear of potential regret from taking an action prevents people from changing the status-quo. This fear of regret is one of the primary reasons for omission bias. According to status-quo bias, in the example above, people who were endowed with the anonymized card are less likely to give up their privacy for an additional $2 because (1) they are loss averse, so they don't want to give up their privacy and (2) they are more apprehensive about the negative outcomes associated with disclosing their personal information.

Making any decision about the importance (or lack thereof) of security and privacy in a decision, it is necessary to have information about the security and privacy of the product. In the next subsection, we describe the previous work that informed our design decisions on the icons used to communicate the privacy and security choices.

2.2 Labels and Indicators

An interview of people considering IoT products found that consumers not only did not have privacy information, they did not know where to find it [11]. In this work, our goal was to compare WTP/WTA for IoT devices. A comparison of WTP/WTA based on the privacy and security of the devices requires awareness of the privacy implications of a purchase. In order to do this, we built upon previous research on the conditions under which participants in research and in markets include privacy in their decision-making.

There is a significant body of research that supports the power of simple indicators that communicate aggregate information exposure risks as effective decision-making aids on multiple platforms [14]. An early comparative study by Tsai et al. showed that the provision of indicators for aggregate privacy risk decreases the discrepancy between users' expressed privacy concerns and their behavior [30]. Tsai and her coauthors recruited people who expressed high levels of concern about online privacy then asked them to select vendors for two products using a search engine. When the search engine provided participants with only links to the vendors' webpages and the price of the product, there was no evidence that participants' decisions were influenced by the vendors' privacy policies. Participants in this case systematically chose the vendor that offered the lowest price. However, when the product listing was augmented 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. Following on that pilot two later studies on the web were Privacy Bird and Privacy Finder. The former displayed a red, yellow, or green bird to indicate if a website's privacy policy met the users' stated preferences, the latter was a privacy-enabled search engine that generated privacy ratings on a 5-point scale for each of its search results. Repeated studies showed that both Privacy Bird and Privacy Finder resulted in users' choosing those websites with stronger privacy policies [31, 30, 7, 9].

Studies conducted on mobile platforms have shown similar results [27, 31, 4] using a range of interactions and indicators. In each case the presence of the indicators themselves influenced privacy behaviors but to different degrees. For example, a study comparing icons, brief text, and long text found that only the indicator impinged decision-making in the app marketplace [5]. Later work verified that text permissions not only did not influence decision-making, but also were rarely read and not well understood by consumers in the app marketplace [13].

For this study, the choice of icons and framing was significantly motivated by the findings of a study conducted by Rajivan et al. [25]. In that study, the authors compared stars, locks, and eyes (based on [27]) using both positive and negative framing conditions. The results of this quantitative evaluation on the efficacy of multiple framing mechanisms and icons was that positively framed risk information presented on a 5-point scale using the padlock icon was most effective. Therefore, here we use the padlock icon and a five point scale to communicate the aggregate privacy risk. An illustration of the privacy indicator used in this experiment is shown in Figure 1.

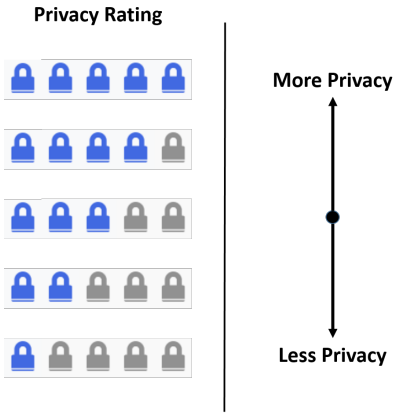

Fig. 1. Privacy rating communicated using the lock icon. More locks imply more privacy.

If the existence of the privacy icon alone significantly changes decision-making, then the different WTP/WTA framings should not have a substantive effect on participant choices.

2.3 Ratings

Previous experiments in willingness to pay above compared similar products: web pages and apps. In the IoT domain, evaluation of the privacy and security provided by the devices with the corresponding apps and services is itself an active research area. For consistency between devices and based on previous work above [31, 30, 7, 9] we calculated the risk ratings 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 validity of our work is influenced by the generation of consistent believable ratings, we include the following additional research on the creation of ratings.

In the past decade, Machine Learning has also been proposed to automate the generation of privacy policies [33, 16, 22, 24, 34]. One of the machine learning approaches that directly informed our work was the framework proposed by Harkous et al. [16]. 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 efficacy of their framework with an application that accepts a link to a privacy policy and then generates appropriate disconnect icons[3] based on preset rules.

3 Marketplace Design

For this study, we created two versions of the IoT Marketplace: WTA Version and WTP Version. In both versions of the marketplace, devices that offered a higher level of privacy were priced higher than those that offered a lower level of privacy. We did this to see if people would purchase devices that offered better privacy even if they charged a premium. In a study conducted by Emami-Naeini et al., participants were asked to rank different factors based on the influence they had on their purchase of IoT devices. The results from this study showed that "privacy and security were ranked

[3] Disconnect icons are privacy icons that were developed a Mozilla-led working group. These icons were designed to make privacy policies understandable, and to communicate data collection and usage practices [8].

among the most important" after price and features [11]. In this study, we control for the impact of features on participants' choices by making sure that all devices offered the same features. The only distinguishing factors between the devices was the privacy rating and price.

The *WTA version* provides participants with indicators for aggregate privacy risk. This version of the marketplace emulates the WTA scenario. The *WTP version* also provides participants with indicators for aggregate privacy risk. These indicators are the same as the ones used in the *WTA version*. However, this version of the marketplace emulates the WTP scenario. Therefore, we expect fewer participants using the *WTP version* of the marketplace to make privacy-preserving purchases when compared to the participants using the *WTA version*.

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

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

Fig. 2. (a)Positively framed privacy rating illustrated using the padlock icon. (b) A short description explaining the privacy rating.

3.1 WTA Version

As noted in Section 2, the lack of indicators for aggregate privacy risk could lead to people not considering privacy while making their purchase choices [30]. 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. [25]. 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 2. The generation of the aggregate privacy rating is discussed in the Appendix. Please note that the generation of privacy ratings is not the focus of this work. The focus of this study is on the behavioral aspects of purchase decisions for IoT devices.

In addition to providing participants with aggregate risk ratings, this version of the marketplace also emulates the WTA scenario. To emulate the WTA scenario, we employ status-quo bias in our design. Status-quo bias consists of two primary components:

1. Loss Aversion: When performing an action that leads to a 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 state change, then people are more likely to stick to the status quo to avoid the regret caused by the potential negative outcomes.

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

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. Participants were provided with the ability to switch between categories by clicking on the respective tabs. A screenshot of our design is shown in Figure 3 (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

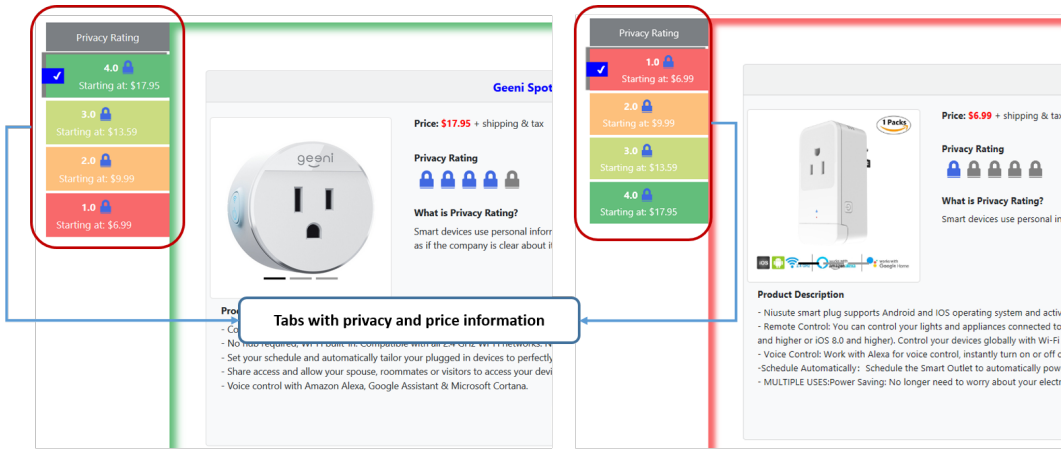

(a) WTA Version (b) WTP Version

Fig. 3. (a) *WTA version* 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) *WTP version* has the category with privacy rating 1 as the default and the tabs are ordered in the ascending order of privacy (or descending order of price)

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 attribute a higher weight to their loss in privacy when compared to their gain in monetary saving. Furthermore, loss of 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 to the same outcome occurring due to inaction [26]. We hypothesize that people will exhibit *omission bias* and would avoid purchasing a device from a lower privacy category to avoid potential regret.

3.2 WTP version

Similar to the *WTA version*, in this version of the marketplace all devices are categorized by their privacy rating. Users could also switch between categories by 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 shown in Figure 3 (b)). By changing the default and the order of categories, we change the 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 this interface emulates the WTP scenario. More participants using this version of the marketplace are likely to purchase a device with a lower privacy rating to save money.

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 two experimental groups. The two groups are WTA group and WTP group. While the participants in the WTA group used the *WTA version* of the marketplace, participants in the WTP group used the *WTP version* of the marketplace.

Everyone that agreed to participate in the study, irrespective of the group they were assigned to, was 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. For the first two categories, participants were asked to select products that they were most likely to purchase. These two categories were used as a way to help participants familiarize themselves with the interface of the marketplace. The results for these categories can be found in the Appendix. 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 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 through ads on university classifieds, posting flyers on campus, and through email blasts sent to university students. The recruiting material was designed to target people who would be interested in purchasing and utilizing a smart plug. Specifically, we solicited participants by stating that they would receive a free smart plug. The first half of the ad also contained different use cases for a smart plug. A digital copy of the flyer used in classifieds ads and email blasts can be found in the Appendix. By recruiting participants who had an intention to use the smart plug we were able to record purchase decisions that are comparable to the ones they make when purchasing a device on an actual real-world marketplace.

We conducted power analysis using the R *MKpower* package to determine the appropriate sample size for our study. For the power analysis, we used the following estimations for mean and standard deviation: est. mean WTA = 3, est. sd WTA = 1, est. mean WTP = 2, and est. sd WTP = 1. The analysis showed that we would need 13.09 participants per group. The power was set to 0.8.

For this study, we recruited 20 participants per group which is a little over the computed sample size. So we had a total of 40 participants.

5 Results

We found that participants in the WTA group were more likely to purchase products with a higher privacy rating when compared to participants in the WTP conditions. This shows that it is possible to emulate WTA and WTP scenarios through interface design. Furthermore, despite having the same visual indicators for privacy, a significant number of participants in the WTP group purchased products with the lowest privacy

rating. This demonstrates that the design of the interface can play a significant role in nudging people towards privacy preserving decisions.

5.1 Demographics

All participants were over 18, and the sample skewed younger. Fifty-five percent of the participants were between 18-25 years old; 30% were between 25-35 years old, and 15% of the participants were older than 35. Out of the 40 participants, 50% were men, and 50% were women.

Group	Mean	Standard Deviation	Median
WTA	2.8	1.24	3
WTP	2.15	1.18	2

Table 1. The table presents the mean, standard deviation, and median privacy rating for the two experimental groups.

5.2 Descriptive Statistics

As you can see in Figure 4, the distribution of Smart Plug purchases made by the participants in the WTA group is skewed towards a higher privacy rating. For participants in the WTP groups, the distributions are skewed towards a lower privacy rating. The mean, standard deviation, and median for the two experimental groups can be found in Table 1.

Figure 5 compares the distribution of privacy ratings for the products purchased by participants in the two experimental conditions. The median privacy rating for the WTA group is higher than that of the WTP group. This indicates that a lot more participants in the WTA group purchased products with the highest privacy rating when compared to the participants in the WTP condition. Alternatively, more people in the WTP condition purchased products with the lowest privacy rating when compared to the WTA group.

In the following subsections, we determine if the observed differences between the WTA group and the WTP group are statistically significant.

Fig. 4. Bar chart comparing the distribution of purchases made by participants using the *WTA* and the *WTP* versions of the marketplace.

5.3 WTA vs WTP

The goal of the study was to see if it is possible to emulate WTA and WTP scenarios without explicitly asking the participants to either give up their privacy for a monetary benefit or pay to keep their information private. Specifically, we wanted to see if more participants in the WTA group purchased products with a higher privacy rating when compared to the WTP condition. Once again we performed the single-tailed Wilcoxon rank-sum test. The results from our tests show that H1 is true (w = 260.5, p-value = 0.044, r = 0.270). The p-value has been adjusted for multiple testing. People in the WTA condition are more likely to purchase devices with a higher privacy rating compared to people in the WTP condition.

Despite having the same indicators for privacy, participants in the WTA and the WTP groups made significantly different purchase decisions. Participants who were endowed with the highest privacy (WTA group) 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 (WTP group) were less likely to spend a few more dollars to protect their privacy i.e.

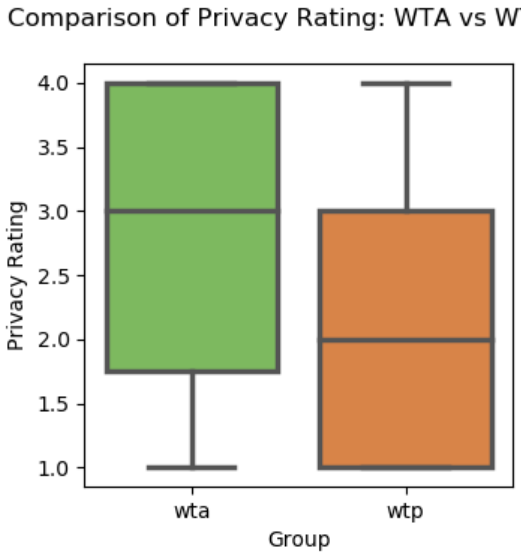

Fig. 5. A comparison of box plots representing the distribution of privacy ratings for purchases made by participants using the *WTA* and the *WTP* versions of the marketplace.

they were not willing to pay to protect their privacy. This indicates the interface of the marketplace can play a significant role in nudging participants towards or away from making privacy-preserving decisions.

5.4 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 WTA category purchased devices with a higher privacy rating when compared to the participants in other experimental groups. It follows that more participants in the WTA 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 6. 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.

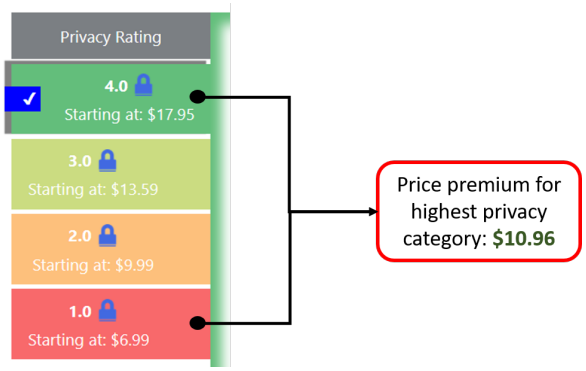

Fig. 6. The price premium when a participant purchases a device from the highest privacy category is $10.96. This is calculated by computing the difference in starting prices between the highest and lowest privacy categories.

On average the participants in the WTA condition paid a premium of $6.18 with $6 being the median. The participants in the WTP condition on average paid a premium of $3.74 with $3 being the median. A comparison of means shows that the participants in the WTA condition on average paid $2.44 more than the participants in the WTP condition. A comparison of box plots representing the distribution of the premium paid by participants in different groups is shown in Figure 7.

We conducted a single-tailed Wilcoxon's rank-sum test to evaluate the statistical significance between groups. The results from the test show that the price differences between the WTA condition and the WTP condition were statistically significant ($W = 260.5$, p-value $= 0.044$, r=0.301).

In a study conducted by Emami-Naeini et al., participants reported that they would be willing to pay a premium of 10%-30% of the base price of an IoT device for privacy and security [11]. However, the value attribute to privacy can vary based on whether people are asked to pay for privacy or how much they would accept for disclosing their private information [3, 15]. In the past, this variance in the value attributed to privacy was shown by conducting studies where they explicitly asked participants to

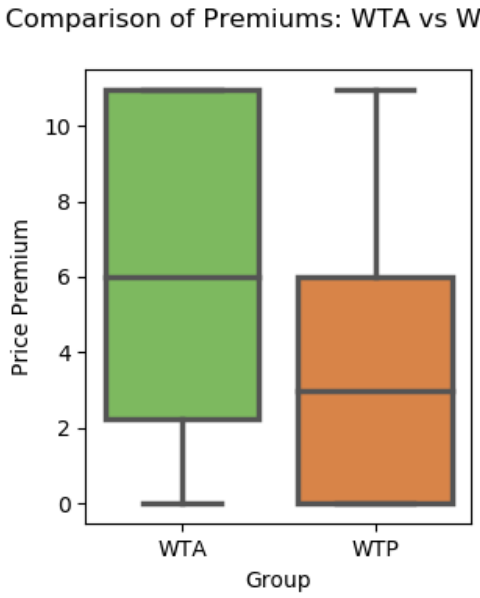

Fig. 7. A comparison of box plots representing the distribution of premium paid by participants using the *WTA*, and the *WTP* versions of the marketplace.

specify their WTP or WTA for privacy. Here, we show that the design aspects of the interface can be used to emulate WTA and WTP scenarios. These results are promising and indicate that people can be nudged to make more privacy-preserving decisions through the interface design.

Finally, we would like to state that the difference in premiums between the *WTA* and the *WTP* conditions can vary based on the privacy sensitivity of the IoT device in consideration. Further research is needed to elicit consumers' behavior in those conditions.

5.5 Tabs Viewed

For both the WTA and the WTP 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 WTP 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 are visible on the tabs. This could indicate an overall 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 WTA 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 while price was still a consideration, participants in the WTA category were less likely to purchase devices with a lower privacy rating.

A significant portion of participants within the WTA and the WTP 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.

5.6 Time to Decision

For each participant, we recorded the time taken to select a smart plug device. On average, the decision time for participants in the WTP conditions was 3.56 minutes. The average decision time for participants in the WTA condition was 4.68 minutes. This was approximately 1 minute more than what was observed for the WTP group. The median for the WTA condition was also higher than that of WTP group.

The results from the one sided t-test show that the decision time between the WTA and WTP is not statistically significant (Cohen's d = 0.463, t = 1.466, df = 32.564, p-value = 0.076).

6 Discussion

The results from our experiment show that more participants in the WTA condition purchased devices from the highest privacy category when compared to the WTP group. Conversely, those participants in the WTP condition made more purchases from the lowest privacy category. Recall that the goal of the study was to emulate WTA and WTP scenarios through interface design. The results from for experiment i.e. the gap in privacy premiums and privacy ratings show that we have successfully achieved this goal.

These results are promising and indicate that people can either be nudged towards or away from making privacy-preserving choices through interface design. This has significant implications. For example, if Amazon were to categorize devices based on their privacy and show people devices in the highest privacy category by default, a lot more people would be willing to pay a premium to purchase a device with the highest privacy rating. Furthermore, this would give companies a monetary incentive to design and manufacture devices with higher privacy and security standards.

We selected participants who self-identified as planning to purchase home IoT devices. If those with high levels of contextual privacy concerns are therefore self-excluding by their rejection of IoT privacy risks, then the differences between groups will represent only those in the current market. Previous research has focused on obtaining either subjects who value privacy (e.g., [30] or those with representative privacy preferences (e.g., [10, 23]). We choose to recruit individuals who sought the devices in order to obtain more ecologically valid purchase decisions.

The possibility that participants in both groups simply made the first choice is mitigated by the fact that participants in the WTA group had to choose to pay more money. Were all the device prices the same, we could not distinguish these purchases from a simple status-quo decision. The difference in price indicates a difference in perceived benefit. Spending more money indicates that the differences are substantive as well as significant.

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 the

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.

7 Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. CNS 1565375 CNS 1814518; support from a Cisco Research Award No. 1377239, and funds from the Comcast Innovation Fund. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF, Cisco, Comcast, Indiana University, or University of Denver.

References

1. Acquisti, A.: Privacy in Electronic Commerce and the Economics of Immediate Gratification. In: Proceedings of the 5th ACM Conference on Electronic Commerce. p. 21–29. EC '04, Association for Computing Machinery, New York, NY, USA (2004). https://doi.org/10.1145/988772.988777
2. Acquisti, A., Brandimarte, L., Loewenstein, G.: Privacy and human behavior in the age of information. Science **347**(6221), 509–514 (2015). https://doi.org/10.1126/science.aaa1465, `https://science.sciencemag.org/content/347/6221/509`
3. Acquisti, A., John, L.K., Loewenstein, G.: What Is Privacy Worth? The Journal of Legal Studies **42**(2), 249–274 (2013). https://doi.org/10.1086/671754
4. Agarwal, Y., Hall, M.: 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. pp. 97–110. MobiSys '13, ACM, New York, NY, USA (2013). https://doi.org/10.1145/2462456.2464460, `http://doi.acm.org/10.1145/2462456.2464460`
5. Benton, K., Camp, L.J., Garg, V.: Studying the Effectiveness of Android Application Permissions Requests. In: 2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops). pp. 291–296 (2013). https://doi.org/10.1109/PerComW.2013.6529497
6. Blue, V.: Strava's Fitness Heatmaps are a 'Potential Catastrophe'. Engadget (Feb 2018), `https://www.engadget.com/2018/02/02/strava-s-fitness-heatmaps-are-a-potential-catastrophe/`
7. Cranor, L.F., Arjula, M., Guduru, P.: Use of a P3P User Agent by Early Adopters. In: Proceedings of the 2002 ACM Workshop on Privacy in the Electronic Society. pp. 1–10. WPES '02, ACM, New York, NY, USA (2002). https://doi.org/10.1145/644527.644528, `http://doi.acm.org/10.1145/644527.644528`
8. Disconnect: Disconnect Privacy Icons. `https://web.archive.org/web/20170709022651/disconnect.me/icons`, [Online; accessed 28-June-2019]

9. Egelman, S., Tsai, J., Cranor, L.F., Acquisti, A.: Timing is Everything? The Effects of Timing and Placement of Online Privacy Indicators. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. p. 319–328. CHI '09, Association for Computing Machinery, New York, NY, USA (2009). https://doi.org/10.1145/1518701.1518752, https://doi.org/10.1145/1518701.1518752
10. Emami-Naeini, P., Dheenadhayalan, J., Agarwal, Y., Cranor, L.F.: Which Privacy and Security Attributes Most Impact Consumers' Risk Perception and Willingness to Purchase IoT Devices?
11. Emami-Naeini, P., Dixon, H., Agarwal, Y., Cranor, L.F.: Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. pp. 534:1–534:12. CHI '19, ACM, New York, NY, USA (2019). https://doi.org/10.1145/3290605.3300764, http://doi.acm.org/10.1145/3290605.3300764
12. Ericsson: Internet of Things forecast -– Ericsson Mobility Report. https://www.ericsson.com/en/mobility-report/internet-of-things-forecast (Nov 2018)
13. Felt, A.P., Ha, E., Egelman, S., Haney, A., Chin, E., Wagner, D.: Android Permissions: User Attention, Comprehension, and Behavior. In: Proceedings of the Eighth Symposium on Usable Privacy and Security. SOUPS '12, Association for Computing Machinery, New York, NY, USA (2012). https://doi.org/10.1145/2335356.2335360, https://doi-org.proxyiub.uits.iu.edu/10.1145/2335356.2335360
14. Garg., V.: A lemon by any other label. In: Proceedings of the 7th International Conference on Information Systems Security and Privacy - ICISSP,. pp. 558–565. INSTICC, SciTePress (2021). https://doi.org/10.5220/0010295205580565
15. Grossklags, J., Acquisti, A.: 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), http://weis2007.econinfosec.org/papers/66.pdf
16. Harkous, H., Fawaz, K., Lebret, R., Schaub, F., Shin, K.G., Aberer, K.: Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In: 27th USENIX Security Symposium (USENIX Security 18). pp. 531–548. USENIX Association, Baltimore, MD (2018), https://www.usenix.org/conference/usenixsecurity18/presentation/harkous
17. Horowitz, J.K., McConnell, K.: Willingness to Accept, Willingness to Pay and the Income Effect. Journal of Economic Behavior Organization **51**(4), 537–545 (2003). https://doi.org/https://doi.org/10.1016/S0167-2681(02)00216-0, https://www.sciencedirect.com/science/article/pii/S0167268102002160
18. Horowitz, J.K., McConnell, K.E.: A Review of WTA/WTP Studies. Journal of Environmental Economics and Management **44**(3), 426–447 (2002). https://doi.org/https://doi.org/10.1006/jeem.2001.1215, https://www.sciencedirect.com/science/article/pii/S009506960191215X
19. Hsu, J.: The Strava Heat Map and the End of Secrets. WIRED (Jan 2018), https://www.wired.com/story/strava-heat-map-military-bases-fitness-trackers-privacy/
20. Knetsch, J.L.: The Endowment Effect and Evidence of Nonreversible Indifference Curves. The American Economic Review **79**(5), 1277–1284 (1989), http://www.jstor.org/stable/1831454
21. Lecher, C.: It's not easy to opt out of the Strava heat map that's revealing secret locations. The Verge (Jan 2018), https://www.theverge.com/2018/1/29/16945866/strava-heat-map-opt-out-guide
22. Liu, F., Wilson, S., Story, P., Zimmeck, S., Sadeh, N.: Towards Automatic Classification of Privacy Policy Text (12 2017). https://doi.org/10.1184/R1/6626285.v1
23. Momenzadeh, B., Gopavaram, S., Das, S., Camp, L.J.: Bayesian Evaluation of User App Choices in the Presence of Risk Communication on Android Devices. In: Clarke, N., Furnell, S. (eds.) Human Aspects of Information Security and Assurance. pp. 211–223. Springer International Publishing, Cham (2020)
24. Mysore Gopinath, A.A., Wilson, S., Sadeh, N.: 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. pp. 850–855. Association for Computational Linguistics, Brussels, Belgium (Oct-Nov 2018), https://www.aclweb.org/anthology/D18-1099

25. Rajivan, P., Camp, J.: Influence of Privacy Attitude and Privacy Cue Framing on Android App Choices. In: Twelfth Symposium on Usable Privacy and Security (SOUPS 2016). USENIX Association, Denver, CO (2016), https://www.usenix.org/conference/soups2016/workshop-program/wpi/presentation/rajivan
26. Ritov, I., Baron, J.: Status-quo and Omission Biases. Journal of Risk and Uncertainty **5**(1), 49–61 (Feb 1992). https://doi.org/10.1007/BF00208786, https://doi.org/10.1007/BF00208786
27. Schlegel, R., Kapadia, A., Lee, A.J.: Eyeing Your Exposure: Quantifying and Controlling Information Sharing for Improved Privacy. In: Proceedings of the Seventh Symposium on Usable Privacy and Security. pp. 14:1–14:14. SOUPS '11, ACM, New York, NY, USA (2011). https://doi.org/10.1145/2078827.2078846, http://doi.acm.org.proxyiub.uits.iu.edu/10.1145/2078827.2078846
28. Templeman, R., Korayem, M., Crandall, D.J., Kapadia, A.: PlaceAvoider: Steering First-Person Cameras away from Sensitive Spaces. In: NDSS. pp. 23–26. Citeseer (2014)
29. Thaler, R.: Toward a Positive Theory of Consumer Choice. Journal of Economic Behavior & Organization **1**(1), 39 – 60 (1980). https://doi.org/https://doi.org/10.1016/0167-2681(80)90051-7, http://www.sciencedirect.com/science/article/pii/0167268180900517
30. Tsai, J.Y., Egelman, S., Cranor, L., Acquisti, A.: The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. Information Systems Research **22**(2), 254–268 (2011). https://doi.org/10.1287/isre.1090.0260, https://pubsonline.informs.org/doi/abs/10.1287/isre.1090.0260
31. Vu, K.P.L., Chambers, V., Creekmur, B., Cho, D., Proctor, R.W.: Influence of the Privacy Bird® user agent on user trust of different web sites. Computers in Industry **61**(4), 311 – 317 (2010). https://doi.org/https://doi.org/10.1016/j.compind.2009.12.001, http://www.sciencedirect.com/science/article/pii/S0166361509002097, human-Centered Computing Systems in Industry - A Special Issue in Honor of Professor G. Salvendy
32. Whitten, A., Tygar, J.D.: Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. In: Proceedings of the 8th Conference on USENIX Security Symposium - Volume 8. p. 14. SSYM'99, USENIX Association, USA (1999)
33. Wilson, S., Schaub, F., Dara, A.A., Liu, F., Cherivirala, S., Leon, P.G., Andersen, M.S., Zimmeck, S., Sathyendra, K.M., Russell, N.C., 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). vol. 1, pp. 1330–1340 (2016)
34. Zimmeck, S., Bellovin, S.M.: Privee: An Architecture for Automatically Analyzing Web Privacy Policies. In: 23rd USENIX Security Symposium (USENIX Security 14). pp. 1–16. USENIX Association, San Diego, CA (2014), https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/zimmeck

Appendix

A. Privacy Ratings

Please note that the generation of privacy ratings is not the focus of this work. To show that the generating such ratings is possible, we cited the relevant past work in Section 2.3. Additionally, the participants were only provided information about what the rating represented. They were not provided any information about how these ratings were generated. So the procedure used to generate these ratings did not influence participants' decisions.

The privacy ratings were generated manually 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. The ratings were positively framed and represented using a padlock icon. An illustration of this can be found in Figure 1.

B. Results for Fitness Tracker and Security Cameras

	w	p-value	r
Security Camera	99	0.0016	0.734
Fitness Tracker	120.5	0.0099	0.551

Table 2. The table presents the results for the single tailed Wilcoxon rank-sum tests.

C. Recruitment

For this study, we wanted to recruit participants who would be interested in purchasing and utilizing a smart plug device. We believed that by recruiting participants with such inclination, we would be able to record purchase decisions that are comparable to the ones people make when they are purchasing a device on an actual real-world

marketplace. To achieve this, we designed our recruiting material to target people who would be interested in purchasing and utilizing a smart plug device. Specifically, we solicited participants by stating that they would receive a free smart plug. The first half of the flyer also contained different use cases for a smart plug. In the second half of the flyer, we explain to them that they will be given a "$25 Amazon gift card to purchase a Smart Plug of their choice from a list of options presented to them". We also state that as compensation for participating in the study, they will get to keep the smart plug they purchased and the amount left on the gift card after the purchase has been made. A digital copy of the flyer used in classifieds ads and email blasts can be found in Figure 8.

Fig. 8. A digital copy of the flyer used in classifieds ads and email blasts.