Comparing Hypothetical and Realistic Privacy Valuations

Joshua Tan
Carnegie Mellon University
jstan@cmu.edu

Mahmood Sharif
Carnegie Mellon University
mahmoods@cmu.edu

Sruti Bhagavatula
Carnegie Mellon University
srutib@cmu.edu

Matthias Beckerle
Carnegie Mellon University
beckerle@cmu.edu

Michelle L. Mazurek
University of Maryland
mmazurek@cs.umd.edu

Lujo Bauer
Carnegie Mellon University
lbauer@cmu.edu

ABSTRACT

To protect users' privacy, it is important to understand how they value personal information. Prior work identified how framing effects alter users' valuations and highlighted the difficulty in eliciting real valuations through user studies under hypothetical circumstances. However, our understanding of users' valuations remains limited to specific entities, information types, and levels of realism. We examined the effects of realism and purpose of use on users' valuations of their personal information. Specifically, we conducted an online study in which participants (N=434) were asked to assign monetary value to their personal information in the context of an information marketplace involving different receiving parties, while we experimentally manipulated the level of realism of the scenario and the timing of eliciting valuations. Among our findings is a nuanced understanding of valuation biases, including when they may not apply. For example, we find that, contrary to common belief, participants' valuations are not generally higher in hypothetical scenarios compared to realistic ones. Importantly, we find that while absolute valuations vary greatly between participants, the order in which users prioritize information types (i.e., users' relative valuations of different attributes) remains stable across the levels of realism we study. We discuss how our findings inform system design and future studies.

CCS CONCEPTS

• **Security and privacy → Economics of security and privacy**; *Privacy protections*; • **Human-centered computing** → Empirical studies in ubiquitous and mobile computing;

KEYWORDS

Privacy economics; human factors; online study

ACM Reference Format:
Joshua Tan, Mahmood Sharif, Sruti Bhagavatula, Matthias Beckerle, Michelle L. Mazurek, and Lujo Bauer. 2018. Comparing Hypothetical and Realistic Privacy Valuations. In *2018 Workshop on Privacy in the Electronic Society (WPES'18), October 15, 2018, Toronto, ON, Canada.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3267323.3268961

WPES '18, October 15, 2018, Toronto, ON, Canada

ACM ISBN 978-1-4503-5989-4/18/10.
https://doi.org/10.1145/3267323.3268961

1 INTRODUCTION

A growing trend in online services is the pervasive collection and use of users' personal information—attributes like age, gender, and location—for a range of purposes, from website customization and service personalization to targeted advertising. While users sometimes appreciate the benefits resulting from the use of their information, they often express concern about sharing it [43, 61]. Hence, understanding how users value their personal information has become an important question, including for system design (e.g., will users be outraged by what information a system is using and for what purpose? [13]) and legislation and public policy (e.g., how should users be compensated as part of data-breach lawsuits [37]).

Measuring users' valuations of personal information, however, is challenging due to the many factors that affect privacy preferences and behavior. For example, experimental elicitation of users' valuation of their information is often complicated by the endowment effect: namely, users may undervalue personal information that they have already shared [18]. More importantly, experimentally elicited valuations are commonly inconsistent with real behavior—users who report being concerned about their privacy often share their personal information in real life [56], a phenomenon often called *the privacy paradox* [41, 64].

Prior work, however, has not precisely quantified the difference in hypothetical and real valuation of personal information. Hence, it remains unclear whether the privacy paradox is a uniquely strong phenomenon, or whether it simply reflects the commonly observed *hypothetical bias*—humans' tendency to respond differently to hypothetical scenarios than to similar real scenarios—that has been well studied for other kinds of goods [33, 34]. Further, it is unclear to what extent hypothetical valuations may still be useful, even if they do not perfectly capture real-world behavior; for example, are relative valuations derived from hypothetical questions valid even if absolute valuations are not? In addition, while various studies have considered both the effect of contextual information on privacy preferences and some of the economic biases inherent in privacy valuation, little work has considered how these factors interact.

To fill these gaps, we conducted an online between-subjects user study in which we elicited users' valuations for seven personal attributes when shared with six different receiving parties. Our study was framed in the context of an information market that aggregates information about users, sells it to interested parties, and compensates users whose information is sold at a price the users set. Participants were assigned to one of five conditions that varied in terms of realism—to measure hypothetical bias—and in

whether endowment effects were introduced. Thus, in total, we controlled for four factors, each of which has been individually found to affect users' preferences and behavior: the type of personal information [32], with whom information is shared [46], the existence of endowment effects [18], and whether the scenario is realistic or hypothetical [56]. By including all four factors, we can explore how well results about the economic biases of privacy valuation do (not) generalize to different information-sharing contexts.

Our study improves our understanding of valuation biases in ways that can benefit both system design and future studies. Contrary to what one would expect given the privacy paradox, for example, we found that participants' valuations for the attributes we studied were generally not significantly higher in the hypothetical conditions than in the realistic conditions. We also observed a nuanced endowment effect: Rather than decreasing the valuations of attributes overall, it appears to affect only the valuations of certain attributes (e.g., phone number and home address).

One of our goals was to determine whether we could establish specific relationships between the valuations of attributes based on the conditions we varied. We found that participants' rankings of attributes remained stable even when changing the level of realism, introducing endowment effects, or changing the entities to whom the information was sold. Thus, we were able to train machine-learning models on data from hypothetical conditions to predict rankings in realistic conditions with high accuracy.

We next discuss related work (Sec. 2), followed by an explanation of our methodology (Sec. 3). We then present our results (Sec. 4), discuss how they inform system design and future studies (Sec. 5), and conclude (Sec. 6).

2 RELATED WORK

Our research lies at the intersection of three different lines of work: work studying how context affects privacy concerns and behavior, privacy-economics research studying how users valuate privacy, and general economics research studying hypothetical bias. We next discuss related work in each of these subfields.

2.1 Privacy in context

Prior work showed that users are concerned about sharing personal data in different applications, such as search result personalization, single-sign-on authentication (SSO), recommender systems, and others (e.g., [8, 15, 27, 28, 43, 55, 59, 61]). For example, researchers discovered that although users generally preferred personalized search results, for certain sensitive topics this benefit was outweighed by privacy concerns [43].

As supported by Nissenbaum's theory of *privacy as contextual integrity* [40], context is a key factor in users' decisions to share personal data online. Contextual integrity identifies five parameters that affect data-sharing decisions: the data subject, sender, recipient, information type, and transmission principle (e.g., retention time). As discussed in Sec. 3, in this work we explore the effects of the last three parameters jointly.

In the light of users' privacy concerns and motivated by the contextual integrity theory, researchers have studied what factors affect users' data-sharing decisions, finding that factors such as the type of data being shared, with whom it is being shared, and its retention time affect users' willingness to share information (e.g., [2, 8, 11, 14, 17, 25, 32, 35, 42, 48, 54]). Leon et al. performed a large-scale survey in which participants were asked about their willingness to share 30 types of information with online advertisers [32]. They found some types of information that most participants would share (e.g., gender) and other classes that about half of users would prefer to keep private (e.g., phone number). Acquisti et al. found evidence of endowment and order effects on privacy valuation; whether participants were asked how much they would sell their data for versus how much they would pay to protect it and the order of different offers for their data both significantly affected the values they provided [2]. Similar endowment effects were found by researchers studying users' willingness to pay for privacy protection in smartphone applications [14]. In recent work, Chanchary and Chiasson found that the presence of mechanisms to control what data is shared and with whom increased participants' willingness to share [8]. Differently from us, prior research efforts focused on studying one dimension of contextual integrity at a time. In contrast, we examine interactions among the different dimensions.

Prior work also emphasized the importance of trust on sharing decisions. Costante et al. used a general trust perception model to quantify the user's trust in different websites [11]. Joinson et al. found that the trust that users place in the entity with whom they share data affects their willingness to share and that trust can compensate for privacy-invasive sharing: If a user trusts an entity, they may share otherwise-sensitive data [25]. These findings further support our design choice to study sharing decisions with a variety of entities whom users may trust variably.

2.2 Valuating privacy

Some studies relied on conjoint analysis—a technique borrowed from the economics literature that can be used to indirectly extract the monetary values of products' features by asking users to rank different products—to find the monetary value users assign to privacy controls and private information [19, 29, 50, 51]. For example, Pu and Grossklags found that participants sharing a friend's personal information with a smartphone app would sell it for $1.01 when it is irrelevant to the app's functionality, and for $0.68 otherwise. While conjoint analysis is useful in some cases, the mechanics of using it (it relies on users either ranking all options or making a series of choices about which of two options is preferred) make it difficult to design a plausibly realistic condition for valuating privacy. As such, it cannot easily be used to measure hypothetical bias in our context.

The previously described studies examined privacy behaviors using hypothetical surveys. However, privacy-sensitive disclosure often happen in more realistic settings (e.g., when filling in membership forms at stores) in which participants may gain tangible benefit from the disclosure—which may change their behavior. In fact, prior work showed that users' inclination to share their attributes in real settings is often higher than reported in hypothetical studies [22, 56]. The difference between reported privacy attitudes and actual behavior is often referred to as the *privacy paradox* [41, 64].

Many research efforts have studied users' data-sharing behavior in realistic settings (e.g., [5, 6, 12, 16, 18, 20, 23, 46, 47, 58, 60]).

To observe actual behavior, researchers in this line of work try to ensure *incentive compatibility*, namely, that participants achieve the best outcome by acting according to their true preferences. Notably, Tsai et al. found that participants shopping for sensitive products online are ready to pay a $0.60 premium to purchase from a website that protects their privacy [60]. To ensure incentive compatibility, the money used for shopping was deducted from the amount participants received as compensation for participating in the study. Huberman et al. used a second-price auction as an incentive-compatibility mechanism in a study where participants were asked to share their age or weight with others in return for a monetary gain [20]. They found that participants with less socially acceptable traits (older or more overweight) valued their information more highly. Grossklags et al. presented participants with two types of offers: one involving protecting themselves from information release in exchange for money and the other involving releasing their information in exchange for money [18]. They found signs of *endowment effects*: participants were willing to pay lower amounts in exchange for protecting their information from being released than they were willing to accept for selling their information.

In contrast to prior work, we elicit our participants' valuations for several personal attributes and measure how they change between hypothetical and realistic conditions, as well as when the data is shared with different entities. We show that the privacy paradox may not be as universal as once believed to be. Moreover, our study design allows us to quantify people's *hypothetical bias* when evaluating different attributes (i.e., to quantify the actual difference in valuations between hypothetical and real settings).

2.3 Hypothetical bias

Hypothetical bias for public and private tangible goods has been extensively studied (e.g., [9, 10, 33, 34, 38]). For example, one meta-analysis combined results from 29 different experimental designs, finding that hypothetical bias causes participants to over-value all goods, but more so for public goods than private goods [33]. As another example, in the domain of buying tangible goods, realistic scenarios have been found to better predict actual behavior than hypothetical scenarios [9].

We extend prior work by systematically exploring hypothetical bias specifically in the context of sharing private personal attributes.

3 METHODOLOGY

We conducted a large-scale between-subjects online study with 434 participants, which we split between five conditions. To administer the study, we used Prolific [49], a crowdsourcing platform developed at the University of Oxford. We analyzed the collected data via quantitative methods to get insight into how different factors—information type, receiving party, and realism—affect participants' absolute and relative valuations of their information. Below we present our study design, the analysis methods, and the limitations of each.

Entity	Description
Ad networks	Finding potential consumers to advertise products or special deals
Federal agencies	Producing census data about American people
Insurance companies	Customizing and advertising insurance plans
Market research companies	Providing guidance to companies about consumer preferences
Political parties	Conducting political surveys and polls
Research pools	Recruiting participants for academic research studies

Table 1: Entities we asked about and their descriptions (as shown to participants).

3.1 Study design

In all study conditions, participants first completed a distraction task in which they were asked to distinguish between images of real objects and objects generated by an artificial intelligence algorithm. Because the distraction task was mainly intended to prevent suspicion in the deception conditions (see below), it was always presented as the main task in the study. The distraction task also served to avoid bias during recruitment; by avoiding advertising the study as one concerning privacy, we hoped to avoid bias in our sample toward the less privacy-concerned. Moreover, including the distraction task helped control for social-desirability effects, which may lead participants who know their privacy-related behavior is being observed to over-value it [44].

Then, to elicit participants' valuation of their personal information, we asked participants to assign a dollar value to each of seven personal attributes as remuneration for sharing them with six different receiving parties. The attributes we asked about were age, email address, gender, home address, occupation, phone number, and relationship status. We chose these attributes because previous work showed that different users find them to be differently sensitive [32]. Further, these attributes can be requested via Google Single Sign-On (SSO), which (as described below) adds to the realism of our realistic conditions. Table 1 lists the six receiving entities, along with descriptions provided to participants. We developed this list based on entities with which people often share personal information in real life; our pilot studies confirmed that these entities elicited sufficient variation in responses.

This valuation activity was performed in the context of a (fictional) information market operated by our institution. As in the work of Laudon and Varian [31, 62], the market was presented as a central entity that gathers personal information about individuals and sells it to interested entities. For each combination of attribute and party, participants could either assign a price at which to sell the attribute or opt not to sell the attribute. We designed our pricing explanation to mimic eBay's explanation of their auction mechanism (except in reverse, because users are selling rather than buying). We explained to participants that attribute buyers would use their limited budgets to purchase only the lowest-priced attributes available; however, analogously to eBay, all sellers would be paid at or

above their specified price for any attribute sold.[1] By modeling our explanation on eBay's, we expected that participants would have sufficient understanding to effectively price their attributes. We note that this value-elicitation mechanism is incentive compatible (see [30]); namely, it incentivizes participants to provide valuations that reflect how much they believe various attributes are actually worth.

To introduce them to the market, participants were required to read several paragraphs of explanatory text, including examples. To help ensure participants read carefully, we broke the text into multiple pages and included two attention-check questions assessing the participant's understanding of the market. The first attention check concerned the overall goals of the market, and the second how buyers would select lowest-priced offerings. In each case, if the participant got the question wrong once, we returned them to the explanation and asked them to try again. We exclude any participant who failed either question both times from our analysis.

Each participant was assigned to one of five conditions; two realistic and three hypothetical (and increasingly less realistic):

Realistic with endowment ($Real_{End}$). Participants in this condition were (deceptively) told that the information market was operational and that they could potentially earn additional money by opting to sell their information. We provided them with a URL to a mock-up website describing the information market. Participants were asked at the beginning of the study to sign in with their Google SSO accounts in order to share their information. The reason for this was twofold: to convince participants that they would need to sell real data, rather than making up fictional attributes to sell; and to establish an endowment effect by having participants provide us with personal data before valuating that data (i.e., the potential to undervalue information that has already been shared) [2]. Before the valuation task, we showed participants the information already collected from them (as a reminder that they had already shared it with us). We also told participants that if any information that was not available on their Google account was sold, the market would contact them to collect it.

A crucial part of the design of this condition (and $Real_{NoEnd}$, described next) was to make the user experience fully realistic up to the point where participants would have potentially been additionally paid for one or more of their attributes. Specifically, for participants who might have had some of their attributes purchased, the interactions they experienced were exactly the same (including actually sharing their attributes with us via SSO login) as they would have been if the marketplace had been real; the only difference is that instead of receiving additional payment for attributes, they received a debrief, described below. Similarly, for participants who did not bid a sufficiently low value for their information to be purchased, the process they experienced was identical to what it would have been if the marketplace had been real, with the debrief tacked on at the end. Hence, although we did not actually purchase attributes participants offered for sale, the study design minimizes or eliminates the impact of this on the realism of the condition as experienced by participants.

Participants were debriefed about the deception at the end of the study. We then asked them which information collected from their Google account was accurate, stressing that their answers would not affect their compensation. We also deleted all information we gathered about participants when they logged in through Google SSO. Prolific approved our collection of participants' personal information via Google SSO. In contrast, the collection of personal information is against the terms of service in more popular and well-studied platforms (e.g., Amazon Mechanical Turk [26]).

Realistic without endowment ($Real_{NoEnd}$): This condition was nearly identical to $Real_{End}$. The key difference was that participants were not asked to log in with Google SSO at the beginning of the study (thus eliminating endowment effects). After we collected their valuations, we did ask participants to log in; we then gave them the option to revisit their valuations.

As in condition $Real_{End}$, participants were debriefed about the deception at the end and were asked about the accuracy of the data collected via Google. Then, all Google SSO data we gathered was deleted. We also asked participants who changed their initial valuations about their reasoning.

Less realistic (Hyp_{Low}): In this hypothetical condition, participants were told that they would not be earning money through the market or sharing their information with it. Instead, participants were told that they would be helping to evaluate a market that was on the verge of becoming operational.

Even less realistic (Hyp_{Medium}): This condition is even more hypothetical than Hyp_{Low}. Participants were told that they would be helping to evaluate the idea of the information market, which researchers were considering implementing.

Least realistic (Hyp_{High}): In the most hypothetical condition, participants were merely instructed to imagine they were selling their information to different parties via a market.

Participants were never asked to log in with Google SSO during any of the hypothetical conditions.

In the traditional economics literature, the valuations elicited in our realistic conditions ($Real_{End}$ and $Real_{NoEnd}$) are often referred to as *revealed preferences* and the valuations elicited in the hypothetical conditions (Hyp_{Low}, Hyp_{Medium}, and Hyp_{High}) as *stated preferences* or *contigent valuations* [7]. We use the terms *realistic* and *hypothetical valuations* for the sake of simplicity and clarity to a non-economics audience.

As the last task in the study, participants from all conditions were asked whether they think their valuations would differ outside of a user study, ten questions from the Internet Users' Information Privacy Concerns (IUIPC) scale [36], and demographic questions. The question about whether valuations would change outside the study was intended to gauge whether participants in realistic conditions believed the presented scenario. The IUIPC scale is used to gauge users' privacy concern and consists of three sub-scales: control, awareness, and collection. The sub-scales range between 1 and 7, with higher values indicating higher privacy concern. The protocols for all five conditions can be found in App. A.

We used Prolific controls to confine the sampling of our participants to people living in the United States who were at least 18 years old and whose approval rate was above 90%. Participants in all conditions were compensated $2.50 (corresponding to $10/hr compensation and above the $7.25/hr minimum wage), and participants in the two realistic conditions received an additional $1.50

[1] https://ocsnext.ebay.com/ocs/sr?&query=337

to compensate for the deception. Our study was approved by our institution's ethics board.

3.2 Analysis

Here we discuss the quantitative methods and metrics that we used in our analyses.

Examining absolute valuations. To understand the factors impacting participants' valuations, we implemented a mixed-effect linear model [45]. In this model, the dependent variable was the amount of dollars an attribute was sold for.[2] Values were classified as outliers using the interquartile range (IQR) method; log-transformed values falling outside three times the IQR (often denoted extreme outliers) were excluded from our primary analysis. This method has been used by other researchers studying valuation data (e.g., [63]) and was important for two reasons. Firstly, outlier removal was necessary to satisfy assumptions of normality for the regressions we performed. Secondly, it allowed us to exclude instances in which participants gave impractically large values for attributes (e.g., $10 billion for a phone number),[3] a behavior observed in other privacy valuation studies (e.g., [18, 57]).

To understand participants' likelihood of selling attributes, we also implemented a mixed-effect logistic regression model. This model helped us measure how the likelihood of selling information changed as a result of changing the factors we considered. A random participant intercept was incorporated into both models to account for participants providing values for multiple attributes and parties.

Model selection was performed according to a backward selection approach by comparing the Bayesian information criterion (BIC) for different model variations [53]. Model fit was quantified using both marginal R^2 and conditional R^2 [4]. Marginal R^2 describes the variance explained by fixed effects alone, while conditional R^2 describes the variance explained by the combination of fixed and random effects [39].

To examine whether individual factors in our model had a statistically significant effect, we perform Wald tests. Hypothesis tests for non-zero effects of model factors for attribute type, receiving party, and level of realism constitute a priori hypotheses, as these hypotheses were formulated at the beginning of the study, informed by prior work. Given the exploratory nature of our examination of factor interactions, for Wald tests involving interaction terms we apply the Holm-Bonferonni multiple testing correction, within each regression model.

Examining rankings. To compare rankings we use Kendall's τ correlation and the Normalized Discounted Cumulative Gain (NDCG).

Kendall's τ is a correlation statistic that ranges between -1 and 1, where 1 indicates that two sets of attributes have exactly the same rankings and -1 indicates completely opposite rankings. This metric considers all pairs of attributes; the more pairs that are ordered the same (e.g., both participants rank attribute A above attribute B), the higher the τ. For example, $\tau = 0.25$ means that 62.5% of the attribute pairs had the same internal order in the two participants' rankings, while $\tau = 0.75$ means that 87.5% of pairs did.

Kendall's τ is agnostic to the location of the mismatch between rankings; as a result, this metric is insufficient by itself when some mismatches are more important than others. Therefore, we complement it using NDCG [21], which is often used to evaluate the performance of information-retrieval and recommendation systems. Given a predicted ranking $\hat{r}$ and a true ranking r, NDCG measures how well $\hat{r}$ estimates r, where a perfect match has NDCG = 1 and poorer matches have NDCG approaching 0. NDCG takes position in the rankings into account; higher-ranked attributes must be predicted accurately to achieve a high gain. In contrast, mispredictions in lower-ranked attributes have less impact on the gain.

Examining dropout rates. To test whether dropout rates differed between conditions, we used Pearson's χ^2 omnibus test. For pairwise comparisons between conditions, we used Fisher's exact test. Post-hoc comparison p-values were corrected using the Holm-Bonferroni method.

A concern when dropouts rates differ between conditions is whether this introduces bias into experimental groups; e.g., if a higher dropout rate in *Real* conditions led to a smaller proportion of privacy-sensitive participants in *Real* compared to *Hyp* conditions. To test this, we examined the IUIPC scores of participants who did not drop out.[4] Specifically, we use the Mann-Whitney U test to examine whether *Real* participants differed from *Hyp* participants in any of the three IUIPC privacy subscales.

3.3 Limitations

Our results should be interpreted in the context of several study limitations. First, eliciting valuations directly may not exactly match situations in which users typically disclose information. In systems like SSO and Android permission systems, users more typically derive a non-monetary benefit, such as saving time or receiving personalized service, from taking an action that results in the disclosure of personal information. Nonetheless, because this effect would equally modify all conditions and attributes—and because there are many non-monetary benefits and contexts that could be studied—we believe that our straightforward method of preference elicitation is both appropriate and resulted in useful insights.

Second, some participants in $Real_{End}$ and $Real_{NoEnd}$ may not have have been convinced that the information marketplace was real despite the deception. We tried to address this in part through the use of Google SSO. By incorporating Google SSO, nearly all participants shared at least one attribute (email address) and many shared their birthdate, gender, and profile picture (see Sec. 4). The goal of this sharing was not to collect all data the marketplace would collect, but to improve realism by bolstering belief that we would be collecting that data. Participants' answers also evidence their belief in the realism of the presented scenario, as fewer than 13% of participants in the realistic conditions reported they would change their responses outside of our study. Further, the 11 participants in $Real_{NoEnd}$ who changed their initial valuations after signing in with Google made comments suggesting that the act of signing in added to the study's realism.

[2] We applied a log(amount +1) transformation to satisfy assumptions of normality and to account for 0 values.

[3] Inspection of these instances suggests that some participants voiced discontent with being asked to price personal data by setting very high prices.

[4] Our analysis could not directly examine the IUIPC scores of dropouts, since this information was not available.

Online studies involving deception can suffer if potential participants share information before taking the study. For our study, we believe it is unlikely that participants discussed the study's deceptive component with future participants, as (unlike for Amazon's Mechanical Turk) we were unable to find online communities discussing Prolific tasks and the data was collected within roughly five hours.

Third, as in many studies, our participants are not completely representative of the population. As shown in prior work, participants recruited through crowdsourcing platforms are often younger, more educated, and more privacy-sensitive than the general population [26]. While we are unaware of studies comparing Prolific users with the general population, we have no reason to believe they would be more representative than those of other crowdsourcing platforms.

Finally, it may be possible that participants in conditions $Real_{End}$ and $Real_{NoEnd}$ used fake (or "burner") Google accounts to participate in our study. We did not encounter signs of such behavior.

4 RESULTS

We analyzed the data from our survey quantitatively, along several axes. First we examined how participants' valuations and likelihood of selling were affected by the factors we considered. Second, we studied how the rankings of attributes differed among conditions. Finally, we studied the reasons that led some participants to drop out from the study. We report on the results after presenting our participants.

4.1 Participants

A total of 457 participants completed our survey. We excluded from our analysis 22 participants who failed either attention check after two attempts. Of the remaining 434 participants, 61 were assigned to $Real_{End}$, 53 to $Real_{NoEnd}$, 104 to Hyp_{Low}, 112 to Hyp_{Medium}, and 104 to Hyp_{High}. Participants were randomly assigned to conditions, and the number of participants assigned to $Real_{End}$ and $Real_{NoEnd}$ (92 and 98, respectively) was comparable to that of the hypothetical conditions; however, as we discuss in Sec. 4.4, $Real_{End}$ and $Real_{NoEnd}$ had higher dropout rates.

Participants' ages ranged from 18 to 76 years, with a median of 29 years. The gender distribution was 41% female and 57% male, with the remainder specifying other or choosing not to answer. The majority (58%) of participants held an associate's degree or higher. 60% of participants reported an annual income of between $25,000 and $99,999, 18% reported $24,999 or less, and 14% reported $100,000 or more. Their most common primary occupation types were college student (13%); service (e.g., retail clerk, server; 11%); unemployed (10%); and computer engineer or IT professional (10%). Their average IUIPC factor scores were 6.0 (sd=0.9) for control, 6.4 (sd=0.9) for awareness, and 5.8 (sd=1.2) for collection, indicating that our participants are concerned about their privacy.

According to our outlier criteria, 42 participants (408 valuations) were considered outliers: 11 participants (118 valuations) in $Real_{End}$, 2 participants (28 valuations) in $Real_{NoEnd}$, 7 participants (70 valuations) in Hyp_{Low}, 13 participants (122 valuations) in Hyp_{Medium}, and 9 participants (70 valuations) in Hyp_{High}. Outlier values ranged from $250 to $10,000,000,000. We excluded participants with outlier values from our primary analyses, which removed less than 10% of our data. The proportion of outliers removed did not significantly differ between conditions, according to a χ^2 test of independence (p=.07).

For most *Real* participants, we collected their Google profile picture, email address, and date of birth (Figure 1). Besides profile picture, most of the collected info was reported by participants as accurate.[5]

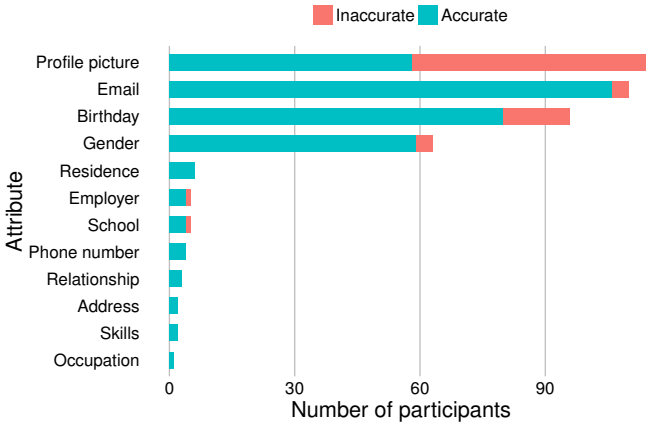

Figure 1: Attributes collected from *Real* participants' Google profiles.

4.2 Attribute valuations

To understand how participants value their personal attributes, it is necessary to consider both the dollar values they assigned to their attributes and instances in which they chose to forgo selling an attribute altogether. We analyze both of these outcomes using separate regression models, trained on the values provided by the 392 participants not classified as outliers. We provide examples of predictions from our regression models for illustrative purposes. We also make conclusions about individual factor significance (e.g., about $Real_{End}$, i.e., the endowment effect) based on hypothesis tests. While predictions from our trained models may indicate a factor has a non-zero effect, our conclusions about non-zero effects with respect to the population are based on statistically sound hypothesis tests.

An overview of participants' attribute valuations is provided in Figure 2. The percentage of participants that chose to sell each attribute (to different third parties and in different conditions) is summarized in Figure 3.

4.2.1 Dollar values. We first attempted to determine which factors affected, and to what extent, participants' assignment of dollar values to personal attributes. We investigate this via a mixed-factor linear regression analysis. We began model selection from a model that included fixed effects for condition, attribute, receiving party, and all two-way interactions between condition, attribute, and party. In addition, the model included fixed effects for the following

[5] We suspect the reason for profile pictures not always being accurate is that Google returns placeholder profile pictures when the user has not explicitly set one in their profile.

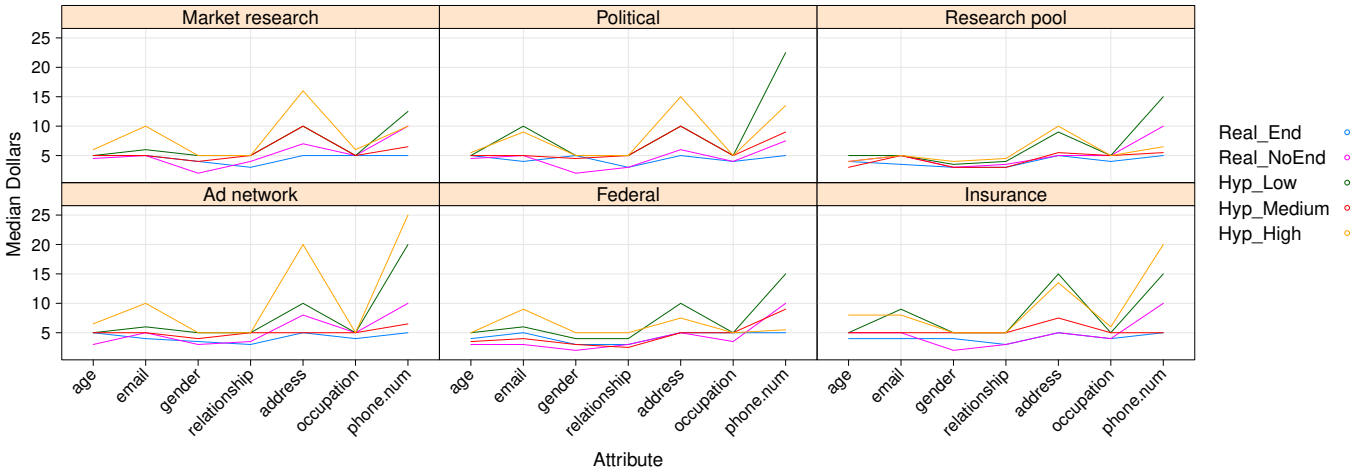

Figure 2: Median values assigned by participants. Shown for different attributes, receiving parties, and levels of realism for the elicitation scenario.

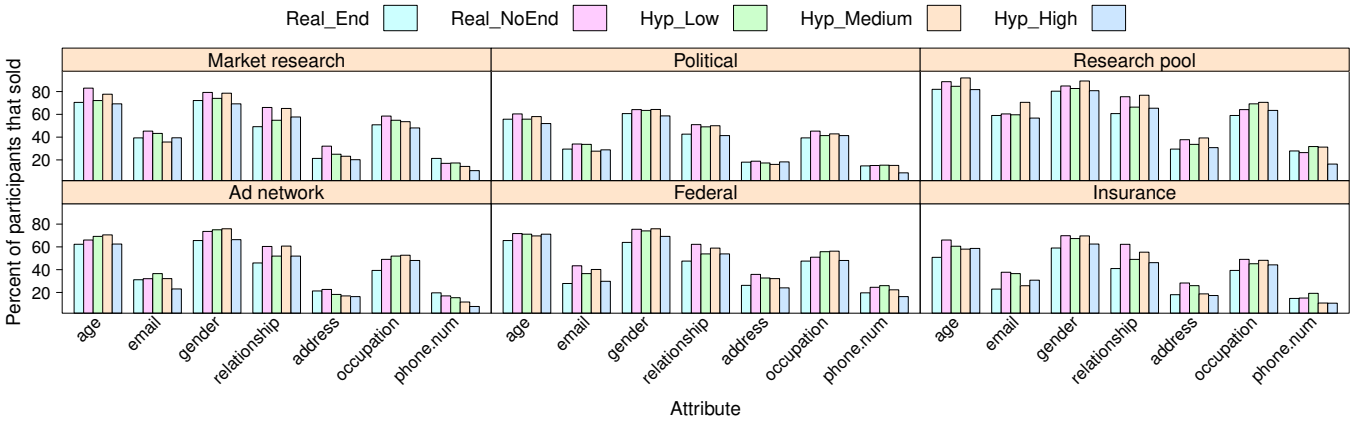

Figure 3: Percent of participants that chose to sell an attribute. Shown for different attributes, receiving parties, and levels of realism of the elicitation scenario.

demographic variables: age (centered by subtracting the mean), gender, highest level of education, primary occupation, income range, ethnicity, and the three IUIPC subscales (control, awareness, and collection, each centered). Table 2 lists the variables present in the model after performing model selection.[6] Wald tests for each fixed effect in the final model other than condition[7] were statistically significant ($p \leq .04$).

The results of our regression model are summarized in Table 2. For this model, 74.8% of the outcome variance was explained by the combination of fixed effects and the random participant intercept, while 13.3% of the variance was explained by the fixed effects alone. This suggests that knowledge of attribute type, receiving party, realism with which the elicitation scenario is presented, and the demographic variables we considered is insufficient to accurately predict an arbitrary individual's dollar valuations. Instead, dollar valuations more strongly depend on a combination of these factors and latent factors relating to each individual's preferences that we did not measure (i.e., the random participant intercept). Our model includes interactions between condition and receiving party and between condition and attribute, suggesting values differ based on the combination of scenario realism and either party or attribute. However, our final model does not include an interaction between attribute and party, suggesting their combination does not greatly impact values. The factors present in our final model were consistent with earlier models that we trained on separate pilot data.

For individuals with IUIPC-collection scores equal to the average in our sample (5.7), our model predicts the highest valuation ($24.22) for phone number in the least realistic scenario (Hyp_{High}) when selling to political parties. Given the same IUIPC-collection score, the lowest predicted value ($2.97) is for gender when selling to

[6] In Tables 2, 3, and 4, variables corresponding to interaction terms are shown only if statistically significant.
[7] Since interaction terms involving condition were included in the model, condition was also retained.

a research pool in the scenario where elicitation occurs prior to attribute collection ($Real_{NoEnd}$).

Attributes that could be used to contact participants—phone number, home address, and email address—were sold for higher amounts than other attributes. For receiving parties, participants listed higher amounts when selling to political parties, advertising networks, market research companies, and insurance companies; and lower for federal agencies and research pools. Based on our model, the type of attribute tends to have a larger effect on selling price than receiving party.

Comparing $Real_{End}$ to $Real_{NoEnd}$, we do not find a general trend of valuations being less if participants had already given us their Google attributes (the coefficient for $Real_{End}$ is not significantly different than 0), which one might expect because of the endowment effect. However, we do find instances where the interaction between $Real_{End}$ and certain attribute types is statistically significant: on average, participants sold their home address and phone number for significantly less in $Real_{End}$ (phone number: \$9.20, home address: \$8.01) than in $Real_{NoEnd}$ (phone number: \$13.82, home address: \$10.60).

We expected to find that valuations decreased as the realism of the elicitation scenario increased. Such a trend would also have been consistent with previous results on hypothetical bias effects [22, 56]. We did not find evidence of such a trend in general, however. The only potential exception to this is the abovementioned finding that specific attributes in $Real_{End}$were sold for significantly less than in $Real_{NoEnd}$. Although the primary difference between these conditions is in their handling of endowment rather than realism, one could consider $Real_{End}$to be more "realistic" than $Real_{NoEnd}$because in practice the attributes that users are selling are often already technically available to third parties (as participants are explicitly reminded in $Real_{End}$but not in $Real_{NoEnd}$).

Last, we checked whether participants would update their initial valuations if given the chance. Recall that as part of our study design, we presented participants in $Real_{NoEnd}$ with the option to revisit their initial valuations after logging in to Google. We hypothesized that participants would become more conscious about their privacy once shown the information they shared with us, and, in consequence, would increase their valuations or decide not to sell some of their attributes. Our observations contradicted this hypothesis: we found that 42 out of the 53 participants in $Real_{NoEnd}$ (~80%) did not update their initial valuations, while three participants decided to sell more attributes. Of the remaining eight participants, two decided not to sell some of their attributes, and six increased a subset of their valuations by 1.25× to 5.25×. We speculate that endowment effects may have affected most participants' decisions not to change their initial valuations after sharing their attributes with us.

To understand how our outlier exclusion criteria affected our results, we performed the same regression on our data with outliers included. All factors retained in the outliers-excluded model were also present in the outliers-included model. In our outliers-included model, gender, *IUIPC-control*, and the interaction between attribute and party were additionally retained.

Parameter estimates for the outliers-included model can be found in App. B.

Parameter	Est.	95% CI	p-value
(Intercept)	3.90	[2.88, 5.19]	<**.01**
condition ($Real_{NoEnd}$)			
$Real_{End}$	0.14	[-0.19, 0.6]	.46
Hyp_{Low}	0.09	[-0.19, 0.45]	.57
Hyp_{Medium}	-0.04	[-0.28, 0.28]	.79
Hyp_{High}	0.28	[-0.04, 0.72]	.10
attribute (age)			
email	0.43	[0.29, 0.59]	<**.01**
gender	-0.19	[-0.26, -0.11]	<**.01**
relationship	-0.05	[-0.13, 0.04]	.28
address	1.09	[0.85, 1.36]	<**.01**
occupation	0.22	[0.11, 0.35]	<**.01**
phone.num	1.67	[1.32, 2.08]	<**.01**
party (research pool)			
ad.network	0.15	[0.04, 0.26]	<**.01**
federal	0.04	[-0.06, 0.14]	.45
insurance	0.21	[0.1, 0.33]	<**.01**
market	0.20	[0.09, 0.31]	<**.01**
political	0.21	[0.09, 0.34]	<**.01**
iuipc.collection.centered	0.10	[0.03, 0.17]	<**.01**
$Real_{End}$:address	-0.33	[-0.45, -0.18]	<**.01**
$Real_{End}$:phone.num	-0.40	[-0.52, -0.26]	<**.01**

Table 2: Parameter estimates for our values mixed model (baselines are reported in parentheses). Estimates are back-transformed into dollar amounts. Interaction terms are shown if statistically significant. Bold p-values are statistically significant. Predictions for specific factor values can be computed by transforming parameter estimates using log(x+1), summing the transforms, and retransforming to dollars.

4.2.2 Declining to sell attributes. Another way to characterize how much participants value their personal information is by the number of attributes they would prefer not to sell at any price. For each receiving party and condition, Fig. 3 presents the percentage of participants that agreed to sell their attributes. While participants were less likely to sell certain attributes (e.g., phone number), they were likely to sell others (e.g., age).

We built a mixed-effects logistic regression model to understand the factors affecting participants' likelihood of selling personal attributes. Model selection began from a model that included fixed effects for condition, attribute, party, and all two-way interactions between condition, attribute, and party. Additionally, it included the only demographic variable present in the final regression model on dollar values (IUIPC collection score, centered). After model selection, the interaction between condition and party was not included; all other variables were retained in the final not-sold model. In particular, in contrast to the regression model for values, the interaction between attribute and party was retained. Tests of fixed effects for model comparison were performed using likelihood ratio tests; *IUIPC-collection* and the two retained interactions terms were found to be statistically significant ($p < 0.001$).

The results of this model are shown in Table 3. For this model, 81.2% of the outcome variance was explained by the combination of fixed effects and the random participant intercept, while 28.4% was explained by the fixed effects alone. Similarly as with attribute values, the small variance explained by the fixed effects suggests that, on their own, they may be insufficient to predict a particular individual's decision to sell. While a mixed-effect model can account for a significant degree of variance in participants' decision to sell (where the model is trained on those participants' valuation data), other factors may need to be considered for accurate predictions for individuals not appearing in the training data.

For an individual with IUIPC-collection score equal to the average in our sample (5.7), our model predicts the probability of selling to be lowest for phone number when selling to an political party in the least realistic scenario (Hyp_{High}), at 0.4%. In contrast, our model predicts the highest odds of selling for age when selling to a research pool, in a scenario where personal information has already been provided ($Real_{End}$), at 98%.

Attribute types that can be used to contact a person were generally the least likely to be sold, particularly phone number and address. This is similar to he trend we observed for dollar valuations (Section 4.2.1). However, the extent to which an attribute was less likely to be sold compared to our baseline attribute type and receiving party (age and research pool, respectively) depended on the particular combination of the two. For example, for many (but not all) attributes, our model predicts increased odds of selling when selling to federal agencies, insurance companies, or political parties.

Again, as was the case for dollar values, we did not find a general difference in selling decisions between $Real_{End}$ and $Real_{NoEnd}$ participants.

We had expected that hypothetical bias would lead to more participants refusing to sell attributes in the hypothetical conditions than in the realistic conditions. However, we did not observe statistically significantly higher rates of refusal either in general or for particular attributes.

Parameter	Est. Odds	Odds 95% CI	p-val
(Intercept)	63.71	[23.05, 176.11]	**<.01**
condition ($Real_{NoEnd}$)			
$Real_{End}$	0.35	[0.09, 1.32]	.12
Hyp_{low}	0.78	[0.25, 2.45]	.67
Hyp_{Medium}	0.66	[0.21, 2.06]	.48
Hyp_{High}	0.38	[0.12, 1.19]	.10
attribute (age)			
email	0.05	[0.03, 0.11]	**<.01**
gender	0.72	[0.34, 1.54]	.40
relationship	0.18	[0.09, 0.36]	**<.01**
address	0.01	[0.00, 0.01]	**<.01**
occupation	0.07	[0.03, 0.13]	**<.01**
phone.num	<0.01	[0.00, 0.01]	**<.01**
party (research pool)			
ad.network	0.09	[0.05, 0.14]	**<.01**
federal	0.12	[0.07, 0.2]	**<.01**
insurance	0.05	[0.03, 0.08]	**<.01**
market	0.19	[0.11, 0.32]	**<.01**
political	0.04	[0.02, 0.06]	**<.01**
iuipc.collection.centered	0.55	[0.43, 0.71]	**<.01**
relationship:ad.network	3.42	[1.78, 6.57]	**.01**
address:federal	5.66	[2.9, 11.03]	**<.01**
phone.num:federal	4.29	[2.13, 8.64]	**<.01**
relationship:insurance	4.53	[2.37, 8.65]	**<.01**
address:insurance	4.89	[2.46, 9.74]	**<.01**
occupation:insurance	3.52	[1.85, 6.7]	**<.01**
relationship:political	4.46	[2.34, 8.52]	**<.01**
address:political	4.12	[2.04, 8.34]	**<.01**
occupation:political	3.60	[1.89, 6.86]	**<.01**
phone.num:political	4.15	[1.97, 8.76]	**<.01**

Table 3: Parameter estimates for the logistic regression model estimating how factors affect the likelihood of selling (baselines are reported in parentheses). Bold p-values are statistically significant. Interaction terms are shown only if statistically significant.

4.3 Attribute rankings

We found that rankings of attributes were stable across both condition and receiving parties. On average, the attributes had the following ranking (from least to most important): gender, age, relationship status, occupation, email address, home address, and phone number. The same average ranking was observed when considering specific conditions and receiving parties.

As a consequence of this stability, average rankings are reasonably good predictors of participants' actual rankings, including in the realistic conditions. For example, the average rankings exactly matched 55.12% of the combined $Real_{End}$ and $Real_{NoEnd}$ participants' rankings, with mean accuracy (measured via Kendall's τ) of 89.22% and mean NDCG of 95.85%. (We note that exact matching is quite difficult, given that there are 7! = 5040 possible rankings.) Because the average rankings are the same across conditions, this means that hypothetical studies can potentially be sufficient to learn about how users prioritize sharing their data.

To improve the accuracy of our predictions, we further tested whether full attribute rankings in realistic conditions can be predicted via users' input on a subset of anchor attributes. Essentially, this simulates a system that asks users to specify their preferences for sharing a few attributes (e.g., age and occupation) in order to predict the users' preferences for all seven attributes. As a proof of concept of such a system, we apply a trivial principle to predict rankings: $C(N, 2)$ comparisons can define an order among N items.[8]

We first selected the *anchor* attributes (i.e., the subset of attributes we simulated asking users about). In the experiments, we varied the number of anchor attributes between two and seven. For simplicity, we refer to the number of anchor attributes used in any particular experiment as x. To choose the x anchors, we divided the average ranking into x groups of (almost) equal sizes and chose the attribute with the maximum variance in ranking from each group. The intuition behind this selection method is that anchor attributes

[8] $C(N, K)$ denotes the binomial coefficient $\binom{N}{K}$.

with high variance contain more information than attributes with low variance about the rankings of other attributes.

Once the x anchor attributes were selected, we represented rankings as a $C(x, 2)$ feature vector: one feature for each pair of anchor attributes a_i and a_j, where a value of 1 indicated that $a_i > a_j$, 0 indicated that $a_i = a_j$, and -1 indicated that $a_i < a_j$. We then used a training set of such feature vectors to train $C(7, 2)$ classifiers, or 21 classifiers, one for each pair of attributes. Each classifier attempts to predict (based on the input anchor attributes) which of its two targeted attributes should be ranked higher. To predict overall ranking, we ordered the attributes according to the number of times they were predicted to be more important than others. We used Gentle AdaBoost [52] to train the classifiers that compare attribute pairs. Prior work that used similar ranking algorithms (especially in the area of information retrieval) was often constrained by the need to be computationally efficient, which necessitated the use of less intuitive ranking algorithms (e.g., [24]).

We evaluated the performance of our prototype by training models on data from either hypothetical or realistic conditions, and testing the models on data from real conditions. We performed five cross-validation rounds, each time selecting 90 random participants for training the models, and evaluating on the rankings of 24 the remaining participants from the realistic conditions. Figure 4 shows the prediction performance. The horizontal dashed lines represent the performance (in exact match, Kendall's τ accuracy, and NDCG) of predicting using only averages, without anchor attributes; the solid and dot-dash lines represent performance in the same metrics when using x anchors, for hypothetical and realistic training respectively.

Overall, our results suggest that it is possible to predict rankings with high accuracy while only asking users a small number of potentially disruptive questions. For example, by asking a user to rank three attributes, one can predict the full rankings better than when using the average rankings alone. It is also important to note that the performance of models which were trained on hypothetical data was comparable to the performance of models trained on realistic data, providing further evidence that hypothetical scenarios may be sufficient to learn about users' real rankings.

4.4 Study dropouts

A number of participants who began our study ultimately dropped out. To investigate why they dropped out, we compared dropout rates between conditions. We considered a participant as beginning the study once they had accepted the study consent form, and considered a participant as dropping out of the study if they did not successfully enter the completion code on Prolific (the completion code was shown at the end of the survey). Based on this criteria, we had 31 dropouts in $Real_{End}$ (5.5%), 45 dropouts in $Real_{NoEnd}$ (8.0%), 11 dropouts in Hyp_{Low} (2.0%), 14 dropouts in Hyp_{Medium} (2.5%), and 6 dropouts in Hyp_{High} (1.1%). Differences in dropout rates by condition were statistically significant ($\chi^2(4) = 47.2, p < .01$). Comparing pairs of conditions, all *Hyp* conditions significantly differed in dropout rate relative to $Real_{NoEnd}$ (Holm-Bonferroni-corrected Fisher's Exact Test, $p \leq .01$). We did not find statistically significant differences in dropout rates between $Real_{End}$ and $Real_{NoEnd}$, Hyp_{Low} and Hyp_{Medium}, nor between Hyp_{Medium} and Hyp_{High}.

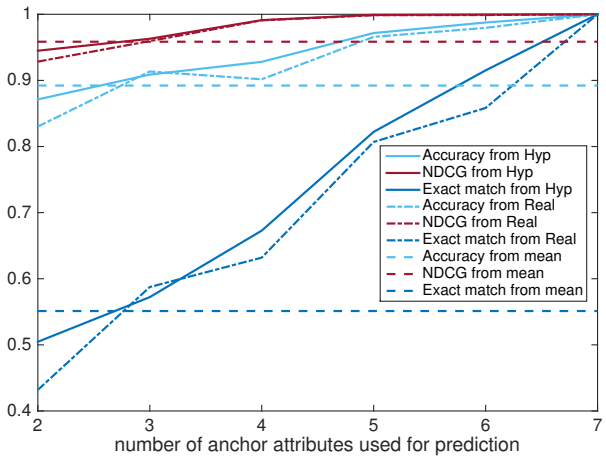

Figure 4: The performance of ranking prediction as a function of the number of anchor attributes. The horizontal dashed lines show the performance when using the average rankings for prediction.

We examined whether dropout rate differences between *Real* and *Hyp* participants biased our experimental groups in such a way that, among those who did not drop out, *Real* participants were less privacy-sensitive than their *Hyp* counterparts. We did not find significant differences in IUIPC scores (or any measured demographic variable) between *Real* and *Hyp* participants (all p>.08). Furthermore, our valuation results support the idea that *Real* dropouts did not lead to less privacy-sensitive participants in *Real* than in *Hyp* conditions. If *Real* dropouts tended to drop out for privacy reasons, we would expect *Hyp* participants to have overall higher selling prices or increased likelihood of not selling attributes (since privacy-sensitive *Hyp* participants would likely be retained). Thus, the expected effect on *Real* values due to dropouts would be in the same direction as the expected change on *Real* values due to hypothetical bias. The fact that we did not observe statistically significant differences in attribute valuations for the majority of cases suggests that dropouts did not introduce a bias in privacy-sensitivity between experimental groups, and that our experimental manipulation did not affect valuations for those cases.

5 DISCUSSION

We discuss the implications of our findings for designing studies to measure privacy preferences and for designing systems that incorporate privacy valuation.

5.1 Hypothetical bias

Hypothetical bias is typically measured via *calibration factors*, or the hypothetical value of a good divided by its true value. For most goods, people typically overestimate their selling price in hypothetical scenarios, resulting in calibration factors that are greater than 1.

In this study, we examined hypothetical bias for privacy valuation. We did not find a general trend of higher hypothetical values than realistic values for attributes. Based on predictions from our

fitted regression model, the degree of hypothetical bias we observed was relatively small. For example, the largest average attribute calibration factor predicted by our model (when selling to advertising networks in Hyp_{High}) was 1.61, much smaller than what List and Gallet found for public and private goods (4.44 and 8.41, respectively) [33]. We did not find significant differences in the relative ordering of attributes between hypothetical and realistic conditions. These perhaps counterintuitive findings reinforce how nuanced and dependent on context people's privacy preferences are, while also illustrating a way of roughly capturing these preferences, despite the challenges involved.

Lastly, we observed a different kind of biasing effect of real or realistic studies, which may be important to consider for valuations of privacy goods. We found significantly increased dropout rates among our realistic-scenario participants compared to hypothetical-scenario ones. We speculate that dropouts in our realistic scenarios may have been more privacy-concerned, which is supported by pilot data in which dropouts were more likely to not sell attributes during initial valuations (i.e. before the Google SSO request). The potential for the most privacy-concerned to drop out should therefore be be taken into account when putting study results to practical use or deciding how to design future studies.

In particular, if the goal of a study is to understand not only selling prices but also decisions about whether to sell, then the study design should use a real or realistic scenario. Further, if the goal is to understand real behavior in settings where participation is mandatory (or necessary in practice, e.g., due to network effects), then researchers should work to reduce dropout rates. This might require increased compensation or assurances about data-handling practices. On the other hand, if the focus is purely on valuation, the level of realisim is less important, as neither the level of realism (as expressed through the five conditions) nor whether participants dropped out appeared to strongly influence dollar valuations.

5.2 Endowment effect

Although typically used to explain differences between people's willingness to accept compensation for a good and their willingness to pay for that good, the endowment effect can be applied to other contexts (e.g., perception of a company being more trustworthy if a person owns its product [3]).

We observed evidence of the endowment effect on the selling prices of phone number and home address, which were on average priced lower when information had already been shared. This effect was largest for phone number: an average of \$4.62 less in $Real_{End}$ than $Real_{NoEnd}$. One potential explanation for not observing this effect across all attributes may relate to the low selling prices of some attributes, which may have limited our ability to observe significant differences.

5.3 Implications for system design

Understanding how users value personal attributes can help a system designer to decide which attributes to ask a user to share with the system itself and with other users, or to choose default sharing and visibility settings. Understanding privacy valuation may also help with decisions about when and how to use customer data for targeted advertising.

Our findings suggest that system designers should carefully consider the impact of endowment effects; if users have already shared information, then their reported values may underestimate their unbiased values. Further, if the system only requires understanding priorities among attributes (e.g., to select which attributes to share in a single-sign-on setting), then it may be sufficient to collect valuations for a subset of anchor attributes.

5.4 Other considerations

Prior work suggests that users have difficulty valuating privacy in part due to uncertainty and malleability [1]. Our results provide further evidence for these ideas.

Participants' responses discussing the difficulty of assigning values to attributes demonstrated uncertainty about the potential consequences of selling attributes. Some potential consequences, like receiving spam, appeared well understood, but several participants provided comments such as, "It is unclear if any of the entities I am potentially selling my information to would be subject to non-disclosure and/or privacy/hacking policies." Participants were also uncertain about how to appropriately value their personal attributes. (This is perhaps unsurprising, since personal attributes are not typically thought of in terms of price in everyday situations.) Many participants mentioned wanting to see other people's values for context in setting their own.

Finally, our experimental results reinforce that notion that privacy valuations are malleable. For certain attributes, whether participants had already shared some personal information had a significant impact on their privacy valuations.

6 CONCLUSION

We conducted a between-subject online survey with 434 participants to study how users valuate personal information. Our study enhanced our understanding of prior work's findings by providing nuanced insight to how different factors (e.g., level of realism and endowment effects) work in concert to affect users' valuations. For example, we find that, for the attributes we studied, valuations elicited under realistic conditions are not overall lower than ones elicited under hypothetical conditions—contrary to what one would expect given the privacy paradox. We also found that, in contrast to the valuations, rankings of different types of personal information were independent of the factors we studied. We interpreted our findings and discussed how they shed light on the design of systems and user studies.

ACKNOWLEDGMENTS

This work was supported in part by CyLab at Carnegie Mellon University via a CyLab Presidential Fellowship; by NSF grant DGE-0903659; by a grant awarded to the University Corporation for Advanced Internet Development (Internet2) under the sponsorship of the U.S. Department of Commerce, National Institute of Standards and Technology; and by a gift from Google.

REFERENCES

[1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
[2] Alessandro Acquisti, Leslie K. John, and George Loewenstein. 2013. What is privacy worth? *The Journal of Legal Studies* 42, 2 (2013), 249–274.

[3] Sadia Afroz, Aylin Caliskan Islam, Jordan Santell, Aaron Chapin, and Rachel Greenstadt. 2013. How privacy flaws affect consumer perception. In *Proc. STAST*.

[4] Kamil Bartoń. 2016. MuMIn: Multi-model inference. https://cran.r-project.org/package=MuMIn. (2016).

[5] Alastair R. Beresford, Dorothea Kübler, and Sören Preibusch. 2012. Unwillingness to pay for privacy: A field experiment. *Economics Letters* 117, 1 (2012), 25–27.

[6] Juan Pablo Carrascal, Christopher Riederer, Vijay Erramilli, Mauro Cherubini, and Rodrigo de Oliveira. 2013. Your browsing behavior for a big mac: Economics of personal information online. In *Proc. WWW*.

[7] Richard T. Carson, Nicholas E. Flores, Kerry M. Martin, and Jennifer L. Wright. 1996. Contingent valuation and revealed preference methodologies: Comparing the estimates for quasi-public goods. *Land economics* (1996), 80–99.

[8] Farah Chanchary and Sonia Chiasson. 2015. User perceptions of sharing, advertising, and tracking. In *Proc. SOUPS*.

[9] Jae Bong Chang, Jayson L. Lusk, and F. Bailey Norwood. 2009. How closely do hypothetical surveys and laboratory experiments predict field behavior? *American Journal of Agricultural Economics* 91, 2 (2009), 518–534.

[10] Kay Connelly, Ashraf Khalil, and Yong Liu. 2007. Do I do what I say?: Observed versus stated privacy preferences. In *Human-Computer Interaction*. Springer, 620–623.

[11] Elisa Costante, Jerry Den Hartog, and Milan Petkovic. 2011. On-line trust perception: What really matters. In *Proc. STAST*.

[12] Dan Cvrcek, Marek Kumpost, Vashek Matyas, and George Danezis. 2006. A study on the value of location privacy. In *Proc. WPES*.

[13] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated experiments on ad privacy settings. *Proc. PETS* 2015, 1 (2015), 92–112.

[14] Leyla Dogruel, Sven Joeckel, and Jessica Vitak. 2017. The valuation of privacy premium features for smartphone apps: The influence of defaults and expert recommendations. *Computers in Human Behavior* (2017).

[15] Serge Egelman. 2013. My profile is my password, verify me! The privacy/convenience tradeoff of Facebook Connect. In *Proc. CHI*.

[16] Serge Egelman, Adrienne Porter Felt, and David Wagner. 2012. Choice architecture and smartphone privacy: There's a price for that. In *Proc. WEIS*.

[17] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. 2009. Timing is everything?: The effects of timing and placement of online privacy indicators. In *Proc. CHI*.

[18] Jens Grossklags and Alessandro Acquisti. 2007. When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In *WEIS*.

[19] Il-Horn Hann, Kai-Lung Hui, Sang-Yong Tom Lee, and Ivan P.L. Png. 2007. Overcoming online information privacy concerns: An information-processing theory approach. *Journal of Management Information Systems* 24, 2 (2007), 13–42.

[20] Bernardo A. Huberman, Eytan Adar, and Leslie R. Fine. 2005. Valuating privacy. *IEEE Security & Privacy* 3, 5 (2005), 22–25.

[21] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.

[22] Carlos Jensen, Colin Potts, and Christian Jensen. 2005. Privacy practices of Internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies* 63, 1 (2005), 203–227.

[23] N. Jentzsch, S. Preibusch, and A. Harasser. 2012. Study on monetising privacy. *ENISA* (2012).

[24] Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *Proc. KDD*.

[25] Adam N. Joinson, Ulf-Dietrich Reips, Tom Buchanan, and Carina B. Paine Schofield. 2010. Privacy, trust, and self-disclosure online. *Human–Computer Interaction* 25, 1 (2010), 1–24.

[26] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara B. Kiesler. 2014. Privacy attitudes of Mechanical Turk workers and the US public. In *Proc. SOUPS*.

[27] Bart P. Knijnenburg and Alfred Kobsa. 2013. Making decisions about privacy: Information disclosure in context-aware recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3, 3 (2013), 20.

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

[29] Hanna Krasnova, Thomas Hildebrand, and Oliver Guenther. 2009. Investigating the value of privacy in online social networks: Conjoint analysis. In *Proc. ICIS*.

[30] Vijay Krishna. 2009. *Auction theory*. Academic press.

[31] Kenneth C. Laudon. 1996. Markets and privacy. *Commun. ACM* 39, 9 (1996), 92–104.

[32] Pedro G. Leon, Blase Ur, Yang Wang, Manya Sleeper, Rebecca Balebako, Richard Shay, Lujo Bauer, Mihai Christodorescu, and Lorrie Faith Cranor. 2013. What matters to users? Factors that affect users' willingness to share information with online advertisers. In *Proc. SOUPS*.

[33] John A. List and Craig A. Gallet. 2001. What experimental protocol influence disparities between actual and hypothetical stated values? *Environmental and Resource Economics* 20, 3 (2001), 241–254.

[34] Joseph Little and Robert Berrens. 2004. Explaining disparities between actual and hypothetical stated values: Further investigation using meta-analysis. *Economics Bulletin* 3, 6 (2004), 1–13.

[35] Miguel Malheiros, Sacha Brostoff, Charlene Jennett, and M. Angela Sasse. 2013. Would you sell your mother's data? Personal data disclosure in a simulated credit card application. In *The Economics of Information Security and Privacy*. Springer, 237–261.

[36] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.

[37] Matt Miller. 2017. Lawsuits rolling in over Equifax data breach. https://goo.gl/wcU16m. (2017). Accessed: 2018-02-28.

[38] James J. Murphy, P. Geoffrey Allen, Thomas H. Stevens, and Darryl Weatherhead. 2005. A meta-analysis of hypothetical bias in stated preference valuation. *Environmental and Resource Economics* 30, 3 (2005), 313–325.

[39] Shinichi Nakagawa and Holger Schielzeth. 2013. A general and simple method for obtaining R^2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution* 4, 2 (2013), 133–142.

[40] Helen Nissenbaum. 2011. A contextual approach to privacy online. *Daedalus* 140, 4 (2011), 32–48.

[41] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100–126.

[42] Judith S. Olson, Jonathan Grudin, and Eric Horvitz. 2005. A study of preferences for sharing and privacy. In *Proc. CHI*.

[43] Saurabh Panjwani, Nisheeth Shrivastava, Saurabh Shukla, and Sharad Jaiswal. 2013. Understanding the privacy-personalization dilemma for web search: A user perspective. In *Proc. CHI*.

[44] Delroy L. Paulhus. 1991. Measurement and control of response bias. (1991).

[45] Jose Pinheiro, Douglas Bates, Saikat DebRoy, Deepayan Sarkar, and R Core Team. 2015. *nlme: Linear and nonlinear mixed effects models*. R package version 3.1-122.

[46] Aarathi Prasad, Jacob Sorber, Timothy Stablein, Denise Anthony, and David Kotz. 2012. Understanding sharing preferences and behavior for mHealth devices. In *Proc. WPES*.

[47] Sören Preibusch. 2013. The value of privacy in Web search. In *Proc. WEIS*.

[48] Sören Preibusch, Kat Krol, and Alastair R Beresford. 2013. The privacy economics of voluntary over-disclosure in Web forms. In *The Economics of Information Security and Privacy*. Springer, 183–209.

[49] Prolific. 2014. Prolific (crowdsourcing platform). https://www.prolific.ac/. (2014). Accessed: 2018-02-28.

[50] Yu Pu and Jens Grossklags. 2015. Towards a model on the factors influencing social app users' valuation of interdependent privacy. *Proc. PETS* 2016, 2 (2015), 61–81.

[51] Yu Pu and Jens Grossklags. 2015. Using conjoint analysis to investigate the value of interdependent privacy in social app adoption scenarios. In *Proc. ICIS*.

[52] Robert E. Schapire and Yoram Singer. 1999. Improved boosting algorithms using confidence-rated predictions. *Machine learning* 37, 3 (1999), 297–336.

[53] Gideon Schwarz. 1978. Estimating the dimension of a model. *The Annals of Statistics* 6, 2 (1978), 461–464.

[54] Fuming Shih, Ilaria Liccardi, and Daniel Weitzner. 2015. Privacy tipping points in smartphones privacy preferences. In *Proc. CHI*.

[55] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proc. CHI*.

[56] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. 2001. E-privacy in 2nd generation E-commerce: Privacy preferences versus actual behavior. In *Proc. EC*.

[57] Sarah Spiekermann, Jana Korunovska, and Christine Bauer. 2012. Psychology of ownership and asset defense: Why people value their personal information beyond privacy. *Available at SSRN 2148886* (2012).

[58] Jacopo Staiano, Nuria Oliver, Bruno Lepri, Rodrigo de Oliveira, Michele Caraviello, and Nicu Sebe. 2014. Money walks: A human-centric study on the economics of personal mobile data. In *Proc. Ubicomp*.

[59] San-Tsai Sun, Eric Pospisil, Ildar Muslukhov, Nuray Dindar, Kirstie Hawkey, and Konstantin Beznosov. 2011. What makes users refuse web single sign-on? An empirical investigation of OpenID. In *Proc. SOUPS*.

[60] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research* 22, 2 (2011), 254–268.

[61] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proc. SOUPS*.

[62] Hal R. Varian. 1996. Economic aspects of personal privacy. https://goo.gl/fiAWk1. (1996). Accessed: 2018-02-28.

[63] Lukasz Walasek, Rebecca J. Wright, and Tim Rakow. 2014. Ownership status and the representation of assets of uncertain value: The balloon endowment risk task (BERT). *Journal of Behavioral Decision Making* 27, 5 (2014), 419–432.

[64] Allison Woodruff, Vasyl Pihur, Sunny Consolvo, Lauren Schmidt, Laura Brandimarte, and Alessandro Acquisti. 2014. Would a privacy fundamentalist sell their DNA for $1000... if nothing bad happened as a result? The Westin categories, behavioral intentions, and consequences. In *Proc. SOUPS*.

A PROTOCOL

[All conditions: Consent form]
[$Real_{End}$: Login with Google SSO]
[All conditions: Distraction task]

$Real_{End}$: Introducing the information marketplace:

Thank you for your feedback on our eyeglasses designs!

We have partnered with other researchers at Carnegie Mellon University who are creating an Information Market. **You have been selected to receive an offer to participate in this market.**

Please carefully read about how this market works:

The Information Market aggregates information about individuals and sells it to interested 3rd parties, such as research pools or advertising networks. Individuals (like you) can set a price for each attribute they are interested in selling (e.g., city where live, level of education) or decide not to sell an attribute at all. In addition, individuals can set a different price for each kind of 3rd party. For example, individuals can decide to sell information about the city they live in to a research pool, but never to sell this type of information to political parties.

For more information, please visit: [URL]

$Real_{NoEnd}$: Introducing the information marketplace:

Thank you for your feedback on our eyeglasses designs!

We have partnered with other researchers at Carnegie Mellon University who are creating an Information Market. **You have been selected to receive an offer to participate in this market.**

Please carefully read about how this market works:

The Information Market aggregates information about individuals and sells it to interested 3rd parties, such as research pools or advertising networks. Individuals (like you) can set a price for each attribute they are interested in selling (e.g., city where live, level of education) or decide not to sell an attribute at all. In addition, individuals can set a different price for each kind of 3rd party. For example, individuals can decide to sell information about the city they live in to a research pool, but never to sell this type of information to political parties.

For more information, please visit: [URL]

Hyp_{Low}: Introducing the information marketplace:

Thank you for your feedback on our eyeglasses designs!

We have partnered with other researchers at Carnegie Mellon University who are creating an Information Market. **You have been selected to receive an offer to help evaluate this market, which will begin operating soon.** As the marketplace is not yet operating, your answers will only be used for research and no actual exchange will be performed with the marketplace (i.e., you won't need to actually share your information, and you won't earn any money).

Please carefully read about how this market works:

The Information Market aggregates information about individuals and sells it to interested 3rd parties, such as research pools or advertising networks. Individuals (like you) can set a price for each attribute they are interested in selling (e.g., city where live, level of education) or decide not to sell an attribute at all. In addition, individuals can set a different price for each kind of 3rd party. For example, individuals can decide to sell information about the city they live in to a research pool, but never to sell this type of information to political parties.

For more information, please visit: [URL]

Hyp_{Medium}: Introducing the information marketplace:

Thank you for your feedback on our eyeglasses designs!

We have partnered with other researchers at Carnegie Mellon University who are exploring the concept of an Information Market. **You have been selected to receive an offer to help evaluate this market concept.** As the marketplace is imaginary and not actually operating, your answers will be used only for research and no actual exchange will be performed with the marketplace (i.e., you won't need to actually share your information, and you won't earn any money).

To help us evaluate this new marketplace concept, please imagine the following scenario:

The Information Market aggregates information about individuals and sells it to interested 3rd parties, such as research pools or advertising networks. Individuals (like you) can set a price for each attribute they are interested in selling (e.g., city where live, level of education) or decide not to sell an attribute at all. In addition, individuals can set a different price for each kind of 3rd party. For example, individuals can decide to sell information about the city they live in to a research pool, but never to sell this type of information to political parties.

Hyp_{High}: Introducing the information marketplace:

Thank you for your feedback on our eyeglasses designs!

We have partnered with other researchers at Carnegie Mellon University who are studying buying and selling preferences for personal information (e.g., age, education level) . **You have been selected to receive an offer to participate in this research.**

Please imagine the following scenario:

The Information Market aggregates information about individuals and sells it to interested 3rd parties, such as research pools or advertising networks. Individuals (like you) can set a price for each attribute they are interested in selling (e.g., city where live, level of education) or decide not to sell an attribute at all. In addition, individuals can set a different price for each kind of 3rd party. For example, individuals can decide to sell information about the city they live in to a research pool, but never to sell this type of information to political parties.

[All conditions: First attention question (if the participant fails to answer correctly the first time, return to the introduction and reask the question)]

All conditions: Market's means of operation:

How the Information Market operates:

There is no direct contact between the 3rd parties that purchase information and the individuals selling that information. All transactions occur within the marketplace.

A 3rd party can ask the Information Market for data about individuals with specific profiles, such as those in a particular income range. For a given 3rd party budget, the Information Market will spend that budget by purchasing the lowest priced information first. If your information is sold, you will be compensated an amount higher or equal to the price you assigned to it. (This is similar to how eBay works, but in reverse: on eBay you would have to pay an amount equal to or lesser than your highest bid.)

Here is an example:

An ad network with a $10 budget can ask the Information Market for home address information. Let's assume that Dan and Carol are the only individuals willing to sell their home addresses. Dan is willing to sell his home address for $5 and Carol is willing to sell hers for $6. Since the ad network cannot afford to purchase both Dan and Carol's information, it will be instead only purchase Dan's.

Please note that the numbers used here are only for illustration. A seller in the Information Market may assign any price he/she sees suitable on his/her information (e.g., $1, $10, $100).

[All conditions: Second attention question (if the participant fails to answer correctly the first time, return to Information Market's means of operation and reask the question)]

$Real_{End}$: Reminder:

In what follows, please remember: You have already shared some information with us by logging into Google.

Here is some of the data you have shared with us: [show information.]

If you decide to sell information about yourself and it is bought by a 3rd party, we will automatically share your information with that party. If additional information about you is required, we will contact you with instructions on how to provide it and payment will be contingent upon that information being provided. For example, we may ask you to install a browser extension in order to share your browsing history with us.

$Real_{NoEnd}$: Reminder:

In what follows, please remember: If you decide to sell information about yourself and it is bought by a 3rd party, we will automatically share your information with that party. If additional information about you is required, we will contact you with instructions on how to provide it and payment will be contingent upon that information being provided. For example, we may ask you to install a browser extension in order to share your browsing history with us.

Hyp_{Low}: Reminder:

In what follows, please remember: As the Information Market is still not operating, your answers will be used only for research and no actual exchange will be performed with the marketplace (i.e., you won't need to actually share your information, and you won't actually be paid for your attributes).

Hyp_{Medium}: Reminder:

In what follows, please remember: As the Information Market is imaginary and not actually operating, your answers will be used only for research and no actual exchange will be performed (i.e., you won't need to actually share your information, and you won't actually be paid for your attributes).

Hyp_{High}: Reminder:

In what follows, please remember: The Information Market is completely hypothetical. Your answers will be used only for research and no actual exchange will be performed (i.e., you won't need to actually share your information, and you won't actually be paid for your attributes).

All conditions: Valuation task:

For each pair of an attribute and a third party, please mark whether you would sell the attribute to the 3rd party for a certain price, or refuse to sell. Either enter a dollar amount or check "do not sell."

[The following input dialogue was used for each of the seven attributes (note that javascript was used to highlight missing rows and prevent the insertion of dollar amounts when the "Do Not Sell" option was selected)]

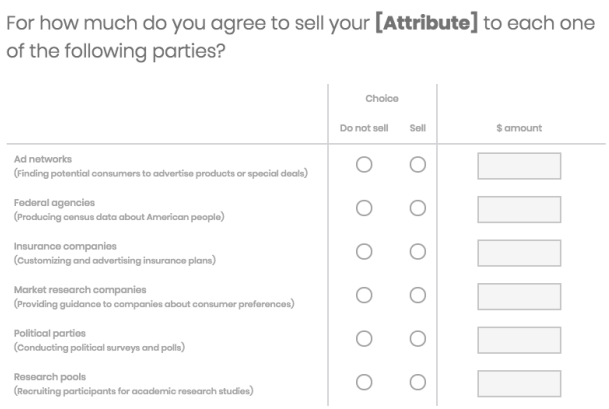

[$Real_{NoEnd}$: Google SSO login]
[$Real_{NoEnd}$: Present data that we were able to collect from the participant's Google account]

$Real_{NoEnd}$: Revisit valuations:

You are given the option to update the answers you've previously provided, if you like.

For each pair of attribute and third party, please specify a selling price or refusal to sell. Either enter a dollar amount or check "Do not sell."

[Show same dialogue as before]

[Ask participant's in $Real_{NoEnd}$ who changed their initial valuations about what led to their decision]

Real$_{End}$ and *Real*$_{NoEnd}$: Debrief:

Thank you for participating in our research.

One challenge in studying privacy decision making is that if participants are aware that researchers are studying privacy behavior, they may change their behavior in response. As a result, to capture more natural behavior, it is sometimes necessary for researchers to deceive study participants. In this study, we told you that we would sell your information to an ad network, but that was not true. We will not sell your information to anyone. However, you will receive an extra $1.50 in compensation instead. As researchers, we take the privacy of our participants very seriously, and we will use your data only for research and protect it carefully.

By allowing you to believe that your information could be sold for real money, we hoped to capture a more realistic valuation of your privacy than if we had just asked you to imagine a hypothetical scenario. We believe this approach will provide important information about how people value the privacy of their data.

As part of our obligation to protect the safety of our participants, we submitted our study for review by Carnegie Mellon University's Institutional Review Board (also known as an ethics board), which approved our research. However, if you have any concerns about the study, please share them with us below.

[Ask participants about whether they would provide different answers outside of a study]
[*Real*$_{End}$ and *Real*$_{NoEnd}$: Ask about attribute validity (highlight that participants' answers would not affect their compensation)]
[All Conditions: Ask IUIPC and demographic questions]

B VALUES REGRESSION MODEL INCLUDING OUTLIERS

The parameter estimates for the mixed-effects regression model when including outliers are reported in Table 4.

Parameter	**Est.**	**95% CI**	**p-value**
(Intercept)	4.33	[2.5, 7.1]	< **.01**
condition (Real$_{NoEnd}$*)*			
Real$_{End}$	0.27	[-0.26, 1.19]	.38
Hyp$_{Low}$	0.10	[-0.32, 0.78]	.69
Hyp$_{Medium}$	-0.01	[-0.38, 0.58]	.97
Hyp$_{High}$	0.05	[-0.35, 0.7]	.84
attribute (age)			
email	0.37	[0.13, 0.65]	**.01**
gender	-0.11	[-0.25, 0.04]	.14
relationship	0.04	[-0.12, 0.24]	.65
address	1.05	[0.65, 1.56]	< **.01**
occupation	0.34	[0.12, 0.61]	**.01**
phone.num	1.38	[0.84, 2.07]	< **.01**
party (research pool)			
ad.network	0.15	[-0.04, 0.37]	.13
federal	0.11	[-0.07, 0.31]	.26
partyinsurance	0.29	[0.08, 0.54]	< **.01**
market	0.26	[0.06, 0.49]	< **.01**
political	0.26	[0.05, 0.51]	**.01**
gender (female)			
gender.male	0.15	[-0.12, 0.5]	.30
gender.other	-0.20	[-0.82, 2.59]	.77
gender.prefer.na	43.24	[6.08, 275.59]	< **.01**
iuipc.control.centered	0.07	[-0.09, 0.25]	.41
iuipc.collection.centered	0.16	[0.02, 0.31]	**.02**
Real$_{End}$:address	0.81	[0.39, 1.36]	< **.01**
Real$_{End}$:phone.num	2.54	[1.63, 3.76]	< **.01**
Hyp$_{High}$:phone.num	0.74	[0.31, 1.32]	**.01**
Hyp$_{High}$:ad.network	0.40	[0.16, 0.68]	**.03**
Hyp$_{High}$:political	0.40	[0.16, 0.7]	**.04**

Table 4: Parameter estimates for our values mixed model with outliers included. Estimates are back-transformed into dollar amounts. Bold p-values are statistically significant. Interaction terms are shown only if statistically significant. Marginal R^2 = 0.13. Conditional R^2 = 0.79.