Computers in Human Behavior 61 (2016) 443–453

Contents lists available at ScienceDirect

Computers in Human Behavior

journal homepage: www.elsevier.com/locate/comphumbeh

Full length article

A methodology for estimating the value of privacy in information disclosure systems

Ron Hirschprung*, Eran Toch, Frank Bolton, Oded Maimon

Department of Industrial Engineering, The Iby and Aladar Fleischman Faculty of Engineering, Tel Aviv University, P.O. Box 39040, Tel Aviv, 6997801, Israel

ARTICLE INFO

Article history:
Received 2 December 2015
Received in revised form
10 February 2016
Accepted 10 March 2016

Keywords:
Online privacy
Value of privacy
Utility and privacy
Information disclosure
User study
Privacy in electronic commerce

ABSTRACT

In many types of information systems, users face an implicit tradeoff between disclosing personal information and receiving benefits, such as discounts by an electronic commerce service that requires users to divulge some personal information. While these benefits are relatively measurable, the value of privacy involved in disclosing the information is much less tangible, making it hard to design and evaluate information systems that manage personal information. Meanwhile, existing methods to assess and measure the value of privacy, such as self-reported questionnaires, are notoriously unrelated of real −world behavior. To overcome this obstacle, we propose a methodology called VOPE (Value of Privacy Estimator), which relies on behavioral economics' Prospect Theory (Kahneman & Tversky, 1979) and valuates people's privacy preferences in information disclosure scenarios. VOPE is based on an iterative and responsive methodology in which users take or leave a transaction that includes a component of information disclosure. To evaluate the method, we conduct an empirical experiment (n = 195), estimating people's privacy valuations in electronic commerce transactions. We report on the convergence of estimations and validate our results by comparing the values to theoretical projections of existing results (Tsai, Egelman, Cranor, & Acquisti, 2011), and to another independent experiment that required participants to rank the sensitivity of information disclosure transactions. Finally, we discuss how information systems designers and regulators can use VOPE to create and to oversee systems that balance privacy and utility.

1. Introduction

In many information systems, the user faces decisions that trade privacy with benefits. A typical exchange occurs in electronic commerce where the user receives a discount for providing private information (Ackerman, Cranor, & Reagle, 1999; Acquisti, 2004; Hann, Hui, Lee, & Png, 2002). Another common example is online social networks, such as Facebook, where the user gains social capital benefits thorough the disclosure of information but is also exposed to a loss of privacy (Dwyer, Hiltz, & Passerini, 2007; Min & Kim, 2015). We may define any application that allows the user to make decisions about information disclosure in return for some benefit as relevant to this tradeoff between tangible benefits and the loss of privacy. Several theories, such as the Privacy Calculus, assume that users make decisions by rationally maximizing their expected benefits against the possible cost of disclosing information (Culnan & Bies, 2003; Dinev & Hart, 2006; Li, Sarathy, & Xu, 2010). However, the intangibility of the value of privacy, the inherent uncertainty in privacy decisions (Jensen, Potts, & Jensen, 2005), its context-dependence, and its sensitivity to various biases make it a challenge to understand the utility and cost of privacy (Acquisti, Brandimarte, & Loewenstein, 2015; Baek, 2014; Kehr, Kowatsch, Wentzel, & Fleisch, 2015). In this context, uncertainty is prevalent due to limited transparency and the fact that users do not always know how their data will be used, or even how authentic is the electronic service (Featherman, Valacich, & Wells, 2006).

Experimental results show that customers are willing to pay a premium of about 5% to buy from more privacy-protecting vendors (Tsai et al., 2011). However, these results are contrasted against methods that aimed to assess the value of privacy through surveys, in which the premium is consistently higher. For example, participants requested a declared value of $39.83 to $49.78 (Hann et al., 2002) compared to a value of about $3.5 when analyzing actual

* Corresponding author.
E-mail addresses: ronyh@post.tau.ac.il (R. Hirschprung), erant@post.tau.ac.il (E. Toch), frankbolton@post.tau.ac.il (F. Bolton), maimon@eng.tau.ac.il (O. Maimon).

http://dx.doi.org/10.1016/j.chb.2016.03.033
0747-5632/

behavior (Tsai et al., 2011). Assessing the premium people will be willing to spend on privacy is notoriously difficult because of the abstract nature of privacy and the contradictory characteristics of users' privacy behavior (Acquisti, Friedman, & Telang, 2006; Longpre & Kreinovich, 2006). This discrepancy indicates the methodological challenges of estimating the value of privacy and its sensitivity to the method and context in which the value is estimated.

Understanding the value of privacy provides a basis for estimating the utility gains from a transaction and allows for optimizing information disclosure processes. To balance the benefits of information disclosure and the cost of privacy, both of these notions should be quantified using the same unit of comparison. For example, the same "currency", which allows it to be measured against other aspects of sharing and standardized between users and between scenarios. Technology designers can use this information to design better privacy controls and to adjust incentive systems. Governments can use this information to regulate privacy through markets, which require understanding of the subjective value of privacy (Spiekermann, Acquisti, Böhme, & Hui, 2015).

This paper proposes a methodology for estimating the value of privacy as the Willingness to Accept (WTA), that is, the lowest monetary reward users are willing to accept to divulge personal information, rather than the Willingness to Pay (WTP), that is, how much users are willing to pay to protect their personal information (Acquisti, John, & Loewenstein, 2013). The methodology, called VOPE (Value of Privacy Estimator), is an iterative process that recovers a quantified estimation for the value of privacy in an ecommerce transaction. We demonstrated our methodology in an empirical study ($n = 195$) and validated our results by analyzing the correlation with another independent experiment that ranks the values of privacy ($n = 118$) and by comparing our results to the results of Tsai et al. (Tsai et al., 2011). Current methodologies for estimating the value of privacy are rather biased (e.g., they address the users directly by asking about their preferences), or they do not provide intrinsic value (e.g., when ranking only hidden values of different scenarios). Our methodology, by contrast, bypasses those two obstacles, providing a reliable value of privacy that can be accommodated in economic models.

2. Related work

The value of privacy is context dependent, and a specific user may consent to divulge the same personal information for different minimal rewards in different contexts (Acquisti et al., 2015). In the case of electronic commerce, the value of privacy is correlated with users' willingness to pay a premium price for not disclosing some information against given benefits in a particular transaction and in a particular context (Chiu, Wang, Fang, & Huang, 2014). Tsai et al. (2011) showed that individuals would pay more for goods on the Internet if they perceived that their personal information would be kept by retailers with better privacy guarantees. In the case of mobile applications, smartphone users were willing to pay premiums for applications that were less likely to request access to personal information (Egelman, Felt, & Wagner, 2013). Huberman, Adar, and Fine (2005) showed that users adjust the exact value of privacy according to the social context. Svensson (2003) showed that the value of privacy is a composition of the probability for a privacy breach and the actual cost of the damage when the breach occurs.

Because there is a tradeoff between information sharing and benefits (Price, Adam, & Nuseibeh, 2005), the value of privacy may be defined as: *the value of the benefits at the equilibrium point, when an individual is indifferent to the information disclosure*. However, as simple as the definition of the value of privacy is, growing evidence indicates the challenges of estimating the value of privacy (Acquisti et al., 2006; Hann et al., 2002). Wathieu and Friedman (2007) showed that even when users are aware that they do not know how their information will be used, they tend to base their decisions on speculation. Models from behavioral economics have common ground with the psychology behind privacy decision-making: Usually, individuals have incomplete information when performing privacy decisions. Even if the information is complete, users are not always aware of the consequences of their decisions (Baek, Kim, & Bae, 2014; Cho, Lee, & Chung, 2010), and even if information and consequences are known, users' decisions can be biased (Acquisti & Grossklags, 2007).

In order to better design systems that introduce a tradeoff between privacy and utility, the value of privacy estimation method should comply with four criteria: the value must be explicitly numeric and in measurable units (e.g., U.S. Dollar); the value must be reliable in the sense that it should be robust enough to withstand basic manipulations; the value must be available in common transactions and not only at extreme ones; the value must reflects the preferences of individuals. Existing approaches that aim to estimate the value of privacy can be classified into 6 categories: a) Direct surveys, which ask individuals about the price they would be willing to pay for their personal information (Hann et al., 2002). However, surveys are considered to be easily biased by the wording of the questionnaires (Braunstein, Granka, & Staddon, 2011), and are known to be inaccurate when relying on reported behavior that is infrequently and irregularly (Staddon, Acquisti, & LeFevre, 2013); b) Indirect surveys, which ask people to indicate a general scale of willingness to share a piece of information but do not ask for the value of sharing (Braunstein et al., 2011); c) States of privacy transition, which ask individuals to rank privacy decisions according to their sensitivity (that reflects the ordinal value of privacy), but do not offer a model that quantifies the value and the importance of these preferences (Kosa, El-Khatib, & Marsh, 2000; Preibusch, 2013); d) Worst case scenario analysis, which measures the maximal financial loss as a result of information disclosure and thus may not indicate a useful utility value, particularly when the loss is a continuous function (Longpre & Kreinovich, 2006); e) Second-price auctions, in which participants are asked to state their value of privacy, but only those with the n lowest values will win (gain the lowest value not participating). This methodology introduces a relationship between participants, which usually does not exist in privacy decisions, like in ecommerce, and is subjected to manipulations (Danezis, Lewis, & Anderson, 2005; Staiano et al., 2014); and f) Effect of privacy breech on company valuation, which is relevant only to public corporations and can be initiated only when a breech occurs (Acquisti et al., 2006; Garg, Curtis, & Halper, 2003). There is no single methodology that address the four criteria, i.e., takes into account the bias effect, yields numeric and continuous values, and can be applied to common situations that are relevant to the individual's privacy decision.

3. Value of Privacy Estimation (VOPE)

There are several competing theories that describe how people make decisions about privacy. Privacy Calculus theories follow the rational model in privacy decision-making, assuming a rational agent that will choose the alternative with the highest expected value when a risk of privacy violation is introduced (Culnan & Bies, 2003; Dinev & Hart, 2006; Li et al., 2010; Von Neumann & Morgenstern, 2007). However, the fact that users face uncertainty prevents us from assuming rationality in users' decisions (unlike the expected utility theory), thus, those decisions deviate from the optimum. Our methodology relies on Prospect Theory (Kahneman & Tversky, 1979) to model privacy preferences when users face

uncertainty regarding the outcome of information disclosure. Prospect Theory aims to explain the relationship between the increase in the probability of privacy violation (information disclosure, in our case) and the increase in the perceived cost. Hogarth and Kunreuther (1985) found that ambiguity affects users' choices by diverting them from optimality towards alternatives with less uncertainty. The implications of information disclosure are vague to the average user, and ambiguity aversion in decision-making strategies increases when the user is less familiar with the odds (Fox & Tversky, 1995).

Liu and Colman (2009) showed that repetitive decisions decrease aversions from ambiguity. Therefore, a phase of learning that precedes the actual sampling of the value of privacy may increase the reliability of the results. We take this approach as the basis for our method, applying it to privacy decision-making. We avoid asking the user about the value of privacy. Instead, relying on a repetitive game, in which in each iteration the user is asked to accept or reject a transaction with a given discount that includes some component of information disclosure. After each iteration, the user is presented with feedback that states whether her private information was disclosed. The process allows users to identify their preferences and to analyze their prospects of privacy. If the decisions converge, it can be used to deduct the distribution of the value of privacy. In our method, the user faces a tradeoff between the loss side and the gain side. Research shows that people's behavior is not symmetric between the gain and loss sides (Cohena, Jaffray, & Said, 1987), thus, one cannot deduct the other, and both have to be combined in the same game.

This approach is implemented in our VOPE methodology through the use of an iterative game in which the user is not asked explicitly about the perceived cost. We measure the perceived cost of privacy, which is the loss side in the tradeoff between benefits and information disclosure. In the experiment that we conducted, the user is experiencing a fictitious ecommerce site. On the main screen of the game, as shown in Fig. 1, the users are alerted that their ecommerce transaction of a specific item (e.g., an asthma inhaler) was completed. They are offered a discount of $10 in return for their consent to allow ACME (the company that operates our fictitious ecommerce site) to use the information about the transaction by possibly showing information about the user to other customers. Because well-known products were used in the scenario, we assumed that the participant was familiar with its approximate price. The following message is shown to the user, reflecting a hypothetical transaction: "When someone shows interest in purchasing an asthma inhaler, they might be informed that a 32-year-old male who lives in the USA and earns about $40,000 a year also bought an asthma inhaler". The personal information in the message is taken from the demographic details provided by the participants. By the end of each iteration, if the user agreed to the offer (a 'yes' answer), she is alerted whether the information was actually disclosed. This indication is randomly assigned according to a pre-set probability that is hidden from the user.

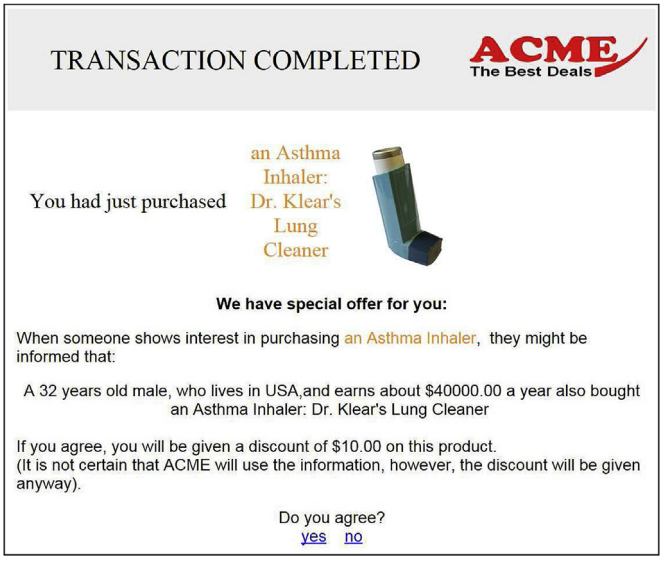

Fig. 1. The main screen of the experiment. The user is informed about the details of a transaction that was completed and is offered to receive a discount in return for the consent to let ACME (the fictive ecommerce site) use the participant's personal details in future hypothetical transactions. If the participant chooses a 'yes' answer, the user will be alerted as to whether the information was used, according to a pre-set probability that is hidden from the user.

3.1. Definition, protocol and algorithm

The protocol of VOPE aims to measure the value of privacy in a given transaction that includes a component of privacy disclosure. We assume that for a specific item and under a given scenario (e.g., asking user to share some information about an ecommerce transaction they have just carried), users are consistent, a property that is defined in the following way: Let $B(u,i,s)$ be the benefit that is offered to user u for disclosing some information on purchasing item i under scenario s. Given two benefits, $B_1(u_1,i_1,s_1)$ and $B_2(u_2,i_2,s_2)$, if $B_1 > B_2$, given that it is the same user ($u_1 = u_2$), the same item ($i_1 = i_2$), and the same scenario ($s_1 = s_2$), if the user is willing to disclose some information for this item under this scenario for the benefit B_2, we assume that the user will also agree to do so for the same item and under the same scenario for benefit B_1. On the other hand, if the user does not agree to disclose the information for some item under a given scenario for benefit B_1, she will refuse to do so for the same item and under the same scenario for benefit B_2. Relying on this assumption, we can deduced that: if the answer to the question about the willingness to disclose information in return to a given discount (as depicted in Fig. 1) is positive, the user's value of privacy is lower or equal to the discount; and if the answer is negative, the value of privacy is higher or equal to the discount. This rule is implemented as part of the algorithm, which uses a binary search with the refinement of the initial parameters, as will be described henceforth.

The game includes two phases, a learn mode and a run mode. In the learn mode, the user gets to know the environment and can learns the odds of disclosure and the system behavior; in the run mode, the user's decisions are recorded and analyzed to estimate the value of privacy. The learn mode phase of the process emulates the editing stage described in Prospect Theory's decision-making process, in which the user heuristically evaluates the outcomes of the decision.

The VOPE algorithm is depicted in Algorithm 1. VOPE has to be initialized with the following parameters: a) Id_j – the initial discount for item j, that is, the discount that will be offered to the user on the first iteration (relative to the item's price, e.g., $10); b) $I2\_Ld_j$ – the discount for item j for the second iteration if user's answer is positive (in this case, the price will be lower, e.g., $5); c) $I2\_Hd_j$ – the discount for item j for the second iteration if the user's answer was 'no' (in this case, the price will be higher, e.g., $15); d) L_j – the number of iterations in learn mode (e.g., 10); e) Pd – probability of disclosure (e.g., 0.6); and f) NG – number of iterations in run mode (e.g., 6).

The protocol includes a finite number of iterations. By the end of the game, two situations are possible: a) Both lower and upper value boundaries were found (if the user accepted at least one value and declined at least one value); and b) only one of the

boundaries was found (if the user accepted or declined for all values). Thus, VOPE can yield three types of results: If lower and upper boundaries were found, an explicit value of privacy is calculated (which is the average of the lowest upper boundary and the highest lower boundary); Otherwise, VOPE indicates that the value of privacy is higher (if only lower boundary found) or lower (if only upper boundary found) than the bound. The gap between the lower and upper boundaries, if both are found, can be controlled by the number of iterations. To set up initial parameters and to manage the whole process, we offer a protocol in which initial values are estimated. Then, the game is launched for an initiation run for a small batch with those parameters. The parameters are fine-tuned according to results of the initiation run to center values and minimize gaps. Then, the game is run again for a small batch to test whether further tuning is required. The sample size can also be calculated to allow for statistical significance. VOPE is then launched for the main run with a full sample size, and all data, including the tuning phase, can be used to calculate the final results. An optional validation phase can be added to verify reliability. The complete protocol flow chart is depicted in Appendix A.

```
min_discount ← Infinity; max_discount ← -1
Step_Factor ← 2

for game_number=0 to NG
    for j=1 to Num_of_items            % calculate discounts for the current game
        case game_number
            % learn mode  (and also apply to 1st iteration)
            0:   discount[j] ← Id[j]  % learn mode
            % second iteration
            2:   if User's_preference[j]  == "agree" then % second iteration
                     max_discount[j] ← discount[j]
                     discount[j] ← I2_Ld[j]
                 else
                     min_discount[j] ← discount[j]
                     discount[j] ← I2_Hd [j]
            % 3rd iteration and above
            >2:  if User's_preference[i]  == "agree" then
                     max_discount[j] ← discount[j]
                     if  min_discount[j] < discount[j]
                         discount[j] ← (discount[j] +
                                  min_discount[j])/Step_Factor
                     else
                         discount[j] ← discount[j] / Step_Factor
                 else
                     min_discount[j] ← discount[j]
                     if  max_discount[j] > discount[j]
                         discount[j] ← (discount[j] +
                                       max_discount[j]) / Step_Factor
                     else
                         discount[j] ← discount[j] * Step_Factor
    if game_number == 0                % iterations of learn mode (only for 1'st game)
        for iteration =1 to L
            j ← random(UNIFORM_DISCRETE,1,No_of_items)
            display (GAME_SCENARIO, j, discount[j])
            get User's_preference[j]
            if(random(UNIFORM_CONTINUOUS,0,1) < Pd) &
                User's_preference[j]=="agree") then
                display "Disclosed'
            else
                display 'Not Disclosed'
    else
        for j =1 to No_of_Items        % test mode
            display (GAME_SCENARIO, item, discount[item])
            get User's_preference[item]
            if (random(UNIFORM_CONTINUOUS,0, 1) < Disclosure_ProbabilitPdy) &
              (User's preference[j]=="agree") then
                display 'Disclosed'
            else
                display 'Not Disclosed'

record results
```

Algorithm 1. VOPE algorithm pseudo-code

In the learn mode phase, the algorithm iterates through the scenario L times, displaying the offer to the users to receive a discount of Id_j in return for their consent and (if the user's answer was 'yes') presents the hypothetical outcome of the transaction according to the random uniform distribution based on Pd.

In the run mode, the scenario starts with the discount Id_j. If the user agrees to disclose, VOPE concludes that Id_j is the upper boundary and displays $I2\_Ld_j$ on the second iteration of this item. If the participant declines to share, VOPE concludes that Id_j is the

Table 1
List of items included in the experiment, and the amounts of presented discounts.

#	Item name	Item market price ($) (not displayed to the user)	Discount amount initial parameters ($)		
			Initial discount	2nd Iteration lower discount	2nd Iteration higher discount
			(*Id*)	(*2I_Ld[i]*)	(*2I_Hd[i]*)
1	A notebook: MacBook Air 13.3	1954	10	5	15
2	A book: First Person by Vladimir Putin	11	5	4	6
3	An adult toy: Lelo Lily Massager	84	30	25	35
4	Rechargeable batteries: Energizer AA size	16	3	2	4
5	An asthma inhaler: Dr. Klear's Lung Cleaner	28	10	5	15
6	A smartphone: SAMSUNG GALAXY SIII	347	13	10	16

lower boundary and displays $I2\_Hd_j$ on the second iteration of this item. From now on, if VOPE finds lower and upper boundaries, the discount in the next iteration will be the average of the lowest upper boundary and the highest lower boundary. Otherwise, $Idj+1 \leftarrow Idj/2$ if the user accepts and $Idj+1 \leftarrow Idj^*2$ if the user declines. VOPE repeats the iterations *NG* times for each item.

4. Empirical evaluation

4.1. Experiment design

To evaluate our approach, we conducted a user study with n = 195 participants. As depicted in Fig. 1, the participants were asked to accept or reject a set of hypothetical proposals to receive a discount in a purchase transaction, in return for the user's consent to let the ecommerce site expose some data about the transaction. All items were initiated by discounts according to a subjective selection and were displayed with their pictures. An initial batch of n = 26 participants was launched to fine-tune the initial parameters Then, two more batches were launched, one with n = 78 participants and a probability of disclosures of 0.3 and another one with n = 91 participants and a probability of disclosures of 0.6. The item list and the discounts (after fine-tuning), which were the initial parameters for the experiment, are depicted in Table 1. We recorded users' responses to the proposals, which were the dependent variable of the experiment. A successful result is when both the upper and lower boundaries were found so that the value of privacy can be deduced to be inside this range.

The game included 10 steps in the learn mode (in which the assessed items were selected at random) and 5 rounds of the full set of all 6 items in the run mode. The game started with an ethics consent form that contained a short explanation of the process and an estimation of the time required to complete the task. This was followed by demographic questions (about gender, age and income). The study was approved by a university internal ethics board. The software for the study was implemented using HTML,

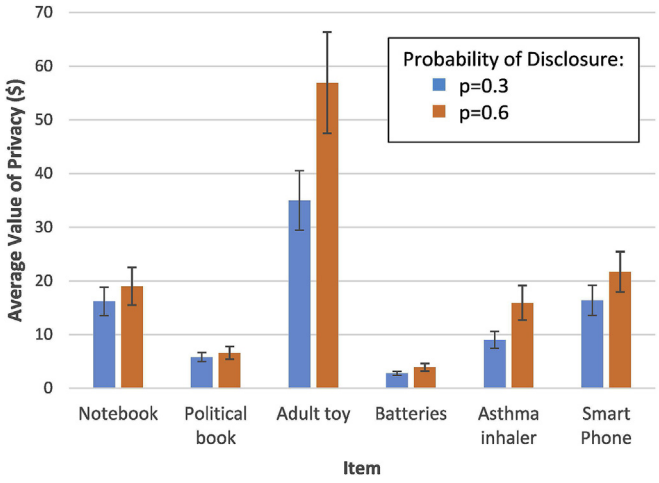

Fig. 2. Values of privacy for all items with different probabilities of disclosure. The X-axis describes the items, whereas the Y-axis describes the value of privacy in US dollars ($). The blue bars stand for the value of privacy when the probability of disclosure p = 0.3, and the orange bars p = 0.6. The whiskers indicate the standard error of the mean (SEM). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2
The value of privacy for all items with probability of disclosure = 0.3 and 0.6. For each item and probability, the minimal value, maximal value, average value and the standard deviation are indicated. Also for each item, the rate of users that had only lower boundary, and of those who had only upper boundary are indicated.

Item name:		Notebook	Political book	Adult toy	Batteries	Asthma inhaler	Smart-phone
VOPE settings: Probability of disclosure: 0.30, Learning steps: 10, Games: 5							
Number of users		78	78	78	78	78	78
Value of privacy	Min	0.9	0.8	4.7	0.4	0.9	3.1
	Max	90.0	26.5	210.0	17.5	65.0	96.0
	Average	16.2	5.8	35.0	2.8	9.0	16.4
	Std.	20.6	5.9	37.2	2.8	11.2	19.6
VOPE settings: Probability of disclosure: 0.60, Learning steps: 10, Games: 5							
Number of users		91	91	91	91	91	91
Value of privacy	Min	0.9	0.8	4.7	0.4	0.9	1.9
	Max	90.0	36.0	210.0	24.0	90.0	96.0
	Average	19.0	6.6	56.9	3.9	15.9	21.7
	Std.	24.3	8.0	59.5	4.7	21.1	24.9
Average for both VOPE settings of the rate of participants with no two boundaries							
Lower boundaries only		11%	16%	22%	14%	18%	15%
Higher boundaries only		10%	19%	20%	15%	18%	22%
Confidence level for the hypothesis that the value of privacy when probability=0.6 is greater than the value of privacy when probability=0.6							
Confidence level		0.97	0.73	0.99	0.99	0.97	0.97

JavaScript, Python and the Flask micro framework and hosted on Amazon Web Services.

The participants were recruited using Amazon Mechanical Turk (or MTurk for short). Mturk is a commonly used tool in information systems, human-computer interaction, privacy studies and behavioral economics studies (Kelley, 2010; Mason & Suri, 2012). It was also validated as a tool for conducting economic behavior experiments (Horton, Rand, & Zeckhauser, 2011). In the domain economic behavior, the reliability and validity of MTurk studies were empirically evaluated, resulting in the conclusion that: "MTurk especially is suitable to conduct survey research if Internet users are the intended population" (Schaarschmidt, Ivens, Homscheid, & Bilo, 2015). American MTurk workers, who were the population of our study, have similar amount of personal information online as the general American population, and have higher levels of awareness of privacy threats than the general population (Kang, Brown, Dabbish, & Kiesler, 2014). Thus, if VOPE is implemented in real ecommerce transaction for example, we can assume that a common user can handle the task as well as the experiment's user. Since VOPE is parametric (as described in section 3.1), the initial values and the steps can be tuned to address different populations with different distributions of the value of privacy.

The reward per assignment was 1.25–1.5 USA dollars, reflecting an hourly wage of approximately 5.5 USA dollars, a standard hourly compensation in Mechanical Turk studies (Ross et al., 2010). Participants were required to be over 18 years old, have an Amazon MTurk HIT rate of 90% or higher, and be from the US to ensure that they understood the survey on a native-tongue level. The study was authorized by the institutional ethics committee (IRB). We followed standard experimental practices in MTurk, making sure all participants had the required hit approval rate (Mechanical Turk inherent trust score) of greater than or equal to 90%. We collected responses from a total of over 212 participants; 195 of them were approved (8 were incomplete and 9 had repetitive answers that hinted of arbitrary responses). Approximately 44% of the participants were males, 55% were female, and less than 1% were 'other' or 'preferred not to report'. Thirty percent of the participants were in the age range of 18–30 years old, 23% were in the range 31–40 years old, 19% were in the range 41–50 years old, 21% were in the range 51–60 years old, and 7% were 60 years old or older. Thirty–five percent of the participants had an annual income of 0–20,000 US dollars per year; 55%, 20,001 to 70,000; 9%, 70,001 to 100,000; and the rest, above 100,000. No significant correlation was found between the experimental variables and the demographic data items.

4.2. Results

For each participant, the experiment produced a list of replies to the offers. We processed the data with a MATLAB script that identified for each user the lowest upper boundary and the highest lower boundary. If both lower and upper boundaries existed, the value of privacy was set to be in that range, and we conveniently set it to their average. Table 2 describes the results of the final price calculated, and Fig. 2 depicts the distribution of the values of privacy for all items when the probability of disclosure is 0.3.

It can be seen that when the probability of disclosure is 0.6, the averages of the value of privacy are higher than with probability of 0.3 for every item, with a confidence level shown at the bottom of Table 2 $\left(\overline{X}-\overline{Y}\sim N\left(\mu_x-\mu_y,\frac{\sigma_x^2}{n_x}+\frac{\sigma_y^2}{n_y}\right)\right)$. This difference in the value of privacy as a function of the probability of disclosure is depicted in Fig. 2. The expectancy of the value of privacy may expected to be doubled if probability is changed from $p = 0.3$ to $p = 0.6$, however, it can be seen that in most items, the value of privacy is changing in a more moderate way. This phenomenon lines up with Prospect Theory because people are not calculating their exact utilities but are giving weights to the probabilities (Kahneman & Tversky, 1979).

When both boundaries found the price converge, which allows us to deduce the value of privacy. We measure the rate of price convergence as the number of iterations required to archive both lower and upper boundaries. Fig. 3 depicts the average convergence for the experiment with a $p = 0.3$ probability of disclosure. The X-axis describes the number of the iterations, and the Y-axis the portion of users who had not reached convergence at this iteration. It can be seen that the average values converge across all products

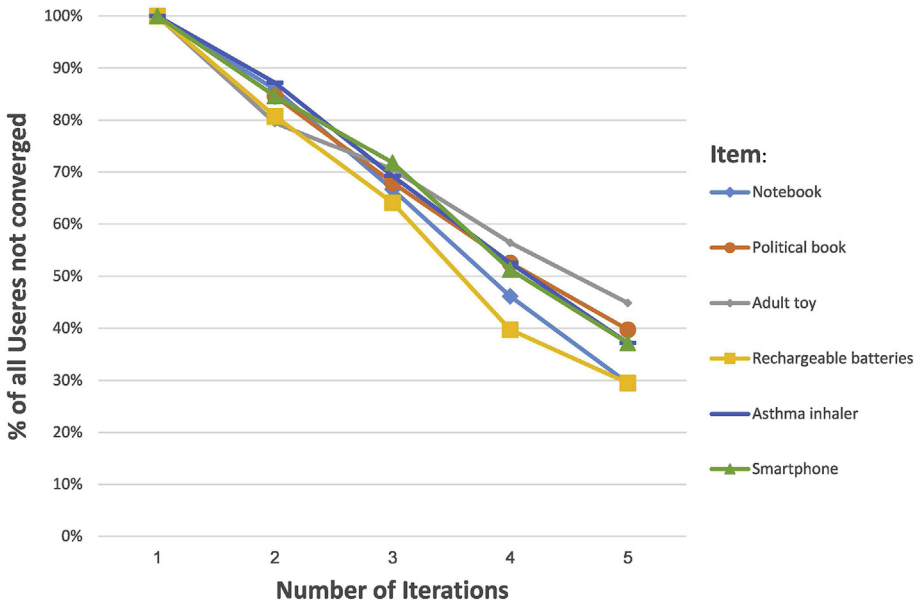

Fig. 3. The rate of convergence for the experiment with a $p = 0.3$ probability of disclosure. The X-axis describes the iteration number, and the Y-axis describes the portion of users who had not reached two boundaries on this iteration. The right point of each graph indicates the portion of users for whom only lower or upper boundaries found for this item respectively (no convergence achieved, and the value of privacy cannot be deduce).

at a rate of approximately 16% per iteration, i.e., with each iteration approximately 16% of the users reaching convergence, having upper and lower boundaries. The "adult toy" has the slowest convergence rate, whereas "rechargeable batteries" have the fastest. We found a negative correlation of $\rho = -0.59$ between the average value of privacy and the average rate of convergence. As the value of a product's privacy increases, its convergence rate decreases (e.g., the notebook, with value of privacy of $16.20, has a convergence rate of 2.5 steps, whereas the adult toy, with value of privacy of $35.00, has a convergence rate of 1.8 steps).

4.3. Validation

To validate the results, we conducted two types of tests. The first test involves comparing our empirical results with another independent study by Tsai (Tsai et al., 2011). We compared the rank of the items' values of privacy as received in our experiment, with the rank of 4 products by Tsai. The products in our experiment were rechargeable batteries, a notebook, a political book, and an adult toy; those in Tsai's experiment were office supplies, a laptop, a book, and sex toy. We see that the ranking matches our results with a correlation coefficient of $\rho > 0.8$. Tsai's experiment does not provide an intrinsic value for privacy; thus, it cannot derive a calculation of users' utility, but it can support the reliability of our work.

The second test involves conducting a different independent online study and crossing the results to perform external validity. The process also started with an ethics consent form that contained a short explanation of the process and an estimation of the time required to complete the task. We sampled a different population in which the participants are asked to rank the same items in the empirical study (as shown in Table 1) according to their willingness to disclose some non-personal data about the purchase in return for a discount. The results of the validation study are directly transformed into a two-dimensional distribution of each item versus its rank. The same distribution can be indirectly deduced from the first game, and those two distribution can be tested for the goodness of fit.

Let V be a discrete distribution matrix that can be derived directly from the validation study, v_{ij} is the rank reported in the validation study experiment for item j by user i ($j > 0$, $J \in \mathbb{N}, j < j' \Rightarrow v_{ij} < v_{ij}$), i and j are unique identifiers of the user and item respectively, and have no quantification meaning. For example, if an Asthma inhaler is item number 3, and user 7 ranked it as the sixth in its value of privacy, than $v_{7,3} = 6$.

Let VP be the probability matrix, vp_{mn} is the probability of item m to get ranked n out of the m items ($0 < n < m$, $n \in \mathbb{N}$). In the above example, if we had 10 users, and an Asthma inhaler (item number 3) was ranked as the sixth in its value of privacy by four users, than $vp_{3,6} = \frac{4}{10} = 0.4$. The formal definition of VP is:

$$vp_{mn} = \frac{\left|\{v_{ij} \in V |\ j = m\ ,\ \ v_{ij} = n\}\right|}{\left|\{v_{ij} \in V |\ j = m\}\right|}$$

The same distribution as vp_{mn} can be deduced indirectly from the experimental study in the following manner: Let G be the result matrix of the experiment phase, g_{ij} is the value of privacy of item j for user i. g_{ij} can have a value if both boundaries were found, a *lower* definition if only lower boundaries were found, and a *higher* definition if only higher boundaries were found. To assure that elements without two boundaries would be placed on the margins when ranking, let us convert *lower* and *higher* indications as follows:

$$g'_{ij} = \begin{cases} gij & \text{if } g_{ij} = value \\ -1 & \text{if } g_{ij} = Lower\ boundary \\ \infty & \text{if } g_{ij} = Higher\ boundary \end{cases}$$

In the above example, if the value of privacy for an Asthma inhaler was found for user number 5 and is $8.5, than: $g'_{5,3} = 8.5$, if only lower boundary found, $g'_{5,3} = -1$, and if only higher boundary was found $g'_{5,3} = \infty$.

Now, Let us rank the items according to their prices among each user i independently:

$$gr_{ij} = RANK\left(\left\{g'_{i1}, g'_{i2}, \ldots g'_{iJ}\right\},\ g'_{ij}\right)$$

so that: $j \geq 1$, $J \in \mathbb{N}$, $gr_{ij} = gr_{ij'} \Rightarrow g'_{ij} = g'_{ij'}$, $gr_{ij} > gr_{ij'} \Rightarrow g'_{ij} > g_{ij'}$. For example, if we had 4 items (indexed 1,2,3,4), and user number 7 had values of privacy of $5, $3, $15, and $14 for the four items respectively, then: $gr_{7,1} = 2$, $gr_{7,2} = 1$, $gr_{7,3} = 4$, $gr_{7,4} = 3$.

We can now construct a two-dimensional probability matrix GP for the game, which has the same meaning as vp_{mn} in the validation. gp_{mn} is the probability of item m to get ranked n out of the m items in the game, and can be calculated from the ranked matrix:

$$gp_{mn} = \frac{\left|\{gr_{ij} \in GR |\ j = m\ , gr_{ij} = n\}\right|}{\left|\{gr_{ij} \in GR |\ j = m\}\right|}$$

Both matrixes, gp_{mn} and vp_{mn} can be compared by applying a two-dimensional X^2 test (Bishop, Fienberg, & Holland, 2007) to complete the validation process.

G_{ixj}	value of privacy of item j for user i	GP_{mxn}	gp_{mn} is the probability of item m to get ranked at the n-th place at the game
G'_{ixj}	Gij with lower & upper boundaries replaced with -1 & ∞, respectively	V_{ixj}	v_{ij} is the rank of user i for item j
GR_{ixj}	The rank of each item issued by the current user	VP_{mxn}	vp_{mn} is the probability of item m to be ranked at the n-th place in the validation

For example, let us assume we have 2 items, medicine and glass, and we obtained the probability matrixes in the game and the validation phase as follows:

$$vp = \begin{array}{c} \text{rank\item} \\ 1 \\ 2 \end{array} \begin{array}{c} \text{Medicine} \\ \begin{pmatrix} 0.2 \\ 0.8 \end{pmatrix} \end{array} \begin{array}{c} \text{Glass} \\ \begin{pmatrix} 0.7 \\ 0.3 \end{pmatrix} \end{array}$$

$$gp = \begin{array}{c} \text{rank\item} \\ 1 \\ 2 \end{array} \begin{array}{c} \text{Medicine} \\ \begin{pmatrix} 0.3 \\ 0.87 \end{pmatrix} \end{array} \begin{array}{c} \text{Glass} \\ \begin{pmatrix} 0.6 \\ 0.4 \end{pmatrix} \end{array}$$

Medicine, in this example, has a probability of 0.2 to be ranked first according to the value of privacy in the game phase and 0.3 in the validation phase. The two-dimensional chi-square test yields $X^2\ (1,\ N = 2X2) = 0.089,\ p = 0.76$; thus, we can accept the null hypothesis that the distributions are identical.

For this test, we conducted another independent study in which we asked participants to rank the sensitivity of the items (the main interface of the ranking process is shown in Appendix B). The validation study has also been conducted using Amazon Mechanical Turk (n = 118), with a protocol, ethics approval and demographics similar to the main study. The validation test can be applied to each item separately ($X^2_m = \sum_n \frac{(vp_{mn} - gp_{mn})^2}{vp_{mn}}$ with $(6-1) =$ 5 degree of freedom). In this test, we obtained *P-value > 0.85* for "notebook computer", "political book" and "smartphone", *P-value > 0.60* for "asthma inhaler", *P-value > 0.30* for "adult toy", and a low *P-value* ($P < 0.05$) only for "rechargeable batteries" (which is the only item for which the null hypothesis that both distributions are identical cannot be accepted). We applied a χ^2 2-dimensional test, with a $(6-1)^2 = 25$ degree of freedom ($\chi^2 = \sum_m \sum_n \frac{(vp_{mn} - gp_{mn})^2}{vp_{mn}}$) and received $X^2\ (25,\ N = 6X6) = 22.43,\ p = 0.61$, which strongly support the null hypothesis that the distributions are identical.

5. Discussion

This research suggests a method for estimating people's privacy preferences in financial terms. Posner (1981) argues for symmetry between "selling oneself and selling a product". However, unlike evaluating a product, estimating the value of a person's information privacy preferences remains an open problem. Our study aims to fill this gap in a particular context, suggesting a way to estimate the value of privacy in scenarios in which information is disclosed as part of a transaction that involves financial benefits. By suggesting such a method, we can analyze people's decision-making processes when privacy is managed in a transaction, reflecting many real-world situations. Valuation can be used to measure the weight people give to different facets of privacy management and the interaction in which information is disclosed. For example, experimenters can quantify the effect of different types of user interaction and privacy notices on the weight users give to privacy.

The ability to valuate privacy has wide applications in many areas in which privacy is managed in an economic setting. Beyond our example of the ecommerce tradeoff between discounts and privacy, the methodology can be applied to behavioral advertising, mobile data collection and so on. The quantification of the price a user "pays" by disclosing information can help create privacy policies that represent the weighted preferences of the community. In the security field, it is well-known that the security level is correlated with the investment in the security system. Organizations adopt an economic approach based on a cost-benefit analysis (Stoneburner, Goguen, & Feringa, 2002), which trades off the cost of the information security mechanism vs. the cost of unwilling information disclosure. In many cases, particularly when the information concerns individuals, the cost of information disclosure cannot be evaluated if the value of privacy is unknown. Another application of this methodology is optimizing the choice configuration of privacy preferences. A configuration set can be design to maximize users' utilities or to maximize equity between users' utilities. This methodology may be extended to estimate the gain side, for example, in disclosing medical information when both "profit" and "loss" are illusive (Rindfleisch, 1997).

Our results highlight the place of uncertainty which has significant effect on people's privacy decision-making (Otim & Grover, 2012). The fact that people's valuation converges with the number of iterations indicates that the behavior of most users is consistent with the amount of information they have. This result calls into question some elements of theories, such as Privacy Calculus, that include uncertainty as part of their models but do not adopt an iterative model of privacy decision-making (Dinev & Hart, 2006). We argue for models that rely on the convergence of privacy preferences based on the availability of feedback. When framing the results of this study, it is important to stress that our results cannot be extended to reducing privacy to a straightforward financial value. Privacy is a normative concept, and in most legal frameworks, it is considered a basic human right (Smith, Dinev, & Xu, 2011). We believe that in many cases, privacy cannot be bought or sold. However, valuating privacy is possible in several important and realistic scenarios, and specifically in scenarios that include a component of active information disclosure against some financial reward. In these scenarios, privacy is already a part of the transaction, though sometimes invisible to the individuals, eroding the trust users have for information systems (Spiekermann et al., 2015). Our work may help in exposing the hidden financial value of privacy, making these transactions more transparent, and improving the control that users have over their information.

To understand how to apply our method, it is important to understand its context and limitations. First, it measures the value of the privacy of a given scenario (Bergström, 2015). The accuracy of the results obtained by VOPE is a function of the number of iterations and the initial parameters. It is difficult to forecast how this value will change when parameters like the probability of disclosure change, not to mention a change in the whole scenario. To find a mathematical correlation between the value of privacy and the experimental parameters (e.g., the probability of disclosure or the price of the product), further research is required. Second, VOPE methodology measures the value of privacy under a specific scenario and for each item separately. Since ecommerce covers a wide variety of items, this characteristic of the methodology may introduce a difficulty in the implementation. This problem may be solved by categorizing the items (e.g., placing paper clip and pen refill in the same category, expecting they both have similar value of privacy), and categorizing the scenarios (e.g., placing two well-known ecommerce sites under the same category, expecting both have same perception from the user point of view). Naturally, this categorization requires further research. Another limitation is that VOPE assumes price is one-dimensional index. However, the model may be upgraded to accommodate a more complicated pricing method, such a stochastic distribution of the price of a specific item for a specific user under a given scenario.

6. Conclusions

In this paper, we describe a method to evaluate the value of privacy. We introduced an algorithmic evaluation methodology called VOPE (Value of Privacy Estimation), and we offered a methodology to validate the reliability of the results. We conducted an empirical experiment in which an ecommerce scenario was simulated. The experiment included 6 items and was conducted with n = 26 participants in the initialization phase; n = 78 participants with a disclosure probability of 0.3, and with n = 91

participants with a disclosure probability of 0.6 in the main phase. The experiment included 5 iterations, after which an average of 74% of the results converged (i.e., the value of privacy was found). We validated our results by crossing them with a different independent experiment, and we received a *P-value* = *0.61*, which strongly supports the reliability of our results. We also compared our results to those conducted by another researcher and again received convincing correlations.

The value of privacy can be used as a parameter in evaluating and designing systems that introduce a tradeoff between the benefit a user may gain and the cost of information disclosure. The advantage of our method is the flexibility to design the utility function according to a chosen policy and to express this policy in economic terms.

Acknowledgements

This work was supported by the Israel Bureau Cyber Grant, no. 3-9758. We thank Michael Birnhack, Tal Zarsky and Niva Elkin-Koren for their comments on early drafts.

Appendix A. The protocol of VOPE

The following is a flow chart of VOPE methodology process that evaluates the value of privacy. The process is divided into three stages: stage 1- tuning of parameters; stage 2- evaluating the value of privacy; stage 3- validations.

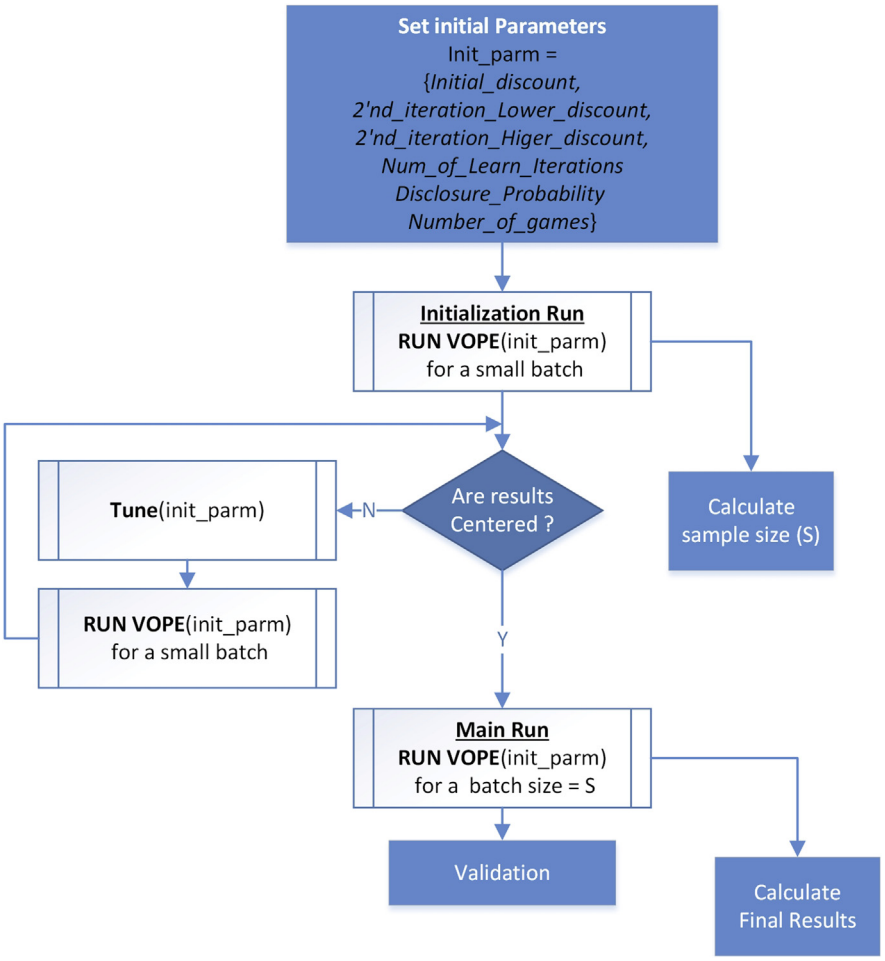

Appendix B. The main interface of the validation of VOPE

The following image depicts the main interface of the validation phase. The participants were asked to rank the items according to the willingness to disclose some non-personal data about the purchase of each item. The item list appears on the left side, and the user was asked to drag and drop the items to the right side according to her rank. In the displayed example, the user has already ranked two items (the Rechargeable Batteries and the Book) and has still four items to rank (Adult Toy, Notebook, Asthma Inhaler, and a Smartphone).

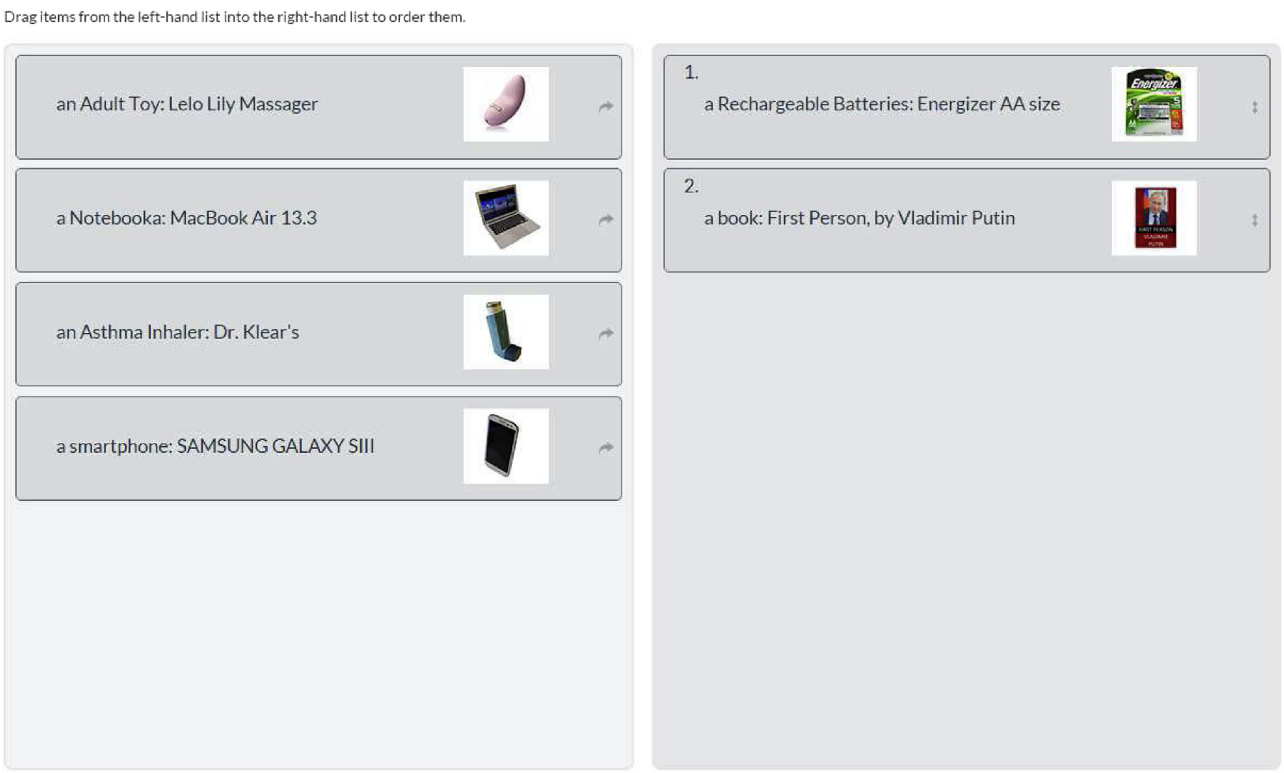

References

Ackerman, M. S., Cranor, L. F., & Reagle, J. (1999). *Privacy in e-commerce: Examining user scenarios and privacy preferences*. NY: ACM.

Acquisti, A. (2004). *Privacy in electronic commerce and the economics of immediate gratification*. s.l. (pp. 21–29) ACM

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science, 347*(6221), 509–514.

Acquisti, A., Friedman, A., & Telang, R. (2006). *Is there a cost to privacy breaches? An event study*. s.l. ICIS

Acquisti, A., & Grossklags, J. (2007). What can behavioral economics teach us about privacy. In *Digital privacy: Theory, technologies, and practices* (pp. 363–377). Auerbach Publications. s.l.

Acquisti, A., John, L., & Loewenstein, G. (2013). What is privacy worth? *The Journal of Legal Studies, 42*(2), 249–274.

Baek, Y. M. (2014). Solving the privacy paradox: a counter-argument experimental approach. *Computers in Human Behavior, 38*, 33–42.

Baek, Y. M., Kim, E.-m., & Bae, Y. (2014). My privacy is okay, but theirs is endangered: why comparative optimism matters in online privacy concerns. *Computers in Human Behavior, 31*, 48–56.

Bergström, A. (2015). Online privacy concerns: a broad approach to understanding the concerns of different groups for different uses. *Computers in Human Behavior, 53*, 419–426.

Bishop, Y. M., Fienberg, S. E., & Holland, P. W. (2007). *Discrete multivariate analysis: Theory and practice*. s.l. Springer Science & Business Media

Braunstein, A., Granka, L., & Staddon, J. (2011). *Indirect content privacy surveys: Measuring privacy without asking about it*. s.l. (p. 15) ACM

Chiu, C., Wang, E. T. G., Fang, Y., & Huang, H. (2014). Understanding customers' repeat purchase intentions in B2C e-commerce: the roles of utilitarian value, hedonic value and perceived risk. *Information Systems Journal, 24*(1), 85–114.

Cho, H., Lee, J.-S., & Chung, S. (2010). Optimistic bias about online privacy risks: testing the moderating effects of perceived controllability and prior experience. *Computers in Human Behavior, 26*(5), 987–995.

Cohena, M., Jaffray, J.-Y., & Said, T. (1987). Experimental comparison of individual behavior under risk and under uncertainty for gains and for losses. *Organizational Behavior and Human Decision Processes, 39*(1), 1–22.

Culnan, M. J., & Bies, R. J. (2003). Consumer privacy: balancing economic and justice considerations. *Journal of Social Issues, 59*(2), 323–342.

Danezis, G., Lewis, S., & Anderson, R. J. (2005). *How much is location privacy worth?*. s.l., s.n.

Dinev, T., & Hart, P. (2006). An extended privacy calculus model for E-Commerce transactions. *Information Systems Research, 17*(1), 61–80.

Dwyer, C., Hiltz, S. R., & Passerini, K. (2007). *Trust and privacy concern within social networking sites: A comparison of Facebook and MySpace*. s.l., s.n. (p. 339)

Egelman, S., Felt, A. P., & Wagner, D. (2013). Choice architecture and smartphone privacy: there's a price for that. *The Economics of Information Security and Privacy*, 211–236.

Featherman, M. S., Valacich, J. S., & Wells, J. D. (2006). Is that authentic or artificial? Understanding consumer perceptions of risk in e-service encounters. *Information Systems Journal, 16*(2), 107–134.

Fox, C. R., & Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *The Quarterly Journal of Economics, 110*(3), 585–603.

Garg, A., Curtis, J., & Halper, H. (2003). Quantifying the financial impact of IT security breaches. *Inforamtion Management & Computer Security, 11*(2), 74–83.

Hann, I.-H., Hui, K.-L., Lee, T. S., & Png, I. P. L. (2002). *Online information privacy: Measuring the cost-benefit trade-off*. s.l., s.n.

Hogarth, R. M., & Kunreuther, H. (1985). Ambiguity and insurance decisions. *The American Economic Review, 75*, 386–390.

Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: conducting experiments in a real labor market. *Experimental Economics, 14*(3), 399–425.

Huberman, B. A., Adar, E., & Fine, L. R. (2005). Valuating privacy. *Security & Privacy, 3*(5), 22–25.

Jensen, C., Potts, C., & Jensen, C. (2005). Privacy practices of Internet users: self-reports versus observed behavior. *International Journal of Human-Computer Studies, 63*(1), 203–227.

Kahneman, D., & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica, 47*(2), 263–292.

Kang, R., Brown, S., Dabbish, L., & Kiesler, S. (2014). *Privacy attitudes of mechanical turk workers and the U.S. public*. Menlo Park, CA, s.n.

Kehr, F., Kowatsch, T., Wentzel, D., & Fleisch, E. (2015). Blissfully ignorant: the effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal, 25*(6), 607–635.

Kelley, P. G. (2010). *Conducting usable privacy & security studies with Amazon's mechanical turk*. Redmond, WA: s.n.

Kosa, T. A., El-Khatib, K., & Marsh, S. (2000). Measuring privacy. *Journal of Internet Services and Information Security (JISIS), 1*(4), 60–73.

Li, H., Sarathy, R., & Xu, H. (2010). Understanding situational online information disclosure as a privacy calculus. *Journal of Computer Information Systems, 51*(51(1)), 62.

Liu, H.-H., & Colman, A. M. (2009). Ambiguity aversion in the long run: repeated decisions under risk. *Journal of Economic Psychology, 30*(3), 277–284.

Longpre, L., & Kreinovich, V. (2006). *How to measure loss of privacy*. s.l. University of Texas at El Paso

Mason, W., & Suri, S. (2012). Conducting behavioral research on Amazon's mechanical turk. *Behavior Research Methods, 44*(1), 1–23.

Min, J., & Kim, B. (2015). How are people enticed to disclose personal information despite privacy concerns in social network sites? The calculus between benefit and cost. *Journal of the Association for Information Science and Technology, 66*(4),

839–857.
Otim, S., & Grover, V. (2012). Resolving uncertainty and creating value from the exercise of e-commerce investment options. *Information Systems Journal, 22*(4), 261–287.
Posner, R. A. (1981). The economics of privacy. *The American Economic Review, 71*(2), 405–409.
Preibusch, S. (2013). Guide to measuring privacy concern: review of survey and observational instruments. *International Journal of Human-Computer Studies, 71*(12), 1133–1143.
Price, B. A., Adam, K., & Nuseibeh, B. (2005). Keeping ubiquitous computing to yourself: a practical model for user control of privacy. *International Journal of Human-Computer Studies, 31*(1), 228–253.
Rindfleisch, T. C. (1997). Privacy, information technology, and health care. *Communications of the ACM, 40*(8), 92–100.
Ross, J., et al. (2010). *Who are the crowdworkers? Shifting demographics in mechanical Turk* (pp. 2863–2872). Atlanta, Georgia: ACM.
Schaarschmidt, M., Ivens, S., Homscheid, D., & Bilo, P. (2015). *Crowdsourcing for survey Research: Where Amazon mechanical turks deviates from conventional survey methods*. s.l. University of Koblenz-Landau
Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: an interdisciplinary review. *MIS Quarterly, 35*(4), 989–1016.
Spiekermann, S., Acquisti, A., Böhme, R., & Hui, K.-L. (2015). The challenges of personal data markets and privacy. *Electronic Markets, 25*(2), 161–167.
Staddon, J., Acquisti, A., & LeFevre, K. (2013). *Self-reported social network behavior: Accuracy predictors and implications for the privacy paradox*. s.l. (pp. 295–302) IEEEE
Staiano, J., Oliver, N., Lepri, B., de Oliveira, R., Caraviello, M., & Sebe, N. (2014). *Money walks: A human-centric study on the economics of personal mobile data*. s.l. (pp. 583–594) ACM
Stoneburner, G., Goguen, A., & Feringa, A. (2002). *Risk management guide for Information Technology systems*. s.l. (30th ed.). National Institute of Standards and Technology (U.S.)
Svensson, A. (2003). *Analysing information systems security*. s.l. School of Economics and Management, Lund University, Department of Informatics
Tsai, J., Egelman, S., Cranor, L., & Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: an experimental study. *Information Systems Research, 22*(2), 254–268.
Von Neumann, J., & Morgenstern, O. (2007). *Theory of games and economic behavior*. s.l. Princeton university press
Wathieu, L., & Friedman, A. (2007). An empirical approach to understanding privacy valuation. In *HBS marketing research paper* (pp. 7–75).