

Empirical Measurement of Perceived Privacy Risk

JASPREET BHATIA, Carnegie Mellon University
TRAVIS D. BREAU, Carnegie Mellon University

Personal data is increasingly collected and used by companies to tailor services to users, and to make financial, employment and health-related decisions about individuals. When personal data is inappropriately collected or misused, however, individuals may experience violations of their privacy. Historically, government regulators have relied on the concept of risk in energy, aviation and medicine, among other domains, to determine the extent to which products and services may harm the public. To address privacy concerns in government-controlled information technology, government agencies are advocating to adapt similar risk management frameworks to privacy. Despite the recent shift toward a risk-managed approach for privacy, to our knowledge, there are no empirical methods to determine which personal data are most at-risk and which contextual factors increase or decrease that risk. To this end, we introduce an empirical framework in this paper that consists of factorial vignette surveys which can be used to measure the effect of different factors and their levels on privacy risk. We report a series of experiments to measure perceived privacy risk using the proposed framework, which are based on expressed preferences and which we define as an individual's willingness to share their personal data with others given the likelihood of a potential privacy harm. These experiments control for one or more of the six factors affecting an individual's willingness to share their information: data type, computer type, data purpose, privacy harm, harm likelihood, and individual demographic factors, such as age range, gender, education level, ethnicity and household income. To measure likelihood, we introduce and evaluate a new likelihood scale based on Construal Level Theory in psychology. The scale frames individual attitudes about risk likelihood based on social and physical distance to the privacy harm. The findings include predictions about the extent to which the above factors correspond to risk acceptance, including that perceived risk is lower for induced disclosure harms when compared to surveillance and insecurity harms as defined in Solove's Taxonomy of Privacy. We also found that participants are more willing to share their information when they perceive the benefits of sharing. In addition, we found that likelihood was not a multiplicative factor in computing privacy risk perception, which challenges conventional theories of privacy risk in the privacy and security community.

CCS Concepts: • **Security and Privacy** → **Human and societal aspects of security and privacy**.

KEYWORDS

Privacy, Privacy Risk Perception, Factorial Vignettes, Multilevel Modeling.

ACM Reference format:

Jaspreet Bhatia, Travis D. Breau, XXX 2017. Title. *ACM Trans. Computer Human Interaction*. 9, 4, Article 39 (March 2010), 6 pages.
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1 INTRODUCTION

Information systems increasingly use personal information in sensitive ways, including recommender systems, personal navigation, and communication over social networks. While the benefits of using these services may outweigh the risks to personal privacy, users can be exposed to privacy harms, such as harms due to automated interferences that affect personal employment, financial and health-related decisions [Solove 2006]. Privacy threats arise when information services are personalized by collecting online behavioral data about users over the course of their service use. This data may include keyword searches, how users click on or scroll through service content, and what actions they take with the service, among others. For specific services, this data can be particularly sensitive as it may describe, for example, user interest in products, their lifestyle choices, and

their locations where they spend their work and free time. Consequently, this poses a critical challenge to service designers because collecting and analyzing user information is necessary to provide personalized and relevant services to users. The privacy paradox describes how users will share their personal information, despite their stated concerns about how the information will be used [Acquisti and Grossklags 2005, Berendt et al. 2005]. An alternative explanation for why users exhibit conflicting behaviors when sharing their personal information is that users perceive privacy harms in terms of a cost-benefit tradeoff, wherein the risk of a privacy harm is reduced when users perceive an increase in the benefits of sharing their personal information with a service. In this paper, we examine this alternate explanation in a series of experiments that present privacy cost-benefit tradeoffs to potential users, also called data subjects. We observed from our surveys that user perceptions of privacy risk decrease when users perceive benefits from sharing their information. Designers and developers can make informed decisions about designing personalized systems, if they have insight into what users care about with respect to their personal privacy, and what tradeoffs they are willing to make under which scenario. Using the framework proposed in this paper and the initial findings from the reported surveys, designers and developers can determine user privacy concerns with respect to different factors and the benefits for which users are willing to make a tradeoff by sharing their information.

In order to design systems that mitigate privacy risk, and that adhere to privacy by design, we need to determine the extent to which different factors effect user perceptions of privacy risk, such that measures can then be taken to reduce that risk. In addition, when companies need to choose one among two data practices, one which calls for deletion of a given data type, and the other which calls for storing the data for future data analytical purposes, companies need a way of determining what user data it can store and use under which scenarios, so as to minimize the privacy risk experienced by the user. The challenges described above can be addressed, if we have a framework to measure privacy risk in a given context and determine the factors that affect privacy risk. Towards this end, we propose an empirical framework that can be used with different factors to measure their effect on perceived privacy risk, and then we report results from the surveys conducted using this framework to measure the effect of six different factors on the perceived privacy risk. Using this framework companies and designers can identify types of user data that are most at risk and take measures to protect these types of data. Regulators can use the framework and the current findings to compute the privacy risk score for a given software system, and also to make sure privacy policies describe in detail practices related to high risk data types and scenarios. In addition, this framework can be used by users to measure the privacy risks posed by a website which would consequently help them make informed decisions about using the website.

The proposed empirical framework consists of a collection of surveys that are tailored to fit an information technology scenario. The surveys can be administered to actual or potential users of a system, to data subjects, or the general public. As shown in Figure 1, the framework consists of pre-tests, one or more factorial vignette surveys, and post-tests. The pre-tests measure participant exposure to risks and how they rank the technological benefits. The exposure survey asks participants to report the frequency of their participation in online activities, such as online shopping or banking, among others. In addition, the exposure survey asks participants about their experiences of privacy harms. The exposure survey is conducted as a pre-test prior to asking participants about their risk tolerances, or as a separate study to inform factorial vignette design. Similar to the exposure survey, the benefits ranking survey identifies benefits with the greatest and least impact on individual risk perceptions. Each vignette consists of a scenario with multiple contextual factors, a risk likelihood scale, and a risk acceptance scale. The scenarios situate participants in the context of a specific cost-benefit tradeoff. Finally, the factorial vignette survey is followed by a post-test demographic survey to compare the sample population against standard demographics, such as age, gender, education level, and income. The post-survey helps determine the extent to which the collected risk measures will generalize to the population of interest.

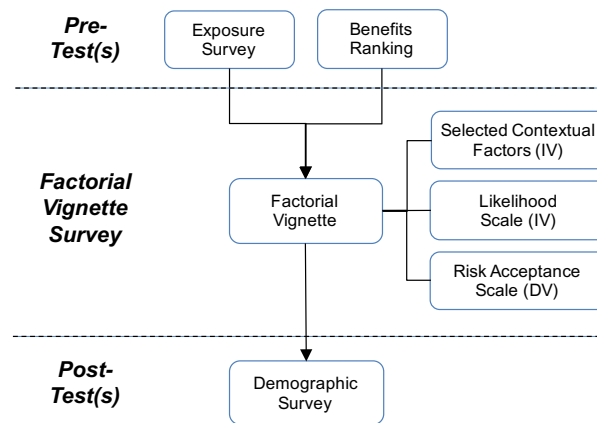


Fig. 1. Empirical framework to measure perceived privacy risk

The contributions of the paper are as follows:

- Empirical framework to measure perceived privacy risk, which is situated in a controlled context with factors that can be experimentally manipulated.
- Likelihood scale, which is based on Construal Level Theory and can be used to manipulate the perceived realism of the privacy harm.
- An evaluation of the framework for different factors, including risk likelihood, privacy harm, data type, computer type, data purpose and demographics, such as age, gender, education, ethnicity, and income.

We observed from the results of applying this framework that privacy risk is impacted by several factors, including the type of data affected, the type of privacy harm, the benefits of sharing information, and in some instances by demographic factors, such as education level and ethnicity. In contrast to earlier definitions of privacy risk, we found that risk likelihood was not a multiplicative factor in predicting privacy risk. Participants in our studies perceived the induced disclosure privacy harm to be minimally harmful as compared to other harms surveyed, and the privacy harm insecurity to be most harmful. In addition, we found that specifying the benefits of sharing information decreases the perceived privacy risk.

The paper is organized as follows: in Section 2, we discuss the related work and background on privacy, risk perception and privacy risk; in Section 3, we describe the factorial vignette survey method and multilevel modeling, which is the statistical method for analyzing the survey data; in Section 4, we present our research questions for evaluating the framework, the study designs to address those questions, and the study results; in Section 5, we discuss our results for each research question; and finally, in Section 6, we present the conclusion and the future work.

2 RELATED WORK

In this section, we review related work on privacy, risk perception and privacy risk.

2.1 Background on Privacy

Over the course of the last century, multiple definitions of privacy have emerged. Westin describes privacy as when a person, group or company can decide for themselves when, how and to what extent information about them is shared with others. Westin defines four states of privacy: (1) solitude, which refers to how one person distances his or herself from others, (2) intimacy, where a person chooses to have a close relationship with a small group of people, (3) anonymity, where a person can move through public spaces while protecting his or

her identity, and (4) reserve, where a person can regulate the amount of information about himself or herself that one wants to communicate to others in order to protect against unwanted intrusion [Westin 1967]. Murphy describes the “right to privacy” as being safe from intrusion, the right to make confidential decisions without government interference, the right to prohibit public use of a person’s name or image, and to regulate the use of personal information [Murphy 1996]. Nissenbaum argues that privacy and data sharing are contextual, meaning that the factors, data type, data recipient, and data purpose among others affect a person’s willingness to share [Nissenbaum 2004, 2009]. Murphy and Nissenbaum’s definition of privacy motivates our choice of the factorial vignette survey design, which allows one to present multiple, co-occurring factors in a single survey while attributing effects to individual factors (and their combinations).

There are different and conflicting views about the importance of privacy. Solove argues that privacy is “a fundamental right, essential for freedom, democracy, psychological well-being, individuality, and creativity” [Solove 2008]. On the other hand, other scholars, such as Moor, argue that privacy is not a “core value” in comparison to the values of life, happiness, and freedom; rather privacy is an expression of the core value of security and asserts that privacy is instrumental for protecting personal security [Moor 1997]. In addition, Solove also notes that with the ardent of new technologies an array of new privacy harms has also arisen [Solove 2006]. To the best of our knowledge there have been no prior studies that measure the differences in privacy risk due to the different privacy harms. This led us to our third research question, “How do different harms affect the perception of privacy risk, in the presence of controlled benefits? (RQ3)” From our surveys to answer RQ3, we observe that some privacy harms are perceived as less and more harmful, notably that the privacy harm *induced disclosure* was least risky, and the harm *insecurity* was most risky among the harms we studied.

Studies have shown differences between a user’s privacy preferences and their actual behavior in similar situations, called the privacy paradox [Acquisti and Grossklags 2005, Berendt et al. 2005]. This paradox could be explained by the argument made by Slovic et al. that people who see social or technological benefits of an activity tend to perceive a reduction in risks associated with that activity [Slovic 2000], which also motivated two of our research questions described in Section 4: how do different benefits affect the perception of privacy risk, in the presence of controlled harms (RQ2), and how do different data types, in the presence or absence of benefits affect the perception of privacy risk (RQ4). The studies reported in this paper further support this argument, that perceived benefits from services will reduce the user’s perception of privacy risk.

2.2 Risk Perception and Privacy Risk

Risk is a multidisciplinary topic that spans marketing, psychology, and economics. In marketing, risk is defined as a choice among multiple options, which are valued based on the likelihood and desirability of the consequences of the choice [Bauer 1960]. Starr defines risk as a function of likelihood and magnitude [Starr 1969]. Risk has also been described as a function of the probability and consequence, where consequence is the measure of damage [Kaplan and Garrick 1981]. More recently, NIST defines risk as the likelihood times the impact of an adverse consequence or harm [Stoneburner 2002]. These definitions of risk motivate our first research question (RQ1) which concerns the extent to which we can manipulate, increase or decrease, an individual’s perception of risk likelihood, and measure its effect on perceived privacy risk. Throughout this paper, the term *risk likelihood* means the likelihood of a privacy harm. It has been shown using construal-level theory from psychology that along four dimensions of increased spatial, temporal, social and hypothetical distances people correlate increased unlikelihood than they do with shorter psychological distances along these four dimensions [Wakslak and Trope 2009]. As correlate measures of likelihood, we chose spatial and social distance as follows: a privacy harm affecting only one person in your family is deemed a psychologically closer and more likely factor level than one person in your city or one person in your country, which are more distal and perceived less likely. The risk likelihood levels used in the framework are as follows, ordered from most likely and least hypothetical to least likely and most hypothetical: Only one person in your family, Only one person in your workplace, Only one person in your city, Only one person in your state and Only one person in your country.

Starr first proposed that risk preferences could be revealed from economic data, in which both effect likelihood and magnitude were previously measured (e.g., the acceptable risk of death in motor vehicle

accidents based on the number of cars sold) [Starr 1969]. In psychology, Fischhoff et al. note that so-called revealed preferences, assume that past behavior is a predictor of present-day preferences, which cannot be applied to situations where technological risk or personal attitudes are changing [Fischhoff et al. 1978]. To address these limitations, the psychometric paradigm of perceived risk emerged in which surveys are designed to measure personal attitudes about risks and benefits [Slovic 2000]. Fischhoff et al. also describe risk as an individual's willingness to participate in an activity [Fischhoff et al. 1978]. For example, one accepts the risk of a motor vehicle accident each time they assume control of a motor vehicle as the driver. In this paper, we measure perceived privacy risk similar to Slovic [Slovic 2000] by estimating a computer user's willingness to share data, including but not limited to personal data, similar to Fischhoff et al. [Fischhoff et al. 1978]. We use the participant's willingness to share their data as the dependent variable in our studies to measure the perceived privacy risk.

Two insights that emerged from the psychometric paradigm of perceived risk and inform our approach and surveys are: (a) people better accept technological risks when presented with enumerable benefits, and: (b) perceived risk can account for benefits that are not measurable in dollars, such as lifestyle improvements, which includes solitude, anonymity and other definitions of privacy [Slovic 2000]. In other words, people who see technological benefits are more inclined to see lower risks than those who do not see benefits. Notably, privacy is difficult to quantify, as evidenced by ordering effects and bimodal value distributions in privacy pricing experiments [Acquisti et al. 2013]. Rather, privacy is more closely associated with lifestyle improvements, e.g., private communications with friends and family, or the ability to avoid stigmatization. Acquisti et al. observed that estimated valuations of privacy were larger when the participants of the study were asked to consider giving up their personal data for money and smaller when they had to pay money for privacy [Acquisti et al. 2013]. Their studies also showed that the participants' decisions about privacy were inconsistent. Finally, the economist Knight argues that subjective estimates based on partial knowledge represent uncertainty and not risk, also known as ambiguity aversion, wherein respondents are unwilling to accept a risk due to uncertainty in the question or question context [Knight 1921].

Some researchers believe that one can measure the "actual" privacy risk, which is a hypothetical, data subject-independent measure of the above-chance probability that any data subject would experience a privacy harm. The concept of an "actual" privacy risk would require continuous surveillance data on data subjects, which details how a system affects those subject's emotional, psychological and physical well-being. This data would include whether data subjects accept a risk by participating in an activity. Fischhoff et al. argue that people's behavior does not reliably reflect an actual risk estimate, if they cannot iterate over the system's design space, including both the possibility of hazards and reliability of safety features [Fischhoff et al. 1978]. In addition, accumulating this surveillance data would introduce a privacy risk paradox, in which the measurement of actual risk would introduce a new, more serious risk by amassing this surveillance data. Finally, the measure of whether a data subject actually experiences a privacy harm, such as whether a data subject's personal information were distorted or mischaracterized, is necessarily a subjective assessment. Fischhoff et al. argue that such assessments are subject to estimator biases and their methods of assessment, if not well documented, can be difficult to reproduce [Fischhoff et al., 1978]. Therefore, while actual privacy risk presents an objective ideal, the concept's general validity and reliability has been criticized in prior work.

2.3 Privacy and Privacy Risk in Human-Computer Interaction

In human-computer interaction (HCI), Palen and Dourish describe privacy as an ongoing process of negotiating boundaries of disclosure, identity and how these concepts evolve over time, in their meaning and interpretation [Palen and Dourish 2003]. They argue that managing privacy involves dealing with ever changing situations rather than just implementing existing rules. They also consider privacy as managing tradeoffs that arise from competing or conflicting needs, and taking into account how technology can break existing barriers and create new barriers [Palen and Dourish 2003]. Lederer et al. describe the "space of privacy" as a non-exhaustive set of interdependent dimensions that define the privacy implications to end users of different phenomenon, such as technical systems, policies, practices and incidents. They cluster these dimensions into three categories: system properties, actor relations, and information types. System properties are the details of the disclosure and

the extent of the user participation in the disclosure. Actor relationship is the relationship between the observer and the subject observed, and how they have been connected in the past, which in turn affects how the subject's information might be used by the observer, and whether the subject trusts the observer. Information types define the extent to which information may be sensitive, and if the disclosure has been made intentionally [Lederer 2003], which led us to study the effect of different data types on privacy risk (RQ4). In addition, discomfort that users experience when participating in privacy-sensitive activities has also been considered in previous research conducted in HCI [Olson et al. 2005 and Wang et al. 2011]. This motivates us to study if predictors like discomfort, identifiability and personal nature of data correlate with the privacy risk, which is our sixth research question (RQ6).

Privacy is defined by Saltzer and Schroeder as “the ability of an individual or organization to decide whether, when, and to whom personal or organizational information is released” [Saltzer and Schroeder 1975]. The concept that organizations have privacy has been used to justify intellectual property as a privacy issue, which is not a mainstream view of privacy. Saltzer and Schroeder note privacy differs from security, which is defined as the “mechanisms and techniques that control who may use or modify the computer or the information stored in it” [Saltzer and Schroeder 1975]. Privacy risk concerns individual users, their behavior and relationships to others, whereas security risks are risks posed by adversaries who attack or threaten a system [Hong et al. 2004].

Risk management has long been used to identify, assess, and prioritize risks and to develop effective risk minimization techniques. While risk analysis is not widely used in HCI design [Iachello and Hong 2007], risk models have been proposed in HCI to address privacy risks. In privacy risk management, designers manage privacy risks by using techniques and strategies, such as categorization, prioritization of risk and the development of interaction techniques to reduce risk [Hong et al. 2004]. Hong et al. introduce a privacy risk analysis consisting of a set of questions, as the first step in their privacy risk model, which aims to encourage system designers to think more deeply about privacy risk concerns [Hong et al. 2004]. The questions are organized into two groups, one concerning the social and organization context in which the system functions (e.g., what kinds of personal information are shared and under what circumstances?), and the second concerning technology used to implement the system (e.g., how long is personal information retained and who has access to it?) Hong et al. provide candidate questions that can be used as a starting point in both groups and can be refined further based on the user base and domain. This risk analysis can be used to understand the average cases in which the application is expected to be normally used, as well as for special cases. The outcome includes potential privacy risks created by the system.

Lederer et al. identified five pitfalls that designers should avoid in interactive design for privacy. These pitfalls are: (1) obscuring potential information flow, (2) obscuring actual information flow, (3) emphasizing configuration over action, (4) lacking coarse-grained control, and (5) inhibiting existing practice [Lederer et al. 2004]. Hilty et al. provide a qualitative approach to risk analysis for pervasive computing that consists of three steps: (1) developing scenarios, (2) screening for potential risks, and (3) applying a risk filter to guide the risk analysis. They developed three kinds of scenarios, which they use in the screening phase: cautious, in which users are cautious of the technology, high-tech, in which users accept the technology, if it is feasible both technologically and economically, and average, in which a tradeoff exists between caution and acceptance. In the screening phase, experts identify the risks associated with a particular application. The experts then prioritize the risks by filtering risks using the following criteria: socioeconomic irreversibility, which concerns whether a user's status can be restored to what it was before the technology came into effect; delay effect, which concerns the delay between the user's use of the technology and the technology's negative effect; potential conflicts, which occurs if the exposure to the risk is intentional or voluntary, and if there are any externalities present (e.g., fairness); and burden on posterity, which concerns whether future generations could be compromised. The authors used this framework to analyze the social and technical risks of ubiquitous computing technologies, including their social and environmental impact [Hilty et al. 2004].

The techniques proposed by Hong et al., Lederer et al. and Hilty et al. above rely on heuristic-based decision making by designers, the success of which depends on designer familiarity with privacy threats and knowledge of how users perceive privacy threats based on multiple factors surrounding the context of information use. In addition, how users perceive risks may depend not only on the designer's system, but more broadly on the

environment in which the system is situated. Consequently, designers perform their risk analysis at design time, whereas privacy risk can change over the course of a system's evolution or due to changes in its environment. To address these challenges, designers need a measure of privacy risk that measures user perception of risk and that can be re-measured over the course of a system's lifetime to check whether the design decisions continue to offer the protections proposed at design-time, which can be accomplished by using the framework proposed in this paper.

In the next section, we introduce our survey designs and the statistical methods that comprise the empirical framework to measure privacy risk.

3 FACTORIAL VIGNETTE SURVEY DESIGN AND MULTILEVEL MODELING ANALYSIS METHOD

The empirical framework to measure perceived privacy risk introduced in this paper consists of multiple pre-survey questions to measure the participant's online behavior, factorial vignette surveys to measure their perceived privacy risk in different contexts, and post survey demographic questions (see Figure 1 in Section 1). In this section, we discuss factorial vignette survey design along with the scale we developed to measure risk likelihood, followed by the statistical method used to analyze the data, called multilevel modeling.

3.1 Factorial Vignette Survey Design

Factorial vignettes provide a method to measure the extent to which discrete factors contribute to human judgment [Auspurg and Hinz 2014]. The factorial vignette method employs a detailed scenario with multiple factors and their corresponding levels, designed to obtain deeper insights, into a person's judgment and decision principles, than is possible using direct questions (i.e., with a prompt "Please rate your level of perceived risk" and a scale). Our factorial vignette survey design measures the interactions between the different independent variables, and their effect on a dependent variable, the person's *willingness to share their personal information*. This includes whether the different independent variables alone, in combination, or none of these factors affect willingness to share.

The factorial vignettes are presented using a template in which factors correspond to independent variables and each factor takes on a level of interest. For each factorial vignette survey (see Section 4), the factor levels replace an independent variable in the survey. The factors are often presented in the context of a scenario, which serves to situate the survey participant in a specific context. For example, a vignette may ask a participant to think about an online shopping experience with a website they routinely use, or to think about applying for a job online at an employment website. While the primary scenario does not change across vignettes, the embedded factors do change. For example, if we are interested in whether privacy risk changes when a person is using a workplace computer versus personal smart phone while shopping online, the survey designer can introduce a new factor $\$CT$ with two levels: workplace computer, and personal smart phone. For a between-subjects variable, a participant only sees and judges one level of the factor, whereas for a within-subjects variable, the participant sees all factor levels. In Figure 2, we present a vignette for an example study with three independent variables, which are data purpose ($\$DP$), computer type ($\CT) and data type ($\$DT$), and a dependent variable, which is willingness to share ($\$WtS$). The variable $\$DT$ is a within-subjects variable, which means that all the participants see and rate all the levels of this variable, whereas the variables $\$DP$ and $\$CT$ are between-subject variables, and each participant sees and rates only one level of this variable. In this vignette, the place holders for the variables are replaced by the values of the levels of these variables for each participant. For instance, for the variable computer type, the variable placeholder $\$CT$ will be replaced by either one of the two levels of this variable, workplace computer, or personal smart phone. The semantic scale for $\$WtS$ consists of eight options starting from Extremely Unwilling (0) to Extremely Willing (8), part of the scale has been omitted for brevity (...).

Please rate your willingness to share your information below with the Federal government, for the purpose of **\$DP**.

When choosing your rating for the information types below, consider the **\$CT** and the purpose above.

	Extremely Willing	Very Willing	Willing	Somewhat Willing	Somewhat Unwilling	...
Age Range	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	
Home Address	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	

Fig. 2. Example Factorial Vignette

In the framework, the vignette survey designer selects multiple contextual factors to include in the scenario. Nissenbaum argues that privacy and data sharing are contextual, and that users are more concerned about their information flowing appropriately, rather than restricting the flow of their information. The author describes appropriate information flow using the framework of contextual integrity, which takes into account the factors that determine whether users will perceive a new technology or system as a risk to their privacy [Nissenbaum 2009]. We evaluate contextual integrity through perceived risk using the following factors in Section 4: the data type shared, the data recipient, the data purpose (also called the benefit of sharing), and the privacy harm.

Kaplan and Garrick define risk as a function of the probability and consequence, where consequence is the measure of damage [Kaplan and Garrick 1981]. More recently, NIST defines risk as the likelihood times the impact of an adverse consequence or harm [Stoneburner 2002]. One approach to measure probability or likelihood is to describe the number of people affected by the adverse consequence: the greater the number of people affected, the greater the probability is that the consequence may affect a randomly selected person. When considering how many people are affected by a consequence, prior research shows that lay people can map ratios (e.g., 1/10,000) to physical people much better than they can map probabilities (e.g., 0.0001%) [Fischhoff et al. 1978]. To evaluate this conclusion, we pilot tested a between-subjects risk likelihood factor with ratio-based likelihood levels. The risk likelihood had four levels, which were the ratios of people who experienced the privacy harm: 1/4, 1/10, 1/100 and 1/1,000. In the pilot study, we found no significant effects among the ratios, which suggests that participants perceive no greater privacy harm when the harm affects 1/4 people versus 1/1,000 people.

As an alternative to ratios, we designed a new risk likelihood scale based on construal-level theory from psychology. Construal-level theory shows that people correlate increased unlikelihood along four dimensions of increased spatial, temporal, social and hypothetical distances, than they do with shorter psychological distances along these four dimensions [Wakslak and Trope 2009]. We chose spatial and social distance as correlate measures of likelihood as follows: a privacy harm affecting only one person in your family is deemed a psychologically closer and more likely factor level than one person in your city or one person in your country, which are more distal and perceived less likely. The risk likelihood levels used in the framework are as follows, ordered from most likely and least hypothetical to least likely and most hypothetical:

- Only one person in your family
- Only one person in your workplace
- Only one person in your city
- Only one person in your state
- Only one person in your country

The evaluation of the likelihood scale is reported later in Section 4.1.

Risk has been described in terms of an individual's willingness to participate in an activity [Fischhoff et al. 1978], for example, one accepts the risk of a motor vehicle accident each time they assume control of a motor vehicle as the driver. To measure privacy risk, we propose to estimate a computer user's willingness to share

data, including but not limited to personal data. The independent variable willingness to share ($\$WtS$) is estimated from survey participant ratings on an eight-point, bipolar semantic scale, labeled at each anchor point: 1=*Extremely Unwilling*, 2=*Very Unwilling*, 3=*Unwilling*, 4=*Somewhat Unwilling*, 5=*Somewhat Willing*, 6=*Willing*, 7=*Very Willing* and 8=*Extremely Willing*. This scale omits the midpoint, such as “Indifferent” or “Unsure,” which can produce scale attenuation when responses are prone to cluster, and which can indicate vague or ambiguous contexts rather than a respondent’s attitude [Kulas and Stachowski 2013].

3.2 Multilevel Modelling Analysis Method

Multilevel modeling is a statistical regression model with parameters that account for multiple levels in datasets, and limits the biased covariance estimates by assigning a random intercept for each subject [Gelman and Hill 2006]. Multilevel modeling has been used to study interactions among security and privacy requirements [Bhatia et al. 2016a, Hibshi et al. 2015].

In our studies, the main dependent variable of interest is *willingness to share*, labeled $\$WtS$. We conducted multiple studies, that have different independent variables of interest that affect our dependent variable $\$WtS$. For the within-subject design, subject-to-subject variability is accounted for by using a random effect variable $\$PID$, which is a unique identifier for each participant. Equation 1 below is our main additive regression model with a random intercept grouped by participant’s unique identifier. The additive model is a formula that defines the dependent variable $\$WtS$, *willingness to share*, in terms of the intercept α and a series of components, which are the different independent variables ($\$IV_1$, $\$IV_2$ and so on). Each component is multiplied by a coefficient (β) that represents the weight of that variable in the formula computed from the survey results that satisfies the multi-level modeling equation. The formula in Equation 1 is simplified as it excludes the dummy (0/1) variable coding for reader convenience.

$$\$WtS = \alpha + \beta_1\$IV_1 + \beta_2\$IV_2 + \dots + \epsilon \quad (1)$$

We analyze the data from our studies in R [R Core Team 2015] using the package lme4 [Bates et al. 2015]. We test the multilevel models’ significance using the standard likelihood ratio test: we fit the regression model of interest; we fit a null model that excludes the independent variables used in the first model; we compute the likelihood ratio; and then, we report the chi-square, p-value, and degrees of freedom [Gelman and Hill 2006]. We performed a priori power analysis for each study using G*Power [Faul et al. 2007] to test for the required sample size for repeated measures ANOVA. In addition, we report the R-squared marginal, which represents the variance explained by the fixed effects in the model, and R-squared conditional, which represents the variance explained by the fixed and random effects in the model [Nakagawa and Schielzeth 2013]. These R-squared measures are implemented in the package MuMIn in R [Barton 2014]. We also report the effect size for each independent variable from the ANOVA analysis, which concerns the extent to which the independent variable affects the dependent variable [Cortina and Nouri 2000]. We perform ANOVA on the multi-level model and report the sum of squared errors for each independent variable and the proportion of variance explained by the independent variable.

4 MEASURING PRIVACY RISK PERCEPTION

We now describe our approach to evaluate the empirical privacy risk measurement framework by answering the following research questions:

- RQ1.** How can we manipulate, increase or decrease, an individual’s perception of risk likelihood, and measure its effect on the perceived privacy risk?
- RQ2.** How do different benefits affect the perceived privacy risk, in the presence of controlled harms?
- RQ3.** How do different harms affect the perceived privacy risk, in the presence of controlled benefits?
- RQ4.** How do different data types, in the presence or absence of benefits affect the perceived privacy risk?
- RQ5.** Does perceived privacy risk vary with a user’s computer setting (e.g., workplace computer or personal smart phone)?

RQ6. Does discomfort, identifiability or the personal nature of data co-vary with, or supplement, perceived privacy risk?

RQ7. How do demographic factors influence the perceived privacy risk?

We designed four factorial vignettes studies, based on the framework introduced in Figure 1, which help us answer the seven research questions (RQ1-RQ7) and three additional surveys to answer research question RQ6. All the studies were conducted independent of each other. Each participant was allowed to participate in each of the factorial vignette studies only once to ensure each level of the between-subjects variable was seen by different participants. Each participant was permitted to respond to one or more of the four factorial vignette studies, as they were conducted independent of each other. On the other hand, each participant could respond to any number of the three additional surveys conducted to measure the discomfort, identifiability and personal nature of data. For all our surveys, we recruited English-speaking participants from Amazon Mechanical Turk (AMT), located in the US, and who had completed ≥ 5000 HITs. We only recruited AMT workers with a $\geq 97\%$ approval rating. The mean time to complete the pilot survey was ~ 20 minutes, thus we allowed 45 minutes for participants to complete the surveys. We paid between \$3 to \$6 per participant for the different surveys, and we published the surveys online using Survey Gizmo.

Before designing the vignette surveys, we conducted a pilot study with an exposure survey of 96 people to understand how often they participate in online activities, and how often they experience privacy harms while using the Internet. In this exposure survey, we asked participants how frequently they perform six activities online: watching television, reading news; sharing medical information with doctors; paying bills, checking bank account balances, or transferring money; shopping for products or services; and using social networking sites. These activities were chosen from the 2015 PEW Internet and American Life Survey of Internet Users [Perrin and Duggan 2015]. The response options for frequency of online behavior are: *a few times a day*, *once a day*, *a few times a week*, *a few times a month*, *a few times a year*, and *never*. Figure 3 shows the frequency of online behavior reported by the participants in the pilot study.




Fig. 3. Exposure Survey on Frequency of Online Behavior

Reading news online was reported as the most frequent activity, and we found that 84% of participants shop online at least a few times a month, and every participant reports shopping online at least once. In addition, shopping online is an activity in which users often provide personal information, such as their shipping address and payment information. Therefore, we chose shopping as the scenario context for Study 1 and Study 2 (see Table 1 for the complete list of studies).

In addition to online behaviors, we surveyed participants to ask how frequently they experience seven different privacy harms from the NISTIR 8062 framework for privacy engineering, which are as follows:

- *Appropriation* is when you feel that your personal information is being used in unexpected ways.
- *Distortion* is when you feel that others are using or disseminating inaccurate, misleading or incomplete information about you.
- *Induced Disclosure* is when you feel the pressure to divulge your personal information to others.
- *Insecurity*, is when you feel that there are lapses in security aimed to protect your personal information.
- *Surveillance* is when you feel that you are being tracked or monitored.
- *Unanticipated Revelation* is where you feel that some information about you is being revealed or exposed.
- *Unwarranted Restriction* is where you feel that you are unable to access or control your personal information.

In Figure 4, we present the reported frequencies of experiencing these harms in our pilot study: notably, the three most frequently experienced harms are surveillance with 37% of respondents with weekly experiences, followed by insecurity and induced disclosure with 39% of respondents reporting that they experience these two harms monthly. We report the effects of these harms on risk perception in Section 4.2.




Fig. 4. Exposure Survey on Frequency of Experiencing Privacy Harms

We designed four factorial vignette studies to address our research questions. The scenario topic for Studies 1 and 2 was chosen from the results of our pilot study as described above. The 2016 Cybersecurity Sharing Act (CISA) codifies portions of Executive Order 13691 into law, which supports the sharing of private-sector incident data with the US government. Due to this sharing of cybersecurity intelligence with the government, companies have reported concerns about violating customer and employee privacy [PWC 2016]. The concern about privacy risk motivates our choice of scenario for Studies 3 and 4, which is sharing personal information with the government to investigate cybersecurity incident.

In Table 1, we summarize the research questions addressed by each study, the independent factor for each study, and if the factor is a within-subjects or between-subjects factor).

Table 1. Study Designs to Evaluate the Empirical Privacy Risk Framework

Research Study	Scenario Topic	Research Questions Answered	Independent Factors
Study 1	Routine sharing with website while shopping online	RQ1, RQ3, RQ7	Risk likelihood (within), data types (within), privacy harms (between)
Study 2		RQ3, RQ7	Risk likelihood (between), data types (within), privacy harms (within)
Study 3	Sharing with government to investigate cybersecurity incident	RQ2, RQ4, RQ5, RQ7	Risk likelihood (between), data types (within), computer type (within)
Study 4		RQ2, RQ4, RQ5, RQ7	Risk likelihood (between), data types (within), computer type (within), data purpose (within)

In the following sections, we describe each study design and report results. We have six sub-sections in this section, and in each sub-section we report design and results from studies for measuring the effect of one or more factors on privacy risk.

4.1 Risk Likelihood and Perceived Privacy Risk

The first research question RQ1 concerns the extent to which we can we manipulate, increase or decrease, an individual's perception of risk likelihood and measure its effect on perceived privacy risk. We designed Studies 1 and 2 to answer this question. As shown in Figure 1 and discussed in Section 3, we introduced a risk likelihood scale based on construal level theory. The risk likelihood levels used in the framework are as follows, ordered from most likely and least hypothetical to least likely and most hypothetical:

- Only one person in your family
- Only one person in your workplace
- Only one person in your city
- Only one person in your state
- Only one person in your country

We designed Study 1 to measure the effect of the different levels of the risk likelihood scale on our dependent variable *willingness to share* ($\$WTS$). In this study, we had three independent variables, including *risk likelihood* ($\$RL$) and *data type* ($\DT), which were both within-subjects factors, and *privacy harm* ($\$PH$), which was a between-subjects factor (see Table 2 for the factor levels).

Table 2. Vignette Factors and their Levels for Study 1 and 2

Independent Factors	Factor Levels
Risk Likelihood ($\$RL$) Study 1: Within-Subject Study 2: Between-Subject	Only one person in your family
	Only one person in your workplace
	Only one person in your city
	Only one person in your state
	Only one person in your country
Data Types ($\$DT$) Study 1: Within-Subject Study 2: Within-Subject	Age range
	Credit Card Number
	Driver's License Information
	Full Name
	Home Address

Privacy Harms (\$PH) Study 1: Between-Subject Study 2: Within-Subject	Phone Number
	Unwarranted Restriction
	Unanticipated Revelation
	Surveillance
	Insecurity
	Induced Disclosure
	Distortion
	Appropriation

The template used for vignette generation for Study 1 is shown in Figure 5. The independent variables \$RL, \$DT and \$PH are each replaced by one level from Table 2 in the distributed survey.

Please rate your willingness to share your \$DT with a shopping website you regularly use, given the following benefits, privacy harm experienced and risks of using that website.

Benefits: Convenience, discounts and price comparisons, anonymous and discreet shopping, certainty that the product is available, wider product variety, and informative customer reviews.

Privacy Harm: \$PH

Given the above benefits and privacy harm, please rate your willingness to share your \$DT. Also consider the following levels of privacy risk:

Privacy Risk Levels	Extremely Willing	Very Willing	Willing	Somewhat Willing	...
Only one person in your family experienced the harm	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	
Only one person in your workplace experienced the harm	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	
...					

Fig. 5. Template used for vignette generation for Study 1 and Study 2 (fields with \$ sign are replaced with values selected from Table 2)

In Study 1, the survey was designed using the framework shown in Figure 1 and consists of three parts: pre-test questions about their exposure to privacy harms (see Section 4.2 for details), factorial vignettes with the independent factors, and the post test questions about their demographics including their gender, age range, education level, household income and ethnicity (see Section 4.6 for demographic results). In Study 1, we had three independent factors (see Table 2) with two within-subject factors: risk likelihood (5 levels) and data type (6 levels); and one between-subjects factor: privacy harms (7 levels). This survey has two variations which were each randomly assigned to half of the participants. In the first variation, all levels of the within-subjects factor risk likelihood were presented on the same page, with one level of the data type factor, and there were six such pages, one for each level of the data type. In the second variation, all the levels of the data type variable were presented on the same page, with one level of the risk likelihood factor. In both variations, the pages in the factorial vignettes were randomized, and so were the questions on each page. The two variations, and the randomization were done to mitigate any ordering effects.

Equation (2) is the main regression equation for Study 1. The regression equation represents the intercept with the baseline levels for the independent variables (α), the coefficients for the independent variables, and the random intercept to account for the subject to subject variability. The formula in Equation 2 is simplified as it excludes the dummy (0/1) variable coding for the reader’s convenience.

$$\$WtS = \alpha + \beta_R \$RL + \beta_{DT} \$DT + \beta_{PH} \$PH + \epsilon \quad (2)$$

Study 1 estimates the effect of the risk likelihood levels on the user's willingness to share when modeled as a within-subjects variable. This means that each survey participant will see each risk level at least once. We next describe Study 2, in which the variable risk likelihood was modeled as a between-subjects variable, meaning that all participants saw the same one factor level across all the vignettes that they were presented.

Consistent with Study 1, the Study 2 was designed to estimate the effects of three independent variables – *risk likelihood* ($\$RL$), *data types* ($\DT), and *privacy harms* ($\$PH$) on the dependent variable user's *willingness to share* ($\$WtS$). Study 2 uses the same factor levels as Study 1, also shown in Table 2. The design of Study 2 is similar to Study 1 as well, except that in Study 2 *risk likelihood* is a between-subjects factor, and *privacy harm* is a within-subjects factor, which changes how many and what levels of each factor were shown to the participants in the survey. In Study 2 we used the same approach as in Study 1 to vary the within- and between-subjects factors.

In Study 1, we found a significant contribution of the three independent factors $\$RL$, $\$DT$ and $\$PH$, for predicting the $\$WtS$ ($\chi^2(41)=2041.7$, $p<0.000$), over the null model, which did not have any of the independent variables, which means that variations in privacy risk are explained by the independent variables risk likelihood, data types and privacy harm. Similarly, for Study 2 we found a significant contribution of the three independent factors $\$RL$, $\$DT$ and $\$PH$, for predicting the $\$WtS$ ($\chi^2(41)=2911.6$, $p<0.000$), over the null model.

In Table 3, we present the Model Term, the corresponding model-estimated Coefficient (along with the p-value, which tells us the statistical significance of the term over the corresponding baseline level), and the coefficient's Standard Error for the factor risk likelihood. We report the coefficients for the factor data type for Studies 1 and 2 in Appendix B, Table B.1. The coefficients are the weights computed from the survey results that satisfy Equation 2 above. In our survey, the semantic scale option *Extremely Unwilling* has a value of 1, and *Extremely Willing* has a value of 8. A positive coefficient in the model signifies an increase in willingness to share and a negative coefficient signifies a decrease in willingness to share. The results in Table 3 show that $\$WtS$ is significantly different and increasing for decreasing levels of $\$RL$, as compared to the baseline level “only 1 person in your family.” For the $\$RL$ level “only 1 person in your workplace,” the $\$WtS$ increases by 0.34 over the baseline level, which denotes an increasing willingness to share, and lowest $\$RL$ level “only 1 person in your country” increases by 1.33 over the baseline, which is four times more likely than the $\$RL$ level “only 1 person in your workplace.” In addition, the $\$WtS$ for the combination workplace, age range, and induced disclosure is $4.662+0.338=5$.

Table 3. Multilevel Modeling Results for Risk Likelihood in Study 1

Model Term	Coeff.	Standard Error
Intercept (Family + Age Range + Induced Disclosure)	4.662***	0.750
Risk – only 1 person in your workplace	0.338***	0.061
Risk – only 1 person in your city	0.846***	0.061
Risk – only 1 person in your state	1.119***	0.061
Risk – only 1 person in your country	1.325***	0.061

* $p \leq 0.05$ ** $p \leq 0.01$ *** $p \leq 0.001$

Table 4 presents the Study 2 results for risk likelihood as a between-subjects variable, which means each participant saw the same one level of the variable across all the vignettes that they viewed. In Table 4, we did not see any significant differences between any of the levels of the variable.

Table 4. Multilevel Modeling Results for Risk Likelihood in Study 2

Model Term	Coeff.	Standard Error
Intercept (Family + Induced Disclosure + Age Range)	2.533**	0.880
Risk – only 1 person in your workplace	-0.104	0.315
Risk – only 1 person in your city	0.436	0.304
Risk – only 1 person in your state	0.149	0.311
Risk – only 1 person in your country	0.485	0.312

* $p \leq .05$ ** $p \leq .01$ *** $p \leq .001$

In Study 1, the risk likelihood variable was within-subjects, that is, all participants saw all the levels of the variable. From Study 1, we conclude that the willingness to share increases as a participant's social and physical distance from the person experiencing the privacy violation increases. This means that the users' perception of privacy risk increases, when they think about a person from their family or workplace experiencing the violation, as compared to the experience of a person somewhere in their state or country. This observation changes, however, when the variable is modeled as a between-subjects variable in Study 2. As a between-subjects variable, participants are less sensitive to the differences between risk levels and we therefore did not see any significant differences between the different levels of the variable.

The difference between the results when risk likelihood is a within-subjects factor in Study 1 and between-subjects factor in Study 2 can be attributed to the observation that between-subjects designs inherently have more noise as compared to within-subjects designs, and can therefore miss real and important differences among the levels of the independent factors [Charness et al. 2012]. Anchoring effects are an influence on the decision maker where the decisions are biased on the initial value presented [Tversky and Kahneman 1974]. To mitigate this bias, in the first study we randomize the order of the levels of the risk likelihood variable for each participant. In addition, in the first study, we have two variations of the survey. In the first variation, all levels of the risk likelihood factor were shown on the same page in a randomized order, and each page in the survey had one level of the other within-subjects factor data type, and the pages of the survey were randomized. Whereas, in the other variation the participants see all levels of the data type factor on the same page in a randomized order, and they see one level of the risk likelihood factor on each page. The page order is also randomized. In our survey, half of the participants were randomly assigned to the first variation, and half were assigned to the second variation, to yield an equal number of responses for both survey variations.

In Study 1, the fixed effects explain 17% of the variance in the dependent variable willingness to share (R-squared marginal=0.17), and both the fixed and random effects explain 56% of the variance (R-squared conditional=0.56). The sum of squares for Study 1 are as follows: risk likelihood (1518), privacy harm (9.7), information type (4038.1), and the total sum of squares for all the fixed effects is 5565.8. Therefore, risk likelihood explains 0.27 (1518/5565.8) of the variance explained by the fixed effects in the willingness to share, and information type explains 0.73 of the variance explained by the fixed effects in the willingness to share. Thus, the majority of the variation in willingness to share that is explained by the fixed effects is explained by the data type variable, followed by the risk likelihood variable.

In Study 2, the fixed variables explain 18% of the variance in the dependent variable willingness to share (R-squared marginal=0.18), and both the fixed and random variables explain 57% of the variance (R-squared conditional=0.57). The sum of squares for Study 2 are as follows: risk likelihood (9.3), privacy harm (713.4), information type (6732.3), and the total sum of squares is 7455. Therefore, privacy harm explains 0.10 (713.4/7455) of the variance explained by the fixed effects in the willingness to share, whereas the information type explains 0.90 of the variance explained by the fixed effects in the willingness to share.

4.2 Privacy Harms and Perceived Privacy Risk

Research question 3 concerns how the privacy risk changes in the presence of different privacy harms and controlled benefits. In Study 1 and 2, we estimate the effect of seven privacy harms ($\$PH$) – *appropriation*,

distortion, induced disclosure, insecurity, surveillance, unanticipated revelation and unwarranted restriction. (see Table 2 for the factor levels used) on a user's *willingness to share* ($SWtS$). The privacy harm definitions are from the NISTIR 8062 framework for privacy engineering, and were presented as follows:

- *Appropriation* is when you feel that your personal information is being used in unexpected ways.
- *Distortion* is when you feel that others are using or disseminating inaccurate, misleading or incomplete information about you.
- *Induced Disclosure* is when you feel the pressure to divulge your personal information to others.
- *Insecurity*, is when you feel that there are lapses in security aimed to protect your personal information.
- *Surveillance* is when you feel that you are being tracked or monitored.
- *Unanticipated Revelation* is where you feel that some information about you is being revealed or exposed.
- *Unwarranted Restriction* is where you feel that you are unable to access or control your personal information.

In Studies 1 and 2, we presented two pre-test questions about exposure to privacy harms. The first question asks participants to report the frequency with which they experience the privacy harms. The second question asks the participants to rank and score the harms based on their severity. In this second question, we present "severity" as "the degree of privacy harm you would experience." For the first question, which concerns the frequency with which the participants experience the privacy harm (see Figure 6), 80% of the participants reported that they experience *Surveillance* a few times a week or less, whereas 92% of the participants reported experiencing *Insecurity* a few times a week or less, and 93% of the participants reported that they experience *Induced Disclosure* a few times a week or less. Figure 6 shows how often the participants experience the harms across Studies 1 and 2.




Fig. 6. Exposure Survey on Frequency of Experiencing Privacy Harms

Figure 7 shows how participants rank the privacy harms. Out of a total 410 participants who were part of Study 1 and Study 2, 57% of the participants ranked *Surveillance* and 54% ranked *Insecurity* as one of the top three most-severe harms. Furthermore, 70% of the 410 participants ranked *Induced Disclosure* as one of the bottom three least severe harms. Unlike *Surveillance* and *Induced Disclosure*, which show a general trend in

the sample population, the harm *Unwarranted Restriction* notably had a near equally likely chance of appearing in any rank order. This suggests that *Unwarranted Restriction* affects different participants in very different ways. To measure participant agreement in severity rankings of privacy harms, we computed Kendall's coefficient of concordance W . Kendall's W measures the communality of judgments for different raters and ranges from 0 to 1 [Kendall 1948]. An increase in W from 0 to 1 signifies an increase in participant agreement. Using the ranks provided by 410 participants in Studies 1 and 2, Kendall's W is 0.096 ($p < 0.000$). Because W is close to 0, there is weak agreement among participants about which harms are relatively more or less severe.

In addition to agreement, one might ask whether harms that people experience frequently are also rated as most severe; in other words, are severity and frequency correlated? While participants clearly disagree about the order of severity, Kendall recommends finding the true ranking of ordinal data by computing an overall sum of ranks [Kendall 1948]. In this computation, the minimum rank sum is the "most severe," whereas the maximum rank sum is the "least severe," and in our data we had no ties. This yields the following true ranking from most to least severe: Surveillance, Insecurity, Unanticipated Revelation, Unwarranted Restriction, Appropriation, Distortion, and Induced Disclosure. To test for correlation between severity and frequency, we computed Spearman's rank correlation coefficient ρ [Spearman 1904]. Spearman's rank correlation coefficient is a non-parametric measure of rank correlation, which determines the extent to which the relationship between two ordinal variables can be described using a monotonic function. The frequencies were obtained from pre-test results shown in Figures 6 and Table A.1 in Appendix A. We found the true ranking for frequencies by computing the overall sum of ranks by assigning a rank number to each response option as follows: Several times a day (6), About once a day (5), A few times a week (4), A few times a month (3), A few times a year (2), Never (1). For *Appropriation*, for example, 13 people reported that they experience this harm *several times a day*, 17 participants reported *about once a day*, 52 reported *a few times a week*, 83 reported *a few times a month*, 147 *a few times a year* and 98 participants reported that they *never* experience appropriation. The sum of ranks for *Appropriation* is $13 \times 6 + 17 \times 5 + 52 \times 4 + 83 \times 3 + 147 \times 2 + 98 \times 1 = 1012$. The true rank order from most to least frequent is: Surveillance, Induced Disclosure, Insecurity, Appropriation, Unwarranted Restriction, Unanticipated Revelation, and Distortion. Spearman's ρ for severity and frequency is 0.32 ($p = 0.49$). Since the p-value is greater than 0.05, we conclude that privacy harm severity is not statistically dependent upon the frequency of experiencing the harm.

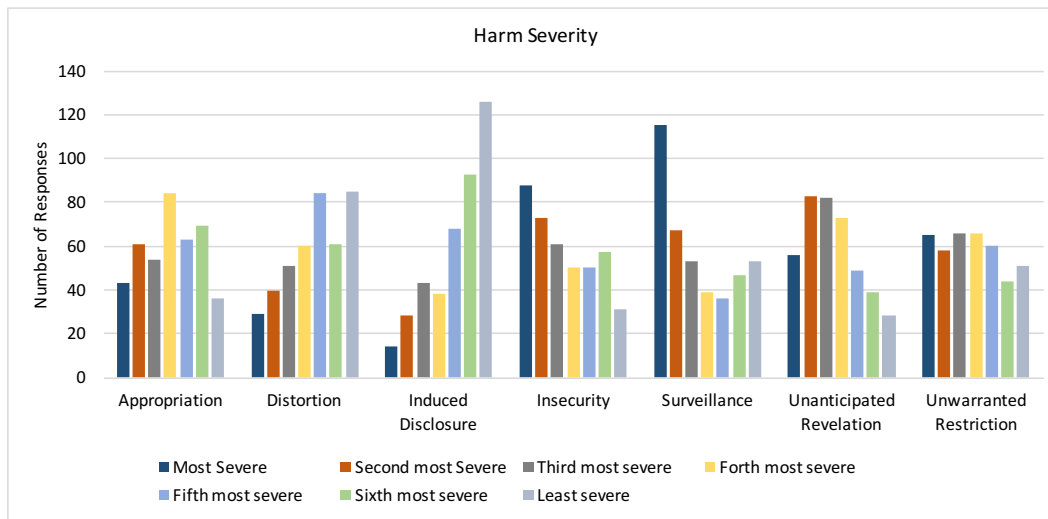


Fig. 7. Severity Based Rank Order for Privacy Harms

In Study 1 and 2, we investigated the extent to which privacy harm predicts changes in perceived privacy risk. Table 5 presents the *Model Term*, the corresponding model-estimated *Coefficient* along with the p-value, which tells us the statistical significance of the term over the corresponding baseline level, and the coefficient's *Standard Error* for $\$PH$ for Study 1 and Study 2 (see Equation 2 in Section 4.1 for the main regression equation). In our survey, the semantic scale option *Extremely Unwilling* has a value of 1, and *Extremely Willing* has a value of 8. A positive coefficient in the model signifies an increase in willingness to share and a negative coefficient signifies a decrease in willingness to share as compared to the baseline level.

Table 5. Multilevel Modeling Results for Privacy Harms for Study 1 and 2

Model Term	Study 1		Study 2	
	Coeff.	Standard Error	Coeff.	Standard Error
Intercept (Family+Age Range+Induced Disclosure)	4.662***	0.750	2.533**	0.880
Privacy Harm – Unwarranted Restriction	-0.647	0.382	-0.543***	0.060
Privacy Harm – Unanticipated Revelation	-0.338	0.381	-0.885***	0.060
Privacy Harm – Distortion	-0.199	0.382	-0.591***	0.060
Privacy Harm – Surveillance	-0.007	0.374	-0.530***	0.060
Privacy Harm – Insecurity	0.068	0.388	-0.928***	0.060
Privacy Harm – Appropriation	0.186	0.372	-0.363***	0.060

* $p \leq .05$ ** $p \leq .01$ *** $p \leq .001$

The results in Table 5 show that $\$WtS$ is significantly different for the levels of the independent variable $\$PH$, as compared to the baseline level, induced disclosure, in Study 2. The negative coefficients in Table 5 for the harms, show that the $\$WtS$ is the maximum for the baseline level *Induced Disclosure*, then decreasing for the other harms, and is the least for the harm *Insecurity*. This is consistent with our finding from the pre-test question about ranking harms based on severity, where 54% of the participants ranked *Insecurity* as the first, second, or third most severe privacy harm. And on the other hand, *Induced Disclosure* was rated as the least severe privacy harm. However, we did not see any significant differences between the harm levels when the privacy harms were presented as a between-subjects variable in Study 1.

In addition, we analyzed data from participants who had ranked *Insecurity* as one of the two most severe harms, and *Induced Disclosure* as the one of the two least severe harm. This yields a total 53 out of 200 total participants in Study 2. We found that the relative difference in the coefficients for $\$WtS$ between *Insecurity* and *Induced Disclosure* increases to 1.3 units (as compared to 0.93 from Table 5), where the $\$WtS$ for *Insecurity* was 1.3 units less as compared to *Induced Disclosure*. This implies that the participants who saw greater differences in the severity of the harms *Induced Disclosure* and *Insecurity*, also reported greater differences in their $\$WtS$ for these harms.

4.3 Benefits, Data Types and Perceived Privacy Risk

The research question RQ2 and RQ4 concern how different benefits affect the perception of privacy risk, in the presence of controlled harms, and how do different data types, in the presence or absence of benefits affect the perception of privacy risk. We designed studies 3 and 4 to answer these two research questions.

Fishhoff et al. suggest that people who see the benefits of performing an activity, tend to perceive the activity as low risk [Fishhoff et al. 1978]. We designed Study 3 and Study 4 to estimate the effects of benefits and data types on the user's willingness to share. The independent variable *data purpose* ($\$DP$) in this study has different levels of societal benefit, including terrorism, imminent threat of death, economic harm, and loss of intellectual property. These data purposes were chosen as benefits to society listed in the Cybersecurity Information Sharing Act of 2015 [S. 754]. In addition, the *data types* ($\$DT$) were chosen from the NIST Special Publication 800-61 guidelines on investigating and reporting cybersecurity incidents to determine which types are frequently used

in a forensic analysis [Cichonski et al. 2012]. We further surveyed security experts to measure the frequency with which they used these data types [Bhatia et al. 2016b]. We partitioned a total of 25 data types into three groups, which were chosen based on the relationship among types in a brief narrative to explain how these types arise in a shared computational setting. As we discuss later, the narrative was used to introduce the data types to participants in short videos. Study 3 and 4 serve to estimate the differences between the dependent variable $\$WTS$ for different data types in the presence and absence of benefits, and to further understand how benefits change the perception of perceived privacy risk.

In Study 3, we measure the effects of four independent variables – the *computer type* ($\$CT$) where the cyber incident occurs, the *data types* ($\$DT$) shared with the US Federal government, the *risk likelihood* ($\$RL$) of a privacy violation, the *privacy harm* ($\$PH$), and their combined effect on the employee’s *willingness to share* ($\$WTS$) their data with the U.S. Federal government (see Table 6 for the factor levels). We surveyed these factors in a single context—sharing cybersecurity incident data with the government—while varying the computer type affected, risk likelihood and the data type. Our sample size for this survey was 80 participants. With 80 responses, we achieved 97% actual power, calculated using G*Power [Faul et al. 2007].

Table 6. Vignette Factors and their Levels for Study 3 and Study 4

Factors	Factor Levels	
Computer Type ($\$CT$) Between-subjects Factor	personal smart phone	
	workplace computer	
Risk Likelihood ($\$RL$) Between-subjects Factor	only one person in your family	
	only one person in your workplace	
	only one person in your city	
	only one person in your state	
	only one person in your country	
Privacy Harm ($\$PH$) (Fixed value)	a privacy violation due to government surveillance	
Data Type ($\$DT$) Within-subjects Factor	Group 1	
	age range	sensor data
	usernames	network information
	passwords	IP address & domain names
	device information	packet data
	device ID	MAC address
	UDID / IMEI	
	Group 2	
	age range	registry information
	OS information	running processes
	OS type & version	application information
	memory data	application session data
	temporary files	
	Group 3	
	age range	contact information
emails	keyword searches	
chat history	keylogging data	
browser history	video & image files	
websites visited		

The surveys for Studies 3 and 4 were designed using the framework shown in Figure 1 and they consist of two main parts: factorial vignettes with the independent factors, and post-test questions about demographics, including their gender, age range, education level, household income, and ethnicity. In addition, in Study 3 we asked participants about their immediate family size, workplace size, and the zip code they reside in (see Section 4.6 for demographic results). All pages with factorial vignettes and the question order are randomized to mitigate any ordering effects.

Equation 3 below is our main additive regression model for Study 3 with a random intercept grouped by participant's unique ID, the independent between-subject measures $\$CT$, which is the computer type, $\$RL$, which is the likelihood of a privacy violation, and $\$DT$, which is the data type (see Table 6). The additive model is a formula that defines the dependent variable $\$WtS$, *willingness to share*, in terms of the intercept α and a series of components, which are the independent variables. Each component is multiplied by a coefficient (β) that represents the weight of that variable in the formula. The formula in Equation 3 is simplified as it excludes the dummy (0/1) variable coding for the reader's convenience.

$$\$WtS = \alpha + \beta_C \$CT + \beta_R \$RL + \beta_D \$DT + \epsilon \quad (3)$$

To measure the effect of different factors and their levels on $\$WtS$, we establish the baseline level for the factor $\$CT$ to be *workplace computer*, $\$RL$ to be *only one person in your family* who experiences the privacy violation, and we set the factor $\$DT$ to *age range*. The intercept (α) is the value of the dependent variable, $\$WtS$, when the independent variables, $\$CT$, $\$RL$, and $\$DT$ take their baseline values.

Study 4 has all the independent variables used for Study 3, in addition to the independent variable for benefits, which are the *data purpose* ($\$DP$) that provide benefits to society (see Table 7 for the factor levels).

Table 7. Independent Variable Benefit and its Levels for Study 4

Factors	Factor Levels
Data Purpose (\$DP)	investigating intellectual property and trade secrets
	investigating economic harm, fraud or identity theft
Within-subjects Factor	investigating imminent threat of death or harm to an individual, including children
	investigating terrorism

In this study, each participant sees and judges a total of three factorial vignettes, one for each data type group. Figure 8 shows the vignette template for Study 4. Each factor in the vignette is replaced by one level from Tables 6 and 7. For Study 3, the variable $\$DP$ was removed from the vignettes, keeping the rest of the vignette the same. The independent variables $\$CT$ and $\$RL$ are between-subject factors, thus participants only see one level of these two factors, and the variables $\$DT$, $\$DP$, and $\$PH$ are within-subject factors, so participants see all combinations of these factors. In the vignette survey design, the $\$DT$ levels were evenly divided into three groups, thus, each participant sees and responds to $3 \times 4 \times 1 = 12$ vignettes combinations. The allocation of $\$DT$ levels to groups was made to ensure that the data types that were technically related are shown together. The data type age range was included in each group as a non-sensitive data type aimed at balancing the $\$WtS$ scale utilization.

Please rate your willingness to share your information below with the Federal government for the purpose of \$DP, given the following risk.

Risk: In the last 6 months, while using this website, only \$RL experienced a privacy violation due to \$PH.

When choosing your rating for the information types below, consider the \$CT, purpose and the risk, above.

	Extremely Willing	Very Willing	Willing	Somewhat Willing	Somewhat Unwilling	...
\$DT	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	

Fig. 8. Template used for vignette generation for Study 3 and Study 4 (fields with \$ sign are replaced with values selected from Table 6 and 7)

The 12 vignette combinations in Study 4 are presented in group-order. First, participants see four vignettes (one for each level of \$DP) for each \$DT group 1-3 in succession, where only the \$DP level changes across each group. Prior to responding to each group of four vignettes, participants watch an approximately 60 second video that illustrates the meaning of each data type, because some data types are technical terms that lay people may not be familiar with, such as running processes or registry information. The \$DT levels were assigned to each group to fit these narratives, thus the groups had to be related in a technical manner. In addition, the videos offer a break between each group of four vignettes. For the transcripts used to narrate these videos please refer to Appendix C.

Before the vignettes, we present a pre-test that asks participants to rank order and score the data purposes based on their benefit to society. Overall, the majority ranked the data purposes as follows: investigating imminent threat of death (68.8%) was most beneficial, followed by terrorism (60.0%), followed by economic harm (63.8%), and ending with intellectual property (68.8%) as least beneficial.

Equation 4 is our main additive regression model for Study 4, with a random intercept grouped by participant's unique ID, the independent between-subject measures *computer type* \$CT and *risk likelihood* \$RL, and the independent within-subject measure *data purpose* \$DP and *data type* \$DT. The additive model is a formula that defines the dependent variable *willingness to share* \$WtS in terms of the intercept α and a series of components, which are the independent variables. Each component is multiplied by a coefficient β that represents the weight of that variable in the formula. The formula in Equation 4 is simplified as it excludes the dummy (0/1) variable coding for the reader's convenience.

$$WtS = \alpha + \beta_c \$CT + \beta_R \$RL + \beta_P \$DP + \beta_D \$DT + \epsilon \quad (4)$$

The effect of different factors and their levels on \$WtS was measured by establishing a baseline level for the factor \$CT to be *workplace computer*, \$RL to be *only one person in your family* who experiences the privacy violation, \$DP to be *investigating intellectual property and trade secrets* and we set the factor \$DT to age range. The intercept (α) is the value of the dependent variable, \$WtS, when the independent variables (\$CT, \$RL, \$DP and \$DT) take their baseline values. The perceived privacy risk is measured by the estimated willingness to share \$WtS on a scale of 1 to 8, wherein 1=*Extremely Unwilling*, 4=*Unwilling*, 5=*Willing*, and 8=*Extremely Willing*, and which estimates an average person's acceptance of the risk.

We now describe our results for the factor data type from Study 3 and Study 4. The results for the data purpose factor are presented in Table 9, and those for the computer type factor are presented in Section 4.4. The levels of the factor risk likelihood did not have any statistically significant differences for either of the Studies 3 nor 4, and hence are reported in the Appendix B, Table B.2. In Study 3, we found a significant contribution of the three independent factors for predicting the \$WtS ($\chi^2(29)=552.62, p<0.000$) over the null model, which did not have any of the independent variables. In Table 8, we present the \$WtS values for select levels of the independent variable \$DT (*data type*) from Studies 3 and 4 (for values of the other data types, see Appendix C). The value of the intercept for Study 3 is 5.616, and for Study 4 the intercept is 6.340. The intercept value in the equation represents the value of the dependent variable when all the factors are set to the baseline level,

which in our case are: only 1 person in your family (risk likelihood), age range (data type), data purpose (intellectual property), workplace computer (computer type) and intellectual property (data purpose). The $\$WtS$ value for each data type in Table 8 is computed by adding the value of the coefficient for each data type to the intercept value. For example, in Study 4, the value of the intercept is 6.340 and the coefficient for the “application session” data type is -1.072, and therefore the value of $\$WtS$ is $6.340 + (-1.072) = 5.268$. The values in Table 8 therefore represent the predicted value of the dependent variable $\$WtS$ for the respective data type, when the other independent factors take their baseline values. As described above, the scenarios and factors in Studies 3 and 4 were identical, with the exception of the factor data purpose (benefits) which appears in Study 4 and not in Study 3. We therefore compute the difference in $\$WtS$ values for the two studies, based on the assumption that since all the other factors and their levels were identical, the change in the values of $\$WtS$ were because of the presence of the benefits.

Table 8. Multilevel Modeling Results for Data Types in Study 3 and Study 4

Term	$\$WtS$		$\Delta \$WtS$ from Study 3 to 4
	Study 3 w/o Benefits	Study 4 w/ Benefits	
Application Session	3.941	5.268	1.327
Browser History	3.466	4.649	1.183
Chat History	2.879	4.378	1.499
Device ID	4.666	5.984	1.318
IPAddresses	4.804	6.093	1.290
OS Type and Version	5.904	6.603	0.699
Registry Information	4.279	5.371	1.093
Usernames and Password	2.591	4.149	1.558
Websites Visited	3.679	4.871	1.193

In Figure 9, we show the $\$WtS$ for a subset of data types from Studies 3 and 4. The $\$WtS$ values are computed by substituting the coefficients from Table 9 for the different levels of the independent variables from each study, respectively, in Equations 3 and 4. The results in Figure 9 have been computed using the following levels of the other independent variables – risk likelihood “only one person in your family,” computer type “workplace PC,” data type “age range,” and for Study 4 the data purpose “investigating intellectual property and trade secrets.”




Fig. 9. Willingness to share of different data types when surveyed with and without benefits

From Studies 3 and 4, we observe that the $\$WtS$ is significantly different across different levels of $\$DT$. The results in Table 8 (and Appendix C) show that $\$WtS$ increases by an average of 1.3 units for factorial vignettes which have explicit benefits, as compared to factorial vignettes which do not show benefits as an independent factor. In Study 4, we found a significant contribution of the four independent factors ($\$CT$, $\$RL$, $\$DT$ and $\$DP$) for predicting the $\$WtS$ ($\chi^2(32)=2415.1, p<0.000$) over the null model, which did not have any of the independent variables. In Table 9 we present the coefficients and the corresponding standard error for the intercept and the different levels of the independent factor $\$DP$ for the regression Equation 4. The baseline level for $\$DP$ in Study 4 was chosen to be “investigating intellectual property and trade secrets.”

Table 9. Multilevel Modeling Results for Benefits in Study 4

Term	Coeff.	Stand. Error
Intercept (family + workplace PC + intellectual)	6.340***	0.421
Data Purpose – economic harm	0.136**	0.044
Data Purpose – terrorism	0.795***	0.044
Data Purpose – imminent death	1.153***	0.044

* $p \leq 0.05$ ** $p \leq 0.01$ *** $p \leq 0.001$

The results in Table 9 show that $\$WtS$ is significantly different and increasing for increasing levels of benefits (from the pre-test) of $\$DP$, as compared to the baseline level “*investigating intellectual property and trade secrets.*” For example, the $\$WtS$ increases by 1.153 over the baseline level for the $\$DP$ level “*investigating imminent threat of death or harm to an individual, including children.*” This $\$DP$ level was also ranked the most beneficial to society in the benefits ranking survey. We therefore conclude that users are relatively more willing to share their information for the purposes that they perceive to be more beneficial to society.

In Study 3, the fixed variables explain 14% of the variance in the dependent variable willingness to share (R-squared marginal=0.14), and both the fixed and random variables explain 60% of the variance (R-squared conditional=0.60). The sum of squares for Study 3 are as follows: risk likelihood (7.18), device type (0.43), data type (1333.83), and the total sum of squares is 1341.44. Therefore, the dependent variable data type

explains 0.99 (1333.83/1341.44) of the variance explained by the fixed effects in the willingness to share, whereas the effect of risk likelihood and device type on the willingness to share is negligible.

In Study 4, the fixed variables explain 16% of the variance in the dependent variable willingness to share (R-squared marginal=0.16), and both the fixed and random variables explain 61% of the variance (R-squared conditional=0.61). The sum of squares for Study 4 are as follows: risk likelihood (3.1), device type (4.6), data purpose (1789.1), data type (3770.3), and the total sum of squares is 5567.1. Therefore, data purpose explains 0.32 (1789.1/5567.1) of the variance explained by the fixed effects in the willingness to share, and data type explains 0.68 of the variance explained by the fixed effects in the willingness to share, whereas the effect of risk likelihood and device type on willingness to share is negligible. These results show that most of the variance explained by the fixed effects is due to the variable data type in Study 3, whereas in Study 4 two-thirds of the variance explained by fixed effects is due to the data type variable, and the rest is explained by the data purpose variable.

4.4 Computer Type and Perceived Privacy Risk

The research question RQ5 concerns whether privacy risk varies with a user's computer setting which we measure in Studies 3 and 4. In order to answer this research question, we measured the effects of the independent factor computer type $\$CT$, which had two levels "workplace computer" and "personal smart phone" in Studies 3 and 4. The analysis of Study 3 and Study 4 data using Equations 3 and 4 did not produce any significant effect of the independent variable $\$CT$ on our dependent variable of interest $\$WTS$. In other words, participants did not appear to perceive any differences between the data stored on their workplace computer versus the data stored on their personal smart phones.

4.5 Correlating Discomfort, Identifiability, and Personal Nature of an Information to the Perceived Privacy Risk

In the human-computer interaction literature, researchers often refer to level of discomfort that users experience when participating in privacy risky activities [Olson et al. 2005 and Wang et al. 2011]. In order to answer research question RQ6 which concerns whether privacy risk correlates with three predictors of interest, discomfort, identifiability and the personal nature of data types, we conducted three additional surveys. The survey results were compared to the risk measures obtained from Study 4. The surveys were conducted on Amazon Mechanical Turk (AMT) with 50 people responding to each survey. The survey presented a set of instructions, and participants were asked to select one level of one predictor of interest for each data type. In each of these surveys, the instructions were followed by five data types for each Human Intelligence Task (HIT) in AMT, and the participants had to respond to each data type based on the instructions shown. Each participant could do one or more of these HITs, across one or more of these surveys.

For the discomfort survey, the following instructions were shown to the participants: "For each highlighted phrase, please select the level of comfort you would experience sharing that information with a website." The options for the discomfort survey and their corresponding numeric value we used for the analysis were: *Extremely Uncomfortable* (0), *Uncomfortable* (2), *Somewhat Uncomfortable* (4), *Somewhat Comfortable* (5), *Comfortable* (7), *Extremely Comfortable* (9).

For the identifiability survey, the following instructions were shown to the participants: "For each highlighted phrase, please select the level of identifiability for the information type. Information that always uniquely identifies a single individual is extremely identifiable, whereas information that never uniquely identifies an individual is extremely anonymous. Consider the information by itself and not in combination with any other information." The options for the identifiability survey and their corresponding numeric value we used for the analysis are: *Extremely Identifiable* (9), *Identifiable* (7), *Somewhat Identifiable* (5), *Somewhat Anonymous* (4), *Anonymous* (2), *Extremely Anonymous* (0).

To measure the extent to which a data type was considered personal, we surveyed participants and showed them the following instructions: "For each highlighted phrase, please choose how personal you believe that information is, when sharing that information with a website." The options for the survey and their

corresponding numeric value that we used for the analysis are: *Extremely Personal* (9), *Personal* (7), *Somewhat Personal* (5), *Somewhat Non-personal* (4), *Non-personal* (2) *Extremely Non-personal* (0).

We analyzed the survey data from these three studies along with the privacy risk measures from Study 4 using simple linear regression, and we also compute the R-squared statistic to measure the proportion of the variance in the dependent variable that is predicted by the independent variable(s). The simple linear regression Equation 5 shows the effect of any of the predictors ($\$P$) of data types on the privacy risk ($\$PR$). In this equation, $\$P$ refers to one of the following: the *discomfort*, *identifiability*, and *personal nature* of the data type, and $\$PR$ is calculated by inverting the respective values of $\$WtS$ of the data types from Study 4, $\$PR=9-\WtS , since we have a total of eight scale options for our risk surveys. This analysis helps us understand if any of our three predictors, discomfort, identifiability, and personal nature of the data type, can predict the associate perceived privacy risk. Notably, surveys for these three predictors are easier to design, thus if correlation is high, the predictors could be used as substitutes for measure perceived privacy risk.

$$\$PR = \alpha + \beta_c \$P \quad (5)$$

Table 10 presents the intercept estimate, the corresponding intercept standard error, the coefficient estimate for the predictor, followed by the standard error of the predictor's coefficient and the R-squared value of the model. In this analysis, our objective was to fit a model that explains the maximum variations in the privacy risk, we therefore experimented with linear regressions with the predictor value, and the square of the predictor value, and calculated the R-squared values for each model.

Table 10. Linear Regression Results for Discomfort, Identifiability, Personal Nature and Privacy Risk

Predictor	Intercept Estimate	Intercept Stand. Error	Predictor Estimate	Predictor Stand. Error	R-squared
Identifiability	2.344***	0.428	0.319**	0.096	0.314
Identifiability ² (quadratic model)	1.400	1.258	-0.054	0.067	0.332
Discomfort	1.694***	0.250	0.421***	0.050	0.747
Discomfort ² (quadratic model)	2.44**	0.416	0.045*	0.021	0.789
Personal nature	0.783	0.435	0.507***	0.074	0.661
Personal nature ² (quadratic model)	0.290	1.816	-0.014	0.051	0.662

* $p \leq 0.05$ ** $p \leq 0.01$ *** $p \leq 0.001$

Based on the linear regression results shown in Table 10, we observe that discomfort was found to be strongly correlated to privacy risk, since 75% of the variations in privacy risk could be explained by the variations in discomfort associated with a data type. Analyzing the quadratic model (Discomfort²), we observe that the square of discomfort values explains 79% of the variations in privacy risk. In the quadratic model, the coefficient of the linear term “discomfort” was not found to be significant, whereas the coefficients of the intercept and the quadratic term were significant. This means that the privacy risk value is directly proportional to the square of the discomfort value. In addition, only 31% of the variations in the privacy risk were explained by the variations in identifiability, even though the linear regression model with identifiability as the independent variable and privacy risk as the dependent variable is significant. The predictor personal nature of the data type explains 66% of the variations in privacy risk, the linear model was the best fit, and the quadratic model was not significant.

Figures 10a-10c break down the relative differences among the estimated value of the privacy risk for each data type from Study 4, and the average value from the survey ratings for each of the three predictors: discomfort, identifiability, and personal nature. The $\$PR$ or *unwillingness to share* for each data type is equal to $uWtS = (9 - \$WtS)$, since we had an eight-point scale for measuring $\$WtS$ and to keep the scale directions the same direction across the predictors and privacy risk. The privacy risk score for the information type is calculated for the baseline levels: *workplace PC* ($\$CT$), *only one person in your family* ($\$RL$), *age range* ($\DT)

and investigating intellectual property and trade secrets (\$DP). For details about the values in the Figures 10a-10c, see Appendix D.

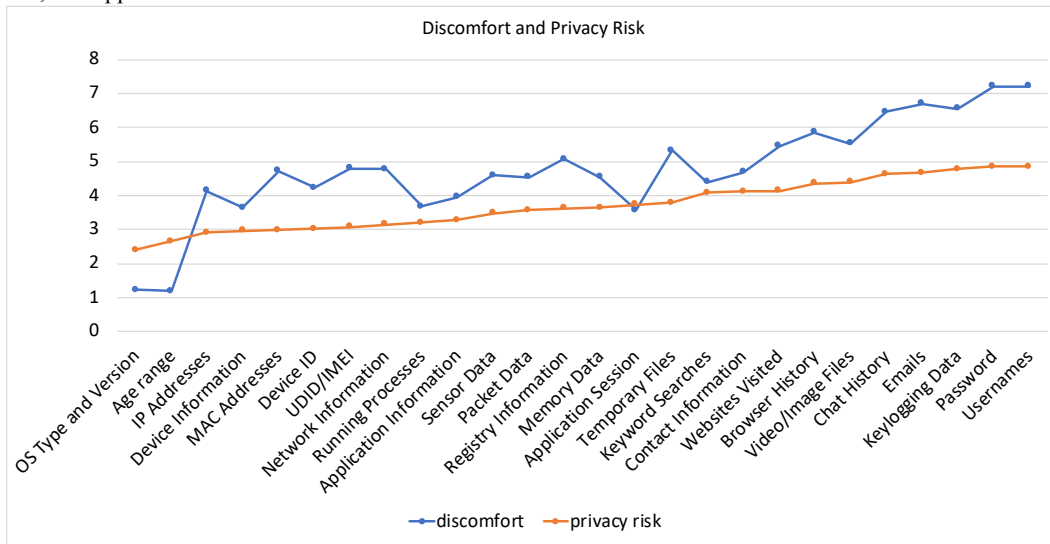


Fig. 10a. Discomfort and Privacy risk

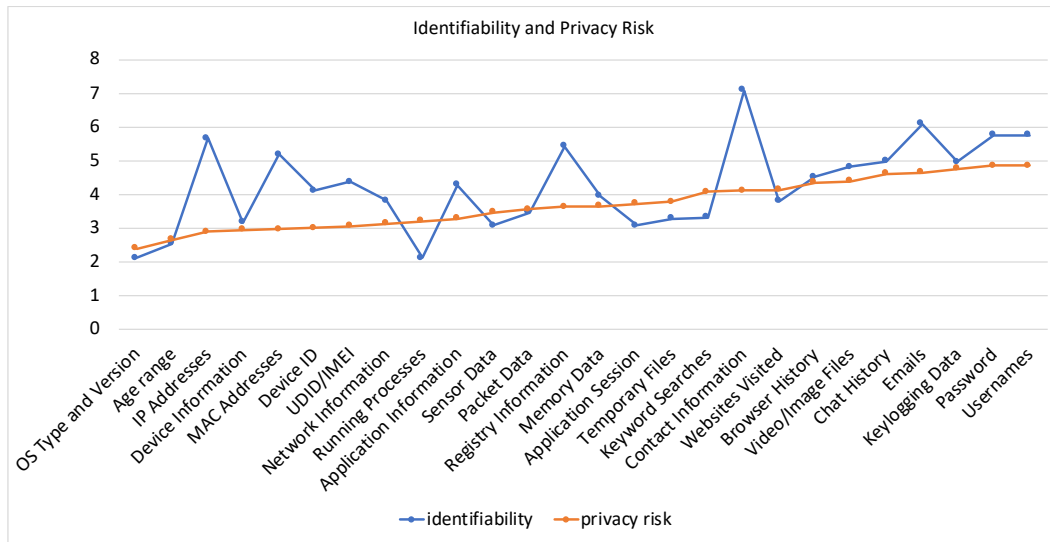


Fig. 10b. Identifiability and Privacy risk




Fig. 10c. Personal Nature of Data Types and Privacy risk

In Figures 10a-10c, we see evidence that privacy risk, discomfort and personal nature follow a general trend. However, we can also note discrepancies, such as unique device identifier (UDID) / international manufacturer’s equipment identifier (IMEI), which respondents on average reported as relatively less risky and less uncomfortable, despite finding this data type as relatively more personal and identifiable. By comparison, IP address is viewed as more identifiable than UDID, but is relatively less risky, uncomfortable, identifiable and personal in nature.

4.6 Estimating the Effect of Demographics on Perceived Privacy Risk

We investigated whether privacy risk is perceived differently by people with different demographic characteristics, to answer research question RQ7. Demographic factors were considered as independent variables and added to the multilevel analysis for each of four factorial vignette studies (Study 1- Study 4). In Table 11 below, we show the demographic factors we studied using post-tests and their corresponding factor levels separated by commas.

Table 11. Demographic Information Factor Levels

Demographic Factor	Factor Levels
Gender	Male, Female
Ethnicity	Asian/Pacific Islander, Black or African American, Hispanic or Latino, Native American or American Indian, Other, White
Education level	Less than a high school diploma, High school graduate no college, Some college or associate degree, Bachelor's degree, Masters degree, Doctoral degree,

Household income	Less than \$25,000, \$25,000 to \$34,999, \$35,000 to \$49,999, \$50,000 to \$74,999, \$75,000 to \$99,999, \$100,000 to \$124,999, \$125,000 to \$149,999, \$150,000 or more,
Age range	18-24 years, 25-29 years, 30-34 years, 35-39 years, 40-44 years, 45-49 years, 50-54 years, 55-59 years, 60 years or older

Table 12 presents the descriptive statistics for demographic information for all four studies.

Table 12. Demographic Information Descriptive Statistics

Demographic		Study 1	Study 2	Study 3	Study 4
No. of Participants		210	200	80	80
Gender	Female	47.14%	48.50%	38.75%	48.75%
	Male	52.86%	51.50%	61.25%	51.25%
Ethnicity	Asian/Pacific Islander	15.71%	13.50%	76.00%	7.50%
	Black or African American	3.81%	8.50%	6.25%	7.50%
	Hispanic or Latino	4.76%	5.50%	10.00%	5.00%
	Native American or American Indian	1.43%	2.00%	6.25%	0.00%
	Other	0.95%	0.50%	1.25%	0.00%
	White	73.33%	70.00%	76.25%	80.00%
Education level	Bachelor's degree	40.95%	44.50%	40.00%	36.25%
	Doctoral degree	0.48%	1.50%	1.25%	5.00%
	High school graduate, no college	14.76%	9.50%	26.25%	15.00%
	Less than a high school diploma	0%	0%	1.25%	0%
	Masters degree	12.86%	8.00%	1.25%	11.25%
	Some college or associate degree	30.95%	36.50%	30%	32.50%
Household income	\$100,000 to \$124,999	6.67%	4.50%	6.25%	5.00%
	\$125,000 to \$149,999	1.90%	1.50%	1.25%	1.25%
	\$150,000 or more	1.43%	1.00%	1.25%	1.25%
	\$25,000 to \$34,999	11.43%	16.00%	22.50%	11.25%
	\$35,000 to \$49,999	19.05%	20.00%	13.75%	13.75%
	\$50,000 to \$74,999	28.57%	22.50%	20.00%	31.25%
	\$75,000 to \$99,999	11.43%	11.00%	15.00%	16.25%
	Less than \$25,000	19.52%	23.50%	20.00%	20.00%
Age range	18-24	3.33%	2.50%	7.50%	5.00%
	25-29	20.48%	20.50%	38.75%	18.75%
	30-34	23.81%	28.50%	28.75%	25.00%
	35-39	18.57%	18.00%	7.50%	17.50%
	40-44	10.48%	6.50%	10.00%	5.00%
	45-49	7.62%	7.00%	3.75%	7.50%
	50-54	5.24%	7.00%	2.50%	8.75%
	55-59	6.19%	6.00%	1.25%	8.75%
	60 or older	4.29%	4.00%	0.00%	3.75%

We tested for an effect of the demographic factors – gender, ethnicity, education level and household income as independent variables across Studies 1 through 4. The effects were calculated by adding the demographic factors as independent variables to the multilevel model for each of the studies. In Table 13 we show the demographic factor levels which were significant, along with their coefficient and the corresponding standard error; Study 4 had no demographic factor levels with significant differences. The baseline levels for these factors were: *18-24 years* (age group), *White* (ethnicity) and *Less than a high school diploma* (education level).

Table 13. Demographic Information

Study	Demographic Factor Level	Coeff.	Stand. Error
Study1	Age range: 60 or older	1.652*	0.735
	Ethnicity: Native American or American Indian	2.124*	0.857
Study 2	Age group: 25-29	2.077**	0.694
	Age group: 30-34	1.949**	0.669
	Age group: 35-39	1.845**	0.674
	Age group: 40-44	2.604**	0.751
	Age group: 45-49	2.089**	0.713
	Age group: 60 or older	1.722*	0.799
	Education: Some college or associate degree	0.455*	0.228
	Income range: \$50,000 to \$74,999	1.099*	0.504
Study 3	Age group: 25-29	-1.602*	0.673
	Age group: 30-34	-1.520*	0.725
	Age group: 35-39	-1.888*	0.893
	Age group: 40-44	-2.529**	0.803
	Age group: 50-54	-2.587*	1.207
	Education Level: Doctoral degree	2.915*	1.462

*p≤.05 **p≤.01 ***p≤.001

In Study 1, we observed that the $\$WtS$ of the participants who identified their ethnicity as “Native American or American Indian” were more willing to share by 2.12 units than participants who identified their ethnicity as “White.” In Study 2, participants who reported their age as greater than or equal to 25 were more willing to provide their information as compared to participants who were between the ages of 18 and 25, except for those between the age of 50-59. Older participants who reported their age from 40-54 years in Study 3 were less willing to share as compared to younger participants reported their age from 18-24. Participants who reported having a doctorate degree were significantly more willing to share than participants reported having less than a high school diploma in Study 3. The differences in willingness to share across different studies can be attributed to the different scenarios. For instance, on average, the participants who were 25 or older were more willing to share their information than participants who were younger given that the scenario topic was shopping, and the benefits of shopping were listed. On the other, participants older than 25 years of age perceived more risk than their younger counterparts when they were asked to share their information with the government, and no benefits were mentioned.

In Study 3, we also collected the participant’s immediate family size, workplace size, and the zip code they reside in. We categorized family size as *small family* (family size <3 members), *medium family* (family size is 3-5 members) and *large family* (family size >5 members); workplace as *small workplace* (number of employees <100), *medium workplace* (number of employees 100-999), and *large workplace* (number of employees >1000); the size of the geographic area they reside in as *rural* (population <2,501), *urban cluster* (population

2,501-49,999) and *urban area* (population >49,999), which are based on the guidance provided by the U.S. Census Bureau; and state size as *small state* (population <1,000,000), *medium state* (population 1,000,000-9,999,999) and *large state* (population >9,999,999). The populations for the zip code area and state were obtained from the US Census Bureau.

We added the factors family size, workplace size, geographic area size and state size as independent factors to our model described in Equation 3 from Section 4.3. In Table 14 below, we show the coefficients and the corresponding standard error for all the terms that were found to be statistically significant.

Table 14. Multilevel modeling results for family, workplace, area and state size for Study 3

Term	Coeff.	Stand. Error
Small workplace	0.984*	0.509
Medium workplace	1.508*	0.631
Medium state	2.937**	1.068
Large state	3.211**	1.077

* $p \leq .05$ ** $p \leq .01$ *** $p \leq .001$

As seen in Table 14, we observe that workplace size and state population size affect the $\$WtS$. The participants who work in a smaller and medium-sized workplace were more willing to share their information as compared to participants who worked in a larger workplace. In contrast, participants who live in a medium and larger-sized state were more willing to share their information as compared to participants who live in a smaller state.

5 THREATS TO VALIDITY

We now discuss threats to validity.

5.1 Construct Validity

Construct validity concerns whether what we measure is actually the construct of interest [Creswell 2014]. For the risk likelihood factor, we focus-grouped the risk likelihood levels and we conducted multiple rounds of pilot testing. The risk likelihood levels ($\$RL$) were further motivated by an established theory of psychological distance, which has been validated in multiple studies [Wakslak and Trope 2009]. The theory predicts that spatial and social distance strongly correlate with perceived event likelihood. While we are measuring perceived risk, similar to Fischhoff [Fischhoff et al. 1978], we assume that a person's willingness to disclose corresponds to their acceptance of the risk; this assumption was used in other study designs by Acquisti and Kobsa to measure privacy-related risk [Acquisti and Grossklags 2013, Knijnenburg and Kobsa 2014]. However, Knight argues that subjective measures are based on partial knowledge and could therefore measure the uncertainty of the situation rather than the perceived risk [Knight 1921].

We did not study the extent to which ambiguity aversion explains the risk: ambiguity aversion, also called Knightian uncertainty [Knight 1921], occurs when a person is inclined to choose familiar risks over unfamiliar risks [Ellsberg 1961]. Personality tests, such as the Ellsberg paradox [Ellsberg 1961], can be used to measure intra-personal ambiguity aversion as a pre-test prior to completing a privacy risk survey.

Finally, the semantic scale anchor labels used for $\$WtS$ in the factorial vignettes could be interpreted differently by participants [Clark and Watson 1995]. To address this threat, we studied these factors as between- and within-subject factors, so that all participants respond to all levels of these variables. During multilevel modeling, we account for subject-to-subject variability using the random effect variable $\$PID$ in the respective regression equation. Another way to address this threat could be to conduct surveys to calibrate the scale options for the dependent variable $\$WtS$ [Furr 2011], which we leave to future work.

5.2 Internal Validity

Internal validity concerns whether the study procedures limit drawing correct inferences from the data [Creswell 2014]. One threat to internal validity is that our survey participants may not understand the meaning of the different data types for which they needed to rate $\$WtS$. In order to mitigate this, we provided instructional videos illustrating the definition of each data type accompanied each group of survey questions for Studies 3 and 4. This included closed captions for the hearing impaired. Additionally, each data type presented with a semantic scale (as seen in Figure 8) had a definition that was displayed when the participant hovered their mouse over the data type. Even though we provided the instructional videos, and definitions we cannot guarantee that the participants watched the videos or read the definitions. However, we do not use their knowledge of the data types as a factor in our analysis. Some of the participants might have prior knowledge about the data types, and some participants may pay attention to the videos and definitions. Our intention is to increase the saliency of the data types before the participants see the data type in the factorial vignette.

We also randomized the order of the vignettes in our surveys. In addition to randomizing the questions order in the survey, and the sub-questions within each question. Despite using randomized question ordering, each study had two or more within-subjects factors, and the choice of the corresponding factor levels when shown on a single page could influence participants' responses. To remove this ordering effect, we created multiple variations of our surveys, where in each variation shows a different within-subjects variable on the same page. Each participant saw a randomly chosen variation of the survey, with each variation being shown to equal number of participants. The data for each survey with all the different variations was analyzed together. This technique should reduce ordering effects that could have been present.

Another factor that affects the internal validity is the fatigue experienced by the participants if the factorial vignette surveys are long [Wallander 2009, Shadish et al. 2002]. In this paper, we report the effect of different factors, such as risk likelihood, data types, data purpose, privacy harms and computer type on the perceived privacy risk, which is in addition to demographic factors and the three predictors of risk: discomfort, identifiability and personal nature of data types. If we designed a single within-subjects study to measure the effect of all these factors and their levels, each participant would have had to respond to more than 7000 vignette questions. We therefore designed multiple studies, each with one or more of these factors, and varying in which studies the factors would appear as within- and between-subjects factors to ensure that the participants experience minimum fatigue. We also randomized the order of the questions and pages in the surveys to ensure that fatigue experienced by the participants does not consistently contribute to the results of particular conditions.

We combined results from Study 4, and the four additional surveys to measure the discomfort, identifiability and personal nature of data types to compute the correlations between privacy risk and these three properties of data types. Our motivation for this measurement is that, while risk involves a number of factors that must be balanced to reduce fatigue, these other measures are relatively easy to measure, independently. The intent of correlating discomfort, identifiability and personal nature of data types with risk was to measure the extent to which any of these factors could explain the privacy risk. If discomfort correlates with the risk, we could substitute risk with a simpler measure, such as discomfort. Overall, these studies were conducted at different times, and participant responses could have been affected by privacy incidents in the news. Alternatively, the studies drew from the same population on Amazon Mechanical Turk. Secondly, our intention was to measure the correlation between the aggregate levels of two variables, which could be measured at different instances.

5.3 External Validity

External validity refers to the extent to which we can generalize the results to other situations [Creswell 2014].

In the privacy risk survey for Study 4, 42.3% of participants reported storing personal data on their workplace computer. This frequency may be higher or lower depending on the sectorial culture and company policies influencing employee behavior, which can affect the level of perceived risk in the ideal population. People who store less personal data on their workplace computer may report lower perceived privacy risk, but

we found no statistical significance to this effect ($p = 0.135$, compared to the null model). In fact, it is possible that people view their workplace computer as storing as much personal data as their smart phone.

Our target population is the average U.S. Internet user. For all four studies, we recruited participants from Amazon Mechanical Turk (AMT) who have a 95% approval rating or higher, and between 1000-5000 HITs completed. Demographically, our participant population deviates from other measures of U.S. Internet users. We recruited less reported Asian, Black and Hispanic participants, and more White participants than were found in 2014 Census data and the 2015 PEW Internet and American Life Survey of Internet users [Perrin and Duggan 2015]. This sample may skew privacy risk perceptions measured in our study that are influenced by race. However, multiple replication studies have shown that the data obtained from AMT is similar to the data obtained from a physical lab [Horton et al. 2011, Paolacci et al. 2010] and that the self-reported information by AMT workers is reliable [Rand 2011]. In addition, it has also been shown that AMT workers are more diverse and representative of the US population as compared to other convenience samples, such as college students [Buhrmester et al. 2011].

In this paper, we evaluated six factors that could affect the user's perception of privacy risk, including risk likelihood, type data, privacy harm, data benefits, computer type and demographic factors. Other factors such as a user's privacy experiences, privacy awareness, and the culture climate can also effects user's perceptions about privacy, which we have not evaluated in this paper [Smith et al. 2011, Dinev et al. 2015].

6 DISCUSSION

In this paper, we introduce an empirical framework that can be used to measure privacy risk as perceived by the user. We also conducted multiple studies using this framework, to measure the effect of different levels of factors such as data types, risk likelihood, privacy harms, data purposes and computer types on perceived privacy risk. We now discuss our results in the context of our research questions presented in Section 4, followed by some potential applications where the framework can be used. We reproduce the research questions below for convenience.

RQ1. *How can we manipulate, increase or decrease, an individual's perception of risk likelihood?*

Risk likelihood is known to be a factor that affects risk [Bauer 1960, Starr 1969, Stoneburner 2002]. In this paper, we introduce a new risk likelihood scale that is based on Construal level theory. Study 1 presented in Section 4.1 shows that risk likelihood can be manipulated using the participant's social and physical distance from the person experiencing the harm to the survey participant. We observed that as the social and physical distance increases from the participant (e.g., from only one person in your family or workplace to only one person in your state or country), then the participant's willingness to share increases. This effect was significant when the likelihood was a within-subjects factor, whereas the differences among the different levels were not significant when likelihood was a between-subjects factor. We did not observe the effect of likelihood to be multiplicative with the willingness to share or acceptance of the risk. In other words, we could not demonstrate that perceived privacy risk is the product of likelihood and impact as recommended by NIST.

RQ2. *How do different benefits affect the perception of privacy risk, in the presence of controlled harms?*

Studies 3 and 4 reported in Section 4.3 appear to confirm Slovic's argument that people increasingly accept technological risks when they are presented with enumerable benefits, and they perceive those benefits to be increasing [Slovic 2000]. On average, the *willingness to share* ($\$WtS$) increases by 1.3 units for factorial vignettes when benefits are explicit, as compared to when benefits are not present. In addition, we observe from Study 4 that $\$WtS$ increases as the perceived benefit of the data purpose increases. The data purpose *investigating imminent threat of death* was ranked by 68.8% of participants as the most beneficial purpose and in turn had the highest $\$WtS$, followed by *terrorism* (60.0%) with the second highest $\$WtS$, by *economic harm* (63.8%) with the third highest, and by *intellectual property* (68.8%) as the least beneficial with the lowest $\$WtS$ among all data purposes surveyed. In Study 4, data purposes accounted for 32% of the variance in willingness to share explained by fixed factors.

Our findings show that participants are more willing to share their information in the presence of societal benefits as compared to when they are not shown any benefits. In 2016, a PEW study found that, in the event of a high-profile national security event, such as the 9/11 terrorist attacks, a majority of American adults were in favor of increased security measures, such as requiring American citizens to carry a national ID at all times, and government monitoring of credit card purchases, emails and personal phone calls, among others, when thinking about the tradeoffs between security and privacy. In our study, we chose the societal benefit to be investigations of cyber security incidents, while the likelihood scale controls for perceived hypotheticality of the privacy harm. Our findings appear to corroborate the PEW results by showing that, despite an increased likelihood, participants perceive a lower risk as the benefits increase.

RQ3. *How do different harms affect the perception of privacy risk, in the presence of controlled benefits?*

New privacy harms have come into picture due to the ardent of new technology [Solove 2006] and therefore there is a need to determine how these harms affect privacy risk. In Study 2, we found statistically significant differences among the privacy harms we surveyed and found that the $\$WtS$ is the highest for the privacy harm *Induced Disclosure*, and the least for the privacy harm *Insecurity*. Unlike the benefits of technology that can easily change when considering different vignette scenarios, we believe a standard catalogue of rank-ordered privacy harms should be used to standardize privacy risk measurement and improve reliability.

RQ4. *How do different data types, in the presence and absence of benefits, affect the perception of privacy risk?*

Data types are known to be a contextual factor that effect user's perception of privacy [Nissenbaum 2004, 2009, Lederer 2003]. The four factorial vignette studies we conducted (Studies 1-4) confirm that data types affect privacy risk and show that 68-99% of the variance due to fixed effects in the willingness to share is explained by the data types. From Studies 3 and 4, we observed significant differences in $\$WtS$ for different data types. Participants were more willing to share information about who they are, e.g., IP address or device information, with the federal government than information about what they do online, e.g., chat history or browser history. Bhatia et al. surveyed 76 security experts to measure the usage of different data types for incidence reporting. In that study, experts reported a mean 8 years of experience in incident analysis and their job titles range from security analyst, security architect to director [Bhatia et al. 2016b]. *Usage* of a data type is the frequency with which the security analyst uses the data type for incident analysis. Figure 11 shows the tradeoff between the security analyst's usage on the y-axis and the public's willingness to share that data type on the x-axis. The willingness to share in the figure is from Study 4. See Appendix E for more details.

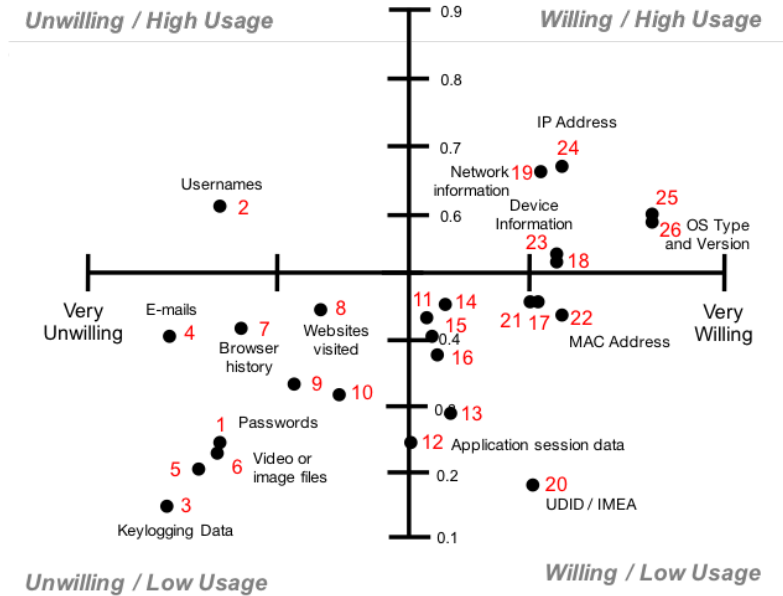


Fig. 11. Trade-off between incident data usage and users' willingness to share their data

The results show that a trade-off exists between data usage and privacy risk, in particular, that few types have high use and high risk (e.g., usernames) and many types have low risk independent of use. For low-risk data types, incident responders may feel comfortable sharing these data types using routine procedures for securing the data. These low-risk data protection procedures could include access control, and disk- and network-based encryption, for example, and security analysts who have access to the data may also be permitted to conduct their investigations by exploring the data, or combining the data with access to a broader set of data in the low-risk categories. For moderate- and high-risk data categories, however, security analysts may need to use data minimization techniques, such as redaction, to remove these data types before sharing. In addition, they may need to restrict access to those security analysts who are investigating specific incidents where the data is needed, and excluding such data from uncontrolled, exploratory data analysis. With respect to the benefits surveyed in Study 4, the willingness to share in Figure 11 shifts each data type to the right as the data purpose increases in benefit to society.

RQ5. Does privacy risk vary with whether a user is using workplace computer or their personal smart phone?

In our studies 3 and 4 (see Section 4.4), we did not find any significance difference among the factor levels for *computer type* (§CT), which were workplace computer and personal smartphone. This could mean that participants do not distinguish between how much personal information they store on their workplace computer versus their personal smartphones. Studies have also shown that employees use workplace computers for personal use, such as chatting, online shopping and checking emails, among other things [Everton et al. 2005, Mastrangel et al. 2006], which could explain why participants perceive the same risk for both types of devices. A follow-up survey could be conducted to understand what types of personal information people store on their work place computers, and how that kind of exposure affects their perception of workplace privacy risks.

RQ6. *Does discomfort, identifiability or the personal nature of data co-vary with or supplement privacy risk?*

Conducting privacy risk factorial vignette surveys with multiple factors is time and resource intensive. Therefore, we investigated whether single factors, such as discomfort, that are frequently referred to in HCI-related research [Olson et al. 2005 and Wang et al. 2011] can be measured in less time and correlated with the more expensive vignette survey.

We conducted three additional surveys (see Section 4.5) to measure the extent to which privacy risk co-varies with predictors, such as discomfort associated with a data type, the extent to which data types are perceived to identify a user, and the personal nature a data type. We found a strong correlation between the discomfort associated with a data type and its corresponding privacy risk, wherein 79% of the variation in privacy risk could be explained by the square of the discomfort value for the data type. In addition, we found a significant linear model with personal nature of the data type that could predict privacy risk, although, only 66% of the variation in privacy risk could be explained by personal nature of the data. We also found a significant linear model for identifiability of the data type and the corresponding privacy risk but found a very weak correlation of 31%. The survey designs to measure discomfort, identifiability and personal nature of data rely on a single semantic scale, which is easier to construct and deploy than a privacy risk survey based on factorial vignettes using pre- and post-tests. Because discomfort shows a strong correlation with privacy risk, a privacy analyst may use discomfort as a low-cost surrogate measure in cases where the harms are perceived to be minimal or the information is not especially personal. Moreover, whether data is identifiable does not predict the risk; however, combining identifiable data with behavioral data may increase the risk as evidenced by Study 4, where behavioral data is higher risk.

RQ7. *How do demographic factors influence the perception of privacy risk?*

In our studies, the demographic factors: age range, education, ethnicity, education level, and household income (see Section 4.6). In Studies 1 and 2, we found that ethnicity is a significant factor in predicting privacy risk. In Study 1, participants who reported their ethnicity as “Native Americans or Indians” were more willing to share their information, as compared to participants who reported “White” as their ethnicity. In addition, participants who reported their age as 60 or above, were more willing to provide their information, as compared to other age groups. In Study 2, with the exception of participants who reported their age as between 50 to 60 years, all the participants of age 25 or more, were more willing to share their information as compared to participants who reported their age as 24 or less. We found that participants in Study 3 who report their age as 25-39 were 1.5-1.8 units less likely to share information than younger participants age 18-24, and older participants age 40-54 were 2.5 units less likely to share than participants age 18-24 years. In Study 3, participants who had a doctorate degree perceived less risk in sharing their information as compared to participants who have just a high school diploma. To our knowledge, little work has been done to understand how privacy affects people with different ethnic backgrounds and in different age brackets, yet our results clearly indicate differences in perceived risk. Consistent with our findings about differences between different demographic populations, Hoofnagle et al. found that young adults in the age range of 18-24 years know less about privacy regulations, as compared to adult populations, and knowledge about privacy is the maximum for adults in the age range 40-55 years [Hoofnagle et al. 2010]. They also found that young adults of 18-24 years of age are less concerned about their privacy than people who are older than 25 years.

In addition, we found that participants who worked in comparatively large workplaces were less willing to share their information, as compared to participants who worked in small or medium-sized workplaces. In contrast, participants who live in a small-populated state were the least willing to share their information, followed by participants who live in a medium-populated state, and the participants who lived in a large city were the most willing to provide their information. These results are important because they suggest different dynamics may be at play in the workplace, versus in the city or state where people live. Small workplaces may benefit from greater degrees of trust and intimacy, whereas larger workplaces may make people feel more vulnerable or may subject people to more policies, procedures and workplace surveillance. Unlike large cities

and higher populated states, which may afford more anonymity due to larger population sizes, people are not anonymous in their workplaces and they rely on their jobs for financial security.

The framework we introduced in this paper to measure privacy risk can benefit system design, engineering, public policy, as well as developers, regulators and users. System developers, including designers, aim to build systems that users feel comfortable and safe using. In privacy, this includes accounting for Privacy by Design (PbD) [Hustinx 2010], wherein the user's privacy is considered throughout the development of the system. To perform PbD, however, developers need a systematic and scalable framework that can help them understand and measure the privacy risk that users experience while using a software system. Using the framework proposed herein, developers can frame their design choices in the context of privacy risk and measure how users perceive those risks, so that designs can be improved to reduce risk. For instance, if a company is deciding between using one of two practices about a given data type, where one practice calls for deleting the user data and not storing it for future purposes, whereas the other practice emphasizes on storing data for future purposes, the company can use the risk scores to inform this decision. If a particular information type or data practice is high risk, designers may introduce risk mitigations to affect the storage and use of that information. This may include limiting collection from the user or encrypting the information before it is stored or not storing it all.

In addition, regulators need a means to identify systems that could put a user's privacy at greater risk. The proposed framework can be used to score data practices for privacy risk, which may help regulators identify potentially high-risk systems. Privacy risk measurements can help privacy policy writers pay special attention to high-risk information types when they describe associated practices in their policies. Furthermore, known high-risk data practices and information can be used to introduce privacy nudges [Acquisti et al. 2017 and Wang et al. 2014] to users in real-time based on user demographics associated with high perceptions of risk. On the other hand, if data subjects misunderstand a technology and consequently perceive it as high risk, public policy could be used to explain the technology and provide additional guidance to reduce the risk in data handling.

7 CONCLUSION AND FUTURE WORK

In this paper, we introduce an empirical framework to measure perceived privacy risk. The framework can be used in multiple ways to predict how users perceive privacy risk for existing and emerging technologies. To illustrate the robustness of the framework, we report a series of surveys performed using the framework to measure the effect of six factors: risk likelihood, data type, privacy harm, data purpose, computer type and demographics on perceived risk. These factors can be used to baseline participants to measure differences among populations, IT services and privacy harms. In addition, we observed from our studies that risk likelihood is not a multiplicative factor when computing perceived privacy risk, which challenges the conventional knowledge that privacy risk is the product of likelihood and impact of a privacy harm [Brooks et al. 2017]. We also observed that participants perceive less privacy risk in the presence of benefits of sharing their information. Our studies showed that privacy risk is lower for induced disclosure harms when compared to surveillance and insecurity harms as defined in Solove's Taxonomy of Privacy. While performing these surveys we noted that within-subject surveys had greater power since all the participants saw all the levels of a given variable and were able to better distinguish between the levels.

The framework allows a privacy analyst, software engineer or user interface designer to tailor the survey to include other factors that may affect privacy risk. This includes understanding how new technologies, such as location-based services and facial recognition, can affect risk perceptions as compared to a more established technology, such as web browsing. Some researchers observed that people's decisions about privacy are inconsistent, and they might behave differently in actual versus hypothetical situations [Acquisti and Grossklags 2005]. This difference in behavior can be explained by how when participants see explicit benefits of an activity, then their perception of associated risk decreases, which in turn changes their behavior.

The framework evaluation shows that the semantics of privacy go beyond mathematical models of information flow and concern how users perceive privacy. In particular, under what circumstances does providing data become risky, and when do mitigations reduce the perceived risk, such as de-identification or avoiding re-purposing? Privacy-by-Design is a proposal that privacy should be considered throughout the entire

engineering lifecycle, and that privacy should be an integral part of system design [Hustinx 2010]. Privacy-by-Design has been endorsed by the European Union as part of their General Data Protection Regulation (GDPR). The GDPR recommends that the default for privacy settings should be kept very high and, in addition, the processing of user personal information should be limited to only when such processing complies with the regulation. However, designers cannot always anticipate a user's perception of privacy risk, nor should designers necessarily treat all personal data as high-risk. In traditional practice, designers may use their personal judgment or user personas to make design decisions related to user privacy, which may also underestimate the risk, e.g., when the designer estimates the risk of a person of a different ethnicity or age range. We believe the proposed empirical framework can be used to better inform design decisions made during software development by surveying prospective users. Understanding how privacy risk varies under different contextual factors can help designers allocate resources more carefully when collecting, sharing, and securing high and moderate risk data. For instance, data redaction techniques could be applied to high and moderate risk data before sharing it with third parties for secondary purposes perceived to be low-benefit to data subjects.

The proposed framework could also help privacy policy authors and regulators understand which data is most at risk, and to take measures to make sure that companies clearly describe their data practices related to those data types. This clarification would help users make more informed decisions about using services, in light of the data practices described in the service's privacy policy. Users could also make use of the predicted privacy risk for other users within their own demographic categories in a given context to decide whether they wish to proceed with a service. Our findings may also be used to augment the privacy personas proposed by Mugan et al. that users can choose to auto-configure their privacy settings [Mugan et al. 2011]. Understanding how individual users perceive privacy risk can inform how users might choose their settings, in turn, supporting broader adaptation of applications.

In addition, the proposed framework could also be used to support the standardization of the privacy impact assessments (PIA). The PIA is a decision tool that helps companies and government agencies detect and reduce privacy risks associated with information systems. This is done by disclosing within an organization or to the public what kinds of personally identifiable information is being collected and for what purpose, and how it will be secured and shared. This is accomplished by making sure that the software complies with the respective legal, regulatory and privacy requirements by identifying the associated risks and their consequences, and by assessing the safeguards in place for the privacy risks. The proposed risk framework could be used by risk analysts or developers to determine the privacy risk for the data types that are collected by companies or the government, and then based on the predicted perceived risk, appropriate privacy controls can be used as recommended by NIST 800-53, Appendix J. These controls are based on legislation and standards and serve to describe basic privacy requirements that can be used across the federal government and critical infrastructure.

An important topic for future work is the effect of combining data types on perceived privacy risk. For example, we did not study the extent to which combining data types of different risk levels can affect overall privacy risk. For instance, the data type *username* might be less risky as compared to the combination of *username* and *password*, or *username* with *home address*. Furthermore, privacy risk can vary with the recipient of the information (shopping websites or government entities), in addition to the benefit provided by sharing the data. This may be due to the level of trust that users place in the recipient regarding whether their information would be used appropriately and only for disclosed purposes. In future work, trust may be worth exploring as a factor that moderates perceived benefits and risks.

PRIOR PUBLICATIONS

We have published the factorial vignettes survey design concept to measure privacy risk [Bhatia et al. 2016a]. In addition, we have previously published preliminary results and analysis from Study 4 [Bhatia et al. 2016b]. In this paper, we use the data set for Study 4 from this previous study to conduct additional analysis such as the analysis to measure the effect of demographic factors on perceived privacy risk, comparison with scenarios with no perceived societal benefits, among other analysis. The remaining studies including the surveys to measure the effect of the independent variables such as *risk likelihood*, *data types*, *privacy harms* and their severity; absence of societal benefits during information sharing and the effect of demographic factors on user's

perceived privacy risk are new. In this paper, we also report our findings from the studies we conducted which were aimed at estimating the effect of predictor variables such as *discomfort*, *identifiability* and *personal nature* associated with information types, and how they correlate with the user's perception of risk, which are also previously unpublished.

ACKNOWLEDGMENTS

We thank Dr. Howard Seltman and Dr. Stephen Broomell for their statistics guidance, the participants of the 10th Annual Privacy Law Scholars Conference (2017) and the CMU Requirements Engineering Lab, for their helpful feedback. We would also like to thank Liora Friedberg and Daniel Smullen for their help with the factorial vignette survey in Study 4, including authoring the transcripts and the video narration describing the information types in Study 4. This work was supported by NSF Award CNS-1330596, NSA Award #141333 and ONR Award #N00244-16-1-0006.

APPENDIX A

In Studies 1 and 2, we asked participants two pre-test questions. The first question asks about the frequency with which participants experience a privacy harm, and the second question asks participants to rank the harms in the order of severity. We show the frequency with which the participants reported to have experienced the harms in the table below.

Table A.1. Results of Frequency of Harm Experience from Study 1 and Study 2

Privacy Harm	Several times a day	About once a day	A few times a week	A few times a month	A few times a year	Never
Appropriation	13	17	52	83	147	98
Distortion	8	7	24	46	113	212
Induced Disclosure	9	20	59	102	133	87
Insecurity	12	20	43	88	181	66
Surveillance	48	33	65	84	111	69
Unanticipated Revelation	7	13	27	62	166	135
Unwarranted Restriction	6	12	32	65	169	126

In the table below, we show the results for the pre-test question from Study 1 and Study 2 which asks participants to rank harms in the order of severity. Each cell for *Rank N* represents the number of participants out of a total of 410 participants who took part in Studies 1 and 2 and who ranked the corresponding harm as “Nth” in severity, where *Rank 1*=most severe and *Rank 7*=least severe.

Table A.2. Results of Rank Order Question from Study 1 and Study 2

Privacy Harm	Number of Participants						
	Rank 1 (most severe)	Rank 2	Rank 3	Rank 4	Rank 5	Rank 6	Rank 7 (least severe)
Appropriation	43	61	54	84	63	69	36
Distortion	29	40	51	60	84	61	85
Induced Disclosure	14	28	43	38	68	93	126
Insecurity	88	73	61	50	50	57	31

Surveillance	115	67	53	39	36	47	53
Unanticipated Revelation	56	83	82	73	49	39	28
Unwarranted Restriction	65	58	66	66	60	44	51

APPENDIX B

In Study 1 and 2, we investigated the extent to which privacy harm predicts changes in perceived privacy risk. Table B.1 presents the *Model Term*, the corresponding model-estimated *Coefficient* along with the p-value, which tells us the statistical significance of the term over the corresponding baseline level, and the coefficient's *Standard Error* for $\$DT$ for Study 1 and Study 2 (see Equation 2 in Section 4.1 for the main regression equation). In our survey, the semantic scale option *Extremely Unwilling* has a value of 1, and *Extremely Willing* has a value of 8. A positive coefficient in the model signifies an increase in willingness to share and a negative coefficient signifies a decrease in willingness to share.

Table B.1. Multilevel Modeling Results for Data Type for Study 1 and 2

Model Term	Study 1		Study 2	
	Coeff.	Standard Error	Coeff.	Standard Error
Intercept (Family+Age Range+Induced Disclosure)	4.662***	0.750	2.533**	0.880
Data Type – Credit card	-1.812***	0.067	-2.342***	0.056
Data Type – Driver's license	-2.480***	0.067	-2.809***	0.056
Data Type – Full name	-0.592***	0.067	-1.172***	0.056
Data Type – Home address	-1.063***	0.067	-1.708***	0.056
Data Type – Phone number	-1.267***	0.067	-1.833***	0.056

*p≤.05 **p≤.01 ***p≤.001

In Table B.2 below, we present the *Model Term*, the corresponding model-estimated *Coefficient* along with the p-value, which tells us the statistical significance of the term over the corresponding baseline level, and the coefficient's *Standard Error* for $\$RL$ for Study 3 and Study 4. In our survey, the semantic scale option *Extremely Unwilling* has a value of 1, and *Extremely Willing* has a value of 8. A positive coefficient in the model signifies an increase in willingness to share and a negative coefficient signifies a decrease in willingness to share.

Table B.2. Multilevel Modeling Results for Risk Likelihood for Study 3 and 4

Model Term	Study 3		Study 4	
	Coeff.	Standard Error	Coeff.	Standard Error
Intercept (Family+Age Range+Induced Disclosure)	5.616***	0.462	6.340***	0.421
Risk – only 1 person in your workplace	-0.128	0.560	-0.611	0.533
Risk – only 1 person in your city	0.660	0.560	-0.519	0.533
Risk – only 1 person in your state	0.100	0.560	-0.355	0.533
Risk – only 1 person in your country	-0.310	0.560	-0.461	0.533

*p≤.05 **p≤.01 ***p≤.001

APPENDIX C

The following are the transcripts for the videos shown to participants in Studies 3 and 4, that describe the meanings of the different data types used in those studies.

Group 1 Transcript: “Device information includes any data about a mobile device, such as telephone numbers and serial numbers. Other identification numbers used by cellular carriers and advertisers are UDID (mobile device unique identifier) and IMEI (international mobile station equipment identity). These are used to track the identity of your device across applications, services, and cell towers. Mobile devices collect data through sensors, including microphones, cameras, and accelerometers, which can record movement, and when you pick up or drop a device. Mobile devices may have other sensors for measuring other phenomena, such as light levels, and temperature. Digital networks use cables and wireless signals to move information (called packets). Network devices rely on addresses to route information between computers. MAC (or Media Access Control) addresses are unique numbers associated with network device hardware. IP (or Internet Protocol) addresses are temporarily assigned numbers associated with MAC addresses used to route information across the Internet. Packets are transmitted as either unencrypted or encrypted. Some networks and applications use encryption to protect the confidentiality of packets. If packets are encrypted, it is difficult for uninvited third parties to view their contents. Applications on a network may require the user to enter a username and password to restrict access to an online account.”

Group 2 Transcript: “Microsoft Windows, Apple MacOS X, and Linux are all examples of an operating system. 0:06 Operating systems are special programs that run on electronic devices that provide a basic set of features enabling other programs (known as applications) to interact with the hardware resources of the device. Desktop and laptop computers and even mobile phones run operating systems. Hardware resources include processors, memory, sensors, and hard drives. Operating systems run applications in a process. Processes are provided permission by the operating system to access hardware. The operating system may also give processes permission to access files or folders on hard drives, or sensors, such as cameras. In operating systems such as Microsoft Windows, applications store and retrieve settings from a special table of values known as the registry, which is stored on the hard drive. In some operating systems, settings and other values are also stored in temporary files for each application. Browsers are examples of applications that store many kinds of data in temporary files. Browsers store session data when a user uses the browser to log in to a website. This data may include login information, cookies, browser history, images, and videos.”

Group 3 Transcript: “Sending emails and participating in chats are common ways to use electronic devices today. These activities are associated with contact information, including address books, telephone numbers, email addresses, and chat usernames. 0:16 Using email and chat applications, personal photos may be sent to others, and may also be sent to you. These photos may be available only temporarily, but are more commonly saved on your device after they have been downloaded. URLs (or Universal Resource Locator) are used to enter the address of a web page into a browser, but may also be used by browsers to transmit information. When browsing the internet, one can perform keyword searches using search engines such as Google, Bing, or Yahoo. In order to transmit your search query to the search website, keywords and other data may be encoded into the URL. Browsers store URLs that you visit throughout the day in a browser history. Malicious applications may be inadvertently downloaded and installed on various devices, operating without the knowledge or consent of the user. Malicious applications, such as keyloggers, can record information about users and transmit it to third parties. Keyloggers collect information by recording all of the keystrokes that users perform when typing on their keyboard. Keyloggers collect login information, data entered into online forms, typed email and chat messages, keyword searches, and other data, which may also be stored in malicious temporary files.”

The table below shows the (\$\text{WTS}\$) for different data types from Studies 3 and 4.

Table C.1. Multilevel Modeling Results for Data Types in Study 3 and Study 4

Term	SWtS		Δ SWtS from Study 3 to 4
	Study 3 w/o Benefits	Study 4 w/ Benefits	
Application Information	4.691	5.721	1.030
Application Session	3.941	5.268	1.327
Browser History	3.466	4.649	1.183
Chat History	2.879	4.378	1.499
Contact Information	3.579	4.874	1.296
Device ID	4.666	5.984	1.318
Device Information	4.779	6.043	1.265
Emails	2.791	4.340	1.549
IP Addresses	4.804	6.093	1.290
Keylogging Data	2.766	4.231	1.465
Keyword Searches	3.654	4.921	1.268
MAC Addresses	4.454	6.028	1.574
Memory Data	4.179	5.353	1.174
Network Information	4.466	5.862	1.396
OS Type and Version	5.904	6.603	0.699
Packet Data	3.954	5.437	1.483
Registry Information	4.279	5.371	1.093
Running Processes	4.579	5.790	1.212
Sensor Data	4.041	5.524	1.483
Temporary Files	4.041	5.209	1.168
UDID/IMEI	4.416	5.928	1.512
Usernames and Password	2.591	4.149	1.558
Video/Image Files	3.366	4.603	1.237
Websites Visited	3.679	4.871	1.193
Average increase in SWtS			1.303

APPENDIX D

In the table below, we show the estimated value of the privacy risk for each data type from Study 4 (see Table 6 and 7), followed by the average value from the survey ratings for each of the three predictors: discomfort, identifiability, and personal nature (see Study 5, Section 4.5). The *privacy risk score* or *unwillingness to share* for each data type is equal to $(9 - \text{SWtS})$, since we had an eight-point scale for measuring SWtS. The privacy risk score for the information type is calculated for the baseline levels: *workplace PC* ($\$CT$), *only one person in your family* ($\$RL$), *age range* ($\DT) and *investigating intellectual property and trade secrets* ($\$DP$).

Table D.1. Correlating Discomfort, Identifiability, Personal Nature with Privacy Risk

Information Type	Privacy Risk	Discomfort	Identifiability	Personal
Age Range	2.660	1.189	2.538	4.406
Application Information	3.279	3.951	4.291	5.608
Application Session	3.732	3.566	3.075	5.418
Browser History	4.351	5.857	4.529	6.961
Chat History	4.623	6.463	4.981	7.654

Contact Information	4.126	4.690	7.093	6.945
Device ID	3.016	4.231	4.125	5.216
Device Information	2.957	3.636	3.175	5.286
Emails	4.660	6.692	6.096	7.167
IP Addresses	2.907	4.138	5.667	6.184
Keylogging Data	4.769	6.552	4.964	7.500
Keyword Searches	4.079	4.383	3.316	5.745
MAC Addresses	2.973	4.731	5.182	5.248
Memory Data	3.648	4.545	3.947	5.679
Network Information	3.138	4.786	3.827	4.873
OS Type and Version	2.398	1.214	2.127	3.358
Packet Data	3.563	4.545	3.464	4.941
Registry Information	3.629	5.070	5.426	5.314
Running Processes	3.210	3.673	2.125	4.362
Sensor Data	3.476	4.596	3.083	5.172
Temporary Files	3.791	5.333	3.280	4.673
UDID/IMEI	3.073	4.788	4.377	5.096
Username	4.851	7.211	5.769	5.685
Password	4.851	7.211	5.769	8.436
Video/Image Files	4.398	5.527	4.827	6.608
Websites Visited	4.129	5.462	3.804	6.216

APPENDIX E

The trade-off analysis compares the incident data usage estimate to the perceived privacy risk for each data type. To estimate the tradeoff between usage and willingness to share a data type, we use the simulation usage method [Bhatia et al. 2016b] and the results from Study 4 for the willingness to share. The simulation usage method aims to simulate the incident cases that a security analyst estimates what percent of cases each data type is used in. The method assumes that, when an analyst estimates the percent of cases in which the data type is used, they are using the same number of cases to describe each estimate. The simulation universe consists of the set of data types D , and the set of incident reports $r \in R$, where $r \subseteq D$ and R can be partitioned into disjoint subsets, one subset for each analyst. We first generate 100 reports per analyst in the simulation. For each analyst's data type estimate, we randomly select a percentage within the reported interval wherein all percentages in the interval are equally likely (e.g., 64% is in 100-50%, and 64% is equally likely as 72% to be the analyst's true estimate). Next, we randomly select a corresponding subset of the 100 reports to match this percentage, and assign that data type to those reports. With this dataset, we can estimate the number of reports affected by removing a set of data types DR from all reports by computing the size of the set of affected reports $\{r \mid \forall r \in R, \exists d \in DR \text{ and } d \in r\}$.

The perceived privacy risk is measured by the estimated willingness to share $\$WtS$ (intercept=intellectual property+1 person in your family+ workplace computer) on a scale of 1 to 8, wherein 1=*Extremely Unwilling* and 8=*Extremely Willing*, and which estimates an average Internet user's acceptance of the risk. In Study 4 we observed that participant's scale use is skewed slightly toward "Extremely Willing," with 65% of all ratings lying between "Willing" and "Extremely Willing." To investigate responses by participants who utilize the full scale, we calculated the standard deviation (SD) for all ratings by participant. We found 19 participants, or 25% of the sample, with a $SD \geq 2$. The table below shows the trade-off of data use versus privacy risk for these 19 participants: as shown, participants are more willing to share information about who they are (e.g., IP address, UDID, MAC address), but they are less willing to share information about what they do (e.g., browser history, e-mails, websites visited, etc.)

Table E.1. Estimates for Incident Data Usage using the Simulated Usage Method and WtS

#	Data Type	Simulated Usage	\$WtS
1	Passwords	0.244	4.149
2	Usernames	0.610	4.149
3	Keylogging data	0.144	4.231
4	E-mails	0.408	4.340
5	Chat history	0.203	4.378
6	Video or image files	0.225	4.603
7	Browser history	0.422	4.649
8	Web sites visited	0.449	4.871
9	Contact information	0.336	4.874
10	Keyword searches	0.319	4.921
11	Temporary files	0.439	5.209
12	Application session data	0.244	5.268
13	Memory data	0.291	5.353
14	Registry information	0.459	5.371
15	Packet data	0.407	5.437
16	Sensor data	0.381	5.524
17	Application information	0.463	5.721
18	Running process information	0.526	5.790
19	Network information	0.667	5.862
20	UDID / IMEI	0.177	5.928
21	Device identifiers	0.464	6.984
22	MAC address	0.440	6.028
23	Device information	0.535	6.043
24	IP addresses / Domain names	0.673	6.093
25	Operating system information	0.600	6.603
26	OS type and version	0.588	6.603

REFERENCES

- [Acquisti and Grossklags 2005] A. Acquisti and J. Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE Security and Privacy*, vol. 3, no. 1, pp. 26–33, 2005.
- [Acquisti and Grossklags 2013] A. Acquisti and J. Grossklags. 2013. An online survey experiment on ambiguity and privacy. *Communications & Strategies*, 88(4): 19-39, 2012.
- [Acquisti et al. 2013] A. Acquisti, L.K. John and G. Lowenstein. 2013. What is the price of privacy. *Journal of Legal Studies*, 42(2): Article 1, 2013.
- [Acquisti et al. 2017] A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L. Cranor, S. Komanduri, P. Leon, N. Sadeh, F. Schaub, M. Sleeper, Y. Wang, and S. Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Comput. Surv.* 50, 3, Article 44 (August 2017). Available at SSRN: <https://ssrn.com/abstract=2859227>
- [Auspurg and Hinz 2014] K. Auspurg and T. Hinz. 2014. *Factorial Survey Experiments*. Vol. 175. SAGE Publications, 2014.
- [Barton 2014] K. Barton, MuMIn: Multi-model Inference, R Package, 1.10.0. 2014.
- [Bates et al. 2015] D. Bates, M. Maechler, B. Bolker, S. Walker. 2015. Fitting linear mixed-effects models using lme4. *J. Stat. Soft.*, 67(1): 1-48, 2015.
- [Bauer 1960] R.A. Bauer. 1960. *Consumer behavior as risk-taking, dynamic marketing for changing world*. American Marketing Association, Chicago, 389, 1960.

- [Berendt et al. 2005] B. Berendt, O. Günther, and S. Spiekermann. 2005. Privacy in e-commerce: Stated preferences vs. actual behavior. *Communications of the ACM*, vol. 48, no. 4, pp. 101–106, 2005.
- [Bhatia et al. 2016a] J. Bhatia, T. D. Breaux, J. R. Reidenberg, T. B. Norton. 2016. A Theory of Vagueness and Privacy Risk Perception. *IEEE 24th International Requirements Engineering Conference (RE'16)*, 2016.
- [Bhatia et al. 2016b] J. Bhatia, T. D. Breaux, L. Friedberg, H. Hibshi, D. Smullen. 2016. Privacy Risk in Cybersecurity Data Sharing. *ACM 3rd International Workshop on Information Sharing and Collaborative Security (WISCS)*, Vienna, Austria, Oct. 2016, 57-64.
- [Brooks et al. 2017] S. Brooks, M. Garcia, N. Lefkowitz, S. Lightman, E. Nadeau. 2017. *An Introduction to Privacy Engineering and Risk Management in Federal Systems*. National Institute of Standards and Technology Internal Report 8062, January 2017.
- [Buhrmester et al. 2011] M.D. Buhrmester, T. Kwang, and S.D. Gosling. 2011. Amazon’s Mechanical Turk: a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 3 (6), 13-5.
- [Charness et al. 2012] G. Charness, U. Gneezy, and M. A. Kuhn. 2012. "Experimental methods: Between-subject and within-subject design," *Journal of Economic Behavior & Organization*, Elsevier, vol. 81(1), pages 1-8.
- [Cichonski et al. 2012] P. Cichonski, T. Millar, T. Grance, and K. Scarfone. 2012. *Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology*, 800-61. Revision 2. NIST Spec. Publ., vol. 800–61, p. 79, 2012.
- [Clark and Watson 1995] L. A. Clark and D. Watson. 1995. Constructing validity: Basic issues in objective scale development. *Psychological Assessment*, 7(3): 309-319, 1995.
- [Cortina and Nouri 2000] J. Cortina and H. Nouri. Effect size for ANOVA designs. Sage Publications, 2000.
- [Creswell 2014] J. Creswell. *Research design: qualitative, quantitative, and mixed methods approaches*. SAGE publications, 2014.
- [Dinev et al. 2015] Tamara Dinev, Allen R. McConnell, and H. Jeff Smith (2015) Research Commentary— Informing Privacy Research Through Information Systems, Psychology, and Behavioral Economics: Thinking Outside the “APCO” Box. *Information Systems Research* 26(4):639-655. [https](https://doi.org/10.1287/isre.2015.264.639)
- [Ellsberg 1961] Daniel Ellsberg. 1961. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75 (4): 643–669, 1961.
- [Everton et al. 2005] W.J. Everton, P.M. Mastrangelo and J.A. Jolton. Personality correlates of employee’s use of work computers. *CyberPsychology and Behavior*, 8, 143–153. 2005.
- [Faul et al. 2007] F. Faul, E. Erdfelder, A.-G. Lang, and A. Buchner. 2007. G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behav. Res. Methods*, 39(2): 175-191, 2007.
- [Fischhoff et al. 1978] B. Fischhoff, P. Slovic, S. Lichtenstein, S. Read, B. Combs. 1978. How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits. *Policy Sci.* 9: 127-152, 1978.
- [Furr 2011] M. Furr. 2011. *Scale construction and psychometrics for social and personality psychology*. SAGE Publications Ltd, 2011.
- [Gelman and Hill 2006] A. Gelman and J. Hill. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge Univ. Press, 2006.
- [Hibshi et al. 2015] H. Hibshi, T. D. Breaux, and S. B. Broomell. 2015. Assessment of risk perception in security requirements composition. *IEEE 23rd Int. Requir. Eng. Conf. (RE)*, pp. 146-155, 2015.
- [Hilty et al. 2004] L. M. Hilty, C. Som, and A. Köhler. 2004. Assessing the human, social and environmental risks of pervasive computing. *Human and Ecological Risk Assessment*, vol. 10, pp. 853–874, 2004.
- [Hong et al. 2004] J. I. Hong, J. D. Ng, S. Lederer, and J. A. Landay. 2004. Privacy risk models for designing privacy-sensitive ubiquitous computing systems. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques (DIS '04)*, ACM, New York, NY, USA, 91-100, 2004. DOI=<http://dx.doi.org/10.1145/1013115.1013129>

- [Hoofnagle et al. 2010] C. Hoofnagle, J. King, S. Li, and J. Turow. 2010. How different are young adults from older adults when it comes to information privacy attitudes and policies. *SSRN Working Paper Series 2010*; 4(19): 1-20. DOI=10.2139/ssrn.1589864
- [Horton et al. 2011] J.J Horton, D.G. Rand, R. J. Zeckhauser. The online laboratory: conducting experiments in a real labor market. *Experimental Economics* September 2011, Volume 14, Issue 3, pp 399–425.
- [Hustinx 2010] Peter Hustinx. 2010. Privacy by design: delivering the promises. *Identity in the Information Society*, Volume 3, Issue 2, pp 253–255, August 2010.
- [Iachello and Hong 2007] Giovanni Iachello and Jason Hong. 2007. End-user privacy in human-computer interaction. *Trends Human-computer Interact.* 1, 1 (January 2007), 1-137, 2007. DOI=http://dx.doi.org/10.1561/1100000004
- [Kaplan and Garrick 1981] S. Kaplan and B. J. Garrick. 1981. On the quantitative definition of risk. *Risk Analysis*, 1 (1): 11-27, 1981.
- [Kendall 1948] M.G. Kendall. *Rank Correlation Methods*. 1948. Charles Griffin and Company Limited.
- [Knight 1921] F.H. Knight. 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin Company, 1921.
- [Knijnenburg and Kobsa 2014] B. Knijnenburg, A. Kobsa. 2014. Increasing sharing tendency without reducing satisfaction: finding the best privacy-settings user interface for social networks. *35th Int'l Conf. Info. Sys.*, pp. 1-21, 2014.
- [Kulas and Stachowski 2013] J. T. Kulas and A. A. Stachowski. 2013. Respondent rationale for neither agreeing nor disagreeing: Person and item contributors to middle category endorsement intent on Likert personality indicators. *J. Res. Pers.*, vol. 47, no. 4, pp. 254-262, Aug. 2013.
- [Lederer et al. 2003] S. Lederer, J. Mankoff, and A. K. Dey. 2003. Towards a Deconstruction of the Privacy Space. *Workshop on Privacy in Ubicomp 2003: Ubicomp communities: privacy as boundary negotiation*. 2003.
- [Lederer et al. 2004] S. Lederer, J. I. Hong, A. K. Dey, and J. A. Landay. 2004. Personal privacy through understanding and action: five pitfalls for designers. *Personal Ubiquitous Comput.* 8, 6 (November 2004), 440-454. DOI=http://dx.doi.org/10.1007/s00779-004-0304-9
- [Moor 1997] J. H. Moor. Towards a theory of privacy in the information age. 1997. *Computers and Society*, vol. 27, no. 3, pp. 27–32, 1997.
- [Mugan et al. 2011] J. Mugan, T. Sharma, N. Sadeh. 2011. *Understandable Learning of Privacy Preferences Through Default Personas and Suggestions*. Carnegie Mellon University's School of Computer Science Technical Report CMU-ISR-11-112, <http://reports-archive.adm.cs.cmu.edu/anon/isr2011/CMU-ISR-11-112.pdf>, August 2011.
- [Murphy 1996] R. S. Murphy. 1996. Property rights in personal information: An economic defense of privacy. *Georgetown Law Journal*, vol. 84, p. 2381, 1996.
- [Mastrangel et al. 2006] P.M. Mastrangel, W. Everto, J.A. Jolton. Personal use of work computers: distraction versus destruction. *CyberPsychology & Behavior*, 9:730–41, 2006.
- [Nakagawa and Schielzeth 2013] S. Nakagawa, H. Schielzeth. A general and simple method for obtaining R² from generalized linear mixed-effects models. *Methods in Ecology And Evolution*. 2013;4(2):133-142.
- [Neal et al. 2003] A. Neal, M. Humphreys, D. Leadbetter, and P. Lindsay. 2003. Development of hazard analysis techniques for human-computer systems. In *Innovation and Consolidation in Aviation*, (G. Edkins and P. Pfister, eds.), pp. 255–262, Aldershot, UK: Ashgate, 2003.
- [Nissenbaum 2004] H. Nissenbaum. 2007. Privacy as contextual integrity. *Washington Law Review*, 79, 2004
- [Nissenbaum 2009] H. Nissenbaum. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford Law Books, 2009.
- [Olson et al. 2005] J. S. Olson, J. Grudin, and E. Horvitz. 2005. A study of preferences for sharing and privacy. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems (CHI EA '05)*, ACM, New York, NY, USA, 1985-1988. DOI=http://dx.doi.org/10.1145/1056808.1057073
- [Palen and Dourish 2003] Leysia Palen and Paul Dourish. 2003. Unpacking "privacy" for a networked world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, New York, NY, USA, 129-136. DOI=http://dx.doi.org/10.1145/642611.642635

- [Perrin and Duggan 2015] A. Perrin, M. Duggan. 2015. *Americans' Internet Access: 2000-2015*. PEW Internet and American Life Project, June 26, 2015.
<http://www.pewresearch.org/fact-tank/2016/02/19/americans-feel-the-tensions-between-privacy-and-security-concerns/>
- [Paolacci et al. 2010] G. Paolacci, J. Chandler, and P.G. Ipeirotis. 2010. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making* 5, 411–419.
- [PWC 2016] PWC, “The Global State of Information Security Survey 2016: Turnaround and transformation in cybersecurity,” 2016.
- [R Core Team 2015] R Core Team. 2015. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. 2015. URL <http://www.R-project.org/>.
- [Rainie and Maniam 2016] L. Rainie and S. Maniam. 2016. *Americans feel the tensions between privacy and security concerns*. PEW Internet and American Life Project, February 19, 2016.
- [Rand 2012] D.G. Rand. (2012). The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, 299, 172–179
- [S. 754] S.754 – 114th U.S. Congress. Cybersecurity Information Sharing Act of 2015.
- [Saltzer and Schroeder 1975] J.H. Saltzer and M.D. Schroeder. 1975. The Protection of Information in Computer Systems. *Proceedings of the IEEE*, 1975. 63(9): p. 1278- 1308
- [Shadish et al. 2002] W. R. Shadish, T.D. Cook, and D. T. Campbell, *Experimental and quasi-experimental designs for generalized causal inference*, Boston, MA, US: Houghton, Mifflin and Company, 2002.
- [Spearman 1904] C. Spearman. 1904. The Proof and Measurement of Association between Two Things. *The American Journal of Psychology*, vol. 15, no. 1, 1904, pp. 72–101.
- [Starr 1969] C. Starr. 1969. Social benefit versus technological risk. *Science*, 165, pp. 1232-1238, 1969.
- [Slovic 2000] P. Slovic. 2000. *The Perception of Risk*. Earthscan Publication, 2000.
- [Smith et al. 2011] H. J. Smith, Tamara Dinev, and Heng Xu. “Information Privacy Research: An Interdisciplinary Review.” *MIS Quarterly* 35, no. 4 (2011): 989-1015. doi:10.2307/41409970.
- [Solove 2006] D.J. Solove. 2006. A Taxonomy of Privacy. *University of Pennsylvania Law Review*, vol 154, no 3, January 2006, p. 477.
- [Solove 2008] Daniel J. Solove. 2008. *Understanding Privacy*. Harvard University Press, 2008.
- [Stoneburner 2002] Gary Stoneburner, Alice Y. Goguen, and Alexis Feringa. 2002. *Risk Management Guide for Information Technology Systems*. SP 800-30, Technical Report, NIST, Gaithersburg, MD, United States, 2002.
- [Tversky and Kahneman 1974] A. Tversky and D. Kahneman. 1974. Judgment under uncertainty: heuristics and biases. *Science* 185, 1124–1131.
- [Wakslak and Trope 2009] C. Wakslak and Y. Trope. 2009. The effect of construal level on subjective probability estimates. *Psychol. Sci.*, vol. 20, no. 1, pp. 52-58, Jan. 2009.
- [Wallander 2009] Lisa Wallander, “25 years of factorial surveys in sociology: A review,” *Social Science Research*, Volume 38, Issue 3, 2009, pp. 505-520.
- [Wang et al. 2011] Y. Wang, G. Norice, and L. F. Cranor. 2011. Who Is Concerned about What? A Study of American, Chinese and Indian Users’ Privacy Concerns on Social Network Sites. *International Conference on Trust and Trustworthy Computing Trust 2011: Trust and Trustworthy Computing* pp 146-153.
- [Wang et al. 2014] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for Facebook. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2367-2376. DOI=<http://dx.doi.org/10.1145/2556288.2557413>
- [Westin 1967] A. F. Westin. 1967. *Privacy and Freedom*. New York, NY: Atheneum, 1967.

Received XXXX; revised XXXX; accepted XXXX