PRELIMINARY DRAFT

Framing and the Malleability of Privacy Choices

Idris Adjerid, Alessandro Acquisti, George Loewenstein

[2014 WORKSHOP ON THE ECONOMICS OF INFORMATION SECURITY]

Policy approaches for addressing emerging consumer privacy concerns increasingly rely on providing consumers with more information and control over the usage of their personal data. In three experiments, we evaluate the efficacy of such mechanisms in the face of subtle but common variation in the presentation of privacy choices to consumers. We find that consumers' decision frames and thus, their propensity to select privacy protective alternatives can be subtly but powerfully influenced by common-place heterogeneity in the presentation of privacy choices. Our results suggest that choice mechanisms alone may not reliably serve policy maker goals of protecting consumers' privacy in the face of emerging data practices by firms.[1]

1. Introduction

In the United States, as well as in other countries, notice and consent mechanisms have become the predominant approach to consumer privacy protection. A recent report by the World Economic Forum (2013), for example, advocates a shift from "trying to control the data itself" to "focusing on the uses of data," and proposes a framework for increasingly granularity of control over the use of data by consumers, but not necessarily limiting data collection by firms. The implicit premise of such policies is simple and, at least superficially, compelling: Armed with information and control, consumers can make informed choices between different market offerings (from social networks to mobile apps) and manage their privacy in accordance with their individual preferences for privacy, without stifling innovation as a

[1] The authors gratefully acknowledge research support from the following organizations: National Science Foundation (Award CNS-1012763), IWT SBO Project on Security and Privacy for Online Social Networks (SPION), U.S. Army Research Office under Contract DAAD190210389 through Carnegie Mellon CyLab, and TRUST (Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422).

result of rigid regulation (FTC, 2012; The White House 2012).The industry is, largely, an active advocate of such an approach, and an active participant in it, with many websites giving users extensive controls over data sharing (Chavez, 2011; Richter, 2011). Facebook, for example, provides members with control over uses of their personal information for advertising purposes, and mobile platforms provide consumers with control over the collection and use of their location information.

Extant research has suggested, however, that providing consumers with control does not guarantee that individuals will be able to use that control to make informed, self-interested, choices about data sharing; indeed one recent paper by two of the current authors showed that giving web-users more control encouraged them to share more information, in effect giving them more rope to hang-themselves through self-revelation (Brandimarte, Acquisti, and Loewenstein, 2012). Furthermore, and the main focus of the current paper, tools and options providing control may focus on irrelevant choices, be confusing, or be presented in a fashion that subtly but powerfully influences the consumer's propensity to pick protective or transparent data sharing settings. For instance, across systems, services, and sites, choices surrounding sensitive collections of personal information may not be consistently presented to consumers as "privacy" choices. As a case in point, the choice to limit the use of personal information for advertising purposes is presented to Facebook users under the header of "Facebook Ads" (rather than under Facebook's "Privacy Settings and Tools"). Similarly, on Android mobile devices, the options to limit the collection and use of location information are presented to users under the "Location Settings" header (not under "Privacy Settings").

While those differences might seem inconsequential, findings from behavioral decision research have shown that subtle changes in a consumer's decision frame (a decision maker's conception of acts, outcomes, and contingencies associated with a particular choice) can significantly affect her actions (Tversky and Kahneman, 1981). Focusing on lesser studied but still common variations in the presentation of privacy choices, in this manuscript we investigate, and find evidence in support of, the prediction that consumer privacy decision making will be significantly impacted by subtle changes in decision frames induced by commonplace heterogeneity in the presentation of privacy choices.

We conducted three randomized experiments on Amazon's Mechanical Turk in which we evaluated the impact on participants' likelihood of selecting privacy protective options of altering (1) the label on privacy relevant choices, (2) the mix of high relevance and low-relevance choices, and (3) the accept vs. reject presentation of privacy relevant choices. We find that simply altering the decision frame of otherwise identical choices can have a substantial effect on individual choice. Changing the label of privacy-relevant choices from "Privacy Settings" to "Survey Settings" resulted in participants being 56% less likely to choose the privacy protective option. We also found that participants were approximately half as likely to select the privacy protective option for a high relevance choice when it was presented together within a set of low relevance choices. Finally, we found that participants were 51% less likely to select the privacy protective option when presented as a choice to allow a use of their personal information (allow frame) than when they were presented the objectively identical setting as a choice to reject a use of their personal information (prohibit frame). In all three studies presented in the paper, the set of choice frame manipulations evaluated were inspired by, and closely modeled on, the framing of privacy-relevant choices in existing contexts. The close modeling of our studies on real choice settings is intended to increase the relevance of the research findings for regulators, firms, and consumers.

This work contributes to three main streams of research. First, an ongoing debate in the policy and behavioral economic literature focuses on how the design (or "architecture") of market choices can ameliorate or impair individuals' decision making and welfare. In this context, Keller et al. (2011) have recently proposed a novel approach termed *enhanced active choice*. They posit that decisions in active choice context (i.e., when consumers are forced to make a choice and no default is present) can be framed in a manner that promotes the desired goals of the communicator. Active choice architectures are particularly relevant in the context of privacy decision making, as consensus on choice defaults has been particularly difficult to reach. For instance, proposals to have browsers' do not track capabilities activated for consumers by default resulted in significant delays for the Do Not Track initiative (Singel, 2012).[3]

[3] Do Not Track is a policy initiative seeking to provide browser functionality that would allow consumers to restrict the collection of data on their online browsing behavior.

Whereas prior research on enhanced active choice focused on manipulations of choice frames that sought to promote behavior intended to be beneficial for individuals (e.g. getting a flu shot), in the current research we demonstrate how choice framing can be leveraged in active choice contexts to arbitrarily induce increased information revelations and data allowances from technology users.

Secondly, this work contributes to the behavioral economics and marketing literature on framing effects. Prior work (Huber, Payne, and Puto, 1982; Simonson and Tversky, 1992) has demonstrated that the addition of irrelevant or strictly dominated options can lead to contrast effects which have the effect of shifting preferences between non-dominated options. In our modified context, in which participants were asked to make a sequence of choices, we find that the mix of low and high relevance choices leads to an *assimilation* effect, in which highly relevant choices are overshadowed by the set of low relevance choices. This finding has broad implications in a number of decision contexts; for example, the manner in which individuals interpret or react to nutritional information or side-effects of prescription medication.

Finally, this work contributes to the policy literature addressing emerging consumer privacy concerns and eliciting consumer preferences for privacy. The extant literature on consumer privacy decision making has traditionally modeled consumers as economically rational agents that make stable and consistent tradeoffs between the utility from data disclosures and privacy risks (e.g., Dinev and Hart, 2006; Fogel and Nehmad, 2009). However, a smaller but growing body of work has started to document some non-normative factors affecting people's privacy decision making (e.g., Moon, 2001, Acquisti, John, and Loewenstein, 2012). This work contributes to this emerging stream of research and is, to our knowledge, the first to highlight that control mechanisms, ostensibly aimed at improving consumer privacy, can themselves be fashioned in a manner that has a subtle, non-normative, and powerful influence on consumers' propensity to pick protective data sharing settings. This suggests these mechanisms can introduce increased consumer privacy risk independent of consumers' subsequent disclosure behavior (which has been the primary focus of prior studies; Brandimarte, Acquisti, and Loewenstein, 2012). These control mechanisms are increasingly relevant in light of the ongoing shift

towards indirect ways of collecting consumer personal information, not through overt requests for information (i.e. explicit consumer disclosures) but by, e.g., monitoring patterns of online browsing.

The central implication of this work is that providing consumers with greater control over privacy options may be a necessary but not sufficient policy mechanism to address privacy concerns, particularly in contexts in which firms have strong incentives to strategically leverage subtle manipulations of choice framing to solicit high rates of information sharing from consumers. These concerns are exacerbated if consumers fail to notice such subtle variation in the presentation of privacy relevant choices, or as past research demonstrates (Liberman, Samuels & Ross, 2004) and the current research reinforces, significantly underestimate their impact on behavior. Such inadvertent and unaware susceptibility to framing effects is of increasing consequence to consumers, given the growing usage of personal information in commercial contexts, some of which may be particularly intrusive or even discriminatory. For instance, Sweeney (2013) finds that black-identifying names were 25% more likely to get an online ad suggestive of an arrest record relative to white-identifying names.

However, we caution that we cannot claim to know whether consumers are providing more or less personal information in the marketplace than is in their personal interest. What we can claim with greater confidence is that, *if* the goal of policy makers is to protect privacy, our findings suggest that the current trend in policy toward increasing choice and notification mechanisms (see, e.g.,World Economic Forum, 2013) may not necessarily achieve that intended goal. Alternative mechanisms should therefore also be considered, such as leveraging OECD fair information practice principles for baseline consumer protection against data practices perceived to be particularly intrusive or harmful. At a minimum, this could result in a more manageable set of privacy contexts that consumers are required to navigate, increasing the likelihood that they invest in understanding the nuances and tradeoffs associated with the privacy choices presented to them. For firm data practices where choice is desirable (e.g. when significant benefit exists for consumers), additional protections could include designating uniform and consistent standards for presenting choice coupled with considerations of choice architecture (e.g. choice framing) that limit firms' abilities to manipulate consumers in their own interests but instead empower consumers

to make choices that reflect their desired balance of personal privacy and benefit from the collection and use of their personal information.

2. Control, Privacy Decision Making, and Choice Architecture

In recent years, policy makers have advocated increased consumer choice for emerging privacy-sensitive technologies and practices such as data aggregation and behavioral advertising (see, e.g., FTC, 2012). Industry has, by and large, responded to calls from policy makers to provide consumers with increased choice in the collection and use of their personal information.[4] This reaction from industry can be largely attributed to the self-regulatory nature of proposed policy interventions, which allows firms to define consumer choice in a manner that fits their particular business context and environment (Solove, 2013). However, critics argue that an increased reliance on consumer choice considerably shifts the burden of understanding and evaluating complex and uncertain privacy tradeoffs to consumers, and that increased choice may not necessarily reduce consumer privacy risks or better align behavior with individual preferences for privacy. For instance, scholars point out that "organizations, as a rule, will have the sophistication and motivation to find ways to generate high opt-in rates" (Solove, 2013) and that "many data-processing institutions are likely to be good at obtaining consent on their terms.." Schwartz (2005). These concerns seem justified, given that other well-intentioned regulatory interventions relying on increased consumer choice have been subverted by the way in which these choices have been presented to consumers. In 2010, for example, regulators required [5] that banks halt practices of levying, by default, exorbitant fees for consumers who overdrafted their accounts.[6] In response to the requirement that consumers be defaulted into a regime in which they would not be able to overdraw their accounts via

[4] For instance, data aggregators and digital advertising groups have provided consumers some degree of control over the collection and use of their personal information (Singer, 2013).

[5] 12 C.F.R. § 205

[6] These fees totaled 37 billion dollars in revenue for banks in 2009 and have been described by critics as a form of predatory lending, given that consumers earning less than 30,000 a year were twice as likely to incur them (Pew Center of the States, 2012).

ATM withdrawals, banks presented the choice to *continue* to be able to overdraft and incur these fees to customers as the option to enroll in "overdraft protection."[7]

These concerns are exacerbated by a growing body of work finding that people are particularly susceptible to deviations from economically rational models of decision making in the context of privacy decision making (Acquisti, John, and Loewenstein, 2012; Adjerid, Acquisti, Brandimarte, and Loewenstein, 2013). This may suggest that privacy decision making is also vulnerable to framing effects that may defy or nullify control mechanisms and thus, conferring on 'choice architecture' the potential to dramatically influence consumer privacy decision making, but not necessarily in directions that benefit service consumers. While we consider specific examples of the former (framing effects that may defy control mechanisms) in the following section (Section 3), we discuss the latter (relevant choice architectures) in the rest of this section.

The behavioral economics literature (e.g., Johnson et al., 2012) has studied manipulations of choice architecture (the manner in which a choice is presented to a decision-maker) in an effort to counteract or even leverage decision biases to improve individual decision making across a range of contexts (e.g. healthcare, saving behavior, organ donation, etc.). Keller et al. (2011) proposed a choice architecture they termed *enhanced active choice* in which choices presented to individuals are *framed* in a manner that promotes the desired goals of the communicator. Keller et al. (2011) argue that when use of default choices is controversial or impractical, active choice coupled with manipulations of *choice framing* can be used to influence individual choice. Framing effects have been studied across an established empirical and theoretical literature (Levin, Schneider, and Gaeth, 1998; Kahneman and Tversky, 1981) and refer to the phenomena of "simple and unspectacular changes" in the presentation of decision problems leading to changes in choice (Kühberger, 1998). Classic framing studies have focused on differences that arise from decision frames that highlight positive vs. negative aspects of a given choice. Kahneman and Tversky (1979) demonstrate that highlighting costs (lives lost) from a medical

[7] A survey of more than 6,000 people from the Pew Center administered following implementation of the regulation found that large numbers of people had fallen for the ruse, despite their preference for having such transactions declined, and that nearly one in five people still incurred overdraft fees (Pew Center of the States, 2012).

intervention vs. the gains (lives saved) can lead to an increased preference for risky options; Levin and Gaeth (1988) finds that perception of the quality of ground beef differ based on whether it is labeled as "75% lean" or "25% fat"; Ganzach & Karsahi (1995) find that framing choices in terms of losses (e.g. losses suffered from not using a credit card) is a more effective at altering behavior relative to framing which highlighting gains.

3. Market Examples of Privacy Frames and Experimental Hypotheses

In order to identify manipulations of decision frames that may be leveraged in an active privacy choice context, we evaluated current approaches employed by firms for soliciting consumer choice. Specifically, we sought subtle variation in these approaches with the potential to differentially highlight consumer privacy concerns. The focus on subtle variations allows us to identify manipulations of decision frames that consumers and policy makers may not be likely to identify as significant influences on their own choices. While we identify anecdotal evidence of popular technology services moving away from choice framing with the potential to highlight privacy considerations, we cannot claim that the significant heterogeneity in control mechanisms we identified, or the changes over time observed in them, is necessarily purposeful (that is, intended to elicit varying data allowances from consumers). In fact, at least some of the variation we identified is likely a consequence of practical limitations in presenting consumer privacy choices (e.g. many choices in the context of social media have privacy relevance, and thus cannot all be presented under the same label). In addition, some of the variation we identified may even drive consumers towards more protective behavior - which would presumably be at odds with any firm motivations to elicit greater data allowances from consumers. As a result, the impact of identified variations on individual decision frames and subsequent consumer privacy decision making is uncertain, and is the object of our experimental testing.

3.1 Label-Framing Effect

We first consider the potential impact of minor changes to the labeling of privacy-relevant choices on consumer decision frames and their subsequent privacy protective behavior. We found that privacy-relevant contexts varied significantly in terms of whether choices are presented to consumers as "privacy" choices. For instance, the Android mobile platform (Figure 1, Android version 4.03) present choices with significant privacy implications using descriptive labels such as "Location Settings" or "Account Settings" and present other choices with arguably less privacy relevance (e.g. backup options) as "Privacy Settings". In contrast, the Apple iOS 7 (see Figure 1) presents similar choices to consumers (including the choice to limit tracking by advertisers) under the general "Privacy" label.

[Figure 1: Mobile Platform Settings]

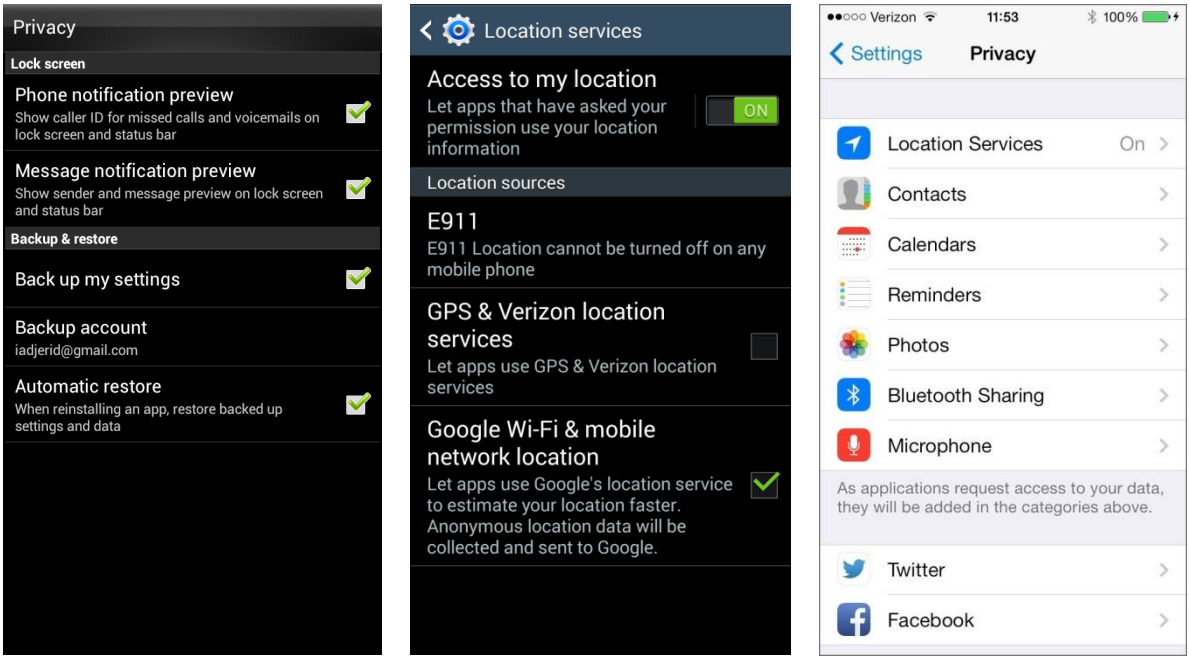

Android OS 4.0.3 ***Android OS 4.3*** ***iOS 7***

Moreover, this variation in the labeling of privacy relevant choices is not unique to mobile platform. For example, Facebook presents some subset of setting to users as "Privacy Settings and Tools" (Figure 2a), while other privacy relevant choices are presented using different labels (Figure 2b).

[Figure 2a: Facebook Privacy Settings]

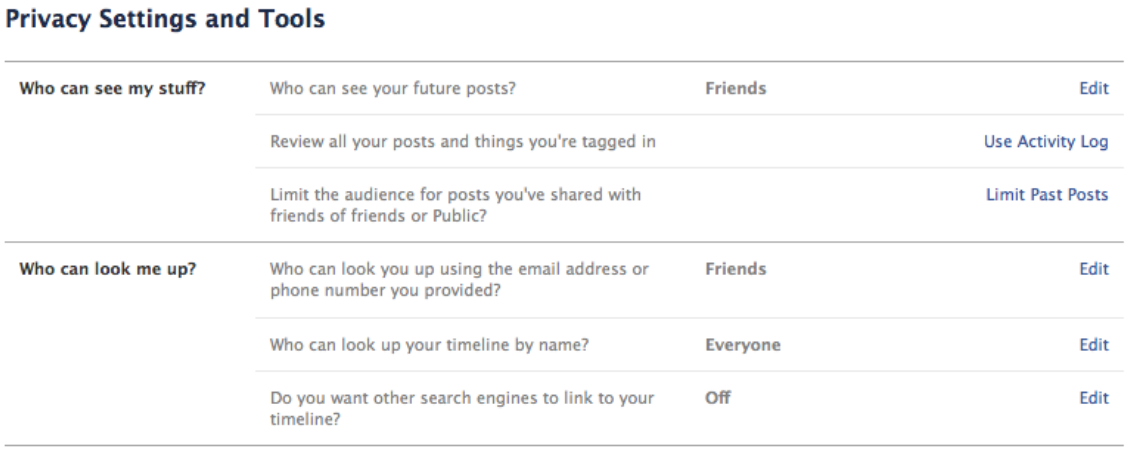

[Figure 2b: Timeline and Tagging Settings]

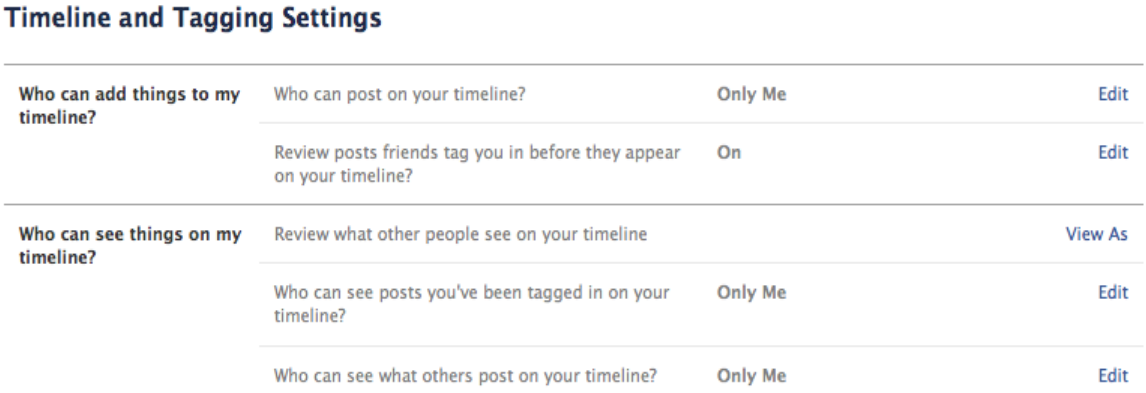

Finally, there is some indication that popular services presenting privacy relevant choices to consumers as "privacy" choices are on the decline. For instance, Facebook recently altered the label of the settings on the main Facebook page from "Privacy Settings" to simply "Settings" (See Figure 3) and the more recent versions of Android have dropped "Privacy" settings altogether.

[Figure 3: Changes in Facebook Privacy Settings]

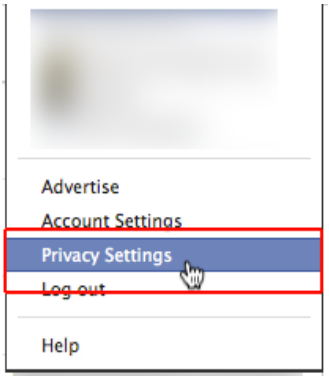

Facebook Main Page Settings 2013

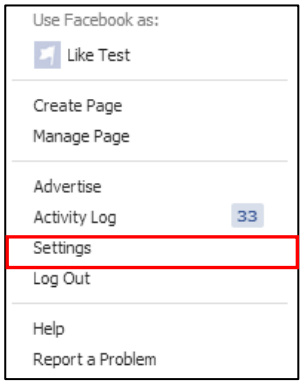

Facebook Main Page Settings 2014

Prior research has found that simple and subtle changes to the labeling of decision problems can significantly alter behavior. Liberman, Samuels, and Ross (2004), for example, find that the labeling of a prisoner dilemma game influenced participants' perceptions of the goal of the game with a "Wall Street Game" label resulting in significantly less cooperation from participants, relative to a "Community Game" label. Moreover, they find that participants (both with and without psychology backgrounds) greatly underestimate the effect of these label manipulations on their own behavior. Further, Burnham, McCabe, and Smith (2000) find a strong impact on cooperation when participants in a two-player reciprocity game are labelled as either "partners" or "opponents." In line with this stream of research, we hypothesize that, whether intentional or justified by other platform constraints, this variation in the labeling of privacy relevant choices may significantly alter consumers' decision making:

H1a: Labeling otherwise identical choices as "Privacy Settings" relative to an alternative descriptive label will alter the decision frame in a manner that highlights privacy considerations, resulting in the choice of more privacy protective options.

The behavioral literature has identified individual engagement as a moderating factor of framing effects. For instance, previous works (e.g. Maheswaran & Meyers-Levy 1990 and Rothman et al. 1993) found

that framing effects are more pronounced for populations which are more engaged in a particular decision context. Thus, we hypothesize:

H1b: The effect of a "Privacy Settings" label on the choice of the privacy protective option will be more pronounced for high relevance choices.

3.2 Mixed Relevance Choice Sets

We also consider the potential for variation in choice sets to alter the decision frame again, in a manner that differentially highlights privacy concerns. Specifically, we consider the impact of presenting privacy choices as either one homogenous set of high relevance choices or a mixture of both low and high relevance choices. Mixed relevance choice sets are common in privacy-relevant contexts. For example, Android users are presented the choice to restrict the use of their location information for Google services while also being offered the option to view a "Compass Warning" (Figure 1). Similarly, iOS 7 (Figure 4) presents consumers, within the same choice set, the option to restrict access to their location information for applications with presumably varying degrees of privacy relevance (e.g. Twitter vs. World Clock).

[Figure 4: iOS 7 Location Services Settings]

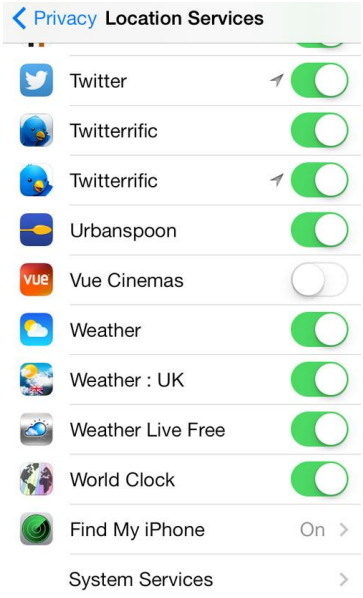

Mixed relevance choice sets also prevail in other contexts. For example, Figure 5 below depicts a subset of the nearly 100 companies that participate in the Digital Advertising Alliance's effort to allow consumers to opt-out of behavioral advertising. In this case, it is possible the tradeoffs (both benefits and costs) associated with opting-out of Google's services (the largest online advertising firm) may greatly outweigh those associated with smaller, lesser known firms.

[Figure 5: Digital Advertising Alliance's interface for opting-out of behavioral advertising]

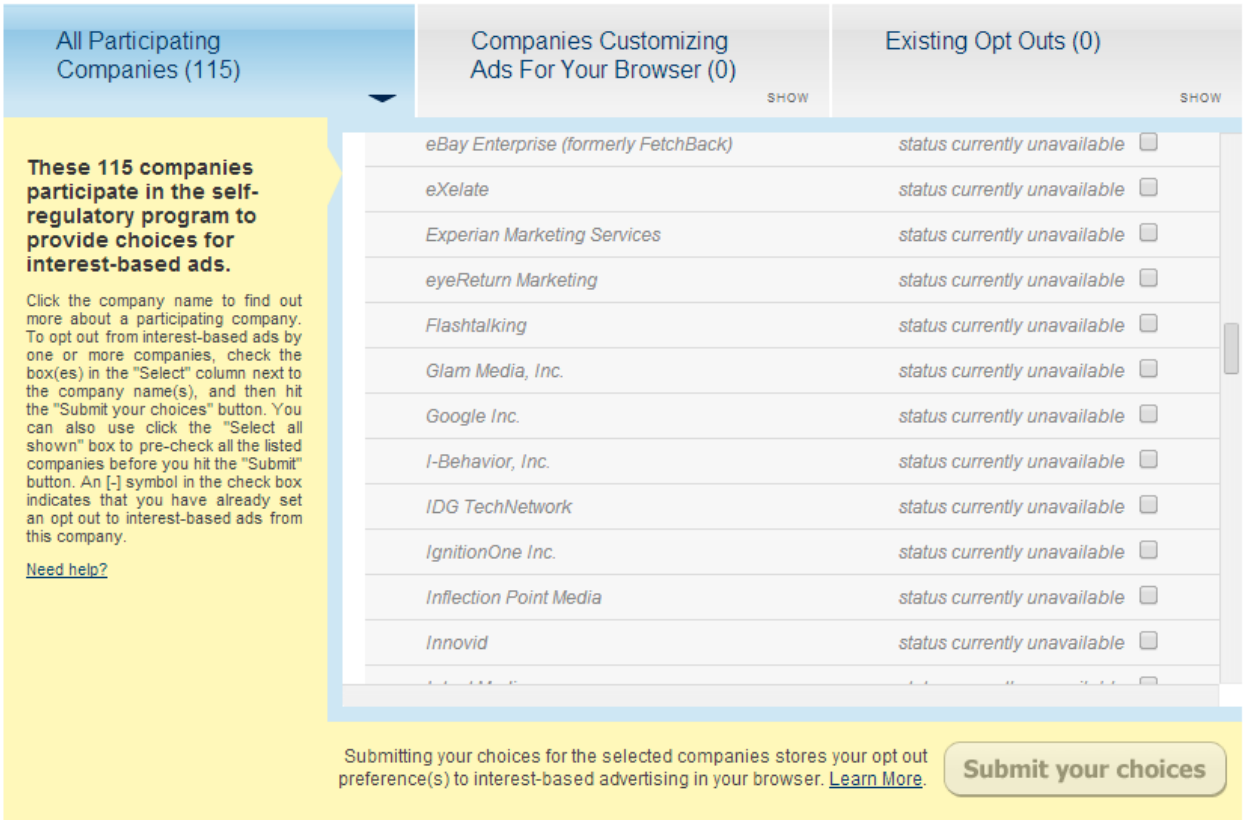

The effect of mixed relevance choice sets can be ambiguous. On the one hand, an established theoretical (e.g. Kahneman & Tversky, 1979) and empirical literature (e.g. Mas, 2006; Genesove and Mayer, 2001) finds that individual judgment and decision making is sensitive to relative differences and comparative effects, including in privacy contexts (e.g. Acquisti, John, and Loewenstein, 2012). Thus mixed relevance choice sets may result in contrast effects, through which a high relevance choice is perceived, in relative terms, as more risky in a mixed set relative to a homogenous set. On the other hand, a mix of low and

high relevance choices may lead to an assimilation effect, through which participants' judgment of the risk associated with high relevance choices is diminished or simply not noticed due to the low relevance choices preceding it.

What factors will determine which of these effects prevails in a specific situation? We conjecture that attention will play a key role. Individuals who quickly peruse a list of privacy protection options, in which the focal option is embedded in a sea of obviously trivial other options, are likely to assume that all of the options are trivial, leading to an assimilation effect. However, individuals who peruse such a list of privacy options more carefully are likely to recognize that the focal option is far more serious than the others, and the focal options is likely to appear especially serious in contrast. We therefore hypothesize that the propensity of participants to exhibit either effect will depend on the time taken by participants to make their selections:

H2a: For participants taking relatively shorter times to make their selections, presenting high relevance choices within a set of low relevance choices will result in an assimilation effect and diminish the focus on the relevant choice, and decrease the likelihood of participants choosing the privacy protective option.

H2b: For participants taking relatively longer time to make their selections, presenting high relevance choices within a set of low relevance choices will result in a contrast effect and increase the focus on the relevant choice, and increase the likelihood of participants choosing the privacy protective option.

Since we have no scientific basis for judging what is a short or long period of time to make such a selection, we operationalize it in the conventional fashion; by taking a median split on decision time.

3.3 Motivations to Share

Finally, we consider the propensity of variation in the presentation of privacy relevant choices to highlight motivation to share personal information. First, we extend our evaluation of the potential impact of labels

on decision frames by considering not only the impact of labels that highlight privacy concerns (e.g. "Privacy Settings"), but also the impact of labels which may highlight motivations to allow uses of personal information. Specifically, we hypothesize that:

H3a: Labeling otherwise identical choices as "Sharing Settings" relative to "Privacy Settings" will alter the decision frame in a manner that highlights motivations to allow uses of personal information, resulting in the choice of less protective options.

Moreover, we evaluate whether presenting a choice to allow use of personal information (an accept frame) vs. a choice to prohibit use of personal information (reject frame) can alter the decision frame to differentially highlight motivations to share personal information. Manipulations of accept/reject framing are common across privacy relevant contexts. For example, Figures 5 and 6 illustrates how privacy-relevant choices can be presented to consumers (sometimes simultaneously) as either as a choice to allow or restrict access to personal information.

[Figure 5: Facebook Custom Privacy]

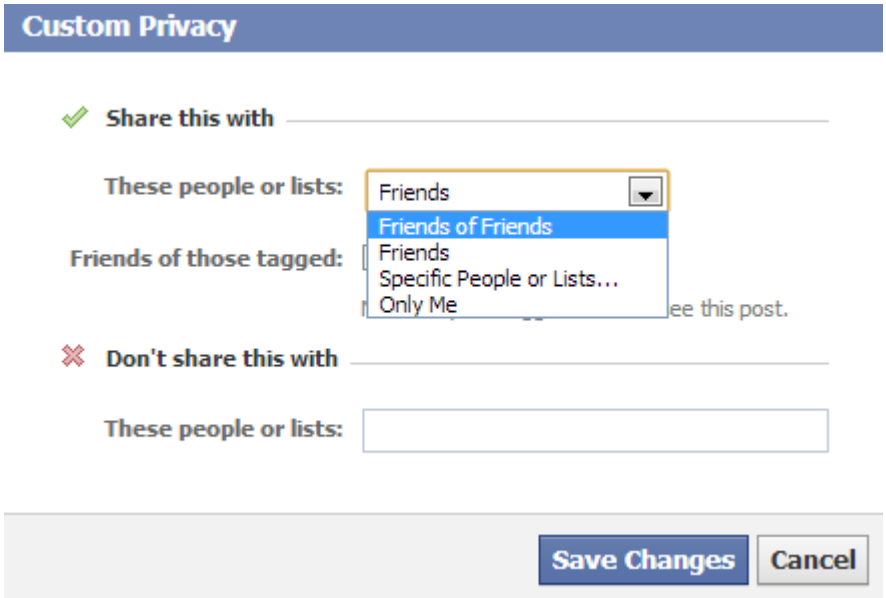

[Figure 6: Google+ Share Interface]

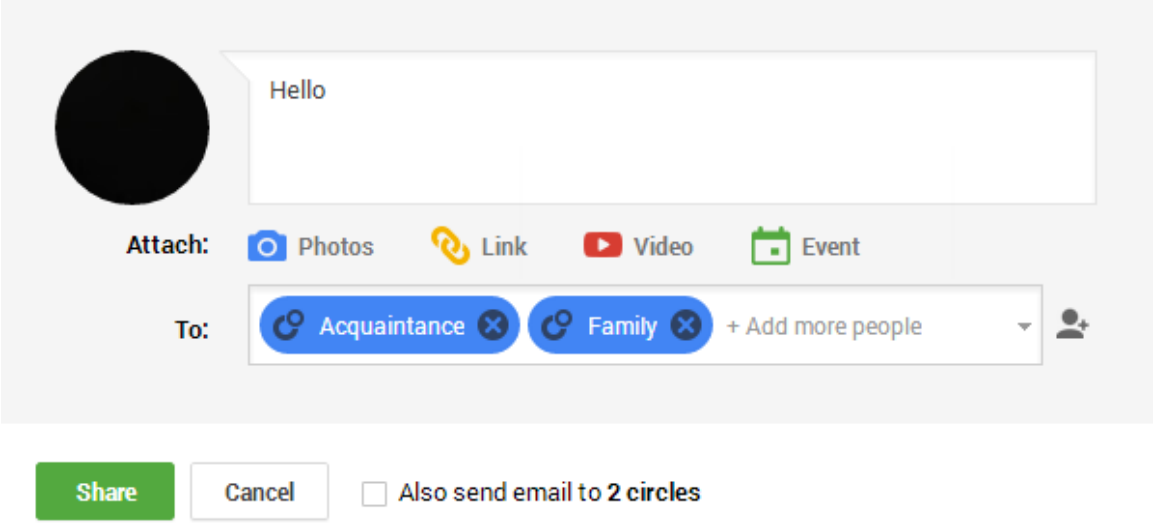

Prior research has found that differing response modes (e.g. an accept vs. reject presentation of a choice) can significantly influence decision making across contexts, particularly when these modes differentially highlight competing considerations or motives in choice contexts (Shafir, 1993). Johnson, Haibul, and Keinan (2007) proposed *Query Theory* as one explanation of why different response modes elicit variation in valuation and judgment. Broadly, they suggest that individuals execute a series of sequential queries (e.g. "What are the advantages of owning this product?" or "What are the disadvantages of owning this product?") to generate judgments, and the propensity of different response modes to impact the valence and ordering of these queries can generate variation in individual judgments of objectively identical options. For instance they demonstrate that the classic endowment effect (Kahneman, Knetsch, & Thaler, 1990) can be explained by differences in valence and ordering of the queries that participants make, and Appelt, Hardisty, and Weber (2011) find support for a Query Theory explanation of asymmetric discounting of gains and losses. More in line with our context, Hardisty, Johnson, Weber (2009) find that query theory can explain attribute framing effects, with differences in the valence and ordering of participants' queries explaining variation in participants' willingness to incur a cost presented as a "tax" vs. as an "offset". Specific to an accept/reject framing, Shafir (1993) posited, generally in line with the later theoretical work by Johnson et al (2007), that positive dimensions of choice weigh heavier

under an accept frame while negative dimensions of that same choice weight heavier under a reject frame. As a result we hypothesize that:

H3b: Framing choices as a choice to prohibit a use of personal information will highlight negative dimensions of such uses resulting in an increased propensity to choose protective options relative to framing choices as a decision to allow uses of personal information.

4. Empirical Approach

In three experiments, we examine how manipulations of choice framing can influence choice of privacy settings. All three experiments consisted in random-assignment online questionnaires conducted using Amazon's Mechanical Turk (AMT), an online service that, among other functions, connects researchers with potential participants and is becoming increasingly popular among social scientists conducting online experiments. All three experiments were advertised as studies on "ethical behavior." Participants were asked to choose between four settings which determined how their answers to the questionnaires would be stored, shared, and used.[8] In Experiments 1 and 3, participants were also asked to answer the sensitive questions about unethical behaviors.[9] All three experiments were two-factor between-subject designs, with the first manipulated factor, in all experiments, being the label of the choices presented. In Experiment 1, the second manipulated factor was the relevance of choices presented to participants. In Experiment 2, the second manipulated factor was whether privacy choices were presented as a homogeneous set of high relevance choices or a mixed relevance set. In Experiment 3, we also manipulated whether participants were presented the settings in an accept or reject frame.

[8] We restricted participants to subjects with over a 95% hit approval rate on AMT. We included attention check questions at the start of the questionnaire following accepted practices in the field (e.g. Oppenheimer, Meyvis, & Davidenko, 2009). We also included a screening survey which both prevented individuals from participants in a given experiment multiple times and prevented individuals from participating in more than one experiment.

[9] Because the central manipulation in experiment 2 involved manipulating choice sets, comparisons of disclosure behavior as a measure of objective risk between conditions was not meaningful. Thus, we opted to reduce participant risk and not ask for sensitive disclosures.

Across all experiments, the primary dependent variable was the choice of privacy protective options by participants. This measure was intended to capture the propensity of individuals to act on privacy sensitivities under different manipulations of choice framing: we analyze the effect of various framing manipulations on individuals' average propensity to select the privacy protective option. For experiments in which participants were asked to make multiple binary selections, we used a panel random effects linear probability model regression estimation approach to estimate the impact of choice frame manipulations on the propensity of participants to select the privacy protective option while correcting standard errors for the correlation between multiple responses from a single participant (Liang and Zeger, 1986).[11] As a secondary dependent variable, we also examined the actual participant admit rates to unethical behavior. We utilized this measure to evaluate whether disclosure behavior differed between manipulations of choice framing. Specifically, we investigated whether participants who chose less restrictive settings counteracted this behavior by sharing less sensitive information.

4.1 Pilot study

We conducted an initial pilot study with the goal of evaluating the perceived relevance of various uses of personal information that could be used in the context of the experiments in this manuscript. We recruited 104 participants (M_{Age} = 31 SD_{Age}=12.4, M_{Female} = .34) from AMT to complete a brief questionnaire that asked respondents to imagine that they were participating in a study on ethical behavior using the same introductory text provided to participants in the subsequent experiments. Participants were then asked to evaluate the extent to which they would want the choice to opt-out (or opt-in) of various uses and handling of their responses (see Appendix, Table 1). Participants were asked to rank each item on a 1-5 scale, with 1 being "Very Important" that they would be provided with the choice and 5 being "Very Unimportant" that they would be provided the choice. The four choices ranked as most important were

[11] Considering our two-factor between subjects design, we opted for a linear probability model estimation in lieu of a non-linear estimation approach (e.g. logit) to avoid troublesome coefficient and standard error estimates for interaction effects in non-linear regression models (Ai and Norton, 2003). Moreover, Angrist and Pischke (2008) have shown little qualitative difference between the Logit and linear probability specification.

considered "High Relevance," while the four choices ranked least important were deemed "Low Relevance." The importance-of-choice ranking by participants for high relevance choices was significantly different from the ranking for low relevance choices (2.47 vs. 2.79, $t(103)=3.98$, $p<.001$).[12]

5. Experiment 1

Experiment 1 evaluated whether the labeling of privacy choices can alter individuals' decision frames and thus affect their choice of privacy protective options and how this effect differed for high relative to low relevance choices.

5.1 Design and Procedure

The design was a 2 (Privacy Label, Descriptive Label) X 2 (High Importance, Low Importance) between subjects design. Between subjects, we manipulated whether a choice was designated with a "Privacy Settings" or "Survey Settings" label. We also manipulated, between subjects, the choice set presented to users. For the "High Relevance" conditions, participants were presented only the four settings ranked in the pilot study as most important (e.g. "Allow my responses to be shared with other participants of the study"), and conversely participants in the "Low Relevance" condition were provided only the four settings ranked least important in the pilot study (e.g. "Allow my responses to be used for academic publications")—See Table 2 in the Appendix. We hypothesized, in Section 3, that a privacy label would highlight privacy concerns for participants, resulting in more self-protective (information concealing) behavior (H1a) and that this effect would be more pronounced for high relevance settings (H1b).

Participants on AMT were invited to take an online study on ethical behavior for a payment of $.25. Participants were first asked demographic questions, which included no directly identifying

[12] The difference between the "High Relevance" and "Low Relevance" sets was more drastic (1.86 vs. 3.XX, p=?) when including only three choices within each set. Because we expected (and found evidence in support) that the inclusion of a fourth choice would only make our results more conservative, we opted to include it to make manipulations of mixed relevance more feasible.

information but asked for their city and zip of residence and other demographic information.[13] They were then provided with four choices that related to the use and storage of their responses in the survey. Finally, participants were presented with eight questions related to ethically questionable activities (See Appendix, Table 5). The questions used were the ones rated as intrusive in Acquisti, John and Loewenstein (2012), and were presented in random order. Finally, participants were asked a set of exit questions which evaluated, among other things, their satisfaction with the privacy protection provided and their perception of whether participating in the study would result in some subsequent harm to them.

5.2 Results

Two-hundred and four individuals (M_{Age} = 29 SD_{Age}=9.6, M_{Female} = .34) participated in the experiment. We found that, on average, participants presented with choices labeled "Privacy Settings" were 56% more likely to choose the more protective choices relative to those presented the same choices as "Survey Settings" (25% vs.16%, $t(202)=2.1729$, $p=.03$). Parsing between conditions with high and low importance settings (see Figure 7), we found support for *H1b:* label framing effects were driven by participants presented high relevance settings (42% vs. 28%, $t(99)=2.212$, $p=.03$). For low relevance setting choices, perhaps due to a floor effect, the effect of label was insignificant ("Privacy Settings": 7% vs. "Survey Settings": 4%, $t(101)=1.039$, $p=.3$). A random effects panel regression (Appendix, Table 6, Column 1) confirmed this finding with a negative and significant coefficient estimate ($\beta_{PrivacyLabel}=-.14$, $p=.017$) for the main effect of the "Privacy Settings" label (*H1a supported*). As expected, the interaction of *PrivacyLabel* and *LowImpt* was negative and of considerable magnitude, but was not statistically significant using a two-tailed test ($\beta_{PrivacyLabel\ *LowImpt}= -.10$, $p =.147$).

[13] These questions were intended to elicit a level of quasi-identifiability, such that participants would not perceive disclosure as being entirely risk-free. In exit questions, several participants commented that disclosing their geographic location did, in fact, make them uncomfortable in answering some of the questions on ethical behavior.

[Figure 7: Experiment 1 Summary Results]

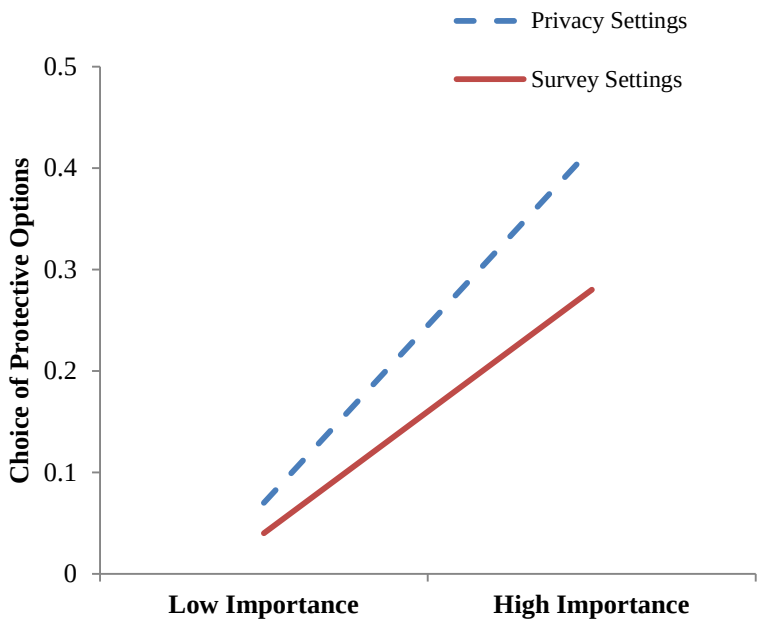

We did not find evidence that participants choosing less protective settings due to the framing manipulations counter-acted this behavior by disclosing less information: participants presented "Survey Settings" had comparable admit rates (percent of unethical behaviors admitted to) compared to those presented "Privacy Settings" (High Relevance Settings: 53.18% vs. 53.65%, $t(97)=.12$, $p=.92$; Low Relevance Settings: 51.25% vs. 55.04%, $t(97)=.893$, $p=.37$).This is again confirmed in our random effects panel regression with a near zero and highly insignificant estimate on the effect of *PrivacyLabel* on admit rates (Appendix, Table 6, Column 2). We also evaluated the impact of actually choosing privacy protection options on subsequent admit rates using the variable *TotalDeny* which captures the percent of settings for which a participant chose the privacy protective options (ranging from 0-1). We found that choosing a privacy protective option correlates with lower admit rates ($\beta_{TotalDeny}=-.10$, $p=.1$) counter to the expected causal effect of more protection resulting in *more* sensitive disclosures (Appendix, Table 8, Column 1) and that this correlation was unaffected by the framing manipulations, with an insignificant estimate on the interaction of *TotalDeny* and *PrivacyLabel* (Appendix, Table 8, Column 2). We suggest

that this result emerges due to unobserved individual difference factors (e.g. risk-aversion or underlying privacy preferences) which drive both lower admit rates and the choice of privacy protective options.

5.3 Discussion

The results of Experiment 1 provide evidence that subtle changes in the labeling of privacy relevant choices can significantly alter an individual's propensity to select protective options, particularly for high relevance privacy choices. Moreover, we do not find evidence that participants choosing riskier settings due to manipulations of choice framing made fewer sensitive disclosures to compensate for this additional risk. This result is consistent with the privacy control paradox identified by Brandimarte, Acquisti, and Loewenstein (2012) where increased control led to a false sense of security and elicited higher levels of disclosure, even when the provided control does not reduce objective risk. Because consumer privacy risk in many contexts (and in our experiment) is a function of both the allowances made via choice mechanisms and the actual information available about them (whether self-disclosed or collected otherwise) this finding has significant implications for consumer risk under privacy choice mechanisms. Specifically, it suggests that manipulations of privacy choice framing, and in particular instances of choice framing which do not highlight privacy concerns, are likely to result in elevated objective risk via a combination of increased allowances and a continued propensity towards sensitive disclosure by consumers. In this experiment, we focused on how choice-set level manipulations (i.e. choice set labels) impacted decision frames for all of the choices within that set. In contrast, Experiment 2 investigates the propensity of changes in the composition of the actual choice set to alter decision frames and impact the propensity of consumers to choose privacy protective options.

6. Experiment 2

Experiment 2 examined whether the mixture of high and low importance privacy choices can influence the choice of privacy protective options by evaluating the baseline effect of mixed relevance choice sets on participant privacy protective behavior for high relevance choices. Experiment 2 also extends our

result from experiment 1 by evaluating the robustness of the label framing effects in the more realistic context of mixed relevance choice sets.

6.1 Design & Procedure

The design was a 2 ("Privacy Settings", "Survey Settings") X 2 (Homogenous Importance, Mixed Importance) between subject design. Between subjects, we again manipulated whether choices were presented to users as "Privacy Settings" or as "Survey Settings". We also manipulated, between subjects, whether participants were presented with a homogenous set of high relevance choices similar to the conditions in Experiment 1 or a mixed set of three low relevance choices and one high relevance choice (Appendix, Table 3). The common choice between all conditions was whether participants would allow their responses to be shared with religious organizations.[14] As explained in Section 3, we hypothesized that mixed relevance choice sets may result in either a contrast (H2a) or assimilation effects (H2b), and that the propensity to identify either effect would be a function of time taken by participants to make their selections.

In Experiment 2, we evaluate the impact of presenting mixed importance settings on the propensity of individuals to choose privacy protective choices for high importance settings and any interaction of framing effect with mixed relevance settings. We also measured the length of time each individual took to make their choice of settings. Because Experiment 2 involved a mix relevance choice set, comparisons of disclosure behavior between conditions with varying framing manipulations was not meaningful. Thus, we opted to reduce participant risk and not ask for sensitive disclosures. The procedure for participants up until the sensitive disclosures (including their choice of settings) was identical to that of Experiment 1. However, in lieu of the page asking for sensitive disclosures, participants were presented

[14] This setting exhibited the strongest individual framing effect in Experiment 1.

an error message indicating that the page failed to load. They were then asked to continue with the study without answering the sensitive questions and without any penalty to them.[15]

6.2 Results

Five-hundred and twenty two individuals[16] (M_{Age} = 28 SD_{Age}=10.8, M_{Female} = .44) participated in the second experiment. We replicated the finding from Experiment 1. Participants presented with choices labeled "Privacy Settings" were significantly more likely to choose the protective option for the shared choice relative to those presented the same choices as "Survey Settings" (32% vs. 22%, $\chi^2(1)= 6.95$, p=.008). We found a roughly equivalent label framing effects for participants presented a homogenous set of high importance choices and those presented a mix of high and low relevance setting (10% vs. 11%). We also found that, on average, the assimilation effect seems to prevails: participants presented with the homogenous set of high importance choices were more than twice as likely to choose the privacy protective option for the shared choice (involving sharing with religious organizations), relative to participants presented the mixed relevance set of choices (36% vs. 17%, $\chi^2(1)= 25.74$, p<.0001). Given our hypothesis that the effect we would observe would vary depending on the time taken to make the choice of settings, we also parsed the data by the median time (a division, as noted earlier, selected only for its neutrality) taken by participants to make their selection. While not exogenously manipulated, the time taken to make the choice of settings was uncorrelated with our randomly assigned manipulations (i.e. label framing manipulations and the relevance of other choices in the set). This parsing reveals a more nuanced effect (see Figure 8). First, the impact of the homogenous set manipulation was driven strongly by participants who took less than the median time to make their selections (*H2a* supported). These participants were six times more likely to choose the protection option relative to their counter-parts presented the mixed importance set of choices (36% vs. 6%, $\chi^2(1)= 36.56$, p<.0001). The mixed set

[15] Participants reported in exit questions that they had received an error during the survey and that they would have wanted to fully complete the survey and answer the necessary questions. Again, this error was presented *after* the choice of settings had been elicited for all conditions.

[16] We collected more data for this study because we had one observation per participant as opposed to 4 in the prior study (only one setting was common across all conditions).

manipulation did not have a significant effect for participants who took longer than the median time to make their selections (37% vs. 30%, $\chi^2(1)= 1.38$, p=.286). While we do find that the assimilation diminishes for participants taking longer than the median time, we did not find support for a contrast effect (*H2b* not supported). Finally, participants in the homogenous set condition were equally likely to choose the protective option regardless of the time taken by individuals (37% vs. 36%).

[Figure 8: Experiment 2 Summary Results]

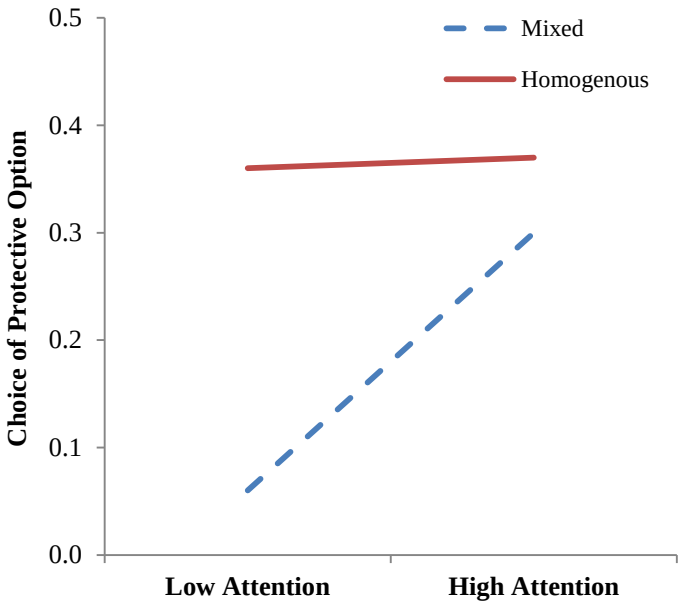

6.3 Discussion

The results of Experiment 2 reinforce the evidence that subtle manipulation of the framing of privacy choices may have a significant effect on the propensity of individuals to choose privacy protective option. First, we are able to replicate the result from Experiment 1 and also demonstrate that the label framing effect persists for high relevance choices even when presented as part of a mixed relevance set. More centrally, the results of Experiment 2 highlight the powerful impact of mixed relevance choice sets on the propensity of individuals to choose privacy protective options. While prior research finds that individual judgment and decision making is sensitive to comparative effects, we instead find evidence of an assimilation effect where participants were significantly *less* likely to select the privacy protective option for a high relevance choice when it was presented together within a set of low relevance choices. This

effect was particularly pronounced for participants who took less than the median time to make their selections. Because the time taken by participants was not exogenously manipulated, it may be the case that participants who took less time to make their privacy choices were simply less concerned about privacy and the sharing of their personal information. However, we find that in the homogenous choice set condition, participants were *equally* likely to select privacy protective option regardless of the time taken to make privacy selections, suggesting that this effect is driven by the choice framing and not unobserved factors. This last result, in particular, also highlights the potential for careful consideration of choice framing to bridge gaps in behavior for consumers who vary in their attention to privacy risks. Thus far, we've evaluated the impact on privacy decision making of variation in the labeling of the overall choice set presented to consumers (Experiment 1) and in this experiment variation in the composition of the choice set (Experiment 2). Finally, in Experiment 3, we restrict our focus to whether variation in the presentation of individual choices can also impact participant decision frames and in doing so the selection of privacy protective options.

7. Experiment 3

Experiment 3 evaluated the effect on choice of protective options of changes to the accept/reject frame privacy choices.

7.1 Design and Procedure

The design was a 2 ("Privacy Settings", "Sharing Settings") X 2 (Allow Settings vs. Prohibit Settings) between subject design. Between subjects, we manipulated whether choices were presented under a "Privacy Settings" or a "Sharing Settings" label. We also manipulated, between subjects, whether participants were presented the settings as a choice to allow a use of personal information or prohibit a use of personal information (settings were objectively identical). Again, the propensity of individuals to choose the privacy protective option was the dependent variable of interest. In this experiment we evaluated the baseline effect of an accept vs. reject frame on participant choice of privacy protective

options. Previously, we hypothesized (H3) that presenting a choice to prohibit a use of personal information will highlight negative dimensions of such uses resulting in an increased propensity to choose protective options relative to framing choices as a decision to allow uses of personal information. To avoid potentially confusing double negatives, we only presented participants the choice to restrict or accept affirmative statements. For instance, in Experiment 1, we provided participants a setting related to passwords. In this experiment, we instead use a setting focused on encryption as the "allow" framing avoided a potentially confusing double negative ("Allow my responses to be stored unencrypted" vs. "Allow my responses to be stored on a drive that is not password protected"). Finally, and in contrast to Experiment 1 and 2 which focused on comparisons of privacy labels and descriptive labels, Experiment 3 instead focuses on labels with the potential to highlight motivations to share personal information. Specifically, in lieu of the descriptive "Survey Settings" label we now label privacy choices as "Sharing Settings".

7.2 Results

Four-hundred and one individuals (M_{Age} = 28, SD_{Age}=10.8; M_{Female} = .44) participated in Experiment 3. We found support for *H3b;* participants in the "allow" condition were 51% more likely than those in the "prohibit" condition (56% vs. 37%) to choose the privacy protective option when presented a choice to prohibit a use of personal information vs. allow that same use ($t(399) = 4.658$, $p<.001$). A random effects panel regression (Appendix, Table 7, Column 1) confirmed this finding with a negative and significant estimate of the effect of the prohibit frame ($\beta_{Prohibit}$=-.16, p=.009). In fact, participants presented the prohibit frame were more likely to choose the privacy protective behavior than were participants in any condition across all three experiments.

However, we found that the "Sharing Settings" label did not exhibit the hypothesized effect (*H3a* not supported). First, participants in the prohibit frame did not exhibit any label framing effect and were equally likely to choose the privacy protective option irrespective of whether they were presented the Sharing or Privacy Settings label (55.8% vs. 56.3%, t(.088), p=.93). Participants in the allow frame

seemed to trend opposite the hypothesized results, being *more* likely (40% vs. 33%, $t(207)= 1.1.69$, $p=.24$) to choose the privacy protective option when presented the choices as "Sharing Settings" relative to "Privacy Settings" (See Figure 9), but the result is not statistically significant.

[Figure 9: Experiment 3 Summary Results]

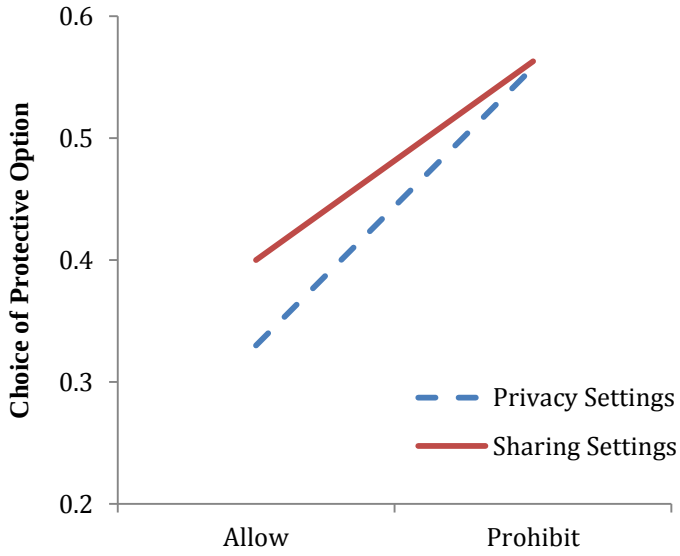

While this result was initially surprising, we suggest that it may point to the counterintuitive effect of attempts to present choices with significant privacy implications in a manner that is perceived as overtly promoting less protective behavior. Participants exposed to the "sharing" label may have felt that we attempted to manipulate them into picking less protective options which resulted in distrust and increased caution from participants. Alternatively, it is possible that the "sharing" label may have activated concerns about sharing because it activates privacy concerns in a fashion that "survey settings" does not. Similar to experiment 1, we find that manipulations of choice framing did not have a significant effect on participant admit rates (Appendix, Table 7, Column 2), that the choice of privacy protective options was again correlated with reduced participant admit rates ($\beta_{TotalDeny}$=-.06, p=.04), and that framing manipulations did not have an effect on the correlation between choice of privacy protective options and sensitive disclosures (Appendix, Table 8, Column 4).

7.3 Discussion

The results of Experiment 3 again highlight the significant role of choice framing in driving individual privacy behavior: presenting privacy settings as a choice to restrict versus allow uses of personal information elicited stark differences in the choice of protective options. In fact, presenting the choice in a manner that highlights the rationale to restrict uses of personal information (i.e. the reject frame) resulted in the highest levels of privacy protective behavior of any manipulation across the three experiments in this manuscript. However, this experiment also highlights the potential bounds for choice framing. First, we find that a label framing effect could be meaningful at the margins and not have an effect when other factors are already eliciting consumers to be protective (participants presented the prohibit frame did not exhibit any label effect). Moreover, we illustrate that unintended consequences may occur from attempts to highlight motivations to share personal information. Specifically, this may result in reduced confidence and trust from consumers and a trend towards more protective choices and less disclosure. This has significant implications for firms that attempt to highlight such motivations when presenting consumers with choices regarding the collection and use of personal information.

8. Discussion: Tradeoffs of Privacy Protection and Choice Framing

Privacy protective behavior is not without its costs: consumers choosing more restrictive data settings may do so at the expense of valuable online services, product customization, or tailored advertising and promotions. As a result, the challenge remains to design choice architectures in privacy contexts that balance consumer privacy considerations against competing utility gains from data allowances. This challenge reflects a broader controversy in the economics literature on choice architecture with considerable disagreement over what criteria are appropriate for identifying when a behavioral intervention is appropriate and the measures for evaluating the effectiveness of these interventions (Loewenstein and Haisley, 2007; Bernheim and Rangell, 2007; Sugden, 2008). In our experimental context, the specific concern is that, while some framing manipulations (e.g. descriptive choice labels) result in less privacy protective behavior from participants, this may, in fact, be desirable if it is better

aligned with participants' own stated preferences. While aligning behavior with stated preferences is one of many imperfect criterion for justifying behavioral interventions[17], the misalignment of consumer stated preferences for privacy and actual behavior, known as the Privacy Paradox, is well documented in the privacy literature (Jensen, Potts, and Jensen, 2005) as a challenge for policy makers and firms in managing consumer privacy concerns. Moreover, Loewenstein & Haisly (2007) suggest that contexts where individuals need help *self-officiating* (i.e. acting on their stated preferences or desired goals) are practical contexts in which behavioral interventions could be more easily justified.

To address this concern, we recruited 96 participants from Amazon's Mechanical Turk (M_{Age} = 30 SD_{Age}=10.6, M_{Female} = .34) to complete a questionnaire which, similar to our pilot study, provided a description of our experimental context and asked them to imagine that they were participating in such a study on ethical behavior. They were then asked to report their propensity to choose the privacy protective option for each use of their data presented (in our pilot study a separate sample of Mechanical Turk users were asked only to report their perceived importance/relevance of each choice in our study). As expected, we found that choice relevance was a good predictor of participants' stated preferences for the privacy protective option; participants were significantly more likely to suggest that they would select the privacy protective option for high relevance choices than for low relevance choices (49% vs. 21%, $t(95)=9.73$, $p<.001$). More importantly, we found that, while our studied manipulations resulted in significant *relative* increases in the choice of privacy protective options, the absolute level of choice for privacy protective options was, in fact, lower than the levels predicted by participants in in two of the three experiments, and was always closest to predicted level relative to other framing manipulations (See Figure 10).

[17] For instance, individuals may not consistently anticipate their future behavior and elicitations of hypothetical or predicted preferences may be subject to the same framing effects as actual behavior.

[Figure 10: Stated Preferences vs. Behavior]

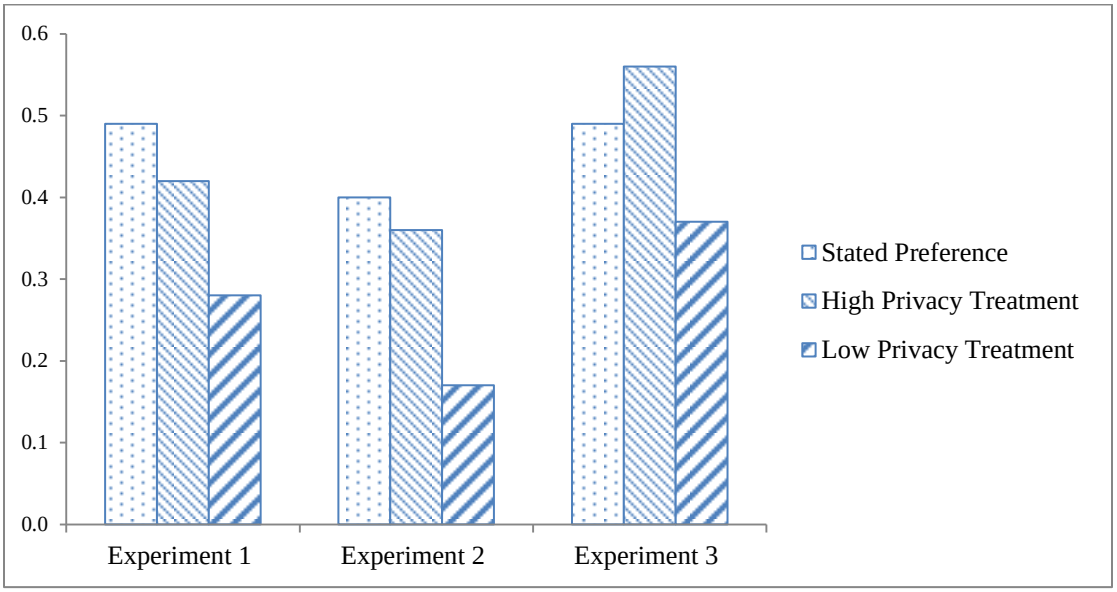

**For ease of comparison we only included high relevance settings for Experiment 1

9. Conclusion

This manuscript presents evidence that subtle real-world heterogeneity in the presentation of privacy-relevant choices can trigger or dull consumer privacy concerns, and significantly impact consumers' choices of privacy settings. Across three experiments, the choice of the privacy protective options for objectively identical choices ranged from 17% to 55% depending on choice framing. These results raise significant concerns about proposed policy approaches to alleviate consumer privacy concerns which center on giving consumers more choice, potentially at the expense of supporting consumer protections (e.g. data collection limitation): a recent World Economic Forum Report titled "Unlocking the Value of Personal Data: From Collection to Usage" suggests that new technological options can give individuals control over their own information while allowing data assets to flow relatively freely (World Economic Forum, 2013) and a senior advisor for a large technology firm (and a contributor to the report) recently stated that "There's no bad data, only bad uses of data."

One may argue, given the considerable value that firms hope to derive from the collection and use of consumer personal information, that it would be, at a minimum, surprising if they did not strategically leverage subtle variations in choice framing (as have firms in other industries) to elicit greater allowances

from consumers via these proposed control mechanisms. Moreover, market forces which have restricted other controversial data practices by firms (e.g. consumer privacy concerns relating to behavioral advertising have limited its use-Ponemon Institute, 2010) are unlikely to counteract this practice. The subtle manipulations we study are unlikely to be noticed by consumers, and, if they are noticed their impact on behavior is unlikely to be appreciated by consumers (Liberman, Samuels, and Ross, 2004). The concerns raised in this manuscript reinforce those raised by prior work demonstrating that increased control can result in a false sense of security leading to risky subsequent disclosure behavior (Brandimarte, Acquisti, and Loewenstein, 2012). This could explain why participants who chose ostensibly riskier options due to framing manipulations continued to disclose at levels comparable to their counterparts who chose more protective options.

Alternate policy approaches would include privacy control mechanisms as a subset of privacy protections afforded to consumers. For instance, regulators may first consider simply restricting data practices perceived to be particularly intrusive or potentially harmful to consumers and uniform standards for soliciting consumer choice in emerging privacy contexts where consumer choice is desired. This latter recommendation has precedent in other contexts (e.g. healthcare or finance) in which regulators have provided standardized formats for soliciting consumer consent. In contrast to the current privacy choice mechanisms available to consumers, privacy choice mechanisms for emerging data practices by firms may be informed by a growing literature in behavioral economics focusing on designing choice architectures that aide consumer in improved decision making. These insights have been applied to other contexts by high level policy units (e.g. the UK behavioural insights team) with considerable success, and could include framing choice to properly highlight both costs and benefits stemming from the collection and use of personal information and manipulation of choice defaults. These suggestions could limit firms' abilities to manipulate consumers in their own interests while empowering consumers to make choices that reflect their desired balance of personal privacy and utility from uses of their personal information.

The poet Robert Frost described "consent in all forms" as the "strongest and most effective force in guaranteeing the long-term maintenance of power" where the "dominated acquiesce in their own

domination". In line with this notion, if choice mechanisms are not carefully crafted and provided alongside supplemental protections, they may have largely the effect of quelling consumer privacy concerns via the *opportunity* to restrict the collection and use of their personal information while, in practice, actually resulting in most consumers continuing to provide broad, and potentially harmful, data allowances to firms.

10. References

1. Acquisti, A., John, L. K., & Loewenstein, G. (2012). The impact of relative standards on the propensity to disclose. Journal of Marketing Research, 49(2), 160-174.
2. Adjerid, I., Acquisti A, Brandimarte, L., and Loewenstein, G. (2013). Sleights of privacy: framing, disclosure, and the limits of transparency. Proceedings of the 2013 Symposium on Usable Privacy and Security (SOUPS 2013), ACM.
3. Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. Economics letters, 80(1), 123-129.
4. Angrist, J. D., & Pischke, J. S. (2008). Mostly harmless econometrics: An empiricist's companion. Princeton university press.
5. Bernheim, B. D., & Rangel, A. (2007). Toward choice-theoretic foundations for behavioral welfare economics. The American economic review, 464-470.
6. Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced confidences privacy and the control paradox. Social Psychological and Personality Science, 4(3), 340-347.
7. Burnham, T., McCabe, K., & Smith, V. L. (2000). Friend-or-foe intentionality priming in an extensive form trust game. Journal of Economic Behavior & Organization, 43(1), 57-73.
8. Chavez, P. (2011). Re: Preliminary FTC Staff Report on "Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers". Google. http://www.ftc.gov/os/comments/privacyreportframework/00417-58065.pdf
9. Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. Information Systems Research, 17(1), 61-80.
10. Fogel, J., & Nehmad, E. (2009). Internet social network communities: Risk taking, trust, and privacy concerns. Computers in Human Behavior, 25(1), 153-160.
11. Ganzach, Y., & Karsahi, N. (1995). Message framing and buying behavior: A field experiment. Journal of Business Research, 32(1), 11-17.
12. Genesove, D., & Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. The Quarterly Journal of Economics, 116(4), 1233-1260.
13. Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. Journal of consumer research, 9(1), 90-98.
14. Jensen, C., Potts, C., & Jensen, C. (2005). Privacy practices of Internet users: self-reports versus observed behavior. International Journal of Human-Computer Studies, 63(1), 203-227.
15. Johnson, E. J., Shu, S. B., Dellaert, B. G., Fox, C., Goldstein, D. G., Häubl, G., ... & Weber, E. U. (2012). Beyond nudges: Tools of a choice architecture. Marketing Letters, 23(2), 487-504.

16. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica: Journal of the Econometric Society, 47(2), 263-291.
17. Keller, P. A., Harlam, B., Loewenstein, G., & Volpp, K. G. (2011). Enhanced active choice: A new method to motivate behavior change. Journal of Consumer Psychology, 21(4), 376-383.
18. Kühberger, A. (1998). The influence of framing on risky decisions: A meta-analysis. Organizational behavior and human decision processes, 75(1), 23-55.
19. Levin, I. P., & Gaeth, G. J. (1988). How consumers are affected by the framing of attribute information before and after consuming the product. Journal of Consumer Research, 15(3), 374-378.
20. Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. Organizational behavior and human decision processes, 76(2), 149-188.
21. Liberman, V., Samuels, S. M., & Ross, L. (2004). The name of the game: Predictive power of reputations versus situational labels in determining prisoner's dilemma game moves. Personality and social psychology bulletin, 30(9), 1175-1185.
22. Loewenstein, G., & Haisley, E. C. (2007). The economist as therapist: Methodological ramifications of'light'paternalism. Available at SSRN 962472.
23. Maheswaran, D., & Meyers-Levy, J. (1990). The influence of message framing and issue involvement. Journal of Marketing research, 27(3), 361-367.
24. Mas, A. (2006). Pay, reference points, and police performance. The Quarterly Journal of Economics, 121(3), 783-821.
25. Moon, Y. (2000). Intimate exchanges: Using computers to elicit self-disclosure from consumers. Journal of Consumer Research, 26(4), 323-339.
26. Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. Journal of Experimental Social Psychology, 45(4), 867-872.
27. Ponemon Institute. (2010). Economic impact of privacy on online behavioral advertising. http://www.ponemon.org/library/economic-impact-of-privacy-on-online-behavioral-advertising
28. Richter, M. (2011). Re: preliminary FTC staff report on "protecting consumer privacy in an era of rapid change: a proposed framework for businesses and policymakers". Facebook. http://www.ftc.gov/os/comments/privacyreportframework/00413-58069.pdf
29. Rothman, A. J., Salovey, P., Antone, C., Keough, K., & Martin, C. D. (1993). The influence of message framing on intentions to perform health behaviors. Journal of Experimental Social Psychology, 29(5), 408-433.
30. Schwartz, P. M. (2005). Privacy inalienability and the regulation of spyware. Berkeley Technology Law Journal, 20, 1269.
31. Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. Cognition, 49(1), 11-36.
32. Simonson, I., & Tversky, A. (1992). Choice in context: tradeoff contrast and extremeness aversion. Journal of Marketing Research, 29(3), 281-295.
33. Singel, R. (2012). IE 10′s 'Do-Not-Track' default dies quick death. Wired. http://www.wired.com/threatlevel/2012/06/default-do-not-track/
34. Singer, N. (2013). A data broker offers a peek behind the curtain. New York Times. http://www.nytimes.com/2013/09/01/business/a-data-broker-offers-a-peek-behind-the-curtain.html?_r=0

35. Solove, D. J. (2013). Privacy self-Management and the consent dilemma. Harvard Law Review, 126, 1879-2139.
36. Sugden, R. (2008). Why incoherent preferences do not justify paternalism. Constitutional Political Economy, 19(3), 226-248.
37. Sweeney, L. (2013). Discrimination in online ad delivery. Queue, 11(3), 10.
38. The Federal Trade Commission (FTC). (2012). Protecting consumer privacy in an era of rapid change: recommendations for businesses and policy makers. http://www.ftc.gov/os/2012/03/120326privacyreport.pdf.
39. The White House. (2012). Consumer data privacy in a networked world. http://www.whitehouse.gov/sites/default/files/privacy-final.pdf.
40. Tversky, A., Kahneman, D., & Choice, R. (1981). The framing of decisions. Science, 211, 453-458.
41. World Economic Forum. (2013). Unlocking the value of personal data: from collection to usage.http://www3.weforum.org/docs/WEF_IT_UnlockingValuePersonalData_CollectionUsage_Report_2013.pdf
42. Zeger, S. L., & Liang, K. Y. (1986). Longitudinal data analysis for discrete and continuous outcomes. Biometrics, 42(1), 121-130.

9. Appendix: Experimental Materials and Regression Results

[Table 1: Relevance Ranking of Uses of Personal Information]

Choice	Description	Mean Relevance
1	Allow my responses to be shown to other participants of this study.	2.37
2	Allow my responses to be shared with religious organizations interested in evaluating personal ethics.	2.46
	Allow my responses to be published on a research bulletin openly available on the internet.	2.47
4	Store my responses only on a password-protect drive.	2.61
_____	Store my responses only on an encrypted drive.	2.64
6	Allow other research groups (beyond the group conducting this study) to access and analyze my responses.	2.65
7	Allow my responses to be shared with various think tanks that focus on ethics.	2.68
8	Allow my responses to be stored beyond the completion of this study. This would allow us to use your responses in future studies and analysis.	2.73
9	Allow research assistants (these are students that aid in research but are not faculty or PhD candidates) to access your responses.	2.83
10	Allow you responses to be used in academic publications.	2.91

[Table 2: High vs. Low Relevance Conditions]

Choice	Description	Condition
1	Allow my responses to be shown to other participants of this study.	High
2	Allow my responses to be published on a research bulletin openly available on the internet.	High
	Allow my responses to be shared with religious organizations interested in evaluating personal ethics.	High
4	Store my responses only on a password-protect drive.	High
5	Allow my responses to be shared with various think tanks that focus on ethics.	Low
6	Allow my responses to be stored beyond the completion of this study. This would allow us to use your responses in future studies and analysis.	Low
7	Allow research assistants (these are students that aid in research but are not faculty or PhD candidates) to access your responses.	Low
8	Allow you responses to be used in academic publications.	Low

[Table 3: Homogenous vs. Mixed Conditions]

Choice	Description	Condition
1	Allow my responses to be shown to other participants of this study.	Homogenous
2	Allow my responses to be published on a research bulletin openly available on the internet.	Homogenous
3	Allow my responses to be shared with religious organizations interested in evaluating personal ethics.	Homogenous
4	Store my responses only on a password-protect drive.	Homogenous
5	Allow research assistants (these are students that aid in research but are not faculty or PhD candidates) to access your responses.	Mixed
6	Allow you responses to be used in academic publications.	Mixed
7	Allow my responses to be shared with religious organizations interested in evaluating personal ethics.	Mixed
8	Allow my responses to be stored beyond the completion of this study. This would allow us to use your responses in future studies and analysis.	Mixed

[Table 4: Allow vs. Prohibit Condition]

Choice	Description	Condition
1	Allow my responses to be shown to other participants of this study.	Allow
2	Allow my responses to be published on a research bulletin openly available on the internet.	Allow
	Allow my responses to be shared with religious organizations interested in evaluating personal ethics.	Allow
4	Allow my responses to be stored unencrypted	Allow
5	Prohibit my responses from being shown to other participants of this study.	Deny
6	Prohibit my responses from being published on a research bulletin openly available on the internet.	Deny
7	Prohibit my responses from being shared with religious organizations interested in evaluating personal ethics.	Deny
8	Only store my responses on an encrypted drive	Deny

[Table 5: Ethical Questions]

Choice	**Description**
1	Have you ever used drugs of any kind (e.g. weed, heroin, crack)?
2	Have you ever let a friend drive after you thought he or she had had too much to drink?
	Have you ever made up a serious excuse, such as grave illness or death in the family, to get out of doing something?
4	Have you ever stolen anything worth more than $100?
	Have you ever had sex in a public venue (e.g. restroom of a club, airplane)?
6	Have you ever fantasized about doing something terrible (e.g. torture) to someone?
7	Have you ever looked at pornographic material?
8	Have you ever downloaded a pirated song from the internet?

[Table 6: Experiment 1 Regression Results]

	Experiment 1	
	(1)	(2)
VARIABLES	Deny	Admit
PrivacyLabel	0.137*	-0.00460
	(0.0614)	(0.0445)
LowRel	-0.242**	0.0140
	(0.0463)	(0.0431)
PrivacyLabel*LowRel	-0.102	-0.0334
	(0.0700)	(0.0612)
Constant	0.280**	0.536**
	(0.0424)	(0.0284)
Observations	816	1,608
Number of Groups	204	201

Robust standard errors in parentheses
** p<0.01, * p<0.05, + p<0.1

[Table 7: Experiment 3 Regression Results]

	Experiment 3	
	(1)	(2)
VARIABLES	Deny	Admit
PrivacyLabel	0.0684	0.0150
	(0.0583)	(0.0303)
Prohibit	-0.158**	0.0375
	(0.0599)	(0.0313)
PrivacyLabel*Prohibit	-0.0737	-0.0249
	(0.0830)	(0.0451)
Constant	0.599**	0.432**
	(0.0425)	(0.0211)
Observations	1,604	3,200
Number of Groups	401	400

[Table 7: Choice of Setting on Disclosure]

	Experiment 1		Experiment 3	
	(1)	(2)	(3)	(4)
VARIABLES	Admit	Admit	Admit	Admit
Deny	-0.0994+	-0.144+	-0.0547*	-0.0621+
	(0.0600)	(0.0842)	(0.0263)	(0.0362)
PrivacyLabel		-0.0271		-0.00517
		(0.0360)		(0.0325)
PrivacyLabel*Deny		0.0796		0.0152
		(0.118)		(0.0527)
Constant	0.553**	0.565**	0.476**	0.479**
	(0.0181)	(0.0252)	(0.0163)	(0.0229)
Observations	1,608	1,608	3,200	3,200
Number of Groups	201	201	400	400