

A Query-Theory Perspective of Privacy Decision Making

Idris Adjerid, Sonam Samat, and Alessandro Acquisti

ABSTRACT

Long-standing policy approaches to privacy protection are centered on consumer notice and control and assume that privacy decision making is a deliberative process of comparison between costs and benefits from information disclosure. An emerging body of work, however, documents the powerful effects of factors unrelated to objective trade-offs in privacy settings. In this paper, we investigate how focusing on the process by which individuals make privacy choices can help explain the impact of rational and behavioral factors on privacy decision making. In an online experiment, we borrow from query-theory literature and measure individuals' considerations (that is, queries) across manipulations of rational and behavioral factors. We find that effects of rational and behavioral factors are associated with differences in the order and valence of queries considered in privacy settings. Our results confirm that understanding how differences in privacy choice emerge can help harmonize disparate perspectives on privacy decision making.

1. INTRODUCTION

In the United States, a long-standing approach to consumer privacy protection has relied on self-regulatory regimes centered on notice and consent mechanisms (FTC 2012; White House 2012). The premise of this approach is that individuals, once provided with sufficient information and

IDRIS ADJERID is an Assistant Professor of Information Technology, Analytics, and Operations at the Mendoza College of Business, University of Notre Dame. SONAM SAMAT is a Ph.D. student at the Heinz College, Carnegie Mellon University. ALESSANDRO ACQUISTI is a Professor of Information Technology and Public Policy at the Heinz College, Carnegie Mellon University. The authors gratefully acknowledge research support from the following organizations: National Science Foundation (awards CNS-1012763, SMA-1327992, and SES-1514192), US 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 (award CCF 0424422). In addition, Acquisti gratefully acknowledges support from the Alfred P. Sloan Foundation and from the Carnegie Corporation of New York via an Andrew Carnegie Fellowship.

[*Journal of Legal Studies*, vol. 45 (June 2016)]

© 2016 by The University of Chicago. All rights reserved. 0047-2530/2016/4502-0020\$10.00

control, will be able to make informed choices between different market offerings, managing privacy and disclosure in accordance with stable, coherent personal preferences. Such a premise is consistent with a widely held view of privacy decision making as a deliberate and rational process (Dinev and Hart 2006; Milne and Gordon 1993).

Over the past decade, the validity of this premise and its underlying assumptions has been questioned in legal, policy, and behavioral circles (Ben-Shahar and Schneider 2014; Solove 2013). In particular, an emerging stream of behavioral work on privacy decision making (Acquisti 2004) has highlighted the decision-making hurdles that consumers face when actuating privacy preferences into market behaviors. This stream of work has suggested that factors independent of variations in consumer preferences and objective trade-offs may still significantly influence consumers' behavior (Acquisti, Brandimarte, and Loewenstein 2015).

Recent research efforts have attempted to harmonize these ostensibly disparate perspectives on consumers' privacy decision making. Dinev, McConnell, and Smith (2015) and Adjerid, Peer, and Acquisti (2016) have proposed analytical frameworks that integrate the rational view of consumers' privacy decision making, which stresses the role of objective trade-offs, and the behavioral view, which stresses the role of heuristics and biases. In this paper, we extend these efforts and investigate how the process by which consumers make privacy choices can provide a common basis for rational and behavioral privacy responses.

We evaluate how changes in privacy choices that result from both behavioral and rational factors also coincide with changes in the prominence of competing considerations in privacy decision contexts. This notion is consistent with query theory (QT), a leading information-processing theory emerging from the psychology and behavioral economics literature (Johnson, Häubl, and Keinan 2007). Query theory posits that, when generating judgments in decision contexts, individuals execute a series of sequential queries (for example, What are the advantages of owning this product? or What are the disadvantages of owning this product?). Query theory suggests that behavioral reactions to different features of decision contexts (for example, choice frames) emerge because of changes in the valence and order of queries that individuals themselves generate in these settings. Applied to privacy decision making, QT would suggest that individuals decompose privacy choices into a series of Why should I disclose personal information? queries versus Why should I not disclose personal information? queries. Critically, QT posits that factors that alter the va-

lence and order of these queries, independent of whether they are rational or behavioral, will likely influence individuals' choices.

In this paper, we present the results of an experiment we conducted to evaluate how individuals' choices regarding personal data uses and permissions can be affected by rational and behavioral factors and whether, in fact, changes in choices due to rational and behavioral factors emerge in a manner consistent with QT accounts of decision making.

In the experiment, we asked participants to make a choice of whether to share sensitive information with an outside entity—a scenario similar to the selection of privacy settings provided by several online services and advertisers. Thus, the experiment focuses on a privacy choice context that is increasingly relevant nowadays, as consumers' information is often collected passively (for example, data collected by browsers via cookies) rather than through active self-disclosure (for example, sharing information on a social network). In these contexts, consumers' choices of privacy settings (for example, choosing “do not track” browser options) are their main recourse against data collections they may deem intrusive. Using this experimental setting, we evaluated whether individuals' privacy choices are susceptible both to rational and to behavioral factors. Using methods common in the QT literature, we also asked participants to autonomously generate their own reasons (that is, queries) either in favor of or against sharing their personal information.

We evaluated the impact of rational factors by altering the entities with which participants were provided the option to share their data that were collected in the experiment (either marketing companies or other research organizations). In other words, we manipulated the risk associated with deciding to disclose personal information. We evaluated the impact of behavioral factors by altering whether data requests were framed as a choice to allow a data use (the accept frame) or a choice to prohibit the same data use (the reject frame). We evaluated whether changes in privacy choices due to these factors are consistent with QT by analyzing the valence and sequence of queries across the rational and behavioral manipulations.

We found that both rational and behavioral factors have an impact on behavior: participants were nearly twice as likely to share their information with other research organizations relative to marketing companies but were also significantly impacted by the framing of the choice (55 percent who agreed to the data use in the allow frame versus 40 percent who permitted the data use in the reject frame). We also found that partici-

pants shown the low risk of data use and those in the allow framing condition reported more queries in favor of sharing personal information relative to those in the high-risk and reject framing conditions, respectively. Similar effects were observed for the order of queries in favor of sharing personal information. Overall, we find evidence that QT can in fact provide one common process for how rational and behavioral effects emerge.

Our results have a number of implications for research and models of individuals' privacy decision making. First, we show that changes in privacy choices caused by ostensibly different factors (rational versus behavioral) are consistent with a QT account of decision making. This bolsters the emerging literature on QT. In particular, to our knowledge, prior work on QT has focused largely on explaining deviations from rational choice models (which we also confirm) but has not directly examined the viability of the theory in explaining rational responses in decision contexts. Second, and more generally, our results highlight the importance of carefully evaluating not only which types of factors (for example, behavioral versus rational) impact privacy decision making but also how these effects emerge from a decision-making-process perspective. In particular, our results suggest that while the effects on privacy decision making can emerge because of diverging factors, how these effects emerge could be reasonably consistent. This is relevant for research and policy efforts focused on nudging individuals toward better privacy decision making by, for example, enhancing individuals' response to rational factors and diminishing their response to behavioral ones.

2. THEORETICAL BACKGROUND

For many years, a widespread view of privacy decision making has been predicated on the notion of a privacy calculus (Dinev and Hart 2006; Milne and Gordon 1993). Under that view, privacy choice is seen as a process driven by rational choices aimed at maximizing the utility an individual derives from information disclosure and information protection. For instance, legal scholars have articulated views of consumers as shrewd privacy balancers who weigh the value to themselves and society of calls for personal information (Westin 2000). Similarly, seminal works on the economics of privacy (Posner 1981; Stigler 1980) have assumed that consumers are deliberative agents who rationally want to disclose positive information about themselves and hide negative information. Under this view, the individual is assumed to have stable and consistent

preferences about privacy—a perspective perpetuated in more recent economic works on privacy, such as the microeconomic analysis of intertemporal trade-offs that arise when merchants acquire information about consumers' preferences (for example, Taylor 2004). Even biologists have suggested that privacy is subject to interpretation in “economic terms” and that it will persist as long as it provides “profitable cost margins” (Klopfer and Rubenstein 1977, p. 64).

In the last decade, however, an emerging body of work has questioned those assumptions. Factors that should logically influence privacy behavior have been found sometimes not to do so. For instance, Kugler and Strahilevitz (2015) find that the duration of surveillance—a factor presumably relevant to individuals' odds of suffering harm from such surveillance—has no impact on individuals' expectation of privacy and has limited effects on the perceived intrusiveness of the surveillance. And Tourangeau (2004) finds that, despite online responses being more likely to be tracked and disseminated, individuals disclose more personal information when solicited through an online form compared with a pencil and paper questionnaire. On the other hand, factors with ostensibly little (and even no) direct impact on objective risks and benefits from disclosure have been shown to predictably and powerfully impact individuals' privacy choices. For example, individuals' intimate disclosures seem to be impacted by a survey's look and feel, subtle variation in the framing of disclosure decisions, choice defaults, and relative judgments (John, Acquisti, and Loewenstein 2011; Acquisti, John, and Loewenstein 2012). Overall, this body of work has highlighted that consumers' privacy preferences may be malleable and that individuals' behavior may be affected by factors with little relationship to objective trade-offs (Acquisti, Brandimarte, and Loewenstein 2015).

With the emergence of robust empirical support for ostensibly disparate views of privacy decision making, scholars have recently attempted to harmonize these perspectives by modifying the assumptions of overarching privacy decision-making models (Dinev, McConnell, and Smith 2015) or by considering how differing empirical approaches in the privacy literature may contribute to simultaneous support for both perspectives (Adjerid, Peer, and Acquisti 2016). In particular, Dinev, McConnell, and Smith (2015) have proposed an expansion of the antecedents–privacy concerns–outcomes model of privacy decision making that incorporates low-effort cognitive responses motivated by frameworks in behavioral economics and psychology. And Adjerid, Peer, and Acquisti (2016) have argued that the

role of rational and behavioral perspectives of privacy decision making can be explained by consumers' overestimation of their response to rational factors in actual choice settings and an underestimation of their response to behavioral factors in hypothetical choice settings.

In this paper, we take a different approach. We focus on how (as opposed to whether) behavioral and rational factors result in changes in privacy choices, in an effort to reconcile the emergence of different effects on individuals' behavior. A focus on how different behavioral responses emerge has proven useful in other contexts. Recent work in economics focuses on how emotional states correspond to well-studied decision biases (and lack thereof). For example, Breaban and Noussair (2013) find that greater positive emotions predict higher prices and larger bubbles in asset markets—generally attributed to the irrationality that is associated with momentum trading. Nguyen and Noussair (2014) find that a more positive emotional state is positively correlated with greater risk taking and that fear, happiness, anger, and surprise are positively correlated with risk aversion. The extant privacy literature has also focused on the process of privacy decision making in order to alleviate privacy concerns. For instance, Petronio (2015) suggests that privacy decision making is a boundary management process by which individuals allot coownership of their personal information. She argues that privacy concerns emerge when coowners fail to negotiate or follow these terms. Hann et al. (2007) take an information-processing theory approach to privacy decision making and conjecture that individuals form expectations and make privacy decisions that are based on how they process information about behavior-outcome relationships in a given context. They find that privacy protections reduce privacy concerns via their effect on the valence of factors that individuals consider (for example, the convenience the service provides versus the potential harms of data disclosure).

We focus on how QT, a leading information-processing theory of behavior, can provide a simplified, process-centric view of privacy decision making that is predictive of ostensibly different behavioral responses in privacy settings. Consistent with an information-processing view of decision making that starts from basic cognitive building blocks and focuses on how “decision making recruits basic processes from memory, attention, and perception” (Oppenheimer and Kelso 2015, p. 283), QT argues that individuals execute a series of sequential queries (for example, What are the advantages of owning this product? or What are the disadvantages of owning this product?) to generate judgments in decision con-

texts (Johnson, Häubl, and Keinan 2007). Leaning on output interference theory (Dempster 1995), Johnson, Häubl, and Keinan (2007) conjecture that the valence and order of queries have a significant impact on observed behavior (retrieval is more successful for earlier queries, which results in a stronger impact on behavior of these queries). The QT literature uses an aspect-listing task that requires participants to list the things they were considering as they made a particular judgment. This task is meant to elicit participants' own considerations about a certain decision context and is treated as an approximation of the implicit queries that arise as respondents make a judgment.

Query theory has been proposed as a generalizable theory of decision making and has already been tested across a series of different decision settings, including willingness to pay for an item, intertemporal trade-offs, and preferences for a carbon tax as an environmental intervention (Johnson, Häubl, and Keinan 2007; Weber et al. 2007; Hardisty, Johnson, and Weber 2010). We argue that QT may be highly relevant to privacy decision contexts. The extant privacy literature has documented powerful, competing motives for both openness and protection in privacy settings (Featherman and Pavlou 2003). As such, it is plausible that privacy choices are decomposed by consumers into Why should I disclose personal information? queries and Why should I not disclose personal information? queries. In practice, queries in favor of disclosure may focus on how disclosure would benefit the participant, other entities (for example, the entity collecting personal information), or society, whereas queries against disclosure might focus on the harm stemming from the same disclosures. Indeed, prior work substantiates that manipulating the salience of competing considerations in privacy settings can significantly impact privacy decision making (John, Acquisti, and Loewenstein 2011). Beyond simply being applicable to privacy settings, we conjecture that QT has the potential to reconcile the emergence of both rational and behavioral privacy perspectives. This is because QT is built on a simple but insightful conjecture: factors that can alter the valence and order of the queries executed in support of a decision are also likely to result in changes in observed behavior. More critically, QT remains agnostic to the factor (that is, either rational or behavioral) that causes the shift in the order and valence of queries in decision settings.

Taken together, the various theoretical perspectives highlighted in this section provide the impetus for our research. The extant literature suggests that consumers likely exhibit both rational and behavioral re-

sponses in privacy settings but has yet to harmonize the emergence of behaviors that assume different models of underlying human behavior. With this gap in mind, we argue—and test—that the process-centric view of decision making provided by QT can help reconcile seemingly disconnected privacy behaviors by providing a common basis for the emergence of rational and behavioral responses in privacy settings.

3. HYPOTHESES

To examine the proposition that QT can account for the impact of both rational and behavioral privacy responses, we sought a privacy decision context where we could test the impact of rational and behavioral factors on behavior. Clearly, many privacy decision settings are available, as are several experimental manipulations of either type. For example, Adjerid, Peer, and Acquisti (2016) focus on the impact on self-disclosure of protections communicated via privacy notices (for example, whether responses are anonymous or identified). They evaluate the impact of a rational factor by manipulating the objective level of protection communicated in a privacy notice, and they manipulate a behavioral factor by varying the relative perception of identical notices (that is, holding the objective information constant).

We focus on a different privacy decision setting that has gained prominence in recent years: consumers' privacy choices made via data use settings (for example, privacy settings). Consumer control through data use settings has become a cornerstone of US policy toward online privacy protection (FTC 2012; White House 2012), policy that has been widely supported by industry (Solove 2013). World Economic Forum (2013) suggests that new technological options can give individuals control over their own information while allowing data assets to flow relatively freely; a senior advisor for a large technology firm (and contributor to the World Economic Forum report) stated that "[t]here's no bad data, only bad uses of data" (Lohr 2013, p. BU3). As a result, consumers' data use decisions now pervade various technology settings, including online social networks, mobile devices, search engines, and web browsers. With this decision context in mind, we sought relevant manipulations of both rational and behavioral factors.

Since our focus is not on demonstrating that rational and behavior factors can impact privacy decision making (this result is reasonably val-

idated in the literature), we leaned on existing literature to identify manipulations of rational and behavioral factors. In effect, we sought manipulations of rational and behavioral factors that we considered likely to impact privacy choices, allowing us to test whether QT is consistent in these impacts.

We first considered factors that would, under rational accounts of privacy decision making, alter observed privacy behaviors. Consistent with Adjerid, Peer, and Acquisti (2016), we manipulated a rational factor by introducing objective differences in the entity with which participants were provided the option to share their data. Leveraging a privacy-calculus perspective of decision making, in which factors that alter the risks and benefits associated with a privacy choice alter behavior (Dinev and Hart 2006), we conjectured that when the decision involved sharing with entities that were perceived as riskier, participants would be less likely to share their information. This conjecture is substantiated by previous empirical studies. For instance, Tsai et al. (2011) find that when participants are presented with a rating for each website's privacy policy in a list of search results, they tend to purchase from websites that offer medium or high levels of privacy, even when the price of the purchased product is higher on these websites. Thus, we make the following hypothesis:

Hypothesis 1. Participants will be less likely to allow a data use when they perceive a higher risk from that permission.

To identify a behavioral factor, we leaned again on prior work in the behavioral economics and privacy literature. We focused on choice framing, or the phenomenon of "simple and unspectacular changes" in the presentation of decision problems, unrelated to their objective costs and benefits, that lead to changes in choice (Kühberger 1998, p. 24). Framing effects have been widely studied in the economics and psychology literature (Levin, Schneider, and Gaeth 1998), and recent work suggests that their effects may be a pronounced choice of data use settings. For example, Adjerid, Acquisti, and Loewenstein (2016) find that individuals were 45 percent more likely to select the privacy-protective option when presented with a choice to prohibit a use of their personal information (the reject frame) than when they were presented with the objectively identical setting as a choice to allow use of their personal information (the accept frame). This finding is consistent with the theory posited by Shafir (1993) that positive dimensions of choice weigh heavier under an accept frame,

while negative dimensions of that same choice weigh heavier under a reject frame. We use a similar manipulation of decision frames and alter whether the choice to share is presented as a choice to allow the sharing of data versus a choice to prohibit the sharing of data. Similar to other framing manipulations, we change only the format of the choice while keeping the objective options constant. We make the following hypothesis:

Hypothesis 2. Participants will be more likely to allow data sharing when the choice is presented in an accept frame relative to a reject frame.

In Section 2, we argued that simultaneous support for hypotheses 1 and 2 introduces some dissonance as to the nature of privacy decision making. We also suggested that a process-centric view of decision making (using QT) has the potential to provide one (but potentially not the only) baseline for the emergence of both effects. Because the QT literature has primarily focused on the potential of QT to explain (and sometimes eliminate) various decision biases, we focus first on how QT can explain the impact of choice framing on privacy decision making (hypothesis 2). In a seminal work, Johnson, Häubl, and Keinan (2007) consider the classic endowment effect in which those endowed with an object (for example, a coffee mug) tend to demand more money in exchange for the object than those who have not been endowed (Kahneman, Knetsch, and Thaler 1991). They argue that individuals in this context decompose their decision into queries such as *Why should I make the trade?* or *Why should I not make the trade?* In their study, participants were randomized into conditions in which they were either endowed with a mug (sellers) or not (choosers), asked to autonomously generate their own reasons either in favor of selling the mug or against exchanging it for money (one reason at a time), and then asked to provide their valuations of the mug. The authors show that sellers generate more queries in favor of keeping the mug and against exchanging it for money relative to choosers. In addition, sellers first generate queries in favor of keeping the mug and against exchanging it for money, while choosers generate queries in the reverse order. Subsequent work finds that QT accounts of behavior are consistent in cases of asymmetric discounting in which individuals discount more heavily when asked to delay rather than accelerate consumption (Weber et al. 2007). Both of these works suggest loss aversion or changes in participants' implicit goals or focus as an explanation for why individuals' queries are shifted. Closer to our context, Hardisty, Johnson, and Weber

(2010) show that QT is consistent with the emergence of attribute framing effects. As a result, we make the following hypothesis:

Hypothesis 3a. Participants in the accept frame will generate more reasons in favor of sharing than participants in the reject frame.

Hypothesis 3b. Participants in the reject frame will generate more reasons against sharing than participants in the accept frame.

Hypothesis 3c. Participants in the accept frame will generate reasons in favor of sharing before generating reasons against sharing, and participants in the reject frame will generate reasons against sharing before generating reasons in favor of sharing.

The potential of QT to be consistent for behavioral responses to rational factors (hypothesis 1) is less explored in the literature but is, we argue, a reasonable extension of the QT framework. First, a striking result of the current QT literature is that decision biases, like the endowment effect and asymmetric discounting, can be eliminated simply by reversing the order of the queries that individuals consider (Johnson, Häubl, and Keinan 2007; Weber et al. 2007). This implies that there exists a sequence of queries that corresponds to rational responses by consumers in these settings. Hardisty, Johnson, and Weber (2010) present some additional evidence that responses to rational factors can be explained by QT. Although they focus on the effect of framing on behavior and the role of QT in explaining the inconsistency in behavior between choice frames, they identify a result that is not the focus of their analysis but that is telling for our purposes. They find that an environmental intervention labeled a tax results in significantly lower approval from participants with Republican leanings (relative to Democrats and Independents) and that this effect can be explained, consistent with other QT studies, by variation in the valence and sequence of queries generated by these participants. In contrast to prior work that focused on behavioral or nonrational drivers of changes in query valence and order, Hardisty, Johnson, and Weber attribute the effects to Republican participants' strong underlying preferences against taxation.¹ Similarly, it is plausible that changes to data use settings that alter the objective risk of harm from a data permission may alter behavior by shifting the valence and order of queries individ-

1. These participants were not entirely consistent in their behavior, since they did not exhibit the same response when the intervention was labeled an offset. In this case, Republican participants' queries did not significantly differ from those of other participants.

uals make. This aligns with prior results by Hann et al. (2007) who use another information-processing theory (expectancy theory) and find that privacy protections alter the valence of considerations in privacy settings. As a result, we make the following hypothesis:

Hypothesis 4a. Participants who are asked to share their sensitive information with a low-risk entity will generate more reasons in favor of sharing than participants who are asked to share their sensitive information with a high-risk entity.

Hypothesis 4b. Participants who are asked to share their sensitive information with a high-risk entity will generate more reasons against sharing than participants who are asked to share their sensitive information with a low-risk entity.

Hypothesis 4c. Participants who are asked to share their sensitive information with a low-risk entity will generate reasons in favor of sharing before generating reasons against sharing, and participants who are asked to share their sensitive information with a high-risk entity will generate reasons against sharing before generating reasons in favor of sharing.

4. EXPERIMENT

To test the hypotheses presented in Section 3, we conducted an experiment using Amazon Mechanical Turk (AMT), an online service that has become increasingly popular among social scientists for conducting online experiments.² Buhrmester, Kwang, and Gosling (2011) demonstrate that AMT samples are just as representative as other Internet samples and are considerably more representative than typical student samples. Steelman, Hammer, and Limayem (2014) find that AMT samples have psychometric properties similar to those of both student and consumer panels. Furthermore, judgment and decision-making experiments using AMT samples have replicated results found in traditional subject samples (Goodman, Cryder, and Cheema 2013).

In the experiment, participants were invited to answer a series of personal questions related to sensitive behaviors and were provided different

2. We restricted participants to subjects from the United States with a human intelligence task approval rate on Amazon Mechanical Turk of over 95 percent. We included attention check questions at the start of the questionnaire following accepted practices in the field (for example, Oppenheimer, Meyvis, and Davidenko 2009).

information—depending on the experimental conditions—concerning the way their responses would be used. This study context provides a number of desirable features for studying privacy decision making. First, it focuses on actual behaviors, not self-reported attitudes or hypothetical behaviors, as participants are required to make disclosures of sensitive information. Second, prior work has found that participants' behaviors in this type of experiment are very responsive to different uses of their sensitive disclosures. For instance, more invasive uses result in participants disclosing less information (Adjerid, Acquisti, and Loewenstein 2016). This suggests that participants do not treat this context as riskless or behave in an arbitrary fashion when making disclosure decisions within it. As such, this context provides a validated framework to evaluate whether QT can predict the impact of rational and behavior factors in privacy settings—as we detail below.

4.1. Procedure and Design

Participants in the study were first shown an introductory screen that described the study context and provided an example of the sensitive questions asked in the study (“Have you ever had a one-night stand?”). They were then asked demographic questions. Participants were next asked to make a decision about whether they would be willing to share their responses to the sensitive questions. We collected queries using the aspect-listing task from Johnson, Häubl, and Keinan (2007): participants were asked to provide their own reasons either in favor of or against sharing their information, one reason a time. Participants were provided with open-ended text fields in which they could write down their reasons in favor of or against sharing their information. They were instructed to keep providing reasons until they could not think of any more reasons to provide, but the system was set up to accept a maximum on 10 reasons.³ Similar to the design of previous QT studies, participants were asked to code their own reasons as either in favor of or against sharing.⁴ For half the participants, reasons were collected before the sharing choice was

3. Participants were not aware that the maximum possible reasons the system could accept was 10, and less than 2 percent of our participants submitted 10 reasons.

4. In addition to the coding performed by participants, we also had two independent coders (who were blind to the experimental conditions) code these reasons as either in favor of or against sharing. There was a high level of agreement between the subjects' codes and each individual coder's codes as well as between the codes of both independent coders (above 80 percent κ agreement in all comparisons). For the analyses, we use subjects' own coding of reasons.

made, while for the other half the reasons were collected after participants made their sharing choice. We manipulated when reasons were collected to confirm whether the same expected pattern in valence and order of reasons is found both before and after participants make the sharing choice. Since we care about two dependent variables (the sharing choice and its reasons), it is important to confirm that the effect of the framing manipulation and risk manipulation on sharing choice is not different depending on when reasons are collected. Finally, participants were asked a set of sensitive questions borrowed from Acquisti, John, and Loewenstein (2012). The text of the question used to elicit reasons and the questions regarding sensitive behaviors are in Appendix Sections A1 and A2.

The experiment employed a factorial 2 (high risk, low risk) $\times$ 2 (accept frame, reject frame) between-subjects design. We tested the impact of a rational factor by altering whether the data use setting presented to participants involved sharing their sensitive disclosures with a low-risk versus a high-risk entity. In particular, we implemented the low-risk condition by asking participants whether they would like to share their responses with other research organizations and the high-risk condition by asking whether they would like to share their responses with a marketing company. Simultaneously, we evaluated the impact of a behavioral factor by varying, between subjects, whether participants were asked to allow the sharing of their information (the accept frame) or to prohibit the same sharing of their information (the reject frame).

4.2. Pilot Study

To ensure that the different data-sharing options used to manipulate a rational factor were perceived as presenting different levels of risk, we ran a prestudy with another set of participants recruited from the same population. These participants were asked to imagine taking a survey on AMT that involved answering ethical questions. They were presented with an example ("Have you ever had a one-night stand?") to provide a sense of the level of intrusiveness of the questionnaire. Then they were asked to imagine different scenarios in which researchers asked them whether they would share their responses to the ethical-behavior questions with other research organizations or a marketing company. These scenarios were presented on a single page, and the order in which they were presented was randomized across participants. Participants were asked to rate how risky they thought it would be to share the information and how likely

they would be to share it, on a 1–7 scale ranging from “not at all” to “very much.”

One hundred twenty participants (mean age = 34.3; 67 percent male) from AMT completed the prestudy. We validated that these different entities represented different levels of risk for our participants: sharing with a marketing company was perceived by participants as significantly riskier compared with sharing with research organizations (4.03 versus 3.24, $t(119) = 5.1, p < .001$), and participants reported being significantly less likely to share with a marketing company relative to other research organizations (3.57 versus 4.70, $t(119) = -6.3, p < .001$).

4.3. Data and Analysis

We recruited 745 individuals (mean age = 32.4; 57 percent male) from AMT. We evaluated the impact of our randomized treatments using the appropriate statistical tests for our variable of interest (for example, t -test or χ^2 test) and supplemented this analysis with the appropriate regression analysis. We estimated the following general model:

$$\text{AllowDataUse}_i = \beta \times \text{Treatment}_i + \alpha \times \mathbf{Y}_i + \mathbf{u}_i,$$

where AllowDataUse_i is a binary measure of whether participant i allows the sharing of his or her sensitive information disclosures and Treatment_i is an indicator variable of our randomized manipulations. In some specifications, we included $\mathbf{Y}_i$, a vector with controls for participant-specific characteristics (for example, age and gender). Estimates on randomly assigned treatments (Treatment_i) are unbiased, as they should be uncorrelated with control variables ($\mathbf{Y}_i$) and the error term $\mathbf{u}_i$. While our controls are not necessary for the unbiased estimation of the effect of our treatments on disclosure behavior, they were included in some specifications to rule out any breaks in randomization and to account for some of the variation in disclosure behavior between participants. We extended this general model to evaluate the effect of our manipulations on the valence and order of queries collected from our participants. Across our analysis, we estimate a combination of ordinary least squares (OLS), probit, and Poisson regressions as appropriate.

4.4. Results

We find the same pattern in sharing choices when queries are collected before or after participants make the sharing choice, with the only difference being an overall decrease in likelihood to share when the queries are

collected first. The Appendix presents graphs for the percentage of participants willing to share their responses when queries are collected first and when the sharing choice is made first (Figures A2 and A3). Given that the pattern across the four conditions is consistent irrespective of when queries are collected, for the rest of the analysis we pool these data.

We find that rational factors (whether a high- or low-risk entity will receive the data) impacted the choice to allow sharing of sensitive information. Participants in the low-risk condition were twice as likely to allow the data use relative to those in the high-risk condition (62 percent versus 34 percent; $\chi^2(1) = 58.08$; $p < .001$). Thus, hypothesis 1 was supported. We also find that behavioral factors had an impact on participants' propensity to allow their sensitive disclosures to be shared. Participants who were presented with an accept frame were significantly more likely to allow sharing of their sensitive information relative to those who were presented with the same choice in the prohibit frame (55 percent versus 40 percent; $\chi^2(1) = 16.76$; $p < .001$). Thus, hypothesis 2 was supported. Estimation of a regression model confirms the presence of the two main effects ($\beta_{\text{Allow}} = .180$, $p < .001$; $\beta_{\text{LowRisk}} = .296$, $p < .001$). The model with interaction shows no significant interaction effect ($\beta_{\text{Allow} \times \text{LowRisk}} = .014$, $p = .835$). Figure 1 presents these results; the regression and probit results are in Table A2.

We analyzed the number of reasons provided in favor and against disclosure and the order in which these reasons were provided across the two framing conditions. To analyze the number of reasons in favor and against, we estimated Poisson regressions with the count of reasons as the dependent variable and the framing condition dummy as the independent variable. We found that participants in the allow condition provided significantly more reasons in favor of sharing than those in the prohibit condition ($\beta_{\text{Allow}} = .288$, $p < .001$), and they provided significantly fewer reasons against sharing than those in the prohibit condition ($\beta_{\text{Allow}} = -.182$, $p = .002$). There are no significant differences between the total number of reasons provided by participants in the allow and prohibit conditions ($\beta_{\text{Allow}} = .01$, $p = .825$). Therefore, we find support for hypotheses 3a and 3b. The Poisson regression coefficients and their standard errors are reported in Table A2.

Next we analyzed the order in which these reasons were submitted by computing the standardized median rank difference (SMRD) for each participant. This score is calculated as $2(\text{MR}_{\text{infavor}} - \text{MR}_{\text{against}})/n$, where $\text{MR}_{\text{infavor}}$ is the median rank of reasons provided in favor of sharing,




Figure 1. Percentages of participants who gave permission to share their information

MR_{against} is the median rank of reasons provided against sharing, and n is the total number of reasons provided. The scores vary from -1 to $+1$, with smaller numbers indicating that in-favor reasons were submitted before reasons against sharing. Using linear regression (OLS), we confirm that SMRD scores were significantly lower for participants in the allow condition than those in the prohibit condition ($\beta_{\text{Allow}} = -.25, p < .001$). Therefore, we find support for hypothesis 3c. The regression coefficients and their standard errors are reported in Table A3.

We conducted the same analysis to test whether the number and order of reasons differed between the high- and low-risk conditions. We found that participants in the low-risk condition provided significantly more reasons in favor of sharing than those in the high-risk condition ($\beta_{\text{LowRisk}} = .726, p < .001$), and they provided significantly fewer reasons against sharing than in those the high-risk condition ($\beta_{\text{LowRisk}} = -.471, p < .001$). There are no significant differences between the total number of reasons provided by participants in the low-risk and high-risk conditions ($\beta_{\text{LowRisk}} = .018, p = .689$); see Table A2. In addition, SMRD scores were significantly lower for participants in the low-risk condition than for those in the high-risk condition ($\beta_{\text{LowRisk}} = -.499, p < .001$); see Table A3. Therefore, we find support for hypotheses 4a, 4b, and 4c.

5. DISCUSSION

Our results suggest that consumers can exhibit, in the same decision context, both rational and behavioral responses to privacy settings: objective differences in the risk associated with disclosure and subtle variation in how these decisions were framed both influence how participants acted during the experiment. More important, the results show that QT can provide one common basis for the emergence of both rational and behavioral responses in privacy settings. We find that the high- and low-risk data uses and the various decision frames resulted in significant differences in the number of positive versus negative queries by participants and in the order of these queries.

Our results have implications for privacy research. Various models of decision making in privacy settings exist, many of which have developed independent of one another. The effort to reconcile these different models is nascent in the privacy literature but mirrors the approach—common in the broader economics literature—of modifying extant models of behavior to account for observed deviations from such models (Oppenheimer and Kelso 2015). For example, prospect theory (Kahneman and Tversky 1979) modified axioms of choice defined by Von Neumann and Morgenstern (1944) to allow different value function for gains relative to losses. And theories of hyperbolic time discounting (Laibson 1997) adjusted the assumptions of traditional discount utility theory (Samuelson 1937) to allow a declining discount rate between the current period and next one. This approach has resulted in economic models of decision making that may be more representative of individual behaviors, which suggests that a similar strategy in privacy research may be also fruitful.

That approach, however, has also been criticized, because models thus revised can become highly complex and increasingly less usable as more anomalies emerge in the literature (Oppenheimer and Kelso 2015). Taking a simplified, process-focused approach to decision making has the potential to address some of these critiques: Weber et al. (2007, p. 522) suggest that QT can augment extant economic models of human behavior by providing “a process-model instantiation and explanation of the effects described mathematically” by economic models of behavior. The value of this focus for privacy research is twofold. First, a process view of decision making sidesteps controversial assumptions of consumers’ privacy decision making while still helping to explain important variation in observed decision making. Second, a focus on the process by which privacy choices

are made may provide a stronger basis for improving consumers' privacy decision making, particularly in the face of policy approaches that are heavily reliant on consistent rational choices by consumers. Johnson, Häubl, and Keinan (2007) also highlight this as a key advantage of a QT approach to decision making. They note that a "mechanism-based explanation might suggest interventions that would reduce or eliminate" observed decision biases (Johnson, Häubl, and Keinan 2007, p. 462). This is substantiated by evidence that reversal of query valence and order can eliminate some observed decision biases, such as the endowment effect and asymmetric discounting. A broader focus in the literature on privacy decision-making processes may also uncover other such opportunities to better align consumers' privacy behavior with the assumptions of policy mechanisms intended to improve consumer welfare.

Our results also have policy implications. Consistent with a growing critique of notice and consent privacy mechanisms, our results reinforce the challenges associated with an overreliance on notice and consent mechanisms for privacy protection. These challenges are particularly relevant if those mechanisms are employed in lieu of more substantive approaches that provide baseline consumer data protections. Even more central is the observation that our findings offer one plausible theory of decision making that seems to span expected rational responses and more problematic (from the perspective of the efficacy of notice and consent approaches) behavioral responses. This insight may help to improve (and critically assess) self-regulatory mechanisms that are intended to aid consumers in managing privacy trade-offs. Understanding that one way in which individuals approach privacy choices is via a sequence of queries and that the order and valence of these queries has an impact on decision making highlights the need to think beyond simply notifying individuals of firms' data practices and providing them with some choice in the matter. Rather, effective notice and consent ought to also consider how factors seemingly disconnected from the objective features of privacy contexts may alter the process of privacy decision making and thus consumers' subsequent privacy choices.

APPENDIX: EXPERIMENT MATERIALS AND ANALYSIS

A1. Question Used to Collect Reasons

Allow other research organizations to collect your responses to the ethical behavior questions?

Now, we would like to know all the **reasons** that you thought about while trying to answer the question above, both in favor and against allowing other research organizations to collect your responses to the ethical behavior questions.

We request that you tell us your reasons **one at a time**. In the box below, please enter your **first** reason, **either in favor or against** allowing other research organizations to collect your responses to the ethical behavior questions.

When you are done, hit the “Next” key to submit it, and proceed to the next page where you can enter the next reason.

Reason 1:

0%100%

Next

Figure A1. Survey prompt for sharing reasons

Figure A1 displays the screen showing the question used to collect reasons in favor or against respondents sharing their answers to sensitive questions and is adapted from Johnson, Häubl, and Keinan (2007).

A2. Ethical Questions in the Study

The ethical questions used in the study are from Acquisti, John, and Loewenstein (2012).

1. Have you ever had sex with the current husband, wife, or partner of a friend?
2. Have you ever masturbated at work or in a public restroom?
3. Have you ever had a fantasy of doing something terrible (e.g., torturing) to someone?
4. Have you ever fantasized about having violent nonconsensual sex with someone?
5. Have you ever, while an adult, had sexual desires for a minor?
6. Have you ever neglected to tell a partner about a sexually transmitted disease from which you were suffering?
7. Have you ever had sex with someone who was too drunk to know what they were doing?
8. Have you ever stolen anything that did not belong to you?

9. Have you ever tried to gain access to someone else's (e.g., a partner, friend, or colleague's) email account?
10. Have you ever looked at pornographic material?

A3. Response Patterns When Queries Are Collected before or after the Sharing Choice




Figure A2. Results when sharing choice is made before the queries




Figure A3. Results when sharing choice is made after the queries

A4. Regression Results

Table A1. Choice to Allow Data Sharing

Variable	Ordinary Least Squares: Share			Probit
	(1)	(2)	(3)	(4)
LowRisk	.296** (.035)	.289** (.049)	.294** (.048)	.296** (.034)
Allow	.179** (.034)	.172** (.047)	.172** (.047)	.179** (.034)
LowRisk × Allow		.014 (.069)	.007 (.069)	.014 (.069)
Male			-.034 (.035)	
Age			-.005** (.002)	
Constant	.240** (.029)	.244** (.033)	.443** (.076)	

Note. Standard errors are in parentheses.
***p* < .01.

Table A2. Analysis of Reasons

	Reasons in Favor (1)	Reasons Against (2)	Total Reasons (3)	Reasons in Favor (4)	Reasons Against (5)	Total Reasons (6)
Allow	.288** (.072)	-.182** (.060)	.010 (.046)			
LowRisk				.726** (.075)	-.471** (.062)	.018 (.046)
Constant	-.103* (.055)	.506** (.041)	.941** (.033)	-.361** (.061)	.618** (.037)	.937** (.032)

Note. Standard errors are in parentheses.

Table A3. Standardized Median Rank Difference Score

	(1)	(2)
Allow	-.252** (.071)	
LowRisk		-.499** (.069)
Constant	.221** (.050)	.333** (.049)

Note. Standard errors are in parentheses.
***p* < .01.

REFERENCES

- Acquisti, Alessandro. 2004. Privacy in Electronic Commerce and the Economics of Immediate Gratification. Pp. 21–29 in *Proceedings of the 5th ACM Conference on Electronic Commerce*. New York: ACM.
- Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein. 2015. Privacy and Human Behavior in the Age of Information. *Science*, January 30, pp. 509–14.
- Acquisti, Alessandro, Leslie K. John, and George Loewenstein. 2012. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research* 49:160–74.
- Adjerid, Idris, Alessandro Acquisti, and George Loewenstein. 2016. Choice Architecture, Framing, and Layered Privacy Choices. Unpublished manuscript. University of Notre Dame, Mendoza College of Business, Notre Dame, IN.
- Adjerid, Idris, Eyal Peer, and Alessandro Acquisti. 2016. Beyond the Privacy Paradox: Objective versus Relative Risk in Privacy Decision Making. Unpublished manuscript. University of Notre Dame, Mendoza College of Business, Notre Dame, IN.
- Ben-Shahar, Omri, and Carl E. Schneider. 2014. *More Than You Wanted to Know: The Failure of Mandated Disclosure*. Princeton, NJ: Princeton University Press.
- Breaban, Adriana, and Charles N. Noussair. 2013. Emotional State and Market Behavior. Discussion Paper No. 2013-031. Tilburg University, CentER, Tilburg.
- Buhrmester, Michael, Tracy Kwang, and Samuel D. Gosling. 2011. Amazon’s Mechanical Turk: A New Source of Inexpensive, yet High-Quality, Data? *Perspectives on Psychological Science* 6:3–5.
- Dempster, Frank N. 1995. Interference and Inhibition in Cognition: An Historical Perspective. Pp. 3–26 in *Interference and Inhibition in Cognition*, edited by Frank N. Dempster and Charles J. Brainerd. San Diego, CA: Academic Press.
- Dinev, Tamara, and Paul Hart. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. *Information Systems Research* 17:61–80.
- Dinev, Tamara, 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* 26:639–55.
- Featherman, Mauricio S., and Paul A. Pavlou. 2003. Predicting E-Services Adoption: A Perceived Risk Facets Perspective. *International Journal of Human-Computer Studies* 59:451–74.
- FTC (Federal Trade Commission). 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>.
- Goodman, Joseph K., Cynthia E. Cryder, and Amar Cheema. 2013. Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Sam-

- ples. *Journal of Behavioral Decision Making* 26:213–24.
- Hann, Il-Horn, Kai-Lung Hui, Sang-Yong Tom Lee, and Ivan PL Png. 2007. Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach. *Journal of Management Information Systems* 24(2):13–42.
- Hardisty, David J., Eric J. Johnson, and Elke U. Weber. 2010. A Dirty Word or a Dirty World? Attribute Framing, Political Affiliation, and Query Theory. *Psychological Science* 21(1):86–92.
- John, Leslie K., Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of Consumer Research* 37:858–73.
- Johnson, Eric J., Gerald Häubl, and Anat Keinan. 2007. Aspects of Endowment: A Query Theory of Value Construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 33:461–74.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991. Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. *Journal of Economic Perspectives* 5(1):193–206.
- Kahneman, Daniel, and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47:263–91.
- Klopfner, Peter H., and Daniel I. Rubenstein. 1977. The Concept Privacy and Its Biological Basis. *Journal of Social Issues* 33(3):52–65.
- Kugler, Matthew B., and Lior Strahilevitz. 2015. Surveillance Duration Doesn't Affect Privacy Expectations: An Empirical Test of the Mosaic Theory. Coase-Sandor Working Paper in Law and Economics No. 727. University of Chicago Law School, Chicago.
- Kühberger, Anton. 1998. The Influence of Framing on Risky Decisions: A Meta-Analysis. *Organizational Behavior and Human Decision Processes* 75:23–55.
- Laibson, David. 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112:443–77.
- Levin, Irwin P., Sandra L. Schneider, and Gary J. Gaeth. 1998. All Frames Are Not Created Equal: A Typology and Critical Analysis of Framing Effects. *Organizational Behavior and Human Decision Processes* 76:149–88.
- Lohr, Steve. 2013. Big Data Is Opening Doors, But Maybe Too Many. *New York Times*. <http://www.nytimes.com/2013/03/24/technology/big-data-and-a-renewed-debate-over-privacy.html>.
- Milne, George R., and Mary Ellen Gordon. 1993. Direct Mail Privacy-Efficiency Trade-offs within an Implied Social Contract Framework. *Journal of Public Policy and Marketing* 12:206–15.
- Nguyen, Yen, and Charles N. Noussair. 2014. Risk Aversion and Emotions. *Pacific Economic Review* 19:296–312.
- Oppenheimer, Daniel M., and Evan Kelso. 2015. Information Processing as a Paradigm for Decision Making. *Annual Review of Psychology* 66:277–94.
- Oppenheimer, Daniel M., Tom Meyvis, and Nicolas Davidenko. 2009. Instructional Manipulation Checks: Detecting Satisficing to Increase Statistical Power.

- Journal of Experimental Social Psychology* 45(4):867–72.
- Petronio, Sandra. 2015. Communication Privacy Management Theory. Pp. 1–9 in *The International Encyclopedia of Interpersonal Communication*, edited by Charles R. Berger and Michael E. Roloff. Chichester: John Wiley & Sons.
- Posner, Richard A. 1981. The Economics of Privacy. *American Economic Review* 71:405–9.
- Samuelson, Paul A. 1937. A Note on Measurement of Utility. *Review of Economic Studies* 4:155–61.
- Shafir, Eldar. 1993. Choosing versus Rejecting: Why Some Options Are Both Better and Worse than Others. *Memory and Cognition* 21:546–56.
- Solove, Daniel J. 2013. Introduction: Privacy Self-Management and the Consent Dilemma. *Harvard Law Review* 126:1880–1903.
- Steelman, Zachary R., Bryan I. Hammer, and Moez Limayem. 2014. Data Collection in the Digital Age: Innovative Alternatives to Student Samples. *MIS Quarterly* 38:355–78.
- Stigler, George J. 1980. An Introduction to Privacy in Economics and Politics. *Journal of Legal Studies* 9:623–44.
- Taylor, Curtis R. 2004. Consumer Privacy and the Market for Customer Information. *RAND Journal of Economics* 35:631–50.
- Tourangeau, Roger. 2004. Survey Research and Societal Change. *Annual Review of Psychology* 55:775–801.
- Tsai, Janice Y., Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 22:254–68.
- Von Neumann, John, and Oskar Morgenstern. 1944. *Game Theory and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Weber, Elke U., Eric J. Johnson, Kerry F. Milch, Hannah Chang, Jeffrey C. Brod-scholl, and Daniel G. Goldstein. 2007. Asymmetric Discounting in Intertem-poral Choice: A Query-Theory Account. *Psychological Science* 18:516–23.
- Westin, A. F. 2000. Intrusions: Privacy Tradeoffs in a Free Society. *Public Per-spective* 11(6):8–11.
- White House. 2012. *Consumer Data Privacy in a Networked World*. Washington, DC: The White House. <http://www.whitehouse.gov/sites/default/files/privacy-final.pdf>.
- World Economic Forum. 2013. *Unlocking the Value of Personal Data: From Col-lection to Usage*. Geneva: World Economic Forum. http://www3.weforum.org/docs/WEF_IT_UnlockingValuePersonalData_CollectionUsage_Report_2013.pdf.