

Choice Architecture, Framing, and Cascaded Privacy Choices

Idris Adjerid^{*}, Alessandro Acquisti⁺, George Loewenstein⁺

^{*}University of Notre Dame, ⁺ Carnegie Mellon University

For consumers, managing privacy online is a continuous, complex process of interrelated choices. Increasingly, this process may be conceived of as “cascaded,” in that a combination of upstream and downstream choices together determine some ultimate privacy outcome and its associated benefits and costs. For example, on social networks, individuals select profile visibility settings (an upstream choice) and then, given those settings, decide what information to post (a series of downstream choices). In combination, upstream and downstream choices determine consumers’ privacy risks and trade-offs. In a series of experiments, we investigate how changes in consumer choice architecture (specifically, choice frames) implemented by service providers can influence consumers’ upstream privacy choices (specifically, their choices of “disclosure settings”), and whether those upstream effects are “mitigated” through changes in levels of downstream self-disclosure after consumers are steered towards choosing riskier or safer upstream disclosure settings. We find, first, that various manipulations of decision frames, common in privacy contexts, significantly impact participants’ upstream choice of disclosure settings. Second, we do *not* find evidence that the impact of choice architecture upstream is mitigated downstream: participants’ self-disclosure rates do not adjust or change across manipulations of decision frames that induced differences in protective disclosure settings upstream. Our results contribute to an emerging behavioral perspective on privacy decision making, and highlight the importance of accounting for the cascaded nature of privacy decision making in both policy and managerial settings.

Published Version Can Be Found:

Management Science

65(5), 2267-2290 (2019)

<https://doi.org/10.1287>

1. Introduction

Consumer privacy decision making is often *cascaded*: a combination of early “upstream” and subsequent “downstream” choices together determine some ultimate privacy outcome and its associated benefits and costs. For example, a person might first decide whether to get on a particular platform, such as Twitter, Facebook or Snapchat, then, if the decision is affirmative, what to post on the site. A patient may first decide whether to sign a form consenting to various uses of his health information, and then how much health information to disclose to a provider. An individual solicited for a marketing survey may decide whether to respond (and whether to provide her full name to the survey company), and then (perhaps in part dependent on whether she provided her name) what to disclose in her responses. In all these examples of cascaded choices, the ultimate privacy outcome, along with its associated benefits and risks, depends on a series of upstream and downstream choices. Cascaded privacy choices are particularly common online, due in no small part to industry and policy makers’ reliance on “choice” mechanisms (often designed by service providers themselves) as tools for privacy protection (FTC, 2012). For instance, consumers may first be asked to choose profile visibility settings in an online social network; and then decide how much to disclose through that network. They may first select search engine settings (such as searching under an identified, linked account, or not); and then perform more or less sensitive searches using the engine. They may decide whether to use incognito mode on a browser, or privacy enhancing technologies such as Tor; and then engage in online browsing activities. Or they may pick their preferences for who will view their profile on a dating website; and then choose what to disclose on said site.

The cascaded nature of online privacy decision making has important implications for models of privacy decision making and for consumers’ privacy risks. However, it has been underexplored in the literature. The dominant focus of the extant privacy literature has been on downstream privacy choices, often in the form of individual self-disclosure decisions (Moon, 2000; Dinev and Hart, 2006; Xu et al., 2009; Xu et al., 2012; Acquisti, John, and Loewenstein, 2012; Brandimarte, Acquisti, and Loewenstein, 2013). The comparatively few studies that evaluate upstream privacy choices tend to investigate them

separately from their associated downstream choices (Johnson, Bellman, and Lohse, 2002; Stutzman, Gross, and Acquisti, 2013). The lack of studies that evaluate upstream and downstream choices jointly, and take into account the increasingly cascaded nature of online privacy decision making, represents a significant gap in the literature. This gap motivates the two central aims of this manuscript.

First, cascaded privacy choices involve upstream privacy choices that can vary in important but subtle ways between providers and over time. In particular, upstream decisions in the form of “disclosure settings” that provide consumers the option to restrict how their downstream self-disclosures are shared and used (e.g., privacy settings on social networks) can be highly variable in how they are presented to consumers. This is relevant in light of research in behavioral economics that shows that subtle changes in the design (or “architecture”) of choices can profoundly influence consumers’ behavior and decision making (e.g., Keller et al., 2011; Johnson et al., 2012; Egelman, Felt, and Wagner, 2013; Thaler, Sunstein, and Balz, 2014). These streams of work suggest that the architecture of upstream privacy choices—even when ostensibly aimed at improving consumer privacy—may be subtly molded to nudge individuals towards greater sharing of personal information, in some cases exposing them to risks that they may be unaware of or underweight (due to, e.g., time discounting and/or immediate emotional factors). Thus, the first focus of our manuscript is on whether changes in consumers’ choice architecture implemented by service providers influence the choice of disclosure settings upstream.

Jointly considering upstream and downstream choices raises a second important consideration: Whether choice architectures that steer consumers towards more or less protective disclosure settings upstream ultimately matter, depends on whether their impact is mitigated (or not) by more or less protective downstream choices (e.g., decision frames that lead consumers to choose risky disclosure settings upstream are also associated with lower levels of self-disclosure downstream). That is, the manner in which factors that change upstream decisions will have an actual impact on privacy trade-offs ultimately depends on the extent to which consumers compensate downstream for changes in upstream decisions, which they may or may not be aware they have been influenced to make. The risk compensation and moral hazard literature suggests that this compensation behavior is likely and argues

that downstream choices can undo the effect of upstream ones. For example, Peltzman (1975) famously raised the question of whether safety innovations in automobiles, such as seatbelts, would ultimately have little or no effect because consumers would take advantage of the increased margin of safety by driving more recklessly. Similarly, the moral hazard literature documents instances where insurance subsequently reduces measures that individuals take to mitigate risks, such as, in the domain of health, preventative care and unhealthy behaviors (Klick and Stratmann, 2007; Kelly and Markowitz, 2009). Such downstream compensatory behavior, which is critical to mitigating the effects of induced changes in upstream behavior, may, however, be less likely when changes in upstream choices are a result of subtle manipulation of choice architecture. Consumers may simply not notice subtle changes to their choice architecture, and even if they do notice, they may fail to anticipate the effect that subtle changes in architecture can have on their own behavior (Lieberman, Samuels, and Ross, 2004) or simply fail to react to the knowledge that their architecture has been manipulated (Loewenstein et al., 2015). Because this mitigation is critical to overall consumer privacy outcomes but is largely unaddressed in the privacy literature, our second focus is on whether the effects of framing on upstream choices are in fact mitigated by more or less protective downstream privacy choices.

We address both of these goals of the manuscript in a series of online experiments. Across the experiments, we first evaluate whether the propensity to select protective disclosure settings (e.g., opting into encryption or restricting a use of personal information) is significantly impacted by subtle changes in decision *frames*: “the decision maker’s conception of acts, outcomes, and contingencies associated with a particular choice” (Tversky and Kahneman, 1985). Specifically, we focus on the impact of changes in decision frames induced by common variations in the presentation of these disclosure settings. Then, we evaluate whether impacts on upstream privacy choices due to exogenous changes in choice architecture are mitigated by different levels of downstream self-disclosure.

Experiment 1 focuses on the potential for subtle variation in choice architecture to impact consumers’ upstream choices. Specifically, it evaluates the impact of altering decision framings (via choice labels) on participants’ likelihood of selecting protective upstream disclosure settings. Experiment 2 tests the impact

of altering other choice labels on the likelihood of selecting protective upstream disclosure settings, and then evaluates whether levels of downstream self-disclosure differ in ways that mitigate upstream effects of framing on disclosure settings (e.g., frames that lead to choosing risky disclosure settings upstream are associated with lower levels of self-disclosure downstream). Experiment 3 evaluates the impact on downstream self-disclosure of simply assigning participants to protective versus risky disclosure settings. Experiment 4 evaluates the impact of a distinct framing manipulation (whether disclosure settings are presented as a choice to “allow” or “prohibit” a use of data) on the likelihood of selecting protective upstream disclosure settings, and again whether any impacts upstream are mitigated by different levels of downstream self-disclosure. At the same time, Experiment 4 again evaluates the impact on downstream self-disclosure of simply assigning participants to protective versus risky disclosure settings. Across the four experiments, we based framing manipulations on examples that are common in online privacy contexts. For example, choices to provide location information are sometimes presented using a “Privacy” label and sometimes using a more descriptive “Location” label. We also asked participants to make actual privacy choices and self-disclosures rather than hypothetical or self-reported choices, increasing the ecological validity of the experiments.

Across the experiments, we consistently found that various manipulations of decision frames that are common to privacy contexts and online services can significantly alter individual upstream choice of disclosure settings. In Experiment 1, participants were 58% more likely to choose protective disclosure settings (upstream choice) when setting labels were changed from “App Settings” to “Privacy Settings.” These results were confirmed in Experiment 2 where we compared a “Privacy Settings” label to a “Survey Settings” label. In Experiment 4, participants were 45% more likely to select more protective disclosure settings when they were presented with settings as a choice to prohibit a use of their personal information (reject frame) than when they were presented with the objectively identical setting as a choice to allow a use of their personal information (accept frame). However, while framing manipulations induced significant changes in the choice of protective disclosure settings upstream, downstream self-disclosure rates were nearly identical between these conditions. That is, participants *did not* mitigate

framing-induced effects on their upstream disclosure settings through different levels of downstream self-disclosure. In contrast, downstream self-disclosure levels did vary when, instead of being indirectly induced to choose different disclosure settings, participants were *assigned* to either protective or risky disclosure settings (see Experiments 3 and 4). Specifically, subjects assigned to less protective disclosure settings were significantly less likely to disclose sensitive information. This last result confirms that our observed lack of mitigation downstream is not because self-disclosure decisions are simply not sensitive to the disclosure settings we provided to participants.

This work makes a number of contributions. One contribution is to the literature on privacy decision making. In recent years, this literature has grown significantly (Acquisti, Brandimarte, and Loewenstein, 2015), reflecting the increasing attention towards privacy in a digital society and the economic trade-offs produced by data protection and disclosures (Adjerid et al., 2015; Goldfarb and Tucker, 2011). Within this literature, recent work (Acquisti, John, and Loewenstein, 2012; Brandimarte, Acquisti, and Loewenstein, 2013) has started to question the dominant normative perspective of privacy decision making which assumes that consumers have stable preferences for privacy, and make privacy choices solely based on their associated benefits and costs. Our results bolster this emerging behavioral perspective on privacy decision making which reveals the often incoherent and self-destructive decisions that consumers make. In particular, we highlight the importance of considering privacy decision making holistically as a complex process of heterogeneous and cascaded privacy choices. This allows us to extend the literature in various ways.

The first relates to the importance of framing effects in privacy decision making. A recent focus of privacy research is on understanding the impact of granular control provided by disclosure settings on privacy decision making (Xu et al., 2012; Brandimarte, Acquisti, and Loewenstein, 2013). Prior research has documented significant effects of providing consumers granular control (e.g., through disclosure settings) on privacy concerns and self-disclosure behavior, and suggests that these effects are independent of whether providing control actually decreases risk. For example, Xu et al. (2012) suggest that providing consumers control through disclosure settings can reduce privacy concerns even when control is

“illusory,” and Brandimarte, Acquisti, and Loewenstein (2013) show that the propensity of control provided through disclosure settings to reduce privacy concerns and increase self-disclosure persists even when the objective risks that consumers face are actually elevated. We highlight a different but related insight: the actual options chosen when an individual is presented with control via disclosure settings are malleable to changes in decision frames that are common in privacy contexts and that are largely independent of the objective costs and benefits of the setting.

A second way in which this work extends the behavioral privacy literature relates to cascaded privacy choices, and produces a more overarching insight. When consumers’ upstream choices are subtly influenced by choice architecture, they don’t seem to perceive a shift in their risk level and don’t adjust in downstream choices to mitigate this change in upstream risk. However, when consumers are simply assigned protective or risky upstream choices, they do seem to perceive a shift in their risk level and do adjust their downstream choices; this indicates both that they do care about their ultimate level of privacy, and suggesting directions that privacy regulation could take. Together, these results underscore the potential of an emerging behavioral perspective on privacy decision making to inform our understanding of privacy decision making in complex and heterogeneous online choice settings. Finally, these insights may also extend beyond the privacy context to other cascaded decision making contexts, particularly when choice architecture is likely to have powerful impacts on consumers’ decision making.

The research we report in this paper also contributes to a small but growing body of research aimed at addressing the gap in research at the intersection of behavioral economics, choice architecture, and technology-mediated choices. This gap has been noted by both decision sciences and information systems scholars. Specifically, Johnson et al. (2012) highlight the need to consider the intersection of choice architecture and technology contexts since “this interaction with decision technology is likely to increase in future years as computing devices become more unobtrusively integrated into our daily environment.” They conjecture that choice architecture could have powerful and persistent impacts in highly customizable and dynamic online decision settings. Goes (2013) echoes this sentiment and suggests that applying insights from behavioral economics to technology settings may yield important insights. Our

results illustrate how services online can easily manipulate the choices made by consumers through subtle changes in how options are presented, and how these manipulations can have persistent effects on behavior, but ambiguous effects on welfare.

2. Theory: Choice Architecture and Cascaded Privacy Choices

New information technologies have affected the nature of privacy choice. Increasingly, consumers' privacy decision making is cascaded: especially in online settings, privacy risks and benefits are determined by a combination of interdependent choices that consumers often make at different points in time. The scholarly literature on privacy has mainly investigated different types of privacy decision making in isolation from each other. For instance, significant efforts have focused on understanding self-disclosure behavior as a measure of privacy concern and decision making (e.g., Dinev and Hart, 2006). With the advent of the Internet and e-commerce, the literature has also evaluated consumers' willingness to engage in transactions with merchants under various degrees of privacy risk and concern (e.g., Kim, Ferrin, and Rao, 2008), their interest in using privacy-enhancing technologies (e.g., Fischer-Hübner, 2001), and their choice of privacy settings (Stutzman, Gross, and Acquisti, 2013). However, the cascaded nature of online privacy decision making—for instance, how upstream choices influence downstream ones—has been underexplored in the literature.

In addition, an emerging body of work on privacy decision making has focused on the role of heuristics and decision biases in consumers' privacy decision making. The dominant assumption in the literature that studies consumer privacy behavior has been that consumers are utility-maximizing, rational agents who possess reasonably stable, and thus predictable, preferences for privacy. These economically rational consumers are assumed to weigh costs and benefits in a rational fashion when deciding (for instance) what to disclose and what to keep private (Klopper and Rubenstein, 1977; Milne and Gordon, 1993; Dinev and Hart, 2006). In contrast, an emerging behavioral perspective has highlighted how factors arguably disconnected from the objective calculus of benefits and costs of information sharing can have a powerful influence on actual privacy choices (Moon, 2000; Acquisti, John, and Loewenstein, 2012). In fact, the behavioral decision research literature has highlighted how the design (or “architecture”) of

decision settings can influence, ameliorate, or even impair individuals' decision making, and therefore their welfare. This literature highlights how apparently insignificant changes in the design of services, policies, and interfaces can have significant impacts on behavior (Thaler, Sunstein, and Balz, 2014). For instance, seminal work on choice architecture finds that simply changing the default option for organ donation from opt-in to opt-out can dramatically change individuals' expressed willingness to donate (Johnson and Goldstein, 2003). In the context of privacy, researchers have investigated how privacy choice architecture can significantly affect individual privacy decision making (Almuhimedi et al., 2015). While the privacy literature as well as the broader choice architecture literature has focused on design changes that ostensibly improve consumer decision making (Johnson et al., 2012), manipulations of choice architecture do not always promote consumer welfare. Thaler, Sunstein, and Balz (2014) point out that "choice architects do not always have the best interests of the people they are influencing in mind" and that "wily but malevolent" architects can have devastating effects on the people who are influenced by them. This can also be the case when it comes to privacy settings, where choice architecture can be employed to shift consumers towards behaviors that primarily benefit data collection organizations (Acquisti, Brandimarte, and Loewenstein, 2015).

Together, the literature on behavioral privacy and that on choice architecture suggest that variation, whether naturally occurring or deliberately orchestrated, in the complexity and presentation of privacy decision making can significantly impact consumers' choices, and that these effects may have ambiguous welfare implications for consumers. Welfare implications are complicated by the cascaded nature of online choices, since evaluating the impact of alternative choice architectures operating at a single level of privacy decision making may not be sufficient to draw conclusions about impacts on overall consumer privacy risks and welfare. For example, effects of upstream choices that might seem destructive to consumers might be much less destructive if protective downstream decisions are taken that mitigate this upstream effect (see Table 1).

[Table 1: Examples of Cascaded Privacy Decision Making and Mitigation]

Upstream Choice	Downstream Choice	Mitigating Behavior
Selecting browser privacy settings (e.g., turn on private browsing)	Web browsing behavior	Less sensitive sites are visited when private browsing is chosen relative to when it is not chosen.
Deciding whether to log-in to a retail website	Purchasing behavior	Less sensitive purchases are made when user purchases anonymously as opposed to when logged in.
Selecting search engine settings (e.g., whether search history is stored)	Search behavior	Less sensitive searches are made when storage of search history is prohibited relative to when it is allowed.

In the remainder of section 2, we consider how consumers make privacy decisions in cascaded privacy decision settings, focusing on (1) the impact of varying choice architecture on upstream choices and (2) whether downstream privacy choices vary in ways that mitigate effects of choice architecture on upstream choices.

Because many combinations of upstream and downstream privacy choices are possible, we narrow our focus to a specific operationalization of the terms. For upstream privacy choices, we focus on disclosure settings, similar to what is presented to individuals on popular Internet services (e.g., Facebook privacy settings). For downstream choices, we focus on self-disclosure of personal information that is presumably used or shared in accordance with the disclosure settings selected upstream.

Our focus on disclosure settings is based on the critical role that both industry and policy makers have given to notice and consent mechanisms to manage online privacy (FTC, 2012). Under such mechanisms, providers of Internet services inform end-users about possible uses of their data, and provide them with diverse disclosure settings that offer some degree of control over the uses and visibility of their data. We consider the choices made via disclosure settings “upstream” because their overall impact can be mitigated or exacerbated by downstream privacy choices. For example, the risks of choosing to share a social media profile publicly (e.g., through a disclosure setting) can be mitigated or exacerbated by the downstream choice of whether to include sensitive information on that profile. For downstream choices, we focus on individual self-disclosure decisions, because they are a key measure of privacy decision making in both the rational and the behavioral privacy literature (Dinev and Hart, 2006; Acquisti, John,

and Loewenstein, 2012), and because the risk of upstream privacy choices is increasingly tied to subsequent self-disclosure decisions (again because these settings alter the visibility and use of these disclosures).

In the rest of the sections we propose a general framework of analysis and present three hypotheses, which we test in our experiments. The manipulations we utilize to test these hypotheses vary between experiments and are detailed preceding the experiment in which they are introduced.

2.1. *Upstream Choices: Framing*

With respect to upstream disclosure settings, we focus on whether framing—a central tool of choice architecture—can influence consumer selections of disclosure settings. Framing refers to the phenomenon of “simple and unspectacular changes” in the presentation of decision problems leading to changes in choice (Kühberger, 1998). Framing effects are generally attributed to changes in how individuals construe a decision context or stimulus they engage with. Such changes can include highlighting a particular dimension of a decision context or altering its perceived norms or goals. A substantial empirical and theoretical literature (Levin, Schneider, and Gaeth, 1998) has evaluated framing effects, focusing mostly on differences in behavior that arise from decision frames that highlight positive versus negative aspects of a given decision. In a task referred to as the “Asian disease task,” Kahneman and Tversky (1979) demonstrated that highlighting costs (lives lost) from a medical intervention versus the gains (lives saved) can lead to an increased preference for risky options; Levin and Gaeth (1988) found that perception of the quality of ground beef differs based on whether it is labeled as “75% lean” or “25% fat”; and Ganzach and Karsahi (1995) found that framing decisions in terms of losses (e.g., losses suffered from not using a credit card) is more effective at altering behavior relative to framing which highlights gains.

Framing effects are particularly relevant to disclosure settings since these settings can be highly complex, vary across services and even within a service over time, and are provided by firms and services, both online and offline, that maintain significant discretion over which options to present to end-users and how to present them. Recent privacy research casts doubt on whether these disclosure settings will actually be used by consumers in self-interested ways. In fact, extant research has highlighted the

paradoxical effect that mere *perceptions* of control (absent objective changes to risk) can have on consumer behavior—namely, increasing consumers’ willingness to disclose more, and more sensitive, information to others (Brandimarte, Acquisti, and Loewenstein, 2013). Solove (2013) argues that “organizations, as a rule, will have the sophistication and motivation to find ways to generate high opt-in rates” and Schwartz (2005) suggests that “many data-processing institutions are likely to be good at obtaining consent on their terms.” These concerns are 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.¹ However, this literature does not consider whether highly varied disclosure settings can introduce subtle changes in decision frames that can systematically and powerfully influence consumers’ choice of disclosure settings, potentially to their detriment.

Varying decision frames in privacy settings can have these powerful effects if brought to bear on privacy considerations for consumers who are distracted by other, often more salient and immediate, benefits of disclosure (e.g., accessing content, installing a useful application, etc.). Whether framing effects will materialize and be relevant in privacy contexts is, however, open to debate. Although the general phenomenon of framing is well established, it is clear from this literature that the strength of framing effects can vary by problem domain and context (Levin, Schneider, and Gaeth, 1998). For example, even seminal framing effects (e.g., those identified in the Asian disease task) have been shown to differ as the probability and size of potential losses or gains vary (Schneider, 1992).

The behavioral literature on framing offers additional relevant insights concerning the general features of decision contexts that exacerbate framing effects. First, the framing literature finds that these effects are especially prevalent when relevant information about the decision contexts is missing and/or when consumers lack relevant expertise (Schoorman et al., 1994). Privacy contexts are often

¹ Consider an example unrelated to privacy. In 2010, regulators required that banks halt practices of levying, by default, exorbitant fees for consumers who overdrafted their accounts. In response to the requirement that consumers be defaulted into a regime in which they would not be able to overdraw their accounts via ATM withdrawals, banks presented to their customers the choice to continue to be able to overdraw and incur these fees as the option to enroll in “overdraft protection.” A survey of more than 6,000 people administered by the Pew Center following implementation of the regulation found that large numbers of people had fallen for the ruse, despite their preference for having such transactions declined (Pew Center on the States, 2012).

characterized both by a lack of information about the risks and benefits of the decision setting and by asymmetries between what consumers know (or believe they know) and actual data use and privacy risks in popular online contexts (Balebako et al., 2012). This suggests that information about privacy decision settings which can reduce framing effects may often be lacking. Second, framing effects are exacerbated when there are low levels of deliberation by consumers. Takemura (1994) found that participants asked to carefully document their decision making process did not exhibit framing effects. Although numerous surveys (e.g., Turow and Hennessy, 2007) indicate that privacy is an issue that arouses significant concerns among consumers, this level of deliberation is unlikely in privacy settings since privacy considerations can be secondary to more salient decision points or goals in online contexts (e.g., downloading an application, buying a product, etc.). Prior research suggests that such considerations are only likely to be activated when contextual factors raise the salience of privacy concerns (Acquisti, John and Loewenstein, 2012), and identifies a common *disconnect* between stated privacy concerns and the level of actual concern revealed by behaviors (see, for instance, Jensen, Potts, and Jensen, 2005). This suggests that framing effects are likely to be substantial in the more common situation where privacy contexts fail to elicit high levels of concern.

Taking these overlaps into account, we present the follow hypothesis:

H1: Variation in the framing of upstream disclosure settings that has arguably no effect on the objective benefits and costs of data allowances will influence the selection of protective disclosure settings.

2.2. Downstream Privacy Choices

The impact of framing on choice of protective disclosure settings captures only part of the sequence of cascaded privacy decision making. The choice of disclosure settings alone does not represent the ultimate privacy costs and benefits experienced by an individual who is using an online service. These ultimate privacy costs and benefits also depend on subsequent downstream privacy decisions, including, commonly, how much and what to disclose. The propensity of individuals to alter their downstream

behaviors in response to upstream privacy choices and risk is a critical assumption of both policy makers (FTC, 2012) and industry advocates of these regulatory approaches for protecting privacy. In its response to FTC regulation, for example, Facebook suggested that consumers who use its service are able to engage in “corrective measures” if they are dissatisfied with the privacy options they are provided with. Specifically, Facebook suggests that consumers can start by “sharing less information” or that consumers can even “terminate their relationship altogether” (Richter, 2011). Google similarly highlights that its customers can “liberate” their data (meaning that they can download and export all of their data from Google services) and switch to other providers if they are dissatisfied with the protections and options provided to them (Chavez, 2011).

To understand the importance of cascaded privacy decision making, and the relevance of framing effects to this discussion, consider an analogy: the decision context of a driver renting a car from an airport and navigating a difficult winter road. The driver has to make a number of related choices that, together, will influence various outcomes, such as her overall costs, safety, the damage to her vehicle, time to destination, and so forth. Those choices might begin with which car to rent and whether to take out optional insurance, continue with the consideration of which route to take, and then proceed with how to drive when on the chosen route. The ultimate outcomes of the trip are a function of all of these related decisions (e.g., the driver’s safety will be impacted by the choice of vehicle, of route, and of care in driving). In this driving scenario, considering a single choice (such as her choice of a route) would fail to capture the complete picture of the benefits and risks the driver faces during the trip. Similarly, turning back to privacy, focusing only on the effects of framing on consumers’ upstream selection of different protections (H1) would fail to capture the complete picture of the consumer’s privacy benefits and risks. Rather, we need to also consider whether downstream privacy decisions—in our case, what the consumer will choose to disclose via a service—will mitigate effects of framing-induced changes in upstream privacy choices. Can we expect framing manipulations that lead to less protective disclosure settings upstream to be associated with lower levels of self-disclosure downstream?

Whether we observe this mitigation hinges critically on the extent to which consumers compensate in their downstream behavior as a result of changes in risks upstream. The extant privacy literature has evaluated behavioral reaction to changes in privacy protections, and suggests that such compensation is likely. The privacy literature has shown time and time again that individuals compensate in relevant privacy behaviors when their level of protection is altered. Culnan and Armstrong (1999) found that the use of fair information practices by firms can engender trust from consumers, reducing privacy concerns and perceived risks of disclosure; Miyazaki and Krishnamurthy (2002) and Hui, Teo, and Lee (2007) found a significant effect of privacy seals on consumer perception of firm privacy practices and their willingness to disclose personal information. These results, if applied to our setting, imply that differences in chosen protections would also result in different levels of downstream self-disclosure.

Perspectives from other fields (for instance, economics) also argue that consumers will compensate for more or less risky upstream choices in related downstream choices. In particular, a number of works suggest that individuals will be more likely to engage in a risky downstream behavior if they can reduce or shift the cost of that behavior via earlier choices. Prior studies have found that consumers who purchase health insurance may subsequently increase their healthcare utilization and their risky behavior. For example, Dave and Kaestner (2009) found that obtaining health insurance reduces prevention and increases unhealthy behaviors among elderly men; Kelly and Markowitz (2009) found that having higher body mass is associated with also having insurance; and Klick and Stratmann (2007) found that legal mandates requiring that insurers cover diabetes care increased the body mass index of diabetes patients. These impacts are attributed to the effects of moral hazard, since initial choices by consumers (such as purchasing health insurance) shift the cost of risky downstream behaviors to third parties.

Relatedly, theories of risk compensation and homeostasis theory predict that individuals will alter their behavior in the face of factors or choices that make a particular decision context more or less risky (Peltzman, 1975; Wilde, 1981). For instance, Peltzman (1975) argues that standards that force manufacturers to make safer automobiles, or regulation that requires drivers to wear seatbelts, would be ineffective, in aggregate, because drivers would simply compensate by driving faster or more recklessly.

Peltzman found confirmatory evidence for this theory: automotive safety standards intended to reduce traffic deaths caused 10,000 additional deaths, primarily to non-vehicle occupants. Subsequent work also found evidence of the compensatory reactions theorized by Peltzman, but found that benefits from safety regulation outweighed the harm from compensatory reckless driving (Crandall and Graham, 1984; Levy and Miller, 1999). Theories of moral hazard as well as risk compensation and homeostasis would predict, similar to the extant privacy literature, that individuals will compensate for higher upstream privacy risk in their downstream self-disclosure behavior (i.e., disclose less) and, in doing so, be likely to mitigate the impact of choosing riskier disclosure settings upstream.

Our focus, however, is somewhat different from the situations described by prior works, in that we are interested in whether compensatory reactions occur when changes in upstream choices are induced via subtle manipulations in the framing of decision settings. We conjecture that, unlike what has been documented in the prior literature, changes in risk induced by subtle manipulations of framing of upstream decisions may not result in the same mitigating behavior by consumers downstream.

One reason for this is that implicit in all of these prior theories is the assumption that individuals are cognizant that their risk has been altered, such that the risk is now perceived as out of sync with their desired levels. For example, regulation that requires all drivers to wear a seatbelt (and which is then subsequently enforced by law enforcement), represents a salient shift in their risk. Drivers who would, absent this regulation, drive without a seatbelt are cognizant of the change in their risk level and may have compensatory reactions that result in mitigating behavior downstream (e.g., faster driving). In contrast, subtle manipulations of choice architecture may not even be noticed by consumers, leading them to attribute their shift in behavior to changes in their own preferences or a rational response to their context or situation. Even if noticed, consumers may fail to anticipate the effect of changes in choice architecture on their own behavior (Lieberman, Samuels, and Ross, 2004) and may not alter their behavior even when told their architecture has been manipulated (Loewenstein et al., 2015). Lieberman, Samuels, and Ross (2004) found that individuals consistently underestimate the impact of subtle variation in decision contexts on their own behavior even when it is pointed out to them and they are asked to deliberate on its

potential impacts. This suggests that, unlike Peltzman's context, individuals exposed to framing manipulations may simply not be cognizant that their chosen protections have actually been shifted, making compensatory reactions less likely. In fact, Loewenstein et al. (2015) found that the impact of choice architecture (in their case, a manipulation of defaults) persists even when consumers are either pre-alerted that the architecture has shifted, or informed after the fact and then allowed to change their choices. These results suggest that framing effects may persist even when individuals are aware that their decision frame has been manipulated—in which case, it seems plausible, they would be less likely to attempt to compensate for such effects. If these arguments hold and consumers fail to compensate for framing effects that result in riskier disclosure settings chosen upstream, it is unlikely that we will observe lower levels of downstream self-disclosure that mitigate this additional risk.

Even if some (potentially weak) compensatory reactions do occur, we may still fail to observe differences in downstream self-disclosure that mitigate the risk from framing-induced changes in upstream choices; if framing manipulations also have a direct effect on downstream behaviors, this may counteract any compensating behavior by consumers. Specifically, we argued previously that framing effects likely influence decision making via their impacts on consumers' construal of decision context and the considerations that are highlighted (or not). In privacy settings, different frames may differentially highlight privacy considerations, thus influencing initial upstream choices. If these considerations remain highlighted (or subdued) as participants transition to downstream self-disclosure choices, this may result in sustained protective (or open) self-disclosure decisions. This may neutralize the effects of compensatory reactions since consumers who choose less protective settings (as a result of framing) may actually go on to disclose in a more open, less protective, fashion. To make these arguments more concrete, consider again the context of the driver renting a car. Assume she had not intended to purchase additional insurance, but is compelled to purchase it by a subtle manipulation of her decision frame that, for example, highlights the potential risks of not purchasing such insurance. Although she is aware that she now has additional car rental insurance, she may have little to no compensatory reactions in her downstream choices (e.g., how carefully she drives) if she doesn't perceive a deviation from her desired

level of risk from the change in her upstream choices. In addition, if framing manipulations impact our driver's decision to purchase additional insurance (i.e., more protection) by highlighting potential risks from driving in an unknown area, this elevated focus on driving risks could result in more careful driving (counteracting any compensatory reactions to having insurance that do occur).

Either a lack of compensatory reaction by consumers, or a direct effect of framing on self-disclosure that counteracts these compensatory reactions, will support our conjecture that choice architecture impacts upstream will not be mitigated by different levels of downstream self-disclosure.

Thus, we hypothesize:

H2: Individuals will disclose sensitive information downstream at similar levels between manipulations of framing, despite objective shifts in protective disclosure settings chosen upstream.

While disentangling the independent effect of the mechanisms behind H2 is beyond the scope of this manuscript and would make our empirical investigation unwieldy, we do posit a counter-factual hypothesis to help validate our arguments in support of H2. In particular, we argued previously that the lack of mitigation in downstream self-disclosure levels theorized in H2 is because framing-induced shifts in upstream choices (and thus risk) differ from the shifts in risk that were the focus of prior literature (e.g., risk compensation and moral hazard literature, privacy literature). If this is the case, then shifting the same upstream privacy risks impacted by framing in a manner more in line with what the prior literature has considered (e.g., mandatory regulation) should result in the predicted mitigation by consumers downstream. In the case of our driver renting a car, this suggests that state regulation or a workplace policy requiring her to purchase additional insurance would, in contrast to the case where she purchases insurance because of changes in her choice architecture, result in differences in her downstream behavior. In privacy settings, a comparable scenario is one where risky or protective disclosure settings upstream are not shifted via framing manipulations but directly assigned to consumers such that their risk is being more directly shifted. In this scenario, the change in risk upstream would more likely be perceived as

deviating from a consumer's desired level of risk, making strong compensatory reactions by consumers that result in differences in downstream self-disclosure more likely.

Thus, we hypothesize:

H3: Individuals will disclose sensitive information downstream at lower levels when they are directly assigned to less protective disclosure settings upstream.

3. Methodology

In four randomized experiments, we evaluate the impact of framing on individuals' upstream disclosure settings (H1), whether impacts of framing upstream are mitigated by different levels of downstream self-disclosure (H2), and whether direct assignment to different upstream disclosure settings is mitigated by different levels of downstream self-disclosure (H3).

Experiment 1 evaluates how the propensity of individuals to select protective disclosure settings (upstream privacy choices) is impacted by different manipulations of choice framing (H1). Participants in this experiment were told that they were being recruited for a pilot study that involved the installation and use of a mobile application focused on financial management. They were then asked whether they wanted to grant permission for this mobile application to use various personal data on their mobile devices. Participants were told that their decisions would be implemented in the actual app they would need to install and use as part of the study. The primary dependent variable in this experiment is the decision to grant a particular data permission.

Experiments 2 tests both whether framing has an impact on upstream choices (H1) and whether downstream self-disclosure differs as a function of manipulations of upstream choice frames (H2). Participants in this experiment were asked to take an online study that would require them to select a number of disclosure settings that governed how their responses in the study would be shared and used. These choices were deliberately influenced using experimental manipulations similar to those used in Experiment 1. Participants were then asked to provide responses to several questions related to personal,

including sensitive, behaviors. Although the ostensible goal of the experiment was to investigate participants' engagements in various behaviors (e.g., "Have you ever looked at pornographic material?"; see examples in Appendix A4), we were not interested in the behaviors per se, but rather in participants' willingness to disclose information about engaging in them. Because all of our experiments use random assignments to the different conditions, we can assume the distribution of participants' actual past engagement in these activities to be similar across conditions. Thus, higher or lower admission rates across conditions signal an impact of our experimental manipulations. Prior research has successfully used this approach to examine privacy-sensitive behaviors (e.g., Acquisti, John, and Loewenstein, 2012; Moon, 2000).

Experiment 3 focuses on H3 and examines whether levels of downstream self-disclosure differ when participants are randomly assigned to different levels of protective disclosure settings upstream. This experiment uses disclosure settings identical to the upstream choices presented to participants in Experiment 2. Experiment 4 simultaneously tests H1 and H2 but uses a different framing manipulation than Experiments 1 and 2 to test whether our findings are generalizable and not unique our specific manipulations. We simultaneously test again whether assignment to different upstream protections will influence downstream self-disclosure (H3). A summary of our experiments is in Table 2 below.

The four experiments were conducted using Amazon's Mechanical Turk (AMT), an online crowdsourcing service that has become increasingly popular among social scientists for conducting online experiments.² The validity of AMT samples in behavioral experiments has by now been investigated in a multiplicity of studies. Buhrmester, Kwang, and Gosling (2011) demonstrate that AMT samples are just as representative as other Internet samples, and considerably more representative than typical student samples. Steelman, Hammer, and Limayem (2014) found that AMT samples have psychometric properties that are similar to those of both student and consumer panels, and that using U.S. AMT

² We restricted participants to subjects from the United States with a hit approval rate on AMT of over 95%. We included attention check questions at the start of the questionnaire following accepted practices in the field (e.g., Oppenheimer, Meyvis, and Davidenko, 2009). We also included a screening survey which both prevented individuals from participating in a given experiment multiple times and prevented individuals from participating in more than one experiment.

samples replicated validated results from the technology acceptance model. Furthermore, judgment and decision making experiments using AMT samples have replicated results found in traditional subject samples (Goodman, Cryder, and Cheema, 2013).

To select manipulations of decision frames to use in our experiments, we surveyed approaches currently employed by online services for soliciting consumer privacy choice, specifically seeking subtle variation in these approaches that had the potential to highlight or downplay consumer privacy concerns. This focus on subtle variations among existing privacy choice mechanisms allows us to identify manipulations of decision frames that consumers and policy makers may not be likely to identify as significant influences in their own choices, and that are of immediate relevance to the design of these mechanisms and consumer outcomes.

Table 2: Overview of Experiments

Exp	Framing Manipulation	Upstream / Downstream Choices	Purpose
1	Label Frame: “Privacy Settings” vs. “App Settings”	Privacy Settings / N/A	Evaluate the effect of choice framing on upstream privacy choices (H1).
2	Label Frame: “Privacy Settings” vs. “Survey Settings”	Privacy Settings / Self-Disclosure	Evaluate the effect of choice framing on upstream privacy choices using a different label framing manipulation (H1). Evaluate whether levels of downstream self-disclosure differ as a result of manipulations of choice framing upstream (H2).
3	N/A	N/A / Self-Disclosure	Test whether downstream self-disclosure differs when participants are randomly assigned to different disclosure settings upstream (H3).
4	Accept vs. Reject Frame	Privacy Settings / Self-Disclosure	Evaluate the robustness of the effect of choice framing on upstream privacy choices using different manipulation (accept/reject framing) (H1); compare whether levels of downstream self-disclosure differ as a result of manipulations of choice framing upstream (H2); and simultaneously evaluate the impact of assignment to protective/risky disclosure settings upstream on downstream self-disclosure (H3).

3.1. Measures and Estimation

The two dependent variables captured in our experiments are (1) the selection of protective disclosure settings (upstream choice) and (2) downstream self-disclosures (downstream choice). Since both variables

are captured as repeated measures (e.g., participants made multiple choices of settings and multiple self-disclosures), we use a panel random effects regression as a primary estimation approach, while correcting standard errors for the non-independence of multiple responses from a single participant (Zeger and Liang, 1986).³

$$[ProtectiveSetting_{ij}, Admit_{ij}] = \beta * Treatment_i + \alpha * Y_i + \theta_i + u_{ij}.$$

The first dependent variable (*ProtectiveSetting_{ij}*) is a binary measure capturing whether a participant *i* chose the protective option (e.g., denying an intrusive use of their self-disclosures or opting into encryption) for disclosure setting *j*. The second dependent variable (*Admit_{ij}*) is a binary measure of whether a participant *i* admitted to engaging in a sensitive behavior *j* (i.e., self-disclosed sensitive personal information). In some specifications, we also include Y_i , a vector with controls for participant-specific characteristics (e.g., age, gender, etc.). θ_i is the participant-specific random effect, and u_{ij} is the error term. Estimates on randomly assigned treatments (*Treatment_i*) are unbiased, because they are uncorrelated with observed (Y_i) and unobserved (θ_i) individual differences and the error term u_{ij} . Although our controls are not necessary for the unbiased estimation of the effect of our treatments on disclosure behavior, we include them in some specifications to rule out breaks in randomization and account for some of the variation in disclosure behavior between participants.

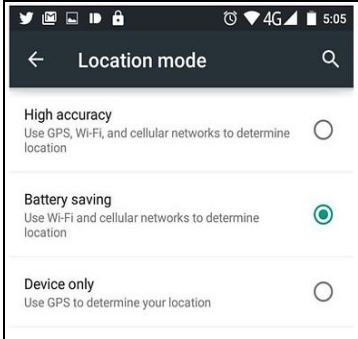
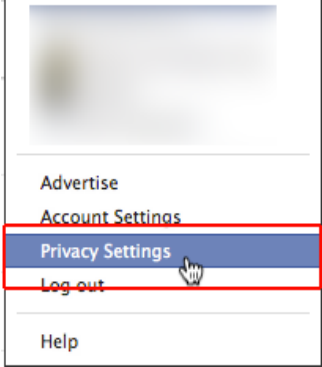
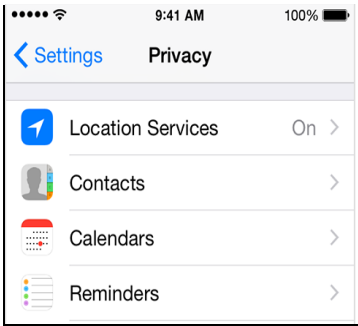
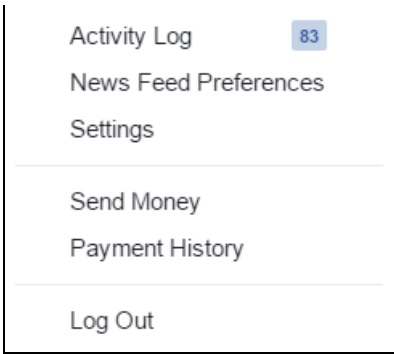
4. Experiment 1

Experiment 1 focuses on Hypothesis 1. It evaluates the impact of framing changes to the labeling of privacy-relevant choices on consumer choice of privacy protections. Online services often obfuscate privacy-relevant choices under frames that are not presented to consumers as “privacy” choices. For instance, consider Figure 1. The Android mobile platform presents choices with significant privacy implications using descriptive labels such as “Location Settings” (at the time of this manuscript, Android mobile platforms do not have any settings labeled “Privacy Settings”), whereas Apple iOS 7 presents

³ We opt for a linear probability model estimation in lieu of a non-linear estimation approach (e.g., logit) to avoid inaccurate coefficient and standard error estimates for interaction effects in non-linear regression models (Ai and Norton, 2003). Angrist and Pischke (2008) have shown little qualitative difference between the logit and linear probability specification.

similar choices to consumers under the general “Privacy” label (Figure 1). Similarly, Facebook altered the label of the settings on the main Facebook page from “Privacy Settings” in 2013 to simply “Settings” in 2014; this change has persisted into 2016 (see Figure 1). Similar variation persists in the detailed settings they offer (e.g., “Privacy Settings and Tools” versus “Timeline and Tagging Settings”). Although such changes in labels may appear to be subtle, extant behavioral literature has posited that minor changes in the labels can have powerful effect on consumer decision making. We extend the findings in this literature to the context of privacy. We investigate how changes in the labeling of privacy choices may impact individual choice of protective disclosure settings.

Figure 1: Variation in Setting Labels

Mobile	Social Media
<p data-bbox="500 863 672 892"><i>Android Lollipop</i></p> 	<p data-bbox="894 863 1170 892"><i>Facebook Main Page (2013)</i></p> 
<p data-bbox="558 1325 613 1354"><i>iOS 8</i></p> 	<p data-bbox="894 1325 1170 1354"><i>Facebook Main Page (2016)</i></p> 

The potential for variations in labels to influence decision making is substantiated by an extensive body of behavioral research. Burnham, McCabe, and Smith (2000) find a strong impact on cooperation in

a two-player reciprocity game when labeling participants as either a “partner” or an “opponent.” They attribute these effects to subconscious priming of trust in the other player in the game and suggest that “people form expectations about each other’s intentions using mental modules which process contextual information.” Epley, Caruso, and Bazerman (2006) use the more generic labels of “strategic competition game” and “cooperative alliance game,” and also find differences in participant behavior. They suggest that “changing the way the game is described is likely to change participants’ beliefs about the normative responses of other participants and to alter participants’ own behavior in turn.” Liberman, Samuels, and Ross (2004) reinforce these results and find that participants are more likely to cooperate when playing the “Community Game” as opposed to the “Wall Street Game.” They attribute this effect to a slightly different mechanism however, and suggest that differences in the description of the tasks or games evokes variation in the perceived norms or goals of the game (e.g., business dealings versus ethical dilemmas).

In privacy contexts, changes in the label of choices may influence behavior if they alter individuals’ subjective predictions about the probability that their data would be used in a privacy-invasive manner by the data requestor. Also, changes in the label of choices may alter the individual’s construal of the decision goals that are most relevant in a decision context. For example, certain labels may highlight privacy protection goals that may otherwise be subdued by other competing goals or norms in a setting (e.g., attaining some immediate benefit from data disclosure). Prior work substantiates the potential for these effects to emerge: Acquisti, John, and Loewenstein (2012) found that subtle contextual cues can influence self-disclosure behavior by priming or subduing privacy considerations. Thus, we leverage variation in the labeling of upstream privacy choices to evaluate H1 and conjecture that, whether intentional or justified by other platform constraints, labeling otherwise identical choices as “Privacy Settings” relative to an alternative descriptive label will alter the decision frame, resulting in the choice of more privacy-protective options.

4.1. Design

Experiment 1 consisted in a 2-condition between-subject design in which participants were asked to choose between different data permission settings, and were randomly assigned to either a condition in

which the permissions were presented under a “Privacy Settings” label, or one that used an “App Settings” label. Participants were recruited via AMT to take a screening survey for (ostensibly) identifying pilot-testers of a new mobile application. Participants were told that this application organizes users’ financials and identifies opportunities for savings by analyzing monthly bills and credit card statements. They were informed that, if chosen for the pilot study, they would be asked to install and use the app for a four-week period as well as provide some of their financial information (e.g., bank account numbers, etc.). The particular context of a financial management application was chosen to ensure a relevant privacy dimension of app usage. In order to minimize participant suspicions regarding our deception, we created a AMT requestor account named “AAG Mobile Technologies.” We also ran a short survey prior to the main study asking descriptive questions about mobile phones and apps usage (participants from this initial survey were excluded from the experiment). Overall, our study was understandable and our deception was effective: following the debrief, participants admitted in their free text response to being fooled by our deception, and indicated that they would be genuinely interested in the app described.

4.2. Procedure

Participants were first asked questions about their demographics (gender, age, race, education, etc.), their mobile carrier, the phone they currently own, and their usage of applications to manage their finances. Participants were then presented with the option to allow or deny three permissions randomly selected from a subset of six total permissions modeled after current categories of permission types on mobile platforms (for instance, they were asked whether they would allow our mobile application to collect their location information; the complete list of permissions can be found in A1 in the Appendix).⁴ To make these selections relevant to participants, participants were told that their selections in the survey would be used by the app if they were selected for the pilot study. To introduce a cost for choosing restrictive settings, participants were also told that, while their choice to limit some permissions would not influence

⁴ We selected a subset of three permissions to avoid the potential of a high number of permissions triggering privacy concerns for individuals. Elevated baseline concerns may impact framing effects, lead participants to drop out of the study, or cause participants to make their choices arbitrarily (e.g., accept or reject all permissions).

their chances of being selected for the pilot, it would influence the functionality available in the app. After making their choices, participants were provided a debrief to address the deception used in the study and were allowed to opt-out of the study if they desired. They were also given an additional bonus payment to account for the fact that there wasn't an actual pilot study (which they would have been additionally compensated for).

4.3. Results and Discussion

One hundred and five participants completed the study, four of whom opted out of the study after being given the debrief at the end of the study, leaving 101 usable responses ($M_{Age} = 33$, $SD_{Age} = 9.9$, $M_{Male} = .56$). As noted, the experiment had two conditions manipulated between subjects: the permissions provided to participants were presented using either a "Privacy Settings" label or the more neutral "App Settings" label. Every other feature of the experiment was identical between conditions. For each participant, we calculated the proportion of denied permissions relative to total requests and found support for H1: participants presented with choices labeled "Privacy Settings" were 58% more likely to choose the more protective choices relative to those presented with the same choices as "App Settings" (47% vs. 30%, $t(99) = 2.4573$, $p = .02$).

Table 3: Experiment 1 Results

VARIABLES	<i>Upstream Disclosure Settings</i>	
	(1) <i>ProtectiveSetting</i>	(2) <i>ProtectiveSetting</i>
PrivacyLabel	0.173* (0.0699)	0.165* (0.0682)
Age		-.00003 (0.00387)
Male		-0.0515 (0.0704)
White		-0.0148 (0.0804)
College		0.0467 (0.0711)
Setting Controls	NO	YES
Constant	0.299** (0.0487)	0.0706 (0.141)
Observations	303	303

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

A random effects panel regression confirms this finding with a positive and significant coefficient estimate on *PrivacyLabel* (Table 3, Column 1). A second regression confirms that the results are consistent when including controls for demographics and dummies for the specific setting presented to participants (Table 3, Column 2).⁵ 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 (H1 supported). To the extent possible, we provided a context similar to the real world, and provided participant decisions that involved actual risks and benefits to them. Overall, our results from Experiment 1 provide evidence that subtle variation in framing can produce substantial shifts in the propensity to choose protective disclosure settings.

5. Experiment 2

In the first experiment, we found that subtle manipulations of choice framing had a significant effect on individual choices of privacy protection levels, supporting the first hypothesis in our manuscript. Similar to Experiment 1, in Experiment 2 we consider disclosure settings as the upstream choice and seek to reaffirm that subtle variation in the presentation of these settings can alter individual choice of protective disclosure settings. In addition, in Experiment 2 we also evaluate the second hypothesis of the manuscript, that downstream decision making will not vary following changes in initial choices of privacy protections. Specifically, we asked participants to make downstream self-disclosures of sensitive information about themselves. We considered these choices downstream (relative to upstream disclosure settings) because participants first had to choose whether to allow or deny particular uses of these sensitive disclosures (e.g., whether other participants of the study could view their sensitive disclosures) and then had to decide what disclosures to make. As such, the risk associated with upstream choices can be significantly attenuated by modifying downstream self-disclosure behavior. This allows us to assess

⁵ Including dummies for actual settings presented accounts for any breaks in randomization for the subset of permissions shown to participants.

H2 by evaluating whether downstream self-disclosure decisions vary following framing-induced changes in chosen in disclosure settings upstream.

5.1. Design, Pre-Study, and Procedure

Participants on AMT were invited to participate in an online task advertised as a study on ethical behavior which would require them to first select a number of disclosure settings that govern how their responses in the study would be shared, and then answer several questions related to sensitive behaviors (a sample question was presented to participants in the introductory text). The advertised study context of ethical behavior is a shift from Experiment 1 and is a validated context to evaluate framing effects while also offering an opportunity to study related self-disclosure behavior (Acquisti, John, and Loewenstein, 2012). Also, it avoids the need to use deception with participants.

When choosing which disclosure settings to provide to participants in the study, we were concerned that disclosure settings that were either too irrelevant or too invasive might result in little variation in what participants decide, irrespective of any experimental manipulation. To address this concern, we recruited 104 participants ($M_{\text{Age}} = 31$, $SD_{\text{Age}} = 12.4$, $M_{\text{Female}} = .34$) 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 Experiment 2. Participants were then asked to evaluate the extent to which they would want the choice to opt-out of (or opt-in to) various uses and handling of their responses. 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 with the choice (see results in A2 in Appendix).

Using the data from this pre-study, we finalized the design of Experiment 2. Specifically, Experiment 2 consisted of a 2x2 between-subject design in which we manipulated the label and also the set of disclosure settings provided. We manipulated whether choices were presented to users using either a “Privacy Settings” or a “Survey Settings” label; in addition, we assigned participants to either “High Importance” choice sets (the four highest ranked disclosure settings) or “Low Importance” choice sets (the four lowest ranked disclosure settings; see A3 in Appendix).

Participants were first shown an introductory screen that described the study context and provided an example of the sensitive questions asked in the study. Participants were then asked demographic questions, which included no directly identifying information but asked for their city and zip of residence as well as other demographic information.⁶ They were then provided with four choices that related to the use and storage of their responses in the survey. Depending on the condition, participants were shown either high importance or low importance settings which were presented as either “Privacy Settings” or “Survey Settings.” Participants were then presented with eight questions representing sensitive disclosures (see A4 in Appendix). The questions used were the ones rated as intrusive in Acquisti, John, and Loewenstein (2012), and were presented in random order.

5.2. Results and Discussion

Two hundred and four individuals ($M_{\text{Age}} = 28$, $SD_{\text{Age}} = 9.7$, $M_{\text{Male}} = .66$) participated in the experiment. We replicated the results of Experiment 1 and 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%, $t(202) = 2.1729$, $p = .03$). Comparing conditions with high and low importance settings, we found that this effect was driven by participants presented high importance settings (42% vs. 28%, $t(99) = 2.212$, $p = .03$). For low importance setting choices, perhaps due to a floor effect, the effect of the label was insignificant (“Privacy Settings” 7% vs. “Survey Settings” 4%, $t(101) = 1.039$, $p = .3$). Because we did not find significant differences in choice of settings for low importance settings, we focused our subsequent analysis on framing conditions of those participants provided with high importance settings. A random effects panel regression (Table 4a, Columns 1 and 2) confirms our initial finding with a negative and significant coefficient estimate ($\beta_{\text{PrivacyLabel}} = -.14$, $p = .026$) for the main effect of the “Privacy Settings” (H1 supported).

⁶ 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.

[Table 4a: Experiment 2 Results – Upstream Choices]

VARIABLES	Upstream Disclosure Settings	
	(1) <i>ProtectiveSetting</i>	(2) <i>ProtectiveSetting</i>
PrivacyLabel	0.137* (0.0615)	0.119* (0.0599)
Age		0.00148 (0.00287)
Male		0.00129 (0.0673)
White		-0.0512 (0.0629)
College		0.132* (0.0663)
Constant	0.280** (0.0425)	0.216+ (0.111)
Observations	404	404

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

Second, we evaluated whether the impact on protective disclosure settings of randomly assigned framing manipulations (where we can expect privacy concerns and underlying propensity for self-disclosure to be comparable between conditions) is mitigated by different levels of downstream self-disclosure. We found that, despite the differences in disclosure settings, participants in different conditions did not exhibit different levels of self-disclosure. Participants who were presented with a “Survey Settings” label admitted to behaviors at a rate comparable to those presented with “Privacy Settings” (53.18% vs. 53.65%, $t(97) = .12, p = .92$).⁷ This is again confirmed in our random effects panel regression (Table 4b, Columns 1 and 2) with a near zero and insignificant estimate on the effect of *PrivacyLabel* on admission rates (H2 supported).

Combined, these results have a number of important implications. First, they reinforce our first hypothesis that subtle variation in the framing of disclosure settings can significantly alter an individual’s propensity to select protective disclosure settings, particularly for choices that are of high importance to individuals. More significantly however, we also find support for our second hypothesis that participants will disclose sensitive information downstream at similar levels between manipulations of framing, despite objective shifts in protective disclosure settings chosen upstream.

⁷ Unsurprisingly, participants presented with low importance settings also did not exhibit differences in their self-disclosure.

[Table 4b: Experiment 2 Results – Downstream Self-Disclosure]

Downstream Self-Disclosure		
VARIABLES	(1) <i>Admit</i>	(2) <i>Admit</i>
PrivacyLabel	-0.00460 (0.0446)	0.00380 (0.0413)
Age		-0.00214 (0.00245)
Male		-0.00234 (0.0423)
White		0.0902+ (0.0531)
College		-0.0620 (0.0468)
Constant	0.536** (0.0284)	0.563** (0.0809)
Observations	792	792

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

We note that evaluating the direct relationship between the protectiveness of disclosure settings organically chosen by participants and their downstream self-disclosure is not particularly informative in our context (e.g., as a way to test H3). This relationship is endogenous in practice—individuals who are more privacy conscious will choose more protection and also disclose less than other participants. If the effect of unobserved privacy concerns dominates, then choosing more protection would correlate with less self-disclosure, and conversely if impact of the objective levels of privacy protection dominates. If these effects cancel each other out, there would be no observable correlation between chosen protection level and downstream self-disclosure.

As we noted in our theory, the lack of mitigation downstream may be due a combination of different phenomena. For instance, the impact of choice architecture on upstream disclosure settings is not sufficiently heavy handed to elicit the compensatory reactions identified by the extant privacy literature as well as some of the economics literatures. As a result, we do not observe mitigation of upstream effects and no differences in self-disclosure. The second possibility is that individuals do, in fact, compensate when framing induces shifts in their upstream disclosure settings, but this compensatory reaction is counteracted by direct effects of framing on downstream self-disclosure (see Section 2.2 for more detail on this). Although, we cannot distinguish between these two potentially concurrent mechanisms, the

finding that upstream impacts of framing are not mitigated downstream by different levels of self-disclosure still has significant implications for consumers navigating complex and cascaded privacy decision settings. Specifically, it suggests both that individual privacy risks can be impacted by subtle variability in the presentation of upstream decision settings, and that this effect persists due to a continued propensity towards data allowances and disclosure in downstream privacy decision making.

6. Experiment 3

Our first two experiments provide evidence that a subtle manipulation of decision frames impacts initial choice of protective disclosure settings (H1 supported), but that downstream self-disclosure decisions fail to mitigate this effect and are nearly identical between manipulations of choice frames (H2 supported). Experiment 3 focuses on H3 and evaluates whether simply assigning participants to protective versus risky disclosure settings would in fact impact downstream self-disclosure (H3).

In addition to the theoretical relevance of H3, testing this hypothesis addresses a more practical concern in our empirical setting. Namely, an alternate explanation for why we do not observe differences in downstream self-disclosure that mitigate the impact framing has on upstream disclosure settings is that participants may simply not be sensitive to any level of the disclosure settings provided to them in the experiment. That is, participant self-disclosure in our experiment may not vary, no matter how protective or risky the disclosure settings chosen upstream are. This may occur if, for example, participants assume some degree of anonymity in our experimental setting. We sought to preempt this concern in the first two experiments by not assuring participants that they would be anonymous. In fact, a notice of confidentiality protections was purposefully excluded from the consent form (with IRB approval) to avoid this particular concern. Rather, we provided a debrief at the end of our experiment which assured them that their data would be handled in the most protective manner possible. In addition, it is known in the AMT community that they are not anonymous, since their AMT activities can be linked to their Amazon

account.⁸ Nonetheless, Experiment 3 alleviates this concern further by evaluating whether participants' self-disclosure behavior is sensitive to changes in the protectiveness of upstream disclosure settings.

6.1.1. Design and Procedure

Participants from AMT were invited to again take a study using the same “ethical behavior” context as in Experiment 2. For each of the three disclosure settings rated as most important in Experiment 2, we randomly manipulated (via a graphical notice) whether the disclosure setting was set to be protective (e.g., their responses would not be shared with other participants of the study) or not (see A5 in Appendix). This resulted in participants being shown privacy notices where, depending on the condition, all three disclosure settings were assigned to be protective all the way down to where none of the disclosure settings were assigned to be protective. We then asked participants to answer the same eight sensitive questions used in Experiment 2. Recall that because of random assignments to the different conditions, we can assume the distribution of participants' actual past engagement in these activities to be similar across conditions. Thus, higher or lower admission rates across conditions signal an impact of our experimental manipulations.

6.1.2. Results and Discussion

One hundred and eighty nine participated completed Experiment 3. We found that the addition of a risky disclosure setting decreased the probability of participants' admitting to sensitive behavior. Specifically, we found that each additional risky disclosure setting resulted in a 3.3% decrease in probability of admitting to a sensitive behavior ($\beta_{\text{RiskySetting}} = -.033, p = .02$; Table 5, Column 1). These results are consistent when including demographic controls for participants (Table 5, Column 2). These results provide support for H3 and show that, when disclosure settings are randomly assigned to be more or less protective, downstream self-disclosure does vary to mitigate this risk. This is contrast to when the same protections are shifted via manipulations of choice frames and downstream self-disclosure does not shift. We are, however, cautious to draw definitive conclusions from comparisons between what we find in this

⁸ https://www.reddit.com/r/mturk/comments/2hqfvp/your_turk_id_is_not_anonymous/;
<http://turkernation.com/archive/index.php/t-17525.html>

experiment and the results of Experiment 2; for example, differences in the samples between experiments (which were conducted at different times) may explain why we observe mitigating behavior here but not Experiment 2. We address this concern in Experiment 4.

[Table 5: Experiment 3 Results]

VARIABLES	Downstream Self-Disclosure	Downstream Self-Disclosure
	(1) <i>Admit</i>	(2) <i>Admit</i>
RiskySetting	-0.0332* (0.0139)	-0.0297* (0.0136)
Age		-0.00343* (0.00159)
Male		-0.00574 (0.0343)
White		-0.0460 (0.0560)
College		0.0239 (0.0334)
Constant	0.435** (0.0268)	0.592** (0.0832)
Observations	1,512	1,496
Number of id	189	187

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

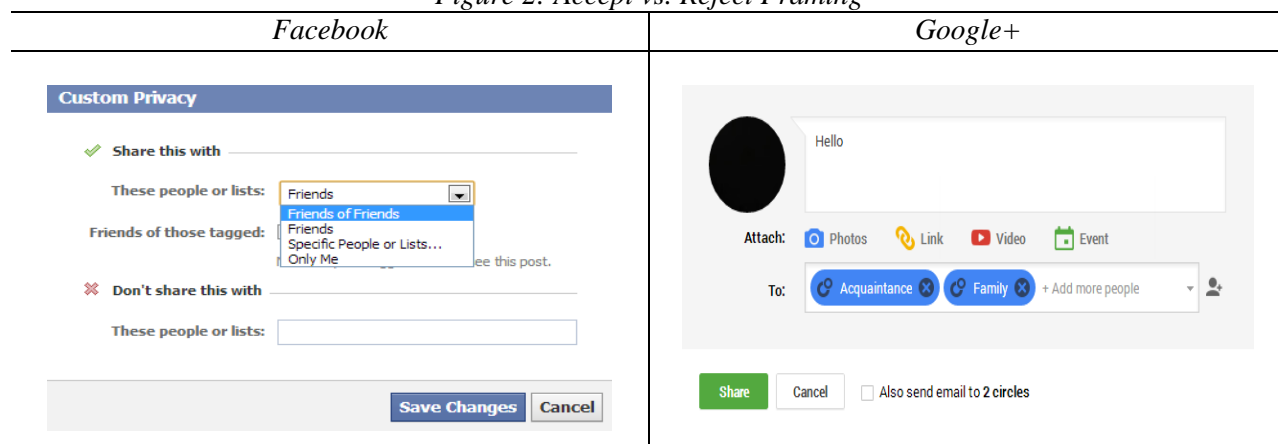
7. Experiment 4

Experiment 4 tests the three hypotheses in the manuscript simultaneously. Specifically, it evaluates differences in downstream self-disclosure when upstream disclosure settings choices are deliberately influenced using framing manipulations (H1, H2) and when upstream disclosure settings are simply assigned to be more or less protective (H3).

Experiment 4 shows robustness of our results in two ways. First, examining all hypotheses simultaneously helps rule out concerns related to differences in the samples between prior experiments. Second, we modify our experimental manipulations to show robustness of our prior results. In particular, Experiments 1 and 2 used a similar manipulation of choice framing, limiting our ability to draw broad conclusions about the effect of choice architecture and framing in cascaded privacy choice settings. It is

possible, for example, that the effects identified thus far are idiosyncratic to our chosen manipulation of choice framing. Thus, we consider in Experiment 4 framing manipulations involving whether privacy choices are presented as a choice to *allow* a use of personal information (an accept frame) versus a choice to *prohibit* the same use of personal information (reject frame). Similar to our prior manipulation, variation in accept/reject framing is common across privacy-relevant contexts. For example, Figure 2 illustrates how privacy-relevant choices can be presented to consumers (sometimes simultaneously) as either a choice to allow or restrict access to personal information. Prior research substantiates the potential of an accept versus reject presentation of a choice to significantly influence decision making across contexts, attributing these effects to the potential of these manipulations to differentially highlight competing considerations or motives in choice contexts (Shafir, 1993).

Figure 2: Accept vs. Reject Framing



Johnson, Häubl, and Keinan (2007) propose *Query Theory* as one explanation of why an accept versus reject frame elicits 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. They suggest that if an accept versus reject frame influences the valence and ordering of these queries, it can generate variation in individual judgments of objectively identical options. Consistent with this theory, although substantially predating it, Shafir (1993) posited that positive dimensions of choice weigh heavier under an accept frame while negative dimensions of that same choice weigh heavier under a reject frame. In line with this theory, a

reject choice frame in privacy contexts may highlight the negative dimensions of data allowances (e.g., potential repercussions and risks from data allowance) relative to an accept choice frame of the same objective choices. Thus, we use an accept/reject presentation of upstream privacy choices as our framing manipulation in Experiment 4.

7.1.1. Design and Procedure

Experiment 4 consisted in a 2 (“Accept Frame”, “Reject Frame”) + 2 (“Non Risky Setting”, “Risky Setting”) between-subject design. The first set of conditions mimic the design of Experiments 1 and 2 and allow us to evaluate the impact of manipulations of choice frames on, first, the choice of disclosure settings (similar to prior experiments), and whether differences in disclosure emerge between framing manipulations. In these two conditions, we manipulated, between subjects, the framing of the privacy settings by altering whether participants were presented with the settings as a choice to allow a use of personal information (accept frame) or prohibit the same use of personal information (reject frame). Since, in Experiment 2, framing effects only materialized for high importance disclosure settings, we utilized only the set of high importance settings for this experiment (see A6 in Appendix). Afterwards, participants were presented with the same eight sensitive disclosure questions as in Experiment 2.

In our second set of conditions, participants were not provided with any upstream privacy choices (i.e., no data settings presented) and were instead provided with a notice informing them how their self-disclosures would be used. Whereas the notice in Experiment 3 introduced up to three uses, we introduced only a single use so as to be more consistent with the size of framing effects in prior experiments. As a result, this manipulation is binary. In one condition (“Non Risky Setting”), participants were provided with a standard research assurance that their responses would be kept confidential and only used for research purposes. In the other condition (“Risky Setting”), participants were instead informed that a risky disclosure setting would be applied to their self-disclosure decisions (e.g., that their responses will be shared with other participants of the study)—see A7 in Appendix for details on these conditions. The disclosure setting was chosen at random from the four provided in the first two conditions. Participants were again presented with the same eight sensitive disclosure questions as in Experiment 2.

7.1.2. Results

Three hundred individuals ($M_{\text{Age}} = 34$, $SD_{\text{Age}} = 10.6$, $M_{\text{Male}} = .46$) participated in Experiment 4. We found support for H1; participants in the “reject” condition were 45% more likely than those in the “accept” condition (58% vs. 40%; $t(144) = 2.655$, $p = .008$) to choose the privacy-protective option. A random effects panel regression (Table 6a, Column 1) confirmed this finding with a positive and significant estimate of the effect of the reject frame ($\beta_{\text{RejectFrame}} = .17$, $p = .012$). Moreover, we again found that downstream self-disclosure did not vary between manipulations of choice frames: participants presented with the accept framing had comparable admission rates (percent of unethical behaviors admitted to) compared to those presented with the reject framing (51% vs. 49%, $t(144) = .65$, $p = .52$). This is again confirmed in our random effects panel regression with a near zero and insignificant estimate on the effect of *RejectFrame* on admit rates ($\beta_{\text{RejectFrame}} = -.013$, $p = .743$; Table 6a, Column 2)—H2 supported.

[Table 6a: Accept/Reject Framing Effects]

VARIABLES	Upstream Disclosure Settings	Downstream Self-Disclosure
	(1) <i>ProtectiveSetting</i>	(2) <i>Admit</i>
Reject Frame	0.174* (0.0687)	-0.0127 (0.0386)
Age	-0.00289 (0.00341)	-0.00380* (0.00182)
Male	0.00673 (0.0702)	0.0168 (0.0396)
White	-0.0239 (0.0830)	0.0895+ (0.0458)
College	0.0273 (0.0674)	0.0112 (0.0381)
Constant	0.645** (0.102)	0.468** (0.0525)
Observations	560	1,120

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

In addition to this analysis, we evaluated participants’ admission rates when they were assigned a risky disclosure setting selected at random from the four used in the accept/reject conditions (*RiskySetting* is a binary indicator of participants in this condition) versus when they were provided with a standard

research confidentiality assurance (see A7 in Appendix). We found a significant and negative effect of informing participants that a risky disclosure setting is applied to their responses ($\beta_{\text{RiskySetting}} = -.057, p = .097$; Table 6b, Column 1). Excluding those who took less than 3 seconds to read the information provided (13% of participants) we found even stronger effects ($\beta_{\text{RiskySetting}} = -.077, p = .045$; Table 6b, Column 2)—H3 supported.

[Table 6b: Assigned Risky Disclosure Settings and Self-Disclosure]

	Downstream Self-Disclosure	
	(1) <i>Admit – Full Sample</i>	(2) <i>Admit – No Low Attention Participants</i>
RiskySetting	-0.0577+ (0.0348)	-0.0777* (0.0387)
Age	-0.00248 (0.00166)	-0.00238 (0.00174)
Male	0.00873 (0.0367)	0.0189 (0.0410)
White	0.0117 (0.0480)	0.0201 (0.0546)
College	-0.0546 (0.0353)	-0.0518 (0.0383)
Constant	0.560** (0.0587)	0.546** (0.0655)
Observations	1,200	1,040

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

7.2. Discussion

With these results, we reinforce support for the main hypotheses of the manuscript. First, we again identify a significant impact of choice framing on privacy decision making: presenting privacy settings as a choice to reject versus accept uses of personal information elicited significant differences in the choice of protective disclosure settings (H1 supported). In addition, again, we do not find differences in downstream self-disclosure behavior following shifts in privacy protection driven by accept/reject framing manipulations (H2 supported). In contrast, we find that self-disclosure is adjusted to mitigate risk when changes in the same disclosure settings are directly manipulated via privacy notices (H3). Together the results demonstrate that subtle manipulations of the architecture of privacy choices can meaningfully

impact consumers' choice of privacy-protective options. These results also demonstrate that while the impacts of framing upstream are not mitigated by downstream privacy choice, random assignment to risky disclosure settings is mitigated by different levels of downstream privacy choices.

8. Future Research and Policy Implications

This manuscript presents evidence that subtle heterogeneity in the presentation of privacy-relevant choices, that occurs in real decision contexts, can trigger or quell consumer privacy concerns, and significantly impact consumers' initial privacy choices. Moreover, we found that significant changes in chosen protection levels upstream did not result in changes in behavior in downstream self-disclosure decisions. In contrast, assigning the same disclosure settings to be protective or risky does result in significant differences in downstream self-disclosure. Our work has a number of important implications for privacy research.

First, the research has implications for a body of work evaluating the impact of granular control in privacy settings. This research has identified powerful impacts of providing granular control on consumer privacy concerns and decision making, and suggests that these effects may be paradoxical: Xu et al. (2012) suggest that the effect of control would likely persist even if control was "illusory" and Brandimarte, Acquisti, and Loewenstein (2013) find that the effects of control persist even when objective risk is elevated. We highlight that choice architecture that subtly influences participants' choice of actual privacy options via control mechanisms is also highly relevant to privacy decision making and consumer welfare.

Second, our work highlights the importance of studying privacy decision making as a process involving interrelated decisions over time; doing so may allow for more holistic privacy research and theoretical models of privacy behavior. While we focus on behavioral models of privacy decision making, future research efforts could also profit from studying cascaded privacy choices using more traditional approaches similar to those used to study moral hazard in insurance markets. The most productive approaches to the topic may involve a melding of behavioral and more traditional economic perspectives.

In addition, we consider only two levels of privacy choice that flow sequentially. Future work may relax this and consider cascaded choices that are circular or iterative in nature. For example, participants may choose privacy protections, then choose what to disclose, and then go back and readjust their privacy protections. Studying how choices at one level of choice influence subsequent downstream choices does introduce some endogeneity concerns. There are unobserved factors (e.g., latent privacy concerns) that are difficult to account for empirically but may simultaneously influence many levels of privacy decision making. To reduce these concerns, we chose to evaluate differences in behavior across exogenously assigned groups. An alternative approach may be to rigorously measure endogenous factors (e.g., privacy concerns) and control for them, but this introduces the challenge of accounting for other unanticipated sources of endogeneity (e.g., personality characteristics).

A final research implication of our work relates to why we don't observe an adjustment in downstream privacy decision making. We identify this effect in a fairly specific instance (i.e., following framing manipulations in an upstream choice) and it would be of interest to see the extent to which this phenomenon holds more generally. Relatedly, our manuscript is not conclusive on the mechanism through which this occurs. We suggest that it could be driven by the subtle nature of framing effects. Future research may seek to disentangle the role of this mechanism, or other mechanisms we haven't considered, in driving the behavior we observe in the downstream privacy choices. Along the same logic, there may be conditions under which random assignment may drive stronger or weaker compensatory reactions (e.g., when the shifts in protection are made more or less salient).

There are, naturally, some limitations to these investigations. First, the use of the specific variation that motivates the framing manipulations in our experiments may or may not persist over time. However, we suggest that the relevance of our findings extends beyond any single manipulation of choice frames: as long as consumers manage their privacy via cascaded and heterogeneous privacy choices, designed largely at the discretion of online service providers, similar heterogeneity with the propensity to impact consumer decision frames will likely persist.

There is no direct evidence that the methods that Internet providers use to elicit choices of disclosure settings, or the changes over time observed in these settings, are intended to elicit maximal data allowances from consumers. Some of the variation we identified is, no doubt, accidental and some may reflect practical limitations in presenting consumer privacy choices (e.g., there may be too many privacy-relevant settings to include all under the same label). Given the considerable value that firms hope to derive from the collection and use of consumer personal information, however, it would be 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.

Another limitation arises from the experimental nature of the work, with constraints in terms of external validity, due to context and sample selection. We sought to address these concerns by collecting actual versus hypothetical choice of settings and self-disclosures, modeling experimental manipulations on actual variation in privacy settings, and varying the experimental context of study. If anything, however, we believe that these differences make our experiments more conservative. In real-world settings, where downstream choices are typically made long after initial choices, compensatory downstream decisions seem even less likely than in our experiments, in which downstream choices immediately followed upstream ones.

These limitations aside, these results raise significant concerns about proposed policy approaches to alleviate consumer privacy concerns. Currently, these approaches 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). A senior advisor for a large technology firm (and contributor to the report) stated that “There’s no bad data, only bad uses of data.”⁹ Our results suggest that providing consumers with greater control over

⁹ Lohr, S. (2013). Big data is opening doors, but maybe too many. *The New York Times*.
<http://www.nytimes.com/2013/03/24/technology/big-data-and-a-renewed-debate-over-privacy.html>

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 manipulations of choice framing to elicit higher rates of information sharing from consumers (Acquisti, Adjerid, and Brandimarte, 2013). These concerns are exacerbated by the lack of adjustment in downstream privacy choice. Such inadvertent susceptibility to framing effects and the lack of adjustment downstream is of increasing consequence, given the growing usage of personal information in commercial contexts, some of which may be particularly intrusive or even discriminatory. For instance, Sweeney (2013) found that black-identifying names were 25% more likely to get an online ad suggestive of an arrest record relative to white-identifying names.

Alternate policy approaches might include privacy control mechanisms as only one component 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, as well as introducing 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 aid consumers 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.

9. Conclusion

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. Understanding these trade-offs and how consumers approach privacy choices is critical, and our understanding continues to evolve. As a result, it remains a challenge to design policies that lead choice architects to create privacy contexts that balance consumer privacy considerations against competing utility gains from data allowances. This challenge is even more daunting when one considers the complexities introduced by heterogeneous and cascaded choices susceptible to subtle manipulations of decision framing. What is clear, however, is that relying on control, provided at the discretion of service providers, as the predominant mechanism for privacy protection may not suffice. 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, considered in light of cascaded privacy choice, and provided alongside supplemental protections, they may have largely the effect of quelling consumer privacy concerns by providing the *opportunity* to restrict the collection and use of their personal information while, in practice, actually implementing policies that result in consumers continuing to provide broad, and potentially harmful, data allowances to firms.

10. References

1. Acquisti, A., Adjerid, I., & Brandimarte, L. (2013). Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security & Privacy*, (4), 72-74.
2. Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509-514.
3. 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.
4. Adjerid, I., Acquisti, A., Telang, R., Padman, R., & Adler-Milstein, J. (2015). The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Science*, 62(4), 1042-1063.
5. Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, 80(1), 123-129.
6. Almuhiemedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., ... & Agarwal, Y. (2015). Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 787-796). ACM.

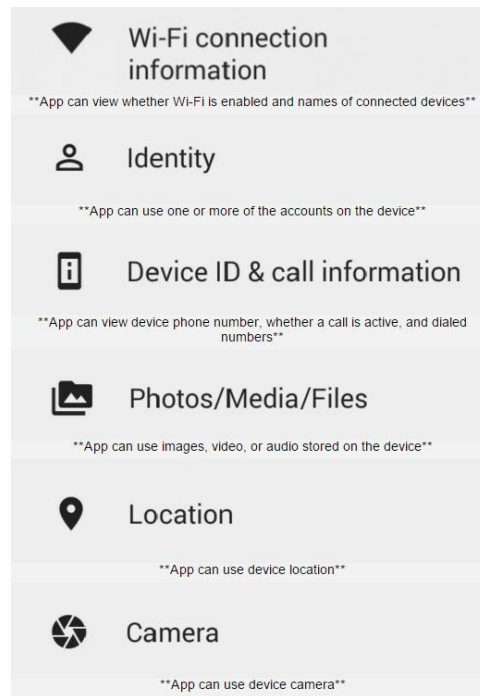
7. Angrist, J. D., & Pischke, J. S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
8. Balebako, R., Leon, P., Shay, R., Ur, B., Wang, Y., & Cranor, L. (2012). Measuring the effectiveness of privacy tools for limiting behavioral advertising. Web.
9. Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced confidences privacy and the control paradox. *Social Psychological and Personality Science*, 4(3), 340-347.
10. Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1), 3-5.
11. 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.
12. 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.
13. Crandall, R. W., & Graham, J. D. (1984). Automobile safety regulation and offsetting behavior: Some new empirical estimates. *American Economic Review*, 74, 328-31.
14. Culnan, M. J., & Armstrong, P. K. (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science*, 10(1), 104-115.
15. Dave, D., & Kaestner, R. (2009). Health insurance and ex ante moral hazard: Evidence from Medicare. *International Journal of Health Care Finance and Economics*, 9(4), 367-390.
16. Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1), 61-80.
17. Egelman, S., Felt, A. P., & Wagner, D. (2013). Choice architecture and smartphone privacy: There's a price for that. In *The economics of information security and privacy* (pp. 211-236). Springer Berlin Heidelberg.
18. Epley, N., Caruso, E., & Bazerman, M. H. (2006). When perspective taking increases taking: Reactive egoism in social interaction. *Journal of Personality and Social Psychology*, 91(5), 872.
19. 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>.
20. Fischer-Hübner, S. (2001). *IT-security and privacy: Design and use of privacy-enhancing security mechanisms*. Springer-Verlag.
21. Ganzach, Y., & Karsahi, N. (1995). Message framing and buying behavior: A field experiment. *Journal of Business Research*, 32(1), 11-17.
22. Goes, P. B. (2013). Editor's comments: Information systems research and behavioral economics. *MIS Quarterly*, 37(3), 3-8.
23. Goldfarb, A., & Tucker, C. (2011). Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3), 389-404.
24. Goodman, J. K., Cryder, C. E., & Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making*, 26(3), 213-224.
25. Hui, K.-L., Teo, H. H., and Lee, S. Y. T. (2007). The value of privacy assurance: An exploratory field experiment. *MIS Quarterly*, 31(1), 19-33.
26. 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.
27. Johnson, E. J., Bellman, S., & Lohse, G. L. (2002). Defaults, framing and privacy: Why opting in-opting out I. *Marketing Letters*, 13(1), 5-15.
28. Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives?. *Science*, 302(5649), 1338-1339.
29. Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: A query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(3), 461.
30. 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.

31. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 47(2), 263-291.
32. 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.
33. Kelly, I. R., & Markowitz, S. (2009). Incentives in obesity and health insurance. *Inquiry*, 418-432.
34. Kim, D. J., Ferrin, D. L., & Rao, H. R. (2008). A trust-based consumer decision-making model in electronic commerce: The role of trust, perceived risk, and their antecedents. *Decision Support Systems*, 44(2), 544-564.
35. Klick, J., & Stratmann, T. (2007). Diabetes treatments and moral hazard. *Journal of Law and Economics*, 50(3), 519-538.
36. Klopfer, P. H., & Rubenstein, D. I. (1977). The concept privacy and its biological basis. *Journal of Social Issues*, 33(3), 52-65.
37. Kühberger, A. (1998). The influence of framing on risky decisions: A meta-analysis. *Organizational Behavior and Human Decision Processes*, 75(1), 23-55.
38. 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.
39. 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.
40. Levy, D. T. & Miller, T. (1999). Review: Risk compensation literature, the theory and evidence. *Journal of Crash Prevention and Injury Control*.
41. 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.
42. Loewenstein, G., Bryce, C., Hagmann, D., & Rajpal, S. (2015). Warning: You are about to be nudged. *Behavioral Science & Policy*, 1(1), 35-42
43. Milne, G. R., and Gordon, E. M. (1993). Direct mail privacy-efficiency trade-offs within an implied social contract framework. *Journal of Public Policy and Marketing*, 12(2), 206-215.
44. Miyazaki, A. D., & Krishnamurthy, S. (2002). Internet seals of approval: Effects on online privacy policies and consumer perceptions. *Journal of Consumer Affairs*, 36(1), 28-49.
45. Moon, Y. (2000). Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of Consumer Research*, 26(4), 323-339.
46. 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.
47. Peltzman, S. (1975). The effects of automobile safety regulation. *The Journal of Political Economy*, 677-725.
48. Pew Center on the States (2012). Overdraft America: Confusion and concerns about bank practices. http://www.pewtrusts.org/uploadedFiles/wwwpewtrustsorg/Fact_Sheets/Safe_Checking/Overdraft_America_Final.pdf.
49. 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.
50. Schneider, S. L. (1992). Framing and conflict: Aspiration level contingency, the status quo, and current theories of risky choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(5), 1040.
51. Schoorman, F. D., Mayer, R. C., Douglas, C. A., & Hetrick, C. T. (1994). Escalation of commitment and the framing effect: An empirical investigation. *Journal of Applied Social Psychology*, 24(6), 509-528.
52. Schwartz, P. M. (2005). Privacy inalienability and the regulation of spyware. *Berkeley Technology Law Journal*, 20, 1269.

53. Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory & Cognition*, 21(4), 546-556.
54. Solove, D. J. (2013). Privacy self-management and the consent dilemma. *Harvard Law Review*, 126, 1879-2139.
55. Steelman, Z. R., Hammer, B. I., & Limayem, M. (2014). Data collection in the digital age: Innovative alternatives to student samples. *MIS Quarterly*, 38(2), 355-378.
56. Stutzman, F., Gross, R., & Acquisti, A. (2013). Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of Privacy and Confidentiality*, 4(2), 2.
57. Sweeney, L. (2013). Discrimination in online ad delivery. *Queue*, 11(3), 10.
58. Takemura, K. (1994). Influence of elaboration on the framing of decision. *The Journal of Psychology*, 128(1), 33-39.
59. Thaler, R. H., Sunstein, C. R., & Balz, J. P. (2014). Choice architecture. *The Behavioral Foundations of Public Policy*.
60. Turow, J., & Hennessy, M. (2007). Internet privacy and institutional trust insights from a national survey. *New Media & Society*, 9(2), 300-318.
61. Tversky, A., & Kahneman, D. (1985). The framing of decisions and the psychology of choice. In *Environmental Impact Assessment, Technology Assessment, and Risk Analysis* (pp. 107-129). Springer Berlin Heidelberg.
62. Wilde, G. J. (1981). Objective and subjective risk in drivers' response to road conditions: The implications of the theory of risk homeostasis for accident aetiology and prevention. Queen's University.
63. 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
64. Xu, H., Teo, H. H., Tan, B. C., & Agarwal, R. (2009). The role of push-pull technology in privacy calculus: The case of location-based services. *Journal of Management Information Systems*, 26(3), 135-174.
65. Xu, H., Teo, H. H., Tan, B. C., & Agarwal, R. (2012). Research note. Effects of individual self-protection, industry self-regulation, and government regulation on privacy concerns: A study of location-based services. *Information Systems Research*, 23(4), 1342-1363.
66. Zeger, S. L., & Liang, K. Y. (1986). Longitudinal data analysis for discrete and continuous outcomes. *Biometrics*, 42(1), 121-130.

Appendix: Experimental Materials

[A1: Experiment 1: Mobile Permissions]



[A2: Importance Ranking of Uses of Personal Information]

Choice	Description	Mean Importance
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
3	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
5	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 my responses.	2.83
10	Allow my responses to be used in academic publications.	2.91



[A3: High vs. Low Importance 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
3	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 my responses.	Low
8	Allow my responses to be used in academic publications.	Low

[A4: Sensitive 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?
3	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?
5	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?

[A5: Experiment 3 Example Graphical Notice]

	More Protective
	Less Protective

Shared with religious organizations interested in evaluating personal ethics	No
Shown to other participants of this study	Yes
Published on a research bulletin openly available on the internet	Yes

[A6: Accept vs. Reject Condition]

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

[A7: Experiment 4 Materials]

Condition	Text
Non-Risky Condition	Your responses will be kept confidential and only used for purposes of this study.
Risky Condition*	Your responses in this study will be stored unencrypted.
	Your responses in this study will be shown to other participants of the study.
	Your responses in this study will be shared with religious organizations interested in evaluating personal ethics.
	Your responses will be published on a research bulletin openly available on the internet.

*One risky disclosure setting was selected at random to present to participants in this condition