

Carnegie Mellon University
H. JOHN HEINZ III COLLEGE
School of Public Policy and Management

DISSERTATION
By
Idris Adjerid

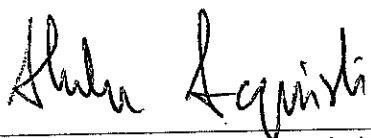
Submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
In
Information Systems and Management

Uninformed Consent: The Benefits and Limits of Transparency and
Choice in Privacy Decision Making

May 31, 2013

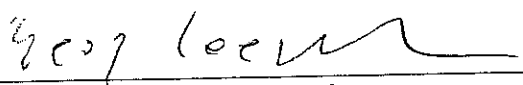
Accepted by the
Dissertation Committee:



Professor Alessandro Acquisti, Chair

7/22/13

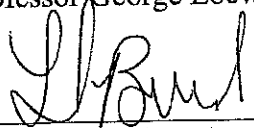
Date



Professor George Loewenstein

7/22/13

Date




Professor Linda Babcock

7/22/13

Date

Approved by the
Dean



Ramayya Krishnan

7/22/13

Date

**Uninformed Consent: The Benefits and Limits of Transparency and Choice in Privacy
Decision Making**

**A dissertation submitted to the Heinz College – School of Information Systems
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Information Systems and Management by
Idris Adjerd**

Dissertation Committee:

**Alessandro Acquisti (Chair), George Loewenstein, and Linda Babcock
Carnegie Mellon University Pittsburgh, Pennsylvania**

April, 2013

Abstract:

Solutions to privacy concerns centered on notifying consumers about (transparency), and granting them control over the collection and use of their personal information (choice) are pervasive. Policy makers posit that these measures will aid consumers in improved privacy decision making. Conversely, scholars argue that these protections may have a negative impact on market efficiency and firm technology innovation and adoption. Chapter 2 evaluates the impact of regulation providing consumers transparency and choice on technology adoption by hospitals and finds, in contrast to prior results, evidence for a beneficial role of privacy regulation. I also find evidence that these gains may be a result of reduced barriers to adoption stemming from consumer privacy concerns. In Chapters 3 and 4 I shift my focus to evaluate the premise proposed by policy makers that increased transparency and choice will improve consumer privacy decision making. In Chapter 3, I first find that simple privacy notices communicating lower privacy protection can, under some conditions, result in less disclosure from participants, in-line with the policy aims for increased transparency. However, I also find that simple and common changes in those same notices, exploiting individual heuristics and biases, can result in the effect of even straightforward and accessible privacy notices being predictably manipulated (Experiment 1) or entirely thwarted (Experiment 2). Finally, in chapter 4 I find substantial malleability in individual privacy decision making in response to changes in choice framing. Specifically, the labeling of settings, the mix of setting relevance, and the presentation of choices as a choice to reject all impacted the decision frame for participants in a manner that significantly influenced participants' choice of privacy protective settings. Taken together, these results suggest that while privacy solutions centered on transparency and choice may alleviate barriers to technology adoption stemming from consumer privacy concerns, the implicit assumption that they will reduce consumer privacy risks may be questioned. Implications for policy makers include a persistence, and perhaps increase, in consumer privacy risks despite increased transparency and control.

Dissertation Outline

I.	Chapter 1: Transparency and Choice in the Context of Privacy Choice.....	4
II.	Chapter 2: The Impact of Privacy Regulation on Technology Adoption: The Case of Health Information Exchanges.....	10
III.	Chapter 3: A Sleight of Privacy: The Limits of Transparency.....	47
IV.	Chapter 4: Why Choice May Not Suffice: Framing and Malleable Preferences for Privacy.....	71
V.	Chapter 5: Discussion and Conclusions.....	101
VI.	References.....	105
VII.	Appendices.....	111

I. Transparency and Choice in the Context of Privacy Choice

The advent of the internet alongside widespread adoption of technologies that facilitate nearly continuous connectivity has given rise to a new digital age so sweeping that it is difficult to imagine a facet of life not altered or transformed by it. These include the manner in which we connect and interact with one another (Garton, Haythornthwaite, and Wellman, 1997; Kumar, Novak, and Tomkins, 2010), consume news and entertainment (Deuze, 2001; Bhattacharjee, Gopal, and Sanders, 2003), educate and learn (Allen and Seaman, 2005), and even pursue romantic interests (Ellison, Heino, & Gibbs, 2006). Many of these changes offer likely benefits to society including the broader flow of information and ideas (Lessig, 2002), increased transparency and awareness of world events (Huang, 2011), and more consumer choice and increasingly efficient markets (Peterson, Balasubramanian, and Bronnenberg, 1997; Bakos, 1998). However, these advances have also given rise to significant privacy concerns stemming from an unprecedented level of collection, aggregation, and analysis of personal information about individuals. For instance, companies termed “data brokers” consolidate information from thousands of online (and some offline) sources to create consumer profiles with information ranging from names, addresses, age, and race to pregnancy, new births, and political contributions (Becket, 2013). The largest of these data broker has amassed a database with information on over five hundred million consumers, including almost all American consumers, and processed 50 trillion data transactions in 2011 with a reported revenue of \$1.3 billion (Singer, 2012). In addition, advances in computer science and statistics have resulted in increasingly sophisticated uses of personal information for marketing purposes. For example, behavior or targeted advertising analyzes consumer personal information to make inferences (sometimes sensitive ones) about consumer preferences in order to more effectively tailor advertising.

Policy makers recognized early on the potential of increased digitization of information and improved computing to fundamentally alter the manner in which sensitive personal information is collected, analyzed, and disseminated. They also recognized that this would result in unique and substantial privacy concerns (Ware, 1973). To address these concerns, policy makers have often relied on transparency and

control solutions; or solutions centered on notifying consumers about, and granting them choice in how their personal information is collected, used, and disseminated. This is consistent with early conceptualizations of privacy which highlight transparency and control as central tenants. For example, Alan Westin (1967) suggests that “privacy is the claim of individuals, groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others”. While choice is explicit in Westin’s definition of privacy, implicit in his definition is the existence of sufficient transparency such that an individual is aware of how their information may be collected, used, and shared. From a policy and regulatory perspective, early implementations of transparency and choice in the context of privacy came in the form of a *Code of Fair Information Practices* developed as part of a 1973 report titled “Records, Computers and the Rights of Citizens” (Ware, 1973; Gellman 2012). These were later refined and popularized by the OECD privacy guidelines (OECD, 1980).

Transparency and control privacy solutions are also reflected in numerous legislative efforts intended to address consumer privacy concerns. For example, The Health Insurance Portability and Accountability Act (HIPAA) privacy rule includes provisions requiring covered healthcare entities to provide patients notice of the collection and use personal health information and also provide choice with respect to some uses of personal information. The Graham-Leach-Bliley Act of 1999 requires financial agencies to provide consumers notice of data practices and give them the option to Opt-out of certain data collection and uses. Moreover, policy makers continue to lean on transparency and control to address emerging privacy issues. For example, in the context of online consumer privacy, firms and policy makers have reached a broad consensus on self-regulatory approaches¹ centered on providing consumers with “transparency and control” solutions. Both the FTC white paper on consumer privacy and the White House Consumer Bill of Rights (FTC, 2012; The White House, 2012) presented transparency and notice as central to consumer privacy protection. Industry leaders broadly concurred with the approaches

¹ The impact of regulatory vs. self-regulatory approaches is not a focus of this thesis. A detailed discussion of the implications of self regulatory approaches in the context of privacy are discussed in Solove (2013).

outlined by policy makers. In comments on the FTC privacy framework, Facebook stated that “...companies should provide a combination of greater transparency and meaningful choice...” for consumers, and Google stated that making the “collection of personal information transparent” and giving “users meaningful choices to protect their privacy” are two of their guiding privacy principles (Santalesa, 2011). Privacy advocates have also embraced these approaches (Reitman, 2012).

However, it is possible that privacy regulation intended to protect consumers comes at a cost. Scholars argue that while regulation restricting the flow of information could result in privacy gains for individuals, they may also have costly consequences by reducing the efficiency of markets (Stigler, 1980; Posner, 1981). Thus the challenge remains for policy makers remains to strike a balance between the need for individual privacy and encouraging technology innovations with the propensity to provide significant public and private good (e.g. national security and business and innovation). Along this vein, scholars have started to empirically evaluate the costs and benefits of privacy protection, including those granting consumers transparency and control (Miller and Tucker, 2009). This thesis contributes to this literature and focuses on the impact of transparency and control from two distinct perspectives. First, I consider a firm centric perspective in which I evaluate the impact of transparency and control on firm innovation and adoption of technology (Chapter 2). Secondly, I consider a consumer centric perspective in which I evaluate the potential of transparency (Chapters 3) and control (Chapter 4) solutions to influence consumer privacy decision making, and alleviate privacy concerns and reduce risks.

With respect to the first focus of this dissertation, scholars have differed on the potential impact of transparency and control on firm innovation and technology adoption. Some argue that transparency and control solutions may hamper firms with unnecessary costs that would stifle innovation and technology adoption (Lenard and Rubin 2005, 2006), while others argue that privacy protection will decrease barriers to innovation stemming from consumer privacy concerns and thus, spur innovation and technology adoption (McGraw et al, 2009; Bamberger and Mulligan, 2011). A small but growing body of empirical

work has primarily found evidence in line with a negative impact of transparency and control on technology innovation and adoption by firms. Chapter 2 is a contribution to this stream of work in which I evaluate the impact of regulation requiring transparency and consent in the context of Health Information Exchanges (HIEs). HIEs are innovative healthcare technology initiatives that increase coordination between healthcare providers and are central to the national health IT strategy to improve efficiency and quality of care through enhanced sharing of patient data. However, they primarily involve the facilitated sharing of sometimes sensitive health information and have thus, raised significant privacy concerns. To soothe these concerns, numerous states have enacted laws establishing strict requirements that require patients to be provided notice and choice with regard to the medical data shared through HIEs. I investigate the impact of this regulation on the adoption and success of Health Information Exchanges. I find that among all states with laws intended to promote HIE adoption, those that had requirements for patient consent experienced greater HIE adoption and success. These results suggest that gains we identify may have been due to the propensity of transparency and choice to alleviate consumer privacy concerns associated with attempts to promote technology efforts. These findings contribute to the debate over the impact of privacy regulation on technological progress and provide insights on the delicate balance between privacy concerns and the benefits of technology adoption.

The second focus of this dissertation is on the propensity of transparency and choice to address consumer privacy concerns and reduce risk. In principle, giving consumers more control over, and more information about, how their personal data is used seems an unarguable improvement over a situation in which consumers are left in the dark. In particular, policy makers posit that consumers notified of data practices by firms may in response, alter disclosure behavior in line with their preferences for privacy. Similarly, consumers offered increased choice may opt-out of forms of data analysis and collection deemed intrusive, again in line with their individual preferences for privacy. Unfortunately, the ability of even improved transparency solutions or additional control to better align consumer attitudes towards privacy with actual behavior and reduce regret from over sharing is ultimately questionable. A growing body of

work suggests that many consumers are uncertain about their preferences for privacy, and that observed privacy behavior is often contrary to stated concerns (e.g., Acquisti, 2009) and is affected by contextual factors (John, Acquisti, & Lowenstein, 2012). Scholars have also started to question the effectiveness of transparency and notice alone in protecting consumer privacy. For example, Solove (2013) states that “Despite my early optimism about opt-in, I now believe it will fail. One reason is that organizations, as a rule, will have the sophistication and motivation to find ways to generate high opt in rates.” Other scholars suggest that, “many data-processing institutions are likely to be good at obtaining consent on their terms...” (Schwartz, 2005). Commenting on transparency and control solutions provided in Gramm-Leach-Bliley (GLB) Act, Janger and Schwartz (2001) find very few customers changed their behavior (opted out) after being provided increased transparency.

In chapter 3, I evaluate the impact of transparency on individual privacy decision making. Specifically, I evaluate the propensity of simple privacy notices to consistently influence privacy decisions making. I find that while simple privacy notices communicating lower privacy protection can, under some conditions, result in less disclosure from participants (in line with the policy aims for increased transparency), simple and common changes in those same notices, that exploit individual heuristics and biases, can result in the effect of even straightforward and accessible privacy notices being predictably manipulated (Experiment 1) or entirely thwarted (Experiment 2). In Experiment 1, I demonstrate that the impact of privacy notices on disclosure is sensitive to whether notices are presented as increasing or decreasing in protection, even when the objective risks of disclosure stay constant. In Experiment 2, I demonstrate that the propensity of privacy notices to impact disclosure can be muted by a number of simple and minimal misdirections (such as a mere 15 second delay between notices and disclosure decisions) that do not alter the objective risk of disclosure. It follows that privacy notices can – on the one hand – be easily marginalized to no longer impact disclosure, or – on the other hand – be used to influence consumers to share varying amounts of personal information. Transparency may, therefore, become a “sleight” of privacy.

Finally, In chapter 4, I evaluate the propensity of increased consumer control to consistently impact individual privacy decision making and the reduce privacy risk. Specifically, I evaluate the propensity of individuals to choose more protective privacy settings under different framing of otherwise identical choices. Utilizing common but subtle differences in the manner that privacy choices are presented to consumers online I have designed four studies that seek to evaluate the potential of these differences in presentation to impact the propensity of individuals to select protective settings and also to make personal disclosures. This work leans on a literature on choice framing which finds, consistently, that individual choice can be altered by different presentations of otherwise identical choices and also a priming literature which suggests that subtle manipulations can elicit different concepts and mindsets (e.g. privacy or sharing) which can have significant impacts on individual behavior. Generally, I find considerable malleability in the choice of privacy protective settings under varying decision frames. First, I find that that the labeling of privacy relevant choices as “Privacy Settings” results in the choice of more protective settings relative to labeling the identical choices as “Survey Settings”. Moreover, I find that presenting individuals with an important privacy choice mixed with other settings that were pre-ranked by participants to be less relevant, significantly decreased the likelihood of participants choosing the protective option for the important choice. Finally, I find that participants presented privacy relevant choices as a choice to allow a use of their personal information (accept frame) were significantly *less* likely to choose the privacy protective option relative to those presented the identical choice as a choice to prohibit a use of their personal information (reject frame). Generally, I find that manipulations of framing do not have a significant impact on disclosure, when I control for the choice of settings by participants. These results suggest that common but subtle variation in the framing of privacy choices can predictably influence users to chose less protective privacy options.

II. Chapter 2: The Impact of Privacy Regulation on Technology Adoption:

The Case of Health Information Exchanges

1. Introduction

Numerous consumer services thrive today thanks to the exchange and use of personal, and sometimes sensitive, information. The risks associated with the potential misuse of that information, however, have fueled a debate over the appropriate approach to protecting consumers privacy and the role of regulation in that protection (Solove 2004; Lenard and Rubin 2005, 2006; Goldfarb and Tucker 2011a, 2011b). A recent string of policy changes by high-profile Internet firms such as Facebook and Google has propelled this debate into the national spotlight (Steel and Vascellaro 2010; Horowitz 2011; Angwin and Valentino-Devries 2012). Is the loss of privacy the cost to pay for innovation? Or can, in fact, the protection of consumers' data be compatible with technological progress? This chapter explores the role of privacy regulation in promoting or impeding technology adoption and innovation in the context of sensitive health data. It evaluates the impact of laws imposing patient consent requirements on the number of attempted efforts to electronically exchange patient data between unaffiliated organizations (known as Health Information Exchanges, or HIEs), and, within those, the number of HIEs actively sharing health information.

HIEs are information sharing collaborations between otherwise disconnected health entities. HIEs have spurred significant debate over the appropriate balance of patient privacy and the potential gains to healthcare providers and their patients. HIEs generally rely on the availability of digital health information, such as what would be available via an electronic medical records (EMR) system, and are expected to improve efficiency and quality of care through enhanced availability of patient clinical data at the point of care. At the same time, privacy, security, and consent come up consistently as primary patient concerns associated with HIEs (Simon et al 2009), leading privacy advocates and patient rights groups to advocate for increased privacy regulation (Greenberg, Ridgely, and Hillestad 2009; Milliard 2010).

Policymakers interested in the healthcare industry thus face the same challenge that emerges in other industries: how to address patient privacy concerns without over-regulating the disclosure of health information and stifling the growth and emergence of exchange efforts.

Across different states, US regulators have solved this challenge in different ways. Some states have enacted HIE-specific laws with requirements for patient consent; others have enacted HIE-specific laws that do not have such requirements; and others have not enacted HIE-specific laws at all. This variation allows us to evaluate the impact of the laws – and, in particular, the strictness of privacy protection afforded by the laws - on the number of attempted Health Information Exchanges (HIEs), and within those, the number of HIEs actively sharing health information.

My empirical strategy uses semi-annual data from a six year period (2004 through 2010) to compare the number of *attempted* and *operational* HIEs across states with variation in the extent to which regulation provided patients the option to consent with respect to the use of their data in an exchange (“attempted” HIEs include both HIEs actively sharing health information, and those merely in the planning phases). I disentangle the impact of consent requirements from incentives for HIE adoption using between-state variation in consent requirements and HIE-specific legislation. Including state and time fixed effects and controls for relevant observables (other elements of the laws, differences in state wealth, populations, Health IT adoption, and so forth), I find that among all states with laws intended to promote HIE adoption, those that included explicit requirements for patient consent experienced greater HIE activity (an increase of approximately 2 attempted HIEs and 0.644 more operational HIEs). This result is bolstered by the finding that states with incentives for exchange and explicit requirements for patient consent are almost twice as likely to report that privacy concerns presented a minimal challenge in their development compared to states with no HIE legislation, or states with incentives for exchange but no requirements for consent. These results are robust to evaluations of possible endogeneity. I do not find evidence that HIE laws are passed as a result of increased HIE activity (i.e., reverse causation). Also, I

find consistent results when I consider the impact of state unobservables that may be correlated with the passage of HIE-promoting legislation (such as changes in political attitudes or public opinion towards the importance of health IT).

This work is, to my knowledge, the first empirical study on the role of privacy protection on HIE progress. More broadly, it contributes to the empirical and policy literature evaluating the impact of privacy protections on technological progress. Unlike some earlier findings presented so far in the literature (see Section 2), these results offer evidence for a possible nuanced and sometimes beneficial role of privacy regulations in encouraging technology adoption.

2. Related Literature

2.1 Innovation, Regulation, and Privacy

In the United States, baseline protection of personal information has been explicitly mandated by law in a few sectors, such as healthcare and financial services.² In other sectors, such as online commerce, policymakers' intervention in the market for personal information has been minimal, and self-regulation has been preferred (Mulligan and Goldman 1997; Tang, Hu, and Smith 2007). The debate over the proper way to address privacy protection remains intense; several bills related to consumers' privacy rights are currently under review in Congress.³

When considering whether to enact new privacy legislation, policymakers face a trade-off. Regulating the use of consumer data can protect individuals' privacy and increase consumer welfare (for instance, by making identity theft less likely; Romanosky, Telang, Acquisti 2011). However, it can also increase firms' costs and decrease efficiency (for instance, by imposing additional technological controls or

² Consider, for instance, the Health Insurance Portability and Accountability Act (P.L.104-191) for healthcare providers, and the Gramm-Leach-Bliley Act of 1999 (P.L. 106-102) for financial institutions.

³ They include the Do Not Track Kids Act of 2011 (H.R. 1895), the Do-Not-Track Online Act of 2011 (S. 913), and the Commercial Privacy Bill of Rights Act of 2011 (S.799).

administrative procedures to protect individuals' data). For instance, early analysis of privacy economics by scholars such as Stigler (1980) and Posner (1981) implied that privacy regulation may create market inefficiencies. These scholars posited that privacy regulation limits the availability of information necessary for the efficient operation of markets. One particular drawback of privacy regulation that has caught the attention of scholars and policymakers is the potential for a dampening effect on technology adoption and innovation. For example, Lenard and Rubin (2005, 2006) argue that laws related to notifying consumers of data breaches may impede e-commerce and stifle technological development by discouraging firms from innovating by using consumers' personal information (or stop collecting it altogether).

On the other hand, there are reasons to believe that the effects of privacy regulation on consumers' and firms' welfare may not always be negative: privacy legislation could encourage adoption and acceptance of privacy-sensitive technologies (such as behavioral advertising), increasing consumer welfare (for instance, through customized ads) and simultaneously increasing firm profitability. In fact, some scholars have argued that privacy regulation provides necessary assurances to users and subjects of these systems, encouraging participation, adoption, and acceptance of new services or technologies (McGraw et al 2009; Bamberger and Mulligan 2011).⁴

Notwithstanding the substantial debate between privacy advocates and data industry lobbyists, the empirical work evaluating the impact of privacy regulation on technology innovation and adoption has been limited to date. Current evidence seems to support the view that privacy restrictions have a negative effect on technology innovation, adoption, and effectiveness. Goldfarb and Tucker (2011a) found that display advertising became far less effective at impacting consumers' stated purchase intent after the

⁴ On some occasions, industries have even sought increased privacy regulation in order to push forward potential privacy sensitive technologies. The Electronic Communications Privacy Act of 1986 (ECPA), which updated the 1968 Wiretap Act, was the result of such collaborative public interest/private sector efforts: the industry feared that, without legal protection against eavesdropping and interception, consumers would be reluctant to use emerging electronic media, such as cellular phones and email.

enactment of EU laws restricting advertisers' ability to collect data on web users. Specific to the impact of privacy regulation in healthcare and technology adoption, Miller and Tucker (2009) found that disclosure laws inhibited the adoption of EMR that are thought to have positive externalities. Actions of industry echo this conclusion, with attempts to amend long-standing privacy regulation to ease the burdens on industry. For example, legislators have recently introduced a bill to modify the Video Privacy Protection Act of 1988 to allow the sharing of user viewing behavior subject to consumers' consent. This initiative has been spurred primarily by a Facebook/Netflix collaboration aimed at facilitating the sharing of viewing habits with friends online – and marketers too.⁵

2.2 Health Information Exchanges and Patient Privacy

Health Information Exchanges (HIEs) allow electronic health information sharing between hospitals, health plans, pharmacies, laboratories, physicians, and other relevant entities. They rely on digital health records (e.g. in the form of an EMR system) and often start as regional partnerships of healthcare stakeholders (e.g. several neighboring counties within a state) that agree to share their patient health information via a shared technology platform and also abide by terms of use (covering a wide range of issues, such as patient privacy). HIEs are almost exclusively initiated as independent non-profits, or under the umbrella of a non-profit entity (only three of the 88 operational HIEs that I identified are operating as independent for-profit organizations). While early adopters of health information exchange efforts emerged in the early-nineties (the oldest exchange effort in my dataset was initiated in 1990), only in the last decade has there been significant growth in HIE activity, including the number of HIEs attempted and an increasing number of HIEs actually facilitating electronic exchange of health information (i.e., operational): in 2004 there were only .25 operational HIEs per state, compared to 1.5 as of 2010.

⁵ Amendment to 2710 of title 18, Video Protection Act, United States Code (H.R. 2471).

HIEs promise considerable welfare gains. For instance, increased availability of patient information can lead to more informed treatment decisions, and exchanges may reduce redundant testing by allowing a physician access to test results performed by a patient's prior physicians.⁶ However, the success of HIEs can be hindered by sustainability issues, misaligned incentives from competing healthcare entities, and technological and interoperability constraints (Vest and Gamm 2010). In particular, data privacy and security challenges have elicited concerns from legislators, privacy advocates, and patient rights groups. A significant body of work has looked at patient and physician attitudes towards health privacy and the sharing of personal health information. Sankar et al (2003) found that patients dealing with increasingly sensitive information (e.g. HIV-related information) are less likely to share their medical information. Simon et al (2009) performed a qualitative study of patients' attitudes towards HIE and found that privacy, security, and consent issues consistently came up as critical concerns. Simon et al (2009) also found that when information was provided to the patients regarding both HIE benefits and privacy/security protections, the overwhelming majority stated that they would consent to inclusion of their personal health information.

Experts hold conflicting views on the impact and applicability of existing legislative solutions to protect patient privacy in the context of HIE. While Greenberg et al (2009) and McDonald (2009) agree that federal protections need to be revisited in light of the emergence of a National Health Information Network (which is envisioned to ultimately link regional and state-level HIEs), they differ on the need for updating state protections. McDonald (2009) suggests that new restrictions beyond the protection afforded by HIPAA would interfere with efficient and safe care. Greenberg et al (2009) advocates updates to state legislation to better address privacy issues specific to HIEs.

⁶ Jha et al (2009) suggests that nearly 8 billion dollars could have been saved in 2004 from eliminating redundant tests.

Various states have attempted to address HIE privacy concerns by passing legislation specifying privacy and security requirements specific to HIE efforts. Across states, there has been significant variation in the extent to which regulation provided patients the option to consent prior to the use of their data in an exchange. These laws form the basis of my analysis, and are discussed further in section 4.

3. The Role of Regulation in HIEs' Adoption and Success

To understand the role of privacy protection in the adoption and success of HIEs, I analyzed the impact of privacy-relevant HIE regulation on the number of attempted HIEs, as a measure of HIE adoption, and operational HIEs (i.e., those actually exchanging data), as a measure of successful HIEs. Because the success of an HIE is heavily dependent on the participation of potential adopters, in this section I first consider the broader motivations for exchange (in terms of benefits and costs) of an HIE for its potential adopters (such as hospitals, health plans, pharmacies, laboratories, and physicians), and then discuss how privacy regulation can affect their trade-offs.

3.1 Adopters' Benefits and Costs

Aggregate benefits of HIEs can be categorized into two broad categories: cost savings and quality of care gains. These benefits are enjoyed in various proportions by HIEs stakeholders - their adopters, but also their patients. Health information exchanges have the potential to significantly decrease the costs of providing healthcare. Walker et al (2005) estimate that, when fully implemented, health information exchange could yield approximately 78 billion dollars in annual savings. Jha et al (2009) estimate that, in the US, eliminating avoidable instances of injury to a patient resulting from a medical intervention, such as administering the wrong medication, and redundant medical tests would save over 24 billion dollars in a single year. Because of the substantial potential cost-savings, federal and state governments, as the largest medical payers in the United States, have undertaken an array of policy activities in support of

health information exchange, including funding states to encourage HIE growth.⁷ Potential cost savings, however, may not only be reaped by medical payers (such as insurers). For example, savings from reducing redundant testing may improve hospital profitability under a prospective payment system, where Medicare and many state Medicaid programs reimburse hospitals a flat amount per admission based on the diagnosis group. This is also the case for emergency rooms, where in addition to Medicare/Medicaid, many private insurers pay a fixed fee. Gains in quality of care may also be realized due to the increased availability of comprehensive health information, which should allow clinicians to make better treatment decisions and fewer mistakes. This benefit would be especially salient in the emergency care context, in which the patient may not be able to report pre-existing conditions or drug allergies. To date, the benefits with respect to HIE are still uncertain: little empirical evidence exists to support (or refute) claims of such benefits. Establishing an HIE requires, however, substantial capital investment to support the technology and administrative infrastructure (Vest and Gamm 2010). The cost of technology efforts intended to facilitate health information exchange was a key barrier to early exchange efforts (Vest and Gamm 2010). In recent years, the emergence of large vendors providing HIE solutions has reduced technology costs for exchange.⁸

Because patient health information is generally regarded as confidential and sensitive, and is governed by an array of state and federal laws, entities that choose to participate in exchange may also incur a cost of data disclosure. As entities expand the sharing of their data (particularly to entities outside of their organizational purview), they run the risk of other entities improperly handling or disclosing personal health information. In terms of direct costs, inappropriate disclosure of sensitive health information in some states is treated as a criminal offense, and other states provide legal avenues to seek damages based on inappropriate disclosure of sensitive health information. These liabilities become even more severe in

⁷ The Health Information Technology for Economic and Clinical Health (HITECH) established the “State Health Information Exchange Cooperative Agreement Program” which provides incentives for states to pursue HIE efforts (42 U.S.C. § 3013).

⁸ This downward trend in cost of HIE technology solutions is captured in the model by the time fixed-effect term.

the context of highly sensitive health information, such as HIV or mental health data. Disclosure costs may also include less tangible costs such as reputational degradation and a decline in trust from patients (outside the healthcare industry, Acquisti, Friedman, and Telang 2006 showed that companies that suffered data breaches experienced a significant short-term drop in market value). Moreover, exchanges may face costs in responding to concerns or pressure from patients, regulatory bodies, or the advocacy community stemming from the exchange of sensitive patient data.

3.2 The Impact of Regulation

Adopters' benefits and costs associated with the development of HIEs are also impacted by regulation. While I am not able to observe in my data benefits and costs associated with individual HIEs for their various stakeholders (and in particular their potential adopters), I can observe the total number of exchanges attempted and the ones that become operational (see Section 4). These observed variables are a direct function of the expected benefits and costs of potential HIE adopters, and are therefore impacted, in turn, by HIE-relevant legislation. In my empirical analysis I consider two types of legislative initiatives: 1) regulation that may promote the adoption of HIEs in a state, and 2) privacy regulation, and in particular initiatives establishing requirements for patient consent in the context of exchange.

The potential impact of HIE promoting legislation is fairly obvious. Some states have passed legislation that actually creates an exchange in the state or reduces the cost of implementing exchange efforts by awarding grants to entities interested in initiating exchange efforts. Other states have also moderated HIE costs or reduced uncertainty surrounding attempting an exchange by establishing government bodies that define a vision for state-wide health information exchange and provide guidance on instituting a governance structure or developing a sustainable business model. The role of *privacy* legislation is comparatively ambiguous. As noted earlier, the literature is ambivalent about how privacy regulation could affect technology adoption. In short, privacy regulation may increase benefits if it reassures potential adopters by assuaging the concerns of the patients they serve, increasing adoption of exchange

efforts; it may also hinder these efforts, if privacy regulation increases administrative costs and decreases patient participation.

Central to the privacy debate in general, and to the HIE privacy debate in particular, is the issue of patient consent (consent, or informed consent, are cornerstones of the OECD's privacy guidelines⁹ and the Federal Trade Commission's Fair Information Practice principles). HIE adopters and patients seem to have somewhat opposing views towards patient consent. HIEs have expressed concerns that imposing consent requirements, particularly those that require patients to provide consent prior to the inclusion of their data in an exchange (i.e. opt-in), would add administrative costs and limit patient participation (National eHealth Collaborative 2011, Pritts et al 2009). In contrast, Simon et al (2009) found that patients felt that they should have to provide consent for health information exchange (i.e., an opt-in system); a system that assumed their willingness to participate without obtaining explicit consent (i.e., an opt-out system) would not be acceptable to them. Anecdotal evidence on the impact of obtaining patient consent on HIE success and adoption has been mixed.¹⁰ There have been claims that overly restrictive consent models stifle the growth and information sharing capabilities of some HIEs (Pritts et al 2009), and yet HIEs have successfully implemented methods of obtaining patient consent with relatively high rates of opt-in by patients (eHealth Initiative 2007). As a result, different states have responded differently to patient consent requirements. Some states enacted legislation promoting the adoption of HIEs (for instance, providing funding for HIE activities) but also imposing requirements for patient consent; other states enacted similar pro-HIE legislations, but did not impose similar consent requirements; and yet other states did not enact any HIE-specific legislation. However, the role of such requirements is not obvious and requires additional discussion.

⁹ Organisation for Economic Cooperation and Development (OECD), OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (Sep. 23, 1980).

¹⁰ In a number of discussions, we found that experts and HIE administrators recognized the importance of preserving patient privacy but also had some concerns about over regulating the use of personal health information to the extent that it would make HIE overly costly.

As consent requirements are likely to impact the amount and type of patient data available for exchange, they could have a significant impact on the benefits. On the negative side, requirements for patient consent may result in limited or patchy patient agreement to have their data included in the HIE (Lai and Hui 2006), in which case the potential benefits to HIE adopters may be hindered. For instance, in a recent report (National eHealth Collaborative 2011), a number of HIEs suggested that requiring patients to opt-in to an HIE was a barrier to achieving the critical mass of patient records needed to accelerate adoption. However, it is not clear that consent requirements will necessarily lead to lower levels of patient health records in an exchange. While patients in states *without* consent requirements may not have the choice to participate in an exchange, the same is not true for healthcare providers. Healthcare providers may in fact decide *not* to initiate or join HIE efforts in those states in the face of patient privacy concerns or to avoid the threat of privacy-related lawsuits. McGraw et al (2009) provides support for this view and argues that a comprehensive framework that implements core privacy principles can bolster trust from patients and medical providers, thus promoting adoption. Lastly, Simon et al (2009) suggests that levels of patient consent could be relatively high: when study participants were provided information about health information exchange, 88% stated they would consent to the use of their health information.

Requirements for consent are also likely to impact the technology and administrative costs associated with implementing an HIE. On one hand, HIEs operating in states with more stringent consent protections may require additional investment in technical and administrative controls to meet regulatory requirements (e.g. clerical time by staff or technical controls to garner and track patient consent decisions). However, it is not clear that consent requirements will always lead to greater costs from the adoption of technological and administrative controls to protect patient privacy. It may be the case that more protective legislative environments (those dealing with consent and other privacy issues) may in effect force the “privacy issue,” resulting in HIEs that are foresighted in terms of handling privacy concerns and developing mitigating technologies and policies. This increased foresight and attention to

privacy may help HIEs avoid expensive and time consuming retrofitting and other roadblocks in the future as a result of patient privacy concerns. In terms of disclosure costs, consent requirements may increase the liability of healthcare providers in the event of a privacy intrusion due to an exchange (e.g. inappropriate disclosure of health information). Conversely, clear guidelines for protections governing the use and exchange of sensitive health data, such as consent requirements, may result in reduced disclosure costs by making privacy intrusions or litigation following a privacy invasion less likely. For example, patients that are properly informed about the benefits and risks associated with exchange but still provide their consent for the use of their data in exchange efforts may be less likely to seek damages in the event of a privacy intrusion. Lastly, exchanges operating in states with consent requirements may face lower disclosure costs in terms of pushback from privacy advocacy or patient rights groups.

A related issue is the possible detrimental effects of *regulatory ambiguity* on technological adoption (Marcus 1981). For instance, Leonardi (1978) argued that uncertainty in the interpretation of clean air laws hindered the development of fuels made of pulverized coal and other coal-related technologies. Similar uncertainty is likely to pervade the HIE environment. A recent report on state disclosure requirements suggests that “states differ greatly in the way their statutes address personal health information (PHI) types, PHI holders, PHI receivers, different treatment scenarios, consent processes and forms, and requirements for HIPAA’s minimum necessary standard.” (Pritts et al 2009). As a result, privacy regulation with explicit and broad-reaching requirements for consent may mitigate costs faced by exchange simply from identifying and disambiguating disclosure requirements and their applicability to exchange.

Finally, future privacy regulation in the context of exchange is still under debate in many states and on the federal level (Greenberg et al 2009; McDonald 2009), which generates further uncertainty about potential changes to legislation that may impact disclosure costs. Extensive work has focused on the role of *regulatory uncertainty* on the impact of firm investment and innovation. A number of scholars (Luo 2004;

Ishii and Yan 2004; Bittlingmayer 2001) argue that postponing investment to gather more information and assurances regarding future regulatory changes is a common reaction to regulatory uncertainty. Regulatory uncertainty is especially relevant in the context of HIEs, given that states that passed HIE legislation without considerations for consent often relied on pre-existing legislation, and thus made no special considerations to reduce the uncertainty surrounding future privacy requirements for exchange or inform the current consent debate. In fact, some states that introduced HIE legislation without consent requirements explicitly introduced uncertainty with respect to future legislative action regarding patient privacy. For example, Maryland state law (Md. Code Ann. § 19-143) calls upon future legislators to “[e]valuate any changes in state laws that are necessary to protect the privacy and security of health information stored in electronic health records or exchanged through a health information exchange in the State.” As a result, it may be the case that states that pass regulation specifying consent requirements reduce uncertainty for potential entities interested in pursuing exchange efforts thus encouraging HIE attempts and success.

3.3 Modeling approach

In summary, arguments can be made for both positive and negative impacts of consent requirements on HIE’s net benefits, and therefore the likelihood of healthcare organizations forming them, and their continued operation. This is summarized in a conceptual model (extending the one in Miller and Tucker 2010) that illustrates the relationship between my observed dependent variables and the underlying costs and benefits for potential HIE adopters:

$$\text{AttemptedHIE}_{jst}, \text{OperationalHIE}_{jst}^* = f(\text{NetRegionBenefit}_{jst} \mid \text{PatientConsentRegulation}_{jst}, \text{HIEPromotingRegulation}_{jst})$$

The model suggests that the existence of legislative initiatives promoting HIEs and regulating patients’ consent to the usage of their data will impact the net benefit of an exchange. This, in turn, will impact the

likelihood of an HIE being attempted, and the likelihood of its continued operation. Note that this model assumes a latent variable construct where region j (e.g. a county or group of neighboring counties) in state s at time t attempts an exchange if the net benefit of an exchange in the region is positive, and the exchange in this region continues to operate after inception if the net benefit of exchange remains positive. Because HIEs have emerged as regional efforts, I define my conceptual model at the regional level. However, legislative initiatives related to HIEs were enacted at the state level. Hence, my empirical analysis aggregates the counts of *AttemptedHIEs* and *OperationalHIEs* at the state level. Due to the relatively short period of my analysis, systemic differences between both regions within a state and across states (such as hospital market structure or socioeconomic factors) are likely constant in the time-period of my analysis. Individual HIE characteristics are also considered in my analysis, even though they are not explicitly represented in the conceptual model.

In the rest of the chapter, I seek to understand the net effect of instituting consent requirements on HIEs attempted and HIEs reaching operational status.

4. Data

My analysis uses 6 years of data to assess the impact of the different legislative approaches to HIEs on the number of attempted and operational HIEs across states.

4.1 HIE Data

Consistent with the literature, I define an HIE as any project or initiative focused on electronic patient health data exchange between two or more disparate organizations or stakeholders.

In order to identify attempted HIEs across states, I utilized publicly available data from the eHealth Initiative's (eHI) annual compilation of state, regional, and local HIE efforts (eHealth Initiative 2005-2010). This data is based on yearly surveys of HIEs completed by the eHealth Initiative and provides

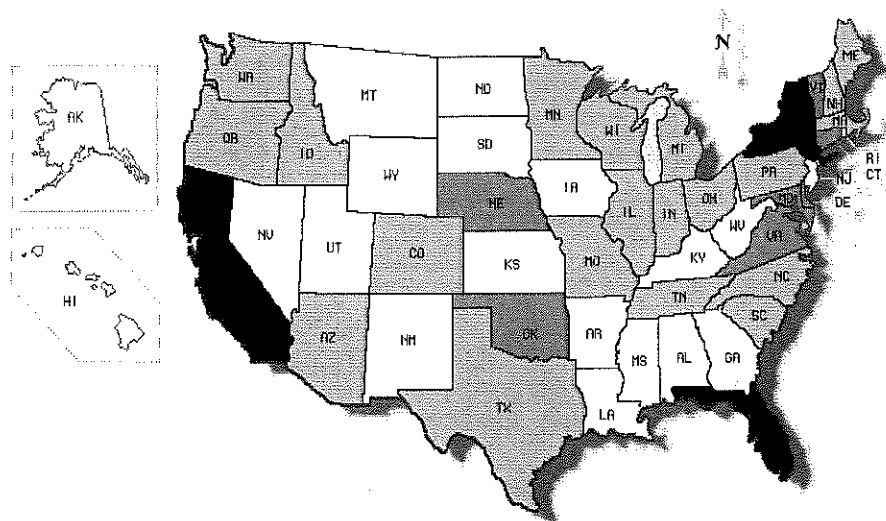
longitudinal information about planning and operational HIEs in the 2004-2010 period. I also used various online resources provided by health organizations and individual exchanges to determine their status as of 2010 and collect any additional information on characteristics of these exchanges (e.g. for-profit status). As noted earlier, at the beginning of 2004, there were only a handful of established exchanges. As of 2010, I identified 312 HIE initiatives that were in one of three stages:

1. Planning: The exchange has been initiated but continues to be in the early stages of development and is not actively sharing health information (n=132).
2. Operational: The exchange has been initiated and is actively sharing health information between providers (n=88).
3. Failed: The exchange was initiated but had ceased operations as of 2010 (n=92).

I complemented the eHI data with a national survey of Health Information Exchanges collected in 2010 (Adler-Milstein, Bates, and Jha 2011). The survey has responses from 73 planning exchanges and 75 operational exchanges (summary statistics for key variables are provided in Table 1 below). Through various online resources and correspondence with exchanges, I was able to obtain detailed information for an additional six exchanges, for a total of 154 exchanges in my dataset (78 Planning and 76 Operational) or 70% of the 220 planning and operational exchanges and 87% of the 88 operational exchanges. This survey data also provided details on individual characteristics of these exchanges and more granular information on when they were formed and became operational. The resulting dataset follows these exchanges from 2004 to 2010 - the period when HIEs came into prominence and saw significant growth (in 2004 there were only .25 active HIEs per state compared to 1.5 active HIEs as of 2010) and also the period in which the eHealth Initiative annual surveys were administered. I coded my data semi-annually to more finely capture changes in HIE status (i.e. whether they were planning or operational) and the time that my laws of interest were enacted. From this dataset I generate my two dependent variables:

- The number of operational HIEs in a state as of 2010 is reported in Figure 1 and summary statistics on the HIEs that responded to the HIE survey are provided in Table 1.

- ◆ - One
- ◆◆ - Between Two and Four
- ◆◆◆ - Five or More



25

[Table 1: Summary Statistics on HIEs in Survey Data]

Variable	Description	Mean	S.D.	Min	Max	Obs.
Planning and Operational Exchanges						
YearsPursuing	Years an HIE has been pursuing exchange efforts	4.07	3.18	.33	20	146
IndependentOrg	Percent of exchanges functioning as independent organizations	.38	.488	0	1	148
FormalGov	Percent of exchanges with formal governance structure	.85	.357	0	1	147
In-Kind Resour.	Reliance on Time-In-Kind Resources	1.48	.674	1	3	139
OneTimeContrPay	Reliance on one time contribution from payers	1.75	.78	1	3	131
RecurringFeePay	Reliance on recurring fee from payers	1.76	.85	1	3	130
HighGov	Percent of HIEs indicating heavy reliance on government funding	.56	.49	0	1	130
Operational Exchanges						
Years Operational	Years an HIE has been sharing health data	3.45	3.69	.08	18.3	75
Lab_Provide	Number of labs providing data	6.8	13.5	0	100	60
Pharm_Provide	Number of pharmacies providing data	39.36	19.4	0	1000	49
Hospital_Provide	Number of hospitals providing data	9.66	19.4	0	118	65
AmbPrac_Provide	Number of Ambulatory Practices providing data	71.7	286	0	2184	60
PubHealth_Provide	Number of public health agencies providing data	2.02	12.4	0	90	52
PrivPay_Provide	Number of private payers providing data	1.73	3.32	0	14	53
PubPay_Provide	Number of public payers providing data	.6	1.49	0	10	51
PercAmbProv	Percent of Ambulatory Practices in region providing data	.237	.309	0	1	53
PercBedsProvide	Percent of hospital beds providing data	.57	.357	0	1	54
PercBedsReceive	Percent of hospital beds receiving data	.596	.368	0	1	54
Results	Percent of HIEs sharing lab results	.91	.28	0	1	71
Inpatient	Percent of HIEs sharing Inpatient data	.848	.36	0	1	66
OutPatient	Percent of HIEs sharing Outpatient data	.88	.32	0	1	68
PatientNumb (1000)	Number of patients in an exchange	208	227	0	>500	75
Sustainable	Percent of operational exchanges financially sustainable	.33	.474	0	1	75
MonthsSustainable	Number of Months an HIE has been financially sustainable	18.2	16.6	0	60	24

4.2 Legislation

HIE activities are governed by a combination of federal and state laws that regulate the disclosure of health information by various healthcare entities (e.g. physicians, hospitals, insurers, etc.). At the federal level, sharing and use of health data is governed primarily by the Health Insurance Portability and Accountability Act (HIPAA)¹² and associated regulation, which lay out requirements that address, among other things, patient consent, patient access to health records, use of de-identified health data, and security standards for health data. HIPAA was amended in 2009 by the Health Information Technology for Economic and Clinical Health (HITECH) Act, which added some additional privacy requirements, including breach notification requirements for HIPAA covered entities.¹³ While HIPAA laws are likely to have an impact on the disclosure of health information by medical providers, HIPAA applies to all states (my analysis relies on between-state variation) and was passed prior to the time-period of my analysis (See Figure 2 below). HITECH was passed in my period of analysis, but any effect should be picked up by the time fixed-effects element of the model.

At the state level, I considered two categories of legislation: 1) general privacy health laws, not HIE-specific, that were enacted before the significant emergence of HIEs; 2) HIE-specific legislation, aimed at promoting HIE activities and/or focusing on disclosure of patient data and their consent

General health privacy laws (i.e. not HIE-specific) have historically been in place to deal with various aspects of health privacy, such as patient consent, right to access, and patient privilege. I identified state health privacy regulation using the recent compilation of state disclosure laws by Pritts et al (2009), the earlier compilation of general state privacy laws by Pritts et al (2002), and the annual Privacy Journal's Compilation of State Privacy Laws (Smithe & Ryder 2002). However, similar to HIPAA, I found that the state health privacy laws that were not HIE-specific were passed prior to my period of analysis and are

¹² Health Insurance Portability and Accountability Act of 1996, 42 U.S.C. § 1320d-9 (2011).

¹³ Health Information Technology for Economic and Clinical Health Act of 2009, U.S.C. § 3013 (2011)

thus captured by the state fixed effects in the model (some states, such as New Mexico, have passed health privacy regulation *during* the time period of my analysis that explicitly mentioned HIE or exchange. Because I consider these laws as HIE-specific, they are captured in the extended model by the *DisclosureOnly* measure, described below). Moreover, patient consent requirements provided in the majority of non-HIE specific health privacy laws included exceptions to garnering patient consent (specifically those relating to data disclosures for treatment purposes) that effectively precluded the majority of exchange activities. According to Pritts et al (2009), only two states (Minnesota and New York) appear to generally require patient permission to disclose all types of health information and only a few states (Guam, Puerto Rico, New York, Minnesota, and Vermont) usually require hospitals and physicians to obtain patient permission before disclosing health information to other providers. In other words, general health privacy laws that are not HIE specific were passed prior to my period of analysis and their requirements for consent have limited applicability to HIEs; as a result, they are not independent variables of interest in my analysis.

Conversely, HIE-specific legislation was passed in the period of my analysis and has direct applicability to exchange efforts. It is therefore the independent variable of focus for my analysis. Within the last decade, various states enacted legislative initiatives specific to HIEs and/or specific to HIEs disclosure activities. These laws were enacted in the same period as the emergence of HIEs, and are predominately designed to promote exchange within a state, or to impose requirements specifically in the context of exchange. Given that most states¹⁴ do not have requirements to obtain patient consent prior to disclosing

¹⁴ We say most here because New York, Minnesota, and Vermont have some requirements that do generally require consent for disclosure between providers. These states were considered as having consent requirements and are ProHIE & Consent states as they would all subsequently pass HIE-specific legislation (we did not have HIE data for Guam and Puerto Rico).

health information¹⁵ to other providers (which are also the majority of HIE adopters), requirements for consent in HIE-specific laws are especially relevant to the disclosure of health information by exchanges.

I identified HIE-specific laws primarily through various legal search services (e.g. LexisNexis academic and Westlaw) and supplemented these searches with recent reports on disclosure laws and Health Information Exchange (Goldstein and Rein 2010). Most of these laws were enacted after 2007 and include various incentives and requirements for exchange (See supplemental document for some examples of these legislative actions). I identified three primary HIE-specific legislative approaches for my analysis:


- ProHIE and Consent: States with laws intended to promote HIE (such as creating an exchange, providing funding for HIE activities, or designating a government entity to facilitate health information exchange) and also had requirements for consent. (8 states)
- ProHIE and No Consent: States with laws intended to promote HIE, make some mention of privacy protections but do not include requirements for consent (i.e. they rely on the status quo of no consent requirements for the exchange of health information between providers). (11 states)
- No HIE Legislation (No Consent): States without any HIE-specific legislation, and therefore also no explicit requirements for consent. (25 States)

A few remaining states exhibited other variants of HIE-specific legislation: three states and the District of Columbia enacted legislations that included incentives for HIEs but no mention of patient privacy (ProHIE Only); three additional states had some privacy protections (although not necessarily consent requirements) in the context of exchange but no matching incentives. I classified these states as

¹⁵ States have passed more stringent laws for some specific and sometimes sensitive health data (e.g. mental health or HIV data). Because this type of data is generally not the focus of HIEs we focus only on laws restricting the exchange of general health information.

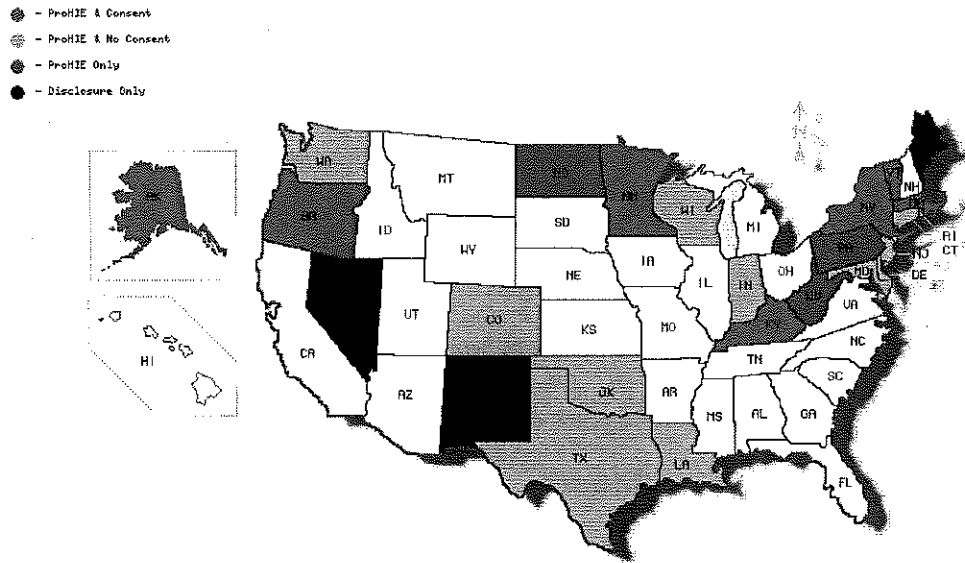
DisclosureOnly states and included them in my extended model. This leaves 25 states without HIE-specific legislation (Figure 3 below illustrates which states have enacted HIE-specific legislation). The various laws and their passage relative to my analysis is summarized in Figure 2.

[Figure 2: Timeline of Health Privacy Legislation]



Given the history of lax enforcement of privacy requirements in HIPAA (Nahra 2008), one may be concerned with similar issues regarding the enforcement of privacy requirements in HIE laws. The case of HIEs is likely a special case, in which enforcement is less of a concern. First, there are often only a handful of highly visible HIEs in a given state, and they tend to be heavily reliant on state and federal funding for support (which almost always come with regular reporting requirements). Second, privacy advocacy groups have identified HIEs as potential sources of privacy invasions, and as such continue to carefully monitor and scrutinize HIE activities and privacy procedures (Miliard 2010).

[Figure 3: States with HIE-Specific Laws]



5. Methods

I first specify a basic fixed effect model to evaluate the impact of HIE promoting legislation in the aggregate (i.e. without including controls to capture states that also had requirements for consent) on *AttemptedHIEs* and *OperationalHIEs*.

Model 1: Basic Fixed Effects Model

$$AttemptedHIE_{st} \text{ \& } OperationalHIE_{st} = \beta_0 + \beta_1 * ProHIE_{st} + \sum \delta_k * X_{st} + \theta_s + \lambda_t + \mu_{st}$$

AttemptedHIE and *OperationalHIE* have been discussed in some detail in the prior section and are defined again in Table 2 below. Similar variables have been used as key measures of HIE adoption and success in other evaluations of HIE progress (eHealthInitiative 2005-2010; Adler-Milstein et al 2009,2011). Although I was able to obtain data on other relevant measures of HIE progress (failed exchanges, the breadth of sharing, participation of various entities, etc.), this data was only available as of

2010. It is therefore not included in my primary longitudinal analysis, but I do utilize it to examine possible endogeneity and differences across different legislative approaches. $ProHIE_{st}$ is a dummy variable (I consider differences across states with consent requirements in the following model) indicating whether a state s had any HIE promoting legislation at time t (as a reminder, t here represents semi-annual intervals). I include a vector of control variables, X_{st} , that accounts for other relevant factors. For example, because these laws were passed to encourage HIE adoption and growth in a state, X_{st} includes controls for other major provision of the HIE law. Specifically, I include dummies to control for provisions in the laws providing funding for HIE activities and designating/creating a state sponsored HIE. Additionally, HIE efforts require that regional players have some minimum level of patient record digitization and health IT infrastructure in order to make for meaningful exchange. As a result, I control for state HIT adoption by including $EMRA_{st}$ and $CPOE_{st}$ to capture hospital adoption of electronic medical records (EMR) and Computerized Provider Order Entry (CPOE) adoption (respectively).¹⁶ I also include a dummy variable indicating if a state has an established HIE (defined in Table 2 below) to capture any possible effects of having a long-standing HIE on additional entry or other HIEs becoming operational. For example, it may be the case that having an established HIE inhibits future entry, thus leading to less operational and attempted HIEs. Lastly, I include controls to capture state population and wealth effects. State and time fixed effects are represented by θ_s and λ_t (respectively) and μ_{st} is the familiar error term. This state, time fixed effect model has been used in the literature to examine the effect of a policy intervention (Bertrand, Duflo, and Mullainathan 2004; Romanosky, Telang, Acquisti 2011). State fixed effects allow us to control for unobserved state specific factors and time dummies allow us to control for time trends. Thus, the unbiased effect of various HIE-specific legislation can be identified from variation across state and time.

¹⁶ From the Health Information and Management Systems Society (HIMSS) Analytics™ Database (HADB), we derive measures of hospital adoption of EMR and CPOE technologies normalized by hospital size measured by number of beds.

Model 2: Extended Fixed Effects Model

$$\text{AttemptedHIE}_{st} \text{ \& OperationalHIE}_{st} = \beta_0 + \beta_1 * \text{ProHIE}_{st} + \beta_2 * \text{ProHIE\&Consent}_{st} + \sum \delta_k * X_{st} + \theta_s + \lambda_t + \mu_{st}$$

I extend my basic model by adding *ProHIE&Consent_{st}*, in order to disentangle the impact of states with HIE promoting legislation and requirements for patient consent from those states that have HIE-promoting legislation, but did not have requirements for patient consent. Note that with the addition of the *ProHIE&Consent* variable, *ProHIE* in this model now captures the impact of states with HIE promoting legislation without requirements to obtain patient consent.

Model 3: Full Fixed Effects Model

$$\text{AttemptedHIE}_{st} \text{ \& OperationalHIE}_{st} = \beta_0 + \beta_1 * \text{ProHIE\&Consent}_{st} + \beta_2 * \text{ProHIE\&Consent}_{st} + \beta_3 * \text{ProHIEOnly}_{st} + \beta_4 * \text{DisclosureOnly}_{st} + \sum \delta_k * X_{st} + \theta_s + \lambda_t + \mu_{st}$$

Lastly, the full fixed effects model parses out all variants of HIE-specific legislation with the addition of *ProHIEOnly_{st}* and *DisclosureOnly_{st}*. This allows us to evaluate the impact of HIE promoting legislation that had requirements for patient consent relative to all other forms of HIE-specific legislation.

Specifically, I can parse out any differences between states that had had some mention of privacy requirements but did not institute consent requirements (*ProHIE and No Consent*), states that only had HIE promoting legislation and no mention of patient privacy (*ProHIEOnly*), and states with only disclosure requirements in the context of exchange (*DisclosureOnly*).

Across these three models, the omitted variable is states with no HIE-specific legislation. All of the measures in my analysis are summarized in Table 2 below.

[Table 2 Summary of Key Measures]

Variable	Description	Source
Dependent Variables		
TotalHIE_{st}	The Number of Planning + Operational HIEs in state s at time t that had not failed as of 2010.	HIE Survey
OperationalHIE_{st}	The total number of HIEs actively exchanging data in state s at time t.	HIE Survey
Independent Variables		
ProHIE_{st}	Dummy variable indicating whether a state s at time t enacted any law intended to promote HIEs.	WestLaw / LexisNexis
ProHIE&Consent_{st}	Dummy variable indicating whether a state s at time t enacted laws intended to promote HIE and also explicitly require HIEs to obtain patient consent.	WestLaw / LexisNexis
ProHIE&NoConsent_{st}	Dummy variable indicating whether a state s at time t enacted laws intended to promote HIE, make some mention of privacy protections, but do not include explicit requirements to obtain patient consent.	WestLaw / LexisNexis
ProHIEOnly_{st}	Dummy variable indicating whether a state s at time t enacted laws intended to promote HIE but made no mention of privacy protections.	WestLaw / LexisNexis
DisclosureOnly_{st}	Dummy variable of indicating whether a state s at time t enacted laws with some privacy requirements for HIEs but do not encourage HIE.	WestLaw / LexisNexis
Control Variables		
Funding_{st}	A dummy variable indicating whether HIE-specific legislation at time t explicitly provides funding opportunities for HIEs in state s.	WestLaw / LexisNexis
StateDesignated_{st}	A dummy variable indicating whether HIE-specific legislation in state s at time t identifies a specific entity or HIE to receive state support	WestLaw / LexisNexis
Population_{st} (1M)	Number of inhabitants in a state	U.S. Census
Per Capital GDP_{st} (1000)	The per capital gross domestic product of a state s at time t	Bureau of Economic Analysis
EMRA_{st}	Percent of Hospitals adopting Electronic Medical Record systems normalized by hospital size of state s at time t.	HADB
CPOE_{st}	Percent of Hospitals adopting computerized provider order entry systems normalized by hospital size of state s at time t.	HADB
EstablishedHIE	A dummy variable indicating if state s at time t has an HIE that has been operational for more than 3 years and had > 50,000 patients as of 2010.	HIE Survey

6. Results

My results are presented in Table 3. Estimates of the basic model for the two dependent variables are presented in columns 1 and 5 respectively. I find that that when HIE promoting legislation is considered in the aggregate, it has a positive but not significant impact on *AttemptedHIEs* and *OperationalHIEs*.

Estimates for the extended model, which parses out the impact of states that had HIE promoting legislation and requirements for consent from states with ProHIE legislation and no requirements for consent (including ProHIE Only states), are presented in columns 2 and 6. I find a large and positive coefficient on ProHIE and Consent for *AttemptedHIE* ($P < .05$) and a similarly positive effect (although insignificant) coefficient for *OperationalHIE*. When I estimate the full model, I find large, positive, and significant ($P < .05$) coefficients on *ProHIE and Consent* for *AttemptedHIE* and *OperationalHIE* (columns 3 and 7). Estimates from the full model suggests that ProHIE and Consent legislation resulted in two additional HIEs attempted and .677 more operational exchanges when compared to no HIE legislation and also resulted in more attempted and operational HIEs than any other form of HIE legislation ($P < .05$). Given that the mean number of HIEs in a state as of 2010 is three, this represents a sizable increase in HIE attempts. In summary, across all models where I differentiate between states that have consent requirements and those without (excludes the basic model), I find a large, positive, and significant coefficient on *ProHIE and Consent* for *AttemptedHIE*. For *OperationalHIE*, I also find large and positive coefficients on *ProHIE and Consent*, but the estimate is significant in the full model specification and not in the extended model. The insignificant coefficients on *ProHIE with No Consent* and *ProHIE only* in the full model (columns 3 and 6) suggest that promoting HIE without requirements for consent may not have been effective in encouraging HIE growth. While not the focus of my analysis, these results suggest that legislation identifying a state-designated HIE had a stifling effect on HIEs attempted (a large and significant negative impact on *AttemptedHIE*). To help interpret the magnitude of these findings, consent

requirements may result in exchange for an additional 130,000 patients in a state with an operational HIE (using a conservative estimate of 210,000 patients covered by an exchange).¹⁷

Lastly, I find a negative but not significant ($P < .10$) coefficient on *Disclosure Only* legislation in the full model. Initially, this trend seems to contradict my prior results where states that had consent requirements exhibit better HIE outcomes. However, my definition of *Disclosure Only* states only required that they had privacy regulation specific to the exchange of health information but not necessarily that they require patient consent. For example, both Nevada and New Mexico included clauses in their laws to preclude HIPAA covered entities (i.e. majority of medical providers and payers) from the requirements of such laws. In any case, this result suggests a negative impact of HIE-specific privacy requirements without matching incentives and without explicit consent requirements.

I also explored the role of the various HIE regulatory approaches on HIE privacy concerns. The survey data on HIEs also included information on the extent to which privacy challenges have been an impediment to their progress. Figure 4 below provides a breakdown of the responses provided by the exchanges in the survey with respect to relevant HIE laws. I find that 58% of HIEs in states adopt a *ProHIE and Consent* approach reported that privacy concerns presented minor or no challenge in their development, compared to 35% of HIEs in states with No HIE law and states with *ProHIE Laws and No Consent* ($P < .05$). I see no difference between HIEs in states with *ProHIE and No Consent* and *No HIE Law*. This is not particularly surprising, given that states with *ProHIE and No Consent* laws generally point to pre-existing federal and state legislation to address privacy concerns. These results offer some additional evidence in support of the positive impact of consent requirements discussed earlier. This impact may reflect the role of legislative privacy protections in assuaging privacy concerns from exchange participants, patients, and regulatory bodies, or the increased early attention to privacy issues that more stringent regulatory approaches elicit.

¹⁷ This estimate is generated from 2009 Survey Data collected by Adler-Milstein et al (2009).

[Table 3: Panel Analysis of HIE Legislation and HIE Activity]

	AttemptedHIEs				OperationalHIEs			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ProHIE	0.315 (0.84)	-0.155 (0.59)			0.152 (0.84)	- 0.00003 (0.00)		
ProHIE & Consent		2.824 (2.01)**	1.958 (2.66)**	3.078 (2.19)**		0.915 (1.50)	0.641 (2.08)**	0.869 (1.63)
ProHIE & No Consent			-0.375 (1.15)				-0.084 (0.31)	
ProHIE Only			-0.276 (0.67)				-0.050 (0.16)	
DisclosureOnly			-2.866 (1.96)*				-1.105 (1.92)*	
State Designated	-1.068 (1.99)*	-2.092 (2.08)**	-1.836 (2.66)**	-2.461 (2.02)**	0.119 (0.32)	-0.212 (0.46)	-0.115 (0.30)	-0.285 (0.49)
Funding	0.335 (0.61)	-0.198 (0.52)	-0.203 (0.55)	-0.432 (1.10)	0.051 (0.18)	-0.122 (0.41)	-0.123 (0.35)	-0.233 (0.65)
Population	-1.962 (1.51)	-1.915 (1.68)*	-1.848 (1.80)*	-4.399 (1.52)	-1.720 (1.48)	-1.704 (1.54)	-1.679 (1.58)	-4.318 (1.72)*
Population Squared	0.00006 (1.64)	0.00006 (1.87)*	0.00006 (2.00)**	0.00001 (1.70)*	0.000 (1.55)	0.000 (1.61)	0.000 (1.63)	0.000 (1.86)*
PerCapitaGDP	-4.204 (0.71)	0.433 (0.06)	-1.179 (0.17)	8.543 (1.38)	-7.802 (2.49)**	-6.299 (1.56)	-6.909 (1.85)*	-3.410 (0.93)
EMRAoption	0.437 (0.45)	-0.001 (0.00)	-0.231 (0.28)	-0.570 (0.62)	0.032 (0.07)	-0.110 (0.27)	-0.199 (0.50)	0.169 (0.24)
CPOEAdoption	1.932 (1.43)	1.928 (1.46)	2.167 (1.58)	4.227 (1.79)*	0.902 (1.61)	0.901 (1.64)	0.992 (1.67)	1.730 (1.51)
EstablishedHIE	-0.113 (0.34)	-0.030 (0.09)	-0.217 (0.69)	-0.202 (0.41)	0.498 (1.96)*	0.526 (2.06)**	0.453 (1.84)*	0.407 (1.12)
Observations	612	612	612	235	612	612	228	235
Number of Groups	51	51	51	23	51	51	51	23
State F.E.	YES	YES	YES	YES	YES	YES	YES	YES
Time F.E.	YES	YES	YES	YES	YES	YES	YES	YES
R-squared	0.51	0.54	0.50	.50	0.43	0.45	0.47	.47
Robust t statistics in parentheses * significant at 10%; * significant at 5%; ** significant at 1%								

6.1 Robustness and Endogeneity of HIE-Legislation

My results in the prior section focus on the number of attempted and operational HIEs and may not necessarily allow us to conclude that a particular regulatory environment leads to “better” HIE outcomes.

For example, requirements for patient consent may spur HIE attempts but, in fact, HIEs in these states


may share less data and cover fewer patients and healthcare entities. Therefore, in Table 4, I use survey data capturing HIE activity as of 2010 to evaluate differences in the variety of data shared (results, inpatient, outpatient), the number of patients covered by an exchange, and the penetration of an HIE with respect to the hospital beds in its region. I find no significant differences in measures that capture the reach of an HIE in terms of participating healthcare entities or in the number of patients covered by the exchange. In summary, I find that states with consent requirements have more attempted and operational exchanges and that these exchanges have comparable levels of participation by healthcare entities, patients covered, and data shared as exchanges in states without consent requirements. I do not have sufficient data for the *Disclosure Only* and *ProHIE Only* states as they had five combined operational exchanges.

[Table 4: Measures of HIE Sharing (No significant differences across groups)]

Measure	Description	ProHIE & Consent	ProHIE & No Consent	NoHIE Law
PatientNumb (thousands)	Number of patients covered by an exchange	146.6	156.7	153.2
LabResults	Percent of exchanges sharing lab results	.833	.866	.97
Inpatient	Percent of exchange sharing inpatient data	.88	.92	.80
Outpatient	Percent of exchange sharing outpatient data	.823	1	.878
% Hospital Beds in Region Providing	Percent of Hospital Beds in a Region providing health information to the exchange	.67	.62	0.56
% Hospital Beds in Region Receiving	Percent of Hospital Beds in a Region Receiving health information from the exchange	.772	.568	.584
Lab_Provide_Yr	Number of labs added to an exchange yearly	1.4	2.56	.70
Pharm_Provide_Yr	Number of pharmacies added to an exchange yearly	1.46	2.12	2.77
Hospital_Provide_Yr	Number of hospitals added to an exchange yearly	6.9	8.13	2.09
AmbPrac_Provide_Yr	Number of Ambulatory Practices added to an exchange yearly	13.2	9.8	3.2
Obs.		18	15	34


The analysis of HIE laws also raises the concern of reverse causality: rather than HIE laws driving HIE activity, they could instead be passed as a result of increased HIE activity. In order to address this concern, I plotted *AttemptedHIEs* and *OperationalHIEs* for the main HIE legislative approaches I identified.¹⁸ Figure 4a and 4b below show that states that ultimately pass consent requirements did not have elevated levels of HIE activity. In fact, they had the lowest level of HIE activity when compared to other legislative approaches. More generally, prior to the period in which most HIE laws were passed (prior to 2007), there were minor differences in absolute levels of both *AttemptedHIE* and *OperationalHIEs*. However, as the data move into 2007 (most of HIE laws I identified were passed after 2007), states with *No HIE Laws* and *ProHIE and No Consent* maintain a roughly constant rate of growth, while states with *ProHIE and Consent* legislation see a sharp increase in both total HIEs attempted and operational exchanges.

[Figure 4a: Trend of Attempted HIEs Relative to ProHIE Legislation]



¹⁸ For clarity we do not plot “ProHIE Only” and “Privacy Only” states.

[Figure 4b: Trend of Operational HIEs Relative to ProHIE Legislation]



I further evaluate possible reverse causality using one- and two-year lead variables of all of the categorizations of HIE laws in my analysis. This allows us to evaluate whether the trends identified in my initial analysis were there prior to enactment of HIE laws. In Table 5 below, I find insignificant impacts for all of the HIE laws, and specifically for *ProHIE and Consent* and *ProHIE and No Consent* states. In terms of possible endogeneity, I am also concerned that the passage of these laws may be correlated with state unobservables that change during the time-period of my analysis and also impact HIE outcomes. For example, the passage of HIE legislation may be correlated with changes in political attitudes or public opinion towards the importance of Health Information Technology which is likely to also have an impact on HIE emergence and success. To address this concern, I estimated the model using only the subset of states that have passed HIE-promoting legislation (4 and 8 in Table 3 above). I see results consistent with those generated from the full model with a sizable and significant ($P < .05$) impact of *ProHIE and Consent* on *AttemptedHIE*. While directionally consistent with my prior results, the coefficient on operational exchanges is no longer significant ($P = .147$).

[Table 5 Model Estimation Using Lead Variables]

	One Year Lead		Two Year Lead	
	AttemptedHIE		OperationalHIE	
	(1)	(2)	(3)	(4)
ProHIE and Consent	0.411	0.073	0.052	-0.087
	(0.75)	(0.10)	(0.26)	(0.41)
ProHIE and No Consent	-0.615	-0.504	-0.217	-0.229
	(1.43)	(1.14)	(0.72)	(0.86)
ProHIE Only	-0.736	-0.653	-0.277	-0.126
	(1.76)	(1.68)	(1.18)	(0.52)
Disclosure Only	-3.032	-2.854	-1.084	-0.848
	(1.96)	(1.86)	(1.75)	(1.36)
State Designated	-1.245	-1.038	0.110	0.199
	(3.15)**	(2.69)**	(0.35)	(0.62)
Funding	0.343	0.315	0.099	0.090
	(0.76)	(0.67)	(0.38)	(0.36)
Population	-2.071	-1.422	-1.757	-1.574
	(1.77)*	(1.37)	(1.56)	(1.47)
Population Squared	.000066	0.000051	.000068	0.000064
	(1.83)*	(1.46)	(1.59)	(1.54)
Per Capita GDP	-5.188	-7.300	-8.144	-8.678
	(0.77)	(1.04)	(2.45)*	(2.66)*
Established HIE	-0.212	-0.185	0.459	0.484
	(0.66)	(0.57)	(1.72)	(1.85)*
CPOE	2.406	2.233	1.087	1.035
	(1.61)	(1.56)	(1.72)*	(1.75)*
EMR	-0.406	-0.487	-0.268	-0.263
	(0.47)	(0.56)	(0.65)	(0.59)
Constant	3.418	3.376	1.731	1.687
	(4.19)**	(3.99)**	(3.22)**	(3.09)**
Observations	612	612	612	612
Number of ID	51	51	51	51
State Fixed Effects	YES	YES	YES	YES
Time Fixed Effects	YES	YES	YES	YES
R-squared	0.49	0.47	0.42	0.40
Robust t statistics in parentheses * significant at 10%; * significant at 5%; ** significant at 1%				

An additional concern is potential differences across states along dimensions that are correlated with the enactment HIE legislation and may also impact HIE progress. I examine this claim in Table 6, comparing a number of state dimensions across states with ProHIE and Consent, ProHIE and No Consent, and No HIE laws and find no significant differences across states. Although the differences are not significant, states with consent requirements did trend higher on some characteristics, such as the propensity to be a

democratic state or have a general health privacy law. These characteristics are fairly static across time and are captured by the state fixed effects element of the model.

[Table 6: Evaluation of Endogeneity of HIE Legislation (No significant differences across groups)]

State Characteristics					
Variable	Description	Cross-Section (2010)			
		Consent	No Consent	No HIE	Source
Failed Exchanges	Number of failed exchanges as of 2010 (normalized by population)	2.49	2.92	3.71	HIE Survey / eHealth Initiative Survey
Young Exchanges	Percent of exchanges pursuing exchange for less than 1.5 Years	.12	.13	.15	HIE Survey
General Health Privacy Law	Percent of states with Law that govern general health disclosure	.75	.45	.4	Georgetown Privacy Project
Population	Number of inhabitants in a state	5.2032	6.77	6.658	U.S. Census
Broadband Pen	Broadband penetration in a state	54.28	51.27	48.81	U.S. Census
Democratic	Indicator if a state generally votes democrat (last five pres. elections)	.75	.54	.32	U.S. Census
Per Capita GDP	Per capita state gross domestic product	44.9	46.3	39.1	Bureau of Economic Analysis
Advanced Degree	Percent of population with advanced degrees	11.42	10.59	9.0	U.S. Census
PopulationOver65	Percent of population over 65	12.23	12.32	12.76	Area Resource File
Managed Care Penetration	Percent of population enrolled in Managed Care	23.5	18.24	19.639	Area Resource File
CPOE Adoption	Hospital adoption of computerized provider order entry systems	.375	.4104	.3865	HADB
EMR Adoption	Hospital adoption of Electronic Medical Record systems	.75	.70	.667	HADB

Top 20 Medical School	Indicator if state has top twenty medical school	.63	.45	.2	U.S. News
HIE Characteristics (Operational and Planning)					
YearsPursuing	Years an HIE has been pursuing exchange efforts	3.70	4.75	3.917	HIE Survey
YearsOperational	Years an HIE has been sharing health data	2.48	4.517	2.482	HIE Survey
IndependentOrg	Percent of exchanges functioning as independent organizations	.322	.47	.38	HIE Survey
FormalGov	Percent of exchanges with formal governance structure	.833	.882	.833	HIE Survey
In-Kind Resources	Reliance on Time-In-Kind Resources	1.333	1.548	1.486	HIE Survey
OneTimeContrPay	Reliance on payer one time payments	1.81	1.733	1.68	HIE Survey
RecurringFeePay	Reliance on payer recurring fees	2	1.67	1.76	HIE Survey
HighGov	Percent of HIEs indicating heavy reliance on government funding	.692	.54	.5	HIE Survey
Sustainable	Percent of operational exchanges that are financially sustainable	.315	.47	.28	HIE Survey

I was also concerned that states with more *AttemptedHIE_{st}* simply attempted more exchanges, and thus also have more failed exchanges. While I do not have longitudinal data on failed exchanges, using a 2010 cross-section I found no significant difference in failed exchanges between states with consent requirements and those without (reported in Table 6). I do not have these concerns for the longitudinal measure of *OperationalHIE_{st}* because, in my time period of analysis, no HIEs became operational and then subsequently failed.

Lastly, there may be concern that actual exchanges in states with consent requirements may differ in relevant ways from the exchanges in states without these laws. For example, states with consent requirements have more non-failed HIE attempts and operational exchanges because of the organizational positioning or structure of the HIEs themselves, the funding sources they receive, or some other characteristics of the exchange themselves (as opposed to their operating environment). I looked at a

number of HIE characteristics and found no statistically significant difference between states across a number of HIE dimensions, including participation of private payers, having a formal governance structure, and reliance on various types of financial support.

7. Limitations

The dependent variables presented in this work may not cover the full breadth of potential measures of success for HIEs. While reaching operational status is a significant milestone for HIEs, prior research on HIEs has also noted that sharing by HIEs has been limited in breadth and scope (Adler-Milstein et al 2009). Future work may evaluate, in more substantive terms, aspects of HIE sharing rather than just whether they are sharing or not. Moreover, I may consider that a higher number of exchanges in a state may not necessarily be a positive outcome. For example, it may be the case that a better outcome is to have only one exchange that facilitates exchange for all providers in the state. In fact, the move towards state and national level exchange efforts suggests that indeed this is a desirable end result. However, the current national strategy for health information exchange involves spurring small regional efforts and then linking them as building blocks of a state and national exchange (Vest and Gamm 2010). I can thus consider (as have other prior work) that a higher number of successfully attempted and operational exchanges in a state as a positive indicator of HIE progress.

8. Discussion and Conclusions

I evaluated the impact of patient consent requirements on the emergence and success of HIEs using a fixed effects model over a span of six years. I find that only states adopting regulatory approaches that provide strong privacy protections to patients through consent requirements, alongside incentives for HIE adoption experienced a higher number of HIE attempts and actual operational exchanges. I also find that exchanges in states with these requirements encountered lower levels of concern due to patient privacy issues. The inclusion of state and time fixed effects controls for relevant observables, and I find no evidence of endogeneity. I believe that the most plausible explanation for the results is that increased

assurances provided by stronger privacy protections may bolster trust, promote adoption, resulting in fewer challenges from patient privacy concerns. Alternatively, the enactment of strong privacy requirements may reduce ambiguity about the applicability of current health disclosure requirements or uncertainty regarding future privacy requirements in the context of exchange thus encouraging HIE adoption.

These results are of interest for a variety of reasons. Given that HIEs are an innovative healthcare technology with the potential to alleviate two of the most pressing concerns of the current healthcare system; rising costs and inconsistent quality, the results of this chapter may inform current and future efforts to incentivize HIE growth while balancing patient privacy concerns. The results in this chapter may also provide more general insights into promoting other innovative technologies that promise significant public benefits but are inherently privacy sensitive. More specifically, it suggests that weak or lacking privacy protections in the context of privacy sensitive technologies may dampen the effectiveness of incentives for adoption. These results do not necessarily contradict previous work suggesting that privacy regulation in some context can result in undesirable outcomes (e.g. lower levels of EMR adoption). For example, privacy concerns in the context of exchange may be significantly more salient, thus the dynamics that impact the pursuit and success of HIEs may be quite different from the adoption of electronic medical records, and as such may have dissimilar interactions with health disclosure laws.

These results also inform the broader debate on the role consumer consent as a privacy protective mechanism, as a number of emerging technology efforts (including behavioral advertising or location sharing) share similar characteristics as HIEs (i.e. focus on information sharing of personal information). Namely, numerous government and corporate entities in the United States have advocated self-regulatory “choice and consent” models of privacy protection that, essentially, rely on users’ awareness and control. However, recent research demonstrates that making individuals feel more in control over the release of personal information may carry the unintended consequence of eliciting riskier disclosures (Brandimarte,

Acquisti, and Loewenstein 2010). As a result, it is not clear that patients in states with consent requirements actually experience greater privacy protections. While I do not evaluate this question specifically, my results actually suggest that states with consent requirements actually have more patient data being shared via HIEs.

The results presented in this chapter have implications for healthcare providers, as well as for policy makers at both the state and federal level. Often, technological progress and privacy protection sit on opposite ends of the table negotiating terms seeking to balance the two. These concerns may be increasingly salient in the case of HIEs given their direct privacy implications, and the considerable attention that has been given to various privacy and security concerns. The results in this chapter suggest that at least in the context of HIEs, a potentially transformative healthcare technology designed to enhance efficiency and quality of care, stronger protections seem to go together with incentives for the development and success of these efforts. More generally, this chapter provides some evidence that it may be a balanced combination of carrot and substantive stick that works most effectively to promote innovation and the adoption of privacy-sensitive technologies.

III. Chapter 3: A Sleight of Privacy: The Limits of Transparency

1. Introduction

In response to persistent consumer privacy concerns and high profile privacy incidents (e.g. see Krazit, 2010), US policy makers have primarily resorted to two strategies. One strategy has consisted of imposing fines on organizations that used consumer data in manners deemed invasive. The other strategy has consisted of self-regulatory efforts to increase transparency about firms' data handling practices (for instance, through simple, accessible privacy policies or notices), as well as to increase consumer control over their personal information (FTC, 2012). As of recent, such "transparency and control" solutions (or choice and notification regimes, as they are also called) seem to have become the object of a surprisingly broad consensus between policy makers, industry, and privacy advocates. Both the FTC white paper on consumer privacy and the White House Consumer Bill of Rights (FTC, 2012; The White House, 2012) presented transparency and notice as central tenants to consumer privacy protection. Industry leaders, such as Facebook and Google, broadly concurred with the approaches outlined by policy makers. In comments on the FTC privacy framework, Facebook stated that "...companies should provide a combination of greater transparency and meaningful choice..." for consumers, and Google stated that making the "collection of personal information transparent" and giving "users meaningful choices to protect their privacy" are two of their guiding privacy principles (Santalesa, 2011). Privacy advocates have also embraced these approaches (Reitman, 2012). While researchers have highlighted the limitations of *current* privacy policies and notices (Jensen and Potts, 2004; McDonald and Cranor, 2009), the general expectation seems to be that some *new and better* future iteration of privacy notices will solve consumers' make privacy decision making issues. This chapter presents experimental evidence that this approach, alone, may not be sufficient.

In principle, notice can certainly improve consumer disclosure decisions, while avoiding potentially burdensome regulation of firms. In normative terms, giving individuals more information about, how their personal data is used seems an unarguable improvement over a situation in which consumers are left in the dark. In particular, policy makers posit that improved transparency will counter the status quo in which privacy concerns are secondary in online decision making, and most consumers do not read overly complex and lengthy privacy notices. Unfortunately, the ability of even improved transparency solutions or additional control tools to better align consumer attitudes towards privacy with actual behavior and reduce regret from over sharing is ultimately questionable. This chapter investigates the hypothesis that even simple, straightforward, and easily accessible privacy notices may not always be effective aids to privacy and disclosure decisions. Specifically, I argue that well documented and systematic biases or limitations in decision making (such as relative judgments and bounded attention) can hinder the propensity of privacy notices to achieve the desired effect of supporting consumers in navigating disclosure related choices.

In a series of experiments, I find that while simple privacy notices communicating lower privacy protection can, under some conditions, result in less disclosure from participants (in line with the policy aims for increased transparency), simple and common changes in the framing of those same notices, that exploit individual heuristics and biases, can result in the effect of even straightforward and accessible privacy notices being predictably manipulated (Experiment 1) or entirely thwarted (Experiment 2). In Experiment 1, I demonstrate that the impact of privacy notices on disclosure is sensitive to whether notices are framed as increasing or decreasing in protection, even when the objective risks of disclosure stay constant. Of particular interest is that users may effectively be led to disclose more than the level justified by objective privacy protection, and therefore face higher objective risks, if online providers put a strong emphasis on increases in privacy protection. Also, I find evidence of a diminishing propensity of privacy notices to impact disclosure over time, suggesting that notice may have an initial impact but that users may settle back into familiar disclosure habits in a short period of time. In Experiment 2, I demonstrate

that the propensity of privacy notices to impact disclosure can be muted by a number of simple and minimal misdirections (such as a mere 15 second delay between notices and disclosure decisions) that do not alter the objective risk of disclosure. I argue that the sort of manipulations captured by the experimental design mimic (if anything, conservatively so) the sort of hurdles that consumers face when making real privacy decisions online. It follows that privacy notices can – on the one hand – be easily marginalized to no longer impact disclosure, or – on the other hand – be used to influence consumers to share varying amounts of personal information. Transparency may, therefore, become a “sleight” of privacy.

Such findings cast doubts on the ability of policy initiatives and design solutions built around transparency to, alone, address consumer privacy concerns. Note that the main implication of the results presented in this chapter is not that notice should be avoided, or is entirely ineffective. In fact, notice may be a necessary condition for meaningful privacy protection. Instead, my results suggest that, disjointed from the rest of the OECD privacy principles (OECD, 1980) of which they were originally part (such as purpose specification, use limitation, and accountability), transparency may not be *sufficient* conditions for privacy protection. Worse, they may reduce to a case of “responsibilization” – a situation where individuals are “rendered responsible for a task which previously would have been the duty of another [...] or would not have been recognized as a responsibility at all” (Wakefield and Fleming, 2009).

2. Theoretical Background and Hypotheses

Extant privacy research (Tsai et al, 2011) has highlighted that hurdles and inconsistencies in privacy decision making may be due, at least in part, to problems of asymmetric information: consumers who face privacy sensitive decisions may often be unaware of how their data is collected and used, and with what consequences. This challenge has been primarily attributed to privacy policies ineffectively communicate privacy risks to consumers. For instance, prior work has found that many privacy policies are not

readable, with many policies beyond the grasp of the average internet user (Jensen and Potts, 2004) and that privacy policies may be excessively costly to navigate (McDonald and Cranor, 2009). To address this issue, researchers have attempted to improve the readability and “usability” of privacy policies. For example, Kelley et al. (2009) developed a “nutrition-label” style presentation of privacy policies that outperformed standard formats in readability, recall, and comprehension (Kelley et al, 2009).

While evidence suggests that improved privacy notices can better inform consumers about the way their data is used, their ability to actually engage in “better” privacy decision making (that is, decisions that the consumer is less likely to later regret, or that better reflect stated preferences) is unclear. Under rational accounts of privacy decision making (Stigler, 1980; Posner, 1981), predicated on the implicit premise that people can estimate stable trade-offs between privacy and other concerns, increasing the availability and comprehensibility of information should result in some increased consistency in privacy decision making. However, substantial literature in behavioral economics and decision research documents systematic inconsistencies in individuals’ choices. That research shows that choice is sensitive to how choice alternatives are framed, and the salience of available information to consumers. For example, Kahneman and Tversky (1979) find that individuals are much more likely to accept a gamble when the choice is framed as avoiding a loss compared to when the objectively equivalent choice is framed as obtaining a gain. Moreover, Kahneman, Knetsch and Thaler (1990) find significant differences in the amount individuals are willing to pay for an item compared to individuals’ willingness to accept for the same item. Given that privacy’s tangible and intangible consequences are often difficult to estimate, numerous heuristics and biases can influence and distort the way individuals value data protection and act on privacy concerns (Acquisti, 2004, 2009). A growing body of empirical research has started highlighting the role of such systematic inconsistencies in privacy decision making. In a similar manner, heuristics may affect how consumers read, and react to, privacy notices. In this chapter, building on the existing body of behavioral and decision research, I use two experiments to evaluate the impact of framing and bounded rationality on the propensity of privacy notices to impact disclosure.

2.1 Framing and Reference Dependence: Experiment 1

Research has highlighted that privacy concerns, and therefore propensity to disclose, are sensitive to relative judgments, which could be explained by “herding effects” (individuals being more willing to divulge sensitive information when told that others had also made sensitive disclosures (Acquisti, John, and Loewenstein, 2012); or by reference dependence, a concept introduced by Kahneman and Tversky in 1979 in which they posited that outcomes are not only evaluated on their absolute value but also on their deviation from a reference point.

Framing, relative judgments, and reference dependence may also impact how individuals react to privacy notices. I argue (and test in a first experiment) that reference dependence may have a significant role in privacy decision making, since a space where consumers now make a considerable amount of privacy decisions – the online marketplace – consists of a constantly changing array of disclosure policies and privacy risks. For example, Facebook privacy settings have undergone several changes which have been presented to consumers as being increasingly protective of their privacy. Thus, some consumers may perceive their privacy protection on Facebook as improving over time. Conversely, consumers that view Facebook’s changes to default settings as less privacy protective, or encounter articles identifying Facebook’s various privacy infractions, may perceive their privacy protection as decreasing over time. Under rational accounts of privacy decision making, if consumers are concerned about their personal data, privacy notices that offer low protection should elicit, on average, lower levels of disclosure relative to notices that offer sufficiently higher protection. Moreover, identical privacy notices should result, on average, in comparable levels of disclosure irrespective of relative changes in privacy notices (i.e. whether they have been increasing or decreasing over time in their level of protection). However, under an alternative account of decision making that incorporates reference dependence, consumers would evaluate privacy notices relative to their deviation from a reference point, such as the level of protection they had in the recent past or they currently have (i.e. status quo). More specifically, consumers presented

privacy notices that are framed as increasing in protection (i.e. preceded by notices that are less protective) would disclose more relative to those that experience no change in privacy protection, and the converse for those presented notices that are framed as decreasing in protection. As such, I posited the following hypotheses, which I test in Experiment 1:

H1a: The framing of privacy notices as increasing in their protection against privacy risks will result in an increased level of disclosure relative to no change in privacy notices.

H1b: The framing of privacy notices as decreasing in their protection against privacy risks will result in a decreased level of disclosure relative to no change in privacy notices.

Kahneman and Tversky also suggested that individuals are loss averse in that they perceive a greater dissatisfaction from losses as compared to the satisfaction from equivalent gains. Hence, I posited the following additional hypothesis:

H1c: changes in disclosure will be greater in magnitude for decreasing protection relative to increasing protection.

2.2 Bounded Rationality and Salience: Experiment 2

In a second experiment, I consider the impact on disclosure of bounded attention and privacy notices salience. Prior studies have demonstrated that attention is a limited resource and that the salience of stimuli can moderate their impact on behavior (Broadbent, 1958). Economists have also proposed that bounded attention may be a contributing factor to sub-optimal consumer decision making. Simon (1955) suggested that individuals may simplify complex decisions by focusing on a subset of the information provided, and DellaVigna (2007) has suggested that the propensity of costs to impact decisions is moderated by the degree of inattention by consumers. Similarly, Hossain and Morgan (2006) have found

that, holding total cost constant, eBay auctions with lower initial prices (accessible cost) and high shipping costs (opaque cost) price significantly higher than the converse. The authors argue that this difference is driven by the increased salience of product price as a cost relative to shipping. I extend this prior work to the context of an online disclosure experience, during which consumers are often multi-tasking and focusing on many different stimuli at once. I argue that privacy notices and the considerations they elicit from consumers are often disjoint from actual disclosure decisions via various “misdirections” – that is, actions or states that do not alter objective privacy risks but may distract consumers from them.

As a baseline, I initially consider the case in which privacy notices immediately precede disclosure decisions and no such misdirection is present. Given that privacy notices will be salient at the point of disclosure, I argue that notices will have an impact on disclosure. For instance, privacy notices communicating stronger privacy protection may result in higher levels of disclosure relative to privacy notices communicating weaker privacy protection. As such, I posited the following hypothesis, which I test in Experiment 2:

H2a: Absent a misdirection, presenting privacy notices immediately preceding disclosure decisions will have an impact on disclosure behavior, with notices presenting low protection resulting in lower levels of disclosure, on average, relative to notices offering higher protection.

I then consider misdirections that have no relevance to the risks communicated in privacy notices, but simply have the propensity to distract consumers from them. For example, consider a brief delay between the presentation of privacy notices and disclosure decisions: it may allow consumers’ attention to drift away from privacy notices to other items – such as other websites or their email accounts. This distracted state may lead to a diminished impact on disclosure of privacy notices communicating privacy risks to consumers. However, I note that the directional changes in disclosure in the presence of a misdirection may be ambiguous and likely dependent on the misdirection itself. For example, a misdirection that

positively impacts goodwill (e.g. reading an article on charitable organizations) may result in all affected consumers disclosing at some commonly high level, irrespective of risks communicated in privacy notices. Conversely, a misdirection that negatively impacts trust (e.g. reading an article on phishing emails) may result in all affected consumers disclosing at some commonly low level despite privacy notices. The common feature, however, is that privacy notices are no longer the main factor driving disclosure and the differences they elicit absent a misdirection should be diminished. As a result, I posited the following hypothesis:

H2b: Introducing a privacy irrelevant misdirection following the presentation of privacy notices and before disclosure decisions diminishes the propensity of privacy notices to impact disclosure behavior.

Finally, I consider privacy “relevant” misdirections. These are misdirections that relate to the privacy risks communicated in the privacy notices but only focus consumers’ attention on a subset of risks. In effect, they do not alter objective privacy risk but may potentially distract from some dimensions of risk communicated in the privacy notice. An example of this type of misdirection is commonly found on online social networks, when consumers are provided granular notice and control over some dimensions of sharing and privacy preferences (e.g. access to one’s personal information by other consumers of the service), but fairly minimal and less salient notice and controls (if any) over the collection and use of personal information by the service providers (e.g. Google+ or Facebook). I posit that this may result in disproportionate focus on the privacy risks from other consumers of a service and lessened focus on the providers of these services. As a result, I posited the following hypothesis:

H2c: Introducing a privacy relevant misdirection focusing on a subset of privacy risks communicated in a notice diminishes the impact on disclosure of other dimensions of risk communicated in the notice.

Following an approach that is pervasive in the experimental literature on privacy and disclosure (Joinson, Woodley, and Reips, 2007; Phelps, Nowak, and Ferrell, 2000; Weisband and Kiesler, 1996), I tested these hypotheses using two survey-based experiments with random assignment, employing as main dependent variable the propensity of participants to answer personal questions in the surveys as a proxy for privacy concerns (Frey, 1986; Singer, Hippler, and Schwarz, 1992).

3. Experiment 1


Experiment 1 was a four condition between-subjects design in which I manipulated changes in privacy notices as increasing or decreasing in protection, and examined the effect of such changes on disclosure relative to conditions in which privacy notices did not change.

3.1 Procedure

Participants were recruited through Mechanical Turk, an online service that connects researchers with potential participants and is becoming increasingly popular among social scientists conducting online experiments. Participants were invited to take two online studies on ethical behavior each of which paid \$.20. At the end of the first study (Survey 1), they were asked to confirm that they wished to continue to the second (ostensibly unrelated but consecutive) study (Survey 2) for an additional \$.20 (all participants chose to continue to the second study but three were prohibited from completing the study because they failed the attention check). In Survey 1, participants were first asked demographic questions, which included email as a required field. Then, they were provided with a simple (i.e. brief, using mundane language, and dealing only with anonymity of responses) privacy notice about the way their answers to the questionnaire would be used. Finally, participants were presented with six questions related to ethically questionable activities (See Appendix B). If they decided to take it, participants would then start Survey 2, which followed the same structure as Survey 1 (see Figure 1 for the flow of the experiment) but had a different aesthetic design to help convince participants they were participating in a separate study (see Appendix C). In exit questions, participants confirmed that they felt they had participated in

two separate studies and that their responses from the first study could not be linked to their responses from the second study. Participants were again asked for their emails and demographic information; then, they were provided a privacy notice about the way their answers to the ensuing questions would be used; finally, they were presented with six new questions about other ethically questionable behaviors (See Appendix B).

Figure 1. Flow of Experiment One

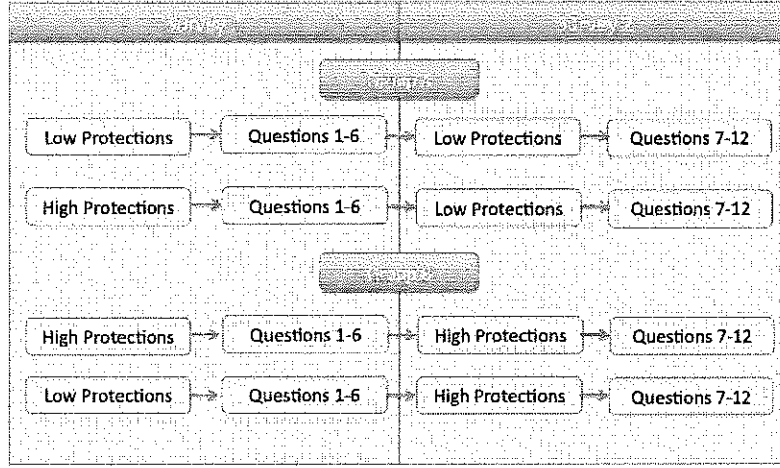


The questions used in both studies were the ones that were rated most intrusive in a 2012 paper by Acquisti, John and Loewenstein (Acquisti, John, and Loewenstein, 2012). The number and type of questions were kept constant across conditions, but the order of questions was randomized within each survey.

3.2 Design

The design was a 2 (high vs. low protection in the first survey) X 2 (high vs. low protection in the second survey). Thus, this study consisted of four groups (randomly assigned) in which privacy either *increased* from the first to the second survey (low protection to high protection: LH), *decreased* (high protection to low protection: HL) or stayed the same (low to low protection: LL or high to high protection: HH). Figure 2 provides an overview of the experimental design. The different levels of protection depended on the degree of possible linkage between the participants' emails and their responses, and therefore different levels of protection on their responses. Specifically, participants offered "low" protection to their responses were informed in the privacy notice that their answers would be linked to their email accounts. Conversely, those offered "high" protection to their responses were informed in the privacy notice that their answers would not be linked to their email accounts (See Appendix A for text of notices).

Figure 2. Design of Experiment 1



This design allowed us to evaluate both the baseline impact of protection using responses in Survey 1 and how the impact of protection may change over time, as some participants were provided identical protection in Survey 2. However, and importantly, the key feature of this design is that in Survey 2 for both the decreasing and increasing protection conditions, participants actually faced identical privacy notices as their respective comparative conditions – thereby allowing us to evaluate the impact of the change in privacy notices on disclosure in Survey 2.

3.3 Analytical model

I used a panel random effects Probit estimation approach to evaluate the overall differences in the propensity to admit to unethical behavior across conditions. I estimated the following model:

$$Admit_{ij} = \beta_0 + \beta_1 * Treatment_i + \beta_2 * Survey1Sharing_i + \beta_3 * Intrusive_j + \beta_4 * Intrusive_j * Treatment_i + \beta_5 * Age_i + \beta_6 * Male_i + \beta_7 * Design1_i + u_{ij}$$

Admit_{ij} measures the propensity to disclose, with a value of 1 if the participant admitted to the behavior and 0 if she denied or skipped, $i = \{1, \dots, N \text{ participants per interaction set}\}$, and $j = \{1, \dots, 12 \text{ questions}\}$.

Treatment is a binary indicator of the presence of my treatment. For example, in the case of decreasing protection, 1 represents a participant that was assigned to a decrease in protection, while 0 represents

participants that were assigned to a condition of no change from Survey 1 to Survey 2. *Survey1Sharing* is a measure of participant sharing levels in Round 1. This was included to control for the possible impact of disclosing more in the first round. *Intrusive* is a binary measure of whether a question is highly intrusive or not. *Intrusive*Treatment* captures any interaction between my treatment and highly intrusive questions. Lastly, *Design1* controls for which survey aesthetic design participants viewed.

The model assumes serial correlation between observations within a panel unit. I allow for the correlation between responses from a single participant when I estimate the variance-covariance matrix of the coefficients, assuming constant correlation between any two answers by the same individual (Liang and Zeger, 1986).

3.4 Results

Four-hundred and thirty-six participants ($M_{\text{Age}} = 30$, $SD=13.5$; $M_{\text{Female}} = .43$, $SD=.49$) completed the study (Survey 1 and 2) There were no significant differences in age or gender across conditions.

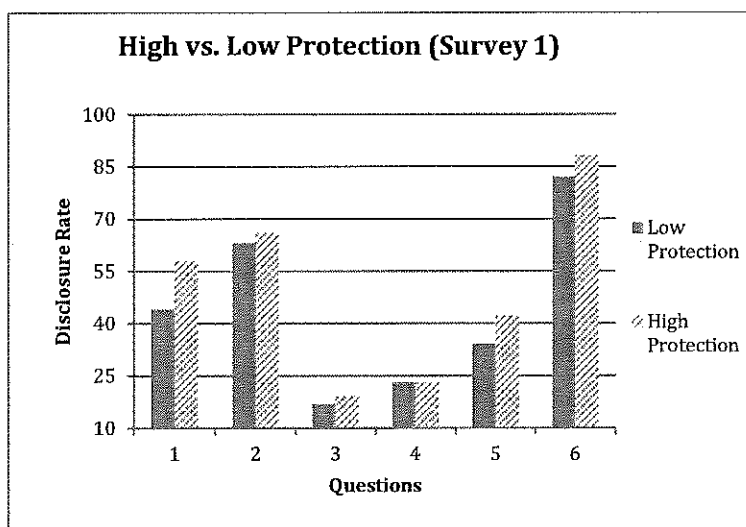
All respondents were presented an attention check question similar to those in a 2009 paper by Oppenheimer, Meyvis, and Davidenko to ensure participants were carefully reading directions (See Appendix C). Lastly, participants responded to exit questions that gauged both their perception of whether privacy protections increased, decreased, or stayed the same (depending on the condition) and their recall of privacy notices in both surveys. A minority (12%) weren't able to accurately recall privacy notices and, thus, disagreed that protections had increased, decreased, or stayed the same. These participants were excluded from the study, leaving 386 usable survey responses.

3.4.1 High vs. Low Protection

I first evaluated the disclosure rates of participants in Survey 1, where participants were randomized into conditions in which they were either presented high or low protection. At this point in the experiment, no participants had been presented the central manipulation of either changing or constant levels of

protection from Survey 1 to Survey 2. Figure 3 shows that participants in Survey 1 were more likely to disclose for 5 of the 6 questions when they were provided high protection ($p < .05$). Normalizing for base rates of disclosure between questions, this translates to a 14% average increase in the propensity to disclose when participants were afforded high protections in Survey 1, with some questions exhibiting more than a 30% increase in the propensity to disclose (see Table 1, Column 1 for the estimated coefficient on *Treatment*).

Figure 3: Differences in Survey 1 Disclosure



Next, I evaluated the impact of high protection relative to low protection when presented in Survey 2. Specifically, I compare participants that had high protection in both surveys to participants that had low protection in both surveys. Figure 4 shows that the impact of high protection does not extend to Survey 2 (see also Table 1, Column 2 for the estimated coefficient on *Treatment*, which is not significant), suggesting potentially some habituation to privacy protection and that users fall into some default mode of disclosure over time. Disclosure in Survey 2 was not systematically impacted by having high protection with only 2 of the 6 questions, demonstrating an increase in the propensity to disclose.

Figure 4: Differences in Survey 2 Disclosure

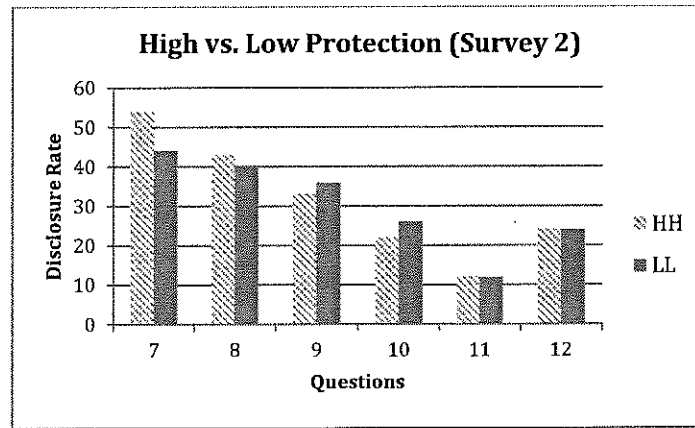


Table 1: Regression Results - High vs. Low Protection

	(1)	(2)
	Admit	Admit
Treatment (High Protection)	0.044	-0.005
	(0.023)*	(0.029)
Intrusive	0.041	-0.111
	(0.019)**	(0.026)***
Age	-0.003	0.002
	(0.001)***	(0.001)**
Male	0.008	0.019
	(0.020)	(0.024)
Survey1Sharing		0.108
		(0.010)***
Design1	0.051	-0.004
	(0.023)**	(0.029)
Constant	-0.029	-0.003
	(0.27)	(0.062)
Observations	2634	1146
** significant at 5%; *** significant at 1%		

3.4.2 Changes in Protection

Next, I evaluated the impact of changes in protection on disclosure. Figures 5 and 6 show relative disclosure rates for each question in Survey 2 of the experiment. Figure 5 displays a trend of higher propensity to disclose (4 of the 6 questions) when participants were presented increasing protection

relative to no change. Conversely, Figure 6 displays a trend of a lower propensity to disclose (4 of the 6 questions) when participants were presented decreasing protection relative to no change. In both cases, differences for questions that did not exhibit the trend were not significant.

Figure 5. Response Rates for Increasing Protection

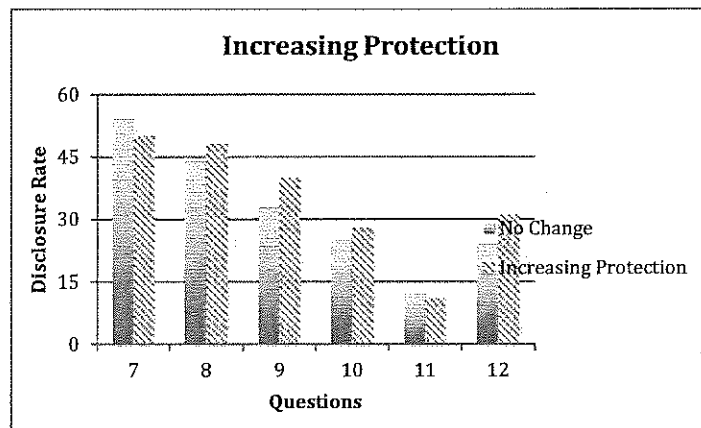
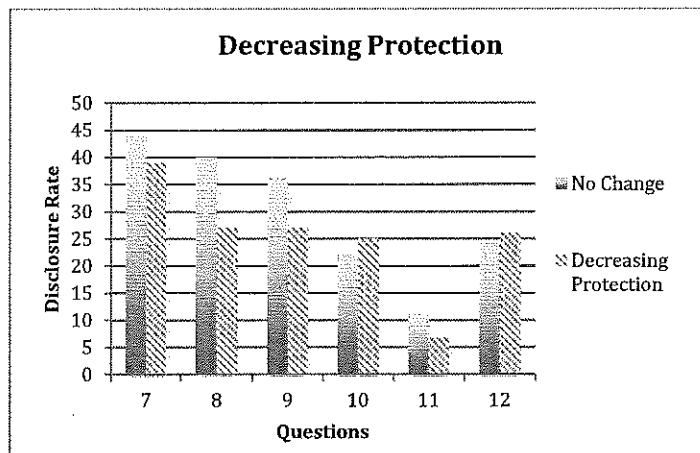


Figure 6. Response Rates for Decreasing Protection



However, base rates of disclosure for questions varied, so I also considered relative differences in the propensity to admit to a particular behavior. I found that, in the increasing protections conditions, participants were, on average, 10% more likely to disclose, with some questions having as high as 30% increase in the propensity to disclose. Similarly, for decreasing protections conditions, I found that participants were, on average, 14% less likely to disclose, with some questions having as high as a 40%

reduction in the propensity to disclose. Table 2 provides estimates of the model described in the prior section. I evaluated differences in disclosure in Survey 2, where participants were provided identical privacy notices, but those in the treatment condition were presented a decrease in protection from the prior round (HL), while those in the control condition (LL) were assigned to no change in protection. Also, I evaluated differences in disclosure in Survey 2 where, again, participants were presented identical privacy notices, but those in the treatment condition were presented an increase in protection from the prior round (LH), with those in the control condition (HH) being assigned to no change in protection. Specifications (1) and (3) estimate a baseline model with only the measure of Round 1 sharing and a dummy variable for the aesthetic design viewed by participants. Specifications (2) and (4) use the full specification described above, with additional controls for the intrusiveness of questions, age and gender. Marginal effects and associated robust standard errors are presented in Table 2.

Table 2: Regression Results – Changing Protection

	Decreasing Assurances		Increasing Assurances	
	(1)	(2)	(3)	(4)
	Admit	Admit	Admit	Admit
Treatment	-0.073**	-0.137***	0.069**	0.114**
	(0.029)	(0.047)	(0.033)	(0.051)
Round11Sharing	0.120***	0.123***	0.106***	0.112***
	(0.011)	(0.011)	(0.012)	(0.012)
Intrusive	--	-0.152***	--	-0.100**
	--	(0.044)	--	(0.043)
Intrusive *Treatment	--	0.115*	--	-0.063
	--	(0.062)	--	(0.058)
Age	--	0.002	--	0.004***
	--	(0.001)	--	(0.001)
Male	--	0.048	--	0.044
	--	(0.026)	--	(0.028)
Design1	-0.021	-0.026	-0.005	0.013
	(0.029)	(0.030)	(0.033)	(0.034)
Observations	1038	1038	1146	1146
** significant at 5%; *** significant at 1%				

I find that, in the basic specification, participants presented decreasing protection disclosed 7% less ($P < .05$) than participants that were presented no change in privacy notices, supporting H1a. Moreover,

participants that were presented increasing privacy protection shared 7% more ($P < .05$) than participants that were presented no change in privacy notices, in support of H1b. In the baseline specification, I did not find support for the loss aversion hypothesis (H1c). In the extended specification, where I teased out the impact of the treatment on the non-intrusive questions separate from intrusive questions, I found a larger baseline effect of the treatment, with a 14% ($P < .01$) decrease in disclosure for participants presented with decreasing protection and a 11% increase in disclosure ($P < .05$) for participants presented with increasing protection (directionally consistent with H1c, although the difference is again not significant).

To put these results in perspective, consider that, according to some sources, Facebook users posted 1.85 million status updates every 20 minutes, or approximately 49 trillion status updates in 2011.¹⁹ Other disclosures (uploading a photo, posting a comment, tagging another user) on Facebook happen at comparable rates. Moreover, Facebook generally advertises changes in privacy settings and practices, often to highlight improvement to user privacy protections.²⁰ Now: in my experiment, I find effects that range from 10% to 14% in terms of influencing disclosure; however, it may be the case that these results are only applicable to a subset of user disclosures (e.g. disclosures with sensitive information). If I take this into account, and assume that this effects apply to only 1% of all status updates on Facebook, a 10% increase in these disclosures translates to an increase of 49 million status updates in 2011. Moreover, if the effects I identified are most applicable to sensitive information, these may in fact be the subset of disclosures most concerning for consumers.

3.5 Limitations

The results in this study rely on a self-selected sample of individuals from the online service used to solicit participants in my survey, as I mandated that users provide their email as a condition of

¹⁹ See summary of Facebook usage statistics. (<http://www.onlineschools.org/visual-academy/facebook-obsession>).

²⁰ See, for instance, Facebook announcement of new privacy options in 2008 (<http://blog.facebook.com/blog.php?post=11519877130>);

Zuckerberg's 2009 open letter about Facebook eliminating "networks" (<http://blog.facebook.com/blog.php?post=190423927130>); and his 2010 announcement of further privacy changes (<http://blog.facebook.com/blog.php?post=391922327130>).

participation in the study. This requirement may have likely excluded potential participants with high sensitivity towards the sharing of their emails. However, those individuals may have reacted even more drastically to changes in privacy notices, and their exclusion may in fact bias the effects I identify downward. Moreover, while anonymity of responses is strongly related to the level of protection of participant responses and anonymization of sensitive data is a common mechanism for ensuring personal privacy, moving from anonymous to identified responses may also involve other changes, besides changes in privacy notices (e.g. impact of disclosures on self-perception of ethicality). I plan to run similar studies looking at other variants of privacy notice that include other privacy dimensions (e.g. the length of retention of responses and breadth of access to responses).

4. Experiment 2

For Experiment 2, participants were invited to create a profile (via an online survey) on a new networking service exclusive to their university. Participants were debriefed after the study and told they were a participant in a research study, so no online social network would be created. Exit questions confirmed that participants did not question the validity of the experimental context.

Experiment 2 was a 2 (access) X 5 (misdirections) mixed design where I manipulated, between subjects, the breadth of access to information disclosures communicated via a privacy notice (profile accessible only to students for the Students Only conditions, or to students and faculty for the Students and Faculty conditions) and the presence of a misdirection between the presentation of privacy notices and disclosure decisions (No misdirection or one of four misdirections: delay, department information pages, student committee, student committee and choice). In the context of this experiment, misdirections are actions or states that do not alter objective privacy risks but may distract consumers from them. Within subjects, I manipulated the intrusiveness of questions: I asked a total of 37 questions on demographics, housing, academics, and social activities, among which were nine questions that, ex ante, I considered disproportionately sensitive to disclose to faculty relative to students. These questions asked students for

their opinions on faculty, departments, and courses (e.g., “Who was your least favorite faculty member?”), dealt with the witnessing and reporting of cheating, and student effort put into academics. The full list of questions in Experiment 2 is proved in Appendix E.

4.1 Procedure

Participants were recruited at the student center of a major North American university and were compensated with a candy bar (approximate value \$1.00). Across all conditions, participants were initially presented with a privacy notice. Depending on the condition, they were either informed that their profile would only be accessible by the university students only (Students condition) or by both faculty and students (Students and Faculty condition, see Table 3). The text for each privacy notice can be found in Appendix D. Thereafter, participants in the no misdirection conditions proceeded immediately to disclosure decisions where they filled out various fields on their profile, while participants in the other conditions were presented with one of four different misdirections before proceeding to fill out the same profile fields.

Table 3. Overview of Conditions in Experiment 2

Notice	Control	Treatment	
Students	No Misdirection	Privacy Irrelevant (Delay, Dept Pages)	Privacy Relevant (Student Comm, Student Comm+Choice)
Students& Faculty	No Misdirection	Privacy Irrelevant (Delay, Dept Pages)	Privacy Relevant (Student Comm, Student Comm+Choice)

I first considered two misdirections without privacy relevance, in that they did not refer to the information provided in the notices nor did they deal with access to participant profiles. The first of these misdirections was a simple 15 second delay between the privacy notice and participant disclosure decisions. The second misdirection presented participants a page where they were asked if they wished to

sign up for departmental information pages. I next considered two additional misdirections which were privacy relevant. The third misdirection informed participants that a student planning committee would be using their profile in order to plan upcoming activities. The fourth misdirection utilized the same student planning committee context, but provided participants control over whether this committee may access their profile. I considered these treatments as privacy-relevant in that they refocus participants' attention on the student access to their profiles. I considered them misdirections as profiles were already accessible by students under all conditions. Table 3 provides an overview of the experimental design and Appendix F provides screenshots of each individual misdirection.

4.2 Analytical model

Similarly to the analysis in the prior experiment, I used a panel random effects Probit estimation. The results are summarized in Table 4. Again, this model assumes that responses from a single participant are serially correlated:

$$\begin{aligned} Admit_{ij} = & \beta_0 + \beta_1 * Student\&Faculty_i + \beta_2 * Academic_j + \beta_3 * Misdirection_i + \beta_4 * Student\&Faculty_i \\ & * Academic_j + \beta_5 * Student\&Faculty_i * Misdirection_i + \beta_6 * Academic_j * Treatment_i + \beta_7 * \\ & Student\&Faculty_i * Academic_j * Misdirection_i + \beta_8 * Identified_i + u_{ij} \end{aligned}$$

$Admit_{ij}$ measures the propensity to disclose, with a value of 1 if the participant answered the question and 0 if she denied or skipped, $i = \{1, \dots, N \text{ participants per interaction set}\}$, and $j = \{1, \dots, 37 \text{ questions}\}$.

$Student\&Faculty_i$ is a binary variable that indicates whether participants were either presented the Student and Faculty (1) or the Students Only privacy notice (0). $Academic_j$ is a binary variable indicating whether a question dealt with sensitive academic issues. $Misdirection_i$ is a binary variable that indicates whether the participant was presented a misdirection, and $Student\&Faculty_i * Misdirection_i$ captures the interaction between the Student and Faculty privacy notice and the misdirection. Lastly, I introduced a variable, $Identified_i$, which captures whether participants chose to identify themselves, and takes a value

of 1 if participants shared both their first and last names or they shared their email, and 0 otherwise. In contrast to the prior experiment, identifying information was optional in this study. I included this measure to adjust for any potential differences in participant propensity to identify themselves.

4.3 Results

Two hundred and eighty participants completed the experiment ($M_{Age} = 21.5$, $SD=3.1$; $M_{Female}=.37$, $SD=.48$), with about 26 to 30 participants per condition. Figure 7 presents disclosure rates for the control conditions (i.e. without a misdirection) and for all conditions with misdirections in aggregate (this pattern is consistent for each individual misdirection as well). In the control condition, participants presented with the “Student Only” notice were 26% more likely to disclose ($P<.05$) relative to participants presented the “Student and Faculty” notice. In the conditions with a misdirection, I see near zero and insignificant differences in disclosure between participants presented “Student Only” and “Student and Faculty” notices.

Figure 7. Control Conditions relative to Aggregated Misdirection Conditions




Table 4 below provides estimates of the above model. Estimates of the aggregate effect of all misdirections are found in column (1). Estimates distinguishing privacy-relevant and privacy-irrelevant misdirections are found in columns (2) and (3) respectively. For ease of interpretation, marginal effects and associated robust standard errors are presented, and less relevant covariates have been dropped in Table 4. For questions that were not sensitive in the context of academics I found insignificant, near 0

coefficients for *Student&Faculty* across all the specifications, indicating that the Student and Faculty privacy notice did not result in differences in disclosure relative to Students only for non-academic questions. However, I found a negative and significant coefficient ($P<.05$) on the interaction *Student & Faculty*Academic*, indicating that, absent a misdirection (i.e. in the control condition), the Student and Faculty privacy notice had a negative impact on disclosure for sensitive academic questions relative to Students Only (H2a supported). To evaluate differences in the disclosure level of sensitive academic questions for participants with the Student and Faculty privacy notice in the control conditions relative to treatment conditions I focus on the interaction *Student&Faculty* Misdirection* Academic*. Across all specifications, I found a positive coefficient on this interaction that offsets the negative impact of the Student and Faculty privacy notice on disclosure of sensitive academic questions. I found this when I considered all of the misdirections as one treatment ($P<.05$). I also found this effect ($P<.05$) when I aggregated privacy relevant misdirections (Column 2), in support of H2c, and privacy irrelevant misdirections (Column 3), in support of H2b.

Table 4. Experiment 2 – Regression Results

	All Misdirection	All Privacy Relevant	All Privacy Irrelevant
	(1)	(2)	(3)
	Admit	Admit	Admit
Student&Fac	-0.007	-0.007	-0.007
	(0.034)	(0.035)	(0.033)
Student&Fac* Academic	-0.109**	-0.110**	-0.111**
	(0.050)	(0.050)*	(0.050)
Student&Fac* Academic* Misdirection	0.097**	0.105**	0.092**
	(0.041)	(0.045)	(0.047)
Observations	10073	6112	5959
* significant at 10%; ** significant at 5%; *** significant at 1%			

4.4 Limitations

The results in Experiment 2 rely on the control condition not being a false positive. I also observe that, in the treatment conditions, disclosure rises disproportionately in the Student & Faculty condition relative to the Students Only. However, the current work does not evaluate underlying processes that may be driving the effect I observe. Initial analysis suggests that misdirections distracted participants from privacy concern and thus these concerns were less primed as they made disclosure decisions. Future work will evaluate this claim more directly via process-oriented studies.

5. Discussion and Conclusions

The findings I presented in this chapter provide evidence of two potential inconsistencies in the impact of privacy notices on disclosure. In my first experiment, I demonstrated that the impact of privacy notices is sensitive to reference dependence, with notices framed as increasing in protection eliciting increased disclosure and notices framed as decreasing in protection eliciting decreased disclosure. In my second experiment, I found that the downward impact of riskier privacy notices on disclosure can be muted or significantly reduced by a slight misdirection which does not alter the objective risk of disclosure.

These results have the most applicability to online services in which user disclosure is a central function (e.g. online social networks), but also have implications for technology settings that attempt to address consumer privacy through privacy notices (e.g. online retailers). For example, Experiment 1 mimics the evolutions of Facebook's and Google's notices with respect to presenting to consumer improvements in privacy protections; Experiment 2 mimics the delays that in real life separate the reading of a privacy notice and later privacy decisions.

Policy makers and firms that deal with the exchange of consumer personal information have advocated the increased readability and usability of privacy policies as improved privacy decision aides for consumers. While these measures may provide some incremental improvements in privacy decision making, inconsistencies in decision making may result in continued disparity in consumer concerns and disclosure behavior, potentially increasing regretful disclosures by users. My results suggest that current

policy and design approaches focusing just on transparency may be limited in their ability to improve consumer privacy decision making. The broad support for self-regulatory approaches focusing on making privacy transparent will likely make privacy notices simpler and more accessible, providing consumers certain benefits: attentive consumers concerned about their privacy may be able to better utilize short, simple, and well-formatted privacy notices to inform disclosure decision. However, the attentiveness of consumers to privacy issues may be sporadic and limited, inhibiting the usefulness of even simple and clear privacy notices. Even worse, attention paid towards self-regulatory approaches with dubious effectiveness may come at the cost of focusing on solutions that get at the heart of the privacy problem. In this regard, the experiments I presented in this chapter illustrate the need to expand the concept of transparency to not only include clarity and ease of comprehension, but also making information communicating privacy risks salient and readily available to consumers when they most require them, at the point of disclosure.

Finally, the findings presented in this chapter may have implications for firms that collect and use personal data, particularly those with consumer personal information at the core of their business models (e.g. online advertising). Firms are likely to have significant long-term ramifications from inadequately communicating information practices to consumers, particularly if failing to do so results in high profile misuses of consumer personal information or breaches of consumers' expectations of privacy. One ramification with major implications for these firms is that consumers' propensity to disclose information may significantly change over time.

IV. Chapter 4: Why Choice may not be Enough: Framing Effects and Malleable Preferences for Privacy

“The strongest and most effective force in guaranteeing the long-term maintenance of power is not violence in all the forms deployed by the dominant to control the dominated, but consent in all the forms in which the dominated acquiesce in their own domination.”

-Robert Frost

1. Introduction

Choice with regards to the collection, use, and disclosure of personal information has long been a central privacy construct (Westin 1967; Miller 1971) and has thus been pivotal in policy guidelines that aim to protect individual privacy (Ware, 1973; OECD, 1980). Policy makers posit that increased choice will empower consumers and allow them to manage their privacy in accordance with their individual preferences for privacy (FTC, 2012; The White House 2012). As a result, individual choice as a form of privacy protection has become pervasive in regulation intended to protect consumers' sensitive personal information. For example, many states require that health information related to sexually transmitted diseases or mental health not be released by a provider except with the permission of the patient. Financial institutions are also required to grant consumers the right to opt-out of certain collections and uses of their sensitive financial information. However, Choice in the context of privacy is not always provided to consumers as a result of regulation. Many providers of online service have also established and continue to update privacy settings which allow consumers to control, to some degree, the flow and access to their personal information. For example, Google provides a privacy dashboard which allows users to control some aspects of their information collection and disclosure; Facebook has provided some form of privacy settings since the early years of the service (Stutzman, Gross, and Acquisti, 2013). Finally, policy makers in large accordance with industry continue to advocate for increased consumer control as a privacy protection against emerging privacy risks. For example, policy makers and firms

broadly agree that providing control to consumers protection necessary to address consumer privacy concerns associated with behavioral advertising (FTC, 2012); consensus on the manner to implement consumer choice in the context of behavioral advertising, however, has been difficult to reach (Bott, 2012).

Similar to transparency, control is seemingly a certain improvement for consumers compared to a state where consumers are not provided any choice with regards to the disclosure and use of their personal information (effectively a 100% opt-in rate). However, the likelihood of choice solutions reducing consumer privacy risks may be questioned. Brandimarte, Acquisti, and Loewenstein found in a 2012 paper that, in practice, an increased feeling of control over the publication of personal data can paradoxically result in increased, and riskier, disclosures. Moreover, scholars argue that “organizations, as a rule, will have the sophistication and motivation to find ways to generate high opt in rates.” (Solove, 2013) and that “many data-processing institutions are likely to be good at obtaining consent on their terms...” Schwartz (2005). Moreover, scholars find that current implementations of privacy choice do not result significant changes in consumer behavior (Janger and Schwartz, 2001). However, scholars arguing that firms will be able to elicit high levels of compliance from consumers refer primarily to heavy-handed approaches for ensuring compliance such as requiring consent as a prerequisite to obtaining a resource or service (e.g. mobile application or online social network) or not providing consumers with sufficient or clear information about the risks associated with the choice. In this chapter, I consider the role of common but subtle differences in the presentation of privacy choices and how these differences may systematically influence consumer choice. Specifically, I conjecture and present evidence that variation in the presentation of choices can presentation of otherwise identical choices can alter the decision frame, or the “decision maker’s conception of acts, outcomes, and contingencies associated with a particular choice “(Tversky and Kahneman, 1981), in a manner that predictably influence choice of privacy protective options. Specifically, I find that that the labeling of privacy relevant choices as “Privacy Settings” results in the choice of more protective settings relative to labeling the identical choices as “Survey Settings”.

Moreover, I find that presenting individuals with an important privacy choice mixed with other settings that were pre-ranked by participants to be less relevant, significantly decreased the likelihood of participants to choose the protective option for the important choice. Finally, I find that participants presented privacy relevant choices as a choice to allow a use of their personal information (accept frame) were significantly *less* likely to choose the privacy protective option relative to those presented the identical choice as a choice to prohibit a use of their personal information (reject frame). Generally, I find that manipulations of framing do not have a significant impact on disclosure, when I control for the choice of settings by participants.

The results in this chapter primarily contribute to two streams of research. First, these results stand to contribute to the stream of research evaluating the intersection of behavioral economics and privacy decision makings. Specifically, this chapter explores the role of limitations in decision making on the propensity of individuals to choose privacy protective options in decisions that expose them to privacy risks. In addition, these results contribute to the framing the decision making and judgment literature on framing and priming effects. Specifically, I find evidence that various factors (e.g. choice sets and accept/reject manipulations) can exacerbate or diminish framing effects. Also, the results in this chapter introduce a number of important implications for policy makers and firms. Given that common and subtle manipulations of the presentation of choices to participants can systematically influence them to choose less restrictive settings for data uses they feel are intrusive, these results highlight the limitations of policy approaches that rely heavily or exclusively on control solutions to address consumers privacy concerns. Even more concerning is that while increased choice may not necessarily result in significant differences in consumer privacy decision making (thus not reducing objective risk), it may alleviate consumer privacy concerns which currently curb intrusive data practices by firms; a case in which a combination of increased choice and exploitation of limitations in individual decision making lead consumers to “acquiesce in their own domination”.

2. Background

I motivate the central arguments in this chapter using a broad empirical and theoretical literature evaluating various forms of framing on individual decision making (Levin, Schneider, and Gaeth, 1998; Kuhberger, 1998; Kahneman and Tversky, 1981). Classic framing studies focused on frames that differentially highlight the positive vs. the negative dimensions of a choice. The seminal demonstration of a framing effect, and one which is widely tested, is categorized by Levin et al. as a form of risky choice framing. It is referred to as the “Asian Disease Problem” and was first presented by Kahneman and Tversky (1979, 1981). In the Asian disease problem, all participants are given the choice of two interventions to a disease outbreak. One intervention was probabilistic with some chance of everyone being saved (dying) and some chance of saving no one (everyone). The other intervention was certain with a sure chance that some number would be saved (die). Between conditions, they manipulate whether the gains of the interventions are highlighted (i.e. how many individuals will be saved) or whether the losses from intervention are highlighted (i.e. how many will die as a result of various interventions). While the two interventions between conditions are objectively identical, they find that participants tend to choose the certain option when the problem is framed in terms of gains and were more likely to gamble (i.e. the probabilistic option) when the intervention is framed in terms of losses. Levin et al (1998) also identify a second category of framing they deem “attribute framing” in which certain dimensions of an object are presented using different frames. For example, a study by Levin and Gaeth (1988) finds that perception of the quality of ground beef differ based on whether it is labeled as “75% lean” or “25% fat” (Levin and Gaeth 1988). Finally, they identify a form of framing they term “goal framing” in which the goal of a context or choice is presented as either pursuing a gain or avoiding a loss. For example, numerous studies find that framing choices in terms of costs (e.g. consequences of not having a breast exam) is a more effective mechanism to influence behavior relative to framing which highlighting gains (Reese et al. 1997; Ganzach & Karsahi 1995).

Framing effects however, do not always influence behavior by highlighting negative or positive dimensions of choice. For example, Liberman, Samuels, and Ross (2004) find that a “Wall Street Game” vs. “Community Game” labeling of a prisoner dilemma game has a significant impact on cooperation. They suggest that this framing influences participants’ perception of the goal of the game. Epley, Caruso, and Bazerman (2006) replicate this result using “strategic competition game” and “cooperative alliance game” labels and again find differences in participant cooperative behavior. Another related stream motivating this chapter is the literature on priming which has also demonstrated some analogous results by manipulating labels to prime various mindsets or concepts that can significantly influence behavior. For example, Burnham et al. (2000) finds a strong impact on cooperation when labeling participants in a two-player reciprocity game as either a “partner” or an “opponent”. Galinsky et al. (2003) found that asking participants to recall instances in which they had power vs. when they were powerless (i.e. priming perceptions of power) has an impact on judgment of risks and optimism. Spencer, Steele, and Quinn (1999) find that priming gender differences in math performance can have a significant negative impact on the performance of qualified women on math evaluations; an effect they deem the “Stereotype Threat”. Jointly these streams of research motivate my manipulations and my predictions in the four subsequent studies.

3. Experiment 1

In a first study, I evaluate whether the labeling of various choices leads to a change in the choice frame and whether this impacts individual behavior. Study 1 is a two factor, between-subjects design in which I manipulate whether (1) choices presented to participants are framed as privacy choices and the (2) importance of choices presented to participants (ranked in a pre-test). I examined the effect of both factors on the propensity of individuals to select protective settings. I hypothesize that the framing of otherwise identical choices as privacy choices will result in more protective decision making by highlighting the privacy dimension of these choices to participants (H1a). Moreover, the literature suggests that magnitude of framing effects can depend on whether participants are engaged in the choice presented (Maheswaran

& Meyers-Levy 1990; Rothman et al. 1993; Krishnamurthya, Carterb, and Blairc 2001). As a result I also hypothesize that this framing effect will be less pronounced for the low importance settings (H1b).

Design

The design was a 2 (“Privacy Settings”, “Survey Settings”) X 2 (High Importance vs. Low Importance). Between subjects, I manipulated whether a particular choice set was framed as a privacy choice or not. First, I manipulated whether choices were presented to users as “Privacy Settings” or as “Survey Settings”. I also manipulated, between subjects, the importance of settings presented to users. In a pre-study, I had participants rank eleven settings which deal with the use and disclosure of user responses, some of which were designed to be more relevant to users than others (see Appendix G for full set of settings tested and results of the pre-test). For the “High Importance” conditions, participants were presented the top four most important settings to participants (e.g. “Allow my responses to shared with other participants of the study”) and conversely participants in the “Low Importance” condition were provided the four least important settings to participants (e.g. “Allow my responses to be used for academic publications”).

Procedure

Participants were recruited through Mechanical Turk, an online service that connects researchers with potential participants and is becoming increasingly popular among social scientists conducting online experiments. Participants were invited to take an online studies on ethical behavior each of which paid \$.25. Participants were first asked demographic questions, which included no directly identifying information but asked for their city and zip of residence and other demographic information. Then, they were provided with a four choices that related to the use and protection of their responses to the survey (See Appendix G). Finally, participants were presented with eight questions related to ethically questionable activities (See Appendix H). The questions used in both studies were the ones that were rated most intrusive in a 2012 paper by Acquisti, John and Loewenstein and were presented in random

order. Finally, participants were then asked a set of exit questions which evaluated, among other things, their satisfaction with the privacy protection provided and their perception of harm from participating in the study.

Analytical Model

I use a linear probability model, random effects estimation approach to evaluate the overall differences in the propensity to limit access and dissemination of individual responses. I estimated the following model:

$$Deny_{ij} = \beta_0 + \beta_1 * PrivacyFraming_i + \beta_2 * LowImp_{jt} + \beta_3 * PrivacyFraming_i * LowImp_{jt} + u_{ij}$$


$Deny_{ij}$ measures the propensity to deny a particular data request, with a value of 1 if the participant denies that use of their responses and 0 if she allowed a particular use, $i = \{1, \dots, N \text{ participants per interaction set}\}$, and $j = \{1, \dots, 4 \text{ settings}\}$. $PrivacyFraming_i$ is a binary indicator of the whether participant i was presented choices as “Privacy Settings” as opposed to “Survey Settings”. $LowImp_{jt}$ is a binary measure of whether the setting j was low or high relevance settings. Finally, I include an interaction between privacy framing and relevance to test differences in the framing effects for low and high importance settings. The model assumes serial correlation between observations within a panel unit. I allow for the correlation between responses from a single participant when I estimate the variance-covariance matrix of the coefficients, assuming constant correlation between any two answers by the same individual (Liang and Zeger, 1986).

Results

I had 204 participants ($M_{Age} = 29$ $SD_{Age} = 9.6$, $M_{Female} = .34$ $SD_{Female} = .48$) take this study. I find that participants presented the choice of settings labeled “Privacy Settings” the privacy framing were generally more likely to choose the more protective choices relative to those presented the same choices


as “Survey Settings”. In the high importance settings condition, participants in the “Privacy Settings” condition were more likely to choose protective setting for three of the four settings provided (See Figure 1) and these differences were significant for settings 2 and 3 ($P < .05$). On average, participants in the “Privacy Settings” condition chose the protective option 47% relative to 35% of the time for participants in the “Survey Settings” condition. Estimation of the random effects panel model (Table 1) confirms this finding with a significant and positive coefficient on *PrivacyFraming* indicating that participants in the “Privacy Settings” condition were 14% more likely to choose the privacy protective option for the high relevance condition (H1a supported). I also find no significant differences in disclosure between conditions with both groups admitting to unethical behavior 53% of the time.

[Figure 1: Summary Results – High Importance Settings]



In the low importance settings condition, participants in the “Privacy Settings” conditions were also more likely to chose privacy protective options relative to those in the “Survey Settings” condition for three of the four settings (see Figure 2) but these differences were not significant.

[Figure 2: Summary Results – Low Importance Settings]



On average, participants in the “Privacy Settings” condition chose the protective option 6% relative to 3% of the time for participants in the “Survey Settings” condition. Overall I find some evidence of a diminished framing effect for low relevance settings. Estimation of the random effects panel model (Table 1) identifies a negative and sizable coefficient on the interaction of *PrivacyFraming* and *LowImpt* which is directionally consistent with my hypothesis H1b suggesting that the framing effect was less pronounced for low importance settings. However, this estimate is insignificant in my analysis ($P=.147$). I find that individuals in the “Privacy Setting” condition were slightly *less* likely (51% vs. 55%) to admit to unethical behavior but this difference was not significant.

Discussion

The results of experiment 1 present some evidence that minor and subtle changes in the presentation of privacy relevant choices can significantly alter individual’s propensity to choose protective options. Moreover, I find that this effect is most pronounced for contexts that individuals felt it was important they were provided choice. Finally, I find that the participants presented choices as “Survey Setting” not only chose less protective settings but were just as likely, if not somewhat more likely, to disclose personal information. Given the commonality of similar decision frames in contexts that require individual privacy

decision making (e.g. social media, mobile privacy, etc.), this experiment suggests that subtle, and maybe accidental manipulation of decision frames may significantly influence individual behavior and subsequent privacy risks.

[Table 1: Study 1 Results]

	(Linear Probability Model)
	Deny
PrivacyFraming	.137
	(0.064)**
LowImpt	-0.242
	(0.047)***
Privacy * LowImpt	-0.102
	(0.070)
Constant	0.28
	(0.042)***
Observations	816
Standard errors in parentheses	
* significant at 10%; ** significant at 5%; *** significant at 1%	

4. Experiment 2

In a second study, I evaluate the factors that may exacerbate or diminish the framing effect identified in study 1. Specifically, I evaluate the potential impact of reference dependence to exacerbate the initial framing effect or, alternatively, the propensity of habituation effects to diminish framing effects. Study 2 is a two factor, between-subjects design in which I again manipulate whether (1) choices presented to participants are framed as privacy choices and the (2) homogeneity of settings in terms of their

importance to participants. I examined the effect of framing and mixture of settings on the propensity of individuals to choose protective settings. Again, I hypothesize that the framing of otherwise identical choices as privacy choices will result in more protective decision making by participants (H2a). I also consider the potential impact of mixed relevance settings on decision making. In the homogenous condition, participants are presented a set of all “High Importance” settings. In the mixed settings condition, participants are presented a mix of 3 low importance settings and one high importance setting.

This study evaluates how framing effects may interact with variance in choice sets in a way that may exacerbate or diminish the relevance of individual settings and thus impact the framing effects. Prior work on choice architecture suggests that adding irrelevant choices or information can alter decision making due to contrast effects or relative judgments that reframe the relevant choice as relatively more or less attractive (Ariely, Loewenstein, and Prelec, 2003; Ariely, 2009). Alternatively, it may be the case that presenting participants with low importance choices may result in low levels of initial concern and lull participants into a false sense of security in which they would be less likely to recognize a high importance choice. Finally, I argue that the expect effect will be associated with attentiveness and deliberation of individuals. Specifically, I argue that attentive individuals will notice the contrast in the mixed relevance settings and thus result in an increased likelihood of choosing protective settings in the mixed settings condition (H2b₁) and that the framing effect will be more pronounced for mixed settings relative to homogeneous settings (H2c₁). Conversely for low attention individuals, I posit that the mixed settings will result in a decreased likelihood to choose protective settings (H2b₂) and that the framing effect will be less pronounced relative to homogeneous settings (H2c₂).

Design & Procedure

The design was a 2 (“Privacy Settings”, “Survey Settings”) X 2 (High Importance Settings vs. Mixed Importance Settings). Between subjects, I again manipulated whether choices were presented to users as “Privacy Settings” or as “Survey Settings”. I also manipulated, between subjects, whether participants

were presented a homogenous set of High Importance settings similar to the condition in study 1 or a mixed set of settings with three Low Importance settings and one important setting shared by the High Importance condition. The individual setting that demonstrated the largest framing effects in Study 1 was whether participants allow their responses to be shared with religious organizations and was used as the shared setting for this study. The propensity of individuals to deny access to religious organizations (the more privacy protective choice) is my dependent variable of interest. I also collect the length of time each individual took to make their choice of settings. I evaluate the impact of presenting mixed importance settings on the propensity of individuals to choose privacy protective choices for high importance settings and the interaction of any framing effect with mixed relevance settings. The procedure is identical to that of Study 1 except that participants are not asked to answer questions about ethical behavior.

Analytical Model

I estimate the model described below using both a linear probability model.

$$Deny_i = \beta_0 + \beta_1 * PrivacyFraming_i + \beta_2 * MixedImp_i + \beta_3 * PrivacyFraming_i * LowImp_i + u_i$$

$Deny_i$ measures the propensity to deny a particular data request, with a value of 1 if the participant denies that use of their responses and 0 if she allowed a particular use, $i = \{1, \dots, N$ participants per interaction set.


$PrivacyFraming_i$ is a binary indicator of the whether participant i was presented choices as “Privacy Settings” as opposed to “Survey Settings”. $MixedImp_i$ is a binary measure of whether participant i viewed the setting as part of a mixed set of settings. Finally, I include an interaction between privacy framing and relevance to test for differential framing effects for mixed and homogenous settings.

Results

I had 522 participants ($M_{Age} = 28$ $SD_{Age} = 10.8$, $M_{Female} = .44$ $SD_{Female} = .49$) take study 2. I collected more data for this study because I had one observation per participant as opposed to 4 in the prior study (only

one setting was common across all conditions). I find a similar framing effect as that found in my previous experiment with the presentation of the setting a “Privacy Setting” resulting in more participants choosing the protective option relative to those presented the same setting as a “Survey Setting”. I find support for this in my estimation of the random effects model with a positive coefficient on *PrivacyFraming* that participants provided homogenous settings were 7% more likely to choose the protective choice ($P=.078$). Moreover, I find that mixed settings had a significant baseline effect on the choice of protective settings with participants in the mixed condition 20% less likely ($P<.01$) to choose the protective option (H2b₂ Supported). Finally, I find a similar framing effect for participants that were presented the setting in a mixed group of settings ($P<.05$) indicating no interaction of mixed settings and framing effects (H2c₂ not supported).

[Figure 3: Second Study Summary Results]




This last result is somewhat surprising because I find a baseline effect of the mixed setting manipulation but this does not seem to influence the framing effects. I considered that this may be due to the fact that only those who pay sufficient attention to the settings noticed the high relevance setting and thus attentiveness of the individual participants may also moderate this effect. To evaluate this claim, I use data collected as part of my survey which captures how long each participant took to make their selection of settings and parsed the data into “High Attention” and “Low Attention” groups depending on whether they were above or below the median time taken.

[Table 2: Study 2 Results]

	(Linear Probability Model)
	Deny
PrivacyFraming	.10
	(0.06)*
Mixed	-0.197
	(0.047)***
Privacy * Mixed	-0.004
	(0.075)
Constant	0.31
	(0.039)***
Observations	522
Standard errors in parentheses	
* significant at 10%; ** significant at 5%; *** significant at 1%	


I find that the effect of mixed settings is driven strongly by “Low Attention” participants with only 5.5% of the participants choosing the protective option relative to 35% of their “Low Attention” counterparts in the homogenous settings condition. For the “High Attention” group, participants in the mixed settings manipulation were somewhat less likely to chose the protective option but this difference was not significant. Finally, I note that the time taken to make the choice of settings does not correlate with the privacy framing or whether participants were presented mixed or homogenous settings.

[Figure 4: Effect of Mixed Settings by Attentiveness]



Restricting my sample to “Low Attention” participants, I also find that the framing effect in the mixed settings condition is less pronounced (5% vs. 10%) when compared to the homogenous settings condition (consistent with H2c₂). Conversely, if I restrict my sample to the “High Attention” group, I see a reversal of the effect with the framing effect now pronounced (17% vs. 10%) for those in the mixed settings condition (consistent with H2c₁). These differences in framing effects, while directionally consistent with my hypotheses were not significant (See Figure 5) . For the homogenous condition, the attentiveness of participants did not alter either the framing effect or the propensity to choose protective options, suggesting that attentiveness is particularly relevant for the mixed settings manipulation.

[Figure 5: Interaction of Framing Effect and Mixed Settings by Attentiveness]



Discussion

The results of study 2 reinforce that subtle manipulation of the framing of privacy choices may have a significant effect on the propensity of individuals to choose protection options. It also highlights the potential impact of presenting users with mixed relevance choices on the propensity of individuals to choose protective options with regard to risky or relevant uses of their personal information. These results suggest that this risk is especially pronounced for low attentiveness individuals. Moreover, I identify a potentially novel theoretical contribution evaluating how reference dependence may differentially highlight the importance of risky choices and exacerbate framing effects. Future experiments would be necessary to explicitly test this finding.

6. Study 3

In a third study, I evaluate the bounds of a framing effect in the context of choice of privacy settings. First I alter the manipulation previously termed "Survey Settings" to a label that frames the goal of the settings as openness or wider disclosure ("Sharing Settings"). Secondly, I manipulate the framing of individual settings to present them as either "accept" or "reject" decisions. Prior literature suggests that presenting a choice as a choice to accept brings to bear the positive features of that choice and thus result in a higher

propensity to accept relative to an objectively identical choice framed as a choice to reject (Shafir, 1993; Tversky and Shafir, 1992). Moreover, they find that positive information features more prominent in the choice to accept and that negative information features more prominently in the choice to reject. Study 3 is a two factor, between-subjects design in which I manipulate whether (1) choices presented to participants are framed in a manner that highlights privacy concerns or a manner that highlights the desire to share personal information and the (2) the framing of the individual settings as either an “allow” or “prohibit” decision. I examine the effect of framing and the allow/prohibit manipulation on the propensity of individuals to choose protective settings, hypothesizing that the framing of otherwise identical choices as privacy choices will result in more conservative decision making by participants relative to the condition in which the choice is framed as choices about sharing or openness (H3a). I also consider the potential impact of altering the language of the individual settings to present the choice as either allowing a particular use of personal information or restricting the use of personal information. Prior work finds that presenting a choice as an allowance highlights positive dimensions of a choice for individuals while presenting choice as a prohibition or rejection highlights negative dimensions for participants (Shafir, 1993). As a result, I hypothesize that framing choices as prohibitions will result in an increased propensity to choose protective settings (H3b). Moreover, Ganzach and Schul (1994) find that negative information figures more prominently in decision making when the choice is framed in terms of a prohibition or a rejection and vice versa for allowances or acceptances. Thus I also posit that the effect of a privacy framing will be pronounced when choices are presented as prohibitions relative to allowances (H3c).

Design and Procedure

The design was a 2 (“Privacy Settings”, “Sharing Settings”) X 2 (Allow Settings vs. Prohibit Settings). Between subjects, I manipulated whether choices were presented to users as “Privacy Settings” or as “Sharing Settings”. I also manipulated, between subjects, whether participants were presented the settings as a choice to allow a use of personal information or prohibit a use of personal information (settings will be objectively identical). Again, the propensity of individuals to choose a protective choice on this shared

settings is my dependent variable of interest. I evaluate the impact of the framing on the propensity of individuals to choose privacy protective, the baseline impacts of presenting choice as a choice to allow vs. a choice to prohibit, and the interaction of any framing effect with the framing of individual settings. The procedure is identical to that of Study 1 except that instead of providing participants a setting related to passwords, I use a setting dealing encryption as the allow framing was easier to understand for that context (“Allow my responses to be stored unencrypted” vs. “Allow my responses to be stored on a drive that is not password protected”).

Analytical Model

I again plan to use a panel random effects estimation approach (both Probit and linear probability model) to evaluate the overall differences in the propensity to limit access and dissemination of participant responses. I estimated the following model:

$$Deny_{ij} = \beta_0 + \beta_1 * PrivacyFraming_i + \beta_2 * Prohibit_j + \beta_3 * PrivacyFraming_i * Prohibit_j + u_{ij}$$

$Deny_{ij}$ measures the propensity to deny a particular data request, with a value of 1 if the participant denies that use of their responses and 0 if she allowed a particular use, $i = \{1, \dots, N \text{ participants per interaction set}\}$, and $j = \{1, \dots, 4 \text{ settings}\}$. $PrivacyFraming_i$ is a binary indicator of the whether participant i was presented choices as “Privacy Settings” as opposed to “Sharing Settings”. $Prohibit_j$ is a binary measure of whether the setting j was presented in a prohibit frame vs. an allow frame. Finally, I include an interaction between privacy framing and the prohibit frame to test differences in the framing effects for settings framed as prohibitions. This models assumes serial correlation between observations within a panel unit. I allow for the correlation between responses from a single participant when estimating the variance-covariance matrix of the coefficients, assuming constant correlation between any two answers by the same individual (Liang and Zeger, 1986).


$$Disclosure_i = \beta_0 + \beta_1 * PrivacyFraming_i + \beta_2 * TotDeny_i + \delta_k * X_k + u_i$$

I also estimate a similar model to evaluate the impact of privacy framing and choice of privacy settings on disclosure. $Disclosure_i$ is a measure of the percent of unethical behaviors that participant i admitted to in the study. $PrivacyFraming_i$ is a binary indicator of the whether participant i was presented choices as “Privacy Settings” as opposed to “Sharing Settings”. $TotDeny_i$ is the number of settings for which of participant i chose the protective option. X_k is a set of demographic controls. Because $TotDeny_i$ is a choice variable and not randomly manipulated it is likely correlated with other factors that bias the estimates in the basic model presented above. As a result, I use an intention to treat approach (Angrist, Imbens & Rubin, 1996) and instrument for $TotDeny_i$ with my $Prohibit_i$ measure above. Because $Prohibit_i$ is a randomly assigned variable I assume that $Cov(Prohibit_i, u_i) = 0$ but predict that $Cov(Prohibit_i, TotDeny_i) \neq 0$. I use a two stage least squares approach to estimate the average effect of choice of settings on disclosure (Angrist and Imbens 1995).

Results


I had 401 participants ($M_{Age} = 28$, $SD_{Age} = 10.8$; $M_{Female} = .44$, $SD_{Female} = .49$) take study 3. First, I find a strong effect of the accept vs. prohibit framing with participants significantly more likely (56% vs. 37%) to choose protective settings when presented a choice to restrict a use of their data vs. allow it (See Figure 6). This difference was significant for all settings provided to participants. Estimation of the random effects model finds a significant ($P < .01$) positive coefficient on *Prohibit* confirming this result.

[Figure 6: Accept vs. Reject Framing]



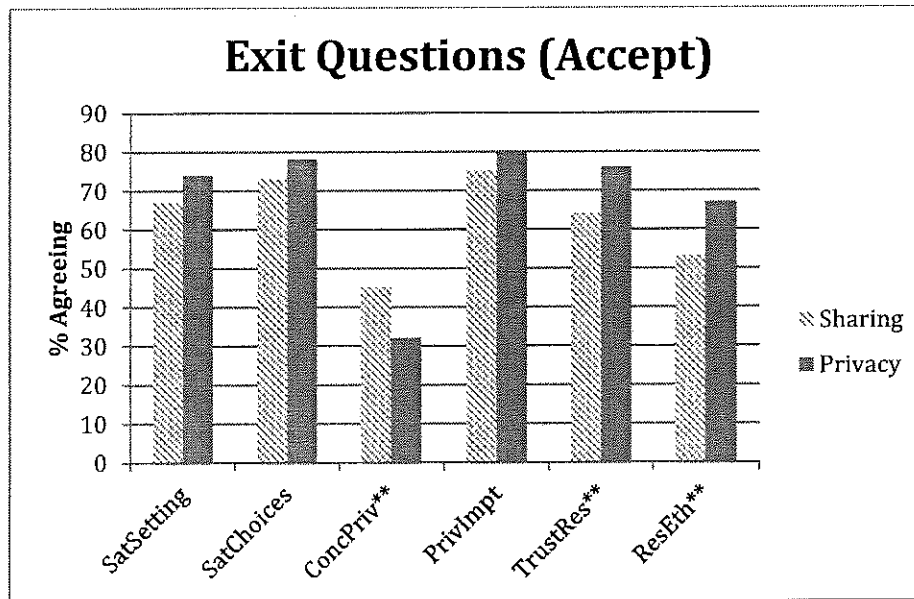
However, I do not find the baseline framing effect identified in the two prior studies in this experiment. In fact, I find a trend in the opposite direction with those presented the “Sharing Settings” framing being somewhat more likely (40% vs. 30%) to choose protective options in the accept frame (participants were equally likely to choose the protective option in the reject frame). Estimation of the random effects model finds a positive coefficient on *PrivacyFraming* but this estimate was not significant ($P=.24$). Moreover, participants in the Sharing Settings condition also trended towards lower levels of disclosure (although this difference was not significant).

[Figure 7: Accept vs. Reject Framing, Accept Frame]



While this result is initially surprising, I suggest that this result may point to the counter-intuitive effect of attempts to present choices with significant privacy implications in a manner that promotes less protective behavior. Namely, I argue that participants may have felt that I was attempting to manipulate them to pick less protective options which resulted in distrust and caution from participants. Evaluating exit questions, I find that despite the fact that participants in the Sharing Settings condition chose more protective options and shared less personal information, they were significantly *more* likely to report being concerned about their privacy when responding to ethical questions (45% vs. 32%, $P < .05$) and were also significantly *less* likely to report that the researchers were ethical (53% vs. 67%, $P < .05$) and trustworthy (64% vs. 76%, $P < .05$).

[Figure 8: Exit Questions, Accept Frame]



An estimation of the model evaluating the impact of privacy framing and choice of settings on disclosure finds a small and insignificant estimate on *PrivacyFraming* suggesting that the framing did not have an impact on disclosure. In the basic estimation I find a negative correlation between the choice of protective settings and disclosure. This is surprising result as I might expect that those who choose more protective options would disclose more, all else constant. I suggest that this is likely due to the unobserved factors correlated both with disclosure and choice of settings. Instrumenting for the choice of settings with my prohibit/accept manipulation, I indeed find a positive effect of more protective settings on disclosure. This estimate however is not significant ($P=.335$).

[Table 3: Study 3 Results]

	Linear Probability Model	Basic Estimation	IV Estimation
	Deny	Disclosure	Disclosure
PrivacyFraming	-.07	.003	.008
	(0.06)	(0.02)	(0.03)

Prohibit	0.16	--	--
	(0.06)***	--	--
Privacy * Prohibit	0.07	--	--
	(0.08)	--	--
TotDeny	--	-.014	.03
	--	(0.006)**	(0.03)
Constant	0.28	.565	.486
	(0.042)***	(0.05)***	(0.08)***
Observations	1604	3208	3208
Standard errors in parentheses			
* significant at 10%; ** significant at 5%; *** significant at 1%			

Discussion

The results of study 3 highlight the substantial role of framing choices as a choice to restrict versus a choice to allow in driving privacy protective behavior from individuals. In particular, it highlights the inconsistency in individual choice with respect to privacy protective options and also demonstrates how manipulations of framing can be used to systematically guide users towards more or less protective settings. Moreover, this study find the unexpected result that attempts to highlight the sharing dimension of choice may backfire and result in reduced confidence and trust from users, thus leading to more protective choices and less disclosure. First, this identifies that framing manipulations attempting to highlight positive dimensions of a choice (in this case sharing) may be bounded by user trust. This may be particularly true when costs (i.e. privacy risks) are relatively obvious to participants.

6. Study 4

In a fourth study, I evaluate the role of framing in the face of opt-in vs. opt-out approaches towards eliciting choice with respect to the collection and use of personal information. Similar to study 1 and 2 I vary between subjects whether choices are presented to individuals as “Privacy Settings” or “Survey Settings”. Secondly, I vary whether the choices are presented to individuals as a choice to opt into a use of their personal information or a choice to opt-out of a use of their personal information. In the context of privacy decision making, *opt-in* refers to the manner of soliciting choice in which individuals have to explicitly consent to the uses of their personal information while in an *opt-out* approach, entities collect and use personal information unless the individual expresses an explicit preference for them *not* do so. Behavioral and decision researchers have also highlighted the important role of a default or status quo bias in shaping individual behavior. For example, Johnson and Goldstein (2003) find individuals in countries where being an organ donor was the default and individuals had to explicitly chose not to be a donor (an opt-out approach) individuals were significantly more likely to be organ donors relative to countries where individuals were asked to explicitly indicate that they wished to be an organ donor (i.e. the default was not being an organ donor). Similar, effects have been identified in a savings context with Choi et al (2004) finding that most individuals chose the default savings rate provided to them. This discussion is relevant to the privacy decision making context with significant variation exists across policy and regulatory contexts with regards to whether privacy relevant choices are presented as opt-in vs. opt-out (Goldstein and Rein 2010), spurring a significant debate over the implications of these approaches for both firms and consumers (Milne and Rohm 2000; Simon et al 2009) including potentially a significant impact on individuals’ likelihood to choose privacy protective options (Solove 2013). Some empirical privacy research has found that these default effects are likely to extend to privacy decision making, with Brandimarte, Acquisti, and Loewenstein (2012) finding that participants are more likely to pick a setting when the choice is selected by default. Similar to Study 1 and 2, I hypothesize (H4a) that the privacy framing will result in the choice of more protective settings. Moreover, I hypothesize (H4b) that an Opt-in (i.e. more restrictive options are the default) manipulation will lead to, on average, the

choice of more protective setting by individuals due to the status quo bias. Finally, I posit (H4c) that the framing effect will be more pronounced for the Opt-out manipulation. Specifically, I conjecture that because the opt-out approach highlights the choice to restrict uses of the data (i.e. say no to uses of your data), the privacy framing may figure more prominently for those individuals.

Design and Procedure

The design was a 2 (“Privacy Settings”, “Survey Settings”) X 2 (Opt In vs. Opt Out). Between subjects, I manipulated whether choices were presented to users as “Privacy Settings” or as “Survey Settings”. I also manipulated, between subjects, whether participants were presented the settings as an opt-in choice vs. an opt-out choice (see figures 9a and 9b). Specifically, I presented participants identically phrased statements relating to how their information would be used, and between conditions, I varied whether participants were asked to indicate which of the uses they allowed (opt-in) or which of the uses they did not allow (opt-out). I evaluate the impact of the framing on the propensity of individuals to choose privacy protective options, the baseline impacts of the default manipulation, and the interaction of any framing effect with default effects. The procedure is otherwise identical to that of Study 3.

[Figure 9a: Opt-In Manipulation]

Below are the ways that this study may use your responses. Check "Yes" for the uses you allow.

	Yes
Share my responses with other participants of the study.	<input type="checkbox"/>

[Figure 9b: Opt-Out Manipulation]

Below are the ways that this study may use your responses. Check "No" for the uses you do not allow.

	No
Share my responses with religious organizations interested in evaluating personal ethics.	<input type="checkbox"/>

Analytical Model

I plan to use a panel linear probability, random effects estimation approach to evaluate the overall differences in the propensity to limit access and dissemination of participant responses. I estimated the following model:

$$Deny_{ij} = \beta_0 + \beta_1 * PrivacyFraming_i + \beta_2 * OptIn_j + \beta_3 * PrivacyFraming_i * OptIn_j + u_{ij}$$


$Deny_{ij}$ measures the propensity to deny a particular data request, with a value of 1 if the participant denies that use of their responses and 0 if she allowed a particular use, $i = \{1, \dots, N \text{ participants per interaction set}\}$, and $j = \{1, \dots, 4 \text{ settings}\}$. $PrivacyFraming_i$ is a binary indicator of the whether participant i was presented choices as “Privacy Settings” as opposed to “Sharing Settings”. $OptIn_j$ is a binary measure of whether the setting j was presented as an choice to opt-in vs. a choice to opt-out. Finally, I include an interaction between privacy framing and the choice to opt-in to test differences in the framing effects for opt-in vs. opt-out choice frameworks. The model assumes serial correlation between observations within a panel unit. I allow for the correlation between responses from a single participant when I estimate the variance-covariance matrix of the coefficients, assuming constant correlation between any two answers by the same individual (Liang and Zeger, 1986).

Results

I had 406 participants ($M_{Age} = 31$ $SD_{Age} = 10.6$, $M_{Female} = .42$ $SD_{Female} = .51$) take study 4. First, I find an effect of privacy framing for the opt-out condition. Participants in the “Privacy Settings” condition were significantly more likely to choose a protective option (41% vs. 28%, $P < .05$) relative to participants in the “Survey Settings” condition. Participants in the opt-in condition did not see an effect of the framing (see Figure 10). Moreover, from figure 10 I can see that this framing effect was driven largely by the strong effect of the opt-out manipulation when a privacy framing was absent (i.e. “Survey Settings” label). In fact, across all manipulations across all studies this group was the least likely to choose protective

options. The choice of protective options when presented with the privacy framing condition was equivalent for both Opt-in and Opt-out conditions. I also find a marginally significant ($P=.058$) positive baseline effect of the Opt-In manipulation. Finally, I find that participants were significant more likely to report that they were satisfied with privacy settings when presented as a choice to opt-out. This is surprising considering this approach actually elicits *lower* levels of protective behavior on average. I suggest that this may be due to individuals reacting to the fact that in the opt-out condition I am extending participants a choice to protect themselves while in the opt-in condition I am asking them to expose themselves to risk.

[Figure 10: Effect of Framing]



[Table 4: Experiment 4 Results]

	(Linear Probability Model)
	Deny
PrivacyFraming	.13
	(0.06)**

Opt-In	0.10
	(0.06)*
PrivacyFraming *	-0.12
OptIn	
	(0.08)
Constant	0.27
	(0.042)***
Observations	1624
Standard errors in parentheses	
* significant at 10%; ** significant at 5%; *** significant at 1%	

Discussion

The results of study 4 highlight the substantial role of the default bias and framing on driving privacy protective behavior from individuals. In particular, it highlights the manner in which an Opt-out approach towards eliciting privacy decision making may be most prominent when privacy dimensions of choice are not highlighted and can be counteracted with privacy framing. Moreover, it suggests that Opt-In approaches towards eliciting choice may be less sensitive to framing manipulations that highlight the costs of such choices and may inform some policy contexts. For example, in contexts where a standardized format may not exist, an opt-in approach towards eliciting privacy preferences may be a method more robust to manipulations of framing. Moreover, in contexts where opt-out approaches are the norm or mandated by policy, it may be the case that subtle manipulation in the framing of choices may increase the probability of protective behavior for those individuals and counter-act the default bias.

7. Discussion and Conclusions

In this chapter I demonstrate significant malleability in individual preferences for privacy under various manipulations of the framing of objectively identical choices. First, I find that the labeling of privacy relevant choices as “Privacy Settings” results in the choice of more protective settings relative to labeling the identical choices as “Survey Settings”. Moreover, I find that presenting individuals with an important privacy choice mixed with other settings that were pre-ranked by participants to be less relevant, significantly decreased the likelihood of participants to choose the protective option for the important choice. Finally, I find that participants presented privacy relevant choices as a choice to allow a use of their personal information (accept frame) were significantly *less* likely to choose the privacy protective option relative to those presented the identical choice as a choice to prohibit a use of their personal information (reject frame). Generally, I find that manipulations of framing do not have a significant impact on disclosure, even when I control for the choice of settings by participants. These results suggest that common but subtle variation in the framing of privacy choices can predictably influence users to chose less protective privacy options.

However, this work offers a largely descriptive view of the world and does not necessarily suggest that all choices with a privacy dimension should, for example, necessarily be framed as a privacy choice or be presented in a homogenous set of relevant choices. Particularly since choices that involve individual personal information are increasing in number, span numerous contexts, and tend to also be associated with some benefit for consumers (e.g. an online service, or better targeted ads). I simply point out that these subtle manipulations in the presentation of privacy relevant decisions can predictably minimize consumer concerns resulting in less privacy protective behavior from consumers.

This work has a number of implications for policy makers. Namely, it suggests that current policy approaches may be relying too heavily on consumer choice to alleviate consumer privacy concerns and alleviate risks. Given that choice may be easily manipulated, it suggests that baseline protections may be

necessary for certain uses of personal information which may be particularly disadvantageous or harmful to consumers. This could be in contexts where there is significant consumer outcry associated with a particular use or collection of personal information or contexts in which the privacy-utility tradeoff is considerably skewed against consumers. The need for baseline regulation is particularly relevant given that the results in this chapter identify only one manipulation of consumer decision biases that can elicit substantial differences in behavior. It is almost certain that others exist as well which may also be effective at predictably swaying consumer decision making. In contexts where policy makers currently rely on consumer choice to address concerns, my results suggest that careful evaluation of the presentation and framing of these choices may be warranted to ensure consumer protection. Finally, it suggests that regulatory or self-regulatory approaches that allow firms considerable leverage in designing choice mechanisms for consumers may result in consumers continuing to face significant privacy risks, particularly if firms continue to have significant incentives to elicit *more* disclosure from consumers.

Chapter 5: Discussion and Conclusions

Transparency and choice mechanism have been a long-standing tenant of privacy protection, are included as central privacy protections in current privacy regulation, and continue to be central policy mechanisms for addressing emerging privacy concerns. This dissertation evaluates the impact of increased transparency and choice on both firm adoption of technology that leverages personal information and consumer privacy risks.

First I show that privacy regulation providing consumers transparency and choice can reduce consumer privacy concerns associated with a particular technology and thus spur technology adoption. Conversely, I find that weaker regulation without transparency and choice failed to alleviate consumer concerns and did not have a similar positive effect on technology adoption. This work contributes to the growing body of empirical investigation of the role of privacy regulation on technology adoption and bolsters the notion that privacy regulation may not exclusively impact technology adoption either positively or negatively, but is likely influenced by a number of contextual and environmental factors and also by the features of the technology. For example, the extent to which a technology is consumer facing, the sensitivity of the data exchanged, the level of control users have over the exchange or use of their information, or the salience/familiarity of the privacy concerns associated with a technology could all be factors that modulate the impact of privacy regulation.

Shifting my focus to the role of increased transparency on individual privacy decision making I find that the impact of privacy notices on disclosure is sensitive to whether notices are presented as increasing or decreasing in protection, even when the objective risks of disclosure stay constant, and that the propensity of privacy notices to impact disclosure can be muted by a number of simple and minimal misdirections (such as a mere 15 second delay between notices and disclosure decisions) that do not alter the objective risk of disclosure. It follows that privacy notices can – on the one hand – be easily marginalized to no

longer impact disclosure, or – on the other hand – be used to influence consumers to share varying amounts of personal information. Transparency may, therefore, become a “sleight” of privacy. Finally, I evaluate the propensity of providing consumers increased control to consistently influence individual privacy decision making and the impact on subsequent privacy risk. I find considerable malleability in the choice of privacy protective settings under varying decision frames including altering the labeling of choices presented to participants, the homogeneity of the importance of the choice set, and presenting choices in accept frame relative to a reject frame. These results suggest that common but subtle variation in the framing of privacy choices can predictably influence users to chose less protective privacy options.

Taken together, the results in this dissertation suggest that a significant challenge exists for policy makers. While transparency and control solutions may be able to reduce barriers to innovation stemming from consumer privacy concerns, in practice, these mechanisms may not consistently or reliably reduce objective risks for consumers. This has important implications for policy makers, firms, and consumers. For instance, transparency and choice mechanisms may simply not be effective in addressing, in the long term, consumer privacy concerns, thus resulting in the persistence of the privacy challenges despite increased transparency and choice. In a more troubling scenario, consumers provided transparency and choice may feel more satisfied or less concerned about their personal privacy despite facing objectively identical or potentially, even higher risks. For example, in experiment 4 in chapter 4, I find that participants provided the choice to opt-out of uses of their data (an approach which resulted in the choice of *less* protective options), reported being *more* satisfied with the settings provided than those provided the choice to opt-in (an approach which resulted in the choice of more protective options). Despite these findings, I argue that increased transparency and control with respect to the collection, use, and dissemination of consumer personal information is likely necessary, but may not be sufficient to address consumer privacy concerns.

As a result, policy makers may consider reducing their reliance on transparency and choice as privacy protective mechanisms. For example, policy makers may consider, alongside transparency and control mechanisms, regulation which restricts collections and uses of personal information particularly disadvantageous or harmful to consumers and implements other OECD privacy principles (OECD, 1980). With respect to regulation, policy makers may also focus on contexts that are technically complex or where the tradeoffs to consumers are not straightforward. The need for baseline regulation is particularly relevant given that numerous decision biases not identified in this work may also be relevant to the context of privacy decision making. This balanced approach of underlying regulation coupled with transparency and choice mechanisms may also help to curb the trend towards increased consumer “responsibilization” – a situation where individuals are “rendered responsible for a task which previously would have been the duty of another [...] or would not have been recognized as a responsibility at all” (Wakefield and Fleming, 2009).

Where policy makers continue to rely on transparency and consumers choice to address concerns, there may be a need to expand the notion of transparency and careful consideration of the presentation of relevant consumer privacy choices. For instance, behavioral economists and decision researchers have identified various strategies to aid consumers in improved decision making that may also be useful for privacy decision making. The 2008 book by Thaler and Sunstein, describes how policy makers can use soft paternalistic interventions or “nudges” to counter-act known limitations in decision making that may inhibit consumers’ ability to make optimal decisions. A nudge utilizes or counteracts a known decision bias (e.g., a default effect) to urge consumers that exhibit limitations in decision making (e.g. limited attention or immediate gratification bias) towards improved decision making, while allowing rational and cognizant consumers to make informed, willful decisions. One example of a nudge is a default choice or setting, be that in the form of savings plans or organ donations. In a similar manner, research may consider providing default settings for consumers that are more protective of consumer privacy, or consider counteracting consumers’ limited attention by intelligently providing relevant parts of privacy

notices to consumers at the points of disclosure. For example, research may intelligently identify disclosures that consumers are likely to regret (e.g. disclosures with vulgarity or mentions of their bosses and coworkers) and remind them at that instant of the various entities that may view this particular disclosure. Similarly, prior to a consumer accepting privacy invasive terms and conditions of a particular application or service, the actual choice may be briefly delayed to allow the time for an evaluation of the trade-offs associated with such a decision, or for a reminder of intrusive uses and exchanges of their data that this particular application may engage in.

Specific to the context of privacy choices, such as privacy settings associated with behavioral advertising or online social networks, I suggest that careful consideration be given to the design of these choices and potentially how framing effects (or lack thereof) may result in predictable differences in choice of protective settings. Specifically, I suggest that high relevance choices with significant impact on consumer privacy outcomes could be presented in a decision frame that highlights the privacy dimension of the choice and elicits protective behavior for those with strong preference for privacy. This however, may be difficult to impose on firms given the current self-regulatory approaches for online consumer protection which grant firms significant flexibility in the mechanisms used to provide consumers transparency and choice. This in combination with considerable incentives for firms to collect increasing amounts of personal information about individuals suggests that consumers may be at risk of strategies designed to elicit high opt-in from consumers. For example, firms today already use some approaches to elicit high opt-in, such as requiring consumers agree to their terms of service as a pre-requisite to having access to the service. While challenges remain with the current approaches towards providing consumers transparency and choice, I suggest that a way forward exists and is feasible. Addressing the concerns identified in this work will likely improve trust between consumers and firms and allow for a healthy balance between technology innovation and personal privacy.

References

1. Acquisti A, Friedman A, Telang R. (2006). Is There a Cost to Privacy Breaches? An Event Study. in Proceedings of the 27th International Conference on Information Systems, Milwaukee, WI, pp. 1563-1580.
2. Acquisti A, John L, and Loewenstein G. (2012). The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research*, 49(2): 160-174.
3. Acquisti A. (2004). Privacy in Electronic Commerce and the Economics of Immediate Gratification. Proceedings of the ACM Conference on Electronic Commerce, New York: Association for Computing Machinery, 21-29.
4. Acquisti A. (2009). Nudging Privacy: The Behavioral Economics of Personal Information. *Security & Privacy*, IEEE 7(6): 82-85.
5. Adler-Milstein J, Bates DW, and Jha AK. (2009). U.S. Regional Health Information Organizations: Progress but Challenges Remain. *Health Affairs*, 28:2, 483-492.
6. Adler-Milstein J, Bates DW, and Jha AK. (2011). A Survey of Health Information Exchange Organizations in the United States: Implications for Meaningful Use. *Annals of Internal Medicine*, 154(10): 666-671.
7. Allen, I. E., & Seaman, J. (2005). Growing by degrees: online education in the United States, 2005. *Sloan Consortium*.
8. Angrist, J. D., & Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American statistical Association*, 90(430), 431-442.
9. Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434), 444-455.
10. Angwin J and Jennifer VD. (2012). Google's iPhone Tracking: Web Giant, Others Bypassed Apple Browser Settings for Guarding Privacy. February 17, 2012. http://online.wsj.com/article_email/SB10001424052970204880404577225380456599176-lMyQjAxMTAyMDEwNjExNDYyWj.html#.
11. Ariely, D. (2009). Predictably irrational, revised and expanded edition: The hidden forces that shape our decisions. Harper.
12. Ariely, D., Loewenstein, G., & Prelec, D. (2003). "Coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, 118(1), 73-106.
13. Bakos, Y. (1998). The emerging role of electronic marketplaces on the Internet. *Communications of the ACM*, 41(8), 35-42.
14. Bamberger K and Mulligan D. (2011). Privacy on the Books and on the Ground. *Stanford Law Review*, Vol. 63.
15. Becket, L. (2013). Big Data Brokers: They Know Everything About You and Sell it to the Highest Bidder. Gizmodo. <http://gizmodo.com/5991070/big-data-brokers-they-know-everything-about-you-and-sell-it-to-the-highest-bidder>
16. Bertrand M, Duflo E, and Mullainathan S. (2004). How Much Should We Trust Differences-in-Differences Estimates?. *Quarterly Journal of Economics*, 119:1, 249-275.
17. Bhattacharjee, S., Gopal, R. D., & Sanders, G. L. (2003). Digital music and online sharing: software piracy 2.0?. *Communications of the ACM*, 46(7), 107-111.
18. Bittlingmayer G. (2001). Regulatory Uncertainty and Investment: Evidence from Antitrust Enforcement. *Cato Journal*, Vol. 20, No. 3.
19. Bott, E. (2012). The Do Not Track standard has crossed into crazy territory. ZDNet. <http://www.zdnet.com/the-do-not-track-standard-has-crossed-into-crazy-territory-7000005502/>
20. Brandimarte L, Acquisti A, and Loewenstein G. (2012). Misplaced Confidences: Privacy and the Control Paradox. *Social Psychological and Personality Science*. <http://spp.sagepub.com/content/early/2012/08/08/1948550612455931>.

21. Broadbent DE. (1958). Perception and Communication. Elmsford, NY, US: Pergamon Press. Volume 340 pp. doi: 10.1037/10037-010.
22. Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2004). For better or for worse: Default effects and 401 (k) savings behavior. In Perspectives on the Economics of Aging (pp. 81-126). University of Chicago Press.
23. DellaVigna S. (2007). Psychology and Economics: Evidence from the Field. NBER Working Paper No 13420.
24. Deuze, M. (2001). Online journalism: Modelling the first generation of news media on the World Wide Web. *First Monday*, 6(10), 1-22.
25. eHealth Initiative. (2005). Emerging Trends and Issues in Health Information Exchange. http://www.ehealthinitiative.org/sites/default/files/file/eHI2005AnnualSurveyofHealthInformationExchange2_0.pdf.
26. eHealth Initiative. (2006). Improving the Quality of Healthcare through Health Information Exchange. <http://www.ehealthinitiative.org/files/eHI2006HIESurveyReportFinal09.25.06.pdf>.
27. eHealth Initiative. (2007). Fourth Annual Survey of Health Information Exchange At the State and Local Levels. <http://www.ehealthinitiative.org/sites/default/files/file/2007 HIE Survey results.pdf>.
28. eHealth Initiative. (2008). Fifth Annual Survey of Health Information Exchange At the State and Local Levels. <http://www.ehealthinitiative.org/sites/default/files/eHI-HIESurveyResultsFinalReport-2008.pdf>.
29. eHealth Initiative. (2009). Migrating Toward Meaningful Use: The State of Health Information Exchange. <http://www.ehealthinitiative.org/sites/default/files/file/2009%20Survey%20Report%20FINAL.pdf>.
30. eHealth Initiative. (2010). The State of Health Information Exchange in 2010: Connecting the Nation to Achieve Meaningful Use. <http://www.ehealthinitiative.org/uploads/file/Final%20Report.pdf>.
31. Ellison, N., Heino, R., & Gibbs, J. (2006). Managing impressions online: Self- presentation processes in the online dating environment. *Journal of Computer - Mediated Communication*, 11(2), 415-441.
32. 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.
33. Frey JH. (1986). An Experiment with a Confidentiality Reminder in a Telephone Survey. *Public Opinion Quarterly*. 50, 267-269.
34. 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>.
35. Galinsky, A. D., Gruenfeld, D. H., & Magee, J. C. (2003). From power to action. *Journal of personality and social psychology*, 85(3), 453.
36. Garton, L., Haythornthwaite, C., & Wellman, B. (1997). Studying online social networks. *Journal of Computer - Mediated Communication*, 3(1), 0-0.
37. Gellman, R. (2012). FAIR INFORMATION PRACTICES: A Basic History. <http://bobgellman.com/rg-docs/rg-FIPShistory.pdf>
38. Goldfarb A and Tucker C. (2011b). Privacy and innovation. Working Paper 17124,
39. Goldfarb A and Tucker C. (2011a). Privacy Regulation and Online Advertising. *Management Science*, Vol. 57, No. 1, pp. 57-71.
40. Goldstein M and Rein A. (2010). Consumer Consent Options for Electronic Health Information Exchange: Policy Considerations and Analysis. Office of Policy and Planning: Office of the National Coordinator for Health IT.

41. Greenberg MD, Ridgely MS, and Hillestad RJ. (2009). Crossed Wires: How Yesterday's Privacy Rules Might Undercut Tomorrow's Nationwide Health Information Network. *Health Affairs* 28:2, 450-452.
42. Horowitz S. (2011). Carrier IQ Faces Federal Probe into Allegations Software Tracks Cellphone Data. Decembr 14, 2011. http://www.washingtonpost.com/business/economy/feds-probing-carrier-iq/2011/12/14/gIQA9nCEuO_story.html.
43. Hossain T and Morgan J. (2006). Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay. *The B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy*. Volume 6, Issue: 2, 1-27.
44. Huang, C. (2011). Facebook and Twitter key to Arab Spring uprisings: report. *The National. Abu Dhabi Media*, 6.
45. Ishii J and Yan J. (2004). Investment under Regulatory Uncertainty: U.S. Electricity Generation Investment Since 1996. Center for the Study of Energy Markets, University of California Energy Institute, UC Berkeley.
46. Janger, E. J., & Schwartz, P. M. (2001). Gramm-Leach-Bliley Act, Information Privacy, and the Limits of Default Rules, *The Minn. L. Rev.*, 86, 1219.
47. Jensen C and Potts C. (2004). Privacy Policies as Decision-making Tools: an Evaluation of Online Privacy Notices. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM Press, New York, NY, 471-478.
48. Jha AK, Chan DC, Ridgway AB, Franz C, and Bates DW. (2009). Improving Safety And Eliminating Redundant Tests: Cutting Costs In U.S. Hospitals. *Health Affairs*, 28:5, 1475-1484.
49. John L, Acquisti A, and Loewenstein G. (2011). Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of Consumer Research*. Volume 37, Issue: 5, 858-873.
50. Johnson, E., & Goldstein, D. (2003). Do defaults save lives?. *science*, 302, 1338-1339.
51. Joinson AN, Woodley A, and Reips UD. (2007). Personalization, Authentication and Self-disclosure in Self-administered Internet Surveys. *Computers in Human Behavior*. 23:275-285.
52. Kahneman D and Tversky A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*. Volume 47, Issue: 2, 263-291.
53. Kahneman D, Knetsch, JL., and Thaler, RH (1990). Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of political Economy*, 98:6, 1325-1348.
54. Kelley PG, Bresee J, Cranor LF, and Reeder RW. (2009). A "Nutrition Label" for Privacy. *SOUPS '09: Proceedings of the 5th Symposium on Usable Privacy and Security*.
55. Krazit T. (2010). Google settles Buzz lawsuit for \$8.5M. CNET. http://news.cnet.com/8301-30684_3-20015620-265.html.
56. Kumar, R., Novak, J., & Tomkins, A. (2010). Structure and evolution of online social networks. *Link Mining: Models, Algorithms, and Applications*, 337-357.
57. Lai Y and Hui K. (2006). Internet opt-in and opt-out: investigating the roles of frames, defaults and privacy concerns. *CPR* 2006: 253-263.
58. Lenard TM and Rubin PH. (2005). Slow Down on Data Security Legislation. *Progress Snapshot 1.9*. The Progress & Freedom Foundation, August 2005.
59. Lenard TM and Rubin PH. (2006). Much Ado about Notification. *Regulation*, Vol. 29, No. 1, pp. 44-50.
60. Lessig, L. (2002). *The future of ideas: The fate of the commons in a connected world*. Vintage.
61. 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.

62. Liang KY and Zeger SL. (1986). Longitudinal Data Analysis Using Generalized Linear Models. *Biometrika*. Volume:73, Issue:1, 13-22.
63. Liang KY and Zeger SL. (1986). Longitudinal Data Analysis Using Generalized Linear Models. *Biometrika*. Volume:73, Issue:1, 13-22.
64. 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.
65. Luo Y. (2004). Building a Strong Foothold in an Emerging Market: A Link Between Resource Commitment and Environment Conditions. *Journal of Management Studies* 41:5.
66. McDonald A and Cranor L. (2009). The Cost of Reading Privacy Policies. *I/S: A J. Law and Policy Inform. Soc.* Volume 4, Issue:3, 543-568.
67. McDonald C. (2009). Protecting Patients In Health Information Exchange: A Defense Of The HIPAA Privacy Rule. *Health Affairs*, 28:2, 447-449.
68. McGraw D, Dempsey JX, Harris L, and Goldman J. (2009). Privacy As An Enabler, Not An Impediment: Building Trust Into Health Information Exchange. *Health Affairs* 28:2, 416-427.
69. Miliard M. (2010). ACLU brings suit against Rhode Island HIE. *Healthcare IT News*. December 01, 2010. <http://www.healthcareitnews.com/news/aclu-brings-suit-against-rhode-island-hie-0>.
70. Miller A and Tucker C. (2009). Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records. *Management Science*, 55:7, 1077-1093.
71. Miller, A. R. (1971). *The Assault on Privacy: Computers, Data Banks and Dossiers*. Ann Arbor: University of Michigan Press.
72. Milne, G. R., & Rohm, A. J. (2000). Consumer privacy and name removal across direct marketing channels: exploring opt-in and opt-out alternatives. *Journal of Public Policy & Marketing*, 238-249.
73. Mulligan D and Goldman J. *The Limits and the Necessity of Self-Regulation: The Case for Both*. U.S. Dept. of Commerce, Privacy and Self-Regulation in the Information Age, 1997.
74. Nahra KJ. (2008). HIPAA Security Enforcement is Here. *IEEE Security and Privacy*. 1540-7993.
75. National eHealth Collaborative. (2011). *Secrets of HIE Success Revealed: Lessons from the Leaders*.
<http://www.nationalehealth.org/ckfinder/userfiles/files/REPORT%20SecretsofHIESuccessRevealed.pdf>.
76. Oppenheimer DM, Meyvis T, and Davidenko N. (2009). Instructional Manipulation Checks: Detecting Satisficing to Increase Statistical Power. *Journal of Experimental Social Psychology*. Volume 45, Issue 4, 867-872.
77. Organisation for Economic Cooperation and Development (OECD), *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data* (Sep. 23, 1980).
78. Peterson, R. A., Balasubramanian, S., & Bronnenberg, B. J. (1997). Exploring the implications of the Internet for consumer marketing. *Journal of the Academy of Marketing Science*, 25(4), 329-346.
79. Phelps J, Nowak G, and Ferrell E. (2000). Privacy Concerns and Consumer Willingness to Provide Personal Information. *Journal of Public Policy & Marketing*. 19:1, 27-41.
80. Posner R. (1981). The Economics of Privacy. *The American Economic Review*, Vol. 71, No. 2. *Papers and Proceedings of the Ninety-Third Annual Meeting of the American Economic Association*, pp. 405-409.
81. Pritts J, Lewis S, Jacobson R, Lucia K, and Kayne K. (2009). *Privacy and Security Solutions for Interoperable Health Information Exchange: Report on State Law Requirements for Patient Permission to Disclose Health Information*. Office of Policy and Research: Office of the National Coordinator for Health IT.

82. Pritts, J, Choy A, Emmart L, and Hustead, J. (2002). The State of Health Privacy: A Survey of State Health Privacy Statutes. Georgetown University, Washington, DC, 2002.
83. Reitman R. (2012). FTC Final Privacy Report Draws a Map to Meaningful Privacy Protection in the Online World. Electronic Frontier Foundation <https://www.eff.org/deeplinks/2012/03/ftc-final-privacy-report-draws-map-meaningful-privacy-protection-online-world>
84. Romanosky S, Telang R, and Acquisti A. (2011). Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management*, 30:2, 256-286.
85. Sankar P, Moran S, Merz JF, and Jones NL. (2003). Patient perspectives on medical confidentiality: a review of the literature. *Journal of General Internal Medicine*, 18:659–669.
86. Santalesa R. (2011). What's Next for the FTC's Proposed Privacy Framework? Information Law Group. <http://www.infolawgroup.com/2011/03/articles/data-privacy-law-or-regulation/whats-next-for-the-ftcs-proposed-privacy-framework>.
87. Schwartz, P. M. (2005). Privacy Inalienability and the Regulation of Spyware. *Berkeley Tech. LJ*, 20, 1269.
88. Simon HA. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*. Volume 69, Issue:1, 99–118.
89. Simon S, Evans JS, Benjamin A, Delano D, and Bates DW. (2009). Patients Attitudes Toward Electronic Health Information Exchange: Qualitative Study. *Journal of Medical Internet Research*, 11:3, e30.
90. Singer E, Hippler H, and Schwarz N. (1992). Confidentiality Assurances in Surveys: Reassurance or Threat? *International Journal of Public Opinion Research*. 4:3, 256-268.
91. Singer, N. (2012). Mapping, and Sharing, the Consumer Genome. *New York Times*. http://www.nytimes.com/2012/06/17/technology/axiom-the-quiet-giant-of-consumer-database-marketing.html?pagewanted=1&_r=1&ref=natashasinger&
92. Smith RE and Snyder KD. (2002). Compilation of Federal and State Privacy Laws. *Privacy Journal*.
93. Solove DJ. (2004). *The Digital Person: Technology and Privacy in the Information Age*. New York: New York University Press.
94. Solove, D. (2013). Privacy Self-Management and the Consent Paradox. *Harvard Law Review*, 126
95. Spencer, S. J., Steele, C. M., & Quinn, D. M. (1999). Stereotype threat and women's math performance. *Journal of Experimental Social Psychology*, 35(1), 4-28.
96. Steel E and Vascellaro J. (2010). Facebook, MySpace Confront Privacy Loophole. May 21, 2010. http://online.wsj.com/article/SB10001424052748704513104575256701215465596.html?mod=W_SJ_business_whatsNews.
97. Stigler GJ. (1980). An Introduction to Privacy in Economics and Politics. *Journal of Legal Studies*. Volume 9, 623–44.
98. 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.
99. Tang Z, Hu Y, and Smith M. (2007). Gaining Trust Through Online Privacy Protection: Self-Regulation, Mandatory Standards, or Caveat Emptor. Heinz Research. Paper 49.
100. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
101. The White House. (2012). Consumer Data Privacy in a Networked World. <http://www.whitehouse.gov/sites/default/files/privacy-final.pdf>.
102. Tsai J, Egelman S, Cranor L, and Acquisti A. (2011). The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research*. Vol. 22, Issue:2, 254–268.

103. Vest JR and Gamm LD. (2010). Health information exchange: persistent challenges and new strategies. JAMIA,17:288-294.
104. Wakefield A and Fleming J. 2009. The Sage International Dictionary of Policing. Sage Publications, London.
105. Walker J, Pan E, Johnston D, Adler-Milstein J, Bates DW, Middleton B. (2005). The Value of Health Care Information Exchange and Interoperability. Health Affairs, 24, w10-w18.
106. Ware, W. H. (1973). *Records, computers and the rights of citizens*. Rand.
107. Weisband S and Kiesler S. (1996). Self-disclosure on Computer Forms: Meta-analysis and Implications. Proceedings of the SIGCHI Conference on Human factors in computing systems, 3-10.
108. Westin, A. R. (1967). Privacy and Freedom. New York Atheneum.

V. APPENDICES

A. Experiment 1: Notices

Privacy Notice	Notice Text
High Protection	<i>The analysis for this study requires that your responses are stored using your email. As such, your responses to the following set of questions may be directly linked back to you.</i>
Low Protection	<i>The analysis for this study requires that your responses are stored using a randomly assigned ID. All other information that could potentially be used to identify you (email, zip code, etc.) will be stored separately from your responses. As such, your responses to the following set of questions cannot be directly linked back to you.</i>

B. Experiment 1: Questions (Highly Intrusive in Bold)

Survey 1 Questions:

1. Have you ever downloaded a pirated song from the internet?
2. While in a relationship, have you ever flirted with somebody other than your partner?
3. **Have you ever masturbated at work or in a public restroom?**
4. **Have you ever fantasized about having violent non consensual sex with someone?**
5. Have you ever tried to gain access to some else's (e.g. a partner, friend, or colleague's) email account?
6. Have you ever looked at pornographic material?

Survey 2 Questions:

7. **Have you ever used drugs of any kind (e.g. weed, heroine, crack)?**
8. **Have you ever let a friend drive after you thought he or she had had too much to drink?**
9. **Have you ever made up a serious excuse, such as grave illness or death in the family, to get out of doing something?**
10. **Have you ever had sex in a public venue (e.g. restroom of a club, airplane)?**
11. **Have you ever while an adult, had sexual desires for a minor?**
12. **Have you ever had a fantasy of doing something terrible (e.g. torture) to someone**

C. Experiment 1: Attention Checks and Study Design

Design 1 and Attention Check:

In the instructions participants were instructed to skip the question. Answering this question would result in an automatic end of the study.

Getting meaningful and useful responses from participants in a study depends on a number of important factors. Thus, we are interested in knowing certain things about you. Specifically, we are interested in seeing whether you take the time to read survey directions and questions carefully prior to providing an answer. So in order to demonstrate that you have read these instructions carefully, please ignore the question below and click the next button without providing an answer. Thank you for your cooperation and participation in this study.

***What is your favorite sport?**

- ☐ Football
- ☐ Soccer
- ☐ Tennis
- ☐ Rugby
- ☐ Don't Play Sports

0% 100%

NEXT

Design 2 and Attention Check:

Getting meaningful and useful responses from participants in a study depends on a number of important factors. Thus, we are interested in knowing certain things about you. Specifically, we are interested in seeing whether you take the time to read survey directions and questions carefully prior to providing an answer. So in order to demonstrate that you have read these instructions carefully, please ignore the question below and click the next button without providing an answer. Thank you for your cooperation and participation in this study.

***What is your favorite sport?**

- ☐ Football
- ☐ Soccer
- ☐ Tennis
- ☐ Rugby
- ☐ Don't Play Sports

0% 100%

Next

Survey Powered By Qualtrics

D. Experiment 2: Notices

Privacy Notice	Notice Text
Students Only	<i>The information you provide will appear on a profile that will be automatically created for you. The profile will be published on a new university networking website, which will only be accessible by university students.</i>
Students & Faculty	<i>The information you provide will appear on a profile that will be automatically created for you. The profile will be published on a new university networking website, which will only be accessible by university faculty and students.</i>

E. Experiment 2: Questions

1. First name
2. Last name
3. Gender
4. Date of birth (MM/DD/YY)
5. Age in years
6. Country of birth
7. Email address
8. Home address
9. Phone number
10. Is your family in Pittsburgh?
11. How often do you see your family?
12. Are you single or married?
13. Do you have a girlfriend/boyfriend?
14. Where do you live?
15. Have you ever had troubles with your roommates?
16. Would you like to move somewhere else?

17. What program are you in? (e.g.: Undergrad Psychology, Grad Math)
18. Which courses are you taking at the moment?
- 19. What was your least favorite class at this university?**
- 20. What was your favorite class at this university?**
- 21. Who was your least favorite professor?**
- 22. Who was your favorite professor?**
- 23. In your experience, which department at this university has the least likable faculty?**
- 24. In your experience, which department at this university has the most likable faculty?**
- 25. Have you ever seen someone cheating?**
- 26. If so, did you inform the instructor?**
27. How many hours a day do you spend studying?
28. Are you working at the same time?
29. Do you receive financial aid from this university or some other non-profit organization?
30. Have you ever attended academic support programs (e.g. Peer Tutoring, Supplemental Instruction) in order to increase your understanding or your grades in a certain subject?
31. Are you a member of any group/community/fraternity/sorority?
32. If so, which group or groups are you a member of?
33. Do you have a Facebook profile?
34. Do you socialize/hang out at a bar at least once a month?
35. Do you have an alcoholic drink at least once a week?
36. In the last three months, have you done any volunteer service?
37. In the last three months, have you made a donation to a Non-Profit Organization?

F. Experiment 2: Misdirections

Delay:

This page takes approximately 15 seconds to load, we appreciate
your patience!



Department Information Pages:

A student activities planning committee, consisting of only CMU students, is requesting access to student profiles in order to better plan this year's upcoming activities. As such, they will have access to your profile.

Student Planning Committee:

We will also be creating CMU college-specific pages that will post relevant activities/lectures occurring in that school. Please select any schools below that you belong to and/or wish to be kept up-to-date on their activities.

[List of CMU Schools and Colleges]

Student Planning Committee + Choice:

A student activities planning committee, consisting of only CMU students, is requesting access to student profiles in order to better plan this year's upcoming activities.

*Do you wish to provide them access to your profile?

G. Full List of Settings Presented to Participants

*The settings are ranking in increasing importance to participants

Setting Number	Description	Condition
1	Allow you responses to be used in academic publications.	Low Importance
2	Allow research assistants (these are students that aid in research but are not faculty or PhD candidates) to access your responses.	Low Importance
3	Allow my responses to be shared with various think tanks that focus on ethics.	Low Importance
4	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 Importance
7	Only store my responses only on an encrypted drive	High Importance
8	Only store my responses only on a password-protect drive	High Importance
9	Allow my responses to be published on a research bulletin openly available on the internet.	High Importance
10	Allow my responses to be shared with religious organizations interested in evaluating personal ethics	High Importance
11	Allow my responses to be shown to other participants of this study.	High Importance

H. Questions on Ethical Behavior

*Questions ranked as very intrusive are in bold

1. **Have you ever used drugs of any kind (e.g. weed, heroine, 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 \$50?**
5. While in a relationship, have you ever flirted with somebody other than your partner?
6. **Have you ever looked at pornographic material?**
7. **Have you ever had a fantasy of doing something terrible (e.g. torture) to someone**
8. Have you ever downloaded a pirated song from the internet?