

The Economics of Personal Data and the Economics of Privacy

Alessandro Acquisti*

November 24, 2010

DRAFT - PRELIMINARY AND INCOMPLETE

Abstract

We provide an overview of the economic analysis of the protection and revelation of personal data. In particular, we (1) describe the evolution of the economic theory of privacy, (2) examine privacy-related trade-offs for data subjects and data holders, and (3) highlight the current economic debate on privacy protection.

*Heinz College, Carnegie Mellon University, Pittsburgh, PA; acquisti@andrew.cmu.edu.

Contents

1	Introduction	3
1.1	The Boundaries of the Economics of Privacy	4
2	The Economic Theory of Privacy	5
2.1	Privacy as Source of Economic Inefficiencies	5
2.2	Critiques of the Chicago School Arguments	7
3	Benefits and Costs of Disclosed and Protected Data	10
3.1	Benefits and Positive Externalities from Disclosed Data	11
3.1.1	Data Holders	12
3.1.2	Data Subjects	16
3.2	Costs and Negative Externalities of Disclosed Data	18
3.2.1	Data Holders	19
3.2.2	Data Subjects	24
4	The Current Debate	32
4.1	Regulation vs. Self-Regulation	32
4.2	Consumers' Valuations of Privacy	35
4.3	Hurdles in Consumer Behavior and Privacy Nudges	37
4.4	The Promise of Privacy Enhancing Technology	40
5	Conclusion: Should We Change the Frame of the Debate?	41

1 Introduction

In modern information economies, the reduction of the cost of storing information has made it possible to capture, save, and analyze increasing amounts of information about the individual. Companies record details of each customer transaction. Websites log their visitors' behavior. Data aggregators link information coming from different sources to compose individual profiles.

The more organizations and individuals embrace digital technologies, the cheaper and faster become the production and processing of personal, and potentially sensitive, data. Thus, privacy concerns grow as well. Several everyday activities can be tracked through information technology. Small pieces of personal data enter databases, whose records may be linked and tracked to form a complete dossier of a person's life. This may happen without the person's consent or even knowledge. In addition to that, hundreds of millions of individuals worldwide have now embraced Web 2.0 technologies (such as blogs and online social networks), through which they willingly broadcast sometimes highly personal information to friends and strangers alike.

Ultimately, the economic consequences of information sharing for all parties involved (the data subject and the actual or potential data holder) can be welfare enhancing or diminishing. In choosing the balance between sharing or hiding one's personal information (and in choosing the balance between exploiting or protecting individuals' data), both individuals and organizations face complex, sometimes intangibles, and often ambiguous trade-offs. Individuals want to protect the security of their data and avoid the misuse of information they pass to other entities. However, they also benefit from sharing with peers and third parties information that makes mutually satisfactory interactions possible. Organizations want to know more about the parties they interact with, tracking them across transactions. Yet, they do not want to alienate those parties with policies that may be deemed too invasive.

But trade-offs are the natural realm of economics. Therefore, economics can help us

understand how individuals and organizations make decisions about the protection and usage of individuals' data, and what are the consequences of those decisions.

In this document, we report on the economic implications of the protection and revelation of personal data. In particular, we present the evolution of the economic theory of privacy (Section 2), we examine current privacy-related trade-offs for data subjects and data holders (Section 3), and we highlight the current economic debate on privacy protection (Section 4).

1.1 The Boundaries of the Economics of Privacy

Before commencing our analysis, we alert the reader of the boundaries inherent to an economic approach to the privacy debate. First of all, in the rest of the document our focus will be, primarily, on *information privacy* - that is, on the issues associated with the collection and usage of individuals' personal information (Westin, 1970). We take this approach because, while privacy is a multi-faceted concept, most of the relevant contemporary economic research focuses on *consumers' data*. Our focus on informational privacy and consumer data, however, should not be construed as a denial of the existence of other dimensions to the privacy debate, which may be more difficult to capture in economic terms (Solove (2006), for instance, distinguishes between privacy issues associated with the collection, processing, or dissemination of personal data, and privacy issues associated with personal *invasions*, such intrusion and decisional interference).

Second, the existence of such trade-offs does not imply that the economic agents are always *aware* of them as they take decisions that will impact their privacy.

Third, the analysis of trade-offs associated with the protection or revelation of individuals' data does not presume that all privacy trade-offs have an explicit monetary dimension. Rather, the economics of privacy tries to understand trade-offs associated with the balancing of one's public and private spheres. Even a broader definition of privacy than the one we use in this report (for instance: a person's interest to keep certain activities, interests, or

thoughts to herself), can still be given an economic dimension: such interest protects the individual's psychological and physical well-being; psychological and physical well-being, in turn, can be interpreted in economic terms as sources of individual utility.

Fourth, there may be privacy dimensions that affect individuals' well-being and are not merely intangible, but in fact immeasurable: for instance, whereas the US legislator has taken an utilitarian approach to data protection, the European legislator has tended to define privacy as a fundamental human right. As Samuelson (2000) notes, those who conceive of personal data protection as a fundamental civil liberty, see it as an interest essential to "individual autonomy, dignity, and freedom in a democratic civil society," independently of the economic considerations we discuss in this report.

2 The Economic Theory of Privacy

In this section we highlight recent economic theories of privacy. We distinguish between theories that stress the welfare-diminishing impact of interrupting the flow of personal information, and studies that arrive at opposite conclusions.

2.1 Privacy as Source of Economic Inefficiencies

Economists have been writing about privacy since, at least, the 1970s. Within the neo-classical economic theory of perfectly competitive markets, "complete" information (the availability of relevant information to all market participants) leads to economic efficiency: for instance, when all consumers know the prices at which every firm is selling its product, competition will drive those prices down to the lowest possible level made possible by the production technology, and will increase consumers' welfare.

Consequently, according to Chicago School's scholar Posner (1978, 1981), the protection of privacy creates inefficiencies in the marketplace, since it conceals potentially relevant information from other economic agents. Consider a job seeker who misrepresents her back-

ground and expertise to an hiring firm: Protecting the applicant's personal information will negatively affect the firm's hiring decision. Therefore, the protection of the former's privacy comes at the cost of the latter's profitability. Hence, removing individuals' personal information from the marketplace through privacy regulation ultimately transfers the cost of that person's possible negative traits on other market players.

Another Chicago School economist, Stigler (1980), believes that governmental interference in the market of personal information is destined, at best, to remain ineffective: since individuals have an interest in publicly disclosing favorable personal information and hiding negative traits, those who decide to protect their personal information (for instance, a debtor who does not want to reveal her credit history) are *de facto* signalling a negative trait. In this case, regulatory interventions blocking the flow of personal information would be redistributive and inefficient: economic resources and productive factors would end up being used inefficiently, or rewarded unfairly, because information about their quality has been removed from the marketplace.

More recently, Calzolari and Pavan (2006) find that the unrestricted sharing of consumers' personal data between two firms may in fact reduce market distortions and increase social welfare, including the consumers'.

Along similar lines, Varian (1996) observes that consumers may suffer privacy costs when too *little* personal information about them is being shared with third parties, rather than too much. The consumer, Varian notes, may rationally want certain information about herself known to other parties: for instance, a consumer may want her vacation preferences to be known by telemarketers, in order to receive from them offers and deals she may be actually interested in.

Also building upon a Chicago School's argument (the so-called Coase theorem), Noam (1997) argues that whether or not a consumer's data will remain protected does not depend on the initial allocation of rights on personal information protection (such as whether or

not that consumer's data is protected by law). Instead, whether data will eventually get disclosed or protected ultimately depends on the relative valuations of the parties interested in that data. If the consumer values her privacy more than the data marketing firm values acquiring that consumer's data, the data *will* remain protected, because – even in absence of a law regulating that protection – the consumer would willingly pay for the right to protect her data.

2.2 Critiques of the Chicago School Arguments

Not all economists, however, have taken the stance that privacy protection inherently causes market inefficiencies, or that consumers who value privacy can simply secure it in the marketplace (Murphy, 1996). Hirshleifer (1980), for instance, criticizing Posner and Stigler's positions on privacy, notes that the assumptions of rational behavior underlying the Chicago School's privacy models fail to capture the complexity of consumers' privacy decision making.

In fact, while the early Chicago School studies of privacy originated in what may be defined a pre-ICT (modern Information and Communication Technologies) era, the development of new information technologies, and Internet in particular, led researchers to formulate more nuanced and granular views of the trade-offs associated with privacy protection and data sharing.

Varian (1996), for instance, notes that the secondary usage of personal data raises particular economic concerns: a consumer may rationally decide to share personal information with a firm because she expects to receive a net benefit from that transaction; however, she has little knowledge or control upon how the firm will later use that data. The firm may sell the consumer's data to third parties at profit, but the consumer may not share any of that profit, or may even bear a cost when the third party abuses her data (for instance, for spam, adverse price discrimination, and so forth; see Odlyzko (2003)). Such negative externality on the consumer is not internalized by the firm (Swire and Litan, 1998). Noam (1997) also

acknowledges that transaction costs, poverty, and other hurdles may not allow consumers to acquire privacy protection under standard market conditions.

Hermalin and Katz (2006) criticize the Chicago School's argument that privacy protection is inherently welfare-diminishing. They note that data protection may have *ex ante* positive effects on economic welfare. For instance, the protection of privacy can make it possible to support insurance schemes that otherwise would not exist. If all potential policy holders had to be tested for potentially fatal health condition, life insurance companies would adjust insurance prices according to the results of those tests. While the outcome would be *ex post* economically efficient (consumers would purchase insurances at actuarially fair rates), from *ex ante* the individual would bear the risks associated with the outcomes of their test results. However, if testing were banned, "then the competitive equilibrium would entail all risk-averse individuals buying full insurance at a common rate." Therefore, "[w]elfare would be greater than under the testing equilibrium both because the (socially wasteful) costs of testing would be avoided and because risk-averse individuals would bear less risk" (Hermalin and Katz, 2006, p. 6). Furthermore, Hermalin and Katz (2006) note that markets may fail to adjust efficiently to additional information, lowering the efficiency of the resulting equilibrium. In their model, two rational agents engage in a transaction in which both are interested in collecting information about the other; privacy protection may actually lead to efficient allocation equilibria, and explicit prohibition of information transmission may be necessary for economic efficiency (as the mere allocation of informational property rights may not suffice).

Similarly, models by Hirshleifer (1971) and Taylor (2003) show that rational economic agents may end up inefficiently over-investing in collecting personal information about other parties (for instance, in order to increase private revenues from sales based on knowledge of the buyer's willingness to pay). Taylor (2004) also finds that, in the presence of tracking technologies that allow merchants to infer consumers' preferences (and then engage in price

discrimination), whether or not the presence of privacy regulatory protection enhances consumer and aggregate welfare depends on the consumers' sophistication. Naive consumers do not anticipate the seller's ability to use past consumer information for price discrimination; therefore, in equilibrium all their surplus is taken away by the firms, unless privacy protection is enforced through regulation. Regulation, however, would not be necessary *if* consumers were aware of how merchants will exploit their data, and strategic enough to adapt their behavior accordingly. We discuss below (Section 4) why there are reasons to believe that consumers, in fact, act myopically when trading off the short term benefits and long term costs of information revelation and privacy invasions.

Similar conclusions are reached by Acquisti and Varian (2005), who study a two-period model in which merchants have access to "tracking" technologies and consumers have access to "hiding" technologies. Internet commerce offers an example: merchants can use cookies to track consumer behavior (in particular, past purchases), and consumers have access to "anonymizing" technologies (deleting cookies, using anonymous browsing or payment tools) that hide that behavior. Consumer tracking will enhance the merchant's profits only if the tracking is also used to provide consumers with enhanced, personalized services.

Other models, in addition to the privacy costs associated with price discrimination and the social welfare implications of sharing of consumer data with third parties, find that the exploitation of personal information for unsolicited marketing can constitute a negative consumer externality (Hui and Png, 2003).

The majority of the theoretical economic work on privacy takes a micro-economic perspective (see also Hui and Png (2005)). However, as we shall discuss further below (Section 3), significant *macro*-economic costs and benefits arise from the protection or trade of individual information. Furthermore, the possibility that Privacy Enhancing Technologies (or PETs) may lead to non-zero sum market outcomes only recently has started being discussed in economic research: the usage of PETs may allow certain personal information to be shared

while other is protected, with common satisfaction of both data subjects and data holders (see Acquisti (2008) and Section 4.4).

3 Benefits and Costs of Disclosed and Protected Data

In this section, we discuss the economic value of personal data and personal privacy by analyzing the individual and social costs and benefits associated with disclosed and protected.

In the context of our analysis, data subjects are consumers, and data holders are firms. We will frame the analysis by presenting the market for personal data and the market for privacy as two sides of a same coin, wherein protected data may carry benefits and costs that are dual, or specular, to the costs and benefits associated with disclosed data for both data subjects and data holders. However, we do not attempt to provide a complete list and exhaustive taxonomy of all the possible types of costs and benefits associated with protected and disclosed data.

By *disclosed* data, we refer, rather loosely, to states in which the data subject may have knowingly or unknowingly shared data with other parties (the data holders), or states in which other parties may have entered in possession of the subject's data, independently of her knowledge or even consent.¹ By *protected* data, we similarly refer to situations in which such disclosures have not take place, independently of whether this may be due to the data subject's intentional protection of personal information, or the potential data holder being unable, or uninterested in, accessing the latter.

Primarily, we are interested in costs and benefits trade-offs that arise as a *consequence* of data having been disclosed or protected. Secondly, we also consider the trade-offs

¹In other words, we use the term “disclosed data” to include situations where the data has been collected by the data subject even without data subject's explicit action. Therefore, “disclosed” refers to a state of the data (its being known to the other party), rather than to the act of disclosing. Since our goal is not to create a rigorous taxonomy, but rather highlight exemplary trade-offs, we will be somewhat loose about distinguishing the costs and benefits associated with the collection, processing, dissemination, or further usage of personal data.

associated with the actual *acts* of *disclosing* (or collecting) data and *protecting* (or not disclosing) data.

Our approach is *dual*: *disclosed* personal information (that is, in our terminology, the lack of data protection) can result in economic benefits for both data holders (savings, efficiency gains, surplus extraction, increased revenues through consumer tracking) and data subjects (personalization, targeted offers, and so forth); at the same time, such disclosures (that is, the *lack* of protection of personal data) can be costly for both firms (costs borne when that data is breached or misused, or collected in ways that consumers deem too intrusive) and consumers (identity theft, price discrimination, stigma or other psychological costs; see also Stone and Stone (1990)). Furthermore, the *act* of collecting data can be costly for data holders (such as the investments necessary to build Customer Relationship Management systems).

Similarly, protected data (that is, *lack* of data disclosure) can be associated with both benefits and costs for data subjects and potential data holders; such benefits and costs are often dual (i.e., the inverse) of the benefits and costs highlighted above: for instance, data subjects and data holders incur opportunity costs when useful data is not disclosed (for instance, they may miss out on opportunities for increased efficiency or increased convenience), although both parties may also benefit in various ways (consumers, for instance, by reducing the expect costs of future identity theft; firms, for instance, by exploiting privacy-friendly stances for competitive advantage). Furthermore, there are costs associated with the *act* of protecting data (investments necessary to encrypt data for the data holders to prevent *further* disclosures; costs of using Privacy Enhancing Technologies for the data subject).

In the rest of this section, we provide examples of some of these trade-offs for both data subjects and data holders.

3.1 Benefits and Positive Externalities from Disclosed Data

Our analysis starts with the economic benefits of disclosed data. We will focus on a) the potential benefits of disclosed data for both data holders and data subjects. However, we will also mention b) the opportunity costs that may be suffered when valuable information is not disclosed, as well as c) the costs of investments necessary to collect and process personal data.

3.1.1 Data Holders

The Benefits of Disclosed Data. In a prescient article published before the advent of the commercial Internet, Blattberg and Deighton (1991) wrote:

It's a marketer's dream - the ability to develop interactive relationships with individual customers. Technology, in the form of the database, is making this dream a reality. Now companies can keep track of customer preferences and tailor advertising and promotions to those needs. For instance, a grocery store system could note that you recently purchased a sample size of dishwashing detergent and could offer you a coupon to buy the large size.

What Blattberg and Deighton (1991) twenty years ago described as the future of interactive marketing in an age of addressability has, today, become reality. Online, the combination of IP addresses, cookies, click-stream data, and deep packet inspection makes it possible to create accurate pictures of consumers' demographic traits and behavior. Offline, credit reporting agencies and data aggregators purchase consumer data from private and public organizations, sanitize it, and combine it, in order to compile rich dossiers of consumers' information - credit and health histories, individual preferences, purchasing patterns - later sold (in both aggregate and individual forms) back to the public and private sectors. Combinations of online and offline *individual* data have also become possible - and

so has the tracking of online behavior across different websites or advertising networks, and the combination of online browsing and behavioral information together with self-disclosed personal information harvested from social media used by consumers. We live in a consumer data-driven and consumer data-focused commercial revolution, in which individuals are at the same time consumers and producers of a most valuable asset: their personal information.

Firms can significantly benefit from the ability to learn so much about their current, or potential, customers. Rich datasets of consumers can improve firms' marketing capabilities, boosting their ability to address specific target markets or customers, and lowering their advertising costs (Blattberg and Deighton, 1991). Firms can therefore increase revenues through targeted offers (Acquisti and Varian, 2005), innovative coupon strategies such (consider, for instance, the recent success of initiatives such as **Groupon.com** (Pitta, 2010)), and improved CRM (Richards and Jones, 2008), as well as increased consumer loyalty (consumers' switching costs increase when a firm is able to use her information for personalized services; Ball et al. (2006)).

By analyzing large amounts of consumer data, firms are able to predict aggregate trends (such as variations in consumer demand) as well as individuals' preferences (Linden et al., 2003), thus minimizing inventory risks and maximizing returns on marketing investment. They can improve their ability to offer useful recommendations to consumers (Bennett and Lanning, 2007), as well as their ability to enforce profit-enhancing price discrimination (Varian, 1985). Furthermore, by observing individual behavior, firms can learn how to improve their services, or re-design it in order to take advantage of the observed behavior.

An example of how consumer information can be leveraged for higher profit is online advertising. E-commerce and online advertising now amount to \$300 billion per year in the US, providing employment to 3.1 million Americans (Deighton and Quelch, 2009). More than their offline counterparts, online ads can be targeted at each individual based on her online behavior (such as her searches, sites visited, clickstream data on a given site) and

inferences made through that data. Such targetability implies that firms reduce the cost of ads wasted on consumers unlikely to be receptive to them. Furthermore, since online ad exposure, click-through behavior, and sometimes even post-exposure online behavior are often measurable, advertisers can monitor and improve the effectiveness of online advertising more than in other marketing channels. Primarily, this allows higher revenues for marketers and merchants (the price of behaviorally targeted advertising is almost 3 times as much the price of untargeted advertising: see Beales (2010)). Secondarily, this can also benefit the consumer: Targeted advertising may give consumers useful information, since the ads are tailored to consumers' interests. Hence, such targeting may reduce the producers' cost of communicating with consumers, and the consumers' cost of obtaining useful information (Lenard and Rubin, 2009; Goldfarb and Tucker, 2010). In turn, revenues from targeted and untargeted advertising may support new services and business models, free content, or low-cost products - benefitting both consumers and firms.

According to Rubin and Lenard (2001), the credit reporting industry offers another example of how the collection and analysis of flows of consumer data can be welfare enhancing. Rubin and Lenard argue that information collected, analyzed, and then resold by credit reporting agencies is used to allocate credit efficiently among potential borrowers - therefore providing value added to the marketplace as well as to consumers themselves.

Organizations also benefit indirectly from consumer data by selling it to other firms. This may be the case even for firms whose primary product is not consumers data, but which nevertheless find in their customers' data a tradable asset of interest to other organizations. It is most naturally the case, however, of Web 2.0 enterprises (such as, for instance, online social networks): for such firms, consumers' data is the primary asset, and therefore their users become, in effect, the product. (The actual customers consist of marketers, advertisers, and data aggregators interested in the behavioral and user-disclosed data generated on the platform.)

The aggregation of individual consumers' data may benefit firms even when the data is not personally identified. Firms may benefit from inferring consumer trends based on the combined analysis of the behavior of many individual agents. Companies such as comScore, for instance, analyze web trends by combining behavioral and survey observations of million of online consumers, and then provide to their clients data which can be used for competitive intelligence, market testing, and segmentation analysis.

The Costs of Undisclosed Data. Conversely, opportunity costs and inefficiencies may arise when potentially welfare-enhancing data disclosures do not take place. For instance, firms without access to consumer data may face significant barriers to entry and competitive disadvantage against firms with larger customer bases, thus limiting competition. Or, mandatory opt-in privacy policies for certain types of data may be costly for firms, when they result in the loss of valuable data (Staten and Cate, 2003). Furthermore, lack of consumer data may make it harder for firms to innovate and offer new services. For the same reason, uncertainty about (or fear of) possible legal reprisals following the collection or processing of consumers data may hinder product innovation.

Similarly, costs of undisclosed data may be suffered by society at large. During the summer of 2010, for instance, the Canadian Ministry of Industry announced that the long-form Census questionnaire would no longer be mandatory. The initiative was motivated by the Government's stance that Canadians "should [not] be forced, under threat of fines, jail, or both, to divulge extensive private and personal information" (even though the Census data is actually never *released* to parties outside Statistics Canada in identifiable form). The transition from compulsory to voluntary, however, could likely result in a drastic decline of total respondents to the long-form questionnaire. The subsequent increase in the risk of non-response bias could, in turn, negatively affect the work of policy makers, researchers, or healthcare providers.²

²"StatsCan head quits over census dispute," CBC News, Wednesday, July 21, 2010.

The Costs of Collecting Data. The benefits from disclosed data we highlighted in this section must be weighted against the cost of the investments necessary to collect and process that data. These costs are economically justifiable when firms expect to gain larger advantages from the analysis of consumer data, while avoiding the costs that may ensue from its misuses. The costs of data gathering and storage have been constantly decreasing thanks to the evolution of ICTs. However, implementing systems that make efficient use of that data is not trivial. As an example, the impact of customer relationship management (CRM) on firm performance remains a debated topic; Krasnikov et al. (2009) find that CRM implementation is associated with an increase in profit efficiency, but a *decline* in cost efficiency.

3.1.2 Data Subjects

The Benefits of Disclosed Data. Data subjects can directly benefit from sharing personal information with firms. A customer might receive immediate monetary compensation for revealing her personal data (e.g., discounts), or she might receive intangible benefits (for example, personalization and customization of information content). In certain cases, the individual might also benefit from her data being given to third parties in the form of improved services, targeted offers, or *less* junk mail (under the assumption that the information provided will, in fact, be used by marketers to screen offers to be sent to consumers: see Varian (1996)). Accordingly, some economists have also proposed a “proptertization” of privacy (see Varian (1996); Laudon (1996); Varian and Shapiro (1997)) where the individual literally sells her own personal information into a marketplace, or attempts to buy back the right to keep that data private.

Better marketing information in the hands of companies may also benefit customers and society in general in an indirect way, by way of positive externalities. For example, better consumer data can allow firms to bring to the market niche products that, without focused

data about potentially interested consumers, might have been too risky to develop (Blattberg and Deighton, 1991).

Prices might, sometimes, be *reduced* as an effect of more targeted (and less wasteful) advertising and marketing. The social waste of efforts spent in building customers data on partial and erroneous information might be reduced in the presence of a well established market for personal data (Laudon, 1996). The proper combination of sharing and hiding different pieces of information could therefore help both firms and consumers, reducing junk and telemarketing on the one hand, and increasing the reliability of gathered data on the other hand.

Furthermore, bargain-hunting consumers may benefit from information-based price discrimination, in the sense that they may be able to acquire goods at lower prices: under certain conditions, microeconomic theory predicts that those consumers benefit from price discrimination if they get offered goods that may not have been produced (or offered to them) in absence of a premium paid by higher-valuation consumers.

Online advertising - and in particular targeted ads - may both inform consumers (providing them better information at a lower search cost), as well as allow other services (for instance, news) to be provided for free to the consumers. Such ads may also be visually less intrusive than non-targeted ads (Goldfarb and Tucker, 2010).

The existence of a secondary market for customer data may also be a potential source of positive externalities for consumers. Such externalities may arise when, for instance, data provided to a website makes the service more convenient or efficient on another site, precisely because of sharing of data between different services (for instance, Facebook Connect enables seamless authentication on third-party Web sites, reducing the user's cost of signing up across different platforms).³

The aggregation of consumers' data may produce other forms of positive externalities.

³See <http://developers.facebook.com/docs/guides/web>.

For instance, consumers may benefit from transactions involving their personal data in the form of easier access to insurance and credit (Jentzsch, 2001).

Furthermore, macro-economic benefits may materialize: the analysis and aggregation of the online behavior, sensor data, and individual decisions of a multitude of separate economic agents may allow the early identification of trends and patterns that would be otherwise hard or impossible to notice, or at least not within a limited period of time. This can benefit society as a whole: the monitoring and aggregation of web searches can allow the rapid identification of an infectious disease outbreak (Wilson and Brownstein, 2009); the combination of inputs from portable devices may be used for traffic and congestion control; data from remote and distributed sensors on consumers' machines may be used for environmental monitoring (Dereszynski and Dietterich, 2007).

As we discuss elsewhere in this report, however, one can argue that such benefits may be enjoyed by consumers without their having to disclose *personally identified* data: the adoption of privacy enhancing technologies can make it possible to satisfy both the need for privacy and the need for sharing data, by selectively protecting and disclosing pieces of personal information.

The Costs of Undisclosed Data. Conversely, some of the highlighted social benefits of disclosed data turn into opportunity costs when a consumer elects not to disclose that data, preventing said data from being used for socially useful purposes (consider, for instance, the case of the voluntary long-form Canadian Census questionnaire discussed above).

Such opportunity costs of undisclosed data become more significant at the individual level, too, as more consumer products and services are made conditional to, or rest on the assumption of, data being disclosed. Consider, for instance, a website that can only be browsed via a Facebook Connect authentication; those individuals who decide not to join the social network because of their privacy concerns will also miss out on the information contained in the third party website. Or, consider a social event that is only announced

through an online social network: the more one's friends (as well as other consumers) get comfortable with disclosing data online, the higher is the opportunity costs for those individuals who do not join a service in order to protect their data. We will further discuss this privacy "externality" in the following section.

3.2 Costs and Negative Externalities of Disclosed Data

In this section we examine the costs and negative externalities of disclosed data. We will focus on a) the costs of disclosed data and privacy intrusions. However, we will also mention b) the costs of *protecting* data and c) the benefits of *protected* data.

3.2.1 Data Holders

The Costs of Disclosed Data. Data holders can suffer tangible and intangible costs from disclosed data. Some of these costs may be associated with the mere collection of that data (for instance, when consumers deem a certain strategy of data gathering too intrusive). Other costs are associated with the actual use (and misuse) of collected data.

Online and offline companies have been punished by the market for data gathering behaviors that, while not necessarily illegal, were perceived as invasive of consumers' privacy. A notorious case was Amazon.com's dynamic price experiment in September 2000. An Amazon.com customer had purchased Julie Taymor's 'Titus' DVD for \$24.49. The following week, he found that the price on Amazon had risen to \$26.24. However, after deleting cookies and stripping his computer of the electronic tags "that identified him to Amazon as a regular customer, [he found that] the price fell to \$22.74" (Streitfield, 2000). As discussions of Amazon.com's price discriminating practices made their way from online discussion forums to national media outlets, Amazon.com suffered what may be arguably described as significant PR damage. The company had to reimburse customers who had paid premium prices for the DVDs and - through a spokesperson - swore off the practice of dynamic pricing, or price discrimination.

Amazon.com is but one in a long list of companies who have attracted consumers' negative attention because of data collection or processing practices deemed objectionable. Consider, for instance, Facebook's Beacon controversy,⁴ or Google's Buzz controversy.⁵ Following similar blunders, other firms have been imposed fines for violating their own privacy policies. For instance, Eli Lilly was required by the Federal Trade Commission to improve its security practices after it identified subscribers email addresses in an email about Prozac (*In re Eli Lilly*, 133 F.T.C. 763, 767, 2002). Microsoft was required to develop a comprehensive information security program - and have it certified every other year by an independent professional - for twenty years after violating its stated privacy policy for the .NET Passport service (*In re Microsoft Corp.*, 134 F.T.C. 709, 742, 2002).

Even longer is the list of companies that suffered costs following data breaches involving their consumers' or employees' data. Data breaches comprise different scenarios - from the mere loss of laptops containing consumers' data (which may or may not have been actually compromised by malicious parties), to the proven exposure of consumers' data following hackers' attacks. Breached organizations can end up paying fines, legal fees, and redress costs (Romanosky and Acquisti, 2009). Following one of the most publicized data breach events in 2005, Choicepoint - a consumer data aggregation company - paid more than \$26 million in fees and fines. Retail giant TJX reported losses above \$250 million after 45 million credit and debit card numbers were stolen from its information systems in 2005. After the theft of a laptop containing 26 million veterans' personal information in 2006, the Department of Veterans Affairs paid \$20 million to veterans and military personnel, even though no evidence conclusively proved that the information had been accessed by malicious third parties. The 2008's Heartland Payment Systems' breach, which affected more than 600 financial institutions (Heartland is one of the largest US credit card processing companies in

⁴Juan Carlos Perez, "Facebook's Beacon More Intrusive Than Previously Thought," PCWorld, December 1, 2007.

⁵"Google Buzz Has Serious Privacy Flaws," Fox News, January 12, 2010.

the United States), cost the company more than \$12 million in fines and fees.

Breached firms may incur significant costs due to consumer redress. Even though most consumer lawsuits against data breaching firms have been dismissed by US courts (Romanosky and Acquisti, 2009), consumers are often offered (or reimbursed the costs of) credit alerts and identity theft insurance services by the breached firms. In addition, the very act of notifying consumers of a breach can be costly: Forrester Research estimates the disclosure costs of unregulated firms at \$90 per consumer and for highly regulated firms at \$305 per consumer (Dignan, 2007).

Consumers *may* also punish firms that they perceive as not adequately protective of their data indirectly. A Ponemon Institute survey suggests that about one consumer out of five terminated their relationships with a company that compromised their data (Ponemon, 2009). The Ponemon report estimates that the costs of data breaches to US firms (combining investigations, legal fees, consumer redress, and actual lost business due to the breach), amounted at \$6.65M per breach in 2008. The amount has been steadily increasing for the past few years.

Privacy concerns may not just adversely affect consumers' *ex post* propensity to purchase (that is, after a merchant has violated a consumer's data). They may also *ex ante* reduce the likelihood that a consumer will engage in certain transactions, precisely because of fears of future privacy costs. In 1999, Forrester Research (1999) estimated that "[t]wo-thirds of online shoppers feel insecure about exchanging personal information over the Internet, affecting the amount of time and money consumers spend online." In 2000, opportunity costs in the order of billions of dollars due to missing sales were estimated by Federal Trade Commission (2000). In 2002, Jupiter Research forecasted that "\$24.5 billion in on-line sales will be lost by 2006 - up from \$5.5 billion in 2001 [because of privacy concerns]." In 2005, similar predictions were reached by Privacy & American Business (2005).

It is hard, however, to precisely estimate these effects. First of all, self-reported attitudes

and intentions of behavior may not necessarily match actual consumers' privacy decisions (see, e.g., Spiekermann et al. (2001)). For instance, self-reported individual claims of behavior (such as terminating relationships with an affected merchant) may not precisely predict consumers' actual actions: data breaches may hurt a firm's "image" without necessarily driving consumers away.⁶ Furthermore, the repeated exposure to privacy invasions (for instance, the increasing number of data breaches reported in recent years) may eventually desensitize consumers through a psychological process of habituation. Similarly, given merchants' ability to effectively obfuscate prices in online transactions (Daripa and Kapur, 2001), and the possibility of linking data across websites in manners that may be unimaginable for the average consumer, price discrimination practices similar to those adopted by Amazon.com in 2000 may go unnoticed, and therefore unpunished.⁷

This suggests that the actual cost that firms bear when they abuse consumer data is still open to investigation. For instance, while evidence exists that the stock market value of companies that suffer data breaches or other privacy blunders is negatively affected, such negative effects may be short lived: Acquisti et al. (2006) found a mean abnormal stock market return of -0.6% for affected companies traded in the NYSE during the day past the breach or intrusion event; however, the authors also found that such abnormal returns reverted to zero a few days after the event.

Similarly daunting is the task of setting the *optimal* level of corporate punishment for abuses of consumers' data. In recent months, a number of privacy regulatory bodies around the world have taken initiatives against companies involved in privacy blunders, and class action lawsuits have been filed against them. Following the failure to block a video showing an autistic boy being bullied by other students on YouTube.com, three Google executives

⁶Ellen Messmer, Data Breaches Hurt Corporate Image but Don't Necessarily Drive Customers Away, *Networked World*, Aug. 29, 2007.

⁷Naturally, consumers actually expect price discrimination to occur in several transactions, such as the purchase on a flight ticket.

were sentenced to six months in prison by an Italian Court (the sentences were suspended).⁸ Following Google’s gathering of private wi-fi data during its capture of images of streets, the UK’s Information Commissioner announced an investigation of the company and the possibility of fining it.⁹ Following its change of terms of service and users’ privacy settings without their assent, the Toronto-based Merchant Law Group filed a class action lawsuit against Facebook.¹⁰ Setting the appropriate level of punishment and liability in these cases is hard: should the punishment be proportional to the consumer harm (which may, itself, be hard to measure), or should be calibrated to create a disincentive to commit engage in similar behaviors in the future? Setting the punishment too high may impede growth and innovation; setting it too low may produce the perverse effect of legitimizing the invasive behavior, transforming it into a mere cost of doing business.

The Costs of Protecting Data. Protecting consumer data can be costly for firms in two senses. First, as noted in Section 3.1, firms may forego potentially lucrative data gathering, mining, and processing in order to avoid future privacy costs. This constitutes, in economic terms, an opportunity cost. Second, in an attempt to avoid *ex post* expected losses due to privacy debacles, firms may incur lower but certain *ex ante* costs. Firms may decide (or be forced by legislative initiatives, such as the Gramm-Leach-Bliley Act) to invest, and perhaps over-invest, in data security and protection: Hoofnagle (2007) reports qualitative findings suggesting that US firms have been increasing security and operational investments following the enactment of data breach disclosure laws.

Additional costs highlighted by Samuelson (2003) comprise the social losses due to “incoherent privacy policies:” amidst a complex array of legislative and self-regulatory initiatives,

⁸Adam Liptak, “When American and European Ideas of Privacy Collide,” *New York Times*, February 27, 2010.

⁹Peter Judge, “Google Could Get Massive UK Privacy Fine Over WiSpy,” *eWeekEurope.com*, October 25, 2010.

¹⁰Emma Woollacott, “Facebook gets more grief in Canada - Class action lawsuit launched,” *TechEye.net*, July 6, 2010.

both consumers and firms are uncertain about the level of protection afforded to, or required for, various types of personal data. This uncertainty is costly in itself, in that it forces data subjects and data holders to invest resources into learning about the admissibility of a given data practice. It also creates costly second order effects, in that it may lead both data subjects and data holders to inefficiently under- or over-invest in data protection.

Similar costs arise for Internet companies that operate worldwide and need to conform their services to differing local standards of privacy protection.¹¹

The Benefits of Protected Data. The issue of whether firms can gain competitive advantage from a pro-consumer privacy stance is still open to debate. While a firm that self-limits its collection of usage of consumer data may forego some of the benefits we have espoused in the Section 3.1, it may gain from limiting its liabilities and costs due to misused data, as well as from attracting privacy-savvy consumers. Whether the latter factor may be a significant driver of sales or consumer loyalty, however, is harder to establish. The relative lack of commercial success in the end-consumer market of privacy enhancing solutions (such as ZeroKnowledge’s Freedom Network, an anonymous browsing application; or PGP, a data encryption application) may signal an absence of a significant demand for those products. On the other hand, Tsai et al. (2007) show that, under certain conditions, consumers try to purchase from more privacy protective merchants even when that may entail paying modest price premia. Hence, privacy protection may be revenue enhancing.

Offering privacy services to consumers might also save costs for merchants in ways that are not directly related to the privacy they provide or through some sorts of economies of scope. For instance, certain anonymous payment systems might have authentication features that decrease the risk of fraud or charge-backs compared to online credit card payments; or, investments aimed at protecting consumers data (such as firewalls and encryption of server

¹¹See, for instance, a Google executive cited in: Adam Liptak, “When American and European Ideas of Privacy Collide,” New York Times, February 27, 2010.

data) may also protect a company's trade secret and information systems.

3.2.2 Data Subjects

The Costs of Disclosed Data. Consumers appear to be sensitive to threats to their personal information. In the United States, market surveys over the years have consistently found that consumers are concerned with the way businesses collect their personal data. In 2000, a Federal Trade Commission study reported that sixty-seven percent of consumers were “very concerned” about the privacy of the personal information provided on-line (Federal Trade Commission, 2000). In 2005, a CBS News survey found that most Americans found their privacy was under “serious threat.”¹² Similarly, in 2009, a survey by Turow et al. (2009) found that a large majority of Americans resist to tailored advertising.

However, the costs consumers actually incur because of disclosed and abused data are complex to categorize, since they comprise tangible and intangible damages that may occur (if at all) long after the data was initially disclosed.

As an example of the nuances of privacy costs, consider Calo (2011)'s distinction between subjective and objective privacy harms. Subjective harms derive from the unwanted perception of observation. They include “unwelcome mental states – anxiety, embarrassment, fear – that stem from the belief that one is being watched or monitored.” Objective harms consist of the unanticipated or coerced use of information concerning a person against that person, and include outcome as diverse as identity theft, the leaking of classified information that reveals an undercover agent, or “the use of a drunk-driving suspects blood as evidence against him.” As Calo notes, the categories represent, respectively, “the anticipation and consequence of a loss of control over personal information.” While no less important, subjective harms are harder to describe in economic terms than objective ones. The latter can often be described in terms of tort (Prosser, 1960). Instead, the former are not usually

¹²“Poll: Privacy Rights Under Attack,” CBS News, 2005, <http://www.cbsnews.com/stories/2005/09/30/opinion/poll>

recognized by US courts as *actual* damage (Romanosky and Acquisti, 2009); furthermore, they often amount to expected (that is, future and probabilistic, as opposed to presently incurred) costs.

An intuitive way of describing the state of uncertainty associated with privacy costs is the “blank check” metaphor. As an individual reveals private information to other parties, she is signing a blank cheque. The cheque may never come back to her, or may come back for an indeterminably small or large price to pay. That price could be a mild embarrassment, an annoying spam, or a devastating case of identity theft. In short, the probability, form, and actual damage from disclosed data are, in Knight (1921)’s terms, *ambiguous* and, up to a point, unknown.¹³

The blank cheque metaphor highlights the fact that information revealed during a transaction might later reappear at unexpected moments, or in new forms, or in a different context. At one extreme of the range of possible harms lies the case of Amy Boyer. In *Remsburg v. Docusearch* (Inc., 816 A.2d 1001, N.H. 2003), the defendant sold personal information about the plaintiff’s daughter to a man who stalked and then killed her. At the other extreme lie cases where the exposure of a person’s personal information does not cause any actual harm other than the discomfort of feeling violated. In between, a spectrum of losses - from minor annoyances to major damages - may occur.

Some of those costs are immediate but intangible: the psychological discomfort with feeling observed or violated; the embarrassment or social stigma when personal data has been disclosed; the chilling effect of the fear that one’s personal sphere will be, in fact, intruded.

Some costs are immediate and tangible: time and efforts spent deleting junk mail; annoyances from telemarketing; higher prices paid due to (adverse) price discrimination.

¹³Knight (1921) distinguished situations characterized by risk (in which the random outcomes of an event can be described with a known probability distribution) from situations characterized by ambiguity (in which those probabilities are unknown).

Some costs are more indirect: for instance, segmentation and profiling (especially in the form of behavioral targeting and advertising) may manipulate the consumer towards services she does not need or cannot afford.¹⁴

Other costs are only probabilistic (that is, expected, rather than occurred, damages): for instance, errors in consumer databases due to poor data handling procedures by firms *may* later cause a consumer's request to be wrongfully denied; or, breached databases *may* later result in identity theft Camp (2007).

Because of their uncertain nature, privacy costs are therefore often hard to assess, and act upon, for the individual, but by no means less real: they often take the form of high-probability events with negligible individual impact (for instance, spam); or, they materialize as high significance events with very low expected probability of occurrence (for instance, being wrongfully denied a mortgage after suffering from identity theft). In either case, because of their low likelihood of occurrence or their limited impact, they may be dismissed as unimportant at the individual level - even if, in the aggregate, they may amount to significant societal damage.

Identity theft due to data breaches offers an example of the intricacies of assessing and compensating privacy costs. Both breaches and identity thefts have been almost monotonically increasing for the past several years: data breaches in 2008 were up 47% from the previous year, while identity fraud victims increased by 8.6% (Federal Trade Commission 2009). In 2005, up to 35 percent of known identity thefts were estimated to be caused by corporate data breaches (Javelin Research, 2006). Javelin Research estimates that corporate and consumer losses due to identity theft amounted to \$56 billion dollars in 2005. Data breaches, however, can have a wide array of consequences for the affected data subjects. When the breach consists simply in the loss of data (for instance, a misplaced laptop), the

¹⁴For instance, some marketing databases explicitly list personal information of individuals suffering from various types of addiction. For an example of a "gamblers database", see <http://www.dmnews.com/media-one-gamblers-database/article/164172/>.

data subject may not suffer any cost. When the breach is due to a deliberate attack by a malicious third party, the compromised personal data is more likely to be used in manners that directly impact the subject: fraudulent unemployment claims, loans, credit card charges, and health insurance charges. The victims can suffer a ruined credit score, inability to access credit or employment, or even criminal charges - in addition to financial damage, psychological costs, and time losses.

However, only a fraction of those costs are currently reimbursed or recovered by consumers. For instance, fraudulent charges on one's credit card are compensated (although credit card companies may pass the charges back to consumers in the form of higher rates), but US court rulings have, in general, not awarded damages for breaches of personal information, due to the plaintiffs' inability to show *actual* damages - as required by negligence tort claims (Romanosky and Acquisti, 2009) - or to show a clear linkage of causation between the breach and the ensuing damage.

In absence of legal liabilities, contractual requirements, or risk of adverse market reactions, the parties that come to control an individual's information may not internalize such privacy costs (Swire and Litan, 1998), and therefore face lower incentives to protect consumer's data. This increases the probability of moral hazard, with the data holder taking risks with the subject's data.

The database of a merchant, for example, might be hacked because of lax security practices implemented by the merchant. The credit card numbers stored there might be stolen and then illegally used. Absent a data breach notification law (and sometimes even notwithstanding its existence), the customers who own those cards may be unable to identify and hold that merchant responsible. Furthermore, absent robust market reaction, adverse selection may cause less reliable merchants succeed in the on-line marketplace (as a purely anecdotal piece of evidence, consider that, following repeated and well publicized data breaches, and notwithstanding millions of dollars in fines and fees, ChoicePoint was purchased in a

cash deal for \$3.6 billion by Reed Elsevier in February 2008).

Indirect Costs. The existence of a secondary market for customer data can also become a source of negative externalities for consumers. Such externalities may arise when the data holding company extracts the full benefit of using the information in its own marketing efforts, or the full price it receives when it sells the information to third parties, but does not internalize the losses that the consumer may derive from the disclosure of private information. Because customers often do not know the sources of data disclosure or the ways in which their data is combined and used, they may not be able to discipline effectively the companies that exploit that data (Swire and Litan, 1998). In economic terms, the company internalizes the gains from using the information (without necessarily sharing a portion of those gains to the consumer), but externalizes some of the losses.

Finally, a more intangible (but not less important) form of indirect consumers' costs arises from the observation that, the more an individual's data is shared with other parties, the more those parties gain a bargaining advantage in future transactions with that individual. Consider, for instance, behavioral targeting. While the consumer receives offers for products she is actually interested in, data holders accumulate data about her over time and across platforms and transaction. This data permits the creation of a detailed dossier of the consumers' preferences and tastes, and the prediction of her future behavior. As the microeconomic models surveyed in Section 2.2 would predict, it is not hard to imagine that, in presence of myopic customers, this information will affect the allocation of surplus of future transactions, increasing the share of the data holder over that of the data subject. In other words, the disclosure of personal data ultimately affects the balance of power between the data subject and the data holder. The long-run effects of such altering of the balance of power are, of course, very hard to predict.

A recent example was provided by an unexpected feature of Microsoft Bing's Cashback. Cashback was designed to save users money while shopping from online retailers. However,

in November 2009, going to certain third-party sellers' site through Bing may have resulted in higher prices displayed to the user than if she had visited the site directly: The third party merchant had clearly engaged in price discrimination based on the visitors' originating url.¹⁵

The Costs of Protecting Data. Consumer privacy costs can be indirect, too. To get informed about risks to their privacy, consumers incur cognitive and opportunity risks. If we take seriously the premise that consumers' privacy relies on knowledge and consent, the costs of getting consumers informed may be astronomical. For the case of online privacy alone, McDonald and Cranor (2008) calculate that, if every US internet users did peruse the privacy policies of the sites she visits, the national opportunity cost for the time needed to read those policies would be on the order of \$781 billion. Similarly, in response to a breach disclosure, consumers must process the information and decide a course of action. This imposes cognitive costs and can raise an unsurmountable burden against making informed decisions.

Protecting one's information also comes at a cost: money spent for an anonymizing service and privacy enhancing technologies, time spent learning to use the protecting technology, or hassles incurred when changing one's behavior and habits. For instance, in order to avoid being tracked or price discriminated, consumers may have to engage in wasteful activities, investing in protective software that otherwise they would have not needed, or experiencing delays or usability costs associated with privacy enhancing technologies such as Tor or PGP.¹⁶

In addition, indirect costs include opportunities lost when the consumers elects not to share data. We have referred to this in Section 3.1 as a privacy externality.¹⁷ The more other consumers get comfortable with data disclosures, and the more firms start to rely on (or in fact require) that data to provide products and services, the higher is the opportunity

¹⁵See <http://bountii.com/blog/2009/11/23/negative-cashback-from-bing-cashback/>.

¹⁶See www.torproject.org/ and www.pgp.com/.

¹⁷In this context, the term was suggested to us by Professor Brad Malin.

costs for consumers who want to protect their data.

By extension, some have used the analogy of (data) pollution (Jean Camp and Wolfram, 2004) to refer to the externalities associated with data disclosures. Peer-pressure to relinquish data in Web 2.0 applications, as well as the erosion of expectations of privacy, may have the perverse effect of making privacy-preserving alternatives to current products and services simply unavailable under prevailing market conditions.

The Benefits of Protected Data. Conversely, some of the highlighted costs of disclosed data turn into benefits when consumer data is protected. For instance, when firms keep consumer data encrypted, they reduce the likelihood that, even if the data is breached, the consumer will suffer from identity theft. Similarly, consumers can benefit when certain personal information is *not* known by the merchant (such as information that the merchant may correlate with that individual's willingness to bargain for a certain good: see Zettelmeyer et al. (2001)).

More notably, numerous of the benefits associated in Section 3.1 with data disclosure may, in fact, still be gained when data is protected.

For instance, Gellman (2002) points out that the privacy-enhancing restrictions in credit reporting brought about by the Fair Credit Reporting Act did *not* impose the impenetrable barrier to beneficial and profitable uses of consumer data that its critics feared before its passage. The welfare-diminishing effects of privacy regulation may have been similarly overestimated in other sectors, as markets find ways of adapting to new restrictions.

Similarly, in a recent paper, Goldfarb and Tucker (2010) find that while the enactment of privacy regulation limiting behavioral targeting did reduce the effectiveness of (and therefore welfare gains from) ads on websites with general content, it had no such impact on ads on sites with specific content, larger ads, or ads with interactive, video, or audio features.

Furthermore, while behavioral targeting reduces consumers' costs of discovery of products that can match their preferences, so can less intrusive technologies: electronic marketplace

in general, as well as search engines and sponsored searches in particular, reduce buyer's search costs (Bakos, 1997) without linking consumer data across platform and transactions.

Acquisti (2008) also points out that privacy enhancing technologies may be used to protect sensitive data while nevertheless allowing the gathering, analysis, and profitable exploitation of non-sensitive, or de-identified, or aggregate data - with shared benefit for both data subjects and data holders.

In fact, many of the costs that firms would incur to increase the protection of consumer data (see Section 3.2.1) can be classified as *fixed* costs rather than *variable* costs. Consider, for instance, the price of investments necessary to implement access control, data encryption, or privacy-preserving mechanisms to collect and/or analyze personal data (see Section 4.4). Once incurred, these costs would be sunk, but they would not necessarily adversely affect the marginal cost of operation of the firm.

4 The Current Debate

4.1 Regulation vs. Self-Regulation

Broadly speaking, under a regulatory framework, the protection of privacy is seen as a desirable outcome *per se* and in itself, regardless of its economic consequences. Under such a framework, privacy becomes an explicit policy goal, and therefore legislation can be passed to - say - promote the development of certain privacy enhancing technologies, stop the deployment of privacy invasive ones, or create liabilities to enforce desirable behaviors in the marketplace.

Under a self-regulatory framework, instead, the policy goal is not the protection of privacy in and of itself. The goal is, rather, the balancing of information protection and information sharing in order to increase aggregate welfare. Under this framework, markets are therefore expected to self-correct through a combination of industry self regulation, consumers

responsibility, and technological solutions; the legislator can steer the marketplace through a combination of incentives, disclosure policies, and - when everything else fails - liability.

Whereas the US legislator has taken an utilitarian approach to data protection, the European legislator has defined privacy as a fundamental human right.

The US legislator's response to the privacy challenges posed by new ICTs has been "utilitarian," inspired by the *free marketplace of ideas* and based on self-regulation and "targeted" standards for specific types of data (such as, for instance, movie rentals and banking data; Reidenberg (1995)) instead of others (such as, for instance, mortgage information). The benefits of this system are its efficiency and adaptability. Its risks lie in the multiplication of codes, approaches, and interpretations, which may leave the consumer unable to protect her personal sphere *ex-ante* or be compensated for errors and violations *ex-post*. Bellotti (1998), p. 67 writes: "[t]his variability tends to result in lengthy court cases to determine whether a legal violation of privacy has occurred."

The European Parliament response, instead, has been more centralized and regulatory. The Parliament established general privacy rights initially with its Directive 95/46/EC. The Directive entitled individuals to receive notice about who is collecting the data and will have access to it, and why the data is being collected. Under the Directive, individuals maintain also the right to access and if necessary correct the information about themselves. Its advantages may be the certainty of what type of protection is afforded. Its risks lie in its costs, enforceability, and inflexibility.

If the social planner's goal was first and foremost the protection of personal information, it is self-evident that self-regulatory mechanisms based on notice and choice models are not working. In the US, consumers' data is routinely gathered, analyzed, and traded without users' consent and even knowledge. Self-regulatory mechanisms such as privacy policies and notices have failed: the vast majority of US consumers do not believe that privacy policies are easy to understand, and as a consequence do not read them - and if they do, they not

understand them (Antón et al. (2004), McDonald and Cranor (2008)), or misinterpret their presence as implying *de facto* privacy protection. Even the proliferation of data breach notification laws at the state level, while succeeding in the goal of reducing some of the information gap that prevents consumers from knowing what happens to their data, have not been particularly effective in countering identity theft: using reported identity theft data from the FTC for the 2002-2008 period, Romanosky et al. (2008) have found that the adoption of disclosure laws reduced identity theft by about 6%.

If, instead, the social planner's policy goal was to increase aggregate welfare, by balancing businesses' needs to mine consumers' data, and consumers' desire for privacy, comparing regulatory and self-regulatory approaches is more difficult.

The debate over the comparative advantages of regulation and self regulation remains intense to this date. On the one hand, Gellman (2002) challenges the view that the unrestricted trafficking in personal information always benefits the consumer, and that privacy trade-offs may merely be evaluated on the basis of monetary costs and benefits. He concludes that an unregulated, privacy-invasive market in personal data can be costly for consumers. F.H.Cate (2002), Cate et al. (2003), Rubin and Lenard (2001), and Lenard and Rubin (2009), on the other hand, claim that legislative initiatives that restrict the amount of personal information available to business would actually penalize the consumers themselves: regulation should be undertaken only when a give market for data is not functioning properly, and when the benefits of new measures outweigh their costs.

It may not be possible to resolve this debate using purely economic tools. Economic theory, as we have discussed above, has brought forward arguments both supporting the view that privacy protection *increases* economic efficiency, and that it *decreases* it. Empirically, the costs and benefits associated with the protection and revelation of consumers' data have not proven easily amenable to aggregation: First, as soon as one attempts an aggregate evaluation of the impact of privacy regulation, one faces the challenge of delimiting the prob-

lem: data breaches, identity theft, spam, profiling, or price discrimination are all examples of privacy problems, yet they comprise very different expected benefits and costs for the parties involved. Second, even within each scenario, it may be hard to statically measure at a point in time the aggregate costs and benefits of data protection and data sharing, since the benefits and costs of privacy happen over time (for instance, data revealed today may only damage the individual years from now). And third, in addition to measurable outcomes (such as the financial losses due to identity theft, or the opportunity costs of spam), other privacy invasions require an estimation of consumers' valuations of privacy. Evaluations in this area, as we discuss in the next section, are far from stable.

Skeptical of the effectiveness of either approach to privacy protection, other scholars have suggested alternative routes. For example, Samuelson (2000) favors an information licensing approach; Litman (2000) supports a tort-law approach; Schwartz (2004) and Laudon (1996) propose (intellectual) property approaches. Laudon (1996), in particular, advocates the creation of national information markets through which individuals can trade the rights on their data (which they own) to others, in exchange for money. This “co-regulative” approach would combine market forces, technologies, and regulation.

Property rights on personal data have not yet been established in the US. National information markets for personal data exists for businesses such as data aggregators and reporting agencies (that collect, mine, combine, and trade consumers' data), but not for the end consumers. However, thanks in particular to the Internet and the evolution of ICTs, consumers today *de facto* engage in transactions involving the implicit “sale” of their personal data for tangible or intangible benefits on a daily basis. Sometimes those transactions involve explicit trades of money for data. For instance, using a supermarket's loyalty card implies selling one's purchase history for a discount on groceries purchased. Other times, the trade is more opaque, in the sense that the disclosure of personal data may not be the primary focus of the transaction. For instance, completing a query on an Internet search engine provides

the searcher with the desired information, in exchange of reveal the searchers' interests and preferences to the search engine.

4.2 Consumers' Valuations of Privacy

Numerous factors influence individuals' privacy concerns (Milberg et al., 1995), and therefore the mental "privacy calculus" that individuals make when deciding whether to protect or disclose personal information (Laufer and Wolfe, 1977; Culnan and Armstrong, 1999; Dinev and Hart, 2006). Researchers from diverse disciplines (such as economics, marketing, information systems, and computer science) have attempted to estimate empirically the value that, in this calculus, individuals assign to privacy and their personal data. The resulting findings suggest that privacy valuations are especially context dependent.

Huberman et al. (2005) used a second-price auction to estimate the price at which individuals were willing to publicly reveal personal information such as their weight. Individuals who thought their information was less desirable and more deviant from the norm for the rest of the group were more likely to exhibit higher valuations. Wathieu and Friedman (2005) found that survey participants were more acceptive of an organization sharing their personal information after having been explained the economic benefits of doing so. Cvrcek et al. (2006) reported large differentials across EU countries in the price EU citizens would accept to share mobile phone location data. Hann et al. (2007) focused on online privacy and, using a conjoint analysis, found that protection against errors, improper access, and secondary use of personal information was worth US\$30.49-44.62 among US subjects. Rose (2005) found that, although most participants in a survey self-reported being very sensitive to privacy issues, less than half of them would be willing to pay roughly \$29 to have their privacy protected by means of property rights on personal information. Both Varian et al. (2005) and Png (2007) estimated US consumers' implicit valuation of protection from telemarketers using data about the Do Not Call list adoptions. They found highly differing values, from a

few cents to as much as \$30. Tsai et al. (2007) found that, when information about various merchants' privacy policies was made available to them in a compact and salient manner, subjects in an experiment were more likely to pay premia of roughly 50 cents to purchase products from more privacy protective merchants.

At the same time, various studies have highlighted a dichotomy between self professed privacy attitudes and actual self-revelatory behavior.

Tedeschi (2002) reported on a Jupiter Research study in which the overwhelming majority of surveyed online shoppers would give personal data to new shopping sites for the chance to win \$100. Spiekermann et al. (2001) found that even participants in an experiment who could be classified as privacy conscious and concerned were willing to trade privacy for convenience and discounts: differences across individuals in terms of reported concerns did not predict differences in self-revelatory behavior. Similar findings were obtained in different settings by Acquisti and Grossklags (2005) and Acquisti and Gross (2006). Coupled with the observation that businesses focused on providing privacy enhancing applications have met difficulties in the marketplace (Brunk, 2002), these results suggest a privacy paradox: people want privacy, but do not want to pay for it, and in fact are willing to disclose sensitive information for even small rewards (for an overview of this area, see Acquisti (2004) and Acquisti and Grossklags (2007)).

In fact, Acquisti et al. (2009) have recently provided evidence of how highly malleable consumers privacy valuations can be: in an application of the endowment effect to the privacy domain, subjects who started an experiment from positions of greater privacy protection were found to be five times more likely than other subjects (who did not start with that protection) to forego money to preserve their privacy.

4.3 Hurdles in Consumer Behavior and Privacy Nudges

A stream of research investigating the so-called privacy paradox has focused on the hurdles that hamper individuals' privacy-sensitive decision making. This line of enquiry has policy implications: The modern microeconomic theory of privacy suggests that, when consumers are not fully rational or in fact myopic, the market equilibrium will tend *not* to afford privacy protection to individuals (Section 2.2). In other words, in absence of regulatory protection of consumers' data, firms will tend to extract the surplus generated in transaction in which consumers' data is used for price discrimination (Acquisti and Varian, 2005; Taylor, 2004).

There is, indeed, evidence that consumers cannot act rationally (in the canonical economic sense) when facing privacy trade-offs. The evidence relies on three set of decision making hurdles: privacy decision making is afflicted by a) incomplete information, b) bounded cognitive ability to process the available information, and c) a host of systematic deviations from theoretically rational decision making, which can be explained through cognitive and behavioral biases investigated by research in behavioral economics and decision research.

Consider, first, the problem of incomplete information. In many scenarios - such as those associated with behavioral monitoring and targeting - the consumer may not even realize the extent at which her behavior is being monitored and exploited. Furthermore, after an individual has released control on her personal information, she is in a position of information asymmetry with respect to the party with whom she is transacting. In particular, the subject might not know if, when, and how often the information she has provided will be used. For example, a customer might not know how the merchant will use the information that she has just provided to him through a website.

Furthermore, the "value" itself of the individual's information might be highly uncertain and variable. The subject and the parties she is interacting with may evaluate differently the same piece of information, and the specific environmental conditions or the nature of

the transaction may affect the value of information in unpredictable ways. For example, a customer might not know what damage she will incur because of her personal information becoming known, she might not know how much profit others will make thanks to that information, or she might not know the benefits she will forego if her privacy is violated. To what, then, is the subject supposed to anchor the valuation of her personal data and its protection?

Second, findings from behavioral economics exhaustively document consumers' inability to exhaustively consider the possible outcomes and risks of data disclosures, due to bounded rationality. Furthermore, the individual will often find herself in a weaker bargaining position than other parties she is interacting with (for instance, merchants). In many transactions, the individual is unable to negotiate a desired level of information protection; she rather faces take-it-or-leave-it offers of service in exchange for personal data.

Third, even if the consumer had access to complete information about all trade-offs associated with data sharing and data protection, she will suffer from cognitive and behavioral biases that are more intense in scenarios where preferences are more likely to be uncertain. One such example is that, if the expected negative payoff from privacy invasions could be estimated, some individuals might seek immediate gratification, discounting hyperbolically (Rabin and O'Donoghue, 2000) future risks (for example of being subject to identity theft), and choosing to ignore the danger. Hence, because of asymmetric information, self-gratification bias, over-confidence, or various other forms of misrepresentation studied in the behavioral economic literature, individuals might choose not to protect their privacy *possibly* against their own best interest. They might be acting *myopically* when it comes to protecting their privacy even when they might be acting *strategically* (as rational agents) when bargaining for short-term advantages such as discounts (Acquisti, 2004).

Consider, for instance, the case of data breaches. As discussed in Romanosky and Acquisti (2009), after being notified of a breach of her financial information, a consumer may

not be able to identify the right course of action: should she, for instance, punish the financial firm that, due to faulty security controls, compromised her data, by changing to a competitor? While this may appear as a risk-reducing behavior, by doing so the consumer would have now disclosed her personal information to another firm - and actually materially increased the probability that another future breach will involve her data. Furthermore, the cost of acting may be significant: calling the breached firm to obtain details about the breach and its consequences, notifying financial institutions of the occurred breach and of potentially compromised accounts, or subscribing to credit alert and insurance services, are all actions which carry perceived cognitive, transaction, and actual costs. Such costs may appear greater to the consumer than the perceived benefit from action. It could also be that, because of psychological habituation due to repeated instances of data breaches report in the media, the consumer may become desensitized to their effects - which counter the desired impact of notifications. Ultimately, the consumer may 'rationally' decide to remain 'ignorant' (following the Choicepoint breach, fewer than 10% of affected individuals availed themselves of the free credit protection and monitoring tools offered by Choicepoint (Romanosky and Acquisti, 2009)). This example suggests how nuanced and full of obstacles is the path that lead from consumer notification of privacy problem to her actually taking action to solve that problem.

An improved understanding of cognitive and behavioral biases that hamper privacy (and security) decision making, however, could also be exploited for normative purposes. Specifically, knowledge of those biases could be used to design technologies and policies that anticipate and counter those very biases (Acquisti, 2009). Such technologies and policies would be informed by the growing body of behavioral economics research on soft or asymmetric paternalism (Loewenstein and Haisley, 2008) as well as research on privacy and security usability. They may help consumers and societies achieve their desired balance between information protection and information sharing.

4.4 The Promise of Privacy Enhancing Technology

Information technologies can be used to track, analyze and link vast amounts of data related to the same individual. However, information technologies can also be used to protect, anonymize, or aggregate those data in ways that are both effective (in the sense that re-identifying individual information becomes either impossible or just costly enough to be unprofitable) and efficient (in the sense that the desired transaction can be regularly completed with no additional costs for the parties involved).

A vast body of research in privacy enhancing technologies suggests, in fact, that cryptographic protocols can be leveraged to satisfy both needs for data sharing and needs for data privacy. Not only is it already possible to complete verifiable and yet anonymous or privacy enhanced “transactions” in areas as diverse as electronic payments (Chaum, 1983), online communications (Chaum, 1985), Internet browsing (Dingledine et al., 2004), or electronic voting (Benaloh, 1987); but it is also possible to have credential systems that provide authentication without identification (Camenisch and Lysyanskaya, 2001), share personal preferences while protecting privacy (Adar and Huberamn, 2001), leverage the power of recommender systems and collaborative filtering without exposing individual identities Canny (2002), or even executing calculations in encrypted spaces (Gentry, 2009), opening the doors for novel scenarios of privacy preserving data gathering and analysis.

In other words, privacy enhancing technologies may make it possible to reach a new economic equilibrium where data holders can still analyze aggregate and anonymized data, while subjects’ individual information stays protected. Arguably, the transition to this new equilibrium could be welfare-enhancing. The challenge, however, is that few privacy enhancing technologies have gained widespread adoption. Many companies based on them have failed in the marketplace Brunk (2002). Several reasons may explain these failures: on the consumers’ side, users’ difficulties and costs in using privacy technologies (see Whit-

ten and Tygar (1999)), switching costs, as well as biases such as immediate gratification, which reduce demands for those products even by privacy sensitive consumers. On the data holders' side, in absence of policy intervention or clear proof that privacy protection acts as a distinctive competitive advantage, it is unlikely that firms will incur the costs to transition to technologies that may, in the short run, limit their access to consumer data relative to their competitors. This calls for a form of what NYU Professor Laudon (1996)'s called "co-regulative" solutions to privacy problems, in which economic forces, cryptographic technologies, and targeted regulatory guidelines conspire to create a system with adequate enforcement and control powers (see also OECD (1997)).

5 Conclusion: Should We Change the Frame of the Debate?

Considering the conflicting analyses we have presented, we conclude that it would be futile to attempt comparing the aggregate values of personal data and privacy protection, in search of a "final" economic assessment of whether we need more, or less, privacy protection. Privacy means too many things, its associated trade-offs are too diverse, and consumers valuations of personal data are too nuanced. Furthermore, economic theory shows that, depending on conditions and assumptions, the protection of personal privacy can increase aggregate welfare as much as the interruption of data flows can decrease it.

In this author's opinion, therefore, solving the privacy problem means to find a *balance* between information sharing and information hiding that is in the best interest of data subjects but also of society as a whole (including *other* data subjects and potential data holders). Current evidence, however, suggests that self-regulatory, market-driven solutions are not, alone, achieving that balance. Similarly, user awareness or education programs, consumer-focused privacy enhancing technologies, and user-controllable privacy solutions are, in our opinion, merely *necessary* but not *sufficient* conditions of privacy balance, because of the numerous hurdles in privacy sensitive decision making highlighted by behavioral decision

research.

However, the promise of current research in cryptographic protocols is that privacy needs may be satisfied without significant damage to those useful flows of personal data. Regulators' interventions aimed at fostering the dissemination and adoption of those technologies, therefore, may help us reach that more desirable economic equilibrium. In such a co-regulatory framework, economics could highlight different trade-offs, technology could help achieve more desirable equilibria, and regulatory intervention could nudge the market to adopt those technologies.

Perhaps, then, the frame of the privacy debate would change as well: the burden of proof for deciding whom and how should protect consumers privacy would go from *prove that the consumer is bearing a cost* when her privacy is not respected, to *prove that the firm cannot provide the same product*, in manners that are more protective of individual privacy.

References

- Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the ACM Conference on Electronic Commerce (EC '04)*, pp. 21–29.
- Acquisti, A. (2008). Identity management, privacy, and price discrimination. *IEEE Security & Privacy*, 46–50.
- Acquisti, A. (2009). Nudging privacy: The behavioral economics of personal information. *Security & Privacy, IEEE* 7(6), 82–85.
- Acquisti, A., A. Friedman, and R. Telang (2006). Is there a cost to privacy breaches? an event study. In *Twenty Seventh International Conference on Information Systems and Workshop on the Economics of Information Security 2006*.
- Acquisti, A. and R. Gross (2006). Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Workshop on Privacy Enhancing Technologies (PET '06)*.
- Acquisti, A. and J. Grossklags (2005). Privacy and rationality in individual decision making. *IEEE Security and Privacy* 3(1), 24–30.
- Acquisti, A. and J. Grossklags (2007). What can behavioral economics teach us about privacy? In S. G. C. L. Alessandro Acquisti, Sabrina De Capitani di Vimercati (Ed.), *Digital Privacy: Theory, Technologies and Practices*, pp. 363–377. Auerbach Publications (Taylor and Francis Group).
- Acquisti, A., L. John, and G. Loewenstein (2009). What is privacy worth? In *Workshop on Information Systems Economics*.
- Acquisti, A. and H. R. Varian (2005). Conditioning prices on purchase history. *Marketing Science* 24(3), 367–381.
- Adar, E. and B. Huberman (2001). A market for secrets. *First Monday* 6, 200–209.
- Antón, A. I., J. B. Earp, Q. He, W. H. Stufflebeam, D. Bolchini, and C. Jensen (2004). Financial privacy policies and the need for standardization. *IEEE Security and Privacy* 2(2), 36–45.
- Bakos, J. (1997). Reducing buyer search costs: implications for electronic marketplaces. *Management science* 43(12), 1676–1692.
- Ball, D., P. Coelho, and M. Vilares (2006). Service personalization and loyalty. *Journal of Services Marketing* 20(6), 391–403.
- Beales, H. (2010). The value of behavioral targeting. Network Advertising Initiative.

- Bellotti, V. (1998). *Technology and privacy: the new landscape*, Chapter Design for Privacy in Multimedia Computing and Communications Environments, pp. 63–98. MIT Press.
- Benaloh, J. C. (1987). *Verifiable Secret-Ballot Elections*. Ph. D. thesis, Yale University.
- Bennett, J. and S. Lanning (2007). The Netflix prize. In *Proceedings of KDD Cup and Workshop*.
- Blattberg, R. C. and J. Deighton (1991). Interactive marketing: Exploiting the age of addressability. *Sloan Management Review* 33(1), 5–14.
- Brunk, B. D. (2002). Understanding the privacy space. *First Monday* 7(10).
- Calo, R. (2011). The boundaries of privacy harm. *Indiana Law Journal* 86.
- Calzolari, G. and A. Pavan (2006). On the optimality of privacy in sequential contracting. *Journal of Economic Theory* 130(1), 168–204.
- Camenisch, J. and A. Lysyanskaya (2001). An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Advances in Cryptology - EUROCRYPT '01*, pp. 93–118. Springer-Verlag, LNCS 2045.
- Camp, L. J. (2007). *Economics of Identity Theft: Avoidance, Causes and Possible Cures*. Springer.
- Canny, J. F. (2002). Collaborative filtering with privacy. In *IEEE Symposium on Security and Privacy*, pp. 45–57.
- Cate, F. H., R. E. Litan, M. Staten, and P. Wallison (2003). Financial privacy, consumer prosperity, and the public good: Maintaining the balance. Federal Trade Commission Workshop on Information Flows: The costs and benefits to consumers and businesses of the collection and use of consumer information.
- Chaum, D. (1983). Blind signatures for untraceable payments. In *Advances in Cryptology*, pp. 199–203. Plenum Press.
- Chaum, D. (1985). Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM* 28(10), 1030–1044.
- Culnan, M. J. and P. K. Armstrong (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science* 10(1), 104–115.
- Cvrcek, D., M. Kumpost, V. Matyas, and G. Danezis (2006). The Value of Location Information. In *ACM Workshop on Privacy in the Electronic Society (WPES)*.
- Daripa, A. and S. Kapur (2001). Pricing on the Internet. *Oxford Review of Economic Policy* 17(2), 202.

- Deighton, J. and J. Quelch (2009). Economic Value of the Advertising-Supported Internet Ecosystem. IAB Report.
- Dereszynski, E. and T. Dietterich (2007). Probabilistic models for anomaly detection in remote sensor data streams. In *23rd Conference on Uncertainty in Artificial Intelligence (UAI-2007)*. Citeseer.
- Dinev, T. and P. Hart (2006). An extended privacy calculus model for e-commerce transactions. *Information Systems Research* 17(1), 61–80.
- Dingledine, R., N. Mathewson, and P. Syverson (2004). Tor: The second-generation onion router. In *Proceedings of the 13th conference on USENIX Security Symposium-Volume 13*, pp. 21.
- Federal Trade Commission (2000, May). Privacy online: Fair information practices in the electronic marketplace: A report to congress. Technical report, Federal Trade Commission.
- F.H.Cate (2002). Principles for protecting privacy. *Cato Journal* 22(1), 33–57.
- Forrester Research (1999). Forrester technographics finds online consumers fearful of privacy violations. <http://www.forrester.com/ER/Press/Release/0,1769,177,FF.html>.
- Gellman, R. (2002, March). Privacy, consumers, and costs - how the lack of privacy costs consumers and why business studies of privacy costs are biased and incomplete.
- Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st annual ACM symposium on Theory of Computing*, pp. 169–178.
- Goldfarb, A. and C. Tucker (2010). Privacy regulation and online advertising. Available at SSRN: <http://ssrn.com/abstract=1600259>.
- Hann, I.-H., K.-L. Hui, T. S. Lee, and I. P. Png (2007). Overcoming online information privacy concerns: An information processing theory approach. *Journal of Management Information Systems* 42(2), 13–42.
- Hermalin, B. E. and M. L. Katz (2006). Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quantitative Marketing and Economics* 4(3), 209–239.
- Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *The American Economic Review* 61(4), 561–574.
- Hirshleifer, J. (1980, December). Privacy: Its origins, function and future. *Journal of Legal Studies* 9(4), 649–664.
- Hoofnagle, C. J. (2007). Security breach notification laws: Views from Chief Security Officers.
- Huberman, B. A., E. Adar, and L. R. Fine (2005). Valuating privacy. *IEEE Security and Privacy* 3, 22–25.

- Hui, K.-L. and I. Png (2003, December). Direct marketing: Privacy and competition. In *Workshop on Information Systems and Economics - WISE '03*, Seattle, WA.
- Hui, K.-L. and I. P. Png (2005, May). The economics of privacy. Technical report, EconWPA.
- Jean Camp, L. and C. Wolfram (2004). Pricing security. In S. Jajodia, L. Camp, and S. Lewis (Eds.), *Economics of Information Security*, Volume 12 of *Advances in Information Security*, pp. 17–34. Springer US.
- Jentzsch, N. (2001). The economics and regulation of financial privacy - a comparative analysis of the united states and europe. Technical report, John F. Kennedy Institute.
- Knight, F. (1921). *Risk, uncertainty and profit*. Hart, Schaffner & Marx; Houghton Mifflin Company, Boston, MA.
- Krasnikov, A., S. Jayachandran, and V. Kumar (2009). The Impact of Customer Relationship Management Implementation on Cost and Profit Efficiencies: Evidence from the US Commercial Banking Industry. *Journal of Marketing* 73, 61–76.
- Laudon, K. C. (1996). Markets and privacy. *Communications of the ACM* 39(9), 92–104.
- Laufer, R. S. and M. Wolfe (1977). Privacy as a concept and a social issue: a multidimensional developmental theory. *Journal of Social Issues* 33(3), 22–42.
- Lenard, T. M. and P. H. Rubin (2009). In defense of data: Information and the costs of privacy. Technology Policy Institute.
- Linden, G., B. Smith, and J. York (2003). Amazon. com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing* 7(1), 76–80.
- Litman, J. (2000). Information privacy - Information property. *Stanford Law Review* 52.
- Loewenstein, G. and E. Haisley (2008). The economist as therapist: Methodological issues raised by light paternalism. In S. A. Caplin A. (Ed.), *Perspectives on the Future of Economics: Positive and Normative Foundations*.
- McDonald, A. M. and L. F. Cranor (2008). The cost of reading privacy policies. *Journal of Law and Policy for the Information Society 2008 Privacy Year in Review issue 4*, 543–897.
- Milberg, S. J., S. J. Burke, H. J. Smith, and E. A. Kallman (1995). Values, personal information privacy, and regulatory approaches. *Communications of the ACM* 38(12), 65–74.
- Murphy, R. (1996). Property rights in personal information: An economic defense of privacy. *Georgetown Law Journal* 84, 2381–2573.
- Noam, E. M. (1997). Privacy and self-regulation: Markets for electronic privacy.

- Odlyzko, A. (2003). Privacy, economics, and price discrimination on the internet. In *Proceedings of the 5th international conference on Electronic commerce*.
- OECD (1997). Implementing the OECD privacy guidelines in the electronic environment: Focus on the internet.
- Pitta, D. (2010). Jump on the bandwagon—its the last one: New developments in online promotion. *Journal of Consumer Marketing* 27(2).
- Png, I. (2007). On the Value of Privacy from Telemarketing: Evidence from the 'Do Not Call' Registry.
- Ponemon, D. (2009). Fourth Annual US Cost of Data Breach Study; Benchmark Study of Companies. *Traverse City, MI: Ponemon Institute..*
- Posner, R. A. (1978). The right of privacy. *Georgia Law Review* 12(3), 393–422.
- Posner, R. A. (1981, May). The economics of privacy. *The American Economic Review* 71(2), 405–409.
- Privacy & American Business (2005). New survey reports an increase in id theft and decrease in consumer confidence. Conducted by Harris Interactive.
- Prosser, W. (1960). Privacy (A Legal Analysis). *California Law Review* 48(3), 338–423.
- Rabin, M. and T. O'Donoghue (2000). The economics of immediate gratification. *Journal of Behavioral Decision Making* 13(2), 233–250.
- Reidenberg, J. R. (1995). Setting standards for fair information practice in the u.s. private sector. *Iowa Law Review* 80, 497–551.
- Richards, K. and E. Jones (2008). Customer relationship management: Finding value drivers. *Industrial Marketing Management* 37(2), 120–130.
- Romanosky, S. and A. Acquisti (2009). Privacy Costs and Personal Data Protection: Economic and Legal Perspectives. *Berkeley Technology Law Journal* 24(3).
- Romanosky, S., R. Telang, and A. Acquisti (2008). Do data breach disclosure laws reduce identity theft? In *Seventh Workshop on the Economics of Information Security - WEIS 2008*.
- Rubin, P. H. and T. M. Lenard (2001). *Privacy and the Commercial Use of Personal Information*. Kluwer Academic Publishers.
- Samuelson, P. (2000). Privacy as intellectual property. *Stanford Law Review* 52(1125).
- Samuelson, P. (2003). The social costs of incoherent privacy policies. Presentation at IBM Almaden Privacy Institute.

- Schwartz, P. (2004). Property, Privacy, and Personal Data. *Harvard Law Review*, 2056–2128.
- Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review* 154(3), 477.
- Spiekermann, S., J. Grossklags, and B. Berendt (2001). E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *3rd ACM conference on Electronic Commerce*.
- Staten, M. and F. Cate (2003). The impact of opt-in privacy rules on retail credit markets: A case study of MBNA. *Duke Law Journal* 52(4), 745–787.
- Stigler, G. J. (1980, December). An introduction to privacy in economics and politics. *The Journal of Legal Studies* 9(4), 623–44.
- Stone, E. and D. Stone (1990). Privacy in organizations: Theoretical issues, research findings, and protection mechanisms. *Research in personnel and human resources management* 8(3), 349–411.
- Streitfield, D. (2000). On the web price tags blur: What you pay could depend on who you are. *The Washington Post*.
- Swire, P. P. and R. E. Litan (1998). *None of Your Business - World Data Flows, Electronic Commerce, and the European Privacy Directive*. Washington, DC: Brookings Institution Press.
- Taylor, C. R. (2003, June). Privacy in competitive markets. Duke University - Economics Working Paper No. 03-10.
- Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics* 35(4), 631–651.
- Tedeschi, B. (2002, June 3rd). E-commerce report; everybody talks about online privacy, but few do anything about it. *New York Times*.
- Tsai, J., S. Egelman, L. F. Cranor, and A. Acquisti (2007). The effect of online privacy information on purchasing behavior: An experimental study. In *6th Workshop on the Economics of Information Security*.
- Turow, J., J. King, C. J. Hoofnagle, A. Bleakley, and M. Hennessy (2009). Americans reject tailored advertising and three activities that enable it. Technical report, Annenberg School for Communication, University of Pennsylvania.
- Varian, H. (1985). Price discrimination and social welfare. *The American Economic Review* 75(4), 870–875.
- Varian, H., F. Wallenberg, and G. Woroch (2005). The demographics of the do-not-call list. *IEEE Security and Privacy* 3(1), 34–39.

- Varian, H. R. (1996). *Economic Aspects of Personal Privacy*. Technical report, University of California, Berkeley.
- Varian, H. R. and C. Shapiro (1997). US government information policy. Presented at Highlands Forum, Department of Defense, June 8, Washington DC.
- Wathieu, L. and A. Friedman (2005). An empirical approach to understanding privacy valuation. In *4th Workshop on the Economics of Information Security*.
- Westin, A. (1970). *Privacy and Freedom*. New York: Atheneum Publishers.
- Whitten, A. and J. D. Tygar (1999). Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *8th USENIX Security Symposium*.
- Wilson, K. and J. Brownstein (2009). Early detection of disease outbreaks using the Internet. *Canadian Medical Association Journal* 180(8), 829.
- Zettelmeyer, F., F. M. S. Morton, and J. Silva-Risso (2001). Cowboys or cowards: Why are internet car prices lower? Technical report, Haas School, UC Berkeley Marketing Working Paper No. 01-1.