Protecting Online Privacy: Self-Regulation, Mandatory Standards, or *Caveat Emptor*

Zhulei Tang[*], Yu (Jeffrey) Hu[†], Michael D. Smith[‡]

This Version: April 2005

ABSTRACT

Information technology-enabled markets enhance retailers' ability to collect, aggregate, and transfer consumer information. These technological capabilities have raised concerns that this information could be used in ways the consumer would not anticipate or authorize.

These concerns have been met with a variety of proposals including approaches placing the onus for protection on consumers, industry self-regulation, and government legislation of mandatory protection standards. However, there has been no research to understand under what circumstances each of these regimes will produce optimal outcomes for customers, retailers, and society. Our research seeks to answer this question using analytic models of asymmetric information.

Our results show that the optimal privacy protection regime depends critically on the characteristics of the market —the number of individuals who face a loss from privacy violations and the size of the loss they face. We find that regimes that place the onus on consumers are socially optimal when few people are sensitive to privacy violations or when the loss they face from privacy violations is low. Conversely, when many people care about privacy protection and the potential loss they face is high, mandatory standards are socially optimal. Finally, for intermediate values, seal-of-approval programs provide socially optimal privacy protection.

(*Privacy; Consumer Surplus; Social Welfare; Internet; Consumer Information*)

Acknowledgements: The authors thank Pei-Yu Chen, Roy Jones, Sandra Slaughter, Tunay Tunca, three anonymous referees at the 2005 Workshop on the Economics of Information Security, and seminar participants at the 2003 Workshop on Information Systems and Economics and Tepper School of Business at Carnegie Mellon University for helpful comments on this research

[*] Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, 15213. email: zhuleit@andrew.cmu.edu.

[†] Sloan School of Management, MIT, Cambridge, MA 02142. email: yuhu@mit.edu

[‡] H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA, 15213. email: mds@cmu.edu

1. Introduction

Information technologies have enabled online retailers to assemble an unprecedented amount of information on consumers. Through direct observation, retailers can record a consumer's on-site browsing behavior, purchase history, and shipping and billing information. Moreover, retailers can add to this information over time, aggregate it across multiple databases, or easily transfer it to third parties.

While a boon to marketers, these capabilities have raised concerns among consumer advocates and regulators that this information will be used in ways that violate consumer privacy. Privacy can be defined in a variety of ways. In this paper we adopt the definition of information privacy as the right to control information about oneself (Westin 1967). A violation of privacy therefore means the use of consumer information revealed during a normal economic transaction that leads to a loss of utility on the part of the consumer.

Several recent surveys suggest that there are widespread consumer privacy concerns in IT-enabled markets.[1] These privacy concerns have led to the application of several privacy protection approaches by governments and various third-parties. These approaches fall into three general categories. The first category is *caveat emptor*, literally "let the buyer beware." Under this approach, retailers are required by law to abide by any agreements they make with consumers to protect their privacy (e.g., through posted privacy policies), but are under no obligation to make such agreements. This approach suffers from the dual problems that not all Internet retailers post their privacy policies, and even fewer consumers invest the time necessary to read and understand privacy policies.[2] The U.S.

[1] For example, Hoffman et al. (1999) find that almost 95% of Web users have declined to provide personal information to Web sites at one time or another due to privacy concerns. The cover story of The Economist's May 1999 issue, "The End of Privacy," cites a survey that shows 80% of Americans worry about what happens to information collected about them and warns that "threats to traditional notions of privacy will proliferate."

[2] For example, Regan (2001) observes that less than one percent of the visitors to six major online travel sites during April 2001 actually read the site's privacy policies.

government has taken a *caveat emptor* strategy toward many common types of consumer transaction data (Baron 2000).

At the opposite end of the spectrum is an approach of *mandatory standards* for privacy protection — where governments intervene and enact strict privacy protection standards for specific types of consumer information. The European Union has adopted mandatory standards for most types of consumer information through their 1996 Directive on Data Privacy (see Smith 2001 for a useful review of this directive). In the United States mandatory standards have been legislated for the use of credit reporting data, health information, some types of financial transactions, and marketing data from minors.

Seal-of-approval programs represent an interesting alternative to *caveat emptor* and mandatory standards regimes, particularly for Internet markets. Under this approach, a retailer can choose — for a fee — to join a seal-of-approval program administered by a seal granting authority. This fee gives the retailer the right to display a logo that certifies that the retailer will follow a certain set of standards to protect consumer privacy, and that the seal granting authority will have the right to monitor the retailer's adherence to these standards. Examples of seal-of-approval programs include the programs offered by TRUSTe and Better Business Bureau.

However, in spite of this wide diversity of potential privacy protection regimes, there has been no systematic research to understand which regimes are optimal for which types of markets. Answers to this question are very important currently as several privacy advocates are arguing that the United States should adopt an overarching set of mandatory standards for all types of customer information, similar to those adopted by the European Union (Ryan 2000).

To address our question, we develop a model of asymmetric information in which the retailer has private information regarding its own cost of protecting consumer privacy. We analyze the retailer's strategy under the three privacy protection regimes: mandatory standards, *caveat emptor*, and seal-of-

approval programs. Our models show that under the mandatory standard regime, retailers will always protect consumer privacy (due to high penalties for privacy violations), while under the caveat emptor regime retailers will never protect consumer privacy (because consumers do not read the retailer's privacy policy). Under seal-of-approval programs, the retailer sends a signal to consumers of its cost of protecting their privacy by joining (or not joining) a seal-of-approval program. Our model shows that if the seal-of-approval program chooses its membership fee and penalty appropriately, consumers can infer the retailer's cost of protecting privacy by observing the retailer's decision to join the seal-of-approval program. In other words, joining the seal-of-approval program can serve as a credible signal that a retailer will protect consumer privacy. Retailers with a low cost of protecting consumer privacy will find it profitable to join the seal-of-approval program, while retailers with a high cost of protecting consumer privacy will find it unprofitable to do so.

We then show that the optimal privacy protection regime will vary depending on the proportion of consumers who care about their privacy and the size of the loss they suffer from privacy violations. When either of these parameters is sufficiently small, *caveat emptor* is the socially optimal privacy protection regime; when both are sufficiently large, the approach of mandatory standards dominates; and for intermediate values, seal-of-approval programs provide the socially optimal solution.

The remainder of this paper proceeds as follows. In Section 2 we review the relevant economic literature on signaling games and the relevant academic literature relating to privacy. In Section 3 we introduce the model for each of three privacy protection regimes, the equilibrium solutions to these models, and discuss their implications. In Section 4 we compare consumer and producer surplus under the privacy protection regimes. In Section 5 we conclude with some broader implications of our work.

2. Literature Review

Our paper primarily relates to two strands of research in online privacy: cost and benefit analysis of consumer information disclosure, and regulation of the market for private information. The first strand of research examines the trade-offs that consumers face when they reveal their information to firms. On one hand, the information can help firms to customize their products, and therefore, better serve the consumers (Ghose and Chen 2003). On the other hand, cross-selling and information risks associated with it may make the risk-averse consumers reluctant to reveal their information (Akcura and Srinivasan 2003). For example, Ghose and Chen (2003) examine these trade-offs in the context of personalization, where firms can adjust the level of personalization through privacy enhancing technology. In the same vein, Akcura and Srinivasan (2003) and Hann et al. (2003) study similar trade-offs in the context of target marketing. Hann et al. (2002) empirically quantifies the value of website privacy protection. Finally, Vila et al (2003) argue that asymmetric information about whether websites will sell private information or not leads to a lemons market for privacy. Even with the privacy policies as signals, the lemons problem cannot be alleviated because there is not enough cost differential between privacy respecting and defecting sites, meaning that government regulation and enforced laws are the only effective methods to make all companies respect consumer privacy. In our model we extend Vila's assumptions to include the possibility that sites may have different costs of protecting a consumer's privacy.

A second related set of papers addresses the problem of regulating a market for private information. For example, Laudon (1996) proposes a national information market where individuals' information is traded in exchange for some form of compensation. Chellappa and Shivendu (2003) also conceptualize privacy as commodity that consumers may trade with vendors. They introduce two regulatory approaches, where the regulator's role is determining whether to allow a private contract between the vendor and the consumers.

The present manuscript differs from these models in that conceptually our model does not treat consumers' personal information as their personal property. Rather, consumers may influence the retailer's action of whether to protect privacy or not through their purchasing decision. Moreover, our paper departs from this strand of research by assuming asymmetric information between consumers and the retailer.

The present paper also differs from the extant literature in that none of the papers we are aware of have examined why different regimes of protecting online information privacy coexist and under what circumstances one regime will be more efficient than others. To explore this question, we draw on the economic literature of games of asymmetric information (Kihlstrom and Riordan 1984, Milgrom and Roberts 1986, Wernerfelt 1988). Our research differs from these models in three ways. First, while the advertising signal in Milgrom and Roberts' is an expenditure that does not improve demand, the signal of whether to join the seal-of-approval program in our model influences demand. Second, Milgrom and Roberts' result rely on repeat purchases, while our results do not. Finally, while Milgrom and Roberts and Kihlstrom and Riordan were attempting to analyzing the signaling value of a particular regime, or focus is on comparing the welfare implications of multiple privacy protection regimes.

Our paper is also related to the broader literature on information regulation in domains other than privacy.[3] Magat and Viscusi (1992) analyze various forms of information regulation, mostly in the context of the use of product labels to reduce the risks of using hazardous chemicals. Breyer (1993) studies the federal regulation of substances that create health risks from a legal prospective. Sunstein (1999) discusses under what circumstances information strategies and information regulation have advantage over command-and-control approaches. Karkkainen (2001) examines the information-

[3] We thank anonymous referees at the 2005 Workshop on the Economics of Information Security for pointing this out.
[5] http://www.ftc.gov/reports/privacy3/fairinfo.htm.

generating mechanisms adopted by the Environment Protection Agency's (EPA) Toxics Release Inventory (TRI) and argues that TRI has an advantage over conventional forms of regulation.

3. A Model of the Retailer, Consumers and Privacy Protection Regimes

3.1. Basic Assumptions of the Retailer and Consumers

We consider a monopolistic retailer who sells a product to consumers at a posted price. Without loss of generality, we assume a zero marginal cost of production. We assume that retailers can either have a high (c_H) or low (c_L, $0 < c_L < c_H$) unit cost of protecting consumer privacy. The retailers' cost of protecting consumer privacy can arise from either the opportunity cost of protecting consumer privacy or the operational cost of protecting consumer privacy through technology systems or personnel. The opportunity cost of protecting consumer privacy in terms of not selling the consumer information to outside marketing companies vary significantly across firms based on the nature of firms' products, the natures of firms' consumers, and the amount of information firms is able to observe about customers. The operational cost of protecting consumer privacy can be quite heterogeneous, as well. For example, firms may employ different privacy protection infrastructure, which cost varies significantly. We refer to a retailer who has a high cost as an H-type retailer and a retailer who has a low cost as an L-type retailer. The levels of c_L and c_H are exogenous and common knowledge to both the retailer and consumers; however, the retailer has private knowledge of its own type.

We assume that if the retailer chooses to protect privacy, consumer privacy will indeed be protected. That is, there are no accidental privacy breaches. Specifically, protecting privacy means that the retailer follows the fair information practice to allow notice, choice, access, participation, security and enforcement of the consumers' private information.[5] If the retailer chooses to protect privacy, the retailer incurs the cost of protecting privacy depending on its type. If the retailer chooses not to protect privacy, the retailer incurs zero cost.

Consumers are divided into two segments depending on whether they care about their privacy (Miyazaki and Fernandez 2001, Pavlou and Chellappa 2001). A proportion of ρ $(0 \le \rho \le 1)$ consumers cares about privacy and incurs a utility loss of L if their privacy is not protected (Hann et al. 2002, Chellappa and Shivendu 2003, Akcura and Srinivasan 2003). The remaining $1-\rho$ consumers do not care about privacy and do not incur a utility loss when their privacy is not protected. We refer to these segments of consumers as *S* (sensitive) and *I* (insensitive), respectively. The *I* consumers are similar to the "unconcerneds" in Akcura and Srinivasan (2003), in that they have no worries and concerns whether their privacy is violated or not. On the contrary, the *S* consumers have such concerns and are willing to take actions to avoid the utility loss from privacy violation. For example, they may want to read the retailer's privacy policy.

Consumers are heterogeneous in their willingness to pay for this product (v~U[0,1]). Each consumer purchases zero or one unit of the product. If a sensitive consumer purchases the product at price p, she obtains a utility of $v-p$ if her privacy is protected and a utility of $v-p-L$ if it is not protected. If an insensitive consumer purchases the product at price p, she obtains a utility of $v-p$ whether her privacy is protected or not. All consumers obtain a utility of zero if no purchase is made.

Finally, we assume that consumers can ascertain whether a retailer claims to protect privacy or not by reading their posted privacy policy. Consumers incur a cost R to read and understand these posted privacy policies.

3.2. Mandatory Standards

Under the mandatory standards regime, the government sets minimum standards for protecting consumer privacy and requires all retailers to follow these standards. For example, in the United States to address privacy protection issues legislation has been passed to require minimum privacy protection standards for consumer credit reporting, health information, marketing data about minors, and

some types of financial transactions (Hahn and Layne-Farrar, 2001). This regime is similar to a conventional approach to regulate economic activity, such as standard setting (Magat and Viscusi 1992) and command-and-control regulation (Sustein 1999).

To model this regime, we assume that if retailers do not protect consumer privacy they are discovered with a probability α $(0 < \alpha \leq 1)$, and if discovered they incur a fine of F.[6] It is straightforward to show that the government can set F sufficiently high so that both types of retailers will find it optimal to protect privacy and that the government has an incentive to do so.[7] Thus, retailers will always choose to follow the law and protect privacy and will incur the respective privacy protection cost of c_H or c_L. Because consumer privacy is protected, both sensitive and insensitive consumers will obtain a utility of $v - p$ if a purchase is made at price p, and a zero utility if no purchase is made.

PROPOSITION 1. Under the mandatory standards approach, all consumers have a belief that the retailer will protect consumer privacy with probability one, and the retailer does indeed protect consumer privacy. In equilibrium, the retailer sets a price of $p_i^* = (1 + c_i)/2$ and obtains a profit of $\pi_i^* = (1 - c_i)^2/4$, where $i = L, H$ and stands for the retailer's type.

For the proof of this and all other propositions, please see the appendix.

3.3. Caveat Emptor

We use a five-stage game to capture the behavior of the retailer and consumers under the *caveat emptor* regime (Table 1). In stage 1, nature decides the retailer's cost of protecting privacy, and this information is revealed only to the retailer. In stage 2, the retailer posts the product's selling price and

[6] For example, in the United States the Federal Trade Commission has the authority to penalize web sites that do not follow their stated privacy policies when dealing with consumer data.

[7] If the retailer protects consumer privacy, its profit function is $\pi = (1 - p_i)(p_i - c_i)$. The optimal price that maximizes the retailer's profit, therefore, is $p_i^* = (1 + c_i)/2$, and this leads to a profit of $\pi_i^* = (1 - c_i)^2/4$. If the retailer does not protect consumer privacy, its profit function is $\pi = (1 - p_i - \rho L)p_i - \alpha F$. The maximum profit is $\pi_i^* = \left(\frac{1 - \rho L}{2}\right)^2 - \alpha F$. When $\alpha F > \frac{(1 - \rho L)^2 - (1 - c_L)^2}{4}$, the retailer will find it optimal to protect privacy.

its policy regarding privacy protection on its website. In the posted policy, the retailer can either claim that it will protect consumer privacy or that it will not protect consumer privacy. In stage 3, consumers observe the posted price and form a belief regarding whether the retailer claims privacy protection in its privacy policy. Let $\delta(p)$ $(0 \le \delta(p) \le 1)$ be the consumers' belief that the retailer does not claim privacy protection, and $1-\delta(p)$ be a consumers' belief that the retailer claims privacy protection. Consumers then decide whether to read the retailer's privacy policy based on their belief. As noted above, if a consumer decides to read the privacy policy, she incurs a reading cost R $(R > 0)$ that represents the time and effort spent by the consumer in order to read and understand the retailer's privacy policy. After incurring this cost, the consumer is informed of the retailer's claim.

Stage 1	***Stage 2***	***Stage 3***	***Stage 4***	***Stage 5***
Nature decides the retailer's cost of protecting privacy. This information is revealed only to the retailer.	The retailer decides whether to claim privacy protection in its privacy policy and its price.	Consumers observe the price, decide whether to read the privacy policy, and whether to purchase.	The retailer decides whether to protect consumer privacy.	The enforcement agency detects deceptive claims with a probability α, and penalizes the retailer if it has made deceptive claims.

Table 1: Timing of the caveat emptor game

Consumers, whether informed or uninformed, then decide whether to make purchases based on the information available to them in stage 3. Next in stage 4, the retailer decides whether to protect the privacy of consumers who have made purchases. Similar to the mandatory standards game above, in stage 5, if the retailer has claimed privacy protection but does not follow through on its promise, this violation is detected with probability α $(0 < \alpha \le 1)$, and penalized by a fine of F.[8] After this stage, the game is over and the retailer's profit and consumers' utilities are realized.

[8] In general, the penalty on deceptive claims (CE case) can be different from the penalty of violating privacy (MS case). We assume the penalty in these two cases are the same for simplicity and without loss of generality.

As noted above, the government can set F sufficiently high so that retailers who claim to protect privacy will, in fact, protect privacy, and we assume that the government desires to do so. Thus, if the retailer has claimed privacy protection, it will follow through on its promise, and if the retailer has made no such claim it will not (since the cost of protecting privacy is positive).

We now focus consumers' decision whether to read the retailer's privacy policy. Lemma 1 shows that the dominant strategy for all consumers is not to read the retailer's privacy policy when the reading cost is sufficiently high.

LEMMA 1. If the reading cost is sufficiently high, i.e., $R \geq L/4$, the dominant strategy for all consumers is not to read the retailer's privacy policy, given any belief $\delta(p)$.

We note that the dominant strategy for an insensitive consumer is not to read the privacy policy for any positive reading cost, because her utility is the same whether her privacy is protected or not, and thus she has nothing to gain from reading the privacy policy. A sensitive consumer, however, can potentially avoid a loss of L by reading the privacy policy and not purchasing when the retailer does not claim privacy protection.[9] Therefore a sensitive consumer will trade-off the benefit from and cost of reading the privacy policy when she decides whether to read it. In Lemma 1 we show that the benefit from reading the privacy policy has an upper bound of $L/4$. Therefore, a sensitive consumer will not read the privacy policy if $R \geq L/4$.

In the analysis that follows, we focus on the case of consumers having high reading cost ($R \geq L/4$). This assumption is consistent with available anecdotal evidence that very few consumers actually read privacy policies. For example, Regan (2001) reports that less than 1 percent of the visitors to six major online travel sites during April 2001 read privacy policies.

[9] Recall that the penalty F is set high enough such that it is unprofitable for the retailer to make deceptive claims in its privacy policy.

An assumption of high cognitive costs to read and understand privacy policies also generally consistent with an observation of the length and complexity of retailers' posted privacy policies. For example, we collected the privacy policies from the top 20 most popular shopping sites listed at Alexa.com. These privacy policies averaged 2,174 words in length, not including links to other supporting pages.[10] This corresponds roughly to 4 pages of single spaced typewritten text. Moreover, many privacy statements tend to use euphemisms for how customer data may be used,[11] which contributes to a reduced level of "opting out." To count the complexity of the privacy policies we examined, we use Microsoft Word's grammar checking feature and found the average *flesch-kincaid* grade level score for these privacy policies was 10.6, well above the recommended grade level for most standard documents.[12]

From Lemma 1 and our assumptions on the reading cost faced by consumers, proposition 2 summarizes these results in the *caveat emptor* game.

PROPOSITION 2. When the reading cost is sufficiently high, i.e., $R \geq L/4$, there exists a unique Perfect Bayesian Nash Equilibrium. In this equilibrium, all consumers have a belief that the retailer does not claim privacy protection in its privacy policy with probability one, i.e., $\delta(p) = 1$, consumers do not read the privacy policy, the retailer does not claim to protect consumers' privacy, and in fact does not protect consumer privacy. The retailer sets a price of $p^* = (1 - \rho L)/2$ and obtains a profit of $\pi^* = (1 - \rho L)^2/4$ in this equilibrium.

3.4. Seal-of-Approval Programs

In this regime, we assume the existence of a third party that allows retailers to display a "seal" logo, in exchange for a commitment to abide by a set of privacy policies in dealing with consumer infor-

[10] For example, while Amazon.com's privacy policy is 2,396 words long, it also includes links to 28 other pages, many of which are important to understand the implications of the privacy policy statements.
[11] For example, notices sometimes refer to *making* customer lists available to others rather than pointing out that the lists were being *rented* (Smith, 2001).
[12] MS Word's help file suggests that a normal document should have a grade level between 6 and 7.

mation. We further assume that the seal granting authority chooses policies to maximize social welfare, charges a fee to members, polices members actions, and imposes a penalty on members who violate the stated privacy policies.

Both the caveat emptor regime and seal-of-approval requires the retailer to post a privacy policy on its site, that is, to disclose their privacy practice through privacy policy. In this sense, both of them are special forms of information regulation (Magat and Viscusi 1992, Breyer 1993, Sustein 1999). The seal-of-approval programs present an interesting supplement to traditional forms of information regulation by aiding consumers in the process of information disclosure. In effect, the logo serves as an easy to understand proxy for the privacy standards adopted by the site, while the penalty imposed by the SOA granting authority on sites that violate these policies guarantees the credibility of these statements. The idea is to make the structure of the information readily processed by the individuals receiving it (Magat and Viscusi 1992) — seals convey the information.

We use a five-stage game to capture the behavior of the retailer and consumers under the seal-of-approval program regime. In stage 1, nature decides the retailer's cost of protecting privacy, and this information is revealed only to the retailer. In stage 2, the retailer posts the product's selling price and decides whether to join a seal-of-approval program. If the retailer joins a seal-of-approval program, it pays a per-transaction fee *t*. In stage 3, consumers observe the posted price and whether the retailer has joined a seal-of-approval program and form beliefs regarding the retailer's type. Consumers then decide whether to make purchases based on the information available to them.
In stage 4, the retailer decides whether to protect the privacy of consumers who have made purchases. In stage 5, if the retailer has joined a seal-of-approval program but does not protect consumer privacy, this violation is detected by the seal granting authority with probability α $(0 < \alpha \leq 1)$. Once

[14] The enforcement mechanism is different under SOA from those of the CE and MS regimes. Under SOA, the SOA authority takes extra care to enforce compliance before taking a violator to court, which includes escalated investigation and terminating the violator's membership.

a violation is detected, the enforcement agency penalizes the retailer with a penalty of M.[14] After this stage the retailer's profit and consumers' utilities are realized. The penalty imposed by the regulation authority differentiates SOA from caveat emptor—it regulates the privacy protection behavior of the licensees. However, unlike mandatory standards, whether a site joins the regulatory regime is purely voluntary. Table 2 summarizes the setup of the game.

Stage 1	***Stage 2***	***Stage 3***	***Stage 4***	***Stage 5***
Nature decides retailer's cost of protecting privacy. This information is revealed only to the retailer.	The retailer decides whether to join a seal-of-approval program and its price.	Consumers observe whether the retailer joins a seal-of-approval program and the price, then decide whether to purchase.	The retailer decides whether to protect consumer privacy.	The seal-granting authority monitors the licensee's protection, and penalizes the licensee if it has not protected privacy.

Table 2: Timing of the seal-of-approval programs game

As above, we assume that the seal granting authority has the incentive to set the penalty M high enough such that joining the seal-of-approval program is a credible signal that a retailer who has done so will choose to protect consumer privacy whether it is an H-type retailer or an L-type retailer. This means that the penalty αM is set higher than both the H-type and L-type retailer's cost of protecting privacy. That is, $c_L < c_H \leq \alpha M$. Obviously, if a retailer has not joined a seal-of-approval program it will choose not to protect privacy because of the positive cost of protecting consumer privacy.

Given the analysis above, consumers will use the retailer's decision to join a seal-of-approval program as a signal of whether her privacy will be protected. If the retailer has joined a seal-of-approval program, both the sensitive and insensitive consumers will obtain a utility of $v - p$ if a purchase is

made at price p, and a zero utility if a purchase is not made. The retailer faces a demand function of $D(p) = \rho(1-p) + (1-\rho)(1-p) = 1-p$.

If the retailer has not joined a seal-of-approval program, a sensitive consumer will obtain a utility of $v-p-L$ if a purchase is made at price p, and a zero utility if a purchase is not made; an insensitive consumer will obtain a utility of $v-p$ if a purchase is made at price p, and a zero utility if a purchase is not made. The retailer faces a demand function of $D(p) = \rho(1-p-L) + (1-\rho)(1-p) = 1-p-\rho L$.

In stage 2 when the retailer decides whether to join a seal-of-approval program, it compares the profit it obtains from joining to the profit from not doing so and chooses an action to maximize its profit. We focus on how the seal-of-approval regime differs from the caveat emptor and mandatory standards regimes. The latter two regimes lead to equilibrium results in which an L-type retailer and an H-type retailer choose the same action of privacy protection — neither type protects privacy under the regime of caveat emptor while both types protect privacy under the regime of mandatory standards. However, the seal-of-approval regime can lead to a separating equilibrium result, in which an L-type retailer and an H-type retailer choose different actions of privacy protection — an L-type retailer chooses to join a seal-of-approval program and to protect privacy while an H-type retailer chooses not to join a seal-of-approval program and not to protect privacy. We assume that the seal granting authority will set up its program in a way such that this separating equilibrium result is obtained.[15] Proposition 3 gives the condition for the existence of this separating equilibrium.

PROPOSITION 3. When $c_L + t < \rho L < c_H + t$, there exists a unique separating equilibrium. In this equilibrium, the retailer will choose to join a seal-of-approval program and to protect privacy if it is

[15] If pooling equilibria are obtained, SOA converts to either MS or CE.

an L-type retailer; it will choose not to join a seal-of-approval program and not to protect privacy if it is an H-type retailer.

It follows from Proposition 3 that a separating equilibrium does not exist when $\rho L \le c_L$. This means that when the expected loss a consumer would face from a privacy violation is less than the cost of protecting privacy to the low type firm, a separating equilibrium cannot be achieved no matter how the seal granting authority sets its membership fee *t*. This proposition also sets upper and lower bounds on the membership fee that the seal granting authority can charge in order to achieve a separating equilibrium. The upper and lower bounds depend on the proportion of sensitive consumers (ρ), the sensitive consumers' loss if their privacy is not protected (*L*), and cost of protecting privacy (c_L, c_H).

From the SOA's point of view, two necessary conditions to guarantee the seals' effectiveness are: first, there ought to be unambiguous privacy protections standards for the licensees to include in their privacy policies; second, the mechanism of enforcement and redress has to be in place. That is, if a licensee violates the SOA's stated privacy policies, the SOA can penalize the licensee, and this penalty, when set appropriately, will ensure that it will never be optimal for the licensees to do so.

4. Welfare Implications

To compare the consumer and producer surplus under each regime, we assume that the retailer can be either L-type with probability $\lambda \in [0,1]$ or H-type with probability $1 - \lambda$. We then calculate the expected consumer and producer surplus under each regime for the three equilibrium regions outlined in proposition 3. In doing so, we first summarize the optimal prices under each regime.

Under caveat emptor, the L-type and H-type retailer charges the same price $p_L^{CE} = p_H^{CE} = (1 - \rho L)/2$; Under seal-of-approval program, $p_L^{SOA} = (1 + c_L + t)/2$ and

$p_H^{SOA} = (1-\rho L)/2$; Under mandatory standards, $p_L^{MS} = (1+c_L)/2$ and $p_H^{MS} = (1+c_H)/2$. Notice that the L-type retailer charges the lowest price under caveat emptor. Because under this regime, the retailer does not protect consumer privacy, it has to compensate consumers through a lower price. The L-type retailer charges the highest price under seal-of-approval programs, because it will transfer a part of the protection cost and membership fee to consumers. Under mandatory standards, the L-type retailer charges a price that is between the price under the other two regimes. The H-type retailer does not join seal-of-approval programs and does not protect consumer privacy, so it charges the same price under caveat emptor and under seal-of-approval programs. The H-type retailer charges a higher price under mandatory standards than under the other two regimes, because it will transfer a part of the protection cost to consumers.

4.1. When $\rho L \le c_L < c_H$

In this case, no separating equilibrium exists under the seal-of-approval regime. Rather, under seal-of-approval, neither type will join the seal program because the protection cost does not justify gain to the retailer from joining. Thus, it is sufficient to only compare two regimes: caveat emptor and mandatory standards. Table 3 shows the calculation of expected consumer surplus and producer surplus under these two regimes.

	Expected Consumer Surplus	Expected Producer Surplus
Mandatory Standards	$[\lambda(1-c_L)^2 + (1-\lambda)(1-c_H)^2]/8$	$[\lambda(1-c_L)^2 + (1-\lambda)(1-c_H)^2]/4$
Caveat Emptor	$(1-\rho L)^2/8 + \rho L^2(1-\rho)/2$	$(1-\rho L)^2/4$

Table 3: Expected consumer surplus and producer surplus when $\rho L \le c_L < c_H$

Because $\rho L \le c_L < c_H$, we have $(1-\rho L)^2 \ge (1-c_L)^2 > (1-c_H)^2$. Thus, consumer and producer surplus are both higher under caveat emptor than under mandatory standards. The optimal regime when $\rho L \le c_L < c_H$ is caveat emptor.

4.2. When $c_L < \rho L < c_H$

In this case, a separating equilibrium exists under seal-of-approval programs. Because, as noted above, the seal granting authority is assumed to be an organization whose goal is to improve social welfare, and the membership fee *t* is negatively correlated with both consumer surplus and producer surplus, we can assume that the seal granting authority tries to minimize the membership fee.[16] When $c_L < \rho L < c_H$, the lower bound on the membership fee is negative. Thus, the membership fee that maximizes either consumer surplus or producer surplus is zero. Table 4 shows the calculation of expected consumer surplus and producer surplus under these three regimes.

Because $c_L < \rho L$, we have $(1-c_L)^2 > (1-\rho L)^2$. Thus, producer surplus is higher under the seal-of-approval programs than it is under caveat emptor. Because $\rho L < c_H$, we have $(1-\rho L)^2 > (1-c_H)^2$. Thus, producer surplus is higher under seal-of-approval programs than it is under mandatory standards, and the seal-of-approval regime maximizes producer surplus when $c_L < \rho L < c_H$. The comparison of consumer surplus is more complex. We can show (see Appendix) that consumer surplus is higher under seal-of-approval programs than under mandatory standards, but whether the seal-of-approval or caveat emptor regime maximizes consumer surplus depends on the proportion of sensitive consumers (ρ) and sensitive consumers' loss if their privacy is not protected (L).

[16] Note that because the membership fee is negatively correlated with both consumer and producer surplus, we would obtain the same result if the seal granting authority's goal was to maximize consumer surplus or producer surplus.

	Expected Consumer Surplus	Expected Producer Surplus
Mandatory Standards	$[\lambda(1-c_L)^2 + (1-\lambda)(1-c_H)^2]/8$	$[\lambda(1-c_L)^2 + (1-\lambda)(1-c_H)^2]/4$
Caveat Emptor	$(1-\rho L)^2/8 + \rho L^2(1-\rho)/2$	$(1-\rho L)^2/4$
Seal-of-approval programs	$\lambda(1-c_L)^2/8 + (1-\lambda)(1-\rho L)^2/8$ $+ (1-\lambda)\rho L^2(1-\rho)/2$	$\lambda(1-c_L)^2/4 + (1-\lambda)(1-\rho L)^2]/4$

Table 4: Expected consumer surplus and producer surplus when $c_L < \rho L < c_H$

4.3. When $c_L < c_H \leq \rho L$

In this case, a separating equilibrium exists under seal-of-approval programs. As above, because the seal granting authority seeks to maximize social welfare and the membership fee *t* is negatively correlated with both consumer and producer surplus, the seal granting authority will try to minimize the membership fee *t* while still maintaining a separating equilibrium. The lower bound on the membership fee is positive in this case. Thus, the membership fee that maximizes either consumer surplus or producer surplus is set at the lower bound $t^* = \rho L - c_H$.

	Expected Consumer Surplus	Expected Producer Surplus
Mandatory Standards	$[\lambda(1-c_L)^2 + (1-\lambda)(1-c_H)^2]/8$	$[\lambda(1-c_L)^2 + (1-\lambda)(1-c_H)^2]/4$
Caveat Emptor	$(1-\rho L)^2/8 + \rho L^2(1-\rho)/2$	$(1-\rho L)^2/4$
Seal-of-approval programs	$\lambda(1-\rho L-c_L+c_H)^2/8$ $+ (1-\lambda)(1-\rho L)^2/8$ $+ (1-\lambda)\rho L^2(1-\rho)/2$	$[\lambda(1-\rho L-c_L+c_H)^2$ $+ (1-\lambda)(1-\rho L)^2]/4$

Table 5: Expected consumer surplus and producer surplus when $c_L < c_H \leq \rho L$

Because $c_L < \rho L$ and $c_H < \rho L$, we have $(1-c_L)^2 > (1-\rho L)^2$ and $(1-c_H)^2 > (1-\rho L)^2$. Because $c_H < \rho L$, we have $(1-c_L)^2 > (1-\rho L-c_L+c_H)^2$ and $(1-c_H)^2 > (1-\rho L)^2$. Thus, producer surplus

is maximized under the mandatory standards regime. The comparison of consumer surplus is more complex. Here any of the three regimes can maximize consumer surplus, depending on the proportion of sensitive consumers (ρ), the sensitive consumers' loss if their privacy is not protected (L), and the probability of the retailer being L-type (λ).

4.4. Summary of producer surplus, consumer surplus, and total welfare

Figures 3, 4, and 5 summarize the optimality of privacy protection regimes when producer surplus, consumer surplus and total welfare are respectively used as the criteria (without loss of generality, in these figures we assume c_L=0.1, c_H=0.4, and $\lambda = 0.5$).[17] In these figures, the x-axis graphs values of ρ, and the y-axis values of L. The regions where the different privacy protection regimes dominate are separated by solid lines. And the dashed lines in Figure 4 and 5 are $\rho L = c_L$ and $\rho L = c_H$. Figure 3 shows that when $\rho L \le c_L < c_H$, the *caveat emptor* (CE) regime maximizes producer surplus. The mandatory standards (MS) regime produces the highest producer surplus for sufficiently high ρ and L. In the intermediate region the seal-of-approval (SOA) regime maximizes producer surplus.

Figure 4 shows the optimal privacy protection regime when consumer surplus is used as the criteria. This figure is more complicated than Figure 3 where producer surplus is used as the criteria. Notably, the line that separates caveat emptor regime and seal-of-approval regime curves back such that the caveat emptor regime can still be the optimal regime even when both ρ and L are moderately high, whereas in Figure 3, seal-of-approval regime is the optimal regime when both ρ and L are moderately high.

Figure 3 is relatively simple, because producer surplus is a decreasing function of both ρ and L. However, consumer surplus is not monotonic with regard to either ρ or L, as shown by the formulae

[17] The proportion of L-type retailer λ will only change the borderlines that divide CE, MS, and SOA when $\rho L \ge c_H$, without changing the relative position of CE, MS, and SOA.

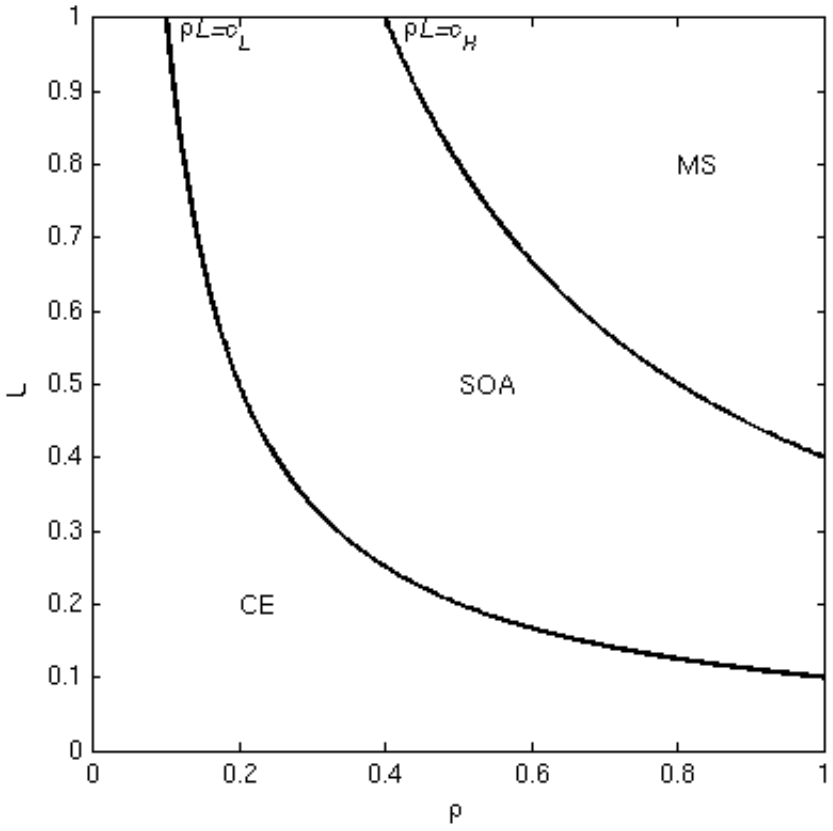

Figure 3: Comparison of producer surplus ($c_L = 0.1,\ c_H = 0.4, \lambda = 0.5$)

of expected consumer surplus under different regimes in Table 4. For example, an increase of L can potentially influence consumer surplus in both ways: it can lead to a lower price charged by the L-type retailer which boosts consumer surplus, but it can also make sensitive consumers less likely to make purchase which negatively affects consumer surplus. As a result of these two competing effects, consumer surplus is not monotonic with regard to L. This non-monotonicity leads to the added complexity of Figure 4.

Figure 5 shows the optimal privacy protection regime when total welfare is used as the criteria. In this figure we see that for sufficiently small ρ or L, caveat emptor produces the highest total welfare; for sufficiently large ρ and L mandatory standards dominates; and for intermediate values seal-of-approval programs produces the highest total welfare. Figure 5 looks more similar to Figure 3 than to

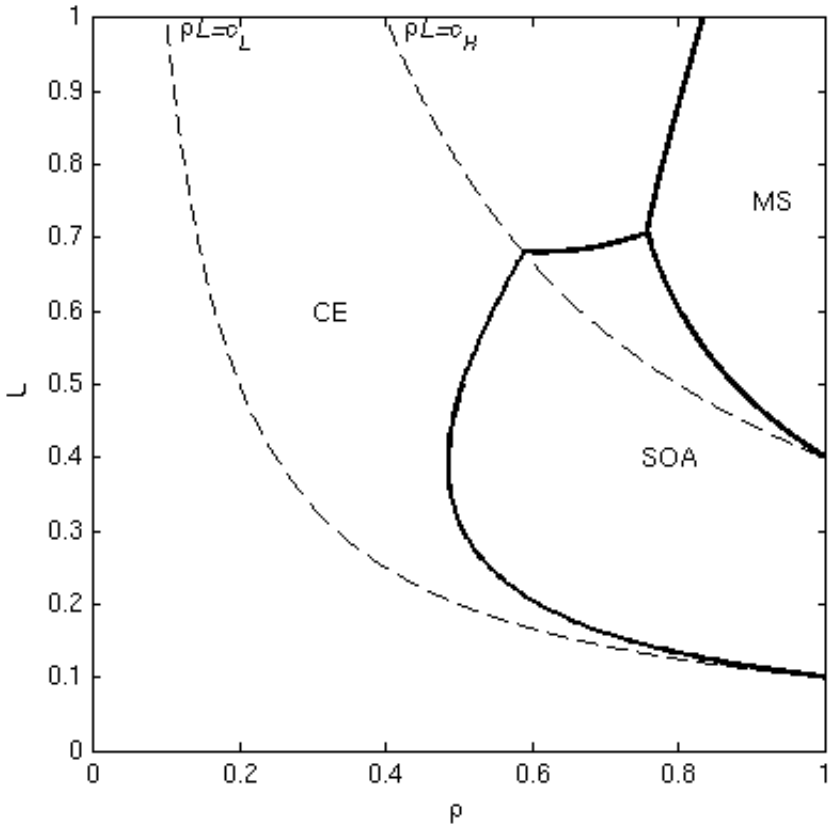

Figure 4: Comparison of consumer surplus ($c_L = 0.1$, $c_H = 0.4$, $\lambda = 0.5$)

Figure 4. This is because producer surplus is a larger component in total welfare than consumer surplus is. We discuss the implications of this finding below.

5. Implications and Concluding Remarks

In this paper we have developed a model of asymmetric information in which an online retailer has private information regarding its own cost of protecting consumers' privacy. We analyze the retailer's strategy under three privacy protection regimes that might be chosen by government regulators or market designers: mandatory standards, *caveat emptor*, and seal-of-approval programs. Under a mandatory standards approach, the retailer protects consumer privacy as a result of high penalty and strong consumer confidence. Under *caveat emptor*, however, the retailer does not protect con

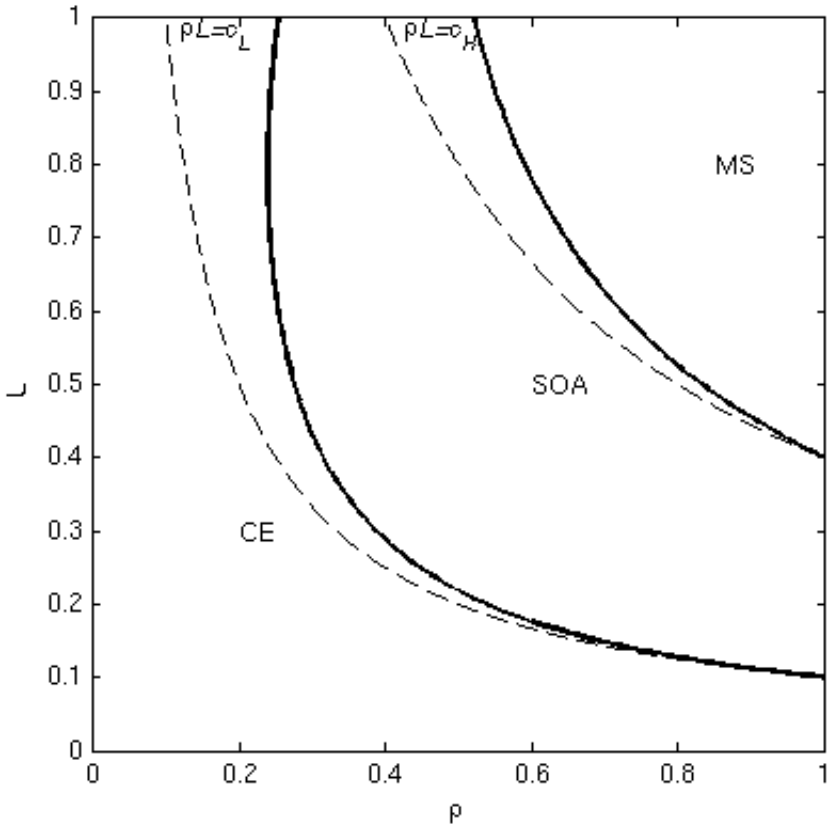

Figure 5: Comparison of Total Welfare ($c_L = 0.1, c_H = 0.4, \lambda = 0.5$)

sumer privacy, because consumers do not read privacy statement, therefore the retailer does not have incentive to protect privacy. Under seal-of-approval programs, the retailer can send a signal to consumers about its cost of protecting their privacy by setting its price and by joining a seal-of-approval program. If the seal granting authority chooses its membership fee and penalty appropriately, joining the seal-of-approval program can serve as a credible signal that a retailer will protect a consumer's privacy. In addition, consumers can infer a retailer's cost of protecting privacy by observing the retailer's price and its decision to join the seal-of-approval program. A retailer with a low cost of protecting a consumer's privacy will find it profitable to join the seal-of-approval program, while a retailer with a high cost of protecting a consumer's privacy will not find it profitable to do so.

We then compare the social optimality of industry self-regulation by analyzing the producer and consumer surplus and total welfare generated under each program for varying levels of sensitive customers and varying levels of consumer loss when privacy is violated. Our results show that for types of information where many consumers care about privacy protection or where consumers face a large potential utility loss from a privacy violation, mandatory standards yield the largest total welfare. This is consistent with the U.S. government enacting mandatory standards of privacy protection for specific types of transaction data such as credit information, health information, and marketing information gathered from minors.

At the other extreme, when few people care about privacy protection or when their potential loss is small, *caveat emptor* is the socially optimal form of privacy protection. This may seem surprising given that this outcome results in a low level of privacy protection. However, it is important to note that consumers benefit in this case through lower prices and retailers benefit by not incurring the cost of protecting privacy. This is consistent with the general *liaise faire* attitude of the U.S. Federal Trade Commission toward regulating privacy protection for many common types of online transaction data.

For intermediate values of ρ and L, seal-of-approval programs can produce the socially optimal privacy protection regime. In this case, seal-of-approval programs allow retailers with a low cost of protecting privacy to protect consumer privacy, catering to the preferences of sensitive consumers. At the same time seal-of-approval programs allow retailers with a high cost of protecting privacy to not protect consumer privacy but charge lower prices, catering to the preferences of insensitive consumers and saving cost of protecting privacy for the retailer. Again, this is generally consistent with the development of seal-of-approval programs for some types of transaction data such as consumer purchase or targeting information.

Our analysis offers strategic insights for a variety of audiences including third party organizations overseeing industry's practice of protecting privacy; businesses upholding guidelines, standards, and practices of privacy protection; and government agencies administering online privacy protection. For third-party organizations, our results suggest that industry self-regulation can be accomplished through seal-of-approval programs if the fees and penalties associated with these programs are chosen effectively. For businesses our results suggest that in the presence of heterogeneous consumer preferences toward privacy protection, differentiation strategies may be effective where some firms provide high levels of privacy protection in exchange for higher prices while other businesses essentially compensate consumers for use of their sensitive data through low prices. Finally for regulatory agencies and other market designers, our results suggest that not all types of data warrant the same approach to privacy protection and thus the model adopted in the United States where privacy protections are enacted on an industry-by-industry basis may be more socially efficient than overarching approaches to privacy protection, similar to those adopted in the European Union, where all types of data are covered by the same set of mandatory standards.

While we do not incorporate the case of duopolistic retailers, our setting of a monopolistic retailer is common in the economic literature of signaling (Kihlstrom and Riordan 1984, Milgrom and Roberts 1986, Wernerfelt 1988). It is also consistent with the literature with respect to the economics of privacy (Chellappa and Shivendu 2003, Akcura and Srinivasan 2003).

References

Acquisti, A., H. R. Varian. 2002. Conditioning Prices on Purchase History. Working Paper, University of California, Berkeley.

Akcura, T. M., K. Srinivasan. 2003. Strategic Role of Customer Intimacy. Working Paper, Purdue University.

Banks, D.T., Hutchinson, J.W., and Meyer, R.J. 2002. Reputation in Marketing Channels: Repeated-Transactions Bargaining with Two-Sided Uncertainty. *Marketing Science*, 21(3): 251-272.

Baron, D. P. 2000. DoubleClick and Internet Privacy. Stanford University Case Number P-32, August.

Benassi, Paola. 1999. TRUSTe: An Online Privacy Seal Program. Communications of the ACM, 42(2): 56-59.

Breyer, S. 1993. Breaking the Vicious Circle : Toward Effective Risk Regulation. Harvard University Press, 1993.

Chellappa, R. K., S. Shivendu. 2003. Managing Piracy: Pricing and Sampling Strategies for Digital Experience Goods in Vertically Segmented Markets. Working Paper, University of Southern California.

Chen, Y., K. Sudhir. 2001. When Shopbots Meet Emails: Implications for Price Competition on the Internet. Working Paper, New York University.

Dreazen, Y. J. 2003. Consumers Are in the Dark on Web-Site Privacy—Study Offers Ammunition To Backers of Curbs on How Personal Data Can Be Used. *Wall Street Journal*, Section D, Page 2, Column 1.

Ellison, G, Ellison, S.F. 2004. Search, Obfuscation, and Price Elasticities on the Internet. NBER Working Paper No. W10570.

Ghose, A., P. Y. Chen. 2003. Personalization and Information Privacy: Retailer Practices, Business Profits and Consumer Welfare. Working Paper, GSIA, Carnegie Mellon University.

Green, E.J. and Porter, R.H. 1984. Noncooperative Collusion under Imperfect Price Information. *Econometrica.* 52(1): 87-100.

Hann, I., K. Hui, T. S. Lee, I. P. L. Png. 2003. Direct Marketing: Privacy and Competition. 2003. Working Paper, National University of Singapore.

Hann, I., T. S. Lee, K. Hui, I. P. L. Png. 2002. Online Information Privacy: Measuring the Cost-Benefit Trade-Off. *Proceedings of 23rd International Conference on Information Systems*.

Hahn, R. W., A. Layne-Farrar. 2001. The Benefits and Costs of Online Privacy Legislation. Working Paper 01-14, AEI-Brooking Joint Center for Regulatory Studies.

Hoffman, D. L., T. P. Novak, M. Peralta. 1999. Building Consumer Trust Online. *Communications of the ACM*. 42(4): 80-85.

Karkkainen, B.C. 2001. Information as Environmental Regulation: TRI and Performance Benchmarking, Precursor to a New Paradigm? 89 *Georgetown Law Journal*, 257.

Kihlstrom, R. E., M. H. Riordan. 1984. Advertising as a Signal. *Journal of Political Economy*, 92(3): 427-450.

Laudon, K. C. 1996. Markets and Privacy. *Communications of the ACM*, 39(9): 92-104.

MacLeod, W.B. and Malcomson J.M. 1988. Reputation and Hierarchy in Dynamic Models of Employment. *The Journal of Political Economy,* 96(4): 832-854.

Magat, W. A. and Viscusi, W. K. 1992. Informational approaches to regulation. MIT Press.

Milgrom, P., J. Roberts. 1986. Price and Advertising Signals of Product Quality. *Journal of Political Economy*, 94(4): 796-821.

Miyazaki, A. D., A. Fernandez. 2001. Consumer Perceptions of Privacy and Security Risks for Online Shopping. *Journal of Consumer Affairs*, 35(1): 27-44.

Noldeke, G. and Van Damme, E. Signalling in a Dynamic Labour Market. *The Review of Economic Studies*, 57(1): 1-23.

Pavlou, P. A., R. K. Chellappa. 2001. The Role of Perceived Privacy and Perceived Security in the Development of Trust in Electronic Commerce Transactions. Working Paper, University of Southern California.

Regan, K. 2001. Does Anyone Read Privacy Policies? E-commerce Times, June 15. (available at http://www.ecommercetimes.com/story/11303.html)

Ryan, J.P. 2000. Privacy, the Common Good, and Individual Liberties in the 21st Century: A Dialogue on Policy, Law, and Values. *American Bar Association Focus*, 15(2).

Seligman, T. J., J. D. Taylor. 2000. FTC Reverses Privacy Policy. *New York Law Journal*, June 19.

Shin, J. 2003. The Role of Selling Cost in Signaling Price Image: Theory and Evidence from the Travel Industry. Working Paper, Sloan School of Management, MIT.

Smith, H. J. 2001. Information Privacy and Marketing: What the U.S. Should (and Shouldn't) Learn from Europe. *California Management Review*, 41(2): 8-33.

Sunstein, C.R. 1999. Informational Regulation and Informational Standing: *Akins* and Beyond. 147 *University of Pennsylvania Law Review* 613.

Taylor, C. R. 2002. Private Demands and Demands for Privacy: Dynamic Pricing and the Market for Customer Information. Working Paper, Department of Economics, Duke University.

Varian,H.R. 1997. Economic Aspects of Personal Privacy. In *Privacy and Self-Regulation in the Information Age*, National Telecommunications and Information Administration: Washington, D.C.

Vila, T., Greenstadt, R., and Molnar, D. 2003. Why We Cannot Be Bothered to Read Privacy Policies: Models of Privacy Economics as Lemons Market. *Proceedings of the 5th international conference on Electronic commerce*. Pittsburgh, Pennsylvania Pages: 403-407.

Wernerfelt, B. 1988. Umbrella Branding as a Signal of New Product Quality: An Example of Signaling by Posting a Bond. *RAND Journal of Economics*, 29(3): 458-466.

Westin, A. F. 1967. *Privacy and Freedom*. Atheneum, New York, NY.

Appendix

Proof of Proposition 1

The retailer's profit function is $\pi = (1 - p_i)(p_i - c_i)$, where $i = L, H$ and stands for the retailer's type. The optimal price that maximizes the retailer's profit is $p_i^* = (1 + c_i)/2$, and this leads to a profit of $\pi_i^* = (1 - c_i)^2/4$.

Proof of Lemma 1

If a consumer who is sensitive to privacy protection reads the privacy policy and knows the retailer claims to protect privacy, her utility is $\max\{v - p, 0\} - R$. Likewise, if this sensitive consumer reads the privacy policy and knows the retailer does not claim to protect privacy, her utility is $\max\{v - p - L, 0\} - R$. Thus, ex ante, a sensitive consumer's expected utility from reading the privacy policy is $\delta \max\{v - p - L, 0\} + (1 - \delta)\max\{v - p, 0\} - R$.

Simplifying, a sensitive consumer's expected utility from not reading the privacy policy is $\max\{v - p - \delta L, 0\}$.

Sensitive consumers can be categorized into four types that are exhaustive and mutually exclusive, based on their valuation for the product.

Type-1: $v - p \geq L$

Such a consumer's expected utility from reading the privacy policy is $v - p - \delta L - R$, while her expected utility from not reading the privacy policy is $v - p - \delta L$. Thus, the dominant strategy when $v - p \geq L$ (i.e., type-1), is not reading the privacy policy, if $R \geq 0$.

Type-2: $\delta L \leq v - p < L$

Such a consumer's expected utility from reading the privacy policy is $(1 - \delta)(v - p) - R$, while her expected utility from not reading the privacy policy is $v - p - \delta L$. Let R^* be the threshold

value of reading cost that makes such a consumer indifferent between reading and not reading. Thus we have $R^* = \delta[L-(v-p)] \le \delta(1-\delta)L \le L/4$. A sufficient condition for a type-2 consumer to not read the privacy policy is $R \ge L/4$.

Type-3: $0 \le v - p < \delta L$

Such a consumer's expected utility from reading the privacy policy is $(1-\delta)(v-p) - R$, while her expected utility from not reading the privacy policy is 0. Let R^{**} be the threshold value of reading cost that makes such a consumer indifferent between reading and not reading. Then we have $R^{**} = (1-\delta)(v-p) < (1-\delta)\delta L \le L/4$. Therefore, a sufficient condition for a type-3 consumer to not read the privacy policy is $R \ge L/4$.

Type-4: $v - p < 0$

Such a consumer's expected utility from reading the privacy policy is $-R$, while her expected utility from not reading the privacy policy is 0. Thus, the dominant strategy for a type-4 consumer is not to read the privacy policy, if $R \ge 0$.

From this discussion, if the reading cost is sufficiently high, i.e., $R \ge L/4$, the dominant strategy for all consumers is not reading the privacy policy, given any belief $\delta \in [0,1]$.

Proof of Proposition 2

Because no consumers will read the privacy policy (Lemma 1), the retailer faces a demand of $\rho(1-p-\delta L) + (1-\rho)(1-p)$ for its product from both insensitive consumers and sensitive consumers, whether the retailer claims to protect privacy or not. If the retailer claims to protect privacy, its profit is $\pi = (1-p-\delta\rho L)(p-c_i)$; if the retailer does not claim to protect privacy, its profit is $\pi = (1-p-\delta\rho L)p$. Thus the dominant strategy for the retailer is not to claim that it protects privacy. This strategy of the retailer in turn supports a consumer belief of $\delta = 1$. Therefore the equilib-

rium in which consumers do not read the privacy policy and the retailer does not claim to protect privacy is a Perfect Bayesian Nash Equilibrium when the reading cost is sufficiently high, i.e., $R \geq L/4$. In this equilibrium, the retailer sets a price of $p^* = (1-\rho L)/2$, and obtains a profit of $\pi^* = (1-\rho L)^2/4$.

Proof of Proposition 3

If the retailer joins a seal-of-approval program and charges a price *p*, consumers, according to the definition of sequential rationality, will infer that the retailer will protect their privacy. Because their privacy is protected, both sensitive and insensitive consumers will obtain a utility of $v-p$ if a purchase is made at price *p*, and zero utility if no purchase is made. The retailer will maximize her profit function of $\pi = (1-p)(p-c_i-t)$. Thus, the optimal price is $p^* = (1+c_i+t)/2$ and the profit is $\pi^* = (1-c_i-t)^2/4$.

If the retailer does not join a seal-of-approval program, consumers in segment *S*, according to the definition of sequential rationality, will infer that the retailer will not protect their privacy. Therefore, a sensitive consumer will obtain a utility of $v-p-L$ if a purchase is made at price *p*, and zero utility if no purchase is made. An insensitive consumer will obtain a utility of $v-p$ if a purchase is made at price *p*, and a zero utility if no purchase is made. The retailer faces a demand of $\rho(1-p-L)+(1-\rho)(1-p)$. The retailer will maximize a profit function of $\pi = (1-p-\rho L)p$. The optimal price the retailer charges is $p^* = (1-\rho L)/2$ and the profit is $\pi^* = (1-\rho L)^2/4$.

The incentive compatibility constraints are

$$(1-c_L-t)^2/4 > (1-\rho L)^2/4 \qquad (IC\text{-}L)$$

$$(1-\rho L)^2/4 > (1-c_H-t)^2/4 \qquad (IC\text{-}H)$$

The two IC constrains are satisfied simultaneously when $c_L + t < \rho L < c_H + t$ holds.

Calculation of consumer surplus and producer surplus under each regime

If the retailer is an L-type, recall that the optimal prices under the three regimes are:

$p_L^{CE} = (1-\rho L)/2$ for the *caveat emptor* game, $p_L^{SOA} = (1+c_L+t)/2$ for the seal-of-approval program game, and $p_L^{MS} = (1+c_L)/2$ for the mandatory standards game. The consumer surplus under *caveat emptor* is

$$CS_L^{CE} = \rho \int_{v \geq p_L^{CE}+L}^{1} (v - p_L^{CE} - L)dv + (1-\rho) \int_{v \geq p_L^{CE}}^{1} (v - p_L^{CE})dv = (1-\rho L)^2/8 + \rho L^2(1-\rho)/2.$$

The producer surplus under *caveat emptor* is $PS_L^{CE} = (1-\rho L)^2/4$.

The consumer surplus under a seal-of-approval program is

$$CS_L^{SOA} = \rho \int_{v \geq p_L^{SOA}}^{1} (v - p_L^{SOA})dv + (1-\rho) \int_{v \geq p_L^{SOA}}^{1} (v - p_L^{SOA})dv = (1-c_L-t)^2/8.$$

The producer surplus under a seal-of-approval program is $PS_L^{seal} = (1-c_L-t)^2/4$.

The consumer surplus under mandatory standards is

$$CS_L^{MS} = \int_{v \geq p_L^{MS}}^{1} (v - p_L^{MS})dv = (1-c_L)^2/8.$$

The producer surplus under mandatory standards is $PS_L^{MS} = (1-c_L)^2/4$.

If the retailer is an H-type, the optimal prices under the three regimes are: $p_H^{CE} = (1-\rho L)/2$ for the *caveat emptor* game, $p_H^{SOA} = (1-\rho L)/2$ for the seal-of-approval programs game, and

$p_H^{MS} = (1 + c_H)/2$ for the mandatory standards game. The consumer surplus under *caveat emptor* is

$$CS_H^{CE} = \rho \int_{v \geq p_H^{CE} + L}^{1} (v - p_H^{CE} - L) dv + (1 - \rho) \int_{v \geq p_H^{CE}}^{1} (v - p_H^{CE}) dv = (1 - \rho L)^2 / 8 + \rho L^2 (1 - \rho)/2 .$$

Producer surplus under *caveat emptor* is $PS_H^{CE} = (1 - \rho L)^2 / 4$.

The consumer surplus under seal-of-approval programs is

$$CS_H^{SOA} = \rho \int_{v \geq p_H^{SOA} + L}^{1} (v - p_H^{SOA}) dv + (1 - \rho) \int_{v \geq p_H^{SOA}}^{1} (v - p_H^{SOA}) dv = (1 - \rho L)^2 / 8 + \rho L^2 (1 - \rho)/2 .$$

The producer surplus under seal-of-approval programs is $PS_H^{SOA} = (1 - \rho L)^2 / 4$.

The consumer surplus under mandatory standards is

$$CS_H^{MS} = \int_{v \geq p_H^{MS}}^{1} (v - p_H^{MS}) dv = (1 - c_H)^2 / 8 .$$

The producer surplus under mandatory standards is $PS_H^{MS} = (1 - c_H)^2 / 4$.