

9-1-2007

Gaining Trust Through Online Privacy Protection: Self-Regulation, Mandatory Standards, or Caveat Emptor

Zhulei Tang
Purdue University

Yu Jeffrey Hu
Purdue University

Michael D. Smith
Carnegie Mellon University, mds@andrew.cmu.edu

Follow this and additional works at: <http://repository.cmu.edu/heinzworks>

Published In

Journal of Management Information Systems, 24, 4, 153-173.

This Article is brought to you for free and open access by the Heinz College at Research Showcase @ CMU. It has been accepted for inclusion in Heinz College Research by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

**Gaining Trust Through Online Privacy Protection:
Self-Regulation, Mandatory Standards, or *Caveat Emptor***

Zhulei Tang*, Yu (Jeffrey) Hu**, Michael D. Smith***

Forthcoming in *Journal of Management Information Systems*

ABSTRACT

Trust is particularly important in online markets to facilitate the transfer of sensitive consumer information to online retailers. In electronic markets, various proposals have been made to facilitate these information transfers. We develop analytic models of hidden information to analyze the effectiveness of these regimes to build trust and their efficiency in terms of social welfare.

We find that firms' ability to influence consumer beliefs about trust depends on whether firms can send unambiguous signals to consumers regarding their intention of protecting privacy. Ambiguous signals can lead to a breakdown of consumer trust, while the clarity and credibility of the signal under industry self-regulation can lead to enhanced trust and improved social welfare. Our results also indicate that although overarching government regulations can enhance consumer trust, regulation may not be socially optimal in all environments because of lower profit margins for firms and higher prices for consumers.

Keywords: (Trust; Privacy; Consumer Surplus; Social Welfare; Internet; Consumer Information)

* Krannert School of Management, Purdue University; email: zhulei@purdue.edu

** Krannert School of Management, Purdue University; email: yuhu@purdue.edu

*** H. John Heinz III School of Public Policy and Management, Carnegie Mellon University; email: mds@cmu.edu

1. Introduction

One area where trust in online markets is particularly important is the ability of retailers to build trust among consumers so that retailers can adequately address consumers' privacy concerns. Consumers' privacy concerns have become more heightened as information technologies have enabled online retailers to collect increasing amounts of consumer information. 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 could be used in ways that violate consumer privacy. Privacy is the ability to control the acquisition and use of information about one's self [42]. A variety of surveys and experiments have shown that privacy concerns are a leading factor impeding electronic commerce (e.g. [16], [25], [36], and [43]). For example, Consumer Reports WebWatch research report indicates that 90 percent of U.S. Internet users over the age of 18 have changed their behavior because of fear of identity theft [25]. Fundamentally these privacy concerns arise because of a lack of trust between businesses and consumers [16], [39].

Trust is the willingness of one party to be subject to the risks brought by another party's actions ([20], [12], and [27]). In the context of online privacy protection, these risks arise from uncertainty or incomplete information about retailers' actions regarding customer information. For example, security breaches may occur due to inadequate data protection or internal controls; retailers may engage in unauthorized secondary uses, such as using information that is collected for one purpose for a different purpose [8]. In this paper we focus on the risks associated with the

latter — secondary use. In this case, retailers’ profit maximizing behavior often leads to privacy-protection practices that would not be optimal for their customers. [32] have pointed out that uncertainty will lead to two information problems: hidden information and hidden action.

Without the ability to credibly signal their trustworthiness in handling consumer information, retailers will be less able to persuade consumers to share sensitive information, hindering welfare creation for both consumers and retailers and impeding the development of online commerce.

As privacy concerns have been identified as a primary barrier to consumer trust online, governments and third parties have proposed various approaches to privacy protection. At the heart of these approaches are fair information practices that aim to empower consumers with more transparency and control over their information. Fair information practices are a set of standards to guide the adequate collection and use of personal information [9], [33]. According to [33], the core fair information principles include: notice, choice, access, and integrity. Notice or awareness requires informing the scope and usage of personal information collection. Choice or consent gives consumers options to opt out secondary uses of information. Access or participation invites people to access their personal information to ensure accuracy and completeness of the information. Integrity or security requires protection against unauthorized access and use of data. These principles need enforcement or redress mechanisms to monitor compliance. Therefore, enforcement or redress is also part of fair information practice principles.

One can categorize the various privacy protection approaches into three general categories according to how and to what degree these fair information practices are implemented. The first category is *caveat emptor*, literally “let the buyer beware.” Under this approach, retailers are under no obligation to post a privacy notice or to obey fair information practices. However, if

they do post privacy policies, they are required by law to abide by them. This approach applies to many common types of consumer transaction data [2].¹

At the opposite end of the intervention spectrum is the *mandatory standards* approach for privacy protection, where governments intervene and enact strict privacy protection standards for broad segments of consumer information. For example, the European Union has adopted mandatory standards for most types of consumer information through their 1996 Directive on Data Privacy (see [35] for a review of this Directive). In the United States, mandatory standards have been legislated more narrowly: for the use of credit reporting data, health information, some types of financial transactions, and marketing data from minors.

Seal-of-approval programs represent a third option, and serve as an interesting alternative to caveat emptor and mandatory standards regimes, particularly for Internet markets. Under this approach, a retailer can choose to join a seal-of-approval program administered by a seal-granting authority. Joining a seal-of-approval program gives the retailer the right to display a logo that certifies that the retailer will follow a set of information practices to protect consumer privacy. The seal-granting authority has the right to monitor the retailer's adherence to these standards. Therefore the seal-of-approval programs provide an industry self-regulatory approach to fostering trust between online retailers and consumers through a trusted third party. Examples include the programs offered by TRUSTe and the Better Business Bureau.

While the relative merits of each of these approaches have been debated in policy and regulatory circles, there has been little systematic research that aims to understand which regime is most effective in enhancing consumer trust and which regime optimizes social welfare. Answers to

¹ For example, the California Online Privacy Protection Act requires web sites that collect personal information on California residents to post a privacy policy on their sites and to comply with their policies [32]. In reality, most retailers post privacy policies. We thank an anonymous referee for this suggestion.

these questions have important policy implications, as some privacy advocates are currently 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 [34].

This research addresses the conflict between enhancing consumer trust in an online environment and promoting society's total welfare. We develop an analytical model with hidden information. Retailers signal their willingness to protect consumer privacy by conforming to certain privacy protection regimes. Because the accuracy of consumers' interpretation of signals varies across different regimes, consumers are exposed to different degrees of risks. This research studies how different regimes can enhance trust in an online environment, where consumer relationships display great social distance, consumers are lacking first-hand experience with retailers [9], and repeated encounters do not necessarily happen. This is one of the first studies to analyze the role of signaling as a mediator to enhancing trust in the context of online privacy protection.

Our model shows that by joining a seal-of-approval program retailers can send an unambiguous signal to consumers regarding privacy protection. Under caveat emptor, retailers can signal its intent to protect privacy by using trust-related arguments such as a privacy policy [22]. However the possibility of misinterpretation of signals due to the length and complexity of typical privacy policies can lead to a breakdown of consumer trust and jeopardize the effectiveness of this regime. Therefore, the accuracy of signals is critical to improving trust under the caveat emptor regime. Increasing the accuracy of signals increases consumers' trust toward privacy protection—leading to higher producer and consumer surplus.

We find that consumers' attitude towards privacy plays an important role in determining the type of privacy protection signal firms choose. Under the caveat emptor regime, when consumers have relatively low sensitivity towards their personal information, firms will only provide

“notice” in their privacy policies. In contrast, when consumers have relatively high sensitivity towards their personal information, firms will post the complete set of fair information practices in their privacy policies; In the intermediate case when consumers have moderate sensitivity towards their personal information, firms with low protection costs will post the complete set of fair information practices while firms with high costs will only post “notice.”

We also find that there is an intrinsic conflict between enhancing trust and achieving optimal social welfare. An *effective* privacy protection regime, such as mandatory standards, is not necessarily the most *efficient* regime in term of social welfare. Mandatory standards undoubtedly increase consumer trust. However, when consumers benefit from uniformly higher standards of privacy protection, they also pay higher prices. This can lead to a social welfare loss, which may outweigh the benefit of a higher standard of privacy protection.

The remainder of this paper proceeds as follows. In Section 2, we review the relevant academic literature. In Section 3, we introduce our model and obtain equilibrium solutions. In Section 4, we conclude with some broader implications of our work.

2. Literature Review

Previous research has shown that trust is essential in exchange relationships due to the embeddedness of actors in an ongoing system of social relations [14], [27] and the uncertainty existing between trading partners [24]. Trust is a crucial enabling factor in relations where uncertainty, interdependence, and fear of opportunism exist [1], [13]. On the Internet, consumers’ privacy concerns represent important sources of uncertainty, interdependence, and fear of opportunism between trading partners.

A variety of papers have analyzed the role of addressing consumer privacy concerns on trust building over the Internet. [4] identify a set of instruments for individuals' concerns about organizational privacy practices, including the collection of personal information, unauthorized secondary use of personal information, errors in personal information, and improper access to personal information — concepts central to fair information practice. [8] show that addressing privacy concerns by using fair information practice can facilitate trust building. [3] and [22] specifically discuss third party privacy seals and privacy statements as trust building tools. Joining seal-of-approval programs can serve as a type of institution-based trust in which consumers perceive that effective third-party institutional mechanisms are implemented to facilitate transaction [44], [31].

A growing body of literature discusses alternative privacy protection regimes (e.g. [18], [38]). Although the merits and limitations of each regime have been discussed extensively in the literature, there is no consensus regarding to which regime should be adopted. [9] study government regulation, industry self-regulation, and technological solutions under a justice framework and argue that the perception of justice shapes consumers' privacy concerns. [40] find that asymmetric information about whether websites will sell private information leads to a lemons market for privacy, and that government regulation and enforced laws are the only effective methods to make all companies respect consumer privacy. [19] demonstrate that the codification by E.U. law and the enforcement by the U.K. government does not improve the disclosure and practice of e-commerce privacy relative to markets in the United States. [15] argue that self-regulation mechanisms can reinforce the reputation of Web sites and provide useful information to consumers.

Although the literature on trust-building finds that privacy statements and privacy seals can enhance trust, there is less discussion of the different roles played by these two institutions. One contribution of this paper is in analyzing how and why these two institutions can influence consumer beliefs differently and achieve different levels of efficiency. While most of the literature focuses on the benefits to trust building, our study presents an integrated analysis of both the costs and benefits of privacy protection and highlights the resulting conflict between promoting trust and social welfare optimization when consumers have heterogeneous evaluations towards losing privacy. We consider both the *effectiveness* and the *efficiency* of different regimes of protecting online information privacy from not only the perspective of retailers, but also the perspective of consumers and society as a whole.

Our paper is also one of the first few studies to introduce the role of signaling in mediating consumer privacy concerns in online trust building. We explore the three most popular privacy protection regimes by considering the problems of hidden information [21], [28], [41]. Further our paper extends prior work by analyzing how the privacy regimes can be compared in terms of the quality of the privacy protection signals.

3. A Model of Privacy Protection and Trust

3.1. Firms and Consumers

Previous research has found that consumers may find it difficult to pre-contractually identify and select firms who have both the ability and willingness to protect consumer privacy (e.g. [32]). A failure to address this problem can lead to reduced consumer trust and a loss of social welfare. Our model studies how various regimes can help address this hidden information problem.

We consider a duopoly model with two competing firms, each of which sells one product to consumers. We use a setting that is similar to [17]’s model of horizontal product differentiation

and consumer taste differentiation. Each firm's product is located at one end of a straight line, and consumers are evenly distributed along this line according to their tastes. A consumer's utility for a product is assumed to be v minus a fit cost that is proportional to the distance between her taste and the product's location.² Fit cost is normalized to be 1 per unit of distance. Without loss of generality, each firm's marginal cost of production is assumed to be zero.

If a firm chooses not to protect consumer privacy, it incurs zero cost. If a firm protects privacy by following fair information practices, it incurs a positive cost. Offering notice, choice, access, security, and enforcement would require additional managerial and technology investment.³ The cost of protecting consumer privacy can vary significantly across firms, depending on the nature of the firm's product or service, the nature of the firm's consumers, and the amount of consumer information the firm can obtain through its transactions with consumers. Moreover, firms incur infrastructure costs associated with protecting privacy that vary across firms because different firms employ different sets of privacy protection infrastructure and personnel. In our model we assume one firm has a low marginal cost of protecting consumer privacy (c_L) and the other firm has a high marginal cost of protecting consumer privacy (c_H), with $0 < c_L < c_H < \infty$.⁴

As shown in a number of surveys (e.g. [16] and [25]), consumers are heterogeneous in how much they care about their privacy. To capture this heterogeneity, we assume that a proportion of θ ($0 \leq \theta \leq 1$) consumers (sensitive consumers) care about privacy and incur a utility loss of L ($0 \leq L \leq 1$) when their privacy is not protected. Likewise, a proportion of $1 - \theta$ consumers

² We assume v is large enough, i.e., $v \geq L + \frac{3}{2}$, so that the market is fully covered. This assumption of full market coverage allows us to study a setting with two competing firms. Otherwise, when the market is not fully covered, we have a case of two local monopolies. The equilibrium results in this case are discussed in Appendix 2.

³ We thank an anonymous referee for this suggestion of how to operationalize the cost of protecting privacy.

⁴ A symmetric case in which two firms have the same cost structure is analyzed in Appendix 2.

(insensitive consumers) do not care about privacy and do not incur a utility loss when their privacy is not protected. The values of θ and L vary depending on the different types of information (financial, medical, demographic, or transaction information) [7].

Without loss of generality, we assume the firm that is located at location 0 (Firm 0) has a low cost of protecting privacy and the firm that is located at location 1 (Firm 1) has a high cost of protecting privacy. Each consumer has unit demand for these two competing firms' products and will choose the product that gives her a higher utility. Let Firm 0's price be p_0 and Firm 1's price be p_1 . Thus, a sensitive consumer at location λ has a utility of $v - \lambda - p_0$ for Firm 0's product if her privacy is protected and a utility of $v - L - \lambda - p_0$ for Firm 0's product if her privacy is not protected. Similarly, a sensitive consumer has a utility of $v - (1 - \lambda) - p_1$ for Firm 1's product if her privacy is protected and a utility of $v - L - (1 - \lambda) - p_1$ for Firm 1's product if her privacy is not protected. An insensitive consumer at location λ has a utility of $v - \lambda - p_0$ for Firm 0's product and a utility of $v - (1 - \lambda) - p_1$ for Firm 1's product, regardless of whether her privacy is protected or not. All consumers obtain a utility of zero if no purchase is made.

3.2. *Caveat Emptor*

Under the caveat emptor regime, each firm can use its privacy policy to signal to consumers its willingness to protect consumer privacy. However, because of the length and complexity of a typical privacy policy posted by a firm, it is hard for consumers to perfectly interpret a firm's signal through its privacy policy. To illustrate this, we collected the privacy policies from the top 20 most popular shopping sites listed at Alexa.com. These privacy policies averaged 2,074 words in length, not including links to other supporting pages. This corresponds roughly to four pages of single spaced typewritten text. We did a readability analysis and found the average Flesch-

Kincaid ([11],[23]) grade level score for these privacy policies was 12.8, well above the recommended complexity for most standard documents. This finding is consistent with [29] and [30]’s findings that consumers found that privacy policies were often too long and confusing.

We use a garbled signal to model how consumers interpret each firm’s privacy policy and obtain their interpretation of whether each firm has signaled post-contractual privacy protection or no protection. We assume a firm can send two types of signals to consumers: a policy (s_L) that signals the firm will follow fair information practices and protect their privacy post-contractually, or a policy (s_H) that signals the firm will not protect privacy post-contractually.⁵

In addition, we assume consumers’ interpretation of a firm’s signal is not 100% accurate. This is consistent with extant empirical findings (e.g. [30]). If a firm sends a low-type signal (s_L), then with probability $\alpha(0 \leq \alpha < 1)$, consumers will obtain an interpretation (r_L) that the firm will protect privacy post-contractually; and with probability $1 - \alpha$, consumers will be confused and will not be able to obtain any interpretation regarding post-contractual privacy protection (r_N). If the firm sends a high-type signal (s_H), with probability $\alpha(0 \leq \alpha < 1)$, consumers will obtain an interpretation (r_H) that the firm will not protect privacy post-contractually; and with probability $1 - \alpha$, consumers will be confused and will not be able to obtain any interpretation regarding post-contractual privacy protection (r_N). Note that when $\alpha = 0$, a firm’s signal becomes completely garbled; and in this special case, privacy policies cannot be used as trust indices to enhance trust [3]. When $\alpha = 1$, a firm’s signal is not garbled at all.

⁵ s_H can be either no privacy policy, or a policy signaling the firm will not follow fair information practices. If providing “notice” does not lead to significant costs for the firm, then all rational firms will provide “notice” and a policy signaling the firm will only provide “notice” will also be s_H . We thank an anonymous referee for suggesting this.

The Post-contractual Holdup Problem

Firms may be unwilling to protect consumer privacy post-contractually, even when they have signaled post-contractual privacy protection. In order for a firm's privacy policy to be a meaningful signal of post-contractual privacy protection, this post-contractual holdup problem that has been identified by previous research (e.g., [32]) must be solved. We argue that effective post-contractual monitoring by the government can help solve this holdup problem under the caveat emptor regime. We assume the government detects deceptive claims with probability μ , and penalizes a firm by F if it has made deceptive claims. If the government sets F sufficiently high, i.e., $\mu F \geq c_H > c_L$, then a firm which has signaled post-contractual privacy protection will indeed protect privacy post-contractually. This is because the firm's post-contractual net profit if the firm protects privacy, which is $(p - c_i)D$, is higher than the firm's post-contractual net profit if the firm does not protect privacy, which is $(p - \mu F)D$, where D is the firm's demand. We assume $\mu F \geq c_H > c_L$ in our model of the caveat emptor regime.

Timing of the Caveat Emptor Game

We use a two-stage game to capture the behavior of firms and consumers under the caveat emptor regime. In Stage 1, both firms post their selling prices (p_0, p_1) and signal post-contractual privacy protection or no protection through their privacy policies (s_0, s_1) . In Stage 2, consumers observe posted prices. Consumers interpret privacy policies and obtain their interpretation of whether each firm has signaled post-contractual privacy protection or no protection (r_0, r_1) .⁶

Consumers form belief $\beta(r_0), \beta(r_1)$ regarding the probability each firm will protect privacy post-

⁶ We assume consumers do not use price as a signal of post-contractual privacy protection. This assumption makes sense because the signals sent by the firms (privacy policies) are backed up by post-contractual monitoring. This monitoring ensures a firm who has signaled protection will protect privacy post-contractually. In our model, the signal sent by the firms on post-contractual privacy protection is a more dominant signal than price is.

contractually. Consumers then decide whether to make a purchase and which firm to purchase from, based on consumers' willingness-to-pay, posted prices, and their belief regarding the probability each firm will protect privacy post-contractually. After this stage, each firm's profit and consumers' utilities are realized and the game ends.

Equilibrium in the Caveat Emptor Game

We solve for the Perfect Bayesian Nash Equilibrium in this game by considering four possible types of equilibrium: a) a pooling equilibrium in which both firms signal no protection (s_H); b) a separating equilibrium in which low-cost Firm 0 signals privacy protection (s_L) and high-cost Firm 1 signals no protection (s_H); c) a pooling equilibrium in which both firms signal privacy protection (s_L); and d) a separating equilibrium in which Firm 0 signals no protection (s_H) and Firm 1 signals privacy protection (s_L). Figure 1 shows the information structure of these four possible types of equilibrium. Proposition 1 summarizes our analyses of the caveat emptor game.

[INSERT FIGURE 1 HERE]

Proposition 1 (Perfect Bayesian Nash Equilibrium in CE game): 1) If $\alpha\theta L \leq c_L < c_H$, there exists a pooling Perfect Bayesian Nash Equilibrium in which both Firm 0 and Firm 1 signal no protection (s_H) and neither protects privacy post-contractually.

2) If $c_L < \alpha\theta L < c_H$, there exists a separating Perfect Bayesian Nash Equilibrium in which low-cost Firm 0 signals privacy protection (s_L) and protects privacy post-contractually, and high-cost Firm 1 signals no protection (s_H) and does not protect privacy post-contractually.

3) If $c_L < c_H \leq \alpha\theta L$, there exists a pooling Perfect Bayesian Nash Equilibrium in which both Firm 0 and Firm 1 signal privacy protection (s_L) and both protect privacy post-contractually.

The proof of this and other propositions can be found in Appendix 1.

Consumer Trust, Demand and Firm Profit in Each Equilibrium

When $\alpha\theta L \leq c_L < c_H$, both Firm 0 and Firm 1 signal no protection and neither protects privacy post-contractually. A consumer's belief regarding the probability of privacy protection, is

$\beta(r_H) = 0, \beta(r_N) = 0, \beta(r_L) = 1$. Firm 0 and Firm 1 will split consumer demand evenly and obtain

the same level of profit. That is, $D_0^* = \frac{1}{2}, D_1^* = \frac{1}{2}, \pi_0^* = \frac{1}{2}, \pi_1^* = \frac{1}{2}$. Figure 2 shows how Firm 0

and Firm 1 divide the market in this and the other two equilibria.

[INSERT FIGURE 2 HERE]

When $c_L < \alpha\theta L < c_H$, the low-cost Firm 0 signals protection and protects privacy post-contractually, while the high-cost Firm 1 signals no protection and does not protect privacy post-contractually. A consumer's belief regarding the probability of privacy protection is

$\beta(r_H) = 0, \beta(r_N) = 0.5, \beta(r_L) = 1$. In this case, more sensitive consumers purchase from Firm 0 than

from Firm 1, while more insensitive consumers purchase from Firm 1 than from Firm 0. We

have $D_0^* = \frac{1}{2}(1 + \frac{\alpha\theta L}{3} - \frac{c_L}{3}), D_1^* = \frac{1}{2}(1 - \frac{\alpha\theta L}{3} + \frac{c_L}{3}), \pi_0^* = \frac{1}{2}(1 + \frac{\alpha\theta L}{3} - \frac{c_L}{3})^2, \pi_1^* = \frac{1}{2}(1 - \frac{\alpha\theta L}{3} + \frac{c_L}{3})^2$.

When $c_L < c_H \leq \alpha\theta L$, both Firm 0 and Firm 1 signal privacy protection and protect privacy post-contractually. A consumer's belief regarding the probability of privacy protection is

$\beta(r_H) = 0, \beta(r_N) = 1, \beta(r_L) = 1$. More consumers purchase from Firm 0 than from Firm 1, because

Firm 0 charges a lower price. We have

$D_0^* = \frac{1}{2}(1 + \frac{c_H}{3} - \frac{c_L}{3}), D_1^* = \frac{1}{2}(1 - \frac{c_H}{3} + \frac{c_L}{3}), \pi_0^* = \frac{1}{2}(1 + \frac{c_H}{3} - \frac{c_L}{3})^2, \pi_1^* = \frac{1}{2}(1 - \frac{c_H}{3} + \frac{c_L}{3})^2$.

When Does The Caveat Emptor Regime Work and When Does It Fail?

All else equal, signaling privacy protection rather than no protection helps a firm obtain a higher demand, while it also increases the firm's marginal cost. A firm would consider this tradeoff when it decides whether to signal privacy protection. Proposition 1 shows that, when $c_L < c_H \leq \alpha\theta L$ and $c_L < \alpha\theta L < c_H$, the caveat emptor regime successfully leads to post-contractual privacy protection for at least some consumers. However, when $\alpha\theta L \leq c_L < c_H$, both firms find the potential demand gain from protecting privacy is outweighed by the increase in marginal cost, and neither signals privacy protection. This case can happen even though θ and L are both large, i.e., a large percentage of consumers care about privacy protection and their utility losses are large when their privacy is not protected. As long as the probability that consumers get confused by the privacy policies is large enough, i.e., α is small enough ($\alpha \leq c_L$), the case of $\alpha\theta L \leq c_L < c_H$ is the only case possible. Figure 3 illustrates when α is large, the caveat emptor regime could lead to post-contractual privacy protection for at least some θ and L . But as α declines, it could fail to lead to post-contractual privacy protection for any θ and L .

[INSERT FIGURE 3 HERE]

Improving Interpretation Accuracy

There are many ways firms can improve the accuracy of consumers' interpretation of the firm's signal through its privacy policy. Currently privacy policies posted by firms have different formats and frequently are long and use confusing legal jargon. If the government or an industry consortium can draft and enforce an industry-wide standard format and template for privacy policies, this could allow consumers to read only key sections of privacy policies and compare different privacy policies side by side. The successful implementation of such a standard could improve the accuracy of consumers' interpretation process.

Another approach to improving the accuracy of consumers' interpretation process is to provide consumers with tools that automatically process the large amount of information in privacy policies, such as the Platform for Privacy Protection (P3P). P3P allows firms to make privacy policies conform to the XML-based P3P standard, and provides consumers with rule-based tools and XML parsers so that they can easily interpret privacy policies [5].

However, although promising, there are significant barriers to the widespread use of P3P protocols [9]. First, in order to implement P3P, both vendor websites and browsers must support P3P. Currently not all websites support P3P and the commercial implementation of P3P on browsers is limited. For example, Microsoft Internet Explorer 6 only implements cookie filtering based on P3P.⁷ Second, the accuracy with which websites' policies are represented by P3P user agents varies, increasing ambiguity and uncertainty for users to understand privacy policies [6].

We conclude that, although both approaches improve the accuracy of interpretation process, they are likely to result in an imperfect interpretation ($\alpha < 1$ in our model). While privacy policies help mitigate the perceived risk of a site and facilitate trust building, uncertainty still remains.

The Effect of Interpretation Accuracy on Consumer Trust, Demand and Firm Profit

Next we will study how the accuracy of consumers' interpretation of firms' signals affects consumer trust, demand and firm profit.

Proposition 2: Consumers' belief regarding the probability each firm will protect privacy post-contractually, given their interpretation of each firm's signal, is a non-decreasing function of the accuracy of the interpretation process (α).

⁷ We thank an anonymous referee for this point.

Proposition 2 shows that consumer trust regarding privacy protection can be enhanced if we can increase the accuracy of consumers' interpretation process (α).

Proposition 3: Firm 0's demand and profit are both non-decreasing functions of the accuracy of the interpretation processes (α). But Firm 1's demand and profit are both non-increasing functions of the accuracy of the interpretation processes (α).

Figure 4 illustrates how demand and firm profit change as the accuracy of the interpretation process changes. When α is very low, neither firm protects consumer privacy, and Firm 0 and Firm 1 compete on prices only. Thus, Firm 0 does not enjoy any competitive advantage when competing with Firm 1.

[INSERT FIGURE 4 HERE]

But as the accuracy of interpretation process increases, Firm 0 and Firm 1 compete on price and privacy protection. Because Firm 0 has a lower cost of protecting consumer privacy, it enjoys a cost advantage when competing with Firm 1. In addition, Firm 0's competitive advantage becomes larger as the accuracy of interpretation process increases. Thus, high-cost Firm 1 *loses* demand and profit and low-cost Firm 0 *gains* demand and profit, as the accuracy of the interpretation process increases. Proposition 3 illustrates that the incentive to improve the accuracy of interpretation process may differ for low-cost and high-cost firms.

The Effect of Interpretation Accuracy on Social Welfare in Caveat Emptor Game:

Next we will study how the accuracy of consumers' interpretation of firms' signals affects social welfare in the caveat emptor game.

Proposition 4: Social welfare in the caveat emptor game is a non-decreasing function of the accuracy of the interpretation process (α) when $\alpha\theta L < c_H$.

The effect of interpretation accuracy on social welfare is less straightforward. We first study the effect of interpretation accuracy on social welfare inside each equilibrium: Under the pooling equilibrium where neither firm protects privacy or the pooling equilibrium where both firms protect privacy, social welfare does not change as α increases. Under the separating equilibrium where only Firm 0 protects privacy, social welfare is increasing as α increases.

However, as Figure 3 illustrates, changes in α will move the boundaries that separate the three types of equilibria in the caveat emptor game. So we also need to compare social welfare across different equilibria at the boundaries. Social welfare is higher under the separating equilibrium than under the pooling equilibrium where neither firm protects privacy at their boundary $\alpha\theta L = c_L$. But comparing social welfare under the separating equilibrium with social welfare under the pooling equilibrium where both firms protect privacy yields an ambiguous answer at their boundary $\alpha\theta L = c_H$. Therefore, we only conclude that social welfare in the caveat emptor game is a non-decreasing function of the accuracy of interpretation process (α) when $\alpha\theta L < c_H$.

3.3. Seal-of-approval Programs

Under the seal-of-approval regime, a firm can send an unambiguous signal to consumers regarding whether it will protect consumer privacy post-contractually by displaying a “seal-of-approval” logo. In order to display the logo, firms need to meet the requirements set by seal-granting authorities. The “seal-of-approval” logo certifies that the firm which displays it will follow a certain set of standards to protect privacy, and that the seal-granting authority will have the right to monitor the firm’s adherence to these standards. The “seal-of-approval” logo serves as a signal that can be easily interpreted by consumers, while the penalty imposed by the seal-

granting authority on firms that display a “seal-of-approval” logo but do not protect privacy, guarantees the credibility of this signal.

The key difference between the seal-of-approval game and the caveat emptor game is that signals are not garbled in the seal-of-approval game while they are garbled in the caveat emptor game. More specifically, we assume the firm can send two types of signals to consumers: displaying a logo (s_L), which signals post-contractual privacy protection, and not displaying a logo (s_H), which signals no protection. In addition, we assume consumers’ interpretation of the firm’s signal is 100% accurate. If the firm signals privacy protection (s_L), with probability 1 consumers would interpret that the firm has signaled privacy protection (r_L). If the firm signals no protection (s_H), with probability 1 consumers would interpret that the firm has signaled no protection (r_H).

We also assume that the seal-granting authority monitors firms’ post-contractual behavior. Deceptive claims are detected with probability μ ($0 \leq \mu \leq 1$) by the seal-granting authority, and penalized by a fine of F . We assume $\mu F \geq c_H > c_L$ in our model of the seal-of-approval regime because setting a sufficiently high penalty is in the best interest of the seal-granting authority. When this assumption does not hold, displaying the seal-of-approval logo becomes a meaningless signal and the market succumbs to the post-contractual holdup problem that is discussed in the caveat emptor game.

Under these assumptions, we use a two-stage game, similar to the caveat emptor game, to model the seal-of-approval program regime. In Stage 1, both firms post selling prices (p_0, p_1) and signal post-contractual privacy protection or no protection through “seal-of-approval” logos (s_0, s_1). In Stage 2, consumers observe posted prices and obtain an interpretation of whether each firm has

signaled protection or no protection (r_0, r_1) . Consumers form belief $\beta(r_0), \beta(r_1)$ regarding the probability of privacy protection given their interpretation of each firm's signal. Consumers then decide whether to make a purchase and which firm to purchase from, based on consumers' willingness-to-pay, posted prices, and their belief regarding the probability of privacy protection. After this stage, each firm's profit and consumers' utilities are realized and the game ends. We note that the seal-of-approval program regime is in fact a special case of the caveat emptor regime in which consumers' interpretation of a firm's signal is 100% accurate, i.e., $\alpha = 1$.

Seal-of-approval Regime Can Enhance Consumer Trust

Under the seal-of-approval regime, the firm can send a signal of whether it will protect privacy by displaying or not displaying a "seal-of-approval" logo (s_L, s_H) . Consumers interpret this signal unambiguously. Consumer trust is higher in the seal-of-approval game ($\alpha = 1$) than in the caveat emptor game ($0 \leq \alpha < 1$). This directly follows from Proposition 2.

Seal-of-approval Regime versus Caveat Emptor Regime

Next we will compare the social welfare under the seal-of-approval regime with the social welfare under the caveat emptor regime.

Proposition 5 Social welfare is no lower in the seal-of-approval game ($\alpha = 1$) than in the caveat emptor game ($0 \leq \alpha < 1$), as long as the θL is low enough, i.e., $\theta L < c_H$.

Proposition 5 shows that the seal-of-approval regime leads to a higher social welfare than the caveat emptor regime, when the percentage of consumers who care about privacy is low and the loss these consumers suffer when their privacy is not protected is low. Figure 5 further explains this. However, when $\theta L \geq c_H$, the caveat emptor regime leads to an outcome of only Firm 0 protecting privacy and the seal-of-approval regime leads to an outcome of both firms protecting

privacy. In this case, the comparison of social welfare under the caveat emptor and seal-of-approval regimes is ambiguous. In all other cases, the seal-of-approval regime leads to the same or higher social welfare than the caveat emptor regime does.

[INSERT FIGURE 5 HERE]

3.4. Mandatory Standards

Under the mandatory standards regime, the government sets minimum standards for protecting consumer privacy and requires all firms to follow these standards. For example, in the United States, 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 [15]. This regime is similar to a conventional regulatory approach such as standard setting [26] and command-and-control regulation [37].

To model this regime, we assume that if a firm does not protect consumer privacy, this violation is detected with probability μ ($0 < \mu \leq 1$) by the government and penalized by F . This parallels enforcement in practice where governments can penalize firms that violate the government's mandatory standards either through litigation, or associated penalties, or legal expenses. It is straightforward to show that the government can set F sufficiently high (i.e. $\mu F \geq c_H > c_L$) such that both low- and high-cost firms find it optimal to protect consumer privacy post-contractually.

Seal-of-Approval Regime versus Mandatory Standards Regime

Under the mandatory standards regime, all firms are forced to follow the standards set by the government and protect consumer privacy post-contractually. Assuming the government can effectively enforce the mandatory standards, this regime leads to the highest level of consumer trust possible—consumers believe that firms will protect privacy post-contractually with

probability 1. In terms of enhancing consumer trust, the mandatory standards regime is obviously *more effective* than the seal-of-approval regime. However, the mandatory standards regime may be *less efficient* than the seal-of-approval regime, in terms of enhancing social welfare.

The key difference between the mandatory standards and the seal-of-approval regimes is that the former forces firms to protect privacy while the latter allows the firm to freely choose. As illustrated by Figure 6, when $c_L < c_H \leq \theta L$, both regimes lead to the same outcome. However, when $c_L < \theta L < c_H$, the seal-of-approval regime leads to an outcome in which only the low-cost Firm 0 protects privacy; it leads to an outcome in which neither firm protects privacy, when $\theta L \leq c_L < c_H$. Under the mandatory standards regime, both firms always protect privacy. Next we will compare the social welfare under these two regimes.

[INSERT FIGURE 6 HERE]

Proposition 6: Social welfare in the seal-of-approval game is no lower than social welfare in the mandatory standards game.

Proposition 6 shows that social welfare is higher under seal-of-approval regime than under the mandatory standards regime, in both of these cases: when $c_L < \theta L < c_H$ and when $\theta L \leq c_L < c_H$.

This proposition shows that forcing firms to protect privacy, as in mandatory standards regime, could lead to an overall decrease in social welfare. This is particularly true when few consumers care about privacy protection and when the potential utility loss from privacy not being protected is small. In these cases, the gain in social welfare from protecting consumer privacy is outweighed by the cost firms incur in protecting consumer privacy.

This result suggests that adopting the mandatory standards regime for nearly all types of consumer information may not always be optimal from a social welfare perspective. For certain

types of purchase and demographic information where few consumers have privacy concerns or where consumers are less sensitive to privacy violations, a mandatory standards regime could lead to a loss of social welfare, because in these cases both consumers who are sensitive to privacy protection and those who are not have to pay for protection through higher prices.

However, we caution that there may be other reasons favoring a mandatory standards regime.

For example, our model assumes that seal-granting authorities can effectively detect and penalize deceptive claims, and hence, the post-contractual hold-up problem is not the focus of this paper.

But if the government can much more effectively and efficiently enforce post-contractual compliance than seal-granting authorities can, especially for certain types of consumer information (such as health and medical information, credit and some types of financial information, and marketing information obtained from minors), then a mandatory standards regime makes more sense than a seal-of-approval regime.

4. Discussion and Concluding Remarks

In this paper, we develop a model of hidden information in which online retailers can signal to consumers regarding their willingness to protect consumer privacy. We analyze retailers' strategies under three privacy protection regimes commonly chosen by market designers or government regulators: *caveat emptor*, seal-of-approval programs, and mandatory standards.

Under a *caveat emptor* regime, retailers can post a privacy policy that imperfectly signals privacy protection or no protection. Under seal-of-approval programs, retailers can send an unambiguous signal by joining seal-of-approval programs. Under mandatory standards, there is no need to send signals because of the high level of government intervention.

This research differs from much of the previous research on trust in that we focus on the role of signaling in enhancing trust in the context of online privacy protection. We find that the extent to

which retailers influence consumer trust depends crucially on the clarity and credibility of the signal retailers send. Seal-of-approval programs increase the credibility of the signal regarding privacy protection, leading to a higher level of consumer trust than the *caveat emptor* regime.

The mandatory standards regime is the most *effective* way of enhancing consumer trust. But we find that it can be less *efficient* than the seal-of-approval programs regime in terms of social welfare, in particular, for cases in which few consumers are sensitive to privacy and when their potential loss is small. This is because mandatory standards regimes lead to higher retailer costs and, as a result, higher prices. This, in turn, leads to a social welfare loss, which may outweigh any benefits from better privacy protection. Effectively, seal-of-approval programs allow customers to self-select whether to deal with a firm that protects privacy or a firm that does not protect privacy (and correspondingly having a lower price). Thus, in general, adopting a mandatory standards regime for nearly all types of consumer information is not a socially optimal approach to protecting consumer privacy.

Despite federal and state governments' efforts to protect privacy, surveys and opinion polls continue to show that some companies do not post privacy policies and are not subject to enforcement. According to the analysis of this paper, this situation falls into the *caveat emptor* regime. Because no signaling mechanism is in place to enhance consumer trust, consumers bear high privacy-related risks when conducting transactions with those companies.

As suggested by our model the existence of consumer heterogeneity also influences the coexistence of different privacy protection regimes. In the U.S. privacy is taken as negotiable historically, while in E.U. it is taken as a human right [35]. Thus, the U.S. currently adopts a sectoral approach to addressing privacy concerns while European countries adopt an omnibus approach to protecting privacy.

However, a limitation of our model is that we do not consider dynamic settings where firms' decisions in one period would influence outcomes in future periods. We do this for the sake of parsimony and clarity of exposition. In a dynamic setting firms would have additional strategies for communicating trust to their customers, including using the value of a brand name as a bond that would be forfeited if trust is violated [41], or signaling trust through potential lost sales in a repeated game setting [10], [19]. But such strategies will not be available to many retailers, particularly in settings with infrequent interaction or short/non-existent purchase histories.

Building on this point, it is important to note that none of the approaches above is a necessary or sufficient condition to build trust, and it would be interesting for future work to discuss the interactions between various privacy trust-building strategies. For example, some of these strategies may work well in combination, such as signaling trust through both repeated interactions and through seal-of-approval programs. Similarly, some strategies may conflict. For example, some established retailers may be reluctant to join any seal-of-approval programs for fear of conflicting with their brand names. In this case, both the caveat emptor regime and the branding effect jointly safeguard consumer privacy.

We also note that alternative approaches for trust-building are especially important when there are no privacy policies to signal protection. Reputation, branding and consumer experience may be used as indices for trust-building and to guide transactions. Moreover, companies can provide monetary payoffs in the form of discounts or better services or a wide product selection to compensate the potential loss of consumer privacy. Still privacy laws that require at least notice or awareness of the fair information practices are needed, so that companies can communicate with consumers about their privacy protection practices meaningfully and effectively.

Companies can potentially offer one price that is associated with privacy protection and another

price that is associated with no protection. This menu of prices allows consumers to self-select whichever option gives them the highest utility. While such a strategy is not at all widespread today, it would be interesting for future research to analyze such versioning strategies.

Our analysis offers strategic insights for a variety of audiences including third-party organizations overseeing the industry's practice of protecting privacy; market designers; 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 retailers' violations of privacy can be caught and penalized effectively. For market designers, our results also suggest that automated solutions to communicating retailer trust, such as the P3P standard, can help improve the efficiency of markets in the presence of caveat emptor regimes. For businesses, our results suggest that the incentive to improve signal accuracy will differ for low-cost firms and high-cost firms. Finally for regulatory agencies, our results suggest that overarching approaches to privacy protection where all types of data are covered by the same set of mandatory standards are not necessarily the socially optimal solution. Instead, government agencies should educate consumers to foster their ability to understand firms' privacy practices, so that consumers and firms can transact more effectively and efficiently.

References

1. Ba, S. and Pavlou. P. A. Evidence of the effect of trust building technology in electronic markets: price premium and buyer behavior. *MIS Quarterly*, 26, 3 (2002), 243-268.
2. Baron, D. P. DoubleClick and internet privacy. Stanford University Case Number P-32, (August, 2000).
3. Belanger F, Hiller J. S., and Smith W. J. Trustworthiness in electronic commerce: the role of privacy, security, and site attributes. *Journal of Strategic Information Systems*, 11, 3-4 (2002), 245-270.

4. California online privacy protection act of 2003. California Business and Professions Code sections 22575-22579, 2003 (available at <http://www.leginfo.ca.gov/cgi-bin/displaycode?section=bpc&group=22001-23000&file=22575-22579>).
5. Cranor, L. F. *Web Privacy with P3P*. Sebastopol, CA: O'Reilly & Associates, 2002.
6. Cranor, L. F., and Reidenberg, J. R. Can user agents accurately represent privacy notices? Discussion paper, Telecommunications Policy Research Conference (2002), (available at <http://web.si.umich.edu/tprc/papers/2002/65/tprc2002-useragents.PDF>)
7. Culnan, M.J. How did they get my name? An exploratory investigation of consumer attitudes toward secondary information use. *MIS Quarterly*, 17, 3 (1993), 341-363.
8. Culnan, M.J. and Armstrong, P. K. Information privacy concerns, procedural fairness, and impersonal trust: an empirical investigation. *Organization Science*, 10, 1 (1999), 104-115.
9. Culnan, M. J. and Bies, R. J. Consumer privacy: balancing economic & justice considerations. *Journal of Social Issues*, 59, 2 (2003), 323-342.
10. Dasgupta, P. Trust as a commodity. In D.G. Gambetta (Ed.), *Trust*. New York: Basil Blackwell, 1988, pp. 49-72.
11. Flesch, R. A new readability yardstick. *Journal of Applied Psychology*, 32, 3 (1948), 221-233.
12. Gambetta, D.G. Can we trust trust? In D.G. Gambetta (Ed.), *Trust*. New York: Basil Blackwell, 1988, pp. 213-237.
13. Gefen, D., Karahanna, E., and Straub, D. Trust and TAM in online shopping: an integrated model. *MIS Quarterly*, 27, 1 (2003), 51-90.
14. Granovetter, M. Economic action and social structure: the problem of embeddedness. *American Journal of Sociology*, 91, 1 (1985), 481-510.
15. Hahn, R. W., Layne-Farrar, A. The benefits and costs of online privacy legislation. (2001). Working Paper 01-14, AEI-Brooking Joint Center for Regulatory Studies.
16. Hoffman, D. L., Novak, T. P., and Peralta, M. Building consumer trust online. *Communications of the ACM*, 42, 4 (1999), 80-85.
17. Hotelling, H. Stability in competition. *The Economic Journal*, 39, 153 (March, 1929), 41-57.
18. Jamal, K., Maier, M., and Sunder, S. Enforced standards versus evolution by general acceptance: a comparative study of e-commerce privacy disclosure and practice in the United States and the United Kingdom. *Journal of Accounting Research*, 43, 1 (2004), 73-96.
19. James, H.S. The trust paradox: a survey of economic inquiries into the nature of trust and trustworthiness. *Journal of Economic Behavior and Organization*, 47, 3 (2002), 291-307.

20. Johnson-George, C., and Swap, W. Measurement of specific interpersonal trust: construction and validation of a scale to assess trust in a specific other. *Journal of Personality and Social Psychology*, 43, 6 (1982), 1306-1317.
21. Kihlstrom, R. E. and Riordan, M. H. Advertising as a signal. *Journal of Political Economy*, 92, 3 (1984), 427-450.
22. Kim, D and Benbasat, I. Trust-related arguments in internet stores: a framework for evaluation. *Journal of Electronic Commerce Research*, 4, 2 (2003), 49-64.
23. Kincaid, J. P., Fishburne, R. P., Rogers, R. L., and Chissom, B. S. Derivation of new readability formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy enlisted personnel. Research Branch Report 8-75, Naval Technical Training, U. S. Naval Air Station, Memphis, TN, 1975.
24. Kollock, P. The emergence of exchange structures: an experimental study of uncertainty, commitment, and trust. *American Journal of Sociology*, 100, 2 (1994), 313-345.
25. Leap of faith: Using the Internet despite the dangers. Research Report, Consumer Reports WebWatch, Yonkers, NY, October 26, 2005 (available at www.consumerwebwatch.org/dynamic/web-credibility-reports-princeton.cfm)
26. Magat, W. A. and Viscusi, W. K. *Informational Approaches to Regulation*. Cambridge: MIT Press, 1992.
27. Mayer, R.C., Davis, J.H., and Schoorman, F.D. An integrative model of organizational trust. *Academy of Management Review*, 20, 3 (1995), 709-734.
28. Milgrom, P. and Roberts, J. Price and advertising signals of product quality. *Journal of Political Economy*, 94, 4 (1986), 796-821.
29. Milne, G.R. and Culnan, M.J. Strategies for reducing online privacy risks: why consumers read [or don't read] online privacy notices. *Journal of Interactive Marketing*, 18, 3 (2004), 15-29.
30. Milne, G.R., Culnan, M.J., and Greene, H. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy and Marketing*, 25, 2 (2006), 238-249.
31. Pavlou, P. A. and Gefen, D. Building effective online marketplaces with institution-based trust. *Information Systems Research*, 15, 1 (2004), 37-59.
32. Pavlou, P.A., Liang, H., and Xue, Y. Understanding and mitigating uncertainty in online environments: an agency theory perspective. *MIS Quarterly*, 31, 1 (2007), 105-136.
33. Pitofsky, R. Privacy online: Fair information practices in the electronic marketplace. Federal Trade Commission Statement before the Committee on Commerce, Science, and Transportation, Washington, DC, May 25, 2000 (available at <http://www.ftc.gov/os/2000/05/testimonyprivacy.htm>).

34. Ryan, J.P. Privacy, the common good, and individual liberties in the 21st century: A dialogue on policy, law, and values. *American Bar Association Focus*, 15, 2 (2000).
35. Smith, H. J. Information privacy and marketing: what the U.S. should (and shouldn't) learn from Europe. *California Management Review*, 41, 2 (2001), 8-33.
36. Smith, H.J., Milberg, S.J., and Burke, S.J. Information privacy: measuring individuals' concerns about organizational practices. *MIS Quarterly*, 20, 2 (1996), 167-196.
37. Sunstein, C.R. Informational regulation and informational standing: *Akins* and beyond. *University of Pennsylvania Law Review*, 147, 3 (1999), 613-675.
38. Swire, P.P. Markets, self-regulation, and government enforcement in the protection of personal information. In *Privacy and Self-Regulation in the Information Age* (pp. 3-19). Washington, DC: U.S. Department of Commerce, 1997.
39. U.S. Public Interest Research Group. Public comment on barriers to electronic commerce. Response to call by U.S. Department of Commerce, 65 Federal Register 15898, April 25, 2000.
40. Vila, T., Greenstadt, R., and Molnar, D. 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, PA, 2003, pp. 403-407.
41. Wernerfelt, B. Umbrella branding as a signal of new product quality: an example of signaling by posting a bond. *RAND Journal of Economics*, 29, 3 (1988), 458-466.
42. Westin, A. F. *Privacy and Freedom*. New York: Atheneum, 1967.
43. Westin, A.F. How to craft effective online privacy policies. *Privacy and American Business*, 11, 6, (2004), 1-2.
44. Zucker, L. Production of trust: Institutional sources of economic structure. *Research in Organizational Behavior*, 8, 1 (1986), 53-111.

Figures




Figure 1: Information structure of possible equilibria in the caveat emptor game




Figure 2: Firm 0 and Firm 1's demand from sensitive consumers and insensitive consumers




Figure 3: When does the caveat emptor regime work and when does it fail?




Figure 4: The effect of interpretation accuracy on demand and firm profit ($\alpha_2 > \alpha_1$)

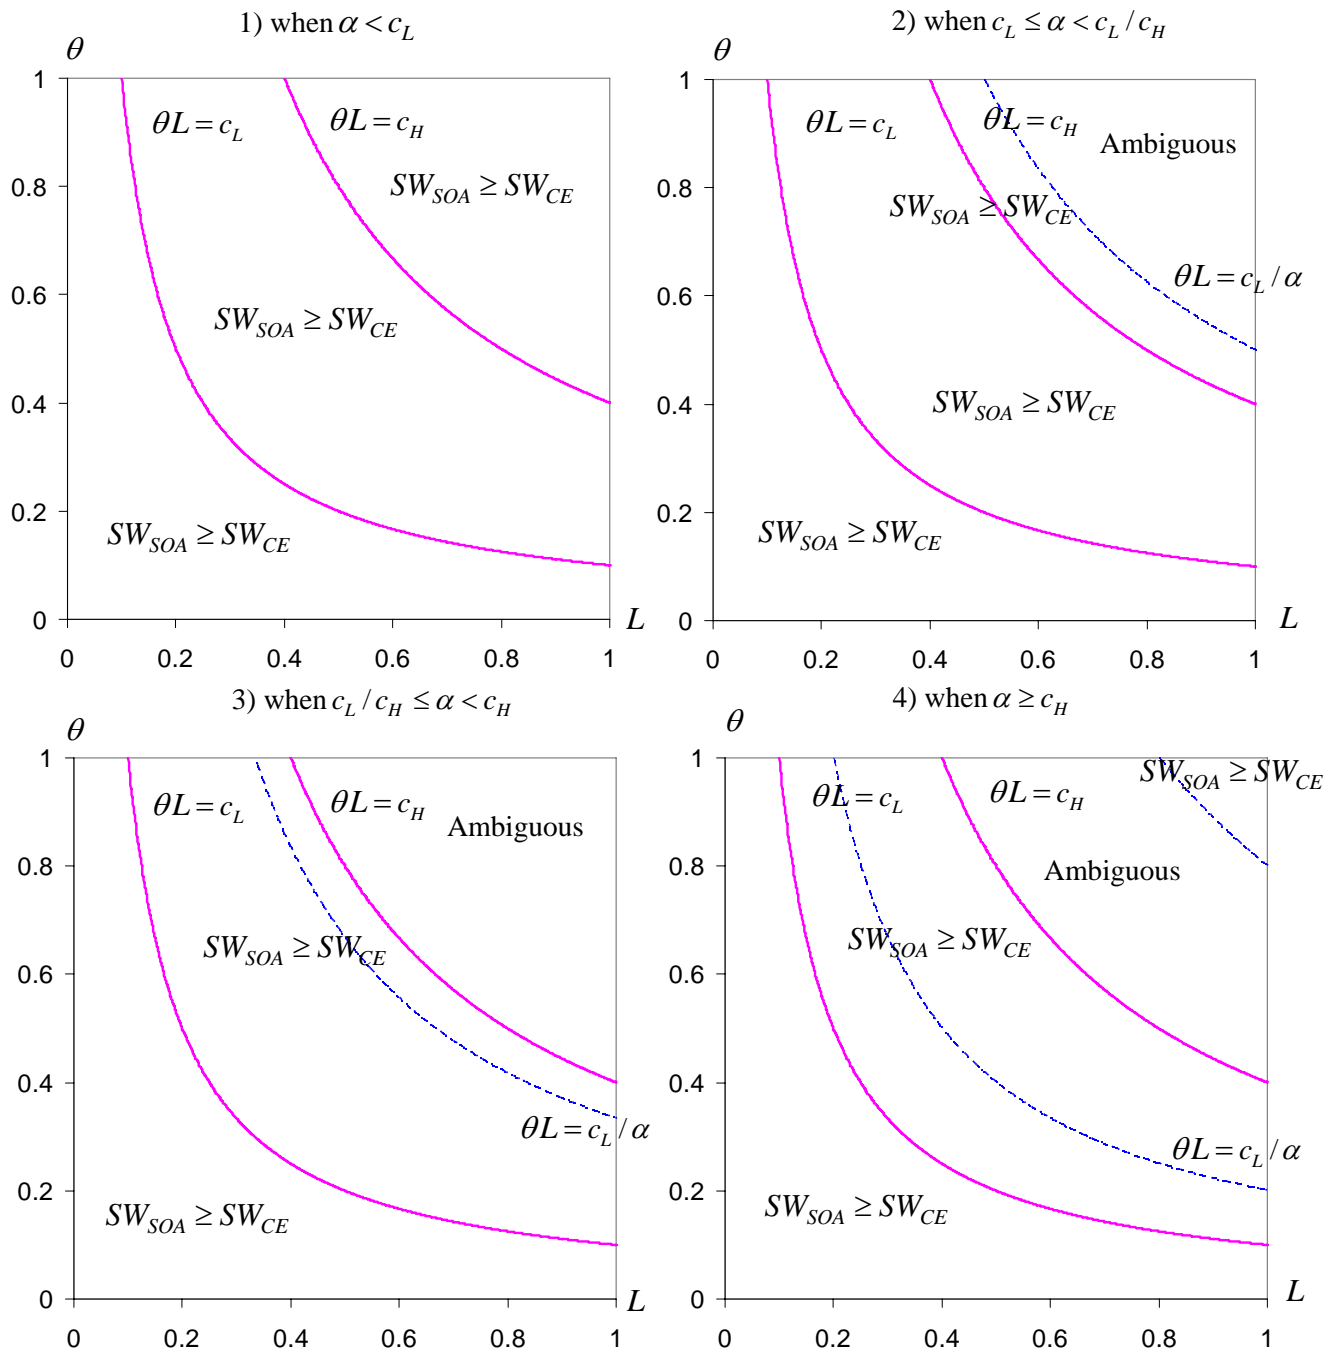


Figure 5: Comparing the seal-of-approval regime and caveat emptor regime

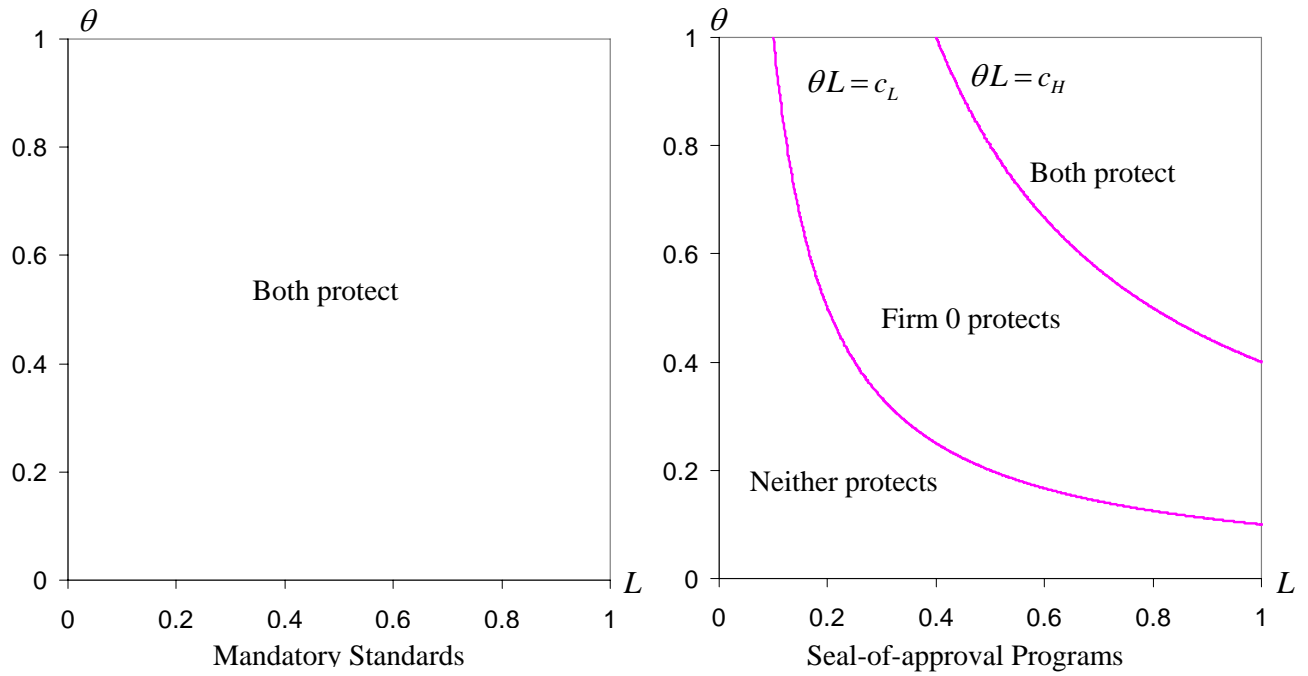


Figure 6: Outcomes under mandatory standards regime and under seal-of-approval regime