Privacy Protection and Consumer Retention*

Bruno Jullien† Yassine Lefouili‡ Michael Riordan§

January 2018
Work in Progress

Abstract

We study the incentives of a website to sell its customers' personal information. Selling information can result in good, bad or neutral experiences for consumers who learn about their vulnerability to bad matches through experience. The cost to the website of selling information is the risk that a bad experience makes consumers end their relationship with the website. The measures adopted by the website to mitigate that cost are neither contractible nor discernible by consumers. Nevertheless, in equilibrium, the website has incentives to be cautious about selling information or spend resources to screen good matches. We characterize the equilibrium privacy policy of the website and its welfare properties and discuss the difficulty of welfare-improving regulations.

Keywords: Privacy, learning, signal jamming, regulation.

*We are grateful to Dirk Bergemann, Jacques Crémer, Gabrielle Demange, Hanna Halaburda, Christian Hellwig, Justin Johnson, Martin Peitz, Georgios Petropoulos, Mike Powell, Roland Strausz, Hal Varian, and our audiences at the MACCI Summer Institute (Schloss Gracht), the TSE Digital Workshop (Toulouse), the 9th Bi-Annual Postal Economics Conference (Toulouse), the Workshop on the Digital Economy (Paris), the EARIE 2016 (Lisbon), the Annual Meeting of the Toulouse Network for Information Technology 2016 (Seattle), the APIOC 2016 (Melbourne), the Northwestern-Toulouse IO Workshop 2017 (Evanston), the 2017 HKUST Workshop on Industrial Organization (Hong-Kong), the Workshop on Institutions, Individual Behavior and Economic Outcomes (Alghero), CRESSE 2017 (Crete), the Gilbert Center for Applied Economics Conference (Berkeley), the Workshop on the Economics of Platforms 2017 (Barcelona), the 10th TOI Workshop on Industrial Organization (Maitencillo), and seminars at the Paris School of Economics, the University of Melbourne and the University of Los Andes, for useful comments and discussions. The financial support of the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 670494) is gratefully acknowledged.

†Toulouse School of Economics, CNRS, Toulouse, France. E-mail: bruno.jullien@tse-fr.eu.

‡Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France. E-mail: yassine.lefouili@tse-fr.eu.

§Columbia University, New-York, USA. E-mail: mhr21@columbia.edu.

1 Introduction

Should a website's privacy policy be regulated? The optimistic view contends that market forces discipline firms: those who do not adopt a privacy policy in line with consumer preferences make less profits. The contrary view is that regulation is necessary to improve consumer welfare: public authorities should define privacy rules that bind firms. We contribute to this debate on privacy regulation by investigating how the prospect of continued interaction affects a website's incentive to protect its customers' personal information. Our middle-of-the-road view is that, while firms generally have positive but imperfect incentives to protect consumer privacy, it is difficult to design rules that reliably improve consumer welfare.

There are many ways to monetize a website.[1] Some do not raise substantial privacy concerns, including banner advertising or direct merchandising at the website. But other ways a website earns revenue use personal consumer information, including behavioral marketing that targets ads with information about consumers online activities. For example, a website might sell a "lead" by connecting a consumer to another company, who is interested in targeting ads to consumers who have expressed an interest in the website's content.[2] Regulators have raised privacy concerns and recommended principles of greater transparency and consumer choice regarding online behavioral marketing activities (FCC, 2009).

Motivated by such issues, we develop a theory of privacy protection for an environment in which consumers learn from experience about their utility of visiting a website, which depends both on a consumer's value of website content and vulnerability to intrusive advertising that might result from the visit. Under these conditions a website's privacy policy can affect consumer retention by altering consumer experience and, thus consumer learning. In our model, a website offering a free service and earning revenue from banner advertising (or another activity that does not raise privacy concerns), gains personal information about consumers who visit, and can sell that information to a third party, (or profit from some other method of behavioral marketing that does raise privacy concerns.) Such information sales could be beneficial, for example, by enabling targeted advertising that informs consumers of desirable products, or intrusive, for example, by increasing exposure to spam, phishing or malware.[3] Those customers experiencing intrusions become

[1] See, for example, https://websitesetup.org/33-ways-to-monetize-website /.

[2] Cost-per-action (CPA) marketing platforms appear to implement something similar to this simple business model for website monetization.

[3] See Kim, Jeong, Kim, and So (2011) for a taxonomy of bad things that happen on the Internet.

more pessimistic about their utility from a return visit to the website, and this learning mechanism gives the website an incentive to adopt a privacy policy that limits third party information sales in order to protect at least partially its customers from intrusion and thereby improve customer retention.

More precisely, we study a simple two-period model. In the first period, a population of consumers enjoy a free service provided by a website. The service is an experience good, for which the consumers have heterogeneous values in the second period. The website's privacy policy is a choice of "precaution", defined as the probability that the website declines to sell customer information to a third party in the first period. Consumer do not directly observe the website's choice of precaution, and instead form equilibrium beliefs. A third-party sale of information results in a consumer experience that may be good, bad, or neutral; a neutral experience is the same as if the website does not sell information. Consumers are unsure of their vulnerability, defined as the probability of a bad experience. In the first period, consumers have identical prior beliefs about vulnerability. In the second period, consumers use Bayes Rule to update their beliefs about vulnerability based on their first-period experiences. The consumers optimize whether to return to the website, given their utility value of the website service and their posterior beliefs of vulnerability. An equilibrium is a profit-maximizing degree of precaution and consumer posterior beliefs (determining their willingness to make a return visit) that are mutually consistent.

Equilibrium is well behaved in this baseline model. We show that precaution is decreasing in the first-period value of third-party sales relative to the second-period value of retaining customers. In a strong privacy regime, this relative value is sufficiently low that the website never sells information. Conversely, in a weak privacy regime, the relative value is sufficiently high that the firm always sells information. There is also an intermediate random privacy region, in which precaution is decreasing in the relative value.[4] Furthermore, equilibrium precaution is greater the more sensitive is consumer retention to beliefs about vulnerability, and the more sensitive consumer beliefs are to experience.

The website's equilibrium incentive for precaution is at best only imperfectly aligned with consumer incentives. This is not surprising because consumers cannot verify the website's privacy policy. We show that if the website were able to commit to privacy policy, e.g. because public authorities enforce a mandatory transparency policy, it would choose a lower degree of precaution than in equilibrium. We also establish that short term

[4]Formally, the website adopts a mixed strategy in this region. The mixed strategy can be "purified" by introducing into the model a vanishingly small amount of incomplete information about the cost of third party sales. See Bagwell and Wolinsky (2002).

(i.e. first period) consumer welfare increases (decreases) with the degree of precaution if consumers' expected utility from third-party sales is negative (positive), while long term (i.e. second period) consumer welfare always decreases with the degree of precaution. Consequently, requiring a transparent privacy policy benefits consumers unambiguously only in those circumstances for which the risks of intrusion are sufficiently small.

Robust welfare-improving regulations are not readily apparent. For example, a tax on information sales increases precaution, but this is detrimental to consumers if the expected benefit of information sales is positive. We also consider an opt-out rule allowing customers to refuse permission for the website to sell their personal information. In the most interesting scenario, in which consumers opt out in the second period if and only if they have a bad experience in the first period, and assuming the website prefers consumers not to opt out, we show that a mandatory opt-out policy leads to more precaution. An opt out-rule necessarily improves consumer welfare in the second-period by revealed preferences, but as with the tax, greater precaution in the first period is not beneficial if, given prior beliefs about vulnerability, the expected consumer benefit of information sales is positive. If, however, intrusions are sufficiently likely or damaging, then either a tax or an opt-out rule improves consumer welfare.

We also study three extensions of the baseline model. These extensions add positive insights about the private incentives for privacy protection, but do no reverse our normative conclusion about the difficulty of designing robust welfare-improving regulations. Firstly, we allow for multiple websites with multi-homing consumers, and find that, if competition between multiple websites reduces the price of information, there is less precaution compared to the single website case. Secondly, we allow for costly verification that third-party uses of personal information are benign, enabling the website to prevent bad consumer experiences. The website's strategy is then given by the degree of precaution and level of verification. We characterize the equilibrium when the website cannot commit to its strategy, and show in particular that the equilibrium level of verification is non-monotonic in the value of personal information. We also show that verification and precaution are strategic substitutes. Thirdly, we study an overlapping generation version of our model in which each generation of consumers lives for two periods, and the website cannot discriminate between new consumers and a share of returning consumers (who can conceal their identity). We show that the website's inability to discriminate leads to less privacy for young consumers and that this effect is stronger the greater the share of non-identifiable consumers.

The economics of privacy literature echoes various themes from the broader information

economics literature (Acquisti, Taylor and Wangman, 2016). For example, the disclosure of personal information can improve the allocation of goods and services via targeted advertising or price discrimination, while secrecy potentially leads to market failure due to adverse selection or costly signaling. We contribute to the literature by developing a related but neglected theme: website privacy policy influences how consumers learn about their tastes for a product attribute. In our model, consumers care both about their direct utility from website services, and their expected utility from third party sales. Website privacy protection in essence is a product attribute, the value of which consumers learn imperfectly from experience. The website chooses privacy protection with the aim of influencing consumer beliefs, but, as is typical of signal-jamming models of product quality, consumers in equilibrium see through these incentives, and correctly predict the firm's actions.[5]

The rest of our paper is organized as follows. Section 2 lays out the baseline model and presents equilibrium and welfare analyses. Section 3 analyzes the effects of tax and opt-out policies, while Section 4 addresses the extensions. All proofs are relegated to the Appendix.

2 Baseline model

Consider a website facing a unit-mass population of consumers for two periods: period 0 and period 1. The service offered to consumers is free and the website obtains an exogenous revenue a per consumer derived from merchandising and/or banner advertising. Moreover, the website acquires personal information on its customers and can sell it to third parties. There is a unit-mass of third parties arriving in each period. For simplicity, we assume that each consumer is of interest to one third party. The third party has to buy customer information from the website if it wants to identify the consumer it is interested in getting

[5] To illustrate signal-jamming incentives for product quality, consider a firm selling an experience good for which a positive experience requires both a high-quality product and a discerning consumer. More precisely, a consumer has a positive experience with probability $q\theta$, where $q \in \{0, 1\}$ is a characteristic of the product and $\theta \in \{0, 1\}$ is a characteristic of the consumer. In response to a positive experience, the consumer forms a posterior belief $r_G = 1$ of being a discerning type; otherwise, the consumer's posterior belief is $r_N = 0$. Thus, even though quality is unobservable, the firm has an incentive for high quality in order to convince a discerning consumer to make a repeat purchase. Our model of equilibrium privacy provision follows a similar logic. The website invests in privacy protection to influence consumers' beliefs about the utility of returning to the website. See Judd and Riordan (1994) and Board and Meyer-ter-Vehn (2013) for more elaborate dynamic signal-jamming model of product quality.

matched with.[6]

A match between a consumer and a third party can lead to three outcomes, observable by the consumer, that we refer to as "experiences". For any match, there is a probability λ that the experience is good (G), which generates a utility $U_G > 0$ for the consumer. However, there is a probability θ that the experience is bad (B), which results in a utility $U_B < 0$ for the consumer. Finally, with the complementary probability $1 - \lambda - \theta$, the experience is neutral (N) for the consumer who gets a utility $U_N = 0$. If a consumer is not matched with a third party, her experience is neutral, and cannot be distinguished from a neutral experience induced by a match.

We assume that parameter θ is a characteristic of the consumer, and is unknown to all parties (including the consumer). Each consumer may be highly vulnerable to a bad experience, i.e., $\theta = \theta_h$, or weakly vulnerable, i.e. $\theta = \theta_l < \theta_h$. We denote by r_0 the *ex ante* probability of being weakly vulnerable. Vulnerability can be interpreted in several ways. One interpretation is that of imperfect targeted advertising where θ is the likelihood that the ad is mistargeted in a way that generates a disutility for the consumer, $1 - \theta - \lambda$ is the probability that the ad is mistargeted in a way that does not affect the consumer, and λ is the probability that the advertising is well-targeted (and thus generates a positive utility). Another interpretation is that of deceptive advertising (e.g. spam or ransomware) where θ measures the probability that a third party who acquires the customer's personal information uses it in a deceptive way. Finally, some consumers may be better protected against aggressive intrusions, for instance because they have a better antivirus or firewall, so that intrusion by harmful third parties is more likely to fail.

Denote

$$\bar{\theta}(r) \equiv r\theta_l + (1 - r)\theta_h$$

the expected vulnerability of a consumer endowed with beliefs r about being of type $\theta = \theta_l$, and define the expected "match utility" as the expected benefit a consumer derives from being matched with a third party, i.e.

$$M(r) \equiv \lambda U_G + \bar{\theta}(r) U_B.$$

The utility u a consumer derive from the service is uncertain at the beginning of period 0: it is distributed with cdf $G(.)$ and mean u_0, which we assume to be large enough for all consumers to participate in period 0. However, each customer learns, upon consuming the

[6] Allowing third parties to obtain consumers' personal information from other sources with an exogenous positive probability leads to qualitatively similar results.

service in period 0, the utility she derives from it.

Moreover, we suppose that all third parties are willing to pay the same amount v_0 for consumers' personal information. Under this simplifying assumption, the strategy of the website in period 0 is to decide whether to sell information at price v_0 or not. In Section 4.2, we will consider an extension in which the website can also verify third party use of customer information to determine whether the match will generate a good experience or not before selling the information.

During period 0, a consumer observes whether she has a good experience, a bad experience, or a neutral experience. This implies that the consumer learns about her vulnerability θ by observing the realized event. At the end of period 0, the consumer revises her beliefs about θ using Bayes Rule. We denote by r_1 the updated probability that $\theta = \theta_l$.

Assume that the situation in period 0 is repeated in period 1, except that the value of personal information to a third party is v_1. It is immediate that the website always sells personal information at price v_1 in period 1 (as there is no future interaction with consumers). The value V_1 of a returning consumer is then equal to $\delta^F (a + v_1)$ where δ^F is the firm's discount factor. The probability to retain a customer depends on her revised beliefs r_1 and is denoted $Q(r_1)$, which is an increasing function of the posterior r_1. More precisely, we have

$$Q(r_1) = \Pr\{u + M(r_1) > 0\}.$$

The expected future utility depends on the posterior r_1 at the end of the period and is given by

$$U_1(r_1) = \delta^C \int_{-M(r_1)}^{+\infty} (u + M(r_1))\, dG(u).$$

where δ^C is the consumers' discount factor. Note that $U_1(r_1)$ is a convex function of the posterior r_1.

Notice that the second period of our model is fully characterized by the value V_1, the retention rate $Q(.)$ and the utility $U_1(.)$. This implies that our model can be easily adapted to alternative assumptions about the second period.

2.1 Strategies and beliefs

A strategy for the website consists of a probability $X \in [0, 1]$ of refusing in the first period to sell personal information to a third party. We will refer to X as (the degree of) *precaution*, and will say that we have a *full privacy* (protection) policy when $X = 1$, and a *no privacy*

(protection) policy when $X = 0$.

The distribution of outcomes observed by consumers at the end of period 0 is what matters from both the consumers' and the website's perspective. The probability of a good experience (event G) and the probability of a bad experience (event B) are

$$p_G(X) = \lambda(1 - X),$$

and

$$p_B(X) = \bar{\theta}(r_0)(1 - X),$$

respectively. Both are decreasing in X. In contrast, the probability of a neutral experience (event N)

$$p_N(X) = 1 - p_G(X) - p_B(X) \tag{1}$$

is increasing in X.

Denote by r_G, r_B and r_N the updated probabilities that $\theta = \theta_l$ at the beginning of period 1, after the events G, B and N are observed, respectively. Beliefs are updated using Bayes Rule. In particular,

$$r_G = r_0$$

and

$$r_B = \frac{\theta_l}{\bar{\theta}(r_0)} r_0$$

Notice that r_G and r_B are not affected by the website's strategy. This is because the website's strategy does not affect the nature of the experience induced by a match (good, bad or neutral) conditional on the information being sold. However, the likelihood that the consumer has a neutral experience depends on the website's strategy. Therefore, the updated probability that $\theta = \theta_l$ when event N is observed depends on the degree of precaution X:

$$r_N = \phi(X) \equiv \frac{1 - (\lambda + \theta_l)(1 - X)}{1 - \left(\lambda + \bar{\theta}(r_0)\right)(1 - X)} r_0. \tag{2}$$

The posterior $\phi(X)$ decreases in X because a higher X reduces the likelihood that a neutral experience results from low consumer vulnerability (rather than the possibility that information was not sold). Moreover, a neutral experience is good news in our model in the sense that for $X < 1$:

$$r_B < r_G < r_N.$$

The intuition behind this result is that a neutral experience can result from the possibility

that a third party obtained the consumer's personal information but the match generated a neutral experience due to the consumer's low vulnerability.

2.2 Equilibrium analysis

Let us first derive the website's optimal strategy for given consumer beliefs r_N. Note that, while the website cannot affect consumers' beliefs, it can affect the likelihood of the three events that can occur (B, G, and N). The benefit of selling customer information depends on the following trade-off. Selling personal information yields an extra revenue v_0 but raises the probability of a good experience by λ and the probability of a bad experience by $\bar{\theta}(r_0)$. This reduces the retention probability by $\lambda[Q(r_N) - Q(r_G)] + \bar{\theta}(r_0)[Q(r_N) - Q(r_B)]$ compared to the situation where personal information is not sold. This cost follows from the change in consumers' beliefs when they have a non-neutral experience (events G and B) compared to a neutral experience (event N).

For given beliefs r_N, the website chooses not to sell customer information (i.e. $X = 1$) if the value of information is such that

$$\frac{v_0}{V_1} < B(r_N) \equiv \lambda[Q(r_N) - Q(r_G)] + \bar{\theta}(r_0)[Q(r_N) - Q(r_B)],$$

while it chooses to sell customer information (i.e. $X = 0$) if the reverse strict inequality holds. The function $B(r_N)$ measures on the sensitivity of consumer retention to beliefs about vulnerability for given beliefs. When $B(r_N)V_1 = v_0$ the marginal benefit of precaution is equal to the marginal cost, and any value of $X \in [0, 1]$ is optimal for the website. Thus, for given beliefs r_N, the website's optimal level(s) of precaution is(are) given by:

$$X^{br}(r_N) \in \arg\max_{X \in [0,1]} X(B(r_N)V_1 - v_0) \tag{3}$$

Hence, an equilibrium in which the website adopts a full privacy policy ($X = 1$, $r_N = \phi(1)$) exists if and only if

$$\frac{v_0}{V_1} \leq \psi^f \equiv B(\phi(1)) \tag{4}$$

Similarly, an equilibrium with no privacy ($X = 0$, $r_N = \phi(0)$) exists if and only if

$$\frac{v_0}{V_1} \geq \psi^n \equiv B(\phi(0)) \tag{5}$$

Notice that, because $\phi(0) > \phi(1)$, precaution is more attractive when consumers expect

no privacy than when they expect full privacy. As a result, we have $\psi^n < \psi^f$, and there is a range of values $v_0/V_1 \in \left(\psi^n, \psi^f\right)$, for which no pure-strategy equilibrium exists. In this range, the website randomizes between selling and not selling customer information and the equilibrium level of precaution X^* is such that

$$\frac{v_0}{V_1} = B\left(\phi(X^*)\right) \tag{6}$$

Thus, in a random privacy regime equilibrium precaution equates the marginal benefit the marginal cost of precaution with endogenous beliefs.

Therefore, we get the following equilibrium characterization:

Proposition 1 *For any values of v_0 and V_1, there exists a unique equilibrium. The equilibrium level of precaution is a non-increasing function of the ratio v_0/V_1 and there exist positive thresholds $\psi^f < \psi^n$ such that:*

- *The website provides full privacy ($X^* = 1$) if $\frac{v_0}{V_1} \leq \psi^f$.*
- *The website's policy is random ($X^* \in (0,1)$) if $\psi^f < \frac{v_0}{V_1} < \psi^n$.*
- *The website provides no privacy ($X^* = 0$) if $\frac{v_0}{V_1} \geq \psi^n$.*

Moreover, from (6) and the fact that $\phi(X)$ is decreasing in X it follows that the equilibrium degree of precaution X^* is non-increasing in v_0 and non-decreasing in V_1.

Also, equilibrium precaution is non-decreasing in the sensitivity of retention to beliefs about vulnerability – measured by the slope of $Q(r)$ for $r \geq r_B$, and in the sensitivity of beliefs to experience – measured by the (absolute value of the) slope of $\phi(X)$ for $X \in [0,1]$.[7]

To illustrate the implied comparative effects, assume for simplicity that $\theta_l = 0$, the retention is always interior (i.e. $0 < Q(0) < Q(1) < 1$) and $G(u)$ is uniform with density α on its support. In this scenario, the relevant formulas simplify to

$$B(r) = \alpha\theta_h[\lambda(r - r_0) + (1 - r_0)\theta r]|U_B|$$

and

$$\phi(X) = (1 - \lambda(1 - X))r_0/(1 - [\lambda + (1 - r_0)\theta_0](1 - X)).$$

In this case, it is straightforward to verify that X^* does not depend on U_G, and is increasing in $|U_B|$, λ, θ_h, and r_0.

[7] Notice that increasing $Q'(r)$ for all $r \geq r_B$ raises the height of $B(r)$ for all $r \geq r_0$, while, as $\phi(1) = r_0$, increasing $|\phi'(X)|$ for all X increases the height of $\phi(X)$ for all $X \in [0,1]$.

3 Policy analysis

In this section we use our model to understand the effects of various public policies aimed at improving consumer privacy.

3.1 Welfare

Before discussing the effects of public privacy policies, we define more precisely our measures of welfare. First, for given beliefs r_N, a website choosing a precaution level X makes an expected profit

$$\Pi\left(r_N, X\right) = \left(1 - X\right) v_0 + \left\{p_G\left(X\right) Q(r_0) + p_B\left(X\right) Q(r_B) + p_N\left(X\right) Q\left(r_N\right)\right\} V_1.$$

We then have the immediate result that the website profit is increasing in r_N (for a given X).[8]

To analyze consumer surplus, we decompose consumer utility into two components: the utility $\bar{U}_0\left(X\right)$ from consumption in period 0 (the "short-term" utility) and the expected utility $\bar{U}_1\left(r_N, X\right)$ from consumption in period 1 (the "long-term" utility) where

$$\bar{U}_0\left(X\right) = u_0 + \left(1 - X\right) M\left(r_0\right),$$

and (recalling that the level of precaution in period 1 is $X_1 = 0$ in equilibrium in our baseline model)

$$\bar{U}_1\left(r_N, X\right) = p_G\left(X\right) U_1\left(r_G\right) + p_B\left(X\right) U_1\left(r_B\right) + p_N\left(X\right) U_1\left(r_N\right).$$

The short-term utility may either decrease or increase with X depending on whether consumers' expected matching utility with prior beliefs, i.e. $M\left(r_0\right)$, is positive or negative. The long-term utility increases with the posterior r_N and with the level of precaution X.[9]

3.2 Transparency and commitment

The first question we address is how a policy that affects X should be conducted. One such policy experiment is to impose full transparency regarding the website's privacy policy. In

[8] More precisely, $\frac{\partial \Pi}{\partial r_N} = p_N\left(X\right) Q'\left(r_N\right) V_1 > 0$.

[9] $\frac{\partial \bar{U}_1(r_N, X)}{\partial X} = \frac{dp_G}{dX}\left(U_1\left(r_G\right) - U_1\left(r_N\right)\right) + \frac{dp_B}{dX}\left(U_1\left(r_B\right) - U_1\left(r_N\right)\right) > 0$ since $\frac{dp_G}{dX}$ and $\frac{dp_B}{dX}$ are negative and $U_1(r_B) < U_1(r_G) < U_1(r_N)$.

our setup this would require the website to reveal the value of X and commit to it. An important aspect of this policy is that it would make the announcement of X credible. This would modify the equilibrium degree of precaution derived in the previous subsection to the one chosen by a website that can publicly commit to a given strategy.[10]

The difference between the case where X is unobservable and the case where it is observed by consumers is twofold. First, under transparency, the website can affect the posterior beliefs by its choice of privacy policy. Since the profit is increasing in consumer beliefs r_N the website will change X in a direction that raises r_N. This is driven by the fact that the website would like consumers to interpret a neutral experience as a stronger signal about their low vulnerability, which is achieved by selling personal information more often. Second, with a transparent privacy policy, it may no longer be optimal for the website to choose a no-privacy regime in period *1*, i.e. $X_1 = 0$, where X_1 denotes the level of precaution in period 1. Indeed, when matches are detrimental to consumers, commitment to some precaution ($X_1 > 0$) in the beginning of period 1 may boost demand. By contrast, when second-period matching is beneficial to consumers on average, the website will choose no privacy in the second period. The next proposition provides conditions under which a regulation mandating transparency leads to a lower precaution level in the first period. It also provides a condition under which such a regulation does not affect the website's second-period privacy policy.

Proposition 2 *A regulation mandating transparency causes the website to choose a weaker privacy policy (i.e. a lower X) in the first period than in the equilibrium with no commitment if at least one of the following conditions holds: (i) the second-period expected value of a match is always positive, i.e. $M(r_B) > 0$ or (ii) $G'(u) + uG''(u) > 0$ for any u. Moreover, if (i) is satisfied then the website commits to a no privacy policy in the second period.*

Let us now consider the effect of a regulation mandating transparency on consumers under the circumstances considered in the proposition above. First, note that consumers benefit in period 1 from the belief-improving effect of a less cautious strategy: decreasing X raises the posterior beliefs $r_N = \phi(X)$ and thus indirectly $\bar{U}_1$, but it also alters the distribution of the posteriors. The total derivative of $\bar{U}_1(\phi(X), X)$ with respect to X is

[10] This scenario corresponds to the case, featured in previous literature (see e.g. O'Brien and Smith, 2014), in which the privacy policy is a publicly observable quality variable.

given by

$$\frac{d\bar{U}_1}{dX} = \underbrace{p_N U_1'(r_N)\frac{\partial \phi}{\partial X}}_{<0} + \underbrace{\frac{dp_G}{dX}(U_1(r_G) - U_1(r_N)) + \frac{dp_B}{dX}(U_1(r_B) - U_1(r_N))}_{>0}.$$

Thus, the overall long-term effect of lowering the degree of precaution on consumers is *a priori* ambiguous. However, this long-term effect can be rewritten as shown in the lemma below.

Lemma 1 *The long-term effect of first-period precaution on consumer surplus is given by*

$$\frac{d\bar{U}_1}{dX} = \frac{dp_G}{dX}(U_1(r_G) - U_1(r_N) + U_1'(r_N)(r_N - r_G)) + \frac{dp_B}{dX}(U_1(r_B) - U_1(r_N) + U_1'(r_N)(r_N - r_B)).$$

Recall that the future utility is convex in the posterior, implying that $U_1(r) - U_1(r_N) + U_1'(r_N)(r_N - r) > 0$. Given that p_G and p_B decrease in X, the long-term effect $d\bar{U}_1/dX$ is always negative. In other words, consumers benefit from a decrease in X in the long run.

If the expected matching utility with prior beliefs is weakly positive, i.e., $M(r_0) \geq 0$, then the short-run effect of lower precaution in the first period is also positive. In this case, the overall effect of lower precaution on consumers is positive.

However, if the expected matching utility with prior beliefs is negative, i.e. $M(r_0) < 0$, then consumers are negatively affected in period 0 by a weaker privacy policy, which creates a tension between the short-term and long-term effects of more/less precaution on consumers.

Moreover, a regulation mandating privacy always leads to a weakly higher precaution level in period 1, which is weakly beneficial (detrimental) to a consumer with updated beliefs r_1 if $M(r_1) \leq 0$ (≥ 0).

Consider first the case in which the second-period expected matching value is always positive, i.e. $M(r_B) > 0$. From Proposition 2 we know that a regulation mandating privacy does not affect the second-privacy policy in this case. Therefore, we can conclude that such regulation raises both short-term and long-term consumer surplus in this scenario.

Consider now the other polar case where the second-period expected matching value is always negative, i.e. $M(\phi(1)) < 0$ and assume that $G'(u) + uG''(u) > 0$ for any $u \geq 0$ such that a regulation mandating privacy leads to a lower precaution level in the first period.[11] In this case, the long-term consumer surplus is (weakly) higher under such a regulation

[11] It is sufficient that condition (ii) in Proposition 2 be satisfied for positive values of u because we are focusing on a case in which $M(r_1)$ is always negative.

because, in the second period, consumers benefit from both lower precaution in the first period and weakly higher precaution in the second period. The effect of such a regulation on short-term is, however, negative in this scenario.

The following proposition summarizes the analysis above.

Proposition 3 - *If the second-period expected value from a match with a third party is always positive then a regulation mandating transparency leads to higher short-term and long-term consumer surplus.*

- *If the second-period expected value from a match with a third party is always negative and* $G(.)$ *is not too convex then a regulation mandating transparency leads to lower short-term consumer surplus and higher long-term consumer surplus.*

This proposition implies that when the expected value from a match for consumers is negative and consumers are sufficiently impatient, consumers do not benefit from a regulation mandating transparency. However, when the expected matching value is positive or consumers are sufficiently patient, they benefit from such a regulation.

3.3 Taxation

One potential way of affecting firms' incentives to sell personal information is to apply a specific tax treatment to transactions involving customer information. Such a tax would not only alter the direct gains v_0 from selling information in period 0 but also the value of retaining a consumer V_1. In our two-period setup, suppose that a proportional tax τ (which may be positive or negative) is levied on transactions involving customer information. Then, the revenue from selling information is $(1-\tau)\,v_0$ while the value of retaining a consumer is $\delta^F\left[a+(1-\tau)\,v_1\right]$. A (positive) tax thus reduces both the revenue from selling customer information and the retention value. However, recall that the equilibrium level of precaution depends only on their ratio. Therefore, we immediately obtain the following result:

Proposition 4 *Denote* τ *a tax which is levied on personal information transactions in both periods. Then the equilibrium level of precaution (weakly) increases with* τ*.*

This result is quite intuitive but not completely obvious. The website compares the current revenues from selling information with the future revenues that include not only future sales of personal information but also other sources of revenues that are not taxed. Hence, the tax affects relatively more the short-run gain from selling information than the

future value of a consumer (irrespective of the value v_1). Notice that an identical tax on all revenues would be neutral.

3.4 Opt-out

An alternative policy to protect consumers is to give them control rights over their personal data. Ideally, a consumer would like to choose which third party can access her personal data and for what purpose. However, contracts are typically incomplete due to private information and lack of verifiability. Here, we assume that whether information is sold or not is verifiable, but not the nature of the match (good, bad or neutral) with the third party buying this information. We allow consumers to opt out, which means they can prevent any sale of personal information (the full privacy regime would then prevail). We assume that in the first-period consumers do not find it optimal to opt out but that they may decide to do so after revising their beliefs about their vulnerability. Thus, at the end of the first period, consumers have three options: they may stop their relationship with the website, they may stay and opt in (i.e., not prevent the website from selling their personal information), or they may stay and opt out.

A consumer's choice between opting in and opting out is governed by her beliefs about her vulnerability to bad experiences. Denote

$$\bar{r} \equiv \frac{\lambda U_G + \theta_h U_B}{(\theta_h - \theta_l) U_B}$$

the (unique) solution to $M(r_1) = 0$. Consumers opt out in the second period if and only if $M(r_1) < 0$, which is equivalent to $r_1 < \bar{r}$. The website's second-period profit per retained consumer is given by $V_1 = \delta^F (a + v_1)$ if the consumer does not opt out and $\bar{V}_1 = \delta^F a$ if the consumer opts out. The probability for the website to retain a consumer who prefers to opt out is $\bar{Q} \equiv Q(\bar{r}) = 1 - G(0)$. Therefore, the retention rate when the opt-out option is available to customers is given by $\max\left\{Q(r_1), \bar{Q}\right\}$.

Allowing consumers to opt out raises the retention rate. The cost for the website is that it deprives it of the opportunity to sell personal data to third parties and, therefore, reduces the expected revenue per customer. Denoting $\bar{V}_1 < V_1$ the value for the website of a customer who chooses to opt out, the retention value of a customer is $\bar{V}_1$ if $r_1 < \bar{r}$, and V_1 if $r_1 \geq \bar{r}$.

Let us consider, for now, the scenario in which the website always prefers that consumers do not opt out:

Assumption $Q(r_B)V_1 > \bar{Q}\bar{V}_1$.

The most interesting scenario is the one where $r_B < \bar{r} < r_0$.[12] This condition ensures that consumers are willing to participate in the first period but opt out in the second period if they have a bad experience in the first period.

Assumption $r_B < \bar{r} < r_0$.

Note that, as far as the website's profit is concerned, the only change introduced by the option to opt out is that the website's future revenues when the experience is bad are reduced from $Q(r_B)V_1$ to $\bar{Q}\bar{V}_1$. Thus, allowing customers to opt out is equivalent to reducing the expected value of retention after a bad experience. Therefore, we now analyze the effect of a marginal change in the expected revenue $Q(r_B)V_1$ from a consumer having a bad experience on the equilibrium level of precaution, and the equilibrium posterior beliefs after a neutral experience.

Lemma 2 *In the baseline model, everything else held equal, a (marginal) reduction in the value of* $Q(r_B)V_1$ *i) lowers the ex-post beliefs* r_N *if* $X^* < 1$*, ii) raises the level of precaution if* $X^* < 1$*, iii) has no effect if* $X^* = 1$.

We can now state the following result regarding the effect of a mandatory opt-out policy on the website's privacy strategy when consumers opt out only after experiencing a bad match.

Proposition 5 *A mandatory opt-out policy leads to (weakly) more precaution.*

Note that the reverse conclusion would hold if opt-out raises the expected revenue of the website, i.e. $Q(r_B)V_1 < \bar{Q}\bar{V}_1$.

An interesting question is whether the website would spontaneously offer an opt-out option to its customers. Recall that the website's profit writes as

$$\Pi(r_N, X) = (1 - X)v_0 + p_G(X)Q(r_0)V_1 + p_B(X)Q(r_B)V_1 + p_N(X)Q(r_N)V_1.$$

Let us change $Q(r_B)V_1$ by $d\varepsilon$. Then, the change in the website's profit is

$$d\Pi = p_B(X)d\varepsilon + p_N(X)Q'(r_N)V_1 dr_N$$

[12] Note that our assumption that consumers do not find it optimal to opt out in the first period imposes a lower bound on $\bar{r}$. However, it is easily checked that this lower bound is less than r_0 both when consumers are myopic (i.e. they only take into account their first-period expected when they decide whether to opt out in the first period) or forward-looking.

Since the profit increases with r_N we have two concurring effects. The website derives a direct benefit from raising future revenues, which gives it incentives to propose an opt-out option when $Q(r_B)V_1 < \bar{Q}\bar{V}_1$. Moreover, when the latter condition holds, the website derives an indirect benefit from offering an opt-out option: higher future profits allows it to commit to making the signal about customer vulnerability more informative, i.e., to raise r_N (because higher future profits lead to lower protection X, which yields higher beliefs r_N). The consequence is that the website would offer the option only when $Q(r_B)V_1 < \bar{Q}\bar{V}_1$.

4 Extensions

4.1 Multi-homing consumers

Let us consider K websites facing a unit-mass population of consumers for two periods: period 0 and period 1. Websites are not competitors on the consumer side. We assume that all consumers multi-home and are active on all websites. This means that all websites have access to all customer information and can potentially sell it to each third party. Each website is as in the basic model.

Given our assumption that there is a one-to-one match between consumers and third parties, each consumer faces the same problem as before: she can either have *one* good experience (G), *one* bad experience (B), or a neutral experience (N). The consumer then revises her beliefs about her vulnerability to bad experiences and decides whether to return to the websites.

The new feature here is that the probability of a non-neutral experience accounts for the fact that many websites can sell the same information. Thus, if x is a symmetric equilibrium probability that a website refuses to sell customer information, the total probability that a third party interested in buying such information does not acquire it is $X = x^K$. With this adjustment in the determination of X, the behavior of consumers is unchanged and, in particular, the equilibrium posterior beliefs for events G, B and N are respectively $r_G = r_0$, r_B and $r_N = \phi(X)$.

On the market for information, multi-homing affects the selling prices in both periods. For the sake of clarity, let us assume that there is perfect correlation between retentions of all websites: a consumer either leaves all websites, which happens with probability $1 - Q(r_1)$, or stays on all websites. In this case, second-period competition for selling information will drive prices of customer information to zero. Thus, the value of retaining a consumer is $V_K = \delta^F a$, strictly less than the value of retaining a single-homing consumer

$V_1 = \delta^F (a + v_1)$.

Let us now turn to competition in the first period. We model competition by assuming that all websites independently and simultaneously decide whether to sell information and, if they do, at which price p they sell it.

Notice that refusing to sell customer information amounts to quoting a price strictly above v_0. Thus, in this game, the strategy of a website can be summarized by a probability distribution over prices with cumulative distribution function $F(p)$. The probability of refusing to sell is then

$$x = 1 - F(v_0),$$

i.e. the probability to quote a price above v_0.

We now characterize a symmetric equilibrium of the game. Let us first derive the website's optimal strategy for given consumer beliefs r_N, assuming that all other websites follow the strategy $F(.)$. Then, the payoff of a website proposing a price $p \leq v_0$ is

$$\begin{aligned} & L_K + p(1 - F(p))^{K-1} \\ \text{where } L_K &= \left[\lambda Q(r_G) + \left(1 - \lambda - \bar{\theta}(r_0)\right) Q(r_N) + \bar{\theta}(r_0) Q(r_B)\right] V_K \end{aligned}$$

The term L_K is the expected continuation payoff when the third party buys the information, and the second term accounts for competition and the probability to win the sale at price p.

The payoff when refusing to sell customer information, or equivalently when quoting a price $p > v_0$, is

$$L_K + (1 - F(v_0))^{K-1} B(r_N) V_K.$$

This expression can be interpreted as follows. Refusing to sell customer information does not imply that access to that information by a third party will not occur because another website may sell the same information. In this case the payoff is L_K. However, with probability $(1 - F(v_0))^{K-1}$, all other websites also refuse to sell customer information and the payoff increases then by $B(r_N) V_K$.

Consequently, the website compares the expected revenue $p(1 - F(p))^{K-1}$ at any price below v_0 with the gain from privacy $(1 - F(v_0))^{K-1} B(r_N) V_K$.

The following lemma shows that one of the possible outcomes is that all websites sell personal information with certainty and competition dissipates fully their profits from data sales.

Lemma 3 *Suppose there are $K > 1$ websites. Then, there always exists a no-privacy*

equilibrium where all websites quote a zero-price for information, and information is always sold.

Thus, as soon as there are multiple websites that can sell the same personal information, there is a risk of a total collapse in the provision of privacy.[13] The following proposition shows that there also exists a symmetric equilibrium with a positive degree of precaution if the value of information v_0 is not too high.

Proposition 6 *Suppose there are $K > 1$ websites. Then there exists a symmetric equilibrium with positive precaution (i.e., $X > 0$) if and only if $v_0 < B(\phi(0))V_K$. It is uniquely defined by the following conditions:*

- If $v_0 \leq B(\phi(1))V_K$, then the websites provide full privacy (i.e., $X = 1$).

- If $B(\phi(1))V_K < v_0 < B(\phi(0))V_K$, then the websites' precaution and price distribution are given by

$$\begin{aligned} v_0 &= B(\phi(X))V_K; \\ F(p) &= 1 - X^{\frac{1}{K}}\left(\frac{v_0}{p}\right)^{\frac{1}{K-1}} \quad \text{for } p \in [X^{\frac{K-1}{K}}v_0, v_0]. \end{aligned}$$

From above we see that when consumers multi-home, the number of websites affects the total level of precaution only through its effect on the value of retaining a consumer. To see that, consider the limiting case where $v_1 = 0$ and thus $V_K = V_1$. In this case, in a mixed-strategy equilibrium, the revised beliefs r_N must be such that each website is indifferent between selling or not the information. When $V_K = V_1$, this indifference condition does not depend on K. However, competition for third parties tends to erode profits so that the value of retention is lower with multiple websites, i.e. $V_K < V_1$, inducing a reduction of the level of precaution.

Corollary 1 *There is (weakly) less precaution in aggregate by multiple websites than by a monopoly website. The profit of a website may increase (if v_0/V_1 is large and X is small) or decrease, compared to the monopoly case.*

Thus, the strategic effect that reduces the aggregate level of precaution may be strong enough to offset the reduction of revenues due to competition in the market for information.

In this section, we have assumed perfect correlation between the retentions of all websites, so that there is Bertrand competition and zero profit in the market for information for

[13] However, we conjecture that if there is an arbitrarily small mass ε of single-homing consumers, the zero price equilibrium exists only if $v_0 \geq B(r_N)V_K$.

second period consumers ($V_K = a$). It follows that when $K > 1$, the aggregate equilibrium level of precaution is independent of K. With imperfect correlation between second-period participations, there would be less competition and V_K should decrease with K. Thus, increasing K above $K = 2$ would have a negative effect on the protection of consumers. Moreover, the value of retention should depend on r_1, which determines the probability that the consumer single-homes (and thus that the website is a monopoly). If we assume for instance that preference shocks are independent, a consumer returning to a given website single-homes with conditional probability $(1 - Q(r_1))^{K-1}$.

Multi-homing has mixed effects on consumer welfare depending on the context. Access to a greater number of websites directly increases the utility consumers get from website content. Moreover, reduced precaution, resulting from competition in the market for consumer information, increases long-run utility, as discussed earlier. Short-term utility, however, increases only if match utility is positive. Regulatory policies, such as mandated transparency or giving consumers control over third-party access to personal information, have the same ambiguous welfare effects as in the single-homing model.

4.2 Verification of third party uses of information

In this section, we assume that the website can verify third party uses of information: the website can incur a cost z, drawn from a distribution with an increasing continuous cdf $H(.)$ over the support $\mathbb{R}_+$, to identify (with certainty) whether a match with a third party will generate a good experience or not.

4.2.1 Strategies and beliefs

The website's strategy now consists of a mapping between the verification cost z and the binary decision to verify or not, as well as the probability X of not selling the information in case there is no verification. Without loss of generality, we assume that the probability X does not depend on the verification cost.

It is straightforward that if the website is indifferent between not verifying and verifying at cost z, it strictly prefers to verify at any cost strictly below z. Then, there must exist a critical level $\hat{z}$ (potentially zero) such that the website verifies the third party's use of information if $z < \hat{z}$ and does not if $z > \hat{z}$. Let us denote by $Y = H(\hat{z})$ the probability of verification. As choosing $\hat{z}$ is equivalent to choosing Y, we use Y as the choice variable.

The expected cost of verification is

$$C_i(Y) = \int_0^{\hat{z}} z dH(z) = \int_0^Y H^{-1}(y)\, dy.$$

The function $\hat{z} = C_i'(Y)$ is thus the marginal cost of increasing the probability of verification. The marginal cost increases in Y. Therefore, we can characterize the website's strategy by a pair

$$(X, Y) \in [0,1]^2.$$

When $Y > 0$, we will say that we have a *strong privacy* (protection) policy when $X = 1$, and a *weak privacy* (protection) policy when $X = 0$. Under a strong privacy policy, the consumer is immune to unwanted intrusions from the sale of personal data, and verification is a way to raise the value to the consumer of visiting the website. The variable Y determines the benefit from allowing access to third parties that provide a good experience and is referred to as the *level of verification*. On the contrary, under a weak privacy policy, verification is the only way to avoid interactions with third parties that do not generate a good experience and, therefore, determines the level of protection against them.

Let us now provide the probability of a each type of second-period experience and determine how it depends on the degree of precaution X and the level of verification Y. The probabilities of a good experience (event G) and a bad experience (event B) are given, respectively, by

$$p_G(X, Y) = \lambda \{Y + (1 - Y)(1 - X)\}$$

and

$$p_B(X, Y) = (1 - X)(1 - Y))\, \bar{\theta}(r_0).$$

Both probabilities decrease with X because a higher level of precaution leads to less sales of personal information to third parties. Moreover, the probability of a good (bad) experience is weakly increasing (decreasing) in the level of verification because more verification decreases (increases) the likelihood that a third party generating a good (bad) experience is denied access to customer information. The probability of a neutral experience (event N) is

$$p_N(X, Y) = 1 - Y\lambda - (1 - Y)\left(\lambda + \bar{\theta}(r_0))\right)(1 - X), \tag{7}$$

This probability increases with X for the same reason why both the probability of a good experience and that of a bad experience decrease with the level of precaution. The way the level of verification affects the probability of a neutral experience depends on the level of

precaution, as shown by

$$\frac{\partial p_N}{\partial Y} = \bar{\theta}(r_0) - \left(\lambda + \bar{\theta}(r_0)\right) X.$$

The reason is that a higher level of verification has two (potential) opposite effects on the probablity of a neutral experience. First, it affects it positively by increasing the probability that a third party generating a bad experiences is denied access to customer information. Second, it affects it negatively by making it more likely that a third party generating a good experience gets access to customer information. The former (latter) effect is the only one in the extreme case of the no-privacy (full privacy) regime. More generally, the effect of verification on the probability of a neutral experience is positive (negative) if the level of precaution is low (high), i.e, if X is less (greater) than $\bar{\theta}(r_0) / \left(\lambda + \bar{\theta}(r_0)\right)$.

The posterior beliefs after a good experience and a bad experience, r_G and r_B, are the same as in the baseline model, while the posterior belief after a neutral experience is now given by

$$r_N = \Phi(X, Y) \equiv \frac{1 - Y\lambda - (1 - Y)(\lambda + \theta_l))(1 - X)}{1 - Y\lambda - (1 - Y)\left(\lambda + \bar{\theta}(r_0)\right))(1 - X)} r_0, \tag{8}$$

A neutral experience is again good news in the sense that, for $X < 1$,

$$r_B < r_G < r_N.$$

The following lemma shows how the posterior belief after a neutral experience depends on the degree of precaution and the level of verification.

Lemma 4 *i)* $\Phi(X, Y)$ *is decreasing in* X*, ii)* $\Phi(X, Y)$ *is decreasing in* Y *for any* $X < 1$*, and constant in* Y *for* $X = 1$.

On the one hand, verification increases the access to personal information of a third party generating a good experience, which makes the signal more informative about vulnerability. On the other hand, a higher level of verification reduces the access to customer information by a third party that does not generate a good experience, which makes the signal less informative about vulnerability. The lemma above shows that the latter effect dominates the former (whenever $X < 1$).

4.2.2 Equilibrium analysis

Let us now describe the website's optimal strategy for given consumer beliefs r_N. Denote

$$\Delta_G(r_N) \equiv \lambda V_1 (Q(r_G) - Q(r_N)) + \lambda v_0$$

the overall benefit from selling customer information when a match leads to a good experience, and

$$\Delta_B(r_N) \equiv \bar{\theta}(r_0) V_1(Q(r_N) - Q(r_B)) - (1-\lambda) v_0$$

the overall benefit from not selling customer information when a match does not lead to a good experience. Verification yields $\Delta_G(r_N)$ when $X = 1$ and $\Delta_B(r_N)$ when $X = 0$.

Consider first the website's decision regarding the level of precaution. Notice that $B(r_N) - v_0 = \Delta_B(r_N) - \Delta_G(r_N)$ which is the payoff of not selling information. The analysis is similar to the one in the baseline scenario (with no verification): the website's optimal level(s) of precaution does not depend on the level Y of verification but only on consumers' beliefs r_N, and is given again by

$$X^{br}(r_N) \in \arg\max_{X \in [0,1]} X(\Delta_B(r_N) - \Delta_G(r_N))$$

implying that X^{br} is non-decreasing in r_N.

Consider now the verification decision. The trade-off faced by the website is different from the one underlying the precaution decision, because verification allows to sell customer information only to third parties that generate a good experience (and would be used only for this purpose since $r_B < r_G$). Let X be a given level of precaution. For this level of precaution, verification raises the probability to sell customer information to a third party generating a good experience from $1 - X$ to 1 and, therefore, yields a benefit $X\Delta_G(r_N)$ from selling customer information to such a third party more often. Verification also reduces the probability to sell customer information to a third party that does not generate a good experience from $1 - X$ to 0, which leads to another expected benefit given by $(1 - X)\Delta_B(r_N)$. The total benefit from verifiying third parties' use of information is then the sum $X\Delta_G(r_N) + (1 - X)\Delta_B(r_N)$, and the website will verify whenever the verification cost is less than this benefit. Thus, some verification occurs (i.e., $Y > 0$) whenever this benefit, evaluated at $X = X^{br}(r_N)$, is positive. Notice that condition (3) implies that

$$\begin{aligned} X^{br}(r_N)\Delta_G(r_N) + \left(1 - X^{br}(r_N)\right)\Delta_B(r_N) &= \Delta_B(r_N) - \max_X X(\Delta_B(r_N) - \Delta_G(r_N)) \\ &= \min(\Delta_B(r_N), \Delta_G(r_N)). \end{aligned}$$

Therefore, the website's optimal verification level is $Y^{br}(r_N)$ where

$$C_i'\left(Y^{br}(r_N)\right) = \max\left\{\min\left(\Delta_G(r_N), \Delta_B(r_N)\right), 0\right\}. \tag{9}$$

Given that $\Delta_G(r_N)$ decreases in r_N while $\Delta_B(r_N)$ increases in r_N, we define r^M as the (unique) solution of

$$\Delta_G\left(r^M\right) = \Delta_B\left(r^M\right) \equiv \Delta^M \tag{10}$$

when it exists in the range $[\Phi(1,0), \Phi(0,0)]$, and set $r^M = \Phi(1,0)$ when $\Delta_G(\Phi(1,0)) < \Delta_B(\Phi(1,0))$, and $r^M = \Phi(0,0)$ when $\Delta_G(\Phi(0,0)) > \Delta_B(\Phi(0,0))$. Notice that $Y^{br}(r_N)$ is single-peaked and maximal at $r_N = r^M$, when $\Delta_G\left(r^M\right) > 0$.

Thus, an equilibrium is characterized by

$$X^{**} = X^{br}(r_N^{**}); \quad Y^{**} = Y^{br}(r_N^{**})$$

along with

$$r_N^{**} = \Phi(X^{**}, Y^{**}). \tag{11}$$

Proposition 7 *There exists a unique equilibrium* (X^*, Y^*), *which varies continuously with the value of information* v_0.

We distinguish three scenarios:

- *Strong protection* $(r_N^{**} > r^M)$: $X^{**} = 1$ and $C_i'(Y^{**}) = \max\left\{\Delta_G(r_N^{**}), 0\right\}$;
- *Random protection* $(r_N^{**} = r^M)$: $X^{**} \in (0,1)$ and $C_i'(Y^{**}) = \max\left\{\Delta^M, 0\right\}$;
- *Weak protection* $(r_N^{**} < r^M)$: $X^{**} = 0$ and $C_i'(Y^{**}) = \max\left\{\Delta_B(r_N^{**}), 0\right\}$.

We now discuss how the equilibrium depends on the value v_0 of personal information. Notice that for v_0 close to 0, the gain $\Delta_G(r_N)$ is negative for all $r_N \in [\Phi(1,0), \Phi(0,0)]$, while the gain $\Delta_B(r_N)$ is positive. Thus, for sufficiently small values of v_0, the website chooses a strong protection policy and no verification (i.e., a *full privacy* regime). Similarly, for sufficiently large values of v_0, the website chooses weak protection and no verification (i.e., a *no-privacy* regime), so that customer information is always sold.

As shown previously, when there is no verification ($Y = 0$), the equilibrium level of precaution is non-increasing in v_0 and the equilibrium moves toward more frequent access of third parties to customer information, leading to higher posterior beliefs r_N. This is, however, less obvious if $Y > 0$ because the level of verification changes as well. We need

the following regularity condition to be able to sign the comparative statics of X^* with respect to v_0.

Condition 1 (C1) *For any $X < 1$ and $Y > 0$ satisfying the conditions for random protection,* $\left(1-\lambda-\bar{\theta}\left(r_0\right)\right) Q'\left(\Phi\left(X,Y\right)\right)\left|\frac{\partial\Phi(X,Y)}{\partial Y}\right|\frac{\lambda V_1}{C_i''(Y)} < 1$.

This condition ensures that increasing the verification level does not reduce too much consumers' participation after a neutral experience, so that the website's incentives to be cautious are not too weakened. More precisely, it rules out the possibility that Y^{**} is reduced so much when v_0 increases that X^{**} must increase to maintain the equilibrium level of posterior beliefs, which could happen because Φ decreases in Y when $X < 1$.

Lemma 5 *The posterior r_N^{**} is non-decreasing in v_0. Moreover, if $C1$ holds, then X^{**} is non-increasing in v_0.*

From this lemma we know that when the value of information v_0 varies from 0 to large, the level of precaution X^{**} varies continuously and monotonically from 1 to 0, while the posterior r_N^{**} varies from $\Phi\left(1,0\right)$ to $\Phi\left(0,0\right)$. Therefore, we get the following first equilibrium characterization:

Proposition 8 *Assume that $C1$ holds.*[14] *There exist positive thresholds $v_s < v_w$ such that the website chooses strong protection if $v_0 \leq v_s$, weak protection if $v_0 \geq v_w$ and a random protection if $v_0 \in \left(v_s\left(V_1\right), v_w\left(V_1\right)\right)$. Moreover, v_s/V_1 is non-decreasing in V_1, and v_w/V_1 is non-increasing in V_1.*

Another preliminary observation is that because $C_i'\left(Y\right) = \min\left\{\Delta_G\left(r_N\right), \Delta_B\left(r_N\right)\right\} \leq \Delta^M$, there is no verification if $\Delta^M < 0$ for all v_0. From condition (10) we get that

$$\Delta^M = \lambda\left[\left(1-\lambda\right)Q\left(r_G\right) - \bar{\theta}\left(r_0\right)Q\left(r_B\right) - \left(1-\lambda-\bar{\theta}\left(r_0\right)\right)Q\left(r^M\right)\right]V_1 \tag{12}$$

which is negative for all r^M if

$$Q\left(\Phi\left(1,0\right)\right) \geq \frac{\left(1-\lambda\right)Q\left(r_G\right) - \bar{\theta}\left(r_0\right)Q\left(r_B\right)}{1-\lambda-\bar{\theta}\left(r_0\right)} \tag{13}$$

More precisely, we have the following result:

[14] For this proposition, $C1$ needs to hold only at $X = 0$.

Proposition 9 *There is no verification of third party use of customer information if (13) holds. Otherwise, there exist* $\underline{\psi}^i < \bar{\psi}^i$ *such that verification occurs if and only if* $\underline{\psi}^i < v_0/V_1 < \bar{\psi}^i$, *with*

i) $\underline{\psi}^i V_1 < v_s < \bar{\psi}^i V_1 < v_w$ *if* $Q(\Phi(0,0)) > \frac{(1-\lambda)Q(r_G)-\bar{\theta}(r_0)Q(r_B)}{1-\lambda-\bar{\theta}(r_0)}$;

ii) $\underline{\psi}^i V_1 < v_s < v_w < \bar{\psi}^i V_1$ *if* $Q(\Phi(0,0)) < \frac{(1-\lambda)Q(r_G)-\bar{\theta}(r_0)Q(r_B)}{1-\lambda-\bar{\theta}(r_0)}$.

Verification occurs under strong privacy if $\Delta_G(r_N) > 0$ while it occurs under weak privacy if $\Delta_B(r_N) > 0$. While the benefit $\Delta_G(r_N)$ increases with v_0, the benefit $\Delta_B(r_N)$ decreases with v_0. Hence, verification occurs only for some intermediate range of values v_0 of the information.

When the posterior r_N is large for any strategy, the website abandons strong privacy at levels of v_0 such that $\Delta_G(r_N) = \Delta_B(r_N) < 0$ so that the benefit of verifying third party use of customer information is negative for all values of v_0. Verification is not useful in this case.

Let us finally consider the way the equilibrium level of verification Y^{**} depends on v_0. Verification allows to restrict sales to third parties generating good experiences, which induces a short-term revenue loss that depends on the price v_0 but also on the level of precaution. Under the strong privacy regime, raising v_0 makes verification more attractive as it generates more sales. In contrast, under the weak privacy regime, raising v_0 makes verification less attractive as it reduces the probability to sell customer information. The next proposition shows that the random privacy regime is similar to the weak privacy regime in this respect.

Proposition 10 *The level of verification* Y^{**} *is non-decreasing in* v_0 *in the strong protection region, decreasing in* v_0 *in the random protection region, and non-increasing in* v_0 *in the weak protection region.*

Hence, we find a non-monotonic effect of the value of personal information v_0 on the level of verification.

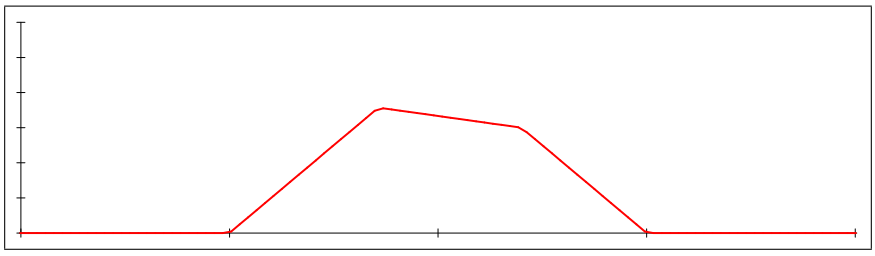

Equilibrium level of verification Y^{**} as a function of v_0.

4.3 Overlapping generations model

To investigate the issue of discrimination between returning and new consumers we consider an overlapping generation version of our model where each generation of consumers lives for two periods. Young consumers visit the website endowed with beliefs r_0 about their vulnerability. Old consumers may or may not return to the website depending on their posterior beliefs. These beliefs are based on experience as before, taking on possible values r_G, r_B or r_N, corresponding to events G, B and N.

When the website can discriminate between returning and new consumers the equilibrium is the one described in the baseline model, where X^* is the level of precaution for new consumers and there is no privacy for old consumers.

Let us assume now that the website cannot discriminate between new consumers and a (sufficiently small) share β of returning consumers.[15] The reason for this inability can be due to the possibility that some consumers remove cookies or otherwise conceal their identity when returning to the website. Let us also assume for simplicity that $v_1 = v_0$.

In each period, the website applies the same policy X to the new consumers and the non-identifiable returning consumers, while it provides no privacy protection to identifiable returning consumers. We consider equilibria in which the website has no incentive to deviate from a stationary policy in which it provides a level of protection $X^*(\beta)$ to new and non-identifiable returning consumers and no privacy protection to identifiable returning consumers. Under such a policy, the value of a non-identifiable returning consumer is

$$V_1(X^*(\beta)) = \delta^F \{a + (1 - X^*(\beta)) v_0\}$$

while the value of an identifiable returning consumer is

$$V_1(0) = \delta^F (a + v_0)$$

A consumer's direct utility u from visiting the website is distributed according to the density function $g(u)$, and the consumer's expected net benefit from third party interaction if she returns to the website is $(1 - X^*(\beta)) M(r_1)$ if the consumer is not identifiable by the website, and $M(r_1)$ if the consumer is identifiable by the website. A non-identifiable old consumer with posterior beliefs r_1 about vulnerability, and who expects $X^*(\beta)$, returns to

[15] Assuming that β is sufficiently small ensures that the website's objective function is concave.

the website with probability

$$Q\left(r_1, X^*\left(\beta\right)\right) = \Pr\left\{u + \left(1 - X^*\left(\beta\right)\right) M\left(r_1\right) > 0\right\} = 1 - G(-\left(1 - X^*\left(\beta\right)\right) M\left(r_1\right))$$

while an identifiable old consumer returns to the website with probability

$$Q\left(r_1, 0\right) = 1 - G(-M\left(r_1\right))$$

The fraction of non-identifiable old consumers who return to the website depends on the distribution of beliefs, which depends on the actual policy used by the website in any period, X, as well as on consumer beliefs in response to no-intrusion, r_N, and the website's stationary policy $X^*\left(\beta\right)$ that consumers expect. If the website departs from $X^*\left(\beta\right)$, and instead chooses X in some period, the retention rate of non-identifiable (resp. identifiable) consumers is $R\left(r_N, X, X^*\left(\beta\right)\right)$ (resp. $R\left(r_N, X, 0\right)$) where

$$R\left(r_N, X, X'\right) = p_G\left(X\right) Q\left(r_G, X'\right) + p_B\left(X\right) Q\left(r_B, X'\right) + p_N\left(X\right) Q\left(r_N, X'\right).$$

In equilibrium, $X = X^*\left(\beta\right)$, so the equilibrium retention rate of non-identifiable consumers is $R\left(r_N, X^*\left(\beta\right), X^*\left(\beta\right)\right)$ and the equilibrium retention rate of identifiable consumers is $R\left(r_N, X^*\left(\beta\right), 0\right)$, where consumers' equilibrium posterior belief in the absence of intrusion is $r_N = \phi(X^*)$.

By Bellman's Principle of Optimality, we only need to consider an isolated one-period deviation in order to confirm a stationary equilibrium. For given consumer beliefs r_N and an expected stationary policy $X^*\left(\beta\right)$, the website's expected marginal net return to privacy (i.e., the expected net benefit from a marginal increase in X) at any date is

$$\begin{aligned} H\left(r_N, X^*\left(\beta\right), \beta\right) \equiv\ & \left(1 - \beta\right)\left[\Delta_B\left(r_N, 0\right) - \Delta_G\left(r_N, 0\right)\right] \\ & + \beta\left[\Delta_B\left(r_N, X^*\left(\beta\right)\right) - \Delta_G\left(r_N, X^*\left(\beta\right)\right) - R\left(r_N, X^*\left(\beta\right), X^*\left(\beta\right)\right) v_0\right] \end{aligned}$$

where

$$\Delta_G\left(r_N, X\right) \equiv \lambda\left[Q\left(r_G, X\right) - Q\left(r_N, X\right)\right] V_1\left(X\right) + \lambda v_0$$

and

$$\Delta_B\left(r_N, X\right) \equiv \bar{\theta}\left(r_0\right)\left[Q\left(r_N, X\right) - Q\left(r_B, X\right)\right] V_1\left(X\right) - \left(1 - \lambda\right) v_0$$

A marginal increase in X increases the website's benefit from privacy protection by $\Delta_B\left(r_N, 0\right) - \Delta_G\left(r_N, 0\right)$ for each identifiable new consumer (as in our baseline model) and

by $\Delta_B(r_N, X^*(\beta)) - \Delta_G(r_N, X^*(\beta))$ for each non-identifiable new consumer. Moreover, because the increase in X applies not only to new consumers but also to non-identifiable old consumers, it leads to a loss v_0 for each non-identifiable old consumer.

We now derive a set of results regarding the way the share of non-identifiable returning consumers affects privacy protection.

Proposition 11 *If β is sufficiently small and $0 \leq \beta' < \beta$ then:*

(i) a no-privacy equilibrium exists in the model with a share β of non-identifiable consumers whenever a no-privacy equilibrium exists in the model with a share β' of non-identifiable consumers;

(ii) a full privacy equilibrium exists in the model with a share β' of non-identifiable consumers whenever a full privacy equilibrium exists in the model with a share β of non-identifiable consumers;

(iii) if $\frac{\partial^2 Q}{\partial r_N \partial X}$ is not too positive, a random privacy stationary equilibrium of the model with a share β of non-identifiable consumers provides less privacy to young consumers than in a random privacy equilibrium of the model with a share β' of non-identifiable consumers.

The condition in part (iii), that greater precaution does not too much increase the sensitivity of demand to beliefs, restricts the curvature of $G(u)$. We are now in position to state the main message of this section.

Corollary 2 *If $\frac{\partial^2 Q}{\partial r_N \partial X}$ is not too positive then a website's inability to discriminate between new consumers and a sufficiently small share of returning consumers leads to less privacy for young consumers. Moreover, this effect is stronger the larger the share of non-identifiable consumers.*

5 Conclusion

Imperfect information creates incentives for a website to protect consumer privacy. Our model demonstrates this in a novel way by assuming that consumers who visit a website learn from experience about their vulnerability to intrusions due to the website selling personal information to third parties, and that consumers who become pessimistic about their vulnerability are less like to return to the website. In response, the website exercises precaution in selling personal information to third parties and verifies third party use of custmer information, in order to profit from better consumer retention. Such a mechanism for privacy protection is tantamount to a signal-jamming theory of product quality.

Our analysis shows how a website's incentive for privacy protection improves with the value of consumer retention relative to the revenue from selling personal information, the sensitivity of consumer retention to consumer beliefs about vulnerability, and the sensitivity of consumer beliefs to experience. Greater privacy protection, however, is a mixed blessing for consumers, who, on the one hand, are better protected from intrusions, but, on the other hand, may be deprived of positive matches with third parties. Consequently, it is difficult for authorities to regulate privacy protection in a way that reliably improves consumer welfare. For example, policies that tax of information sales, improve the transparency of privacy policies, and give consumers more control over their personal information, all have mixed effects on consumer welfare.

There are various interesting directions for further research. One is to assume that consumers have some ability to protect themselves by concealing their identities when returning to a website, e.g. by removing cookies. Another, is to allow websites to charge a subscription fee for continued assess, possibly enabling them to better control their own incentives for privacy protection. Finally, studying alternative models of multi-homing by consumers and competition between websites may yield richer insights.

6 Appendix

Proof of Proposition 1

The pure-strategy equilibria are described above. We have an equilibrium with a random sale probability $X \in (0,1)$ if and only if the website is indifferent between selling and not selling personal information. This happens if and only if

$$Q(r_N) V_1 = v_0 + \left[\lambda Q(r_G) + \bar{\theta}(r_0) Q(r_B) + (1 - \lambda - \bar{\theta}(r_0) Q(r_N)\right] V_1$$

with $r_N = \phi(X)$, which is the same as condition (6). Since $\phi(X)$ is decreasing in X, the latter can only hold when $\psi^n < \frac{v_0}{V_1} < \psi^f$. Conversely, whenever this double inequality holds, there exists a unique $X \in (0,1)$ such that (6) is satisfied.

Proof of Proposition 2

In a context where the second-period precaution level need not be equal to zero, the website's profit function writes

$$\begin{aligned}\hat{\Pi}(r_N, X, X_1) &= (1 - X) v_0 + \\ &\quad \left\{p_G(X) \hat{Q}(r_0, X_1) + p_B(X) \hat{Q}(r_B, X_1) + p_N(X) \hat{Q}(r_N, X_1)\right\} \hat{V}_1(X_1)\end{aligned}$$

where

$$\hat{Q}\left(r_1, X_1\right) = 1 - G\left(-\left(1 - X_1\right) M(r_1)\right)$$

and

$$\hat{V}_1\left(X_1\right) = \delta^F\left[a + \left(1 - X_1\right) v_1\right]$$

First, note that if the second-period expected matching value is always positive, i.e. $M(r_B) = \lambda U_G + \bar{\theta}(r_B) U_B > 0$, then $\hat{Q}\left(r_1, X_1\right)$ is decreasing in X_1 and, therefore, $\hat{\Pi}\left(r_N, X, X_1\right)$ is decreasing in X_1 for any r_N and X. This implies that a website that can commit to its second-period precaution level at the beginning of the second period will find it optimal to choose $X_1 = 0$. Therefore, the website's profit function reduces to $\hat{\Pi}\left(r_N, X, 0\right) = \Pi\left(r_N, X\right)$. As $r_N = \phi\left(X\right)$ is decreasing in X, the marginal gain of the website from increasing X is lower when it can commit to its strategy:

$$\frac{\partial \Pi\left(r_N, X\right)}{\partial X} + \frac{\partial \Pi\left(r_N, X\right)}{\partial r_N} \frac{d\phi}{\partial X} < \frac{\partial \Pi\left(r_N, X\right)}{\partial X}$$

This implies that a full privacy equilibrium exists for a smaller range of values v_0/V_1 while a no privacy equilibrium exists for a wider range. Consider now an equilibrium with no commitment featuring an interior degree of precaution $X^* \in (0, 1)$. Then for any $X > X^*$, we have

$$\Pi\left(\phi\left(X^*\right), X^*\right) > \Pi\left(\phi\left(X^*\right), X\right) > \Pi\left(\phi\left(X\right), X\right)$$

Therefore, the website chooses $X \leq X^*$. Moreover at X^* it holds that $\frac{\partial \Pi}{\partial X} + \frac{\partial \Pi}{\partial r_N} \frac{d\phi}{dX} = \frac{\partial \Pi}{\partial r_N} \frac{d\phi}{dX} < 0$, which implies that the website chooses $X < X^*$.

Let us now show that an alternative sufficient condition for a regulation mandating transparency to lead to lower first-period precaution is that $G'\left(u\right) + uG''\left(u\right) > 0$ for any u.

The above comparison of the optimal degree of precaution with commitment to the equilibrium degree of precaution with no commitment when $X_1 = 0$ extends to any exogenously given $X_1 > 0$. Therefore, a sufficient condition for the website to commit to a first-period degree of precaution lower than X^* is that the equilibrium precaution in the first-period subgame for an exogenously given X_1 be decreasing in X_1. Since $\frac{v_0}{\hat{V}_1(X_1)}$ is increasing in X_1, a sufficient condition for first-period degree of precausion to be decreasing in X_1 is that:

$$\frac{\partial}{\partial X_1} \hat{B}\left(r_N, X_1\right) < 0$$

where

$$\hat{B}\left(r_N, X_1\right) \equiv \left[\lambda + \bar{\theta}\left(r_0\right)\right] \hat{Q}\left(r_N, X_1\right) - \lambda \hat{Q}\left(r_G, X_1\right) - \bar{\theta}\left(r_0\right) \hat{Q}\left(r_B, X_1\right)$$

The latter holds if

$$\frac{\partial}{\partial X_1} \hat{Q}\left(r_N, X_1\right) < \min\left\{ \frac{\partial}{\partial X_1} \hat{Q}\left(r_G, X_1\right), \frac{\partial}{\partial X_1} \hat{Q}\left(r_B, X_1\right) \right\}$$

or if

$$\frac{\partial^2}{\partial r_1 \partial X_1} \hat{Q}\left(r_1, X_1\right) < 0.$$

Since

$$\frac{\partial}{\partial X_1} \hat{Q}\left(r_1, X_1\right) = -M\left(r_1\right) g\left(-\left(1 - X_1\right) M(r_1)\right)$$

then

$$\frac{\partial^2}{\partial r_1 \partial X_1} \hat{Q}\left(r, X_1\right) = \left[-g\left(-\left(1 - X_1\right) M(r_1)\right) + M\left(r_1\right)\left(1 - X_1\right) g'\left(-\left(1 - X_1\right) M(r_1)\right)\right] \underbrace{M'\left(r_1\right)}_{>0}$$

Thus, a sufficient condition for a regulation mandating transparency to cause the website to choose a lower first-period precaution than in the equilibrium with no commitment is that $g\left(u\right) + ug'\left(u\right) > 0$ for any u or, equivalently,

$$G'\left(u\right) + uG''\left(u\right) > 0.$$

Note that in the scenario where $M\left(r_1\right) < 0$ for any r_1 (or, equivalently, $M\left(\phi(1)\right) < 0$) a sufficient condition is that $G'\left(u\right) + uG''\left(u\right) > 0$ for any $u > 0$, which can be interpreted as $G(.)$ not being too convex over $[0, +\infty)$.

Proof of Lemma 1

Denote $\tilde{p}_B\left(\theta, X\right) \equiv \theta\left(1 - X\right)$ and $\tilde{p}_N\left(\theta, X\right) \equiv 1 - p_G\left(X\right) - \tilde{p}_B\left(\theta, X\right)$. Straightforward computations show that

$$\phi\left(X\right) = \frac{\tilde{p}_N\left(\theta_l, X\right)}{\tilde{p}_N\left(\bar{\theta}\left(r_0\right), X\right)} r_0 \Longrightarrow \tilde{p}_N\left(\bar{\theta}\left(r_0\right), X\right) \frac{d\phi}{dX} = \frac{\partial \tilde{p}_N\left(\theta_l, X\right)}{\partial X} r_0 - \frac{\partial \tilde{p}_N\left(\bar{\theta}\left(r_0\right), X\right)}{\partial X} r_N$$

Using

$$\frac{\partial \tilde{p}_N\left(\theta, X\right)}{\partial X} = -\frac{\partial \tilde{p}_G\left(X\right)}{\partial X} - \frac{\partial \tilde{p}_B\left(\theta, X\right)}{\partial X}$$

and

$$\frac{\partial \tilde{p}_B(\theta_l, X)}{\partial X} r_0 = \frac{\partial \tilde{p}_B\left(\bar{\theta}(r_0), X\right)}{\partial X} \frac{\theta_l}{\bar{\theta}(r_0)} r_0 = \frac{\partial \tilde{p}_B\left(\bar{\theta}(r_0), X\right)}{\partial X} r_B$$

we obtain

$$\tilde{p}_N\left(\bar{\theta}(r_0), X\right) \frac{d\phi}{dX} = \frac{\partial \tilde{p}_G}{\partial X}\left(\bar{\theta}(r_0), X\right)(r_N - r_G) + \frac{\partial \tilde{p}_B}{\partial X}\left(\bar{\theta}(r_0), X\right)(r_N - r_G),$$

which can be written as

$$p_N(X)\frac{d\phi}{dX} = \frac{dp_G}{dX}(r_N - r_G) + \frac{dp_B}{dX}(r_N - r_G)$$

because $\tilde{p}_N\left(\bar{\theta}(r_0), X\right) = p_N(X)$ and $\tilde{p}_B\left(\bar{\theta}(r_0), X\right) = p_B(X)$. This yields the result.

Proof of Proposition 4

From the equilibrium analysis, the equilibrium level of X is a non-increasing function of the ratio

$$\frac{(1-\tau)v_0}{\delta^F\left[a + (1-\tau)v_1\right]}$$

Since

$$\frac{\partial}{\partial \tau}\left[\frac{(1-\tau)v_0}{\delta^F\left[a + (1-\tau)v_1\right]}\right] = -\frac{v_0}{V_1 - \tau v_1} + \frac{(1-\tau)v_0 v_1}{(V_1 - \tau v_1)^2} = \frac{-a v_0}{\delta^F\left[a + (1-\tau)v_1\right]^2} < 0$$

the equilibrium level of X is non-decreasing in τ.

Proof of Lemma 2

Reducing $Q(r_B)V_1$ by ε raises uniformly $B(r_N)$ by $\hat{\varepsilon} = \bar{\theta}(r_0)\varepsilon$. We distinguish three cases:

- When $X^* = 1$ we have $B(r_N) \geq v_0$. This remains true as $B(r_N)$ shifts upward.

- When $0 < X^* < 1$ then $B(r_N) = v_0$. An upward shift of $B(r_N)$ reduces the equilibrium posterior beliefs r_N^*. Since ϕ is decreasing in X this leads to an increase in X^*.

- When $X^* = 0$ then $B(r_N) < v_0$. Then an upward shift of $B(r_N)$ results in a value of r_N^* which is lower than, or the same as, before. This implies that X^* remains the same as before or increases.

Proof of Lemma 3

If all other websites announce $p = 0$, then the payoff of a website is constant equal to V_a so that $F(0) = 1$ is a best reply.

Proof of Proposition 6

An equilibrium with full privacy ($x = X = 1$) induces $r_N = \phi(1)$ and exists if and only if

$$v_0 \leq B(r_N) V_K$$

hence the condition.

Consider now an equilibrium with $0 < x < 1$ and thus $X = (1 - F(v_0))^K$, $r_N = \phi(X)$. For any $p \leq v_0$ on the support of the equilibrium strategy we must have

$$p(1 - F(p))^{K-1} = (1 - F(v_0))^{K-1} B(\phi(X)) V_K$$

leading to a cdf

$$F(p) = 1 - (1 - F(v_0)) \left(\frac{B(\phi(X)) V_K}{p} \right)^{\frac{1}{K-1}} \text{ on an interval } [p_{\inf}, p_{\max}].$$

Notice that there cannot be a mass point because it could be undercut profitably. Moreover, we must have $p_{\max} = v_0$ because otherwise setting $p = v_0$ would strictly dominate $p_{\max}$. Thus, we have

$$F(v_0) = 1 - (1 - F(v_0)) \left(\frac{B(\phi(X)) V_K}{p} \right)^{\frac{1}{K-1}}$$

implying that

$$v_0 = B(\phi(X)) V_K.$$

Given that $B(\phi(X))$ is decreasing, the value of X and thus the equilibrium exists and is uniquely defined when $B(\phi(0)) V_K \geq v_0 \geq B(\phi(1)) V_K$.

Then we have

$$F(p) = 1 - X^{\frac{1}{K}} \left(\frac{v_0}{p} \right)^{\frac{1}{K-1}} \text{ on an interval } [p_{\inf}, v_0],$$

which gives

$$p_{\inf} = X^{\frac{K-1}{K}} v_0.$$

Proof of Corollary 1

The first part follows from the equilibrium conditions and $B(\phi(X)) V_K \leq B(\phi(X)) V_1$.

Let V_K be the retention value in a mixed strategy equilibrium, then we have

$$v_0 = B\left(\phi\left(X\right)\right) V_K$$

and differentiating with respect to V

$$\delta^F \left(\lambda + \bar{\theta}\left(r_0\right)\right) Q'\left(r_N\right) V_K \frac{\partial r_N}{\partial V_K} = -\frac{v_0}{V_K}$$

The payoff is

$$L_K + X^{K-1} v_0.$$

$$\begin{aligned} L_K &= \delta^F \left(\left(1 - \lambda - \bar{\theta}\left(r_0\right)\right) Q\left(r_N\right) + \bar{\theta}\left(r_0\right) Q\left(r_B\right) + \lambda Q\left(r_G\right)\right) V_K \\ &\Longrightarrow \quad L_K = \delta^F Q\left(r_N\right) V_K - v_0 \end{aligned}$$

Differentiating with respect to V_k leads to

$$\frac{\partial}{\partial V_K}\left(L_K + X^{K-1} v_0\right) = -\frac{v_0}{V_K} \frac{1}{\lambda + \bar{\theta}\left(r_0\right)} + \delta^F Q\left(r_N\right) + \left(K-1\right) X^{K-2} \frac{\partial X}{\partial V_K} v_0$$

Using $\frac{\partial X}{\partial V_K} = \frac{1}{\phi_X} \frac{\partial r_N}{\partial V_K} = -\left(\frac{1}{\delta^F V_K \left(\lambda + \bar{\theta}(r_0)\right) Q'(r_N) \phi_X}\right) \frac{v_0}{V_K}$, we obtain

$$\frac{\partial}{\partial V_K}\left(L_K + X^{K-1} v_0\right) = -\frac{v_0}{V_K} \frac{1}{\lambda + \bar{\theta}\left(r_0\right)} \left(1 + \frac{v_0}{\delta^F V_K} \frac{\left(K-1\right) X^{K-2}}{Q'\left(r_N\right) \phi_X}\right) + \delta^F Q\left(r_N\right)$$

This is negative for $v_0/V_K > \lambda + \bar{\theta}\left(r_0\right)$ and X is small, in which case the profit is higher with multiple websites. This is positive if v_0/V_K is small and Q is large, in which case the profit is lower with multiple websites.

Proof of Lemma 4

i) $\Phi\left(X, Y\right)$ is a rational fraction in X.

$$\Phi\left(X, Y\right) = \frac{1 - Y\lambda - \left(1 - Y\right)\left(\lambda + \theta_l\right)\right)\left(1 - X\right)}{1 - Y\lambda - \left(1 - Y\right)\left(\lambda + \bar{\theta}\left(r_0\right)\right)\left(1 - X\right)}$$

Therefore, it is monotonic in X (for a given Y) and the sign of $\frac{\partial \Phi}{\partial X}$ is the same as the sign

of the determinant

$$\begin{vmatrix} [\lambda+\theta_l](1-Y) & 1-Y\lambda-(1-Y)(\lambda+\theta_l))(1-X) \\ \left[\lambda+\bar{\theta}(r_0)\right](1-Y) & 1-Y\lambda-(1-Y)\left(\lambda+\bar{\theta}(r_0)\right)\right)(1-X) \end{vmatrix}$$

which is

$$\left(\theta_l-\bar{\theta}(r_0)\right)(1-Y\lambda)(1-Y)<0$$

Therefore, $\Phi(X,Y)$ is decreasing in X.

ii) $\Phi(X,Y)$ is a rational fraction in Y. Therefore, it is monotonic in Y (for a given X) and the sign of $\frac{\partial\Phi}{\partial Y}$ is the same as the sign of the determinant

$$\begin{vmatrix} \theta_l(1-X)-\lambda X & 1-Y\lambda-(1-Y)\lambda(1-X)-(1-Y)\theta_l(1-X) \\ \bar{\theta}(r_0)(1-X)-\lambda X & 1-Y\lambda-(1-Y)\lambda(1-X)-(1-Y)\bar{\theta}(r_0)(1-X) \end{vmatrix}$$

Straightforward computations show that the latter is equal to

$$(1-X)\left(\bar{\theta}(r_0)-\theta_l\right)(-1+\lambda).$$

which is negative for any $X<1$, and equal to zero for $X=1$.

Thus, $\Phi(X,Y)$ is decreasing in Y for any $X<1$, and constant in Y for $X=1$.

Proof of Proposition 7

An equilibrium verifies $r_N^{**}=\Phi\left(X^{br}(r_N^{**}),Y^{br}(r_N^{**})\right)$. We then have:

i) For $r_N<r^M$, $X^{br}(r_N)=0$ and $C_i'(Y^{br}(r_N))=\max\{\Delta_B(r_N^{**}),0\}$ is non-decreasing, implying that $\Phi\left(X^{br}(r_N),Y^{br}(r_N)\right)$ is non-increasing

ii) For $r_N=r^M$, $X^{br}(r_N)\in[0,1]$ and $C_i'(Y^{br}(r_N))=\max\left\{\Delta^M,0\right\}$

iii) For $r_N>r^M$, $X^{br}(r_N)=1$ and $C_i'(Y^{br}(r_N))=\max\{\Delta_G(r_N^*),0\}$ is non-increasing, implying that $\Phi\left(X^{br}(r_N),Y^{br}(r_N)\right)$ is non-increasing

Hence, $\Phi\left(X^{br}(r_N),Y^{br}(r_N)\right)$ is a non-increasing continuous correspondence from $[0,1]$ into itself. This implies that it has a unique fixed point $r_N^*=\Phi\left(X^{br}(r_N^{**}),Y^{br}(r_N^{**})\right)$.

Moreover, the graph of the correspondence is continuous in v_0, implying that r_N^{**} is continuous in v_0.

Proof of Lemma 5

If $\Delta^M\leq 0$, then $Y^{**}=0$ is constant and the results follows from proposition 1. Assume that $\Delta^M>0$.

Suppose that $X^{**}=0$ and $\Delta_G(r_N^{**})>\Delta_B(r_N^{**})$, and let v_0 increase. Then X^{**} is

constant and r_N^{**} cannot decrease because then $C_i'(Y^{**}) = \max\{\Delta_B(r_N^{**}), 0\}$ would not increase and $\Phi(0, Y)$ is decreasing in Y.

Suppose that $X^{**} = 1$, then X^{**} is constant so is r_N^{**} because $\Phi(1, Y)$ is constant in Y.

Suppose $0 < X^{**} < 1$ and $C_i'(Y^{**}) = \Delta^M \geq 0$, and let v_0 increase. We have $r_N = r^M$ and

$$\frac{dr^M}{dv_0} = \frac{1}{\left(\lambda + \bar{\theta}(r_0)\right) Q'(r^M) V_1} > 0.$$

Moreover as $r^M = \Phi(X^{**}, Y^{**})$, we have

$$\frac{dr_M}{dv_0} - \frac{\partial \Phi}{\partial Y}\frac{dY^{**}}{dv_0} = \frac{\partial \Phi}{\partial X}\frac{dX^{**}}{dv_0}$$

implying that the level of precaution X^{**} decreases if

$$\frac{dr_M}{dv_0} > \frac{\partial \Phi}{\partial Y}\frac{dY^{**}}{dv_0}.$$

Using

$$C_i'(Y^{**}) = (Q(r_G) - Q(r_M)) V_1 + \lambda v_0$$

we have

$$C_i''(Y^{**})\frac{dY^{**}}{dv_0} = \lambda\left(1 - Q'(r_M)\frac{dr_M}{dv_0}V_1\right),$$

so that this condition writes as

$$\frac{1}{\left(\lambda + \bar{\theta}(r_0)\right) Q'(r^M) V_1} > \left(1 - \frac{Q'(r_M) V_1}{\left(\lambda + \bar{\theta}(r_0)\right) Q'(r^M) V_1}\right)\frac{\lambda}{C_i''}\frac{\partial \Phi}{\partial Y}$$

and gives condition $C1$.

Proof of Proposition 9

Consider the case where $X^* = 1$ in the baseline model without verification. Then, there is no verification if $\Delta_G \leq 0$ or

$$\frac{v_0}{V_1} \leq \underline{\psi}^i = Q(\Phi(1, 0)) - Q(r_G).$$

This is an equilibrium if in addition $\frac{v_0}{V_1} \leq \psi^f$. Notice that $Q(\Phi(1,0)) - Q(r_G) \geq \psi^f$ if and only if

$$\left(1 - \lambda - \bar{\theta}(r_0)\right) Q(\Phi(1,0)) - (1 - \lambda) Q(r_G) + \bar{\theta}(r_0) Q(r_B) \geq 0$$

Similarly if $X^* = 0$ in the baseline model, there is no verification if $\Delta_B \leq 0$ or if

$$\frac{v_0}{V_1} \geq \bar{\psi}^i = \frac{\bar{\theta}(r_0)}{1-\lambda}\left(Q\left(\Phi\left(0,0\right)\right) - Q\left(r_B\right)\right).$$

We have $\frac{\bar{\theta}(r_0)}{1-\lambda}\left(Q\left(\Phi\left(0,0\right)\right) - Q\left(r_B\right)\right) \leq \psi^n$ if

$$\left(1-\lambda-\bar{\theta}\left(r_0\right)\right)Q\left(\Phi(0,0)\right) - \left(1-\lambda\right)Q\left(r_G\right) + \bar{\theta}\left(r_0\right)Q\left(r_B\right) \geq 0$$

Now suppose that X^* is interior in the baseline model. Then

$$\Delta^M = \lambda\left(Q\left(r_G\right) - Q\left(r_N\right)\right)V_1 + \lambda v_0 = \bar{\theta}\left(r_0\right)\left(Q\left(r_N\right) - Q\left(r_B\right)\right)V_1 - \left(1-\lambda\right)v_0$$

implies that

$$v_0 = \left(\lambda + \bar{\theta}\left(r_0\right)\right)Q\left(r^M\right) - \lambda Q\left(r_G\right) - \bar{\theta}\left(r_0\right)Q\left(r_B\right)V_1$$

Thus, we have $\Delta^M \leq 0$ if and only if

$$\left(1-\lambda-\bar{\theta}\left(r_0\right)\right)Q\left(\Phi(X^*,0)\right) - \left(1-\lambda\right)Q\left(r_G\right) + \bar{\theta}\left(r_0\right)Q\left(r_B\right) \geq 0$$

Notice that the left-hand-side decreases in X^0. Therefore we can distinguish three cases:

1- If $Q\left(\Phi(1,0)\right) \geq \frac{(1-\lambda)Q(r_G)-\bar{\theta}(r_0)Q(r_B)}{1-\lambda-\bar{\theta}(r_0)}$ then this holds for all X and thus $Y = 0$ for all v_0.

2- If $Q\left(\Phi(0,0)\right) \leq \frac{(1-\lambda)Q(r_G)-\bar{\theta}(r_0)Q(r_B)}{1-\lambda-\bar{\theta}(r_0)}$ then the condition holds for no X and verification occurs for $\underline{\psi}^i < \frac{v_0}{V_1} < \bar{\psi}^i = \frac{\bar{\theta}(r_0)}{1-\lambda}\left(Q\left(\Phi\left(0,0\right)\right) - Q\left(r_B\right)\right)$ with $\underline{\psi}^i < \psi^f < \psi^n \leq \bar{\psi}^i$.

3- If $Q\left(\Phi(1,0)\right) < \frac{(1-\lambda)Q(r_G)-\bar{\theta}(r_0)Q(r_B)}{1-\lambda-\bar{\theta}(r_0)} < Q\left(\Phi(0,0)\right)$, there exists a critical value $\bar{X}^i > 0$ such that

$$Q\left(\Phi(\bar{X}^i,0)\right) = \frac{\left(1-\lambda\right)Q\left(r_G\right) - \bar{\theta}\left(r_0\right)Q\left(r_B\right)}{1-\lambda-\bar{\theta}\left(r_0\right)}.$$

and verification occurs for $\bar{X}^i < X^*$. We then define

$$\bar{\psi}^i = \left(\lambda + \bar{\theta}\left(r_0\right)\right)Q\left(\Phi(\bar{X}^i,0)\right) - \lambda Q\left(r_G\right) - \bar{\theta}\left(r_0\right)Q\left(r_B\right)$$

Verification occurs for $\underline{\psi}^i < \frac{v_0}{V_1} < \bar{\psi}^i$ and $\underline{\psi}^i < \psi^f < \bar{\psi}^i < \psi^n$.

Proof of Proposition 10

$C_i'\left(Y^{**}\right) = \max\left\{\lambda\left(Q\left(r_G\right) - Q\left(\Phi\left(1,Y^{**}\right)\right)\right) + \lambda v_0, 0\right\}$ in the strong protection regime which implies that Y^* is non-decreasing in v_0. Since $C_i'\left(Y^{**}\right) = \max\left\{\bar{\theta}\left(r_0\right)\left(Q\left(\Phi\left(0,Y^{**}\right)\right) - Q\left(r_B\right)\right) - (1-\right.$

in the weak protection regime, it follows that Y^*non-increasing in v_0.

In the random privacy region, we have

$$C_i'\left(Y^{**}\right)=\lambda\left(v_0+Q\left(r_G\right)-Q(r^M)\right)=\bar{\theta}\left(r_0\right)Q(r^M)-\bar{\theta}\left(r_0\right)Q(r_B)-\left(1-\lambda\right)v_0$$

which yields

$$Q(r^M)=\frac{v_0+\lambda Q\left(r_G\right)+\bar{\theta}\left(r_0\right)Q(r_B)}{\lambda+\bar{\theta}\left(r_0\right)}$$

and, therefore,

$$C_i'\left(Y^{**}\right)=\lambda\left(\frac{\left(\lambda+\bar{\theta}\left(r_0\right)-1\right)v_0+\bar{\theta}\left(r_0\right)Q\left(r_G\right)-\bar{\theta}\left(r_0\right)Q(r_B)}{\lambda+\bar{\theta}\left(r_0\right)}\right)$$

Hence, Y^{**} is decreasing in v_0 in the random protection region.

Proof of Proposition 11

(i) If β is sufficiently small, then the objective function of a website considering a deviation is concave, and a no-privacy equilibrium exists if and only if

$$H\left(\phi\left(0\right),0,\beta\right)\leq 0.$$

From

$$\frac{\partial H}{\partial\beta}\left(\phi\left(0\right),0,\beta\right)=-R\left(\phi\left(0\right),0,0\right)v_0<0$$

it follows that, for $\beta'<\beta$,

$$H\left(\phi\left(0\right),0,\beta'\right)>H\left(\phi\left(0\right),0,\beta\right)$$

and, therefore, that $H\left(\phi\left(0\right),0,\beta\right)<0$ whenever $H\left(\phi\left(0\right),0,\beta'\right)<0$.

(ii) A full privacy equilibrium exists only if

$$H\left(\phi\left(1\right),1,\beta\right)\geq 0$$

From

$$\begin{aligned}H\left(\phi\left(1\right),1,\beta\right) &= \left(1-\beta\right)\left[\Delta_B\left(\phi\left(1\right),0\right)-\Delta_G\left(\phi\left(1\right),0\right)\right]\\ &\quad+\beta\left[\Delta_B\left(\phi\left(1\right),1\right)-\Delta_G\left(\phi\left(1\right),1\right)-R\left(\phi\left(1\right),1,1\right)v_0\right]\end{aligned}$$

it follows that

$$\frac{\partial H}{\partial \beta}\left(\phi\left(1\right),1,\beta\right)=-\Delta_B\left(\phi\left(1\right),0\right)+\Delta_G\left(\phi\left(1\right),0\right)+\Delta_B\left(\phi\left(1\right),1\right)-\Delta_G\left(\phi\left(1\right),1\right)-R\left(\phi\left(1\right),1,1\right)v_0$$

Since

$$Q\left(r_N,1\right)=Q\left(r_G,1\right)=Q\left(r_B,1\right)$$

we have

$$\Delta_B\left(\phi\left(1\right),1\right)-\Delta_G\left(\phi\left(1\right),1\right)=-v_0$$

Therefore,

$$\frac{\partial H}{\partial \beta}\left(\phi\left(1\right),1,\beta\right)=-\left\{\Delta_B\left(\phi\left(1\right),0\right)-\Delta_G\left(\phi\left(1\right),0\right)+v_0+R\left(\phi\left(1\right),1,1\right)v_0\right\}$$

which is negative because

$$\Delta_B\left(\phi\left(1\right),0\right)-\Delta_G\left(\phi\left(1\right),0\right)+v_0=\bar{\theta}\left(r_0\right)\underbrace{\left[Q\left(\phi\left(1\right),1\right)-Q\left(r_B,0\right)\right]}_{>0}V_1\left(0\right)-\lambda\underbrace{\left[Q\left(r_G,0\right)-Q\left(\phi\left(1\right),0\right)\right]}_{=0}V_1\left(0\right)$$

is positive. Thus,

$$H\left(\phi\left(1\right),1,\beta'\right)\leq H\left(\phi\left(1\right),1,\beta\right),$$

which implies that

$$H\left(\phi\left(1\right),1,\beta'\right)\geq 0\Longrightarrow H\left(\phi\left(1\right),1,\beta\right)\geq 0.$$

(iii) Let us compare $X^*\left(\beta\right)$ and $X^*\left(\beta'\right)$.

Differentiating

$$H\left(\phi\left(X^*\left(\beta\right)\right),X^*\left(\beta\right),\beta\right)=0$$

with respect to β at $\beta=0$ yields

$$\left.\frac{dX^*}{d\beta}\right|_{\beta=0}\underbrace{\left[\lambda+\left(1-\lambda\right)\bar{\theta}\left(r_0\right)\right]\frac{\partial Q}{\partial r_N}\left(\phi\left(X^*\left(0\right)\right),0\right)\phi'\left(X^*\left(0\right)\right)}_{<0}$$

$$\begin{aligned}=\;&\left\{\lambda\left[Q\left(\phi\left(X^*\left(0\right)\right),0\right)-Q\left(r_G,0\right)\right]+\left(1-\lambda\right)\bar{\theta}\left(r_0\right)\left[Q\left(\phi\left(X^*\left(0\right)\right),0\right)-Q\left(r_B,0\right)\right]\right\}V_1(0)\\&-\left\{\lambda\left[Q\left(\phi\left(X^*\left(0\right)\right),X^*\left(0\right)\right)-Q\left(r_G,X^*\left(0\right)\right)\right]+\left(1-\lambda\right)\bar{\theta}\left(r_0\right)\left[Q\left(\phi\left(X^*\left(0\right)\right),X^*\left(0\right)\right)-Q\left(r_B,X^*\left(0\right)\right)\right.\right.\\&+R\left(\phi\left(X^*\left(0\right)\right),X^*\left(0\right),X^*\left(0\right)\right)\end{aligned}$$

Assume that $\frac{\partial^2 Q}{\partial r_N \partial X} \leq 0$. Then

$$Q\left(\phi\left(X^*\left(0\right)\right), X^*\left(0\right)\right) - Q\left(r_G, X^*\left(0\right)\right) \leq Q\left(\phi\left(X^*\left(0\right)\right), 0\right) - Q\left(r_G, 0\right)$$

and

$$Q\left(\phi\left(X^*\left(0\right)\right), X^*\left(0\right)\right) - Q\left(r_B, X^*\left(0\right)\right) \leq Q\left(\phi\left(X^*\left(0\right)\right), 0\right) - Q\left(r_B, 0\right)$$

This, combined with

$$V_1(X^*\left(0\right)) \leq V_1(0),$$

leads to

$$\left.\frac{dX^*}{d\beta}\right|_{\beta=0} \leq \frac{R\left(\phi\left(X^*\left(0\right)\right), X^*\left(0\right), X^*\left(0\right)\right)}{\left[\lambda + \left(1-\lambda\right)\bar{\theta}\left(r_0\right)\right]\frac{\partial Q}{\partial r_N}\left(\phi\left(X^*\left(0\right)\right), 0\right)\phi'\left(X^*\left(0\right)\right)} < 0$$

By continuity, $\left.\frac{dX^*}{d\beta}\right|_{\beta=0}$ remains negative if $\frac{\partial^2 Q}{\partial r_N \partial X}$ is not too positive.

References

[1] Acquisti, A, Taylor, C., and L. Wagman (2016),"The Economics of Privacy," *Journal of Economic Literature* 54 (2), 442-492.

[2] Bagwell, K. and A. Wolinsky (2002), "Game Theory and Industrial Organization," *Handbook of Game Theory with Economic Applications*, 3, 1851-1895.

[3] Bloch, F., and G. Demange (2017), "Taxation and Privacy Protection on Internet Platforms," *Journal of Public Economic Theory*, forthcoming.

[4] Board, S., and M. Meyer-Ter-Vehn (2013), "Reputation for Quality", *Econometrica*, 81(6), 2381-2462.

[5] Federal Trade Commission (2009), "Self-Regulatory Principles for Online Behavioral Advertising," FTC Staff Report.

[6] Judd, K. L. and M. H. Riordan (1994), "Price and Quality in a New Product Monopoly," *The Review of Economic Studies*, 61(4), 773-789.

[7] Kim, W., Jeong, O-R., Kim, C., and J. So (2011), "The Dark Side of the Internet: Attacks, Costs, and Responses," *Information Systems*, 36, 675-705.

[8] O'Brien, D.P., and D. Smith (2014), "Privacy in Online Markets: A Welfare Analysis of Demand Rotations," *FTC Bureau of Economics Working Paper*.