WORKING PAPERS

N° TSE - 947

August 2018

“Privacy Protection and Consumer Retention"

Bruno Jullien, Yassine Lefouili, and Michael H. Riordan

Toulouse School of Economics

Privacy Protection and Consumer Retention[*]

Bruno Jullien[†] Yassine Lefouili[‡] Michael H. Riordan[§]

August 8, 2018

Abstract

A website monetizes information it collects about its customers by charging third parties for targeted access to them. Allowing for third parties who are well-intentioned, a nuisance, or even malicious, the resulting consumer experiences might be good, bad, or neutral. As consumers learn from experience, the website especially risks losing those customers who suffer a bad experience. Customer retention thus motivates the website to be cautious about monetization, or to spend resources to screen third parties. We study the website's equilibrium privacy policy, its welfare properties, competition in the market for information, and the elusiveness of reliable welfare-improving regulations.

Keywords: Privacy Policy, Consumer Retention, Personal Data, Regulation.

JEL Classification: D83, L15, L51.

[*] We are grateful to Simon Anderson, Dirk Bergemann, Yeon-Koo Che, Jacques Crémer, Gabrielle Demange, Hanna Halaburda, Christian Hellwig, Justin Johnson, Martin Peitz, Georgios Petropoulos, Mike Powell, Andy Skrzypacz, Roland Strausz, Hal Varian, Alex White, and our audiences at the MACCI Summer Institute (Schloss Gracht), the TSE Digital Workshop (Toulouse), the 9th Bi-Annual TSE Postal Economics Conference (Toulouse), the Workshop on the Digital Economy (Paris), 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), the UIBE Workhop on IO and Competition Policy (Beijing), ESSET 2018 (Gerzensee), and seminars at the Paris School of Economics, CREST, the University of Melbourne, the University of Los Andes, the University of Montpellier, and the University of Tel Aviv 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. Bruno Jullien and Yassine Lefouili are members of TSE-P (formerly IDEI) whose IT research program is or has been funded by Microsoft, Orange, Facebook and Groupe La Poste.

[†] 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

The internet has transformed how consumers shop. Sophisticated online platforms enable advertisers to target relevant ads to consumers, but, absent reliable safeguards, might also be a vehicle for cybercrime. Malicious advertising (or "malvertising") is particularly nefarious because it degrades the fundamental business model for advertising-supported websites. Successful malvertising places deceptive ads that make disingenuous offers, install malicious code (e.g. ransomware), or phish for personal information (e.g. passwords).[1] Such risks discourage online commerce, or encourage protective measures that reduce website revenue (e.g. ad blockers).

The rise of online commerce and the concurrent rise of cybercrime pose new issues for consumers, businesses, and policymakers. Consumers, while aware of rising cybercrime, nevertheless might be too optimistic about their vulnerability to malicious attacks when visiting websites. This seems especially problematic when the probability of a successful attack is small, but the potential harm is large; intuitively, someone engaging in a risky activity becomes complacent if nothing bad happens. Websites, in turn, balance an incentive to adopt privacy protection measures to protect and retain customers against revenues earned from sharing the information it gathers about them. Meanwhile, publicity about cyberattacks and data breaches force policy makers to weigh the welfare consequences of more strictly regulating website privacy policies, e.g. requiring greater transparency or consumer control. There are different viewpoints on the merits of such regulation. An optimistic perspective contends that market forces discipline firms, while the opposite view contends that binding rules are necessary and desirable to adequately protect consumers. A middle-of-the-road stance - which our analysis supports - takes the view that, while websites may have imperfect incentives to protect consumer privacy, it is difficult to design privacy rules that reliably improve consumer welfare.

There are many ways to monetize a website.[2] Some do not raise substantial privacy concerns, including (untargeted) banner advertising or direct merchandising at the website. But other ways a website might earn revenue use personal consumer information, including behavioral marketing that targets ads using information about consumers' online activities. For example, a website might sell a "lead" by connecting a consumer to another company, who is interested in making an offer to consumers who have expressed an interest in the website's content.[3] Regulators have raised privacy concerns

[1] See RiskIQ (2016) and U.S. Senate Committee on Homeland Security and Governmental Affairs (2014).

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

[3] Cost-per-action (CPA) marketing platforms appear to implement something similar to this simple

and recommended principles of greater transparency and consumer control regarding online behavioral marketing activities (FCC, 2009). The recent European Union General Data Protection Regulation (GDPR) mandates transparency and consent for the storage and processing of personal data. California recently passed a digital privacy regulation mandating greater transparency and consumer control.[4]

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 the consumer's value of website content and on the consumer's vulnerability to intrusions. Under these conditions a website's privacy policy affects consumer retention by altering consumer experience and thus consumer learning. In our model, a website offers a free service and earns revenue from banner advertising (or another activity that doesn't compromise privacy). The website also collects information about its customers that it can use to charge third parties for targeted access to them (or profit from some form of marketing that raises privacy concerns). Such transactions with third parties could be beneficial to consumers, for example, by enabling targeted advertising that informs consumers of desirable products, or intrusive, for example, by increasing exposure to spam, phishing, or malware. Those customers experiencing intrusions become more pessimistic about their overall utility from a return visit to the website, and this learning mechanism gives the website an incentive for a privacy policy that limits third-party transactions 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 the website.[5] The service is an experience good, for which the consumers have heterogeneous values in the second period. The website also sells a matching service to third parties that provides targeted access to consumers. The website's privacy policy is a choice of "precaution", determining the probability that the website does not sell the matching service to an interested third party in the first period. Consumers do not directly observe the website's choice of precaution, and instead form equilibrium beliefs. Matching with a third party results in a consumer experience that may be good, bad, or neutral; a neutral experience is the same as if there is no third-party match. Consumers are unsure of their vulnerability, defined as the probability of a bad experience. In the first period, consumers have identical prior

business model for website monetization.

[4] See e.g. New-York Times, June 28, 2018.

[5] This baseline model can be interpreted as examining a particular cohort in an overlapping generation model that crudely distinguishes between "young" and "old" consumers.

beliefs about vulnerability, and, 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 realized utility value of the website service and their posterior beliefs of vulnerability. An equilibrium is a profit-maximizing level 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. Website precaution in the first period is decreasing in the first-period value of third-party sales relative to the second-period value of retaining customers. In a full-precaution regime, this relative value is sufficiently low that the website does not offer to match third parties with its customers. Conversely, in a no-precaution regime, the relative value is sufficiently high that the firm sells the matching service to all interested third parties. There is also an intermediate partial-precaution region, in which precaution is decreasing in the relative value. The website provides no precaution in the second period, because there is no future relationship with returning consumers.[6]

The website's equilibrium incentive for precaution is at best only imperfectly aligned with consumer welfare. This is not surprising because consumers cannot verify website precaution. Consumer short-term (i.e. first period) welfare decreases with precaution if their expected utility from third-party matching ("match utility") is positive, while long-term (i.e. second period) consumer welfare also decreases with precaution due to less informative learning about vulnerability. If the website could commit to first-period precaution it would choose less precaution than the equilibrium level because of its ability to directly alter consumer beliefs. In this case, the website's incentive for less precaution in the first period appears well-aligned with consumer welfare, assuming no-precaution in the second period remains profit-maximizing. The website, however, would commit to positive second-period precaution if that sufficiently improves customer retention. In this case average consumer welfare might or might not improve, depending on the distribution of marginal consumers, who have heterogenous posterior beliefs about vulnerability.

Robust welfare-improving regulations are not readily apparent. For example, a tax on third-party matching increases precaution, but this is detrimental to consumers if match utility is positive. We also show that a transparency policy that allows the website to commit to a minimum precaution is inconsequential because the website wants to commit to less rather than more precaution. Finally we consider an opt-out rule allowing

[6] Thus the two-period model captures starkly the intuitive idea that the website has a greater incentive to protect patrons with a higher customer lifetime value (CLV), which is the profit attributed to the entire future relationship.

customers to refuse permission for the website to use their personal information for third-party matching. 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, match utility is positive.

We also study two extensions of the baseline model. These extensions add positive insights about website incentives for privacy protection, but do not reverse our normative conclusion about the difficulty of designing robust welfare-improving regulations. First, we allow for multiple websites with multi-homing consumers, and provide a novel model of competition in the market for information. We find that competition reduces the price of the matching service and that there is less precaution compared to the single website case if match utility is positive. Second, 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 level 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 third-party matching. We also show that verification and precaution are substitutes for the website.

The economics of privacy literature echoes various themes from the broader information economics literature (Acquisti, Taylor and Wagman, 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 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 match 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, consumers in equilibrium see through these incentives, and correctly predict the firm's actions.[7]

[7] 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

Our work is related to the paper by Spiegel (2013) on a software producer's choice between charging consumers for the software and offering for free a bundle of the software and ads. In his model, however, consumers are perfectly informed about the utility they derive from an impression and the firm's strategy is observable to them. In a similar vein, O'Brien and Smith (2014) investigate firms' private incentives to offer customers privacy and compare them to the socially optimal incentives. They, however, assume that sellers can commit to privacy policies while we suppose that they cannot (in the absence of privacy regulation). Moreover, there is no scope for learning in their model. Toh (2018) develops a model in which consumers learn gradually about a website's security level. She investigates the website's *ex ante* incentives to invest in security while we investigate the website's ex post incentives to sell access to customers' personal data.

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, transparency and opt-out policies. Section 4 addresses the extensions. Section 5 concludes and all proofs are relegated to the Appendix.

2 Baseline model

Consider a website offering a service to a unit-mass population of consumers for two periods: period 0 and period 1. A consumer derives a utility u from the service, but this utility is unknown at the beginning of period 0. Instead, this utility is perceived to be distributed independently in the population according to a cumulative distribution function $G(u)$ with mean u_0 and support $\mathbb{R}$. Furthermore, the mean utility is large enough that all consumers choose to participate in period 0. Each consumer learns her u upon consuming the service, and this knowledge informs her participation decision in period 1.

The service offered to the consumers is free, and the website has two ways to monetize it. First, the website obtains an exogenous revenue a per visiting consumer, e.g. from banner ads or merchandising that do not target particular consumers. Second, the website collects customer information that it uses to sell a matching service to third parties for a positive price v_t in period $t = 0, 1$.[8] This information, for example, might come from

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 signal-jamming models of product quality. Early models of signal jamming include Riordan (1985), Fudenberg and Tirole (1986), and Holmström (1999).

[8] For a website with scarce advertising space, a can be interpreted as the value of an untargeted banner

tracking consumers' online activity with a cookie, or it might be personal information, such as an email or mailing address, that consumers disclose to the website. The third party might be a data broker creating consumer profiles, or an advertiser targeting a particular consumer group. Recognizing that there are several interpretations of our model, we use the short-hand "selling information" or "information sale" to mean a transaction with a third party, using consumer information collected by the website, that involves a payment to the website.

An information sale can result in three possible consumer experiences, which impact the consumer's utility. There is a probability λ that the experience is good (G) and adds positive utility $U_G > 0$. There is also a positive probability θ that the experience is bad (B), with negative utility $U_B < 0$. In all other events, the experience is neutral (N) yielding $U_N = 0$. Hence a neutral experience occurs with probability 1 in the absence of an information sale, and with probability $1-\lambda-\theta$ if the website sells customer information. A good consumer experience might be beneficial targeted advertising, while a bad experience might come from spam, phishing, or malware.

To fix more general ideas, we consider the following specific targeted advertising scenario. A unit-mass of third-party advertisers arrive each period, and each consumer in a period is of interest to exactly one of these advertisers. In order to target its consumer of interest with an ad, the advertiser is willing to pay a fee v_t to secure the cooperation of the website. With probability λ the targeted ad benefits the consumer and results in a good experience. With probability θ the ad is a nuisance or worse, causing a disutility from a negative experience. We allow that disutility to be small, as for the case of mildly irritating spam, or large, as for the case for ransomware. In all other events, the ad is inconsequential, and the consumer has a neutral experience with no utility consequences.[9]

Critically, consumers are unsure of their preferences over third-party advertising. We model this by assuming that, while λ is a known parameter, θ is an unknown characteristic of 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 weak vulnerability, assumed to be the same for all consumers, and $\theta_0 = r_0\theta_l + (1 - r_0)\,\theta_h$ the ex ante vulnerability of the consumer.

Vulnerability can be interpreted in several ways. One interpretation is mistargeted advertising: θ is the probability of a nuisance ad that generates a small disutility for the

ad and v_t as the incremental value of a targeted ad.

[9] The technology can be interpreted as a special case of a "database" that maps consumer-advertisers pairs into a "“match value" for the advertiser, as posited in Bergemann and Bonatti (2015), assigning v_t along the diagonal and 0 elsewhere. A "query" by a third-party is a request to identify the unique consumer with match value v_t.

consumer, and $1 - \lambda - \theta$ is the probability of innocuous poorly targeted ads. Another interpretation is malvertising: θ is the probability of abusive or criminal use of customer information that imposes a large utility loss, e.g. a denial-of-service attack or identity theft. Finally, some consumers might have a poor understanding of how well they are protected against aggressive intrusions, for instance because of a superior antivirus or firewall, so that intrusion by malicious third parties is more likely to fail.

A website visit thus is an experience good. During period 0, a consumer observes u, and also experiences a good (U_G), bad (U_B), or a neutral (0) utility increment from third party ads. The consumer learns about her θ from the realized experience. At the end of period 0, a consumer knows her value of the website service, and revises her beliefs about vulnerability. We denote by r_1 the updated probability that $\theta = \theta_l$. The same situation repeats in period 1 for returning consumers, except that the price for third party ads is v_1.

In our baseline model, the privacy policy of the website determines the probability $X \in [0, 1]$ that customer information is sold.[10] For instance, the website could sell access to a database the content of which depends on its privacy policy. Each advertiser might find the database useful, or not, for the purpose of targeting customers. Thus the design of the database determines the probability that a given advertiser buys it.[11] We will refer to X as (the level of) precaution, and say we have *full precaution* when $X = 1$, *no precaution* if $X = 0$, and *partial precaution* if $0 < X < 1$.[12]

This probability is unobserved by consumers. In equilibrium, consumers update beliefs about their vulnerability using Bayes Rule and taking as given the website's privacy policy, and the website chooses a profit-maximizing policy given consumer beliefs. Equilibrium and its welfare and policy implications are analyzed next. Later, we consider a richer privacy policy in which the website also can incur a cost to verify third party use of customer information prior to its sale.

[10] More specifically, a privacy policy that restricts what consumer information a website collects and how it is used might reduce the ability of the website to match consumers and interested third parties.

[11] Alternatively, we can intepret $X = 0$ (resp. $X = 1$) as meaning that the website always (resp. never) sells customer information to a third party and interpret $X \in (0, 1)$ as a mixed strategy. Note that the latter can be "purified" by introducing into the model a vanishingly small amount of incomplete information about the value of personal information to third parties. See, for instance, Bagwell and Wolinsky (2002).

[12] Note that, in deriving equilibrium, we can fully characterize the second period with the retention value V_1 and the retention rate $Q(r)$, suggesting our model admits alternative interpretations. For example, setting $\lambda = 0$ and $v_1 = 0$, we could interpret $v_0 X$ as an investment to protect consumer data, $p_B(X)$ as the probability of a security breach, $-U_B$ as possible consumer harm, r as the consumer's beliefs about her vulnerability, and V_1 as the average profit from serving consumers. See Toh (2018) for a more elaborate model of security investment with a similarly structured equilibrium.

3 Analysis

3.1 Equilibrium

Website privacy policy and consumer beliefs are determined jointly in equilibrium. A website has an incentive for privacy protection only if the long-run value of retaining a customer (V_1) is sufficiently high relative to the short-rung gain from exposing the customer to third parties (v_0). The probability of retention ($Q(r_1)$) depends on consumers' beliefs about the utility of returning to the website, which depends on consumers' expectations of privacy. Furthermore, the distribution of beliefs in the consumer population, and therefore average retention, depends on the website's privacy policy. In equilibrium, consumers correctly anticipate privacy, and the website correctly anticipates how privacy affects retention.

Our two-period model gives specific content to this notion of equilibrium. It is immediate that the website sells information to all interested third parties at price v_1 in period 1, as there is no further interaction with the customer. Hence the value of retention is $V_1 = \delta^F (a + v_1)$ where δ^F is the firm's discount factor. The retention probability for a given posterior belief is $Q(r_1) \equiv 1 - G(-M(r_1))$, where $M(r) \equiv \lambda U_G + (r\theta_l + (1-r)\theta_h) U_B$ defines the expected benefit from third party matching when the probability of low vulnerability is r. Posterior beliefs, denoted (r_G, r_B, r_N), are the updated probabilities that $\theta = \theta_l$ at the beginning of period 1 after the events G, B, and N are observed by the consumer. For the most part, we focus on scenarios with $M(r_0) \geq 0$, meaning consumers in period 1 have no objection to third party matching. Heterogeneity of posterior beliefs, however, allows that consumers may disagree on the desirability of third party matching in period 1.

Precaution determines the distribution of consumer beliefs at the beginning of period 1. The probability of a good experience (event G) and the probability of a bad experience (event B) are respectively $p_G(X) = \lambda(1-X)$ and $p_B(X) = \theta_0(1-X)$. Both are decreasing in precaution; it follows that the probability of a neutral experience (event N), $p_N(X) = 1 - p_G(X) - p_B(X)$, is increasing in precaution. The website therefore has an incentive for full (no) precaution if and only if

$$v_0 \leq (\geq) \left[p_G(X) Q(r_G) + p_B(X) Q(r_B) + p_N(X) Q(r_N)\right] V_1.$$

Consumer beliefs at the beginning of period 1 are formed using Bayes Rule, conditioning on realized experience and taking precaution as given. Posterior beliefs after events

G and B are respectively $r_G = r_0$ and $r_B = (\theta_l/\theta_0)\, r_0$. Notice that r_G and r_B are independent of X, and therefore can be treated as parameters.[13] This is because these events occur only if information is sold, and the website's strategy does not affect the conditional probabilities of these events. Event N, however, can occur whether or not information is sold. Consequently, a consumer's posterior belief that $\theta = \theta_l$ after a neutral experience depends on anticipated precaution:

$$r_N = \phi\,(X) \equiv \frac{1-(\lambda+\theta_l)\,(1-X)}{1-(\lambda+\theta_0)\,(1-X)} r_0. \tag{1}$$

Of course, posterior and prior beliefs must be consistent:

$$p_G\,(X)\, r_G + p_N\,(X)\, r_N + p_B\,(X)\, r_B = r_0.$$

A neutral experience is good news in our model in the sense that, for all levels of precaution less than full, it yields the highest posterior belief: $r_B < r_G < r_N$. The intuition behind this result is as follows. Provided $X < 1$, the consumer reasons that a neutral experience could have resulted from the possibility that a third party did gain access to the consumer but the consumer had a neutral experience due to low vulnerability. Thus the consumer becomes more optimistic after a neutral experience. Moreover, $\phi\,(X)$ decreases in precaution because, by reducing exposure to third parties, higher precaution reduces the likelihood that a neutral experience results from low consumer vulnerability rather than from information not being sold. When there is full precaution, however, the consumer is never exposed to third parties and a neutral experience conveys no information: $\phi\,(1) = r_0$.

We are now in a position to characterize equilibrium and provide some comparative static results. The intuition behind the equilibrium is as follows. Selling customer information in period 0 yields extra revenue v_0, but raises both the probability of a good experience by λ and the probability of a bad experience by θ_0, which reduces the retention probability by $\lambda\,[Q\,(r_N) - Q\,(r_G)] + \theta_0\,[Q\,(r_N) - Q\,(r_B)]$, sacrificing future revenue proportionally. We decompose this total effect between the part related to bad experiences and the rest. Toward this end, define

$$\Delta_G\,(r_N) \equiv \lambda V_1\,[Q\,(r_G) - Q\,(r_N)] + (1-\theta_0)\, v_0,$$

as the website's gain from selling information to the third party when it generates a

[13] r_G and r_B are not defined by Bayes rule under full privacy (i.e. $X = 1$). We assume that their values remain equal to r_0 and $(\theta_l/\theta_0)\, r_0$, respectively, in this case.

neutral or a good experience. Define also

$$\Delta_B(r_N) \equiv \theta_0 V_1 [Q(r_N) - Q(r_B)] - \theta_0 v_0$$

as the gain from avoiding selling information when it induces a bad experience. The total gain from being cautious and not selling information is then $\Delta_B(r_N) - \Delta_G(r_N)$. The optimal strategy for the website is thus defined by the following "best response" correspondence:

$$X^{br}(r_N) \in \arg\max_{X \in [0,1]} X(\Delta_B(r_N) - \Delta_G(r_N)). \tag{2}$$

The best response correspondence optimizes the trade-off between avoiding bad experience and selling information. Notice that the gain from precaution decreases with $r_N \in [r_0, \phi(0)]$ and lies in the interval $[\psi^f V_1 - v_0, \psi^n V_1 - v_0]$ where

$$\psi^f \equiv \theta_0(Q(r_0) - Q(r_B)) < \psi^n \equiv (\theta_0 + \lambda) Q(\phi(0)) - \theta_0 Q(r_B) - \lambda Q(r_G). \tag{3}$$

If $\psi^f V_1 - v_0 < 0 < \psi^n V_1 - v_0$, the level of precaution $X^{br}(r_N)$ jumps from full precaution to no precaution when the posterior belief induced by neutral experience crosses a threshold $r^M \in (r_0, \phi(0))$. This threshold is defined (uniquely) as the solution of

$$\Delta_G(r^M) = \Delta_B(r^M). \tag{4}$$

At $r_N = r^M$, the website is indifferent between all levels of precaution. An equilibrium is a level of precaution X^* and a consumer belief r_N^* such that $r_N^* = \phi(X^*)$ and $X^* = X^{br}(r_N^*)$.

Proposition 1 *A unique equilibrium exists. Equilibrium precaution is a non-increasing function of the ratio v_0/V_1 and:*

(i) the website provides full precaution if $v_0/V_1 \leq \psi^f$;

(ii) the website provides no precaution if $v_0/V_1 \geq \psi^n$;

(iii) the website provides partial precaution ($r_N^* = r^M$ and $0 < X^* < 1$) if $v_0/V_1 \in (\psi^f, \psi^n)$.

This result implies that v_0/V_1 is a negative indicator for equilibrium precaution. Rewriting this indicator as $\delta^F (a/v_0 + v_1/v_0)^{-1}$, we see that equilibrium precaution increases with the relative share of income not raising privacy concerns and with the growth rate of the value of information over time. Therefore, everything held equal, we expect more precaution by an e-seller relying extensively on merchandising than by a social network relying extensively on monetization of personal information. Similarly, the website exerts more precaution if a long history record is more valued than a short history record. Furthermore, 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 $X \in [0,1]$.[14]

Illustration. To illustrate the above comparative statics and provide comparative statics with respect to other parameters of the model, assume that $\theta_l = 0$, the retention rate is always interior (i.e. $0 < Q(0) < Q(1) < 1$), and u is distributed uniformly with density α on its support. In this scenario, the relevant formulas simplify to

$$\Delta_B(r) - \Delta_G(r) = \alpha\theta_h[\lambda(r - r_0) + (1 - r_0)\theta_h r]|U_B|V_1 - v_0$$

and

$$\phi(X) = \frac{1 - \lambda(1 - X)}{1 - [\lambda + (1 - r_0)\theta_h](1 - X)} r_0.$$

In this case, we obtain

$$\psi^f = \alpha\theta_h^2(1 - r_0)r_0\,|U_B|\,; \qquad \psi^n = \frac{\alpha\theta_h^2(1 - r_0)r_0\,|U_B|}{1 - [\lambda + (1 - r_0)\theta_h]}$$

and

$$X^* = 1 - \frac{1 - \alpha\theta_h^2(1 - r_0)r_0\,|U_B|\,\frac{V_1}{v_0}}{\lambda + (1 - r_0)\theta_h} \quad \text{for} \quad \frac{v_0}{V_1} \in \left[\psi^f, \psi^n\right].$$

The level of precaution does not depend on U_G and η as they do not affect the sensitivity of retention to consumer beliefs. It is increasing in $|U_B|$, λ, θ_h, and r_0**:** precaution is higher when matches are less likely to generate a neutral experience and bad matches are more detrimental to consumers. The effect of the prior belief r_0 is non-monotonic. For extreme levels of beliefs, the threshold ψ^n becomes very small as the posterior beliefs do not react to information and $\phi(0)$ is close to r_0. Hence, there is no precaution if consumers are very confident they are not vulnerable, or if they strongly believe that they are vulnerable. Precaution emerges only for intermediate prior beliefs and in this case:

$$\frac{\partial X^*}{\partial r_0} = \frac{\theta_h}{(\lambda + (1 - r_0)\theta_h)^2}\left(\alpha\theta_h^2(1 - r_0)r_0\,|U_B|\,\frac{V_1}{v_0} - 1\right) + \frac{1 - 2r_0}{\lambda + (1 - r_0)\theta_h}\alpha\theta_h^2\,|U_B|\,\frac{V_1}{v_0}$$

[14] Notice that increasing $Q'(r)$ for all $r \geq r_B$ raises the height of $\Delta_B(r) - \Delta_G(r)$ for all $r \geq r_0$, while increasing $|\phi'(X)|$ for all X increases the height of $\phi(X)$, except at $X = 1$ since $\phi(1) = r_0$.

is negative for a large prior r_0 but positive for a small prior and a small price of information v_0 (i.e. close to $\psi^f V_1$). The next graph plots the equilibrium level of precaution as a function of v_0/V_1 and r_0 for $\alpha = 1$, $|U_B| = 10$ and $\lambda = \theta_h = 1/3$:

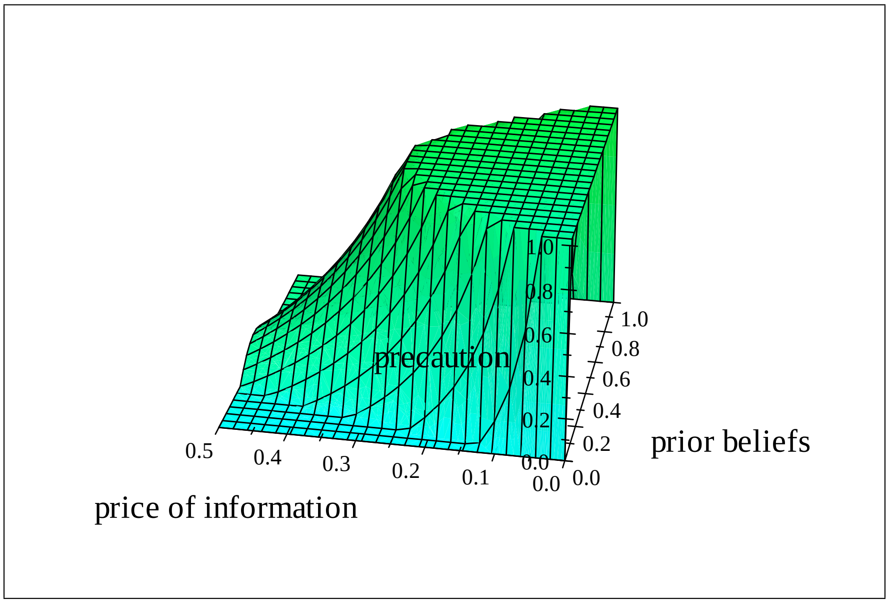

3.2 Welfare

Before turning to public policies intended to enhance privacy, we define our measures of welfare, and then discuss the divergence between social and equilibrium incentives for precaution.

For a given belief r_N (and treating r_G and r_B as parameters), a website choosing precaution X in period 0 makes an expected profit

$$\Pi(r_N, X) \equiv a + (1 - X) v_0 + \mathbb{E}\{Q(r_1)V_1 \mid X\}.$$

where the conditional expectations operator is over realizations of $r_1 \in \{r_G, r_B, r_N\}$. It is immediate that, for any level of precaution X, an improvement in consumer beliefs increases expected profit, i.e. $\partial\Pi(r_N, X)/\partial r_N = p_N(X) Q'(r_N)V_1 > 0$. Furthermore, since the website maximizes $\Pi(r_N, X)$ taking r_N as given, the application of the envelope theorem leads to the following lemma.

Lemma 1 *Marginally lower equilibrium precaution raises expected profit, i.e.* $d\Pi(\phi(X), X)/dX < 0$.

It is useful to decompose consumer expected utility into a short-term component plus a long-term component:

$$U(r_N, X) \equiv U_0(X) + U_1(r_N, X)$$

where $U_0(X) \equiv u_0 + (1-X) M(r_0)$ is the expected utility of consumers in period 0 with a prior belief r_0 and $U_1(r_N, X) \equiv \delta^C \mathbb{E}\{\max(u + M(r_1), 0) \mid X\}$, where the expectations operator is over realizations of both u and r_1, is the expected utility of consumers in period 1. The marginal effect of precaution on short-term utility is equal to the loss of match utility $M(r_0)$, which in general can be either positive or negative. The following lemma, however, establishes that greater equilibrium precaution always decreases long-term consumer utility inclusive of its effect on beliefs.

Lemma 2 *The effect of greater equilibrium precaution on long-term consumer utility is negative, i.e.* $U_1(\phi(X), X)$ *is decreasing in* X*.*

It follows that consumer expected utility is decreasing in equilibrium precaution if $M(r_0)$ is not too negative. In this case, the *ex ante* preference of consumers for marginally less precaution coincides with the preference of the website.

Total expected welfare for given beliefs and precaution is

$$W(r_N, X) \equiv \Pi(r_N, X) + U(r_N, X).$$

Clearly, there is a divergence between equilibrium and social incentives for precaution. In equilibrium, X maximizes $\Pi(r_N, X)$ given $r_N = \phi(X)$, and there is no precaution protection in period 1. This leaves room for several inefficiencies. First, because precaution is unobservable, the website cannot control consumer beliefs. Second, the website ignores the direct effect of X on consumer expected utility. Notice, however, that the marginal divergence between social and equilibrium incentives can be written as

$$\frac{dW(\phi(X), X)}{dX} - \left.\frac{\partial \Pi(r_N, X)}{\partial X}\right|_{r_N=\phi(X)} = \left.\frac{\partial \Pi(r_N, X)}{\partial r_N}\right|_{r_N=\phi(X)} \phi'(X) - M(r_0) + \frac{d\mathbb{E}\{U_1(\phi(X), X)\}}{dX}.$$

The first term is the negative marginal effect of beliefs on expected profit. The second term is the loss of period 0 match utility resulting from marginally greater precaution, which can be either positive or negative. The third term is the marginal effect of precaution on long-term utility, including the effect on beliefs, and is negative according to Lemma 2. Equilibrium precaution in period 0 is deficient when the overall difference is positive. A third possible inefficiency is that it may be desirable to have privacy protection in period 1 if $M(r_1) < 0$ for a sufficiently large segment of consumers in the second period.

We may then distinguish two cases depending on whether match utility is positive or negative. In the "optimistic" case, match utility is positive for all relevant beliefs. In particular, as argued above, if $M(r_0) > 0$, then equilibrium precaution in period 0 is necessarily excessive because both consumers and the website prefer less precaution. Furthermore, if $M(r_1) > 0$ for $r_1 \in \{r_G, r_B, r_N\}$, then zero precaution in period 1 is efficient. Indeed, in this case, reducing precaution would raise consumers' exposure to valuable matches and raise retention due to higher beliefs so that both short-term utility and long-term utility are maximal. In the "pessimistic" case, expected match utility is negative for at least some relevant beliefs, and the social desirability of less precaution is ambiguous for several reasons. First, if $M(r_1) < 0$, then zero precaution in period 1 may be inefficient. Second, if $M(r_0) < 0$, the level of precaution in period 0 may be either insufficient or excessive, depending on whether the short-term loss is outweighed by the long-term utility gain (which depends on the discount factors δ^F and δ^C).

Which case is more salient depends on the nature of customer information and the perceived consumer risks. For example, disclosure of highly sensitive health information may be perceived to cause negative match utility. Alternatively, negative match utility depends on the extent to which information sharing exposes the consumer to malicious attacks. On the one hand, nuisance advertising or innocuous spam might cause only a minor annoyance (small $|U_B|$), that is overshadowed by mutually beneficial targeted advertising (relatively large $|U_G|$). On the other hand, even a small possibility of identity theft due to phishing or malware (very large $|U_B|$) could weigh heavily on match utility. Whether the optimistic case or the pessimistic case is more salient might ultimately depend on policymakers' perceptions of the extent to which online advertising platforms, or other forms of information sharing, increase consumers' risk of becoming victims of serious cybercrime.

3.3 Policy

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

3.3.1 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 $\delta^F\,[a+(1-\tau)\,v_1]$. 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 the ratio $(1-\tau)\,v_0/\delta^F\,[a+(1-\tau)\,v_1]$. As this ratio decreases with the tax rate τ, Proposition 1 implies that a tax which is levied on personal information transactions in both periods would (weakly)[15] increase precaution.

We now turn to privacy regulations.

3.3.2 Transparency and commitment

How should a policy governing X be conducted? One salient policy intervention is to enforce transparency regarding the website's collection and use of personal data. In our setup, transparency may relate to *ex post* disclosure of information sales or to *ex ante* commitment on information that may be sold. For instance, transparency in the European GDPR is of the latter type. A key issue is the extent to which a transparency policy makes credible an announcement of X.

Full commitment Let us assume that a regulation requires a website to reveal and commit to precaution at the beginning of each period, admittedly an exceedingly strong enforcement policy that requires both *ex ante* and *ex post* transparency. We refer to the resulting game as "full commitment"[16] and let X_t denote precaution in period t. The difference between the case where precaution is unobservable and the case where it is observed by consumers is twofold. First, under full commitment, the website can affect posterior beliefs by its choice of precaution. Since profit is increasing in consumer belief r_N, the website will change X_0 in the direction that raises r_N. This is driven by the fact that the website would benefit from consumers interpreting a neutral experience as a stronger signal about their low vulnerability. Second, with full commitment, it may no longer be optimal for the website to choose a no-precaution regime in the second period. Indeed, when third-party matches are detrimental to consumers, commitment to some

[15] Throughout the paper we use "(weakly)" when a given effect is strict unless it is prevented from being so by boundary conditions - here except when $X = 1$.

[16] This scenario corresponds to the case, featured in previous literature (see e.g. O'Brien and Smith, 2014 and Lefouili and Toh, 2018), in which the privacy policy is a publicly observable quality variable.

precaution ($X_1 > 0$) may boost second-period demand. By contrast, when second-period matching is beneficial to marginal consumers on average, the website will choose $X_1 = 0$.

The next proposition provides conditions under which a regulation mandating full commitment results in less precaution in the first period. It also provides conditions under which such a regulation does not affect the website's second-period precaution.

Proposition 2 *Full commitment causes the website to choose (weakly) less precaution in the first period than under no commitment (i.e.* $X_0 \leq X^*$*) if no precaution in the second period is optimal (i.e.* $X_1 = 0$*) or if* $uG'(u)$ *is increasing. Under full commitment,* $X_1 = 0$ *is optimal for the website if either (i)* $M(r_B)$ *is not too negative, or (ii)* $M(r_0) \geq 0$ *and* $uG'(u)$ *is concave.*

Let us now consider the effects of a full-commitment regulation on consumers in light of Proposition 2. The discussion here mimics the discussion of the effects of precaution on welfare as the ability to commit reduces the first period level of precaution. A difference is that the website might also commit to a second period level of precaution. This does not occur if the match utility after a bad experience is not too negative (condition (i)), or if *ex ante* match utility is non-negative and the marginal effect of precaution on participation is concave in beliefs (condition (ii)). We conclude, therefore, that consumers unambiguously benefit in these cases.

Corollary 1 *If* $X_1 = 0$ *is optimal for the website and* $M(r_0) \geq 0$, *then a regulation enforcing full commitment leads to higher short-term and long-term consumer utility.*

There are other scenarios with opposing effects on consumer welfare. First, the website's focus on retaining marginal consumers is not necessarily aligned with the interests of the average consumer. Indeed, if $M(r_0) \geq 0$, then positive second-period precaution detracts from expected consumer welfare.[17] Second, if $M(r_0) < 0$, then consumers are negatively affected by a weaker privacy policy in the first period, creating a tension between short-term and long-term effects on consumer utility. Finally, if second-period

[17] To see why, notice that expected second-period consumer utility can be written as

$$U_1 = E\left\{ \int_{-(1-X_1)M(r_1)}^{+\infty} [1-G(u)]\, du \right\}$$

where expectations operator is over r_1 for a given X_0. Holding X_0 constant, U_1 decreases with positive X_1 if $M(r_0) \geq 0$. Indeed, from $\frac{\partial U_1}{\partial X_1} = E\{[1-G(-(1-X_1)M(r_1))]M(r_1)\}$, $E\{M(r_1)\} = M(r_0) \geq 0$, and the fact that $1-G(-(1-X_1)M(r_1))$ is positive and increasing in r_1, it follows that $\frac{\partial U_1}{\partial X_1} > [1-G(-(1-X_1)M(r_B))]M(r_0) \geq 0$.

match utility is always negative, i.e. $M(\phi(1)) < 0$, and if $G'(u) + uG''(u) > 0$ over the relevant range, then, according to Proposition 2, full commitment reduces precaution in the first period. In this scenario, long-term consumer utility is (weakly) higher, because consumers benefit both from less precaution in the first period and (weakly) more precaution in the second, while the effect on short-term utility is negative.

***Ex post* or *ex ante* transparency** Suppose that the regulator imposes only one type of transparency, either *ex post* or *ex ante.*

Under *ex post* transparency, the website chooses X_t at the beginning of period t and discloses its value once trade with third parties has been completed. Assuming this is aggregate information, each consumer knows the probability X_t to be matched with a third party at the end of period t but not the realized matches. A direct implication is that the website will trade with all third parties in period 1, i.e. $X_1 = 0$, as *ex post* revelation of X_1 does not affect demand. The analysis of the first period is the same as with full commitment as only information that is revealed matters for equilibrium behavior in this period. Hence Proposition 2 implies that a regulation mandating *ex post* transparency leads to a lower precaution level in period 1. Moreover, Corollary 1 applies, showing that when the expected value from a match for consumers is positive or consumers are sufficiently patient, consumers benefit from such a regulation. By contrast, when the expected match utility is negative and consumers are sufficiently impatient, they do not benefit from a regulation mandating *ex post* transparency.

A policy enforcing *ex ante* transparency plausibly enables the website to credibly commit to greater precaution than the equilibrium level, i.e. enforces an announced lower bound on precaution. For example, the website might commit to more precaution by promising (transparently) to collect less data about consumers, which compromises the website's ability to match consumers and interested third parties. We have shown, however, that the website wants to commit to less rather than more precaution in order to improve customer retention after a neutral experience. But it is not clear how the website credibly commits to less precaution. For example, suppose the website announces that it will collect additional personal data about its customers, potentially improving matching with interested third parties. The website would still have the ability and the incentive to refuse to deal with interested third parties, as long as the refusal is unobservable to consumers and did not violate its announced privacy policy. In other words, a commitment to reduce precaution is not credible.[18] In this case, a policy enforcing *ex ante* transparency,

[18] Formally, if the website were to announce $X < X^*$, the equilibrium would be in mixed strategies with the website refusing to sell with positive probability, such that the probability of a match is X^*.

without also sanctioning refusals to deal, would have no effect on equilibrium precaution, under the condition that no-precaution in the second period is optimal for the website, as stated in Proposition 2.

The following proposition summarizes the above discussion.

Proposition 3 *-A regulation enforcing ex post transparency causes the website to choose (weakly) less precaution in the first period, compared to the equilibrium with no commitment, and no precaution in the second period.*

- A regulation enforcing ex ante transparency causes the website to choose the same privacy policy as in the equilibrium with no commitment if no-precaution in the second period is optimal for the website and refusal to sell information is possible.

3.3.3 Opt-out

Another salient policy measure is to give consumers 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 the nature of the match (good, bad or neutral) with the buyer of personal information is not. We allow consumers to opt out, which means they can prevent any sale of personal information (the full precaution regime would then prevail).[19] 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 decision to opt out depends on her beliefs about her vulnerability to bad experiences. Let

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

denote the solution to $M(r_1) = 0$. The most interesting scenario, assumed below, is $r_0 \geq \bar{r} > r_B$, in which case consumers will opt out only after a bad experience.[20] Since opting out assures a match value of $M(\bar{r}) = 0$ instead of $M(r_B) < 0$, the expected retention of those consumers rises from $Q(r_B)$ to $Q(\bar{r})$.

[19] Bloch and Demange (2018) also analyze the effect of an opt-out option on a website's privacy policy (captured by its level of data exploitation). However, they assume that the website's policy is observable to consumers.

[20] Other cases are considered in Section 1 of the Online Appendix.

At the same time, the website's value of retaining consumers who opt out falls from $V_1 \equiv \delta^F (a + v_1)$ to $\bar{V}_1 \equiv \delta^F a$, because the website no longer earns revenue in period 1 from selling information for those customers. Whether or not website profit increases in period 1 depends on the combined effect of greater retention and lower retention value, i.e. on whether $Q(\bar{r})\bar{V}_1$ is greater or less than $Q(r_B)V_1$. The consequences in either case depend on the equilibrium effects of a marginal change in expected profit for customers having a bad experience.

Lemma 3 *In the baseline model, everything else held equal, an increase in* $Q(r_B)V_1$ *leads to a (weakly) higher equilibrium posterior belief* r_N^* *and (weakly) lower equilibrium precaution* X^**.*

We study the website's incentive to offer an opt-out option to its customers (on a voluntary basis) before investigating the impact of a mandatory opt-out policy.

Voluntary opt-out The website's profit is

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

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

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

Since the profit increases with r_N we have two concurring effects. The website derives a direct benefit from higher future revenues from those consumers having a bad experience. Moreover, Lemma 3 implies an additional indirect benefit: more optimistic beliefs about vulnerability increases the retention of those consumers who have a neutral experience. Consequently, the website has an incentive to offer opt-out if $Q(\bar{r})\bar{V}_1 \geq Q(r_B)V_1$. Conversely, the website has no such incentive if the opposite is true.

From Lemma 3 it follows that the website lowers its level of precaution whenever it offers an opt-out option on a voluntary basis. This reduction in the level of precaution leads to an increase in short-term consumer utility if $M(r_0) > 0$ as well as an increase in long-term consumer utility because a neutral experience becomes more informative about vulnerability. Combining this with the fact that if $M(r_B) < 0$ opt-out beneficially shields consumers from bad experiences[21] implies by revealed preferences that voluntary opt-out unambiguously benefits both consumers and the website.

[21] If *ex ante* match utility was negative, consumers might opt out in the first period.

Mandatory opt-out policy The above analysis shows that a mandatory opt-out policy affects the website's behavior if and only if $Q(\bar{r})\bar{V}_1 < Q(r_B)V_1$. In this case, Lemma 3 implies that such a policy leads to more precaution. Website profit of course declines by revealed preference, and greater precaution reduces short-run consumer utility if $M(r_0) > 0$. The long-term consumer welfare effects are generally ambiguous because consumers are shielded from a bad experience if $M(r_B) < 0$ but are negatively affected by the decreased informativeness of a neutral experience about vulnerability; obviously the latter effect dominates if $M(r_B)$ is not too negative.

Assuming an indifferent website offers opt-out, the conclusions are summarized as follows.

Proposition 4 *Assume $M(r_0) \geq 0 > M(r_B)$. (i) If $Q(\bar{r})\bar{V}_1 \geq Q(r_B)V_1$, then the website voluntarily offers opt-out, resulting in (weakly) less precaution, and increasing both short- and long-term consumer utility. (ii) If $Q(\bar{r})\bar{V}_1 < Q(r_B)V_1$, then a mandatory opt-out policy results in (weakly) more precaution. Profits and short-term consumer utility are lower while the effect on long-term consumer utility is negative if and only if $M(r_B)$ is not too negative.*

4 Extensions

4.1 Multi-homing consumers

Let us consider $K \geq 2$ 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 websites are *ex ante* identical so that consumers multi-home and are active on all websites in period 1. This implies 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, except that the consumers derive a utility u_i from visiting website i that follows a joint distribution $G_K(u_1, ...u_K)$ which we assume to be symmetric with marginal cumulative distribution $G(u_i)$.

In a setting with multiple websites, two issues arise: attribution and inference. The first relates to the fact that a consumer having a bad experience may not know which website is responsible for the sale of her personal information. Inference refers to the way a consumer revises her beliefs on each website after a given experience, which depends in particular on the correlation of the vulnerability across websites. Both imperfect attribution and correlation induce a "collective reputation" problem as the sale of personal information by one website affects all websites. To highlight this aspect we assume that consumers do not observe if information is sold or not, and by whom, and that the parameter θ is the same for all websites. We show below that even under these extreme assumptions the market may provide some privacy protection.[22]

Given our maintained assumption that there is a one-to-one match between consumers and third parties, each consumer faces the same potential outcomes 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, observes the realization of $(u_1, ...u_K)$, and decides whether to return to the websites.

A new feature here is that the probability of a non-neutral experience accounts for the fact that many websites can sell information. Thus, if x is a symmetric equilibrium probability that a website does not sell customer information, the total probability that a third party interested in buying such information does not acquire it is $X = x^K$. Notice that we assume independent probabilities that a third-party transacts with a website.[23] 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 $r_N = \phi(X)$.

We model competition by assuming that all websites decide independently and simultaneously on x and the price p for personal information. We assume for simplicity that all websites observe when information is sold but do not observe consumers' experiences. On the market for information, multi-homing affects the selling prices in both periods. Let us first consider period 1. If information is not sold in period 0, the posterior is $r_1 = r_N$ and the consumer returns to a given website i with probability $Q_K(r_N)$ where

$$Q_K(r) \equiv \Pr\left\{u_i \geq 0; \sum_{j=1,K} \max(u_j, 0) \geq -M(r)\right\} + \Pr\left\{0 > u_i \geq -M(r); u_i > \max_{j \neq i} u_j\right\}.$$

The website obtains profit a on each retained consumer and competes with other websites for the sale of personal information. The equilibrium expected period 1 revenue

[22] Our conclusion would extend to the case where vulnerability is website-specific, with θ_i denoting vulnerability on website i, if either the consumer cannot attribute perfectly a bad or good experience to a website or the parameters θ_i are positively correlated.

[23] This is the case if x is interpreted as a mixed strategy, or if websites collect different relevant pieces of information and third-parties need only one piece of relevant information.

of website i from selling information on a returning consumer is then $v_1\hat{Q}_K(r_N)$, where

$$\hat{Q}_K(r) \equiv \Pr\left\{u_i \geq -M(r); \min(u_i, 0) > \max_{j\neq i} u_j\right\}$$

is the probability of unique customer retention, i.e. the probability that a customer returns only to website i. This expression can be obtained as follows. Each website where the consumer is still active knows that the consumer had a neutral experience but does not know on which websites she returned.[24] The equilibrium symmetric strategy of the pricing game (on the market for information) in period 1 is then a mixed strategy on an interval (a standard undercutting argument shows that there are neither mass points nor holes). The website earns the same expected profit for all prices in the interval. Moreover, as the upper bound of that interval wins only if there is no other bid (which happens with probability $\hat{Q}_K(r_N)$), it has to be v_1. This yields the expected payoff $v_1\hat{Q}_K(r_N)$. Thus, if information is not sold by any website in period 0, the period 1 expected revenue as of period 0 is:

$$\delta^F\left[aQ_K(r_N) + v_1\hat{Q}_K(r_N)\right].$$

If information is sold, the equilibrium expected revenue of website i is similar but accounts for the uncertainty on the posterior r_1. Viewed from period 0, it is given by

$$L_K(r_N) \equiv \mathbb{E}\left\{\delta^F\left[aQ_K(r_1) + v_1\hat{Q}_K(r_1)\right] \mid \text{ information is sold in period } 0\right\}.$$

Let us now turn to competition on the market for information in period 0. In this game, the strategy of a website can be summarized by the choice of x and a probability distribution over prices $p \leq v_0$ with cumulative distribution function $F(p)$.[25] We now characterize a symmetric equilibrium of the game. Let us first derive the website's optimal strategy for a given consumer belief r_N, assuming that all other websites follow the strategy $(x, F(.))$. For this purpose, define

$$\begin{aligned} P_K(r_N) \equiv\ & \delta^F a\left\{\lambda\left[Q_K(r_N) - Q_K(r_G)\right] + \theta_0\left[Q_K(r_N) - Q_K(r_B)\right]\right\} \\ & +\delta^F v_1\left\{\lambda\left[\hat{Q}_K(r_N) - \hat{Q}_K(r_G)\right] + \theta_0\left[\hat{Q}_K(r_N) - \hat{Q}_K(r_B)\right]\right\}, \end{aligned}$$

[24] Alternatively, we could assume that websites track consumers and know where they return. In this case, the equilibrium price of information is v_1 if there is a monopoly and 0 if the consumer returns to two or more websites. A website's expected payoff would then be the same.

[25] We could allow the price to be above v_0 without altering the analysis. In this case x would be replaced by $x' = x + 1 - F(p_0)$. Given that putting a mass above v_0 is equivalent to increasing x, we ignore this possibility. See, however, our discussion of transparency and commitment.

which is the expected increase in the payoff from retaining a consumer when no information about the consumer is sold to third parties. Given that $Q_K(r_N)$ and $\hat{Q}_K(r_N)$ are both increasing in r_N, $P_K(r_N)$ is positive and non-decreasing in r_N. Then the payoff of a website choosing x_i and $p_i \leq v_0$ is

$$L_K(r_N) + x_i x^{K-1} P_K(r_N) + (1 - x_i)(x + (1 - x)(1 - F(p_i)))^{K-1} p_i.$$

This expression can be interpreted as follows. Not selling customer information does not imply that access to that information by a third party will not occur because another website may sell the information. In that case the payoff is $L_K(r_N)$. However, with probability x^{K-1}, all other websites also do not sell the information and the payoff is then higher by $P_K(r_N)$. Finally, information is sold by the website if the third party is willing to buy and either no other website sells relevant information or the website's price is the lowest price. Therefore, the comparison of the payoff from selling information at a price p below v_0 and the payoff from not selling information boils down to comparing the expected revenue $(x + (1 - x)(1 - F(p)))^{K-1} p$ with the gain from privacy $P_K(r_N) x^{K-1}$.

Equilibrium is not necessarily unique in a multi-homing environment. The following result shows that one of the possible outcomes is that all websites sell personal information to any interested third party and competition dissipates fully their profits from data sales.

Proposition 5 *Suppose there are at least two websites. Then, there always exists a no-precaution 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.[26] The next proposition shows, however, that there also exists a symmetric equilibrium with a positive level of precaution, denoted X_K^*, if the value of information v_0 is not too high.

Proposition 6 *Suppose there are $K \geq 2$ websites. Then, there exists a symmetric equilibrium with positive precaution (i.e., $X_K^* > 0$) if and only if $v_0 < P_K(\phi(0))$. It is uniquely defined by the following conditions:*

- *If $v_0 \leq P_K(\phi(1))$, then the websites provide full precaution (i.e., $X_K^* = 1$).*
- *If $P_K(\phi(1)) < v_0 < P_K(\phi(0))$, then the websites' level of precaution and price*

[26] However, we conjecture that if there is an arbitrarily small mass ε of single-homing consumers, the zero-price equilibrium exhibited in Proposition 5 exists only if $v_0 \geq \Delta_B(r_N) - \Delta_G(r_N)$.

distribution are given by

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

In a multi-homing context, the number of websites affects the total level of precaution only through its effect on the value of retaining a consumer. The probability of unique retention is smaller and less responsive to beliefs with multiple websites than with a single website. The same holds for retention when the expected value of a match is positive. But when the expected value of a match is negative at relevant levels of r_N, the retention rate is larger with multiple websites as the consumer benefits from multiple visits. In this case the retention may be more or less responsive to beliefs than with a single website, and the level of precaution may be higher if v_1 is small.

The profit of a website may be lower when consumers multi-home for two reasons. First, future competition is detrimental to profits. Second, there is a chance that another website sells customer information, which leads to a free rider problem. However, we have seen that some reduction of the level of precaution may be profitable as it raises average retention. The strategic effect that reduces the aggregate level of precaution could in principle be strong enough to offset the reduction of revenues due to competition in the market for information.

Multi-homing has mixed effects on consumer welfare. Obviously, 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 the long-run utility from any website. Short-term utility, however, increases only if the *ex ante* match utility is positive.

Let us now consider the effect of a transparency policy, focusing on symmetric equilibria (the detailed analysis is provided in Section 2 of the Online Appendix). If websites are subject to *ex post* transparency, i.e. if the information x_i sold by each website i is observed at the end of the period, the analysis is similar to the case of a single website. The websites choose to sell information to all third parties in the second period and may refrain from doing so in the first period. A difference with the case of a non-transparent policy is that it is no longer the case that the maximum price of information v_0 is equal to the incremental value $P_K(\phi(X))$ of privacy. As websites internalize the effect of their policy on consumers' posterior beliefs and retention rates, they may choose a lower level of precaution. We show in the Online Appendix that when all other websites choose $x_j = x$, the equilibrium payoff of a website i is given by

$$L_K\left(\phi\left(x_i x^{K-1}\right)\right) + x_i x^{K-1} P_K\left(\phi\left(x_i x^{K-1}\right)\right) + \left(1 - x_i\right) x^{K-1} v_0.$$

The resulting equilibrium level is then smaller than the level X_K^* under non-transparency because of the effect of precaution on beliefs.

Consider now an *ex ante* transparency policy, where x_i is observed before price competition takes place. Notice first that in the second period, the incentives to refrain from selling information are smaller than for a monopoly as others can also sell the same information. Assuming that information is sold to all third parties in the second period, reducing the level of precaution may not be credible if the website can refuse to sell. In particular, starting from the equilibrium precaution level $x = (X_K^*)^{\frac{1}{K}}$, a website who deviates and announces $x_i < (X_K^*)^{\frac{1}{K}}$ would refuse to sell with probability y such that $y + (1-y)\, x_i = (X_K^*)^{\frac{1}{K}}$. The equilibrium distribution of prices and the market level of precaution would then be unaffected. Hence the equilibrium obtained under *ex post* transparency would not be credible under *ex ante* transparency, suggesting that equilibrium precaution is not affected.

However, with multiple websites competing on the market for information, a website will also care about the effect of its policy on the intensity of competition. In particular, a commitment to a higher level of precaution than under a non-transparent policy, i.e. $x_i > (X_K^*)^{\frac{1}{K}}$, is credible and would induce competitors to price less aggressively. The reason is that firm i's data becomes less attractive, which allows competitors to raise prices without changing the probability of sale. If this effect is strong, which occurs when X_K^* is small, the websites would deviate from $(X_K^*)^{\frac{1}{K}}$, and we show there does not exist an equilibrium in which websites choose pure precaution strategies.

4.2 Verification of third party use of information

In this section, we assume that the website can verify third party use of information. For instance, a website may use Artificial Intelligence and human resources to screen out ads with offensive content, or may verify that requests to access customer information originate from known legitimate third parties. To capture this, we suppose that, by incurring a cost $C_v(Y)$, the website can generate a signal on the user experience that is informative with probability Y. A non-informative signal conveys no information while an informative signal allows to detect perfectly whether the third party will generate a bad experience or not. Thus, an informative signal allows screening third parties generating bad experiences from those leading to good or neutral experiences.[27] For conciseness, we assume that $C_v(Y)$ is convex, $C_v'(0)=0$ and $C_v'(1)=+\infty$.

4.2.1 Strategies and beliefs

The website's strategy now consists of a choice of a level of verification Y as well as the probability X that customer information is not sold in case verification fails. If verification succeeds then the website denies access to information to third parties generating a bad experience and grants other third parties full access to information.

Therefore, we can characterize the website's strategy by a pair

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

Under a full precaution policy ($X=1$), 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 then determines the benefit from allowing access to third parties that do not provide a bad experience. On the contrary, under a no precaution policy ($X=0$), verification is the only way to avoid interactions with third parties that generate a bad experience and, therefore, determines the level of protection against them.

Let us now provide the probability of each type of first-period experience and determine how it depends on the level 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 $p_G(X,Y)=\lambda\left[Y+(1-Y)(1-X)\right]$ and $p_B(X,Y)=(1-X)(1-Y)\theta_0$. 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 experience is increasing in the level of verification unless there is no precaution, because more verification decreases the likelihood that a third party generating a good experience buys customer information. By contrast, the probability of a bad experience decreases with the level of verification unless there is full precaution. The probability of a neutral experience (event N) is then

$$p_N(X,Y)=1-Y\lambda-(1-Y)(\lambda+\theta_0)(1-X). \tag{5}$$

[27] As an illustration, suppose 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 bad experience or not. It is straightforward that there must exist a critical level $\hat{z}$ (potentially zero) such that the website verifies the third party's use of information if $z<\hat{z}$. We can then denote by $Y=H(\hat{z})$ the probability of verification and the cost of verification is $C_v(Y)=\int_0^{\hat{z}} z dH(z)$.

This probability increases with X as in the case with no verification. The effect of the level of verification Y on the probability of a neutral experience depends on the level of precaution X, as shown by $\partial p_N / \partial Y = \theta_0 - (\lambda + \theta_0) X$. The reason is that a higher level of verification has two (potential) opposite effects on the likelihood of a neutral experience. First, it affects it positively by increasing the probability that a third party generating a bad experience 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 effect dominates the latter if the level of precaution is low (so that third parties have an easy access to customer information).

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

$$r_N = \Phi(X, Y) \equiv \frac{1 - Y\lambda - (1 - Y)(\lambda + \theta_l)(1 - X)}{1 - Y\lambda - (1 - Y)(\lambda + \theta_0)(1 - X)} r_0. \tag{6}$$

A neutral experience is again good news in the sense that $r_B < r_G < r_N$ for any $X < 1$. The following lemma shows how the posterior belief after a neutral experience depends on the levels of precaution and verification.

Lemma 4 *i)* $\Phi(X, Y)$ *is decreasing in* X*, ii)* $\Phi(X, Y)$ *is decreasing in* Y *for any* $X < 1$*, iii)* $\Phi(0, 0) = \phi(0)$ *and* $\Phi(1, Y) = r_0$ *for all* Y*.*

Notice that the range of possible beliefs is the same with or without verification. By reducing exposure to a third party generating a bad experience, verification makes the experience less informative about vulnerability and reduces the posterior r_N whenever $X < 1$. Under full precaution, no bad experience can occur and the posterior is not affected by a neutral or a good experience.

4.2.2 Equilibrium analysis

The analysis of the website's decision regarding the level of precaution is similar to the one in the baseline scenario (with no verification). Indeed, as precaution refers to the probability of selling information when there is no verification, the website's optimal level of precaution does not depend on the level Y of verification. The equilibrium level of precaution when the website anticipates a posterior r_N is still $X^{br}(r_N)$, as given by equation (2).

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 or neutral experience (and would be used only for this purpose since $r_B < r_G$). For a given level of precaution X, verification raises the probability to sell customer information to a third party generating a good or neutral 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 generates a bad experience from $1 - X$ to 0, which leads to another expected benefit given by $(1 - X)\Delta_B(r_N)$. The total benefit from verifying third parties' use of information is then the sum $X\Delta_G(r_N) + (1 - X)\Delta_B(r_N)$. When evaluated at $X = X^{br}(r_N)$, the equilibrium condition (2) implies that this gain is equal to $\min(\Delta_B(r_N), \Delta_G(r_N))$. Thus, some verification occurs (i.e. $Y > 0$) whenever this benefit is positive, and the website's optimal verification level $Y^{br}(r_N)$ is given by

$$C_v'\left(Y^{br}(r_N)\right) = \max\left\{\min\left(\Delta_G(r_N), \Delta_B(r_N)\right), 0\right\}. \tag{7}$$

Thus, an equilibrium is characterized by X^{**}, Y^{**} and r_N^{**} that solve equations (2), (6) and (7). Recall from equation (2) that without verification full precaution occurs when $\Delta_G(r_0) \leq \Delta_B(r_0)$, while no precaution occurs when $\Delta_G(\phi(0)) > \Delta_B(\phi(0))$. Hence we distinguish three scenarios:

- *Full precaution*: this occurs when $r_N^{**} = r_0$ and the level of precaution is $C_v'(Y^{**}) = \max\{\Delta_G(r_0), 0\}$.
- *Partial precaution*: this is only possible if the posterior belief after a neutral experience satisfies $r_0 < r_N^{**} = r^M < \phi(0)$, and the level of verification is given by $C_v'(Y^{**}) = \max\left\{\Delta^M, 0\right\}$ where $\Delta^M \equiv \Delta_G\left(r^M\right) = \Delta_B\left(r^M\right)$.
- *No precaution* : this requires a low posterior belief $r_N^{**} < r^M$ and a level of verification given by $C_v'(Y^{**}) = \max\{\Delta_B(r_N^{**}), 0\}$.

Notice that, because $\Delta_G(r)$ decreases in r while $\Delta_B(r)$ increases in r, verification is a single-peaked function of r_N achieving a maximum at posterior belief r^M.

For v_0 close to 0, the gain $\Delta_G(r_N)$ can be made arbitrarily close to 0 or negative (because $r_G = r_0 = \phi(1)$), while the gain $\Delta_B(r_N)$ remains strictly positive. Thus, for sufficiently small values of v_0, the website chooses a full precaution policy and no verification, so that customer information is never sold. Similarly, for sufficiently large values of v_0, the website chooses no precaution and no verification, so that customer information is sold to all interested third parties. The following equilibrium characterization shows that the results of Proposition 1 pertaining the level of precaution extends to a setting with verification:

Proposition 7 *A unique equilibrium exists.*

(i) Equilibrium precaution is non-increasing in the value of information v_0 and there exist a threshold $v^n \in \left(\psi^f V_1, \psi^n V_1\right)$ such that the website chooses full precaution if $v_0 \leq \psi^f V_1$, partial precaution if $v_0 \in \left(\psi^f V_1, v^n\right)$ and no precaution if $v_0 \geq v^n$;

ii) There exists a threshold $\bar{v} > v^n$ such that the level of verification is positive if and only if $v_0 < \bar{v}$. Equilibrium verification is increasing in v_0 in the full- and partial-precaution region and non-increasing in v_0 in the no-precaution region.

Verification always occurs under the full precaution regime because the benefit $\Delta_G\left(r_0\right)$ from selling information when it does not generate a bad experience is positive. It also occurs under the no precaution regime if the benefit $\Delta_B\left(\phi\left(0\right)\right)$ from avoiding sales of information leading to a bad experience is positive, which is the case when v_0 is not too large. Finally, in the partial precaution regime, the gain from verification Δ^M is always positive.

Let us now consider the way the website's equilibrium strategy depends on v_0. As in the case when there is no verification, the equilibrium level of precaution is non-increasing in v_0. When the value of information increases the equilibrium moves toward more frequent access of third parties to customer information, leading to a higher posterior belief r_N. Verification allows to restrict sales to third parties generating good or neutral experiences, which induces a short-term revenue loss that depends on the price v_0 but also on the level of precaution. Under the full precaution regime, raising v_0 makes verification more attractive as it generates more sales. In contrast, under the no precaution regime, raising v_0 makes verification less attractive as it reduces the probability to sell customer information. The proposition shows that the partial precaution regime is similar to the full precaution regime in this respect. Hence, we find a non-monotonic effect of the value of personal information v_0 on the level of verification.

The level of verification Y in our model can be viewed as a form of security investment that reduces the risk of a bad experience from θ to $\left(1-Y\right)\theta$. Introducing such a technology affects the learning process and thus the website's incentives to sell information to third parties that are not proved to generate good or neutral experiences. Since the introduction of the verification technology does not affect the website's incentive to exert precaution X for a given posterior belief r_N but depresses the posterior belief for a given level of precaution, it follows that the website exerts (weakly) less precaution when the probability of a bad experience is lower.

The introduction of the verification technology may have one of the following effects on the level of precaution and the posterior belief r_N.

- First, if $X^* = 1$, they are not affected (this is because $\Phi(1, Y) = r_0$ is not affected by Y).
- Second, if $X^* = 0$ and some verification occurs, the level of precaution remains equal to zero and the posterior belief r_N declines.
- Finally, if there is partial precaution in the absence of a verification technology, i.e. $0 < X^* < 1$, there are two possible scenarios:
  - In the first scenario, the level of precaution declines but remains positive and the posterior belief $r_N = r^M$ is unchanged – this happens when $\Phi\left(0, Y^{br}\left(r^M\right)\right) \geq r^M$;
  - In the second scenario, the level of precaution falls from positive to $X^{**} = 0$ and the posterior belief declines to $r_N^{**} = \Phi(0, Y^{**}) < r^M$ – this happens when $\Phi\left(0, Y^{br}\left(r^M\right)\right) < r^M$.

Notice that in all cases the level of precaution is either unaffected or lower than in the baseline model without verification. In this sense, precaution and verification are substitutes.

Proposition 8 *Verification reduces the equilibrium level of precaution. More generally, a reduction in the marginal cost of verification raises Y^{**} and reduces X^{**}.*

Consider the effect of a uniform reduction in the marginal cost of verification that either makes verification profitable or makes it easier when it is already profitable. Such a reduction of the verification cost could result from technological advances or policy measures. For instance, the law may impose an obligation of transparency on data buyers or policy may promote public or private certification improving information on potential data buyers. From the above analysis it follows that any technology or policy change that would reduce the cost of verification would lead to less precaution, more verification and (weakly) lower posterior belief r_N.

4.2.3 Welfare analysis

Let us now consider the welfare implications of verification. Suppose first that the posterior belief r_N is not affected by the possibility to verify and screen out bad experiences.

Then, a simple revealed preference argument shows that profit increases. From equation (6) we have

$$1 - X^{**} = \frac{1-\lambda Y^{**}}{1-Y^{**}} \frac{r_N - r_0}{(\lambda+\theta_0) r_N - (\lambda+\theta_l) r_0}.$$

The resulting probabilities of events G, B and N are then given, respectively, by

$$\begin{aligned} p_G^{**} &= \frac{\lambda (r_N - r_0)}{(\lambda+\theta_0) r_N - (\lambda+\theta_l) r_0} + \frac{\lambda (\theta_0 r_N - \theta_l r_0)}{(\lambda+\theta_0) r_N - (\lambda+\theta_l) r_0} Y^{**}, \\ p_B^{**} &= \frac{\theta_0 (r_N - r_0)}{(\lambda+\theta_0) r_N - (\lambda+\theta_l) r_0} (1-\lambda Y^{**}), \\ p_N^{**} &= \left[1 - (\lambda+\theta_0) \frac{r_N - r_0}{(\lambda+\theta_0) r_N - (\lambda+\theta_l) r_0}\right] (1-\lambda Y^{**}). \end{aligned}$$

This shows that the availability of a verification technology leads to an increase in the likelihood of a good experience and a decrease in the likelihood of a neutral or bad experience. It follows that the short-term consumer utility increases. Notice, however, that the distribution of the posterior beliefs r_1 undergoes a mean-preserving contraction. Hence, the long-term consumer utility declines reflecting a decrease in the informativeness of the signal after a neutral experience.

Consider now the scenario in which the posterior belief r_N declines as a result of the availability of a verification technology. This occurs only when the resulting level of precaution is zero, thus leading to the following probabilities of experiences: $p_G(0, Y^{**}) = \lambda$, $p_B(0, Y^{**}) = (1 - Y^{**})\theta_0$, and $p_N(0, Y^{**}) = 1 - \lambda - (1 - Y^{**})\theta_0$. As the probability of a bad experience is reduced and the probability of a good experience is either increased or unchanged, the effect of verification on short-term consumer utility is positive. The change in the distribution of posterior beliefs satisfies a single-crossing property so that verification induces a reduction in risk in the sense of second-order stochastic dominance. This implies that the effect on long-term consumer utility is again negative.

We thus reach the following conclusion.

Proposition 9 *The introduction of the verification technology raises short-term consumer utility and reduces long-term consumer utility.*

Introducing verification technology into the baseline model strengthens our perspective that simple regulations that reliably improve consumer welfare are elusive. Consider, for example, a case in which $X^{**} = 0$ and $Y^{**} > 0$; in equilibrium, the website verification partially screens-out bad actors, but otherwise exercises no precaution. In this case, full commitment would create an incentive for the website to reduce verification.

The consequence is to reduce short-term consumer welfare because consumers are more exposed to bad experiences in the first period, but to increase long-term utility because consumer learn about their vulnerability from a neutral experience. Thus, in comparison to the baseline model, full commitment has an ambiguous effect on consumer welfare in the model with verification. Notice that an *ex ante* transparency policy that doesn't constrain verification would again be inconsequential. Similarly a mandatory opt-out policy reduces a website's incentive for verification, harming consumers in the short-run, while improving long-term consumer welfare. Also, opt-out and reduced verification could create incentives for precaution, further reducing short-term consumer utility if $M(r_0) > 0$, and with countervailing effects on long-term utility.

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 sharing personal information with third parties, and that consumers who become pessimistic about their vulnerability are less likely to return to the website. In response, the website exercises precaution in dealing with third parties and verifies third party use of customer information, in order to profit from better consumer retention.

Our analysis shows how a website's incentive for privacy protection improves with the value of consumer retention relative to the revenue from sharing 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 and are less informed about their vulnerability to third-party intrusions. Consequently, it is difficult for authorities to regulate privacy protection in a way that reliably improves consumer welfare. For example, policies that tax information sales, improve the transparency of privacy policies, and give consumers more control over their personal information, all have either mixed or neutral effects on consumer welfare.

There are many 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 endogenously removing cookies. Another is to allow websites to charge a subscription fee for continued access, 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: Proofs

Proof of Proposition 1. Clearly, since $\Delta_B(r_N) - \Delta_G(r_N)$ decreases in v_0 from positive to negative values, full precaution must be an equilibrium for sufficiently small values of v_0/V_1; a necessary and sufficient condition is $\Delta_B(\phi(1)) - \Delta_G(\phi(1)) \geq 0$ which yields the threshold ψ^f. Similarly, no precaution is an equilibrium if and only if v_0/V_1 is sufficiently large that $\Delta_B(\phi(0)) - \Delta_G(\phi(0)) \leq 0$ which yields ψ^n. Thus, there is no pure strategy equilibrium if $\Delta_B(\phi(0)) - \Delta_G(\phi(0)) > 0 > \Delta_B(\phi(1)) - \Delta_G(\phi(1))$. In this range of v_0/V_1, there exists a unique belief that holds the website indifferent about precaution, and this belief pins down equilibrium precaution: $\Delta_B\left(r^M\right) = \Delta_G\left(r^M\right)$ and $\phi(X^*) = r^M$. Monotonicity follows from r^M increasing in v_0/V_1 and $\phi(X)$ decreasing.

Proof of Lemma 2. The function $v(r_1) \equiv \mathbb{E}\{\max(u + M(r_1), 0)\}$ is convex in r_1. Moreover, if $r_N = \phi(X)$, an increase in X induces a mean-preserving contraction of the distribution of r_1; this follows from $\mathbb{E}\{r_1 \mid X\} = r_0$, $p_G(X)$ and $p_N(X)$ both decreasing, and r_N decreasing in X. Therefore, from Rothschild and Stiglitz (1971), $\mathbb{E}\{v(r_1) \mid X\}$ is decreasing in X. The result follows because $U_1(\phi(X), X) = \delta^C \mathbb{E}\{v(r_1) \mid X\}$.

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 $\hat{\Pi}(r_N, X_0, X_1) = (1 - X_0)v_0 + \mathbb{E}\left\{\hat{Q}(r_1, X_1)\hat{V}_1(X_1) \mid X_0\right\}$ where $\hat{Q}(r_1, X_1) \equiv 1 - G(-(1 - X_1)M(r_1))$ and $\hat{V}_1(X_1) = \delta^F[a + (1 - X_1)v_1]$.

Assume first that no precaution in the second period is optimal, i.e. $X_1 = 0$. The website's profit then reduces to $\hat{\Pi}(r_N, X_0, 0) = \Pi(r_N, X_0)$. As $r_N = \phi(X_0)$ is decreasing in X_0, the marginal gain of the website from increasing X_0 is lower when it can commit to its strategy: $\frac{\partial \Pi(r_N, X_0)}{\partial X_0} + \frac{\partial \Pi(r_N, X_0)}{\partial r_N}\frac{d\phi}{dX_0} < \frac{\partial \Pi(r_N, X_0)}{\partial X_0}$. This implies that a full precaution equilibrium exists for a smaller range of values v_0/V_1 while a no precaution equilibrium exists for a wider range. Consider now an equilibrium with no commitment featuring an interior level of precaution $X^* \in (0,1)$. Then for any $X_0 > X^*$, $\Pi(\phi(X^*), X^*) > \Pi(\phi(X^*), X_0) > \Pi(\phi(X_0), X_0)$. Therefore, the website chooses $X_0 \leq X^*$. Moreover at $X_0 = X^*$ it holds that $\frac{\partial \Pi}{\partial X_0} + \frac{\partial \Pi}{\partial r_N}\frac{d\phi}{dX_0} = \frac{\partial \Pi}{\partial r_N}\frac{d\phi}{dX_0} < 0$, which implies that the website chooses $X_0 < X^*$.

Let us now show that an alternative sufficient condition for full commitment to lead to lower first-period precaution is that $uG'(u)$ is increasing in u. The above comparison of the optimal level of precaution under full commitment to the equilibrium level 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 level 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 $v_0/\hat{V}_1(X_1)$ is increasing in X_1, a sufficient condition for first-period level of precaution to be decreasing in X_1 is that the gain in retention probability from being cautious, i.e. $\left\{[\lambda + \theta_0]\hat{Q}(r_N, X_1) - \lambda\hat{Q}(r_G, X_1) - \theta_0\hat{Q}(r_B, X_1)\right\}$ is decreasing in X_1. The latter holds if $\frac{\partial}{\partial X_1}\hat{Q}(r_N, X_1) < \min\left\{\frac{\partial}{\partial X_1}\hat{Q}(r_G, X_1), \frac{\partial}{\partial X_1}\hat{Q}(r_B, X_1)\right\}$ or if $\frac{\partial^2}{\partial r_1 \partial X_1}\hat{Q}(r_1, X_1) < 0$. Since

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

it follows that a sufficient condition for full commitment to cause the website to choose a lower first-period precaution than under no commitment is that $G'(u) + uG''(u) > 0$, i.e. $uG'(u)$ is increasing.

Let us now provide sufficient conditions under which $X_1 = 0$ is optimal. We have

$$\frac{\partial\hat{\Pi}}{\partial X_1} = -\mathbb{E}\left\{G'(-(1-X_1)M(r_1))M(r_1) \mid X_0\right\}V_1(X_1) + \mathbb{E}\left\{1 - G(-(1-X_1)M(r_1)) \mid X_0\right\}V_1'(X_1).$$

The second term is negative. The first term is non-positive if $M(r_B) \geq 0$. Therefore, $\partial\hat{\Pi}/\partial X_1$ is negative - and consequently $X_1 = 0$ is optimal - if $M(r_B)$ is not too negative.

Alternatively, assume that $M(r_0) \geq 0$. Since $\mathbb{E}\{r_1 \mid X_0\} = r_0$, a sufficient condition for the first term to be non-positive is that

$$\mathbb{E}\left\{-G'(-(1-X_1)M(r_1))M(r_1) \mid X_0\right\} \leq -G'(-(1-X_1)M(r_0))M(r_0) \leq 0$$

which holds if $-G'(-(1-X_1)M(r_1))M(r_1)$ is a concave function of r_1, or equivalently (recalling that $M(r_1)$ is linear) that $uG''(u)$ is concave over the relevant range.

Proof of Lemma 3. Raising $Q(r_B)V_1$ by ε lowers $\Delta_B(r_N) - \Delta_G(r_N)$ uniformly by $\hat{\varepsilon} = \theta_0\varepsilon$. Recalling that $\Delta_B(r_N) - \Delta_G(r_N)$ is increasing in r_N, we distinguish between three cases:

- When $X^* = 0$ we have $\Delta_B(\phi(0)) - \Delta_G(\phi(0)) \leq 0$. This remains true when $\Delta_B(r_N) - \Delta_G(r_N)$ shifts downward, which implies that X^* remains equal to 0.

- When $0 < X^* < 1$ then $\Delta_B(r_N^*) - \Delta_G(r_N^*) = 0$. A downward shift of $\Delta_B(r_N) -$

$\Delta_G(r_N)$ raises the equilibrium posterior belief r_N^*. Since ϕ is decreasing in X this implies that X^* decreases.

- When $X^* = 1$ then $\Delta_B(r_0) - \Delta_G(r_0) \geq 0$. In this case, an upward shift of $\Delta_B(r_N) - \Delta_G(r_N)$ results in a value of r_N^* which is greater than, or the same as, before. This implies that X^* remains the same as before or decreases.

Proof of Proposition 5. If all other websites quote $p = 0$, then the future payoff of a website is independent of the price it sets, and is equal to $L_K(r_N)$. Therefore, quoting $p = 0$ is a best reply.

Proof of Proposition 6. An equilibrium with full precaution ($x = X = 1$) induces $r_N = \phi(1)$ and exists if and only if $v_0 \leq P_K(\phi(1))$. Consider now a symmetric equilibrium with $0 < x < 1$ and thus $X = x^K$, and $r_N = \phi(X)$. For any $p \leq v_0$ on the support of the equilibrium strategy we must have

$$p\left[x + (1-x)(1-F(p))\right]^{K-1} = x^{K-1}p_{\max}$$

where $p_{\max}$ is the upper bound of the support. This leads to a cumulative distribution function

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

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 setting $p = p_{\max}$. Thus, we have

$$F(p) = \frac{1 - x\left(\frac{v_0}{p}\right)^{\frac{1}{K-1}}}{1-x} \text{ on the interval } [x^{K-1}v_0, v_0].$$

The equilibrium payoff is then

$$L_K(r_N) + x_i x^{K-1} P_K(r_N) + (1-x_i)\, x^{K-1} v_0,$$

implying that an interior equilibrium verifies $v_0 = P_K(\phi(X_K^*))$. Given that $P_K(\phi(X))$ is decreasing in X, the solution to the equation $v_0 = P_K(\phi(X))$ exists in $(0,1)$ and is unique when $P_K(\phi(1)) < v_0 < P_K(\phi(0))$. This implies that the equilibrium exists and is uniquely defined for this range of values of v_0. Thus, we have

$$F(p) = \frac{1 - (X_K^*)^{\frac{1}{K}} \left(\frac{v_0}{p}\right)^{\frac{1}{K-1}}}{1 - (X_K^*)^{\frac{1}{K}}}$$

which gives $p_{\inf} = (X_K^*)^{\frac{K-1}{K}} v_0$.

Proof of Lemma 4. Straightforward computations show that $\partial\Phi/\partial X$ has the same sign as $(\theta_l - \theta_0)(1 - Y\lambda)(1 - Y)$, which proves (i), and that $\partial\Phi/\partial Y$ has the same sign as $(1 - X)(\theta_0 - \theta_l)(-1 + \lambda)$, which proves (ii). The proof of (iii) is immediate.

Proof of Proposition 7. *Existence and uniqueness of an equilibrium.* An equilibrium verifies $r_N^{**} = \Phi\left(X^{br}(r_N^{**}), Y^{br}(r_N^{**})\right)$. Therefore, a sufficient condition for a unique equilibrium to exist is that the correspondence $\Phi\left(X^{br}(r_N), Y^{br}(r_N)\right)$ has a unique fixed point. We have:

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

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

- For $r_N > r^M$, $X^{br}(r_N) = 1$ and $C_v'(Y^{br}(r_N)) = \max\{\Delta_G(r_N), 0\}$ is non-increasing, implying that $\Phi\left(X^{br}(r_N), Y^{br}(r_N)\right)$ is constant in r_N (recall that $\Phi(1, Y) = r_0$ for any Y).

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, which implies that r_N^{**} is continuous in v_0.

Proof of (i). Let us first show that r_N^{**} is non-decreasing in v_0 and X^{**} is non-increasing in v_0.

Suppose that $X^{**} = 0$ and let v_0 increase. Then X^{**} remains constant and r_N^{**} cannot decrease. To see why the latter part holds, assume that r_N^{**} decreases locally. Combined with $C_v'(Y^{**}) = \max\{\Delta_B(r_N^{**}), 0\}$, this would imply that Y^{**} is non-decreasing locally, while combined with $r_N^{**} = \Phi(0, Y^{**})$ and $\Phi(0, Y)$ decreasing with Y, it would imply that r_N^{**} is increasing locally - a contradiction. It follows that Y^{**} is non-increasing in the no-precaution region.

Suppose now that $X^{**} = 1$ and let v_0 increase. Then X^{**} remains constant and so does $r_N^{**} = \Phi(1, Y^{**}) = r_0$.

Suppose finally that $0 < X^{**} < 1$. We known from Proposition 1 that the result holds if $Y^{**} = 0$. Assume now that $C_v'(Y^{**}) = \Delta^M > 0$, and let v_0 increase. We have $r_N^{**} = r^M$

and

$$\frac{dr^M}{dv_0} = \frac{1}{(\lambda + \theta_0) Q' (r^M) V_1} > 0.$$

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 X^{**} decreases with v_0 if

$$\frac{dr^M}{dv_0} > \frac{\partial \Phi}{\partial Y} \frac{dY^{**}}{dv_0}. \tag{8}$$

Differentiating $C_v' (Y^{**}) = \Delta_B \left(r^M\right) = \theta_0 V_1 \left[Q \left(r^M\right) - Q (r_B)\right] - \theta_0 v_0$ with respect to v_0, we get

$$C_v'' (Y^{**}) \frac{dY^{**}}{dv_0} = \theta_0 \left(Q' (r_M) \frac{dr_M}{dv_0} V_1 - 1\right) = \theta_0 \left(\frac{Q' (r_M) V_1}{(\lambda + \theta_0) Q' (r^M) V_1} - 1\right) > 0,$$

This, combined with $\partial \Phi / \partial Y < 0$ for $X < 1$, implies that condition 8 holds.

We can therefore conclude that X^{**} is non-increasing in v_0. This, combined with the fact that $X^{**} = 1$ if v_0 is sufficiently small and $X^{**} = 0$ if v_0 is sufficiently large implies that there exist thresholds v^f and v^n such that $X^{**} = 1$ if and only if $v_0 \leq v^f$ and $X^{**} = 0$ if and only if $v_0 \geq v^n$. From $\Phi (1, Y) = r_0$ for all $Y \geq 0$, it follows that $X^{**} = 1$ if and only if $\Delta_G (r_0) \leq \Delta_B (r_0)$, which implies that the threshold v^f is the same as the corresponding threshold in the baseline model, i.e. $v^f = \psi^f V_1$. Moreover, from $\Phi (0, Y) \leq \phi (0)$ and the fact that $\Delta_G (r_N) - \Delta_B (r_N)$ is decreasing in r_N it follows that $\Delta_G (\phi (0)) > \Delta_B (\phi (0))$ whenever $\Delta_G (\Phi (0, Y)) > \Delta_B (\Phi (0, Y))$, which implies that v^n is less than or equal to the corresponding threshold in the baseline model, i.e. $v^n \leq \psi^n V_1$. The subsequent proof of part (ii) shows that $Y^{**} > 0$ when $v_0 = v^n$, which implies that the strict inequality $v^n < \psi^n V_1$ holds.

Proof of (ii). Consider first the case where $X^{**} = 1$. Then $r_N^{**} = r_0$, which implies that $\Delta_G (r_0) = (1 - \theta_0) v_0 > 0$. It then follows from $C_v' (Y^{**}) = \Delta_G (r_0)$ that $Y^{**} > 0$ and Y^{**} is increasing in v_0.

Suppose now that $X^{**} \in (0, 1)$. From $\Delta_G \left(r^M\right) = \Delta_B \left(r^M\right)$ it follows that

$$v_0 = \left[(\theta_0 + \lambda) Q \left(r^M\right) - \theta_0 Q (r_B) - \lambda Q (r_G)\right] V_1.$$

Using this and $\Delta^M = \Delta_B\left(r^M\right)$ we get

$$\Delta^M = \theta_0\left[\left(1-\theta_0-\lambda\right)Q\left(r^M\right)+\lambda Q\left(r_G\right)-\left(1-\theta_0\right)Q\left(r_B\right)\right]V_1,$$

which implies that $\Delta^M > 0$ because $Q\left(r^M\right) > Q\left(r_G\right) > Q\left(r_B\right)$. Therefore, $Y^{**} = \left(C_v'\right)^{-1}\left(\Delta^M\right) > 0$.

Let us now turn to the case where $X^{**}=0$. Then $Y^{**}=0$ if and only if $\Delta_B\left(\phi\left(0\right)\right)\leq 0$ which writes as

$$v_0 \geq \bar{v} \equiv \left(Q\left(\phi\left(0\right)\right)-Q\left(r_B\right)\right)V_1.$$

From $Q\left(\phi\left(0\right)\right)-Q\left(r_B\right) > \left(\theta_0+\lambda\right)Q\left(\phi\left(0\right)\right)-\theta_0 Q\left(r_B\right)-\lambda Q\left(r_G\right)=\psi^n$, it follows that $\bar{v} > \psi^n V_1$ and consequently $\bar{v} > v^n$.

Finally, consider the way Y^{**} is affected by v_0. In the full precaution regime, $C_v'\left(Y^{**}\right) = \Delta_G\left(r_0\right)$ is increasing in v_0, which implies that Y^{**} is increasing in v_0. In the no precaution regime, Y^{**} is non-increasing in v_0 because $r_N^{**} = \Phi\left(0, Y^{**}\right)$ is non-decreasing in v_0 and Φ is decreasing in Y. Finally, in the partial precaution region, we have

$$C_v'\left(Y^{**}\right) = \lambda Q\left(r_G\right)-\lambda Q(r^M)+\left(1-\theta_0\right)v_0 = \theta_0 Q(r^M)-\theta_0 Q(r_B)-\theta_0 v_0,$$

which yields

$$Q(r^M) = \frac{v_0+\lambda Q\left(r_G\right)+\theta_0 Q(r_B)}{\lambda+\theta_0}$$

and, therefore,

$$C_v'\left(Y^{**}\right) = \theta_0\left(\frac{v_0+\lambda Q\left(r_G\right)+\theta_0 Q(r_B)}{\lambda+\theta_0}-Q\left(r_B\right)-v_0\right).$$

Hence, Y^{**} is increasing in v_0 in the partial precaution region.

Proof of Proposition 8. The proof for the fact that verification reduces the equilibrium level of precaution is immediate in all cases except when $0 < X^{**} < 1$ and $r_N^{**} = r^M$. However, in that case, straightforward computations show that $1-X^{**} = \left(\frac{1-\lambda Y^{**}}{1-Y^{**}}\right)\left(1-X^*\right) > 1-X^*$.

If $C_v'\ \left(Y\right)$ decreases uniformly, then $Y^{br}\left(r_N\right)$ (weakly) increases uniformly and $\Phi\left(X, Y^{br}\left(r_N\right)\right)$ (weakly) decreases. Suppose first that $Y^{**} > 0$ and $X^{**} = 0$. Then a marginal reduction in C_v' leads to a decrease in $r_N^{**} = \Phi\left(0, Y^{br}\left(r_N^{**}\right)\right)$ and an increase in Y^{**}. Suppose now that $Y^{**} > 0$ and $X^{**} > 0$. In this case, a marginal reduction in C_v' leaves $r_N^{**} = r^M$ unchanged and raises $Y^{**} = Y^{br}\left(r^M\right)$. As $\Phi\left(X^{**}, Y^{**}\right) = r^M$, the level of precaution X^{**} must decrease.

References

[1] Acquisti, A., Taylor, C., and L. Wagman (2016),“The Economics of Privacy,” *Journal of Economic Literature,* 54, 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] Bergemann, D. and A. Bonatti (2015), “Selling Cookies,” *American Economic Journal: Microeconomics*, 7, 259-294.

[4] Bloch, F., and G. Demange (2018), “Taxation and Privacy Protection on Internet Platforms," *Journal of Public Economic Theory*, 20 (1), 52-66 .

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

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

[7] Fudenberg, D. and J. Tirole (1986), “A "Signal-jamming" Theory of Predation,” *The RAND Journal of Economics*, 17(3), 366-376.

[8] Holmström, B. (1999), “Managerial Incentives: A Dynamic Perspective,” *The Review of Economic Studies*, 66(1), 169-182.

[9] 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.

[10] Lefouili, Y. and Y.L. Toh (2018), “Privacy Regulation and Quality Investment,” TSE Working Paper 17-795.

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

[12] Riordan, M.H. (1985) “Imperfect Information and Dynamic Conjectural Variations,” *The RAND Journal of Economics*,16(1), 41-50.

[13] RiskIQ (2016), “RiskIQ’s 2016 Malvertsing Report,”, available at https://www.riskiq.com/infographic/riskiqs-2016-malvertising-report/

[14] Rothschild, M., and J. E. Stiglitz (1971), "Increasing Risk II: Its Economic Consequences," *Journal of Economic Theory*, 3, 66-84.

[15] Spiegel, Y. (2013), "Commercial Software, Adware, and Consumer Privacy," *International Journal of Industrial Organization*, 31, 702-713.

[16] Toh, Y.L. (2018), "Incentivizing Firms to Protect Consumer Data: Can Reputation Play a (Bigger) Role?", working paper.

[17] U.S. Senate Committee on Homeland Security and Governmental Affairs (2014), "Online Advertising and Hidden Hazards to Consumer Privacy and Data Privacy", hearing before the permanent subcommittee on investigations, available at https://www.gpo.gov/fdsys/pkg/CHRG-113shrg89686/pdf/CHRG-113shrg89686.pdf.

Appendix for Online Publication

1 Equilibrium analysis under opt out

Our equilibrium analysis in the scenario in which customers are granted an opt-out option focused on the case where consumers never opt out in the first period and opt out in the second period if and only if they have a bad experience in the first period. In this section, we still focus on the case when consumers never opt out in the first period but allow for all possible scenarios in the second period.

As a preliminary remark, note that when $\bar{r} < r_B$, the equilibrium is not affected by the possibility of opting out as consumers never choose this option. Similarly, in the case where $\bar{r} > \phi(0)$, consumers always opt out in the second period. Therefore, the website always sells customer information to third parties in the first period, i.e. $X = 0$. We focus in what follows on the scenario in which $r_B < \bar{r} \leq \phi(0)$.

Under this assumption, it cannot be the case that all consumers opt out after a neutral experience because this would imply that $X = 0$ and $r_N = \phi(0) < \bar{r}$. Thus, in equilibrium, consumers should opt in with a positive probability, denoted $P_{\emptyset}$, after a neutral experience. As we shall see, this probability may be less than 1. We need to distinguish between two types of equilibria depending on whether consumers opt out or not after a good outcome is observed. In the main text, we considered the scenario in which consumers opt out after a bad experience but not after a good one (i.e. $r_B < \bar{r} < r_0$). We now investigate the existence of an equilibrium in which consumers opt out after a good experience as well. This is the case when $r_0 < \bar{r} < \phi(0)$.

Note first that an equilibrium requires that $r_N \geq \bar{r}$ because some consumers must opt in. This condition is trivially verified if $\bar{r} \leq \phi(1)$, in which case we can replicate the proof of Proposition 1. Defining

$$\begin{aligned}
\bar{v}^f &= (\lambda + \theta_0)\left[Q(\phi(1))V_1 - \bar{Q}\bar{V}_1\right] \\
\bar{v}^n &= (\lambda + \theta_0)\left[Q(\phi(1))V_1 - \bar{Q}\bar{V}_1\right],
\end{aligned}$$

we get the following result.

Proposition 10 *Assume that $r_0 < \bar{r} \leq \phi(1)$. Then, a unique equilibrium exists. Moreover, there exist thresholds $\bar{v}^f$ and $\bar{v}^n$ such that:*

(i) the website provides full precaution ($X^{opt} = 1$)if $v_0 \leq \bar{v}^f$;

(ii) the website provides partial precaution ($0 < X^{opt} < 1$) if $\bar{v}^f < v_0 < \bar{v}^n$;

(iii) the website provides no precaution ($X^{opt} = 0$) if $v_0 \geq \bar{v}^n$.

Proof. Since $r_\emptyset \geq \phi(1) \geq \bar{r}$, it is sufficient to replace $Q(r_G)V_1$ and $Q(r_B)V_1$ by $\bar{Q}\bar{V}_1$ in the proof of Proposition 1 to get this result. ■

The website would not offer the opt-out option on a voluntary basis if $[\lambda Q(r_G) + \theta_0 Q(r_B)]V_1 > (\lambda + \theta_0)\bar{Q}\bar{V}_1$. In this case, notice that $\bar{v}^f$ and $\bar{v}^n$ are greater than $\psi^f V_1$ and $\psi^n V_1$, respectively, and that $X^{opt} > X^*$ in the range of partial precaution. Thus, the effect of a mandatory opt-out policy is to raise (weakly) the level of precaution.

Let us now turn to the scenario in which $\phi(1) < \bar{r} < \phi(0)$. The analysis in the case $r_0 < \bar{r} < \phi(0)$ no longer applies because the level of precaution cannot be too large in equilibrium. Let us define $\bar{X}$ as the unique solution to $\phi\left(\bar{X}\right) = \bar{r}$. Then, in any equilibrium we must have $X^{opt} \leq \bar{X}$. Notice that the equilibrium level of precaution X^{opt} decreases in v_0. Since for sufficiently large values of v_0 we have $X^{opt} = 0$, there must exist some critical level $\bar{v}^o$ such that $X^{opt} < \bar{X}$ if and only if $v > \bar{v}^o$. In this range, the equilibrium is similar to the one when $r_0 < \bar{r} \leq \phi(1)$. However, for lower values of v_0, the equilibrium must be such that $X^{opt} = \bar{X}$ and consumers randomize between opting in and opting out.

Proposition 11 *Assume that $\phi(1) < \bar{r} < \phi(0)$. Then, a unique equilibrium exists. Moreover, there exists a threshold $\bar{v}^o$ such that:*

(i) the website provides partial precaution and $X^{opt} = \bar{X}$ if $v_0 \leq \bar{v}^o$;

(ii) the website provides partial precaution and $X^{opt} \in \left(0, \bar{X}\right)$ if $\bar{v}^o < v_0 < \bar{v}^n$;

(iii) the website provides no precaution ($X^{opt} = 0$) if $v_0 \geq \bar{v}^n$.

Proof. The result is the same as before when $v_0 > \bar{v}^o$ where $\bar{v}^o$ is defined by

$$\bar{v}^o = (\lambda + \theta_0)\left[Q\left(\phi\left(\bar{X}\right)\right)V_1 - \bar{Q}\bar{V}_1\right]$$

For smaller values of v_0, we have $X = \bar{X}$ and $r_\emptyset = \bar{r}$ in equilibrium, and

$$P_\emptyset = \frac{v_0}{\bar{v}^o} < 1.$$

It then suffices to replace $Q(r_G)V_1$ and $Q(r_B)V_1$ by $\bar{Q}\bar{V}_1$ in the proof of Proposition 1. ■

2 Transparency in the multi-homing scenario

Consider first *ex post* transparency and an equilibrium with $0 < x < 1$, $X = x^K$, $r_N = \phi(X)$. For any $p \leq v_0$, replicating the reasoning used to derive the partial precaution equilibrium in Proposition 6, we must have a cumulative distribution function for prices given by

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

Consider now the choice of precaution, and suppose that a website deviates to x_i and $p \leq v_0$. Then the website's expected payoff is $L_K(r_N) + x_i x^{K-1} P_K(r_N) + (1 - x_i) x^{K-1} v_0$ where $r_N = x_i x^{K-1}$. The equilibrium level of precaution x is therefore such that

$$P_K(\phi(X)) - v_0 = (X P_K'(\phi(X)) + L_K'(\phi(X)))\, \phi'(X) > 0,$$

which yields less precaution than X_K^*.

Consider now the case of *ex ante* transparency where websites may refuse to sell. In this scenario, all websites observe all levels of precaution before setting prices. Refusal to sell is possible so that the strategy is a public choice of x_i, followed by a private choice of a probability to refuse to sell y_i and a price p_i. The total probability of not selling is then $z_i = x_i + (1 - x_i) y_i$. We focus on symmetric equilibria $x_i = x$.

Suppose first that $P_K(\phi(x^K)) > v_0$. In this case the payoff of a website choosing $x_i \neq x$ is

$$L_K(r_N) + z_i z^{K-1} P_K(r_N) + (1 - z_i)\, p\, (z + (1 - z)(1 - F(p)))^{K-1} \text{ with } z_i \in (x_i, 1).$$

Assume that the market anticipates $z^K = X_K^* > x^K$. Then it follows from above that the equilibrium obtains at $z = x + (1 - x) y = (X_K^*)^{\frac{1}{K}}$. Hence, when $X < X_K^*$, the equilibrium probability of sale and the distribution of prices is the same as without transparency.

Suppose now that $P_K(\phi(x^K)) < v_0$. Then the symmetric equilibrium distribution of prices is as above. Suppose a website deviates to $x_i \neq x$ while preserving $P_K(r_N) < v_0$. Then all websites set $p \leq v_0$. It can be seen that the lower bound of the support of the price must be the same for all firms (this is because $K - 1$ firms set higher prices than K firms so that there would a contradiction if a smaller number of firms were to charge lower prices). However, one website may have a mass point at v_0 or may not charge high prices.

We consider first the scenario in which the deviating website i puts a mass at v_0.

Then, on the support of prices the following indifference condition must hold for website i: $p\left(x+(1-x)(1-F(p))\right)^{K-1}=x^{K-1}v_0$, and the following condition must hold for the other websites (where we use the fact that other websites can set a price below but arbitrarily close to v_0):

$$p\left(x+(1-x)(1-F(p))\right)^{K-2}\left(x_i+(1-x_i)(1-F_i(p))\right)=\left(x_i+(1-x_i)(1-F_i(v_0))\right)x^{K-2}v_0.$$

This yields

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

and

$$\left(x_i+(1-x_i)(1-F_i(p))\right)=\left(x_i+(1-x_i)(1-F_i(v_0))\right)\left(\frac{v_0}{p}\right)^{\frac{1}{K-1}}.$$

This equilibrium holds provided that the supports are the same, that is, when $x=x_i+(1-x_i)(1-F_i(v_0))$, which holds for $x_i<x$. The payoff of the deviating website with $x_i<x$ is then

$$L_K\left(\phi\left(x_ix^{K-1}\right)\right)+x_ix^{K-1}P_K\left(\phi\left(x_ix^{K-1}\right)\right)+(1-x_i)x^{K-1}v_0.$$

Suppose now that $x_i>x$. If $K=2$ we can apply the previous analysis reverting the role of i and the other website. As the latter sets a mass point at v_0 such that $x+(1-x)(1-F(v_0))=x_i$, the payoff of website i is then given by

$$L_K\left(\phi(x_ix)\right)+x_ixP_K\left(\phi(x_ix)\right)+(1-x_i)x_iv_0.$$

Assume now that $K>2$. In this case we investigate an equilibrium where the deviating website does not charge high prices but only prices between $p_{\min}$ and $\hat{p}<v_0$. The following conditions must hold. First, on the interval $(p_{\min},\hat{p})$ we must have

$$p\left(x+(1-x)(1-F(p))\right)^{K-1}=\hat{p}\left(x+(1-x)(1-F(\hat{p}))\right)^{K-1},$$

and

$$p\left(x+(1-x)(1-F(p))\right)^{K-2}\left(x_i+(1-x_i)(1-F_i(p))\right)=x_ix^{K-2}v_0.$$

Second, on the interval $(\hat{p},v_0)$ we must have

$$p\left(x+(1-x)(1-F(p))\right)^{K-2}x_i=x_ix^{K-2}v_0.$$

Thus, we get:

$$F\left(p\right)=\frac{1-x\left(\frac{v_0}{p}\right)^{\frac{1}{K-2}}}{1-x}\text{ on }\left[\hat{p},v_0\right]$$

Then, we must have $x+\left(1-x\right)\left(1-F\left(p\right)\right)=\left(\hat{p}/p\right)^{\frac{1}{K-1}}\left(v_0/\hat{p}\right)^{\frac{1}{K-2}}x^{K-1}$ on $\left[p_{\min},\hat{p}\right]$, and $x_i+\left(1-x_i\right)\left(1-F_i\left(p\right)\right)=\left(\hat{p}/p\right)^{\frac{1}{K-1}}x_ix^{K-2}$ on $\left[p_{\min},\hat{p}\right]$. For the support to be the same we need that $\left(v_0/\hat{p}\right)^{\frac{1}{K-2}}=x_i/x>1$ or $x_i>x$, which is the case. To complete the equilibrium we verify that the deviating website setting $p>\hat{p}$ would obtain

$$L_K\left(\phi\left(x_ix^{K-1}\right)\right)+x_ix^{K-1}P_K\left(\phi\left(x_ix^{K-1}\right)\right)+\left(1-x_i\right)px^{K-1}\left(\frac{v_0}{p}\right)^{\frac{K-1}{K-2}},$$

which decreases with p and thus is lower than the payoff at $\hat{p}$:

$$L_K\left(\phi\left(x_ix^{K-1}\right)\right)+x_ix^{K-1}P_K\left(\phi\left(x_ix^{K-1}\right)\right)+\left(1-x_i\right)\hat{p}x_i^{K-1}.$$

To summarize we find that the deviation payoff is given by:

$$\begin{aligned}
L_K\left(\phi\left(X_K^*\right)\right)+X_K^*P_K\left(\phi\left(X_K^*\right)\right)+\left(1-\left(X_K^*\right)^{\frac{1}{K}}\right)\left(X_K^*\right)^{\frac{K-1}{K}}v_0\text{ if }x_i &< x\text{ and }x^K\leq X_K^*\ ,\\
L_K\left(\phi\left(x_ix^{K-1}\right)\right)+x_ix^{K-1}P_K\left(\phi\left(x_ix^{K-1}\right)\right)+\left(1-x_i\right)x^{K-1}v_0\text{ if }x_i &< x\text{ and }x^K>X_K^*,\\
L_K\left(\phi\left(x_ix^{K-1}\right)\right)+x_ix^{K-1}P_K\left(\phi\left(x_ix^{K-1}\right)\right)+\left(1-x_i\right)v_0x^{K-2}x_i\text{ if }x_i &> x\text{ and }x^K>X_K^*.
\end{aligned}$$

The left derivative is negative at $x^K>X_K^*$ and is strictly smaller than the right derivative at that point. The reason for this kink in the payoff is that committing to $x_i>x$ induces a strategic effect that leads the other websites to raise their prices. Hence, the only candidate for a pure precaution equilibrium is at $x^K=X_K^*$ and it exists only if

$$\left(X_K^*P_K'\left(\phi\left(X_K^*\right)\right)+L_K'\left(\phi\left(X_K^*\right)\right)\right)\left(X_K^*\right)^{\frac{1}{K}}\phi'\left(X_K^*\right)+\left(1-\left(X_K^*\right)^{\frac{1}{K}}\right)v_0\leq 0$$

and, therefore, only if X_K^* is large enough.