



HAL
open science

Privacy regulation and quality-enhancing innovation

Yassine Lefouili, Leonardo Madio, Ying Lei Toh

► **To cite this version:**

Yassine Lefouili, Leonardo Madio, Ying Lei Toh. Privacy regulation and quality-enhancing innovation. Journal of Industrial Economics, 2024, 72 (2), pp.662-684. 10.1111/joie.12374 . hal-04774302

HAL Id: hal-04774302

<https://hal.science/hal-04774302v1>

Submitted on 8 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WORKING PAPERS

N° 795

July 2023

Privacy Regulation and Quality-Enhancing Innovation

Yassine Lefouili, Leonardo Madio and Ying Lei Toh

Privacy Regulation and Quality-Enhancing Innovation*

Yassine Lefouili[†] Leonardo Madio[‡] Ying Lei Toh[§]

July 2, 2023

We analyze how a privacy regulation taking the form of a cap on information disclosure affects quality-enhancing innovation incentives by a monopolist—who derives revenues solely from disclosing user data to third parties—and consumer surplus. If the share of privacy-concerned users is sufficiently small, privacy regulation has a negative effect on innovation and may harm users. However, if the share of privacy-concerned users is sufficiently large, privacy regulation has a positive effect on innovation. In this case, there is no trade-off between privacy and innovation and users always benefit from privacy regulation.

Keywords: Privacy Regulation, Data Disclosure, Innovation.

JEL Classification: D83, L15, L51.

An earlier version of this paper circulated under the title “Privacy and Quality” and “Privacy Regulation and Quality Investment”. We are grateful to the Editor, Alessandro Bonatti, and two anonymous referees for their insightful comments and suggestions. We would also like to thank Alessandro Acquisti, Rabah Amir, Anna d’Annunzio, Olivier Bonroy, Marc Bourreau, Grazia Cecere, Romain De Nijs, Sebastian Dengler, Jocelyn Donze, Nestor Duch-Brown, Andres Hervas-Drane, Bruno Jullien, Michael Kummers, Marc Lebourges, Matthieu Manant, Bertin Martens, Dimitrios Mavridis, Thierry Penard, Markus Reisinger, Mike Riordan, Rahim Talal, Curtis Taylor, and Jean Tirole for very useful comments and discussions. We also thank our audiences at various seminars and conferences for helpful comments and discussions. Yassine Lefouili gratefully acknowledges the financial support of the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 670494), the Agence Nationale de la recherche (ANR) under grant ANR-17-EURE-0010 (Investissements d’Avenir program), the IDEI-Orange research partnership, and the TSE Digital Center. Leonardo Madio acknowledges financial support from Unicredit & Universities Foscolo Europe fellowship. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

[†]Toulouse School of Economics, University of Toulouse Capitole, 1, Esplanade de l’Université, 31080 Toulouse, Cedex 06, France. Email: yassine.lefouili@tse-fr.eu.

[‡]Department of Economics and Management, University of Padova, Via del Santo, 33, 35123 Padova, Italy. Email: leonardo.madio@unipd.it. Other Affiliation: CESifo.

[§]Federal Reserve Bank of Kansas City. Email: yinglei.toh@kc.frb.org

I. Introduction

In recent years, data protection policies requiring firms to obtain users’ informed consent for the collection and use of their data have emerged in various jurisdictions. Notable examples are the EU General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and the Brazilian General Data Protection Law (LGPD).¹ Stronger obligations are also imposed in the European Union to large “gatekeeping firms” by the EU Digital Markets Act (DMA), which prohibits tracking end-users outside of the actual online service for targeted advertising in the absence of explicit user consent.

A concern raised by many scholars and practitioners is that restrictions to data collection and usage may hamper data-driven investments in quality-enhancing innovation. For example, restricting data collection for ad targeting purposes can lower advertising effectiveness and, consequently, ad revenues [Goldfarb and Tucker, 2011, Johnson et al., 2020, Rafeian and Yoganarasimhan, 2021, Goldberg et al., 2022]. If the firm’s operations are primarily ad-financed, this reduction in revenues may dampen its incentives to invest and innovate [Castro, 2010, Athey, 2014, Kircher and Foerderer, 2022].²

In this paper, we study the potential trade-off between privacy and quality-enhancing innovation. To this end, we consider a setup that resembles closely the business model of many dominant tech firms (particularly social media platforms) in which users “pay with [their] data and [their] attention.”³ In our model, a monopolist offers a free service to users and derives revenues by disclosing user data to third parties (e.g., for targeting ads). Prior to interacting with the users, the monopolist sets its privacy policy (a disclosure level)⁴ and decides how much to invest in innovation (e.g., quality of the sharing tools, new features). The firm’s innovation and disclosure levels are observed by users, who then decide whether to use the service (e.g., whether to create a social media account) and how much information to provide when using it (e.g., how much to reveal about themselves on their social media page; how many posts to write and photos to share). Users derive higher gross utility when the firm’s innovation is higher (e.g., when a social media platform provides more features and better sharing tools) and when they provide more

¹A stream of recent studies has studied the intended and unintended effects of privacy regulations, providing an evaluation of the EU GDPR (see e.g., Johnson [2022]).

²For example, Kircher and Foerderer [2022] show that when Google banned targeted advertising in Android children’s games in 2019, advertisement-dependent firms reacted by reducing the release of feature updates. They explain the reduction in the incentives to provide updates (and fix bugs) as a consequence of reduced revenues from advertising and the lack of an alternative monetization channel.

³The quote was taken from Facebook co-founder Chris Hughes’ opinion piece in the New York Times which appeared on May 9, 2019: <https://www.nytimes.com/2019/05/09/opinion/sunday/chris-hughes-facebook-zuckerberg.html>

⁴For the sake of exposition, the disclosure level is defined in our model as the share of third parties to which the monopolist discloses personal information. However, as discussed later, it can also be interpreted as an inverse measure of other (self-imposed) restrictions on the disclosure of data (e.g., regarding the type or share of data disclosed).

information. Further, they perceive innovation and information as complements (e.g., the better a platform's sharing tools, the more information users share). Although the monopolist's service is free, users may incur idiosyncratic privacy costs, which could be due to their intrinsic preference for privacy or potential adverse market outcomes (e.g., price discrimination and unwanted ads) that they may face as a result of the firm's disclosure of their personal information. We capture heterogeneity in privacy consciousness by considering that users belong to two groups. The first group of users are privacy-unconcerned and do not incur any privacy costs.⁵ The second group of users are privacy-concerned and incur (heterogeneous) privacy costs.

We analyze the desirability for users of a privacy regulation that takes the form of a (binding) cap on data disclosure. The cap could correspond to a set of restrictions on the purposes for which data may be disclosed or the types of third parties with whom data may be shared.⁶ Besides directly reducing users' privacy costs, a disclosure cap also generates a change in the firm's innovation incentives. Thus, the desirability for users of the cap depends (partly) on its impact on the firm's innovation incentives.

A cap on disclosure affects the monopolist's incentives to invest in innovation by altering its disclosure revenues, and therefore its gains from innovation. The monopolist invests in innovation to attract more users (extensive margin effect) and/or to induce current users to share more on the platform (intensive margin effect); both of which increase the stock of information it can monetize. A cap on disclosure reduces the firm's gain from monetizing an additional unit of information. Consequently, for a given number of users, the cap lowers the firm's incentives to induce users to share more data by investing in innovation. However, a cap on disclosure has another effect that goes in the opposite direction. It boosts privacy-concerned users' demand for the firm's service, which increases its marginal benefit from investing in innovation. This effect is stronger the larger the share of privacy-concerned users in the population and dominates the negative effect of the privacy cap on innovation when the share of privacy-concerned users is sufficiently large. Thus, the overall effect of a privacy cap on innovation depends on the size of the share of privacy-concerned users in the market.

In a market with a small share of privacy-concerned users, a disclosure cap will lead to lower innovation. As a result, the regulator faces a trade-off between privacy and innovation. Because of this trade-off, the cap is desirable for users only if the number of privacy-concerned users is not 'too' small as in this case the positive effect of privacy

⁵These users may also represent that by revealed preferences do not consider privacy violations as a cost, although they tend to state otherwise.

⁶U.S. regulations such as the Gramm-Leach-Bliley Act and the Fair Credit Reporting Act forbid the sharing of consumer information with non-affiliated third parties. Many countries (e.g., EU countries) also impose restrictions on international transfers of data, allowing for them only if the foreign-based entity has an adequate level of data protection.

regulation (i.e., the decrease in privacy costs) more than compensates for the reduced innovation level. However, if the share of privacy-concerned users is sufficiently large, the extensive margin becomes more salient to the firm, and the cap increases the firm's incentives to invest in innovation. In this case, there is *no privacy-innovation trade-off* and consumer surplus is unambiguously higher under the cap.

We extend the model along several dimensions and show that our main results carry out in a wide range of circumstances. First, we consider heterogeneous third parties that induce different privacy costs (reflecting differences in how data are used by them). Second, we introduce network effects, whose presence is relevant in the context of social media firms. Third, we consider the case in which disclosure of personal information also delivers direct benefits to users.

Our findings provide insights into the factors that a regulator should consider when deciding whether to set restrictions on a dominant firm's data disclosure policy. If a large share of users are (essentially) privacy-unconcerned,⁷ restrictions on data disclosure are likely to stifle innovation. In this case, the regulator needs to account for the negative effect of the cap on innovation when assessing its desirability for users. On the contrary, if a large share of users is privacy-concerned, the regulator can be more confident about the desirability of the cap for users as it may lead to more, rather than less, innovation.

The rest of the paper is structured as follows. Section II discusses the related literature. Section III presents the model setup. Section IV examines the users' participation and information provision decisions. Section V analyzes the impact of a disclosure cap on innovation and consumer surplus. Section VI presents the extensions. Section VII concludes. All omitted proofs can be found in Appendix A.

II. Related literature

There is a growing pool of literature on the economics of privacy. Our paper lies in the intersection of several issues that researchers have studied: users' information disclosure decisions (e.g., Abrardi et al. [2021], Abrardi and Cambini [2022]), firms' information exploitation/sharing decisions (e.g., Casadesus-Masanell and Hervas-Drane [2015], Dimakopoulos and Sudaric [2018]), and the effects of privacy regulation/taxation on welfare (e.g., Bloch and Demange [2018], Bourreau et al. [2018])

In most of the existing research on consumer privacy, users' decision to share personal data with the firm or otherwise allow the firm to collect their data is modeled as a binary

⁷Users may however respond to changes in innovation and/or privacy levels by changing the amount of information they provide to the firm.

choice, that is, users either decide to share or not to share their personal data. In contrast, we consider a scenario where users decide how much to share with the firm, trading off higher privacy for higher utility from the firm’s service. Our modeling approach is closest to that of Casadesus-Masanell and Hervas-Drane [2015].⁸ Unlike Casadesus-Masanell and Hervas-Drane [2015], we endogenize the firm’s innovation choice. This allows us to derive the impact of a disclosure cap on investment in innovation.

Several papers have studied the trade-off users face when deciding whether/how much personal information to disclose to the firm (or to allow the firm to collect). On the one hand, sharing more information with the firm enables users to obtain higher utility from the firm’s product or service, via better product matching or recommendation (e.g., Bloch and Demange [2018], Hidir and Vellodi [2021], Ichihashi [2020]) or a greater level of product or service personalization (e.g., Bourreau et al. [2018], Casadesus-Masanell and Hervas-Drane [2015]). On the other hand, users incur privacy costs such as higher prices (e.g., in Hidir and Vellodi [2021], Ichihashi [2020]) or higher disutility from the firm’s exploitation of their information (e.g., Bloch and Demange [2018]) when they disclose more information to the firm. We build a model that features a similar trade-off: users experience higher utility when they disclose more information to the firm, but they also face higher privacy costs because of their preference for privacy.

Our analysis throws light on the potential privacy-innovation trade-off and identifies conditions under which a disclosure cap raises the amount of data-driven innovation, implying that the privacy-innovation trade-off may not exist. Our findings echo those of Anderson [2007] for which an advertising cap may either decrease or increase the quality of free-to-air television. However, his findings are driven entirely by the impact of the advertising cap on the extensive margin effect of quality investment, as he considers only the (binary) participation decision of users. Our results depend also on how innovation affects the amount of information provided by users (i.e., the intensive margin effect).

Conti and Reverberi [2021] investigate how a privacy regulation that imposes an opt-in policy affects a firm’s quality in a model featuring price discrimination. The main trade-off users face when deciding whether to share information is that sharing information leads to better consumption experience but allows the firm to offer them personalized prices based on their willingness to pay. Unlike us, they focus on the differences between an opt-in and opt-out privacy regime and find that an opt-out regime may deliver a higher quality level than an opt-in regime if there are strong complementarities between information and quality. We differ from their paper in that we focus on the impact of a different regulatory tool (a cap on information disclosure). Moreover, whereas they find that an increase in

⁸In both models, users decide how much information to share with the firm, but not all information shared by users is disclosed by the firm. In our setting, we consider different revenue models: the firm derives revenues solely from disclosure, whereas it can also charge a positive price in Casadesus-Masanell and Hervas-Drane [2015].

product quality is a necessary condition to raise consumer surplus, in our setting consumer surplus might increase with privacy regulation even if innovation decreases.⁹

The implications drawn from our analysis also complement the findings in the empirical literature that examines the interplay between privacy regulations and data-driven innovation. Goldfarb and Tucker [2012] analyze several empirical studies—in the health-care (see Miller and Tucker [2009, 2011b,a]) and the online advertising (see Goldfarb and Tucker [2011]) sectors—and find that privacy regulations may raise the costs and/or lower the benefits associated with data-driven innovation, hence weakening firms’ investment incentives. Our work shows, in addition, that a privacy regulation can affect the level of innovation even when data is not a direct input for innovation.

More generally, our paper also relates to the literature on the welfare effects of regulating monopolist’s decisions [Sheshinski, 1976, Besanko et al., 1987]. Sheshinski [1976] studies how price regulation alters a monopolist’s decision to invest in quality. He shows that, provided that the regulated price is binding for the firm, price regulation has a negative effect on quality and a positive effect on quantity sold in the market. Unlike Sheshinski, our framework features an *implicit price* that is heterogeneous across users because users have a heterogeneous privacy cost. We show that privacy regulation induces the firm to lower its quality-enhancing innovation (as in Sheshinski’s paper) if the share of privacy-concerned users is sufficiently low. However, we also show that regulation spurs quality-enhancing innovation if the share of privacy-concerned consumers is large enough.

Besanko et al. [1987] study the case in which a monopolist offers a menu of prices and qualities to consumers that are heterogeneous in their willingness to pay for quality. Because of asymmetric information, the monopolist distorts quality to segment the market, even by excluding some consumers from the market. Price regulation reduces the incentive to segment the market and has asymmetric effects on the quality offered to consumers, with some facing a lower quality and others facing a higher quality. As a result, the effect of price regulation is heterogeneous across consumers. Whereas our paper and Besanko et al. show that regulation has ambiguous effects on quality, the main mechanisms are different. In our setting, the monopolist does not engage in second-degree price and quality discrimination. As a result, regulation of the (implicit) price does not alter the incentives of the firm to segment the market. Instead, it generates a demand-boosting effect that may dominate or be dominated by the more standard per-user revenue effect.

⁹Also our analysis is more appropriate for understanding the effects of privacy regulation in the social media market where there is a zero-price for users and therefore there is no potential for price discrimination.

III. Model setup

Consider a firm that offers a service to a unit mass of users at a price of zero. The firm derives revenues from disclosing its customers' personal information to (a subset of) homogeneous third parties (e.g., advertisers or third-party apps) that are uniformly distributed over the interval $[0, 1]$. The firm can choose the level of quality-enhancing innovation $q \geq 0$, and a disclosure level $d \in [0, 1]$, which defines the extent to which personal information is shared with third parties. More precisely, information is disclosed to the third parties located in the interval $[0, d]$ and not disclosed to those in $(d, 1]$.¹⁰

Users. Users (or, interchangeably, consumers) face a trade-off when deciding the amount of personal information they share with the firm, which we denote as $x \in [0, 1]$.¹¹ When users share more information with the firm, they obtain higher utility from its service but also suffer higher privacy-related utility losses arising from the firm's disclosure of consumer data to third parties. These utility losses reflect the users' preference for privacy, which may arise because they value privacy intrinsically (e.g., as a right) or because they may face potential adverse market-mediated outcomes (e.g., price discrimination).¹²

The benefits that a consumer derives from using the firm's service comprise two components. The first component is the utility that the consumer obtains from using the service regardless of whether she shares information about herself. This data-independent benefit component depends on v , the baseline utility level or "intrinsic value" of the firm's services—one can think of this as the utility generated by the firm's minimal viable product—and the firm's level of (quality-enhancing) innovation q . We capture the data-independent benefit of innovation using a linear function βq , with $\beta > 0$. The second component of the benefits a consumer obtains from using the firm's service depends on the amount of data the consumer shares. Sharing data increases the consumer's utility

¹⁰Note that our analysis is equivalent to the case in which the firm discloses a share $d \in [0, 1]$ of the personal information provided by each consumer to all third parties, instead of assuming that it discloses all the personal information provided by users to a subset $[0, d]$ of third parties. More generally, a lower disclosure level can be interpreted either as more (self-imposed) restrictions on the type/share of data that can be disclosed to third parties and/or more restrictions on the set of third parties that the data can be shared with.

¹¹The interpretation of the consumer's input, x , as the amount of information that a consumer provides to the firm is appropriate when considering a matching website, where each user decides on the preferences she reveals, or in the context of services involving user-generated content, where each user decides on the amount of content to share (e.g., on a social media platform). This interpretation may, however, be less suited to other forms of Internet services, such as email and search. In these contexts, users do not directly supply information to the firm; instead, information is generated as a result of their use of the firm's service. The relevant consumer choice variable (and hence the appropriate interpretation of x) is therefore *usage intensity*, rather than information provision level. Correspondingly, α , which formerly captured the marginal cost of providing information, can be interpreted more generally as the marginal opportunity cost of using the firm's service.

¹²Data disclosure can also generate positive effects for users (e.g., better-targeted ads or a more customized news feed). We consider this case in an extension.

from using the firm’s service regardless of whether the firm innovates, though the benefit of data sharing diminishes with the amount of data shared; specifically, we capture this benefit by $x - \frac{x^2}{2}$. Sharing data generates additional benefit γxq to consumers when the firm invests in quality-enhancing innovation, where γ captures the interaction between innovation and information from the users’ perspective. We assume that γ is positive so that innovation and information are complements from the users’ perspective.¹³

When using the firm’s service, a consumer also faces some costs. First, a consumer of type θ who provides an amount of personal information x incurs individual privacy cost θdx .^{14,15} The consumer’s type θ captures their preference for privacy, with a higher θ indicating a stronger privacy preference. We assume that a share λ of consumers are privacy-concerned and have θ that is uniformly distributed in the interval $[\underline{\theta}, \bar{\theta}]$, where $\underline{\theta} > 0$, and the remaining $1 - \lambda$ share of consumers are privacy-unconcerned and have $\theta = 0$. Second, a consumer who joins the service incurs a cost αx for sharing an amount x of personal information (e.g., the effort it takes to fill up different fields on a user profile page or to edit and upload photos), which is homogeneous across users. Third, a consumer also faces a fixed opportunity cost of using the service, which we denote as $K > 0$.¹⁶

Thus, the net utility of a consumer type θ is

$$U(x, \theta, q, d) = v + x - \frac{x^2}{2} + \gamma xq + \beta q - \theta dx - \alpha x - K.$$

Value of personal information. We suppose that all the third parties interested in users’ personal information have the same willingness to pay $r > 0$ for a unit of information. In addition, we assume that the firm is a monopolist in the market for (its customers’) personal information. This simplifying assumption implies that the firm always sets the unit price for access to its customers’ personal information to r , hence fully extracting third parties’ surplus.¹⁷

¹³Assuming $\gamma > 0$ is particularly relevant in the context of social media platforms. For example, in 2020 Instagram improved the quality of its services by introducing Instagram Stories, in an attempt to boost the sharing of original personal content. This move suggests that users view the quality of the social media platform and the amount of information they share as complements.

¹⁴Note also that one can interpret d as the disclosure level above a certain acceptance threshold (assumed to be the same for all users) as in Casadesus-Masanell and Hervas-Drane [2015]. This is justified by the fact that the firm will always choose a disclosure level that is (weakly) above that threshold.

¹⁵Note that the linearity of the privacy cost θdx with respect to d rests on the implicit assumption that sharing personal data with any third party induces the same privacy cost for a given consumer.

¹⁶If $K = 0$, all users would find it (weakly) optimal to use the service (i.e., there would be no extensive margin effect). This would make our model less rich and limit the insights we can draw from it.

¹⁷Note that our results hold qualitatively also if we relax the assumption that the firm extracts all third parties’ surplus. To see why, suppose that third parties are able to capture a positive share of that surplus and assume that the unit price of personal information is ϕr , where $\phi \in (0, 1]$. The firm’s profit function in this variant of our model can be derived from the one in our baseline setting by replacing r with ϕr . This implies in particular that the effect of a disclosure cap on a firm’s choice of innovation is (qualitatively) the same as in the baseline model.

Firm's profit. Denote by $x(\theta)$ the amount of information provided by a consumer of type $\theta \in [\underline{\theta}, \bar{\theta}]$.¹⁸ The firm's profit is the sum of the revenues collected by third parties net of the (fixed) cost $C(q)$ of achieving a level of quality-enhancing innovation q :

$$\Pi(q, d) = rdX(q, d) - C(q),$$

where $X(q, d) \equiv (1 - \lambda)x(0) + \frac{\lambda}{\bar{\theta} - \underline{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} x(\theta) d\theta$ is the total amount of information disclosed by users that join the service.

For tractability, we assume that the cost function for quality-enhancing innovation is quadratic: $C(q) = q^2/2$.

Timing. We consider the following two-stage game:

1. The firm chooses its quality-enhancing innovation level q and its level of information disclosure d .
2. Users observe the levels of innovation and disclosure. They decide then whether to patronize the firm and, if they do, how much personal information to provide.

The equilibrium concept is Subgame Perfect Nash Equilibrium.

Finally, we denote

$$\underline{q} \equiv \frac{\lambda r((1 - \alpha)\gamma + \beta) + \min\{0, \gamma r(\bar{\theta}(1 - \lambda) - \underline{\theta})\}}{\bar{\theta} - \underline{\theta} - \gamma^2 \lambda r},$$

and

$$\bar{q} \equiv \frac{\lambda r((1 - \alpha)\gamma + \beta) + \max\{0, \gamma r(\bar{\theta}(1 - \lambda) - \underline{\theta})\}}{\bar{\theta} - \underline{\theta} - \gamma^2 \lambda r},$$

and make the following assumptions in the remainder of the paper:

Assumption 1. $\bar{\theta} - \underline{\theta} > \gamma^2 r$.

Assumption 2. $\alpha > \gamma \bar{q}$.

Assumption 3. $v > K - \beta \underline{q} - \frac{1}{2}(1 + \gamma \underline{q} - \alpha)^2$.

Assumption 4. $v < K - \beta \bar{q}$.

¹⁸The amount of information $x(\theta)$ can be equal to zero either because the consumer of type θ decides not to use the service, or because she uses the service but decides not to provide personal information.

Assumption 1 ensures that the profit function of the firm is concave in q . Assumption 2 will guarantee that no user finds it optimal to share all their personal information, even if they are privacy-unconcerned.¹⁹ Assumption 3 will ensure that all privacy-unconcerned users find it optimal to join the service. Finally, Assumption 4 will guarantee that no user finds it optimal to join the service and provide no personal information.

IV. Analysis

In this section, we derive the users' optimal choices and the firm's optimal level of (quality-enhancing) innovation. We solve the model by backward induction starting with the user's problem.

i. User's choice

Conditional on patronizing the firm, a user chooses her level of information provided so as to maximize her utility $U(x, \theta, q, d)$. Denoting

$$\tilde{x}(\theta, q, d) \equiv \arg \max_{x \in [0,1]} U(x, \theta, q, d),$$

it is straightforward to show that:

$$\tilde{x}(\theta, q, d) = \begin{cases} 1, & \text{if } \gamma q - \theta d - \alpha \geq 0 \\ 1 + \gamma q - \theta d - \alpha, & \text{if } -1 < \gamma q - \theta d - \alpha < 0 \\ 0, & \text{if } \gamma q - \theta d - \alpha \leq -1 \end{cases}$$

We will focus hereafter on (the most interesting) case in which $\tilde{x}(\theta, q, d) \in (0, 1)$ for all users who join the service. The effect of a marginal increase in quality-enhancing innovation on the amount of information provided by a user is

$$(1) \quad \frac{\partial \tilde{x}(\theta, q, d)}{\partial q} = \gamma > 0, \quad \forall \theta \geq 0.$$

An increase in the level of innovation raises a user's marginal gross utility from providing information; hence, she finds it optimal to provide more information. Note that this effect applies to both types of users and shows that γ captures the extent of complementarity between information disclosure and the level of innovation. Moreover, for all privacy-concerned users, we observe that $\tilde{x}(\theta, q, d)$ decreases in d and θ .

¹⁹The parameter α allows us to have an intensive margin for both types of users.

The following lemma summarizes this discussion.

Lemma 1. *Assume $\gamma q - \theta d - \alpha \in (-1, 0)$ for any θ . Conditional on using the service, the amount of information disclosed by any user is increasing in the level of innovation. Moreover, the amount of information that a privacy-concerned consumer provides to the firm is decreasing in the disclosure level and the idiosyncratic privacy cost parameter.*

Denoting $\tilde{U}(\theta, q, d) \equiv U(\tilde{x}(\theta, q, d), \theta, q, d)$, we get that

$$(2) \quad \tilde{U}(\theta, q, d) = v + \beta q - K + \begin{cases} \frac{1}{2} + \gamma q - d\theta - \alpha, & \text{if } \gamma q - \theta d - \alpha \geq 0 \\ \frac{1}{2} + \gamma q - d\theta - \alpha + \frac{1}{2}[\gamma q - \theta d - \alpha]^2, & \text{if } -1 < \gamma q - \theta d - \alpha < 0 \\ 0, & \text{if } \gamma q - \theta d - \alpha \leq -1 \end{cases}.$$

A consumer of type θ chooses to patronize the firm if and only if $\tilde{U}(\theta, q, d) \geq 0$. Focusing on the case in which there is only partial disclosure (second line) by the users that joined the firm, and denoting

$$\tilde{\theta}(q, d) \equiv \frac{(1 + \gamma q - \alpha) - \sqrt{2(K - \beta q - v)}}{d},$$

we get the following result.

Lemma 2. *Assume $\gamma q - \theta d - \alpha \in (-1, 0)$ for any θ . Then, a user with $\theta \leq \tilde{\theta}(q, d)$ joins the firm and a user with $\theta > \tilde{\theta}(q, d)$ does not join the firm. Moreover, $\tilde{\theta}(q, d)$ is increasing in q and decreasing in d .*

ii. Firm's optimal level of innovation

We now determine the privately optimal level of innovation for any given information disclosure. The profit of the firm, given the participation and information provision of the users, is

$$\tilde{\Pi}(q, d) = rd \left\{ (1 - \lambda) \tilde{x}(0, q, d) + \frac{\lambda}{\bar{\theta} - \underline{\theta}} \int_{\underline{\theta}}^{\tilde{\theta}(q, d)} \tilde{x}(\theta, q, d) d\theta \right\} - C(q).$$

Denoting $D^P(q, d) \equiv \frac{\tilde{\theta}(q, d) - \underline{\theta}}{\bar{\theta} - \underline{\theta}}$ the mass of privacy-concerned users joining the service, the firm's net marginal benefit from innovation is

$$\begin{aligned}
(3) \quad \frac{\partial \tilde{\Pi}(q, d)}{\partial q} &= (1 - \lambda)rd \frac{\partial \tilde{x}(0, q, d)}{\partial q} + \lambda rd \frac{\partial \tilde{x}(\theta, q, d)}{\partial q} D^P(q, d) + \lambda rd \frac{\partial D^P(q, d)}{\partial q} \tilde{x}(\tilde{\theta}(q, d), q, d) - C'(q) \\
&= rd \underbrace{\gamma \left\{ (1 - \lambda) + \lambda D^P(q, d) \right\}}_{\text{intensive margin effect of innovation}} + \lambda rd \underbrace{\frac{\partial D^P(q, d)}{\partial q} \tilde{x}(\tilde{\theta}(q, d), q, d)}_{\text{extensive margin effect of innovation}} - C'(q).
\end{aligned}$$

The above expression shows that the firm accounts for both the intensive and extensive margin effects of quality-enhancing innovation. The intensive margin effect captures how a change in the level of innovation impacts the firms' revenues because of the higher amount of information provided by a *fixed* user base. This effect is positive because of the complementarity between the level of innovation and the level of information provision (by Lemma 1). Note that this effect is generated by the response of both privacy-concerned and privacy-unconcerned users to a change in quality. The extensive margin effect captures how the change in the demand of privacy-concerned users resulting from a change in the level of innovation impacts the firm's profit. This effect is positive as the firm's demand is increasing in the level of innovation (Lemma 2).²⁰

Using (3) and denoting $q^M(d)$ the firm's optimal level of innovation for given information disclosure d , we obtain²¹

$$(4) \quad q^M(d) = \frac{r(\gamma d(\bar{\theta}(1 - \lambda) - \underline{\theta}) + \lambda(\gamma(1 - \alpha) + \beta))}{\bar{\theta} - \underline{\theta} - \gamma^2 \lambda r}.$$

V. Privacy regulation

In this section, we study the impact of a privacy regulation taking the form of a disclosure cap on quality-enhancing innovation and its desirability for consumers.

i. Impact of privacy regulation on innovation

We first examine the effect of a binding disclosure cap on the firm's incentive to innovate.²² From (4), it is clear that $q^M(d)$ is increasing (resp. decreasing) in d if $\bar{\theta}(1 - \lambda) - \underline{\theta} > (<)0$. This leads to the following proposition.

²⁰Note that there is no extensive margin effect for the privacy-unconcerned users.

²¹Note that this expression satisfies the condition $\gamma q - \theta d - \alpha \in (-1, 0)$ appearing in Lemmas 1 and 2 under Assumptions 2-4.

²²We say that a disclosure cap is binding if it is strictly lower than the (unregulated) profit-maximizing level of disclosure, which implies that it leads to a strictly lower disclosure level than the one the firm would choose absent regulation.

Proposition 1. Let $\lambda^q \equiv 1 - \underline{\theta}/\bar{\theta}$. The introduction of any binding cap on information disclosure has a negative (resp. positive) effect on the level of innovation if $\lambda < \lambda^q$ (resp. $\lambda > \lambda^q$).

To explain the result presented in the above proposition, it is useful to examine the effect of a marginal decrease in d on the firm's marginal benefit from quality-enhancing innovation:²³

$$(5) \quad -\frac{\partial^2 \tilde{\Pi}}{\partial q \partial d} = -r\gamma \left\{ \underbrace{\frac{(1-\lambda)}{\text{mass of privacy-unconcerned users}} + \frac{\lambda D^P(q,d)}{\text{mass of privacy-concerned users}}}_{\text{per-unit return effect (-)}} \right\} \underbrace{-\lambda r d \gamma \frac{\partial D^P(q,d)}{\partial d}}_{\text{demand-boosting effect (+)}}.$$

A reduction in d has two opposite effects. First, there is a negative *per-unit return effect* because the disclosure cap reduces the revenues of the firm for each given user (regardless of whether she is privacy-concerned or privacy-unconcerned). Second, there is a positive *demand-boosting effect* because a lower disclosure level leads to an expansion of the user base as more privacy-concerned consumers join the service. This implies that the increase in the total amount of information induced by (more) innovation is larger. Note that the magnitude of the demand-boosting effect is increasing in the share of privacy-concerned users λ , whereas the magnitude of the per-unit return effect is decreasing in λ (because $D^P(q,d) < 1$ for any q and d). This implies that an increase in the share of privacy-concerned users makes it more likely that the positive demand-boosting effect dominates the per-unit return effect. Moreover, rewriting (5) as follows

$$(6) \quad -\frac{\partial^2 \tilde{\Pi}}{\partial q \partial d} = -r\gamma \left\{ (1-\lambda) - \lambda \frac{\underline{\theta}}{\bar{\theta} - \underline{\theta}} \right\},$$

shows that the net effect is negative at $\lambda = 0$ and positive at $\lambda = 1$.²⁴ This, combined with the monotonicity of the R.H.S. of (6) with respect to λ , implies that a disclosure cap has a negative (resp. positive) effect on quality-enhancing innovation if the share of

²³Note that

$$-\frac{\partial q^M(d)}{\partial d} = \frac{\frac{\partial^2 \tilde{\Pi}}{\partial q \partial d}}{\frac{\partial^2 \tilde{\Pi}}{\partial q^2}},$$

which implies that the sign of $-\frac{\partial q^M(d)}{\partial d}$ and $\frac{\partial^2 \tilde{\Pi}}{\partial q \partial d}$ are the same. Moreover, note that the effect of a decrease in d on the extensive margin effect of innovation is zero. This is because $\tilde{x}(\tilde{\theta}(q,d), q, d)$ does not depend on d , and $\frac{\partial \tilde{\theta}(q,d)}{\partial q} + d \frac{\partial^2 \tilde{\theta}(q,d)}{\partial q \partial d} = 0$

²⁴Note that $\frac{\partial \tilde{\theta}(q,d)}{\partial d} = -\frac{\tilde{\theta}(q,d)}{d}$, which implies that $D^P(q,d) + d \frac{\partial D^P(q,d)}{\partial d}$ can be rewritten as

$$\frac{\tilde{\theta}(q,d) - \underline{\theta}}{\bar{\theta} - \underline{\theta}} - d \frac{\tilde{\theta}(q,d)}{d(\bar{\theta} - \underline{\theta})} = -\frac{\underline{\theta}}{\bar{\theta} - \underline{\theta}},$$

which is then used to obtain (6).

privacy-concerned users is small (resp. large) enough.

ii. Desirability of privacy regulation for consumers

We now discuss the desirability for users of a (binding) disclosure cap in light of the strategic effect that it generates on the level of innovation. Consumer surplus is given by

$$CS(q, d) \equiv (1 - \lambda)\tilde{U}(0, q, d) + \frac{\lambda}{\bar{\theta} - \underline{\theta}} \int_{\underline{\theta}}^{\bar{\theta}(q, d)} \tilde{U}(\theta, q, d) d\theta.$$

A (marginal) decrease in disclosure level has the following two effects on consumer surplus:

$$\underbrace{\frac{\partial CS(q^M(d), d)}{\partial q} \frac{\partial q^M}{\partial d}}_{\text{innovation-mediated effect (-/+)}} \quad \underbrace{\frac{\partial CS(q^M(d), d)}{\partial d}}_{\text{direct effect (+)}}.$$

The first term captures the effect that a decrease in disclosure has on consumer surplus via the induced change in the level of innovation. Proposition 1 implies that this term can be either positive or negative. The second term captures the positive direct effect that a decrease in disclosure has on consumer surplus through the induced decrease in privacy costs. If $\lambda > \lambda^q$, a decrease in disclosure leads to more innovation. In this case, both effects are positive and privacy regulation is always desirable for users. If $\lambda < \lambda^q$, however, the two effects have opposite signs and the net effect depends on their relative magnitude. Note that at $\lambda = 0$ the direct effect is zero and, therefore, the net effect is negative. By continuity, this is also the case in the neighborhood of $\lambda = 0$. Moreover, at $\lambda = \lambda^q$, the innovation-mediated effect is zero and, therefore, the net effect is positive. This implies that this effect is positive in the neighborhood of $\lambda = \lambda^q$ (by continuity). Formally, denoting d^M the profit-maximizing disclosure level, let us define

$$\lambda^- \equiv \sup \left\{ \lambda' > 0 \mid - \frac{\partial CS(q^M(d^M), d^M)}{\partial q} \frac{\partial q^M}{\partial d} - \frac{\partial CS(q^M(d^M), d^M)}{\partial d} < 0 \quad \forall \lambda \in [0, \lambda'] \right\}.$$

$$\lambda^+ \equiv \inf \left\{ \lambda' > 0 \mid - \frac{\partial CS(q^M(d^M), d^M)}{\partial q} \frac{\partial q^M}{\partial d} - \frac{\partial CS(q^M(d^M), d^M)}{\partial d} > 0 \quad \forall \lambda \in [\lambda', \lambda^q] \right\},$$

The following proposition provides conditions under which a regulatory cap inducing a disclosure level that is marginally lower than the profit-maximizing level d^M is desirable, or not, for users.


Proposition 2. *A cap leading to a disclosure level that is marginally lower than the profit-maximizing level is desirable (resp., not desirable) for users if $\lambda \in (\lambda^+, 1]$ (resp. $\lambda \in [0, \lambda^-)$).*

The above proposition shows that if the share of privacy concerned users is sufficiently small, then a cap inducing a marginal decrease in the firm's disclosure level (from the unregulated profit-maximizing level) is detrimental to consumers. In this case, the negative impact of the cap on innovation outweighs its direct positive impact (stemming from lower privacy costs). In contrast, if the share of privacy-concerned users is sufficiently large then a cap leading to a disclosure level that is marginally lower than the profit-maximizing level is desirable for users. In this scenario, it is either the case that the disclosure cap has a negative innovation-mediated effect on users that is dominated by its positive direct effect (i.e. when $\lambda \in (\lambda^+, \lambda^q)$) or that both effects are positive (i.e. when $\lambda \in (\lambda^q, 1]$).


Let λ^{CS} denote the share λ such that a reduction in d has a net zero effect on consumer surplus. Through simulations, we show that there are cases for which $\lambda^+ = \lambda^- (\equiv \lambda^{CS})$ and therefore a disclosure cap leading to a disclosure level marginally lower than the profit-maximizing level is (not) desirable for consumers if $\lambda > (<) \lambda^{CS}$. Figure 1 presents a graphical representation of the effect of such a cap on innovation and consumer surplus for the following values: $\beta = 0.3$, $\alpha = 0.5$, $\bar{\theta} = 2$, $\underline{\theta} = 1$, $v = 1$, $r = 0.3$, $K = 1.1$. We focus on the parameter range $\lambda \in [0, 0.5]$ as in this interval the effect of privacy regulation on innovation is negative and, therefore, there is a trade-off between privacy and innovation. We consider two cases: strong complementarity between information and innovation by assuming $\gamma = 0.8$ (panel (a)); weak complementarity between information and innovation by assuming $\gamma = 0.3$ (panel (b)).

The red dotted line identifies the cut-off value λ^{CS} at which the impact of the regulation on consumer surplus is neutral and shows that, for the parameter range considered, the effect of the regulation on consumer surplus—denoted in the y-axis by $CS'(d^M) \equiv -\frac{\partial CS(q^M(d^M), d^M)}{\partial q} \frac{\partial q^M}{\partial d} - \frac{\partial CS(q^M(d^M), d^M)}{\partial d}$ —is monotonically increasing in λ , which implies that $\lambda^- = \lambda^+ = \lambda^{CS}$. On the right of the red dotted line, the regulation is desirable for users even though it reduces the level of innovation. On the left side of the dotted line, the negative effect of the regulation on innovation is large enough to make it undesirable for users.

The two panels only differ in the degree of users' complementarity between innovation and information disclosure. Panel (b) shows that when users perceive innovation and information as strong complements, the cut-off value λ^{CS} is greater than in the case in which users perceive innovation and information as weak complements. In other words, the comparison of the two panels suggests that the area in which the disclosure cap harms users is larger when innovation and information are stronger complements.



(a) Strong complementarity ($\gamma = 0.8$)



(b) Weak complementarity ($\gamma = 0.3$)

Figure 1: Effect of privacy regulation on consumer surplus

VI. Extensions

In this section, we present three extensions. First, we consider the case in which third parties are heterogeneous. Second, we discuss how our analysis would change if we introduce network externalities in the model. Finally, we assume that disclosure generates positive effects on users (and not only privacy costs).

i. Heterogeneous third parties

In this extension, we analyze the scenario where third parties are heterogeneous in the privacy costs that they induce for consumers.²⁵ This heterogeneity could reflect, for instance, differences in data use practices of these third parties. Consider, for example, the

²⁵See Bergemann and Bonatti [2015] and Bergemann et al. [2018] for models in which a monopolist sells personal information to third parties that are heterogeneous along other dimensions.

case of third-party advertisers. Some advertisers may choose to target their advertisements using data at a more aggregated level than others (e.g., based on demographic groups instead of individual characteristics), thereby resulting in lower privacy costs.

Let us assume that disclosing a unit of personal information to a third party of type $s \in [0, 1]$ induces a privacy cost $2\theta s$ for the consumer of type θ (instead of a privacy cost θ in the baseline model). The total privacy cost incurred by a consumer of type θ when an amount x of her personal information is disclosed to third parties located in $[0, d]$ is then given by $\int_0^d 2\theta s x ds = \theta d^2 x$. Correspondingly, the consumer's utility function is

$$U(x, \theta, q, d) = v + x - \frac{x^2}{2} + \gamma x q + \beta q - \theta d^2 x - \alpha x - K.$$

As in the baseline model, we restrict our attention to the case in which there is partial disclosure by the users that joined the firm. Denote

$$\hat{\theta}(q, d) \equiv \frac{(1 + \gamma q - \alpha) - \sqrt{2(K - \beta q - v)}}{d^2} = \frac{\tilde{\theta}(q, d)}{d}$$

the critical value such that a user with $\theta \leq \hat{\theta}(q, d)$ joins the firm and a user with $\theta > \hat{\theta}(q, d)$ do not join the firm.

Note that the equivalent of (6) in this case is given by²⁶

$$(7) \quad -\frac{\partial^2 \tilde{\Pi}}{\partial q \partial d} = -r\gamma \left\{ (1 - \lambda) - \lambda \frac{\hat{\theta}(q, d) + \underline{\theta}}{\bar{\theta} - \underline{\theta}} \right\}.$$

One can immediately verify that the effect of the disclosure cap on the innovation level is negative at $\lambda = 0$, whereas it is positive at $\lambda = 1$.

This confirms the insight that for a sufficiently small (resp. large) share of privacy-concerned users, the effect of a disclosure cap on innovation is negative (resp. positive).

ii. Network externalities

Social media platforms represent an example of firms collecting users' data and investing in innovative features that raise users' incentives to share data. A key aspect of social media platforms is the presence of direct network effects, which have not been considered

²⁶Note that $\frac{\partial \hat{\theta}(q, d)}{\partial d} = -\frac{2\hat{\theta}(q, d)}{d}$, which implies that $D^P(q, d) + d \frac{\partial D^P(q, d)}{\partial d}$ can be rewritten as

$$\frac{\hat{\theta}(q, d) - \underline{\theta}}{\bar{\theta} - \underline{\theta}} - d \frac{2\hat{\theta}(q, d)}{d(\bar{\theta} - \underline{\theta})} = -\frac{\hat{\theta}(q, d) + \underline{\theta}}{\bar{\theta} - \underline{\theta}},$$

which is then used to obtain (7).

in our baseline model. Assume now that the utility of each user increases with the amount of information provided by the other users. To incorporate these network externalities, we augment the consumer's utility in the baseline model by σX , where $X \equiv \lambda \int_{\underline{\theta}}^{\theta^e} x(\theta) d\theta + (1 - \lambda)x(0)$ is the amount of personal information provided by all (other) users,²⁷ θ^e is the expected marginal privacy-concerned user joining the service, and $\sigma \geq 0$ is a parameter that captures the intensity of the externalities.

Let us restrict our attention to the case in which users never disclose all information conditional on joining the service. Then, the optimal level of information provision is given by

$$\tilde{x}(\theta, q, d, \sigma) = 1 - \alpha - d\theta + \lambda\sigma(\theta^e - \underline{\theta} - 1) + \gamma q + \sigma,$$

which is increasing in the value of the network benefit σ . In a fulfilled expectations equilibrium, there exists (as in the baseline model) a threshold $\tilde{\theta}(q, d, \sigma)$ such that only users with type $\theta < \tilde{\theta}(q, d, \sigma)$ use the service. Specifically, assuming that $\sigma < d$, this threshold is given by

$$\tilde{\theta}(q, d, \sigma) = \frac{1 - \alpha - \sigma(\underline{\theta}\lambda + \lambda - 1) + \gamma q - \sqrt{2(K - \beta q - v)}}{d - \lambda\sigma}.$$

We restrict our attention to the case in which not all privacy-concerned users find it optimal to join the service, i.e., $\sigma < d$. The presence of network externalities alters the participation decision of privacy-concerned users. Specifically, more privacy-concerned users would join the service the larger the network benefit, i.e., $\frac{\partial \tilde{\theta}(q, d, \sigma)}{\partial \sigma} > 0$. This intuitive result follows from the fact that users' utility function is increasing in σ and X .

The effect of a marginal decrease in the disclosure level on the marginal benefit of investing in quality-enhancing innovation can again be split into two terms capturing the per-unit return effect and the one capturing the demand-boosting effect. The sign of these effects is the same as in the baseline (see (5)).

Note that the equivalent of (6) in this case is given by²⁸

$$(8) \quad -\frac{\partial^2 \tilde{\Pi}}{\partial q \partial d} = -r\gamma \left\{ (1 - \lambda) - \lambda \frac{\underline{\theta}(d - \lambda\sigma) + \lambda\sigma\tilde{\theta}(q, d, \sigma)}{(d - \lambda\sigma)(\bar{\theta} - \underline{\theta})} \right\}.$$

It is immediate that at $\lambda = 0$ the effect of the disclosure cap on the level of innovation

²⁷Because a consumer is of zero measure in our setting, the total amount of personal information provided by all users and the total amount of information provided by all users but one coincide.

²⁸Note that $\frac{\partial \tilde{\theta}(q, d, \sigma)}{\partial d} = -\frac{\tilde{\theta}(q, d, \sigma)}{d - \lambda\sigma}$, which implies that $D^P(q, d, \sigma) + d \frac{\partial D^P(q, d, \sigma)}{\partial d}$ can be rewritten as

$$\frac{\tilde{\theta}(q, d, \sigma) - \underline{\theta}}{\bar{\theta} - \underline{\theta}} - d \frac{\tilde{\theta}(q, d, \sigma)}{(d - \lambda\sigma)(\bar{\theta} - \underline{\theta})} = -\frac{\underline{\theta}(d - \lambda\sigma) + \lambda\sigma\tilde{\theta}(q, d, \sigma)}{(d - \lambda\sigma)(\bar{\theta} - \underline{\theta})},$$

which is then used to obtain (8).

is negative, whereas it is positive at $\lambda = 1$. Therefore, the insight that a disclosure cap has a positive (resp. negative) effect on innovation if the share of privacy-sensitive users is large (resp. small) enough continues to hold when accounting for the presence of direct network externalities.

iii. Positive effect of disclosure on users

Our baseline model abstracts away from any positive effect of disclosure to third parties on users. Assume now that disclosure brings about benefits to users which we capture through an additional term μdx in the consumer utility function where $\mu \geq 0$.²⁹ Thus, the *net* cost of disclosure for consumer of type θ is $(\theta - \mu) dx$.

Note that the equivalent of (6) in this case is given by³⁰

$$(9) \quad -\frac{\partial^2 \tilde{\Pi}}{\partial q \partial d} = -r\gamma \left\{ (1 - \lambda) - \lambda \frac{\theta - \mu}{\theta - \underline{\theta}} \right\}.$$

The value of expression (9) at $\lambda = 0$ is always positive. This implies that the finding that a disclosure cap has a negative effect on innovation if the share of privacy-concerned users is sufficiently small still holds in this setting. However, the value of expression (9) at $\lambda = 1$ can be either positive or negative depending on the size of the benefit μ . Specifically, it is positive (resp. negative) if μ is lower (resp. higher) than $\underline{\theta}$. This means that the result that a disclosure cap has a positive effect on innovation if the share of privacy-concerned users is sufficiently large carries out to this setting if the benefit from disclosure μ is sufficiently small, but does not otherwise.

VII. Conclusion

In this paper, we study how a privacy regulation taking the form of a cap on data disclosure affects a monopolist's incentives to invest in quality-enhancing innovation and consumer surplus. Our analysis has implications for the regulation of dominant firms' disclosure policy, which attracted policymakers' attention in recent years.

²⁹We assume for the sake of exposition that μ is the same for all users.

³⁰Recall that $\frac{\partial \tilde{\theta}(q,d)}{\partial d} = -\frac{\tilde{\theta}(q,d)}{d}$, which implies that $D^P(q,d) + d \frac{\partial D^P(q,d,\sigma)}{\partial d}$ can be rewritten as

$$\frac{\tilde{\theta}(q,d) - \underline{\theta} + \mu}{\tilde{\theta} - \underline{\theta}} - d \frac{\tilde{\theta}(q,d)}{d(\tilde{\theta} - \underline{\theta})} = -\frac{\underline{\theta} - \mu}{\tilde{\theta} - \underline{\theta}}$$

which is then used to obtain (9).

We find that the impact of a reduction in disclosure level on the firm's optimal choice of innovation depends on the share of privacy-concerned users. Specifically, it shows that the trade-off between innovation and privacy is likely to exist if there is a large share of consumers that are (essentially) privacy-unconcerned. In this case, innovation is likely to decrease with the introduction of stricter privacy regulation and the desirability of regulation for users depends on whether gains from a cap on information disclosure more than compensate for the losses from a lower level of innovation. However, if the share of privacy-concerned consumers is large enough, the trade-off between privacy and innovation may not exist. Therefore, in this case, regulators should be more confident about the desirability of privacy regulation for consumers.

In addition to a disclosure cap, another policy that can be explored using our framework is the taxation of disclosure revenues. The taxation of digital monopoly platforms has been studied, for instance, by Bloch and Demange [2018] and Bourreau et al. [2018]. However, to the best of our knowledge, no paper has considered the impact of taxation on the firm's incentives to invest in innovation. A (unit) tax on the monopolist's disclosure revenues would translate to a reduction in the value of information in our model. Because this policy affects both the innovation and disclosure levels of the firm, the impact of a tax is a priori unclear.

References

- Laura Abrardi and Carlo Cambini. Carpe data: Protecting online privacy with naive users. *Information Economics and Policy*, 60:100988, 2022.
- Laura Abrardi, Carlo Cambini, and Steffen Hoernig. "I don't care about cookies!" platform data disclosure and time-inconsistent users. *Mimeo*, 2021.
- Simon P Anderson. Regulation of television advertising. *The economic regulation of broadcasting markets*, 189:224, 2007.
- Susan Athey. Information, privacy, and the internet. *CPN*, 2014.
- Dirk Bergemann and Alessandro Bonatti. Selling cookies. *American Economic Journal: Microeconomics*, 7(3):259–94, 2015.
- Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, 2018.
- David Besanko, Shabtai Donnenfeld, and Lawrence J White. Monopoly and quality distortion: effects and remedies. *The Quarterly Journal of Economics*, 102(4):743–767, 1987.

- Francis Bloch and Gabrielle Demange. Taxation and privacy protection on internet platforms. *Journal of Public Economic Theory*, 20(1):52–66, 2018.
- Marc Bourreau, Bernard Caillaud, and Romain De Nijs. Taxation of a digital monopoly platform. *Journal of Public Economic Theory*, 20(1):40–51, 2018.
- Ramon Casadesus-Masanell and Andres Hervas-Drane. Competing with privacy. *Management Science*, 61(1):229–246, 2015.
- Daniel Castro. Stricter privacy regulations for online advertising will harm the free internet. *ITIF - Information Technology and Innovation Foundation*, 2010.
- Chiara Conti and Pierfrancesco Reverberi. Price discrimination and product quality under opt-in privacy regulation. *Information Economics and Policy*, 55:100912, 2021.
- Philipp D Dimakopoulos and Slobodan Sudaric. Privacy and platform competition. *International Journal of Industrial Organization*, 61:686–713, 2018.
- Samuel Goldberg, Garrett Johnson, and Scott Shriver. Regulating privacy online: An economic evaluation of the GDPR. *American Economic Journal: Economic Policy*, 2022. Forthcoming.
- Avi Goldfarb and Catherine Tucker. Privacy regulation and online advertising. *Management Science*, 57(1):57–71, 2011.
- Avi Goldfarb and Catherine Tucker. Privacy and innovation. *Innovation policy and the economy*, 12(1):65–90, 2012.
- Sinem Hidir and Nikhil Vellodi. Privacy, personalization, and price discrimination. *Journal of the European Economic Association*, 19(2):1342–1363, 2021.
- Shota Ichihashi. Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–95, 2020.
- Garrett Johnson. Economic research on privacy regulation: Lessons from the GDPR and beyond. *National Bureau of Economic Research Working Paper*, 2022.
- Garrett Johnson, Scott Shriver, and Shaoyin Du. Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1):33–51, 2020.
- Tobias Kircher and Jens Foerderer. Ban targeted advertising in apps? an empirical investigation of the consequences for app development. *Management Science*, 2022. Forthcoming.

Amalia Miller and Catherine Tucker. Encryption and the loss of patient data. *Journal of Policy Analysis and Management*, 30(3):534–556, 2011a.

Amalia R Miller and Catherine Tucker. Privacy protection and technology diffusion: The case of electronic medical records. *Management Science*, 55(7):1077–1093, 2009.

Amalia R Miller and Catherine Tucker. Can health care information technology save babies? *Journal of Political Economy*, 119(2):289–324, 2011b.

Omid Rafeian and Hema Yoganarasimhan. Targeting and privacy in mobile advertising. *Marketing Science*, 40(2):193–218, 2021.

Eytan Sheshinski. Price, quality and quantity regulation in monopoly situations. *Economica*, 43(170):127–137, 1976.

A. Appendix

i. Proof of Lemma 1

The first part of the proof is omitted as it follows from the discussion in the main text. The second part of the proof follows from the fact that

$$\frac{\partial \tilde{x}(\theta, q, d)}{\partial d} = -\theta < 0,$$

$$\frac{\partial \tilde{x}(\theta, q, d)}{\partial \theta} = -d < 0.$$

whenever $\tilde{x}(\theta, q, d) \in (0, 1)$.

ii. Proof of Lemma 2

Omitted as it follows from the discussion in the main text.

iii. Determination of $q^M(d)$

Rearranging and simplifying (3), we obtain the following

$$(10) \quad \frac{\partial \Pi(q, d)}{\partial q} = \frac{d\lambda r \left(\gamma \left(\frac{1-\alpha-\sqrt{2(K-\beta q-v)}+\gamma q}{d} - \underline{\theta} \right) + \frac{\sqrt{2(K-\beta q-v)} \left(\gamma + \frac{\beta}{\sqrt{2(K-\beta q-v)}} \right)}{d} \right)}{\bar{\theta} - \underline{\theta}} + \gamma d(1-\lambda)r - q$$

$$= rd\gamma - \frac{\lambda r(\beta + \gamma(1-\alpha - d\bar{\theta} + \gamma q))}{\underline{\theta} - \bar{\theta}} - q = 0$$

Solving for q , we obtain

$$(11) \quad q^M(d) = \frac{r(\gamma d(\bar{\theta}(1-\lambda) - \underline{\theta}) + \lambda(\gamma(1-\alpha) + \beta))}{\bar{\theta} - \underline{\theta} - \gamma^2 \lambda r}.$$

Proof of Proposition 1

The result easily follows from (4).

Proof of Proposition 2

Omitted as it directly follows from the discussion in the main text.