

“I don't care about cookies!” data disclosure and time-inconsistent users

Original

“I don't care about cookies!” data disclosure and time-inconsistent users / Abrardi, L., Cambini, C., Hoernig, S.. - In: INFORMATION ECONOMICS AND POLICY. - ISSN 0167-6245. - 69:(2024). [10.1016/j.infoecopol.2024.101112]

Availability:

This version is available at: 11583/2995961 since: 2024-12-27T14:11:02Z

Publisher:

Elsevier B.V.

Published

DOI:10.1016/j.infoecopol.2024.101112

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



“I don’t care about cookies!” data disclosure and time-inconsistent users [☆]

Laura Abrardi ^{a,*}, Carlo Cambini ^b, Steffen Hoernig ^c

^a Politecnico di Torino, DIGEP, c.so Duca degli Abruzzi 24, Torino, Italy

^b Politecnico di Torino, DIGEP, Italy

^c Nova School of Business and Economics, Portugal

ARTICLE INFO

JEL classification:

D82
D83
D86
L12
L51

Keywords:

Data disclosure
Time inconsistency
Privacy

ABSTRACT

Time-inconsistent internet users neglect future privacy costs and release too much data to digital firms. We study how regulation that requires user consent for data processing affects firm profits, user surplus, and welfare, depending on the degree of time inconsistency and on firms’ business models. If the firm appropriates sufficiently high profits from data, consent mechanisms increase welfare only if their design facilitates consent refusal and time inconsistency is neither too high nor too low. If firms can make it difficult to opt out, it may be better for society to let the former choose the disclosure level. However, consent policies increase user surplus when time inconsistency is high. Voluntary caps on usage can raise profits by making some users disclose more data.

1. Introduction

If data is the new oil, naive internet users feed an inexhaustible oil field. Individuals often reveal massive amounts of data through their daily online behavior, even though they claim to care about privacy (Acquisti et al., 2015). This behavior partly stems from a widespread pattern of time inconsistency that induces the disclosure of personal information and which lowers future utility in favor of the immediate gratification through users’ browsing activities (Acquisti, 2004; Acquisti and Grossklags, 2005; Tucker, 2019).¹

Current privacy legislation, most notably the European GDPR, relies on the consent principle as the prime mechanism through which people can protect or waive their privacy.² The consent principle is based on the assumption that people, when provided with adequate information, are able to make self-interested and forward-looking choices. However, this assumption fails with time-inconsistent individuals. Because of self-control problems, consent forms are perceived as a nuisance to be dispatched by choosing the default option, which often is to waive privacy. For example, cookie banners may nudge users into disclosure by presenting immediately the option to accept all cookies, while deny-

[☆] We would like to thank Grazia Cecere, Andrea Mantovani, Leonardo Madio, Carlo Reggiani, David E.M. Sappington, Nicolas Soulié and participants at the EARIE 2021 (Bergen), MaCCI 2021 Annual Conference (Mannheim) and the Louvain Economics of Digitization Research Group for helpful comments. This study was carried out within the “Cyber Resilience: Markets, Investment, Regulation” project (CUP E53D23016530001) - funded by European Union - Next Generation EU within the PRIN 2022 PNRR program (D.D.1409 del 14/09/2022 Ministero dell’Università e della Ricerca). This manuscript reflects only the authors’ views and opinions and the Ministry cannot be considered responsible for them. Steffen Hoernig was funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, LISBOA-01-0145-FEDER-022209), POR Lisboa (LISBOA-01-0145-FEDER-007722, LISBOA-01-0145-FEDER-022209) and POR Norte (LISBOA-01-0145-FEDER-022209).

* Corresponding author.

E-mail addresses: laura.abrardi@polito.it (L. Abrardi), carlo.cambini@polito.it (C. Cambini), steffen.hoernig@novasbe.pt (S. Hoernig).

¹ For instance, people often skip merchants’ privacy policies (Tsai et al., 2011) or are willing to reveal personal information in exchange for relatively small rewards (Spiekermann et al., 2001).

² The GDPR applies to all companies having users in the EU, and about 120 countries have now passed privacy laws resembling it (The Economist, February 20th, 2020).

ing disclosure of personal data to third parties requires personalizing a (usually long) list of trackers.

A considerable part of digital firms' profits comes from the exploitation of user data.³ This paper studies how users' time inconsistency about privacy affects digital firms' profits and their decision to exploit user data, and how users and firms react to the introduction of consent policies, such as cookie consent pop-ups. We also investigate the conditions under which introducing these consent policies is preferable for users or for society. Our analysis yields some interesting insights into the effects of consent policies and of their implementation by profit-maximizing digital firms.

One of the main challenges in approaching this problem is the fact that monetizing data is the primary line of business for some online firms, while others obtain a large share of their revenues also from the direct sale of a service. We build a flexible framework that allows us to analyze, in a unified setting, firms with different business models and thus varying degrees of preference for data exploitation. More precisely, we model firms that derive their profits from either one or both of two sources: from the direct sale of digital services, or from monetizing data with third parties, such as advertisers.

In our model, a digital firm provides a service to time-inconsistent users, who might be aware or not of their self-control problem. The usage of the service produces personal data. The firm obtains its revenues from commercial activities related to the usage of the service and from disclosing a share of the collected data to third parties. Data disclosure causes a future privacy loss to users, which they partly neglect due to time inconsistency. We consider two alternative policy options. If a consent policy is in place, the level of data disclosure is chosen by users, otherwise it is chosen by the firm. The consent policy allows users to restrict (at a cost) the disclosure level, by for example the personalization of privacy settings. Unawareness of time inconsistency further dampens users' ability to protect their privacy through restrictions to disclosure, weakening the effectiveness of consent mechanisms.

We have two main results. First, the introduction of consent policies aimed at protecting consumers can backfire and reduce user surplus. When disutility from personalization or time inconsistency are sufficiently high, user surplus is higher if the disclosure decision is left to platforms. In fact, under these conditions, users skip personalization of privacy settings but will subsequently curtail usage, in a belated attempt to protect their privacy. Conversely, platforms have the incentive to limit data disclosure from the outset not to harm future usage revenues, and users are also spared going through privacy popups. Thus, consent policies increase user surplus only if both personalization cost and time inconsistency are sufficiently low - they do not protect users against the effects of time inconsistency if using them is too costly. The scope for welfare-enhancing consent policies is even more limited. Consent policies increase welfare (indicating the social value of the market as the sum of profits and consumer surplus) only when the personalization cost is low and time inconsistency is neither too high nor too low. Users with a very low time inconsistency are more likely to excessively restrict disclosure, preventing the realization of the social value of their data.

Second, users' time inconsistency increases the importance of well-designed consent policies. In fact, although time-inconsistent users are more likely to skip the personalization of privacy settings, firms are also more tempted to exploit users by increasing data disclosure, especially when the value of data is high. Here, the aim of policy intervention should be to raise user awareness and reduce the burden of personalizing privacy settings or opting out directly.

Mandating firms to ask users for consent is a regulatory tool that has been widely adopted, in Europe and elsewhere. However, firms have

³ The exploitation of user data may involve, for example, targeting users based on their revealed personal information for the purpose of advertising or price discrimination.

significant leeway in implementing these consent policies. We thus extend our model to the case in which the digital firm determines the cost for users of personalizing the privacy settings. We find that the firms' incentive to raise rather than reduce the complexity of the consent form depends both on their business model and on users' time inconsistency. When the business model is based on the direct provision of the service or users' time inconsistency is mild, the firm has an interest to boost its commercial activities by reducing the complexity of consent forms, by making sure that users feel that their privacy is protected. Otherwise, firms have an incentive to make it harder for users to restrict disclosure. Interestingly, this firm strategy, while not benefiting users, can at least raise welfare. This feature is important from a policy perspective. In 2020 California enacted a law imposing an opt-out default option, and an increasing number of firms in Europe carry such an option, too. Such a policy protects time-inconsistent users, but at the same time may also reduce the potential social value of data.

Lastly, we further extend the model in two ways. First, we show that voluntary usage caps, such as time limits offered on Facebook, Apple or Google, can increase firm profits in the presence of consent policies. The reason is that a cap provides sophisticated users, i.e. those users who are aware of their time inconsistency, with a means to commit to lower future usage. This in turn reduces their privacy concerns and makes them more willing to choose higher levels of disclosure. The resulting gain in revenues from sharing data can more than outweigh the effects of lower usage. Usage caps also raise welfare if data are sufficiently valuable.

Second, we assume that the privacy cost stems not only from the disclosure of information to third parties, but also from its collection by the firm, thereby making users more exposed to the privacy risk. Results show that in this case users have a higher tendency to choose full disclosure, despite their higher exposure. This counterintuitive result is due to the fact that users now suffer a privacy cost even if they do not share data, so that the decision to restrict disclosure is relatively less advantageous than in the case the privacy loss stems only from data sharing.

In the following section we provide an overview of the literature. The rest of the paper is organized as follows. Section 2 presents our model. Sections 3 and 4 analyze the cases where privacy settings are chosen by the firm or by users under a consent policy, respectively. In Section 5 we extend the model by studying the effect of voluntary usage caps, and of privacy risk stemming also from data collection, and Section 6 concludes. All proofs are contained in the Appendix.

1.1. Related literature

This paper is part of a recent and growing literature on data acquisition, privacy regulation, and their impact on digital businesses.⁴ Two strands are directly related to our concerns. The first focuses on the potential effect of time inconsistency on privacy costs in the data acquisition market. Tucker (2019) suggests that psychological factors strongly influence user behavior on digital platforms. To our knowledge, only three papers consider boundedly rational consumers in this context. Acquisti and Varian (2005) consider repeat sales by a monopolist while assuming that tastes are intertemporally correlated. A firm can benefit from the use of past information, but only if its consumers are

⁴ The literature on privacy is vast. Recent papers have analyzed the impact of competition on privacy (Casadesus-Masanell and Hervas-Drane, 2015), the exploitation of information for personalized pricing (Braulin and Valletti (2016), Montes et al. (2019)), and the impact of taxation on data collection (Bloch and Demange (2018), Bourreau et al. (2018)). Our research is also associated with the literature on data externalities (Choi et al. (2019), Acemoglu et al. (2019)) which explains that excessive data sharing by some users might be consistent with rational behavior when the information disclosed by one consumer reveals information about others.

myopic, because rational consumers would be deterred from consuming in the first place. Using a similar argument, Taylor (2004) provides a theoretical model in which privacy regulations limit the ability of individual merchants to sell customer transaction data to other merchants. Rational consumers anticipate that their present actions may affect the prices they will face in subsequent transactions, hence, they may strategically modify purchasing behavior. Conversely, myopic consumers do not take the future consequences of their present decisions into account and suffer a loss of privacy. Finally, Baye and Sappington (2020) examine a similar setup, but assume that merchants operate inside a platform. Since the platform internalizes the externality of individual transactions, it maximizes its profits by disclosing all transactions data to its merchants. In their model, rational users gain from sharing personal information, as disclosure opens a channel through which they can signal their reservation values. Conversely, myopic users benefit when the platform never shares transactions data with third parties, as this prevents that users are exploited by merchants.

We depart from these studies in the nature of the behavioral bias involved. In the aforementioned literature, users fail to take the future exploitation of their data into account, but do not suffer from a self-control problem. Conversely, in our setting, users may be aware that their future use will be excessive but cannot by themselves prevent it from happening. This time inconsistency affects how seriously users take consent policies that are meant to protect them, which is the focus of this paper. This issue has recently been highlighted by experimental studies. Ek and Samahita (2023) ran an experiment designed to verify whether users really tend to have an undercommitment problem as stressed in the existing literature, or instead an overcommitment problem where users spend resources on commitment devices they do not actually need. They find that while a significant share of users is too pessimistic about their self-control and thus prone to overcommitment, both the fraction and loss of utility of users who are overoptimistic and thus prone to undercommitment dominate in the aggregate. Kummer and Schulte (2019) studied the relationship between willingness to pay for privacy and app design using data from the Google Store, where prior to installing an app users were shown the privacy-related permissions requested by the app. Users thus had a choice between paying for apps with fewer permissions or paying with their data for free apps, and app designers provided both options. Kummer and Schulte show that, at the same price and quality, users demanded less those apps that requested more permissions. This suggests users' awareness of privacy costs, in line with our assumptions.

The second strand of literature studies the effects of privacy regulation, in particular on price discrimination (Montes et al., 2019), targeted advertising (Johnson, 2013), innovation incentives (Lefouilli et al., 2024), and quality investment (Conti and Reverberi, 2021), although these papers do so in a context of full rationality. As in our paper, it is assumed that protecting privacy is costly for consumers, but we assume that users can be time inconsistent and unaware of their own time inconsistency, a combination that to the best of our knowledge has not been analyzed before in the privacy literature. Fainmesser et al. (2022) show that imposing a minimum level of investment in protecting data from leaking out together with a tax on data holdings can lead to an efficient outcome. Our approach complements theirs by considering regulation that involves the user side, specifically the role and complexity of consent forms.

While our paper is part of a stream of work based on the assumption that the collection and disclosure of personal information create privacy costs, a recent tendency in the literature is to start from the opposite baseline assumption that some disclosure of information could raise users' utility. The typical setting here is price discrimination among heterogeneous consumers where it is assumed that without disclosure a seller sets prices at a high level. Ali et al. (2022) show that lower-value consumers can credibly pool into groups to obtain discounts. Galperti and Perego (2023), leveraging on the general model of platform information design of Galperti et al. (2023), show that a platform that knows user types and optimally reveals information to a seller, pools users in a

similar manner if its objective is to maximize consumer surplus; still, if its objective is to maximize the seller's profits, it transmits all the information it has. They show that under-pooling giving users control over data disclosure creates externalities between user groups. No such externalities exist in the latter case (here the authors do not consider users' optimal privacy choices).

Contrary to some previous models in the literature, we do not assume by default that users' choice of disclosure must be at an interior optimum. Rather, we show that the fact that users can reduce their usage to avoid privacy costs creates a fundamental non-convexity in their decision problem, which leads to either no or full disclosure. Miklós-Thal et al. (2024) come to a similar result in a different setting: In their model, users are paid to provide either non-sensitive information or sensitive personal information (the latter at a privacy cost), or both. They show that if the platform learns how to deduct sensitive from non-sensitive information, then users will either disclose nothing (digital hermits) or everything.

We also provide a general framework for analyzing firm behavior, representing different business models of digital firms. Here we follow Bloch and Demange (2018) and Jullien et al. (2020), although their focus is on optimal taxation of digital platforms. There exist many types of platforms, which lead to different data disclosure approaches in order to boost the usage of their service or to extract value from the information, or both (Fainmesser et al., 2022). It is thus crucial to describe their varied choices in a unified setting, so as to provide a taxonomy of the profit-maximizing decisions of each type of firm. Representing the different business models across digital firms is also important for a better understanding of the effects of privacy constraints on firm design and user responses to privacy control tools.

2. The model

We model the interaction between users and a digital firm providing a service. Usage of the service entails the release of personal data, which the firm can monetize with third parties, imposing a privacy cost on users.

A digital firm provides a service to a unit mass of users, who choose the amount of usage $x \in [0, 1]$, e.g., the number of Google searches or Facebook likes. Usage reveals personal information about users' characteristics or preferences. For simplicity, we assume that one unit of usage corresponds to the production of one unit of personal information.⁵ The firm discloses the share $d \in [0, 1]$ of usage-generated information x to third parties, for example advertisers. The share d will be determined by users or by the firm, depending on whether or not there is a regulation in place which requires the consent of users for disclosing data (the consent policy). Users face a privacy risk stemming from the disclosure of the amount dx of information with third parties. For example, third parties may suffer a data breach, or they can exploit this information for purposes that harm consumers.⁶

Actions are divided between two periods, the present and the future. In the present, either the firm or users choose the level of disclosure d , which is observed by all ($t = 1$). Subsequently (at $t = 2$) users choose usage x , and utility from usage and profits are realized. In the future ($t = 3$), users suffer the privacy cost. Users cannot previously commit on their own to a specific usage level, nor to not using the service at all. Fig. 1 represents the timeline of the basic model. In Section 5.1 we extend the baseline setup by an initial period $t = 0$ where the firm offers a usage cap that consumers can voluntarily sign up to.

We now lay out the main features of the model in more detail.

⁵ In practice not all usage may generate personal information. For a consistent modeling of the privacy problem, we focus on the "active" usage that reveals personal information and thus creates a privacy problem.

⁶ In Section 5.2 we explore the case in which privacy risk stems also from the collection of information by the firm.

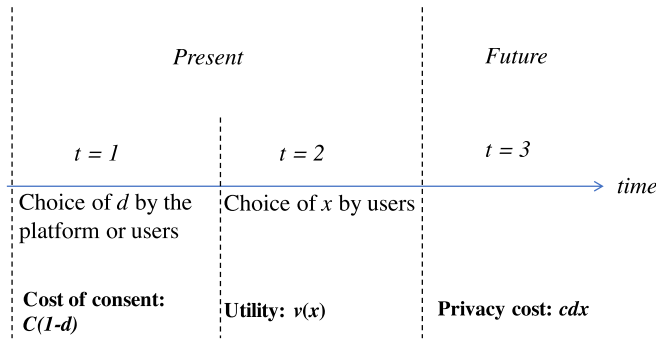


Fig. 1. Timeline of the model.

Users. Usage provides utility $v(x)$ with $v(0) = 0$, $v_x > 0$, $v_{xx} < 0$. For simplicity, we adopt the functional form $v(x) = x(2 - x)$.^{7,8} Users incur a future privacy cost cdx , $c \in (0, 2)$, due to the exploitation of their personal information by third parties, net of any potential benefit obtained from data disclosure.

Users have time-inconsistent preferences with regard to the privacy cost (see the seminal work of Laibson, 1997, and, more recently, O'Donoghue and Rabin, 1999a, O'Donoghue and Rabin, 1999b, and DellaVigna and Malmendier, 2004). In particular, following DellaVigna and Malmendier (2004), users *before using the service* correctly understand that their surplus from using the firm is $v(x) - cdx$, as the utility and the privacy cost are both obtained in the future. Still, *when actually using the service* in $t = 2$, they perceive a surplus equal to $v(x) - \beta cdx$, with $\beta \in [0, 1]$, as the privacy cost in the next period is discounted relative to the utility $v(x)$ received in the present period.

The firm. The firm obtains profits from two different sources, as in Bloch and Demange (2018) and Jullien et al. (2020). Direct provision of the service and the corresponding commercial activities originate profits δx , with $\delta > 0$. For example, these profits incorporate e-commerce commissions and pay-per-click ads. The parameter δ can be interpreted as the incremental value of more data to the firm or alternatively as commissions for clicks. The second source of profits is the monetization of data dx with third parties, for example for targeted advertising or for elaboration by data aggregators, as with search engines and social networks. The total profit obtained from the data shared with third parties is γdx , with $\gamma > 0$, and it is distributed between the firm and the third parties according to the share $\theta \in [0, 1]$ and $1 - \theta$. The parameter γ can be interpreted as the monetary value of the data collected by the firm. Firms such as search engines, social networks, or digital marketplaces, derive high profits from a large variety of information about their customers, their needs and interests, implying a high γ . Conversely, firms specialized in a more specific segment (such as video or music streaming services) have a low γ and thus collect a narrower set of data that they use to tailor their own supply. Third parties, e.g. advertisers, obtain a profit $\sigma(x) = (1 - \theta)\gamma dx$ from the information shared by the firm, where θ can be interpreted as the firm's market power vis-à-vis third parties.

The disclosure decision. We examine two scenarios, which differ in who makes the disclosure decision at $t = 1$. In the first scenario, analyzed in Section 3, the firm chooses the level of disclosure. In the second

scenario, in Section 4, users choose their preferred disclosure level d by personalizing the privacy settings provided to them under a consent policy. Such a personalization imposes a *choice cost* on users equal to $C(1 - d)$, with $C > 0$.⁹ Selecting full disclosure ($d = 1$) implies zero inconvenience, whereas adopting more stringent privacy settings takes an increasing amount of time and attention.¹⁰ We initially assume that the choice cost parameter C is exogenously given and focus on users' disclosure choices (Subsection 4.1). Then we consider the preferred level of choice cost for the firm or from society's point of view (Subsection 4.2).

At $t = 1$, users know that their true net surplus is $U(x) = v(x) - cdx$. Still, when actually using the service ($t = 2$), given a disclosure level d , users select usage x to maximize *perceived* net surplus $v(x) - \beta cdx$, with¹¹

$$x^*(d) = 1 - \frac{\beta cd}{2} > 0. \tag{1}$$

The expression of actual usage in (1) has three important implications. First, more time-consistent users (i.e., with a higher β) restrict usage more, because they better take into account the loss of privacy that usage entails. Second, users respond to higher levels of disclosure d by reducing usage, but this reaction is muted if they are more time inconsistent (lower β). Indeed, in the limit case of $\beta = 0$, usage does not depend on the disclosure level. This fact increases the potential for the exploitation of users' time inconsistency by the firm. Third, and as we will discuss below in Section 4 on the effect of consent policies, the downward adjustment of usage in response to higher d implies that *perceived* net surplus at the chosen usage is a convex function of d , and as a result users either prefer no disclosure ($d = 0$) or full disclosure ($d = 1$) over intermediate privacy settings.¹²

At $t = 1$, users may not be aware of their time inconsistency (DellaVigna and Malmendier, 2004). As a consequence, they may have biased expectations about their actual usage $x^*(d)$. Let us denote by $\hat{x}(d)$ users' belief, at $t = 1$, about their actual usage at $t = 2$. There are two types i of users: *sophisticated* users ($i = s$) and *naive* users ($i = n$), whose weights in the population are given by $\lambda \in [0, 1]$ and $1 - \lambda$, respectively. The two types differ in the degree of awareness of their time inconsistency: At $t = 1$ sophisticated users are aware that at $t = 2$ they will be using the service excessively, i.e. $\hat{x}_s(d) = x^*(d)$. Conversely, naive users at $t = 1$ expect that usage at $t = 2$ is equal to $\hat{x}_n(d) = x^{tc}(d) \equiv 1 - \frac{cd}{2}$, i.e. they believe that they are time consistent ($\beta = 1$). Denote by d_i , x_i^* and U_i^* , with $i \in \{s, n\}$, the levels of disclosure, actual usage and net surplus pertaining to sophisticated or naive users, respectively.

The firm maximizes its total profit, given by

$$\pi = \lambda \pi_s + (1 - \lambda) \pi_n,$$

where $\pi_i = (\delta + \gamma \theta d_i) x_i^*$, $i \in \{s, n\}$.

The social value of the firm's activity, including any data disclosure, is the weighted sum of users' true net surplus and the industry profits, which includes the firm and the third parties' profits:

$$W = \lambda (U_s^* + \alpha (\pi_s + \sigma_s)) + (1 - \lambda) (U_n^* + \alpha (\pi_n + \sigma_n)),$$

where $\alpha \in [0, 1]$ is the weight of the industry profit in the welfare function. A consumer privacy authority whose only concern is maximizing consumer surplus is represented by $\alpha = 0$. Conversely, a social planner

⁷ Choi et al. (2019) assume quasi-concave general functional forms, imposing from the outset that all choices are interior. Yet, users reduce usage under higher disclosure, which leads to non-concavity and potentially to corner solutions such as full or no disclosure. Our setup allows for considering such outcomes.

⁸ While our focus is on firms offering free services to users, note that this surplus can be intended net of any given subscription fee that all users might pay to the firm.

⁹ For example, consent forms might entail long and dry lists of privacy-related options, entailing significant personalization cost for users.

¹⁰ Assuming that consent popups also cause some fixed disutility further reduces the scope for them to protect consumers.

¹¹ Note that the resulting true net surplus becomes negative if $c > 2/[(2 - \beta)d]$, i.e. if privacy cost, disclosure and/or time inconsistency are high, while maximized perceived net surplus is always non-negative.

¹² This is also true with general utility $v(\cdot)$, as long as $v'(\cdot)$ is not too convex.

will normally give a positive weight $\alpha > 0$ to profits. We will refer to the resulting values as “user surplus” or “welfare”, respectively.¹³

3. The firm’s disclosure choice

Let us first assume that at $t = 1$ the firm selects the level of disclosure d . At $t = 2$, users observe the value of d and choose usage $x^*(d)$, which is the same for naive and sophisticated users. In fact, given that users do not choose the disclosure level, their awareness of their time inconsistency plays no role here and does not affect either profits or user surplus.

Let us denote with d_p the level of disclosure chosen by the firm, such that

$$d_p = \arg \max_d (\delta + \gamma \theta d) x^*(d).$$

The firm’s profit is concave in d . In fact, a higher level of disclosure has a positive effect in terms of higher revenue from a given amount of data (through the term $\gamma \theta d$), but also reduces usage $x^*(d)$ and hence the personal information provided by users. This problem has interior or corner solutions, depending on the value of β . Denoting with $\beta_1 = \frac{2\gamma\theta}{c(2\gamma\theta+\delta)}$, and $\beta_2 = \frac{2\gamma\theta}{c\delta}$, we obtain the following result:

Lemma 1. *The profit-maximizing level of disclosure is*

$$d_p = \begin{cases} 1 & \text{if } 0 \leq \beta \leq \beta_1, \\ \frac{1}{\beta c} - \frac{\delta}{2\gamma\theta} & \text{if } \beta_1 < \beta \leq \beta_2, \\ 0 & \text{if } \beta > \beta_2. \end{cases} \quad (2)$$

Both d_p and the firm’s profits $\pi(d_p)$ are decreasing in β .

The firm’s choice of disclosure level in Lemma 1 is continuous and (weakly) decreasing in β , i.e., the firm chooses a higher disclosure level if users are more time inconsistent. By doing so, the firm exploits the behavior of time inconsistent users, who react less to disclosure.

Expression (2) also highlights the role of the firm’s business model on its disclosure choice. In particular, d_p is increasing in γ and θ and decreasing in δ . If the value $\gamma\theta$ of exploited data is low, relative to the direct returns δ , the firm prefers to maximize usage revenues by reducing the disclosure of data. Still, full disclosure is only optimal if $\beta \leq \beta_1$, i.e. if users are sufficiently time inconsistent. For $\delta = 0$, the firm’s profits only arise from sharing data. In this case, the disclosure level is always positive and does not depend on the value of γ . Here the firm’s optimal strategy is analogous to that obtained by Jullien et al. (2020), in a context where users are rational but data sharing causes disutility from advertising.

We now compare the profit-maximizing disclosure level of the firm with the level d_w that would maximize user surplus or welfare given that users are time inconsistent, i.e.

$$d_w = \arg \max_d W(d) = v(x^*(d)) - cd x^*(d) + \alpha(\delta + \gamma d)x^*(d).$$

In determining d_w , if $\alpha > 0$, there is a trade-off similar to that of the firm, owing to the ambivalent effect that disclosure has on profits: A higher disclosure discourages usage but increases the revenues from data sharing. However, contrary to the firm, a social planner also takes into account the utility from usage and the associated privacy cost. Hence, the optimal level of disclosure depends on the social net gains from exploiting data, $\alpha\gamma - c$, i.e. the cost of privacy is not so high relatively to the value generated by data. Notably, when c is sufficiently large relative to $\alpha\gamma$, welfare is maximized at $d_w = 0$, for all β .¹⁴ There is no

¹³ Alternatively, α could be interpreted as the share of national, rather than foreign, firms. While it is helpful to keep this idea in mind, we do not follow it in the text.

¹⁴ Welfare is concave in d when $c > 2\alpha\gamma$ for all β , with $W(0) > W(1)$.

trade-off for the social planner in this case: the cost of privacy is so high, that closing the data business is best for consumers and society, so that the decision to implement or not consent policies in practice becomes pointless. To better highlight the nuanced effect of consent policies, in the rest of the analysis we thus assume that $\alpha\gamma > c$. Denoting with $\beta'_1 = 1 - \frac{\alpha(2\gamma+\delta)}{2c} + \frac{\sqrt{(\alpha\delta+2\alpha\gamma-2c)^2+8(\alpha\gamma-c)}}{2c}$ and $\beta'_2 = \frac{2(\alpha\gamma-c)}{\alpha c\delta}$, the welfare-maximizing level of disclosure is shown in the following Lemma.

Lemma 2. *The socially optimal level of disclosure is:*

$$d_w = \begin{cases} 1 & \text{if } 0 \leq \beta \leq \beta'_1, \\ \frac{1}{\beta c} - \frac{\alpha\delta+1}{2\alpha\gamma-(2-\beta)c} & \text{if } \beta'_1 < \beta \leq \beta'_2, \\ 0 & \text{if } \beta > \beta'_2. \end{cases} \quad (3)$$

Interestingly, time inconsistency has a counter-intuitive effect on socially optimal disclosure, as a higher level of time inconsistency (i.e., a lower β) increases d_w . This is due to the fact that time inconsistency decreases the social cost of raising d , in terms of lower direct revenues $\delta x^*(d)$. There is no gain from protecting users through a lower d_w , since their privacy costs are already fully taken into account.

Lemma (2) also highlights that the increase in the value γ of data is associated with higher socially optimal levels of disclosure, despite the privacy cost imposed on users. However, user surplus is highest with zero disclosure, as it encourages usage and eliminates the privacy cost.

We now compare the socially optimal level of disclosure d_w with the level d_p that the firm chooses. To this aim, let us denote with $\bar{\beta}_1 = \min\{\beta_1, \beta'_1\}$ and with $\bar{\beta}_2 = \max\{\beta_2, \beta'_2\}$. We find the following:

Proposition 1. *The firm’s choice of disclosure is socially optimal if $\beta \leq \bar{\beta}_1$ ($d_p = d_w = 1$) or $\beta \geq \bar{\beta}_2$ ($d_p = d_w = 0$). For $\beta \in (\bar{\beta}_1, \bar{\beta}_2)$, there exists a value $\bar{\theta} \geq 0$ such that $d_p > d_w$ if $\theta > \bar{\theta}$, and $d_p < d_w$ otherwise.*

A first key result from Proposition 1 is that the firm may either overdisclose or underdisclose data relative to the social optimum, depending on how much profits can appropriate from sharing data with third parties. If the firm’s share of profit from data exploitation is sufficiently high, it has the incentive to overdisclose data, as the firm does not internalize the increase in privacy cost caused by higher disclosure. On the contrary, underdisclosure occurs in the opposite case, as the firm fails to internalize the total value of data.

A second result emerging from Proposition 1 is that, under some conditions, the firm’s preferred level of disclosure may still coincide with the welfare maximizing one. In particular, if users’ time inconsistency is sufficiently low, the firm chooses $d_p = 0$ and this choice is optimal both for welfare and user surplus. Moreover, when time inconsistency is sufficiently high, the firm selects full disclosure and this choice is optimal for welfare –though not for user surplus– due to the returns from data. However, for intermediate levels of β , the firm discloses an inefficiently high or low amount of data.

Fig. 2 illustrates the optimal disclosure levels as a function of time inconsistency (on the horizontal axis), for increasing levels of γ from panel (a) to (b), in the case the firm appropriates a high share of profits from data (the overdisclosure case).¹⁵ The red curve represents the firm’s choice d_p , while the blue curve is the socially optimal level d_w . While the firm’s preferred level of disclosure is never below the welfare maximizing level, the increase in the value of data γ (or the decrease of δ) shifts rightwards both curves. Note also that when α decreases, the curve d_w moves towards the left, implying that it is less likely that the firm’s choice of full disclosure remains socially optimal.

¹⁵ The full parameter constellation used in the figure is: $\alpha = 0.3$, $\delta = 40$, $c = 0.9$, $\theta = 0.8$, $\gamma = 10$ in panel (a), and $\gamma = 15$ in panel (b).




Fig. 2. Profit-maximizing and socially optimal disclosure levels, with (a) $\gamma = 10$ or (b) $\gamma = 15$. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Overall, if the firm is allowed to set the level of disclosure and retains a sufficiently high share of profits from data monetization, it will disclose excessively, unless users' time inconsistency is sufficiently high or sufficiently low. It might be interesting to explore how the firm's business model affects the role of time inconsistency in the firm's strategy. We obtain the following Corollary.

Corollary 1. *The thresholds β_1, β_2 are all increasing in γ and decreasing in δ .*

Corollary 1 shows that the interval of β for which the firm overdiscloses data (when θ is sufficiently high) shifts to the right, towards higher values of β , when γ increases or δ decreases. While Proposition 1 finds that the firm does not profit from disclosing data when users are sufficiently time consistent, Corollary 1 highlights also that firms will disclose more data when they rely more on data revenues relative to direct sales.

Due to the firm's tendency to overdisclose data when θ is sufficiently high, and underdisclose them when θ is instead sufficiently low, there is a market failure that regulation (or self-regulation) could address. In the following we explore some common schemes and how well they fare under time inconsistency.

4. Delegating consent to users

Under the privacy legislation in Europe, users are given the right to give or deny their consent to the disclosure of their personal information. In this section, we thus assume that users choose their own disclosure level $d \in [0, 1]$: The firm must present users with options to personalize individual privacy settings. We now analyze how users react to these options, and whether they improve on the outcome when the firm chooses the disclosure level. As we will see, the level of the choice cost is decisive for the outcome.

According to the European GDPR, in order to be valid, consent must be both active and unambiguous, hence pre-checked boxes or bundles of expressions of intention do not constitute valid consent. As a consequence, the action of actively expressing their own preferences regarding privacy is costly for users. In practice, websites typically comply with these provisions by presenting users with the choice of either accepting full disclosure with immediate access to the website, or of going step-by-step through the privacy settings and having to make many choices before accessing the site. Thus, full disclosure is currently often the *de facto* default. We formalize this process by assuming that the personalization of privacy settings imposes a *choice cost* on users equal to $C(1 - d)$, with $C > 0$. Selecting full disclosure ($d = 1$) has lowest inconvenience (normalized to zero), whereas adopting more stringent privacy settings takes an increasing amount of time and attention. The

more stringent privacy settings (i.e., the lower d), the higher the cost $C(1 - d)$ of users for achieving protection.¹⁶

In Section 4.1 we show that users choose extreme disclosure levels for any exogenously given level of C . In Section 4.2 we consider the firm's incentives in selecting the value of C .

4.1. User choice under consent policies

At $t = 1$, users of type i , with $i = s, n$, believe that their future usage will be $\hat{x}_i(d_i)$ as defined above. We also assume that users make their usage decision based on the chosen disclosure level, independently of whether the firm would have preferred less disclosure. Users' expected net surplus is therefore given by

$$\hat{U}_i(d_i) = v(\hat{x}_i(d_i)) - cd_i\hat{x}_i(d_i) - C(1 - d_i).$$

In choosing d_i , users face a trade-off: On the one hand, low disclosure reduces the privacy cost and increases expected usage. On the other hand, achieving low disclosure requires the costly action of going through the privacy settings.

The two types of users differ regarding their expectation about usage \hat{x}_i in the following period. Sophisticated users anticipate that they will overconsume later on and correctly believe that $\hat{x}_s(d_s) = x^*(d_s)$, while naive users are ignorant of their time inconsistency and expect to provide $\hat{x}_n(d_n) = x^{tc}(d_n)$. As a result, users' expectation about future usage affects their present choice of disclosure level.

Proposition 2. *If consent is delegated to users, they either choose zero or full disclosure. Users of type i , $i = s, n$, choose zero disclosure if $C \leq C_i$, and full disclosure if $C \geq C_i$, where*

$$C_s = c - \frac{1}{4}(2 - \beta)\beta c^2 \geq C_n = c - \frac{1}{4}c^2.$$

Users prefer to read the privacy settings and set zero disclosure only if the choice cost C is sufficiently low. In this case their usage is $\hat{x}_i(0) = 1$. Conversely, if the cost of reading the privacy settings is sufficiently high, users choose full disclosure, i.e. $d_s = 1$ and lower expected usage $\hat{x}_i(1)$. The reduction in usage demand as a response to intermediate values of d is so strong that the latter are never optimal choices.

The threshold C_s for sophisticated users increases with time inconsistency (lower β): More time-inconsistent but sophisticated users are less willing to authorize disclosure, in order to protect themselves from their future behavior since they cannot commit to a lower usage level

¹⁶ We could assume here that intermediate privacy settings involve a higher choice cost than simply unclicking all boxes in order to select zero disclosure. As we will see below, this would not change our results.

Table 1
Conditions under which consent policies can improve welfare (W+) or user surplus (U+). Disclosure levels d_u of users, and d_p of firms, respectively.

	$d_p = 0$	$0 < d_p < 1$	$d_p = 1$
$d_u = 0$	-	$W+(\theta > \bar{\theta})$, U+	U+
$d_u = 1$	$W+(\theta < \bar{\theta})$	$W+(\theta < \bar{\theta})$	-

instead. Importantly, time inconsistency implies that users choose less disclosure only if they are aware of it. In fact, naive users have the same choice threshold as time-consistent users (corresponding to $\beta = 1$). Lack of awareness of their time inconsistency makes naive users more likely to stick to the *de facto* default of full disclosure than sophisticated users (i.e., $C_n \leq C_s$). Here the consent policy fails, because it is precisely these users that need more protection.

When disclosure is delegated to users, welfare is defined as $W(d) - C(1 - d)$. As seen in Proposition 1, the firm’s choice may be socially optimal, so it is *a priori* not clear that letting users choose disclosure raises welfare. The following Proposition lays out when consent policies do increase user surplus or welfare - and when they do not. It shows clearly that consent policies are only useful if the choice cost is low enough, so that users find it worthwhile to reduce disclosure.

Proposition 3. *The introduction of consent policies, relative to a situation in which the firm chooses disclosure, strictly increases welfare only if:*

- users restrict disclosure while the firm prefers partial disclosure, when $\theta > \bar{\theta}$;
- users prefer full disclosure while the firm prefers zero or partial disclosure, when $\theta < \bar{\theta}$.

Consent policies strictly increase user surplus only if users restrict disclosure while the firm prefers partial or full disclosure.

The following Table 1 summarizes the results of Proposition 3, highlighting the conditions under which consent policies can increase welfare, depending on the user’s and the firm’s disclosure decision.

The set of values (β, C) for which the consent policy raises welfare has a rather complex shape; it is neither convex nor connected, due to the extreme nature of users’ disclosure choices. Still, a robust conclusion is that the consent policy can have a positive effect on user surplus only if it makes users restrict disclosure (low C) while the firm would choose positive (partial or full) levels of disclosure. Moreover, when θ is sufficiently high, welfare is increased only if users deny disclosure and the firm’s choice of disclosure is partial. Still, if both users and the firm would choose no disclosure (i.e., β is high), user surplus and welfare are lower due to the wasteful choice cost. A different conclusion about welfare emerges when θ is low enough. In this case, consent policies can increase welfare only if users would choose full disclosure due to the high choice cost, while the firm would limit disclosure due to the limited appropriability of profits from data.

These results complement the findings by Conti and Reverberi (2021), who highlight how policies that allow the use of personal data only after the consent of the subjects may decrease consumer surplus due to a negative effect on product quality, provided that the complementarity between information and quality is weak. We show how consent policies can be surplus decreasing for an additional reason, specifically the inefficiency related to burdening consumer with a disclosure choice, when the firm partly internalizes the cost of disclosure.

Fig. 3 and 4 provide more intuition focusing on the case in which θ is sufficiently large. We denote welfare (gross of choice cost) at full disclosure as $W(1)$, with $W(1) = (1 + \alpha\delta + \frac{1}{2}\beta c - c + \alpha\gamma) (1 - \frac{1}{2}\beta c)$, and at zero disclosure as $W(0)$, with $W(0) = 1 + \alpha\delta$, respectively. Fig. 3 provides a graphical representation of the result stated in Proposition 3

when users restrict disclosure (i.e., $C < C_n$).¹⁷ The blue curve represents welfare W_u under the consent policy, while the red curve represents welfare W_p under firm choice of the disclosure level. As a reference level, the Figure also reports the welfare levels under zero disclosure (the dotted line) and full disclosure (the dashed curve). Panel (a) shows that, when θ is sufficiently high (i.e., the firm tends to overdisclose data), there exists an intermediate range of β such that welfare is higher under the consent policy than under firm choice. When the β is high, it is socially preferable to allow the firm to make the disclosure decision. This is because the firm chooses zero disclosure in order to spur usage (in fact, $W_p = W(0)$ for sufficiently high β , where welfare $W(0)$ is represented by the blue line). Users have to pay the choice cost to achieve the same outcome, hence the consent policy lowers welfare. When β is low, the firm exploits users’ time inconsistency and chooses full disclosure (i.e., the function W_p coincides with $W(1)$, represented by the red curve in the figure). Full disclosure is also the socially optimal choice if profits have a sufficiently high weight in the welfare function (panel (a)). Hence, the consent policy increases welfare only if time inconsistency is neither too high nor too low. However, a different conclusion emerges for user surplus (i.e., panel (b)). In fact, the consent policy increases user surplus when the firm discloses enough data, i.e., time inconsistency is sufficiently high.

Fig. 4 provides a graphical representation of the result stated in Proposition 3 when θ is sufficiently high and users do not restrict disclosure because choice cost is too high (i.e., $C \geq C_s$).¹⁸ The curve W_p representing welfare under firm choice is the same as in Fig. 3. However, due to full disclosure, the curve W_u for the consent policy differs from that of Fig. 3 and coincides with welfare under full disclosure $W(1)$. Now the consent policy leads to lower welfare and user surplus if β is sufficiently high: Having opted for full disclosure, contrary to what the firm would have chosen, users restrict usage while the firm fully shares their data. Still, welfare and user surplus do not change if time inconsistency is high, because also the firm chooses full disclosure. As a consequence, firm choice dominates (although not strictly) consent policies with high choice costs.

For intermediate levels of choice cost ($C_n \leq C < C_s$), the curves W_u are combinations of those depicted in Figs. 3 and 4, and results are qualitatively similar to those described in Fig. 3.

Interestingly, when the firm can appropriate a sufficient amount of profits (i.e., θ is high enough), the value of data γ has a counter-intuitive effect on welfare under the firm’s choice W_p . In fact, a higher value of data increases the firm’s incentive to adopt higher disclosure levels. If profits are sufficiently valuable in the welfare function, higher levels of disclosure are also the welfare maximizing choice, as it allows to internalize the value of data. It follows that, when γ increases, the region of β where $W_u > W_p$ shrinks. For sufficiently high γ , $W_p > W_u$ for all β . The opposite result however merges when the user surplus is considered. A higher value of data γ increases the range of β such that user consent is preferred for user surplus over firm’s choice.

4.2. Firm’s incentives to inflate choice cost

The analysis so far has taken the size and shape of the choice cost as given. However, the exact implementation of the consent scheme can itself be a policy variable. In this section we consider some dimensions of this design.

Several observations are in order. The current legal framework of the GDPR provisions and cookie law provide general guidelines for protecting consumer privacy, but the actual design and implementation of

¹⁷ The full parameter set used in the figure is: $\gamma = 10$, $\delta = 40$, $c = 0.6$, $C = 0.2$, $\lambda = 0.5$, $\theta = 1$ and $\alpha = 0.1$.

¹⁸ The full parameter set used in the figure is: $\gamma = 10$, $\delta = 40$, $c = 0.6$, $C = 1$, $\lambda = 0.5$, $\theta = 1$ and $\alpha = 0.1$.




Fig. 3. Welfare (panel a) and user surplus (panel b) when all users restrict disclosure under the consent policy ($C < C_n$) and θ is large.

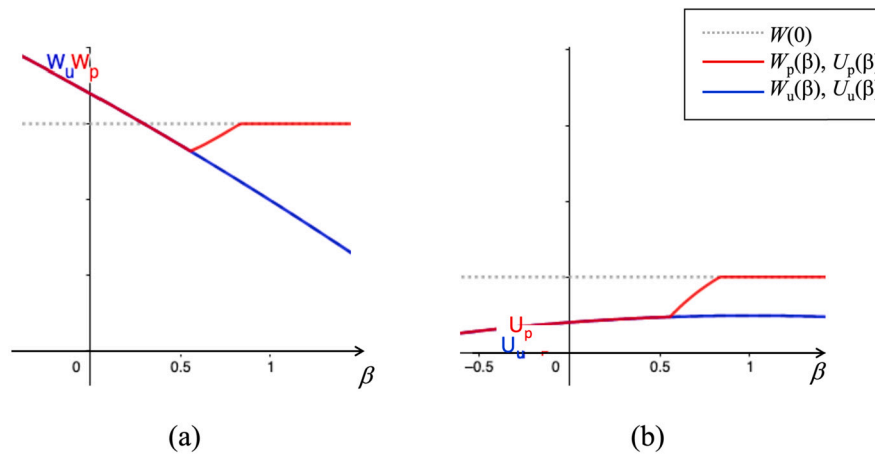


Fig. 4. Welfare (panel (a)) and user surplus (panel (b)) when users choose full disclosure under the consent policy ($C > C_s$) and θ is large.

consent forms is left to the websites' discretion.¹⁹ As a consequence, firms can currently choose how hard it is for a user to personalize the privacy settings. But other approaches can be imagined.

Let us first maintain the shape of the choice cost, with the default option of full disclosure, i.e. $C(1 - d)$. It might then seem obvious that the firm always has an incentive to cause a high (low) level of choice cost C when it tends to overdisclose (underdisclose) data. For example, it might influence the complexity of the choice procedure to make users choose its preferred level of disclosure. However, this may not be true in general. When the firm overdiscloses data, designing a high choice cost implies that it has to trade off higher revenues as a result of more disclosure with lower revenues due to less usage. Analogously, reducing the choice cost when it underdiscloses data means giving up revenues from data, but increasing those from usage. We find the following:

Proposition 4. *Given choice cost function $C(1 - d)$, the firm prefers any $C \leq C_n$ (all users select zero disclosure) if $\beta \geq \tilde{\beta} \equiv \frac{2\gamma\theta}{c(\delta+\gamma)}$, and any $C \geq C_s$ (all users select full disclosure) if $\beta \leq \tilde{\beta}$. For $\beta > 2\frac{\alpha\gamma - C_s}{ca(\delta+\gamma)}$, though, any $C > 0$ reduces welfare as compared to a costless disclosure choice.*

¹⁹ Art. 12 of GDPR states: "The controller shall take appropriate measures to provide any information referred to in Articles 13 and 14 and any communication under Articles 15 to 22 and 34 relating to processing to the data subject in a concise, transparent, intelligible and easily accessible form, using clear and plain language, [...]. The controller shall facilitate the exercise of data subject rights under Articles 15 to 22."

Since $\beta_1 < \tilde{\beta} < \beta_2$, the firm selects a low choice cost only if it would itself choose at most a low level of partial disclosure. As one would expect, the firm prefers that users select full disclosure if its profits from data exploitation are high enough. This will therefore depend on the users' time inconsistency and on the firm's business model. When users are mildly time inconsistent or the direct provision of the service and the corresponding commercial activities are highly profitable, the firm has an interest in helping users to protect their privacy, because by doing so they are more willing to use the service. Hence, the firm has an incentive to make users' choice of full protection easy. Conversely, when the firm obtains high profits from exploiting users' information, either because data have high value or users are severely time inconsistent, it has an incentive to raise the choice cost of users.

A second interesting point is that neither the firm nor society would want to have a choice process that makes sophisticated and naive users choose different levels of disclosure. The reason is that from both points of view their usage and costs have the same value.

More importantly, though, while the firm may not in all cases implement a choice procedure that leads to full disclosure, any procedure that involves a positive choice cost is inefficient if β is high enough. Furthermore, the above also implies that any consent policy with choice cost $C > 0$ is worse for users than one that is costless, or, equivalently, an imposition of zero disclosure. Thus there is a possible market failure that stems from the implementation of the consent provisions themselves, beginning with the *de facto* default option of full disclosure.

This leads to the next step, i.e., changing the design of the choice procedure itself. For example, following the lead of the California Con-

sumer Privacy Act (CPPA), platforms could be asked to facilitate users' decision to opt out by a clear "do not share" option. Indeed, either the firm or the legislator, depending on who has the power to decide, could in principle design a choice process (and therefore implicitly a choice cost function) that has a costless default at the respective preferred value and very high cost for any other level of disclosure. This process would therefore also exhibit the required level of convexity that is needed to make users select this preferred level over zero and full disclosure. While each firm could certainly do so depending on its specific characteristics, as described in our model by the parameters (δ, γ) , it will generally not be possible for governments or regulatory agencies to devise a specific default and process for each firm, because of the immense complexity of doing so and firms' private knowledge about their characteristics and profitability. Therefore a choice process to be described by legislation must first of all be simple to design and applicable to a large variety of firms.

The simplest alternative to a default setting of full disclosure is one with a default of zero disclosure. In our setting, the above discussion shows that the immediate outcome will be that no user would want to change these default settings if service provision was not allowed to depend on this choice, as under the CPPA. The latter condition must be imposed to prevent firms from circumventing the default disclosure option by offering two service tiers, a basic one without disclosure and full service with disclosure. The effect of this no-disclosure default is then to hinder data monetization, even when it would give rise to welfare gains. Thus there is no simple way out of this dilemma using consent policies.

5. Extensions

5.1. Voluntary usage caps

In the case firms have the incentive to overdisclose data, they face a dilemma. Firms are of course aware that they disclose more information than users would like and that this makes the latter limit usage or restrict disclosure. They cannot increase their profits by committing to some disclosure level, since the disclosure levels discussed in Section 3 already maximize these profits.

Rather, in this section we consider a different type of firm self-regulation, where the firm offers its users the option of committing to a maximum level of usage by accepting a cap on future usage. Examples of this type of policies are time limits on Facebook, Screen Time settings on iPhones, and Google Digital Wellbeing on Android phones.^{20,21} We show that the firm can indeed increase its profits by offering such a cap, but only when it is offered to users that can control the level of disclosure and are aware of their time inconsistency.

Suppose that the firm offers a cap on usage and also chooses its level.²² Assume that at $t = 0$ the firm offers users a voluntary cap on usage with level \bar{x} . Users decide at $t = 1$ whether to accept the cap or not, and we assume that they will only accept the cap if their expected utility strictly increases by doing so.²³ At $t = 2$ users choose usage, which cannot exceed the cap agreed to previously.

We consider two scenarios concerning disclosure, either letting the firm choose the disclosure level simultaneously with the usage cap, or

²⁰ See <https://about.fb.com/news/2018/08/manage-your-time>, and <https://wellbeing.google>.

²¹ In the UK, mobile phone operators were required to offer caps on bills or usage, see <https://www.ofcom.org.uk/phones-telecoms-and-internet/advice-for-consumers/costs-and-billing/mobile-bill-limits>.

²² If the firm lets users choose the level of the cap themselves its profits are lower.

²³ This assumption captures a *status quo* bias where taking an action rather than doing nothing has some psychological cost, such as potential future regret of having taken the wrong action.

letting users set the disclosure level under a consent policy after accepting the cap. In the first case we find the following:

Proposition 5. *If the firm sets the disclosure level, it will not offer a binding usage cap.*

This Proposition shows that if the firm sets both the levels of the cap and of disclosure then the former is pointless: The cap simply reduces its revenues without creating any benefits for the firm. Note that the same result would hold if the privacy regulation were to set a disclosure level beforehand, since the level of d has no effect on the outcome (as long as it is a given for users).

In our setting, the firm might want to offer a binding cap to influence users' disclosure decision²⁴. Under a binding cap, users might be more willing to set a high disclosure rate. So let us now consider the case where, after observing the firm's offer \bar{x} , users simultaneously decide whether to accept the cap and set a disclosure level d_i . We find the following:

Proposition 6. *For $C_n < C \leq C_s$ and γ high enough, the firm can increase its profits by offering a binding cap on usage that makes all users adopt full disclosure. For other values of C offering a cap does not increase profits. When θ is sufficiently high, the firm's offer of a cap reduces welfare if γ is relatively low, but otherwise increases it. Conversely, when θ is sufficiently low, the firm offering the cap always increases welfare.*

Thus it is exactly when the choice cost C , for some exogenous reason, has a value that the firm would never choose (see Proposition 4), that there is indeed a potential gain for the firm to offer a cap on usage. This gain arises exclusively from being able to make users that know about their time inconsistency adopt less restrictive privacy settings, while naive users simply ignore the offered cap.

Introducing the cap reduces welfare if the value of data γ is low and θ is high. The reason is that in this case zero disclosure is socially preferable to full disclosure. If γ is high enough, though, then the offer of a cap does increase welfare, as it leads to a better balance between usage and disclosure of sophisticated users.

5.2. Privacy cost of data collection

In our baseline model, users suffer a privacy cost only when their data is disclosed to third parties. In practice, however, users might bear a privacy risk also from the collection of information by the firm, even when it is not shared with third parties. For example, the firm could exploit the information collected to extract surplus from the consumer, or a data breach stealing information stored on the firm's information system causes a loss of privacy for users. In this section we explore the implications of a privacy cost arising both from collection and from disclosure of user information, by assuming that the function representing the privacy cost is $c(1+d)x$, with $c \in (0, 1)$.

When using the service at $t = 2$, users choose usage x to maximize perceived net surplus $v(x) - \beta c(1+d)x$, with²⁵

$$x^{*CD}(d) = 1 - \frac{\beta c}{2}(1+d) > 0, \quad (4)$$

where the superscript CD denotes the result in the case the privacy cost arises both from data collection and data disclosure to third parties. When the privacy loss stems not only from disclosure, but also from data

²⁴ In our model, there are no other behavioral dimensions. Firms might also offer caps to create goodwill with the public and regulators, or to make parents allow their children to use the firm's services, as with Google's Family Link.

²⁵ As in our baseline model, maximized perceived net surplus is always non-negative, while the true net surplus becomes negative if $c > 2/[(2-\beta)(1+d)]$.

collection, usage is lower as users face a greater privacy risk. Nonetheless, the difference between (4) and (1) disappears as users become more time-inconsistent.

The following lemma identifies the firm's optimal level of disclosure when the privacy loss also stems from data collection.

Lemma 3. *Given the privacy cost function $c(1+d)x$, the profit-maximizing level of disclosure is*

$$d_p^{CD} = \begin{cases} 1 & \text{if } 0 \leq \beta \leq \beta_1^{CD}, \\ \frac{1}{\beta c} - \frac{1}{2} - \frac{\delta}{2\gamma\theta} & \text{if } \beta_1^{CD} < \beta \leq \beta_2^{CD}, \\ 0 & \text{if } \beta > \beta_2^{CD}, \end{cases} \quad (5)$$

with $\beta_1^{CD} = \frac{2\gamma\theta}{c(3\gamma\theta+\delta)}$, and $\beta_2^{CD} = \frac{2\gamma\theta}{c(\gamma\theta+\delta)}$. Both d_p^{CD} and the firm's profit $\pi(d_p^{CD})$ are continuous and decreasing in β .

Similar to the results in our baseline model, the firm's optimal disclosure level and profits decrease in β . However, Lemma 3 shows that now the firm has the incentive to further restrict disclosure relative to (2), as disclosure has a stronger crowding-out effect on usage.

We now study the welfare-maximizing level of disclosure, defined as

$$d_w^{CD} = \arg \max_d W(d) \\ = v(x^{*CD}(d)) - c(1+d)x^{*CD}(d) + \alpha(\delta + \gamma d)x^{*CD}(d).$$

The socially optimal level of disclosure is shown in the following lemma.

Lemma 4. *Given the privacy cost function $c(1+d)x$, there exist two values $\beta_1^{CD}, \beta_2^{CD} > 0$, such that: $d_w^{CD} = 1$ if $\beta \leq \beta_1^{CD}$; $d_w^{CD} = \frac{1}{\beta c} - \frac{1}{2} - \frac{\alpha\delta+1+\frac{\beta c}{2}-c}{2\alpha\gamma+\beta c-2c}$ if $\beta \in (\beta_1^{CD}, \beta_2^{CD})$; $d_w^{CD} = 0$ if $\beta > \beta_2^{CD}$.*

As in the case examined in our baseline model, the socially optimal level of disclosure entails full disclosure of information when users have high time inconsistency, while partial or no disclosure for lower levels of time inconsistency. However, compared to a situation in which data collection does not generate a privacy loss, we have that $d_w^{CD} = d_w - \frac{1}{2} \left(1 + \frac{\frac{\beta c}{2} - c}{\alpha\gamma + \frac{\beta c}{2} - c} \right)$, i.e. $d_w^{CD} < d_w$. The socially optimal level of disclosure tends to be lower than in the case the privacy loss stems only from data sharing, because the social planner takes into account the additional risk to privacy.

We now compare the socially optimal level of disclosure d_w^{CD} with the level d_p^{CD} that the firm chooses. Let us define

$$\bar{\theta}^{CD} = \min \left\{ \frac{\delta}{2\gamma} \frac{2\alpha\gamma + \beta c - 2c}{\alpha\delta + 1 + \frac{\beta c}{2} - c}, 1 \right\}.$$

We find the following:

Proposition 7. *The platform's choice of disclosure is socially optimal if β is sufficiently low ($d_p^{CD} = d_w^{CD} = 1$) or β is sufficiently high ($d_p^{CD} = d_w^{CD} = 0$). For intermediate levels of β , $d_p^{CD} > d_w^{CD}$ if $\theta > \bar{\theta}^{CD}$, and $d_p^{CD} < d_w^{CD}$ otherwise, with $\bar{\theta}^{CD} > \bar{\theta}$.*

The results of Proposition 7 are qualitatively similar to those obtained in our baseline model. The reason is that the fundamental trade-off between more usage and more privacy is not altered by the fact that the privacy cost arises not only from sharing data with third parties, but also from the data collection itself. Interestingly, however, Proposition 7 also shows that now there is a higher scope for firm's underdisclosure, as $\bar{\theta}^{CD} > \bar{\theta}$. Intuitively, the fact that the privacy loss stems also from data collection has a larger (negative) effect on the firm's choice of disclosure than on the social planner's. Hence, for any given θ , the firm is

more likely to underdisclose data than in the case the privacy loss only comes from data sharing.

We now evaluate the conditions under which users choose the disclosure.

Proposition 8. *If consent is delegated to users, they either choose zero or full disclosure. Users of type i , $i = s, n$, choose zero disclosure if $C \leq C_i^{CD}$, and full disclosure if $C \geq C_i^{CD}$, where*

$$C_s^{CD} = c - \frac{3}{4}(2 - \beta)\beta c^2 \geq C_n^{CD} = c - \frac{3}{4}c^2.$$

Similar to our baseline model, users disclose their data only if their choice cost is sufficiently high, and naive users are more likely to choose full disclosure than sophisticated users. Interestingly, however, both the thresholds C_s^{CD} and C_n^{CD} are lower than in our baseline model. This implies that users have a higher tendency to choose full disclosure when they are more exposed to the privacy risk, which now stems not only from data sharing, but also from data collection. This counterintuitive result is due to the fact that users now suffer a privacy cost even if they do not share data, so that the decision to restrict disclosure is relatively less advantageous than in the case the privacy loss stems only from data sharing.

6. Conclusions

Our internet browsing activity is constantly interrupted by pop-ups asking us to agree to websites' privacy policies. We often consent immediately to full disclosure because these banners contain too many options and take time to browse. However, this haste to give consent is costly too, in terms of privacy loss. Time inconsistency in neglecting future risks to our privacy makes us more prone to skip the protection offered by privacy law.

How do firms' information disclosure strategies depend on their business model? How does time inconsistency affect these strategies? Are consent-based mechanisms more desirable when users are not fully rational? In this paper, we study the complex interplay between users' time inconsistency, the value of data, and the rules that frame information disclosure. We find two main results.

First, whenever firms disclose information, from a user perspective they disclose too much. They also disclose too much information from the perspective of society as a whole (considering both user surplus and profits) unless the value of data to third parties or time inconsistency are high. Both the level of disclosure and firm profits increase with the degree of users' time inconsistency.

Second, the introduction of consent mechanisms without paying attention to their implementation may be worse for users and welfare than not introducing consent mechanisms at all. The reason is that firms that monetize data may have an incentive to increase the cost (i.e., the complexity) for users to deny their consent. Therefore, from a policy-making point of view, our findings suggest a revision of current privacy legislation to refine the rules under which consent forms are designed. Some countries have already taken steps in this direction. For instance, California recently required platforms to offer users an easy opt-out option. Our analysis suggests that such a solution protects users but reduces total value to society. On the other hand, EU privacy regulation does not provide rules for designing consent forms, resulting in substantial heterogeneity across websites and over-complex privacy settings.

CRedit authorship contribution statement

Laura Abrardi: Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization. **Carlo Cambini:** Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization. **Steffen Hoernig:** Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization.

Declaration of competing interest

None.

Data availability

No data was used for the research described in the article.

Appendix A. Proofs

Proof of Lemma 1. The firm solves the problem

$$\max_{d \in (0,1)} \pi(d) = (\delta + \gamma\theta d) \left(1 - \frac{\beta cd}{2}\right).$$

The interior solution of the problem solves the f.o.c. $\gamma\theta \left(1 - \frac{\beta cd}{2}\right) - (\delta + \gamma\theta d) \frac{\beta c}{2} = 0$, from which we obtain $d_p = \frac{1}{\beta c} - \frac{\delta}{2\gamma\theta}$, continuous and monotonically decreasing in β . Moreover, $d_p \geq 0$ iff $\frac{1}{\beta c} - \frac{\delta}{2\gamma\theta} \geq 0$, i.e. $\beta \leq \beta_2 = \frac{2\gamma\theta}{c\delta}$. Then, $d_p = 0$ for all $\beta > \beta_2$. Finally, $d_p \leq 1$ iff $\frac{1}{\beta c} - \frac{\delta}{2\gamma\theta} \leq 1$, i.e. $\beta \geq \beta_1 = \frac{2\gamma\theta}{c(2\gamma\theta + \delta)}$. Then, $d_p = 1$ for all $\beta \leq \beta_1$. We can thus conclude that d_p is weakly decreasing in β .

By substituting the profit-maximizing level of disclosure d_p in the profit function, we immediately obtain:

$$\pi(d_p) = \begin{cases} (\delta + \gamma\theta) \left(1 - \frac{\beta c}{2}\right) & \text{if } 0 \leq \beta \leq \beta_1, \\ \frac{(2\gamma\theta + \beta c \delta)^2}{8\gamma\theta\beta c} & \text{if } \beta_1 < \beta \leq \beta_2, \\ \delta & \text{if } \beta > \beta_2, \end{cases}$$

which are continuous and decreasing in β .

Proof of Lemma 2. Let us express welfare as $W(d) = x^*(2 - x^*) + \alpha\delta x^* + (\alpha\gamma - c)d x^*$. The first derivative of $W(d)$ is: $W'(d) = x^{*'}(2 - x^* + \alpha\delta + (\alpha\gamma - c)d) + x^*(-x^{*'} + \alpha\gamma - c)$, where $x^{*'} = -\frac{\beta c}{2}$. The second derivative is $W'' = -\beta c \left(\frac{\beta c}{2} + \alpha\gamma - c\right)$, implying that the function is always concave in d given that $\alpha\gamma - c > 0$. To find its maximum, we compute the f.o.c.:

$$-\frac{\beta c}{2} \left(1 + \frac{\beta cd}{2} - cd + \alpha(\delta + \gamma d)\right) + \left(1 - \frac{\beta cd}{2}\right) \left(\frac{\beta c}{2} - c + \alpha\gamma\right) = 0.$$

Solving by d , we obtain:

$$d_w = \frac{1}{\beta c} - \frac{\alpha\delta + 1}{2\alpha\gamma - (2 - \beta)c}, \tag{6}$$

which is continuous and decreasing in β . We can rewrite (6) as $d_w = \frac{1}{\beta c} \left(1 - \frac{\beta c(\alpha\delta + 1)}{2(\alpha\gamma - c) + \beta c}\right)$. Since $d_w \in [0, 1]$, we have that $d_w = 0$ when $1 - \frac{\beta c(\alpha\delta + 1)}{2(\alpha\gamma - c) + \beta c} < 0$, i.e. $\beta > \beta'_2 = \frac{2(\alpha\gamma - c)}{c\alpha\delta}$.

Moreover, $d_w = 1$ when $\frac{1}{\beta c} - \frac{\alpha\delta + 1}{2\alpha\gamma - (2 - \beta)c} \geq 1$, i.e. $\frac{2\alpha\gamma - (2 - \beta)c - \beta c\alpha\delta - \beta c - \beta c(2\alpha\gamma - (2 - \beta)c)}{\beta c(2\alpha\gamma - (2 - \beta)c)} \geq 0$. The previous condition can be rewritten as $\frac{2\alpha\gamma - 2c - \beta c\alpha\delta - \beta c(2\alpha\gamma - (2 - \beta)c)}{\beta c(2\alpha\gamma - (2 - \beta)c)} \geq 0$, which is verified iff the numerator is positive. In turn, the numerator is positive for $\beta \in \left[\frac{2c - \alpha(2\gamma + \delta) - \sqrt{(\alpha\delta + 2\alpha\gamma - 2c)^2 + 8(\alpha\gamma - c)}}{2c}, \frac{2c - \alpha(2\gamma + \delta) + \sqrt{(\alpha\delta + 2\alpha\gamma - 2c)^2 + 8(\alpha\gamma - c)}}{2c} \right]$.

However, $\frac{2c - \alpha(2\gamma + \delta) - \sqrt{(\alpha\delta + 2\alpha\gamma - 2c)^2 + 8(\alpha\gamma - c)}}{2c} < 0$. We thus conclude that $d_w = 1$ for any $\beta \in [0, \beta'_1]$, with $\beta'_1 = \frac{2c - \alpha(2\gamma + \delta)}{2c} + \frac{\sqrt{(\alpha\delta + 2\alpha\gamma - 2c)^2 + 8(\alpha\gamma - c)}}{2c}$.

Proof of Proposition 1. From Lemmas 1 and 2, we immediately obtain that $d_p = d_w = 1$ when $\beta \leq \min\{\beta_1, \beta'_1\}$, and that $d_p = d_w = 0$ when $\beta \geq \max\{\beta_2, \beta'_2\}$.

For $\beta \in (\min\{\beta_1, \beta'_1\}, \max\{\beta_2, \beta'_2\})$, using the expressions in Lemmas 1 and 2 for $d_p, d_w \in (0, 1)$, we can rewrite the condition $d_p > d_w$

as: $\frac{1}{\beta c} - \frac{\delta}{2\gamma\theta} > \frac{1}{\beta c} - \frac{\alpha\delta + 1}{2\alpha\gamma - (2 - \beta)c}$. After straightforward simplifications, the latter can be expressed as $\theta > \frac{\delta}{2\gamma} \frac{2\alpha\gamma + \beta c - 2c}{\alpha\delta + 1} = \bar{\theta}'$.

If $\theta > \bar{\theta}'$, we have that $\beta_1 > \beta'_1$ and $\beta_2 > \beta'_2$. This implies that $\min\{\beta_1, \beta'_1\} = \beta'_1$ and $\max\{\beta_2, \beta'_2\} = \beta_2$. Then, the interval in which the platform makes a socially suboptimal choice is $\beta \in (\beta'_1, \beta_2)$. In particular, for $\theta > \bar{\theta}'$, $d_p > d_w$ for all $\beta \in (\beta'_1, \beta_2)$.

If instead $\theta < \bar{\theta}'$, we have that $\beta_1 < \beta'_1$ and $\beta_2 < \beta'_2$. This implies that $\min\{\beta_1, \beta'_1\} = \beta_1$ and $\max\{\beta_2, \beta'_2\} = \beta'_2$. Then, the interval in which the platform makes a socially suboptimal choice is $\beta \in (\beta_1, \beta'_2)$. In particular, for $\theta < \bar{\theta}'$, $d_p < d_w$ for all $\beta \in (\beta_1, \beta'_2)$.

Proof of Corollary 1. From the expression of $\beta_2 = \frac{2\gamma\theta}{c\delta}$, we immediately find that β_2 is increasing in γ and decreasing in δ .

Consider now the expression of β'_1 . The derivative of β'_1 w.r.t. γ is

$$\frac{d\beta'_1}{d\gamma} = \frac{\alpha}{c} \left(\frac{\alpha\delta + 2\alpha\gamma - 2c + 2}{\sqrt{(\alpha\delta + 2\alpha\gamma - 2c)^2 + 8(\alpha\gamma - c)}} - 1 \right). \tag{7}$$

This expression is positive iff $4 + 4\alpha\delta \geq 0$, which is always verified.

The derivative of β'_1 w.r.t. δ for $\alpha\gamma - c \geq 0$ is

$$\frac{d\beta'_1}{d\delta} = \frac{\alpha}{2c} \left(\frac{\alpha\delta + 2\alpha\gamma - 2c}{\sqrt{(\alpha\delta + 2\alpha\gamma - 2c)^2 + 8(\alpha\gamma - c)}} - 1 \right). \tag{8}$$

This expression is negative iff $\alpha\gamma - c \geq 0$.

Proof of Proposition 2. Sophisticated users choose disclosure d_s to maximize $\hat{U}_s(d_s)$, with $\hat{x}_s = x^*$. Since this utility is convex in d_s , they either choose zero or full disclosure, with

$$\hat{U}_s(0) = 1 - C; \hat{U}_s(1) = 1 - c + \frac{1}{4}(2 - \beta)\beta c^2.$$

Zero disclosure ($d_s = 0$) is optimal for sophisticated users if $\hat{U}_s(0) \geq \hat{U}_s(1)$, or $C \leq C_s = c - \frac{1}{4}(2 - \beta)\beta c^2$, and $d_s = 1$ is optimal for $C \geq C_s$.

Naive users maximize $\hat{U}_n(d_n)$, with $\hat{x}_n = x^{tc}$. Since also this utility is convex in d_n , we obtain $d_n = 0$ if $C \leq C_n = c - \frac{1}{4}c^2$ and $d_n = 1$ for $C \geq C_n$.

Proof of Proposition 3. Welfare levels under zero and full disclosure, before choice cost, are $W(0) = 1 + \alpha\delta$ and $W(1) = \left(1 - \frac{1}{2}\beta c\right) (1 + \alpha\delta + \frac{1}{2}\beta c - c + \alpha\gamma)$, respectively. Under the consent policy, welfare is

$$W_u = \begin{cases} W(0) - C & \text{if } 0 < C \leq C_n, \\ \lambda(W(0) - C) + (1 - \lambda)W(1) & \text{if } C_n \leq C \leq C_s, \\ W(1) & \text{if } C_s \leq C, \end{cases}$$

which is non-continuous in C at the thresholds. When the firm chooses the level of disclosure, welfare becomes,

$$W_p = \begin{cases} W(0) & \text{if } \beta > \beta_2, \\ \tilde{W} \equiv \left(\frac{1}{2} + \frac{\beta c \delta}{4\gamma\theta}\right) \left(\frac{3}{2} - \frac{\beta c \delta}{4\gamma\theta} - \frac{1}{\beta} + \frac{c\delta}{2\gamma\theta} + \alpha\delta \left(1 - \frac{1}{2\theta}\right) + \frac{\alpha\gamma}{\beta c}\right) & \text{if } \beta_1 < \beta \leq \beta_2, \\ W(1) & \text{if } \beta \leq \beta_1, \end{cases}$$

which is continuous. In particular, the function \tilde{W} is convex and non-monotonic in β if $\alpha\gamma > c$. We can now identify the conditions under which $W_p > W_u$. We distinguish between two subcases.

Case 1: $\theta > \bar{\theta}$. This implies that $d_p = d_w = 1$ for $\beta \leq \beta'_1$, $d_p > d_w$ for $\beta \in (\beta'_1, \beta_2]$, and $d_p = d_w = 0$ for $\beta > \beta_2$. In this case, $d_p \geq d_w$ for all β .

If $d_u = 0$, given that $d_w \leq d_p$ for all β , we have that: if $d_p = 0$, $W_p = W(0) > W_u = W(0) - C$; if $d_p = 1$, then also $d_w = 1$, implying $W(1) > W(0)$, and in turn $W_u = W(0) - C < W(1) = W_p$; if $d_p \in (0, 1)$, welfare might be higher or lower under user consent. In particular, when

β is sufficiently close but smaller than β_2 , $W_p = \widetilde{W} < W(0)$. Then, there always is a range of $\beta < \beta_2$ and $C < C_n$ for which $W_u = W(0) - C > W_p$.

If $d_u = 1$ and $d_p = 1$, $W_u = W_p = W(1)$. If $d_u = 1$ and $d_p < 1$, $W_u = W(1) < W_p = \widetilde{W}$, as $d_w < d_p$.

Case 2: $\theta < \bar{\theta}$. This implies that $d_p = d_w = 1$ for $\beta \leq \beta_1$, $d_p < d_w$ for $\beta \in (\beta_1, \beta_2]$, and $d_p = d_w = 0$ for $\beta > \beta_2'$. In this case, $d_w \geq d_p$ for all β .

If $d_u = 0$, given that $d_w \geq d_p$ for all β , we have that $W_p > W_u = W(0) - C$ for all d_p (users under-disclose data relative to the social optimum, and in addition they have to pay the choice cost).

If $d_u = 1$ and $d_p = 1$, $W_p = W_u$. If $d_u = 1$ and $d_p < 1$, delegating consent to users increases welfare only if β is large but sufficiently close to β_1 (as $d_p < 1$ but $d_w = d_u = 1$). When instead $\beta > \beta_2'$ $d_p = d_w = 0$, so that $W_p > W_u$. Then, for all $\beta \in [\beta_1, \beta_2']$, delegating consent to users might either increase or decrease welfare. We conclude that delegating consent to users can increase welfare only if $d_u = 1$ and $d_p < 1$.

Let us now consider the effects on user surplus. User surplus $U(x^*(d)) = v(x^*) - cd x^*$ is maximized for zero levels of disclosure, as $U' = v' x^* - c < 0$ ($U(x^*(d))$ is always decreasing in d). Then, delegating consent to users can improve user surplus only if $d_u = 0$ and $d_p > 0$. If instead $d_u = 1$, $d_u \geq d_p \geq 0$ for all β , so that $U_p \geq U_u$ (with the equality sign for $d_p = d_u = 1$), with U_p, U_u denoting user surplus under the firm's or user's disclosure choice, respectively. Finally, if $d_u = d_p = 0$, $U_u = U(0) - C < U(0) = U_p$.

Proof of Proposition 4. The firm's profits under the consent policy are

$$\pi = \begin{cases} \delta & \text{if } 0 \leq C \leq C_n, \\ \lambda\delta + (1 - \lambda)(\delta + \gamma\theta)x^*(1) & \text{if } C_n < C \leq C_s, \\ (\delta + \gamma\theta)x^*(1) & \text{if } C > C_s, \end{cases}$$

which increases with C if $(\delta + \gamma\theta)x^*(1) > \delta$ or equivalently $\beta < \tilde{\beta} = \frac{2\gamma\theta}{c(\delta + \gamma\theta)}$. Thus for $\beta \geq \tilde{\beta}$ the firm chooses some level $C < C_n$, resulting in $d_s = d_n = 0$ and high usage, and for $\beta \leq \tilde{\beta}$ some level $C > C_s$ with $d_s = d_n = 1$.

The welfare-maximizing level of C is $C = 0$ if $W(0) \geq W(1)$, i.e. $\beta \geq 2 \frac{\alpha\gamma - C_s}{c\alpha(\delta + \gamma)}$.

Proof of Proposition 5. Suppose by contradiction that the cap is binding, i.e. $\bar{x} \leq x^*(d)$. When sophisticated users decide whether to accept the cap at $t = 1$, given \bar{x} and d , they maximize true net surplus $v(x) - cd x$ over $x \in \{\bar{x}, x^*(d)\}$, with

$$v(x^*(d)) - cd x^*(d) = \frac{1}{4}(2 - \beta dc)(2 - (2 - \beta)dc),$$

$$v(\bar{x}) - cd \bar{x} = (2 - \bar{x})\bar{x} - cd \bar{x}.$$

The latter is concave in \bar{x} and decreasing towards $x^*(d)$, therefore sophisticated users accept caps $\bar{x} \in (1 - (1 - \beta/2)dc, x^*(d))$: There are caps with $\bar{x} < x^*(d)$ that strictly increase sophisticated users' net surplus, unless $d = 0$ or $\beta = 1$.

By the same argument, naive users will only accept caps $\bar{x} \in (x^{lc}(d), x^{lc}(d))$, i.e., they will not accept any cap, independently of whether it restricts their usage or not. This is because they believe that they are time consistent and therefore do not need a cap at all. Their actual usage will therefore be $x^*(d)$.

The firm's profits from users of type i , $i = s, n$, are $(\delta + \gamma\theta d)x_i$, increasing in x_i for all levels of d . Naive usage is unaffected by any cap on offer, while sophisticated users only accept caps at some $\bar{x} < x^*(d)$ that lead to lower profits than offering a non-binding cap $\bar{x} \geq x^*(d)$. Thus a profit-maximizing cap is not binding.

Proof of Proposition 6. If sophisticated users do not accept \bar{x} , as before their optimal choices of disclosure are $d_s = 0$ with net surplus $1 - C$ if $C \leq C_s$, and $d_s = 1$ for $C \geq C_s$. The case $C \leq C_s$ is thus the relevant one, because the cap would only be offered to increase d_s .

If sophisticated users accept \bar{x} , then the disclosure level $d_s = 1$ maximizes $v(\bar{x}) - cd \bar{x} - C(1 - d)$ if $\bar{x} \leq C/c$, with net surplus $(2 - \bar{x})\bar{x} - c\bar{x}$. Accepting such a cap \bar{x} and choosing $d_s = 1$ maximizes (and strictly increases) users' net surplus if

$$\bar{x} \in \left(1 - c/2 - \sqrt{C - C_n}, 1 - c/2 + \sqrt{C - C_n}\right),$$

which is both non-empty and below C/c for $C > C_n$, and empty for $C \leq C_n$. Thus if and only if $C_n < C \leq C_s$ there are caps that increase sophisticated users' net surplus while making them switch from zero to full disclosure. Over this range of user cost, naive users choose full disclosure and have usage $x^*(1)$ at $t = 2$.

The firm's profits are $\pi = (\delta + \gamma\theta)[\lambda\bar{x} + (1 - \lambda)x^*(1)]$. Again, this is increasing in \bar{x} , so the firm can choose some \bar{x} close to but below $\bar{x}^* = 1 - c/2 + \sqrt{C - C_n}$ (which is below $x^*(1)$ for $C < C_s$). This increases the firm's profit if $(\delta + \gamma\theta)\bar{x}^* > \delta$, or $\gamma > \frac{\delta}{\bar{x}^*\theta} - \delta/\theta$.

At cap \bar{x}^* and $d_s = 1$ welfare is $v(\bar{x}^*) + \alpha(\delta + \gamma - c)\bar{x}^*$, which is lower than welfare at zero disclosure (i.e., $1 + \alpha\delta$) if

$$\gamma < \frac{1 + \alpha\delta - v(\bar{x}^*)}{\alpha\bar{x}^*} + c - \delta. \tag{9}$$

If θ is sufficiently high, this threshold is higher than $\delta/(\theta\bar{x}^*) - \delta/\theta$, thus for γ in this range the firm's introduction of the cap lowers welfare. Conversely, for θ low enough, the threshold in (9) is lower than $\delta/(\theta\bar{x}^*) - \delta/\theta$, implying that when $\gamma > \frac{\delta}{\bar{x}^*\theta} - \delta/\theta$ the introduction of the cap increases welfare.

Proof of Lemma 3. The firm solves the problem

$$\max_{d \in [0,1]} \pi(d) = (\delta + \gamma\theta d) \left(1 - \frac{\beta c(1+d)}{2}\right),$$

whose interior solution can be obtained from the f.o.c. $\gamma\theta \left(1 - \frac{\beta c(1+d)}{2}\right) - (\delta + \gamma\theta d) \frac{\beta c}{2} = 0$, from which we find $d_p^{CD} = \frac{1}{\beta c} - \frac{1}{2} - \frac{\delta}{2\gamma\theta}$, continuous and monotonically decreasing in β . Moreover, $d_p^{CD} \geq 0$ iff $\frac{1}{\beta c} - \frac{1}{2} - \frac{\delta}{2\gamma\theta} \geq 0$, i.e. $\beta \leq \beta_2^{CD} = \frac{2\gamma\theta}{c(\delta + \gamma\theta)}$. Then, $d_p = 0$ for all $\beta > \beta_2^{CD}$. Finally, $d_p^{CD} \leq 1$ iff $\frac{1}{\beta c} - \frac{1}{2} - \frac{\delta}{2\gamma\theta} \leq 1$, i.e. $\beta \geq \beta_1^{CD} = \frac{2\gamma\theta}{c(3\gamma\theta + \delta)}$. Then, $d_p^{CD} = 1$ for all $\beta \leq \beta_1^{CD}$.

By substituting the profit-maximizing level of disclosure d_p^{CD} in the profit function, we immediately obtain:

$$\pi(d_p^{CD}) = \begin{cases} (\delta + \gamma\theta)(1 - \beta c) & \text{if } 0 \leq \beta \leq \beta_1^{CD}, \\ \frac{(2\gamma\theta + \beta c(\delta - \gamma\theta))^2}{8\gamma\theta\beta c} & \text{if } \beta_1^{CD} < \beta \leq \beta_2^{CD}, \\ \delta \left(1 - \frac{\beta c}{2}\right) & \text{if } \beta > \beta_2^{CD}, \end{cases}$$

which is continuous and decreasing in β .

Proof of Lemma 4. The internal solution of the welfare maximization problem is

$$d_w^{CD} = \frac{1}{\beta c} - \frac{1}{2} - \frac{\alpha\delta + 1 + \frac{\beta c}{2} - c}{2\alpha\gamma + \beta c - 2c}, \tag{10}$$

which is continuous in β . In the limit for $\beta \rightarrow 0$, the r.h.s. of (6) goes to infinity, hence d_w is bounded to 1. By continuity, there exists some (positive) value β_1^{CD} such that $d_w = 1$ for $\beta \leq \beta_1^{CD}$.

The derivative of (10) w.r.t. β is $\frac{dd_w^{CD}}{d\beta} = -\frac{1}{\beta^2 c} + \frac{c(\alpha\delta - \alpha\gamma + 1)}{(c\beta + 2\alpha\gamma - 2c)^2} < 0$. The fact that d_w^{CD} is monotonically decreasing when $d_w^{CD} \in (0, 1)$, jointly with the fact that the value in (10) can be negative (e.g. when δ is high enough) in turn imply that there exists a value of β_2^{CD} sufficiently high such that $d_w^{CD} = 0$ for any $\beta > \beta_2^{CD}$.

Proof of Proposition 7. From Lemmas 3 and 4, we immediately obtain that $d_p^{CD} = d_w^{CD} = 1$ in the limit for $\beta = 0$, and that $d_p^{CD} = d_w = 0$ when $\beta > \max\{\beta_2^{CD}, \beta_2^{CD}\}$.

Consider now intermediate levels of β such that $d_w^{CD} = \frac{1}{\beta c} - \frac{1}{2} - \frac{\alpha\delta + 1 + \frac{\beta c}{2} - c}{2\alpha\gamma - (2-\beta)c} \in (0, 1)$ and $d_p^{CD} = \frac{1}{\beta c} - \frac{1}{2} - \frac{\delta}{2\gamma\theta} \in (0, 1)$. We first verify that the set of β such that $d_w^{CD}, d_p^{CD} \in (0, 1)$ is non-empty. This is true, for example, for $\alpha = 0.3$, $\delta = 40$, $c = 0.9$, $\gamma = 10$ and $\theta = 0.8$, implying $d_w^{CD} = 0.47$ and $d_p^{CD} = 0.70$ when $\beta = 0.3$. From Lemmas 3 and 4, for intermediate levels of β , d_w^{CD} is constant to changes in θ , while d_p^{CD} is monotonically increasing in θ . Using the same parameter set as before, for $\theta = 0.8$ we have that $d_w^{CD} < d_p^{CD}$, while for $\theta = 0.7$ we have that $d_w^{CD} > d_p^{CD}$. This implies that there exists a value $\bar{\theta}^{CD}$ such that $d_p^{CD} > d_w^{CD}$ if $\theta > \bar{\theta}^{CD}$, and $d_p^{CD} < d_w^{CD}$ otherwise.

From the condition $d_p^{CD} = d_w^{CD}$ we obtain $\bar{\theta}^{CD} = \frac{\delta(2\alpha\gamma + \beta c - 2c)}{2\gamma(\alpha\delta + 1 + \frac{\beta c}{2} - c)}$. Moreover, by imposing $d_p = d_w$ we obtain $\bar{\theta} = \frac{\delta(2\alpha\gamma + \beta c - 2c)}{2\gamma(\alpha\delta + 1)}$, which is higher than $\bar{\theta}^{CD}$ for all $c > 0$.

Proof of Proposition 8. Sophisticated users choose disclosure d_s to maximize $\hat{U}_s(d_s)$, with $\hat{x}_s = x^{*CD}$. Since this utility is convex in d_s , they either choose zero or full disclosure, with

$$\hat{U}_s(0) = 1 - C - c + \frac{\beta c^2}{2} \left(1 - \frac{\beta}{2}\right); \hat{U}_s(1) = 1 - 2c + (2 - \beta)\beta c^2.$$

Zero disclosure ($d_s = 0$) is optimal for sophisticated users if $\hat{U}_s(0) \geq \hat{U}_s(1)$, or $C \leq C_s^{CD} = c - \frac{3}{4}(2 - \beta)\beta c^2$, and $d_s = 1$ is optimal for $C \geq C_s^{CD}$.

Naive users maximize $\hat{U}_n(d_n)$, with $\hat{x}_n = 1 - \frac{c}{2}(1 + d)$. Since also this utility is convex in d_n , we obtain $d_n = 0$ if $C \leq C_n^{CD} = c - \frac{3}{4}c^2$ and $d_n = 1$ for $C \geq C_n^{CD}$.

References

Acemoglu, D., Makhdoumi, A., Malekian, A., Ozdaglar, A., 2019. Too much data: Prices and inefficiencies in data markets. NBER Working Paper 26296. NBER.
 Acquisti, A., 2004. Privacy in electronic commerce and the economics of immediate gratification. In: Proceedings of the 5th ACM Conference on Electronic Commerce, EC '04. ACM, New York, NY, USA, pp. 21–29.
 Acquisti, A., Brandimarte, L., Loewenstein, G., 2015. Privacy and human behavior in the age of information. *Science* 347 (6221), 509–514.
 Acquisti, A., Grossklags, J., 2005. Privacy and rationality in individual decision making. *IEEE Secur. Priv.* 3 (1), 26–33.
 Acquisti, A., Varian, H.R., 2005. Conditioning prices on purchase history. *Mark. Sci.* 24 (3), 367–381.

Ali, S.N., Lewis, G., Vasserman, S., 2022. Voluntary disclosure and personalized pricing. *Rev. Econ. Stud.* 90 (2), 538–571.
 Baye, M.R., Sappington, D.E.M., 2020. Revealing transactions data to third parties: implications of privacy regimes for welfare in online markets. *J. Econ. Manag. Strategy* 29 (2), 260–275.
 Bloch, F., Demange, G., 2018. Taxation and privacy protection on internet platforms. *J. Public Econ. Theory* 20 (1), 52–66.
 Bourreau, M., Caillaud, B., De Nijs, R., 2018. Taxation of a digital monopoly platform. *J. Public Econ. Theory* 20 (1), 40–51.
 Braulin, F.C., Valletti, T., 2016. Selling customer information to competing firms. *Econ. Lett.* 149, 10–14.
 Casadesus-Masanell, R., Hervas-Drane, A., 2015. Competing with privacy. *Manag. Sci.* 61 (1), 229–246.
 Choi, J.P., Jeon, D.-S., Kim, B.-C., 2019. Privacy and personal data collection with information externalities. *J. Public Econ.* 173, 113–124.
 Conti, C., Reverberi, P., 2021. Price discrimination and product quality under opt-in privacy regulation. *Inf. Econ. Policy* 55, 100912.
 DellaVigna, S., Malmendier, U., 2004. Contract design and self-control: theory and evidence. *Q. J. Econ.* 119 (2), 353–402.
 Ek, C., Samahita, M., 2023. Too much commitment? An online experiment with tempting youtube content. *J. Econ. Behav. Organ.* 208, 21–38.
 Fainmesser, I.P., Galeotti, A., Momot, R., 2022. Digital privacy. *Manag. Sci.* 69 (6), 3157–3173.
 Galperti, S., Levkun, A., Perego, J., 2023. The value of data records. *Rev. Econ. Stud.* 91 (2), 1007–1038.
 Galperti, S., Perego, J., 2023. Privacy and the value of data. *AEA Pap. Proc.* 113, 197–203.
 Johnson, J.P., 2013. Targeted advertising and advertising avoidance. *Rand J. Econ.* 44 (1), 128–144.
 Jullien, B., Lefouili, Y., Riordan, M., 2020. Privacy protection, security, and consumer retention. TSE Working Paper 18-947, Toulouse.
 Kummer, M., Schulte, P., 2019. When private information settles the bill: money and privacy in Google’s market for smartphone applications. *Manag. Sci.* 65 (8), 3470–3494.
 Laibson, D., 1997. Golden eggs and hyperbolic discounting. *Q. J. Econ.* 112 (2), 443–478.
 Lefouili, Y., Madio, L., Toh, Y.L., 2024. Privacy regulation and quality-enhancing innovation. *J. Ind. Econ.* 72 (2), 662–684.
 Miklós-Thal, J., Goldfarb, A., Haviv, A., Tucker, C., 2024. Frontiers: digital hermits. In: *Marketing Science, Articles in Advance*, pp. 1–12.
 Montes, R., Sand-Zantman, W., Valletti, T., 2019. The value of personal information in online markets with endogenous privacy. *Manag. Sci.* 65 (3), 1342–1362.
 O’Donoghue, T., Rabin, M., 1999a. Doing it now or later. *Am. Econ. Rev.* 89 (1), 103–124.
 O’Donoghue, T., Rabin, M., 1999b. Incentives for procrastinators. *Q. J. Econ.* 114 (3), 769–816.
 Spiekermann, S., Grossklags, J., Berendt, B., 2001. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In: *Proceedings of the 3rd ACM Conference on Electronic Commerce, EC '01*. Association for Computing Machinery, New York, NY, USA, pp. 38–47.
 Taylor, C.R., 2004. Consumer privacy and the market for customer information. *Rand J. Econ.* 35 (4), 631–650.
 Tsai, J.Y., Egelman, S., Cranor, L., Acquisti, A., 2011. The effect of online privacy information on purchasing behavior: an experimental study. *Inf. Syst. Res.* 22 (2), 254–268.
 Tucker, C., 2019. Privacy, algorithms, and artificial intelligence. In: *Agrawal, A., Gans, J., Goldfarb, A. (Eds.), The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, pp. 423–438.