

Carpe Data: Protecting online privacy with naive users

Original

Carpe Data: Protecting online privacy with naive users / Abrardi, L; Cambini, C. - In: INFORMATION ECONOMICS AND POLICY. - ISSN 0167-6245. - ELETTRONICO. - 60:(2022), p. 100988. [10.1016/j.infoecopol.2022.100988]

Availability:

This version is available at: 11583/2971521 since: 2023-03-09T11:52:57Z

Publisher:

ELSEVIER

Published

DOI:10.1016/j.infoecopol.2022.100988

Terms of use:

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

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2022. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.infoecopol.2022.100988>

(Article begins on next page)

Carpe Data: Protecting online privacy with naive users*

Laura Abrardi[†], Carlo Cambini[‡]

June 2022

Abstract

In this paper, we study the optimal design of incentives to induce a digital platform to limit the extraction of data from users, whose privacy loss is further aggravated by their naive use of the platform. We show that caps on the amount of data collected can induce the optimal data-saving effort by the platform. If the platform's effort is not observable, a menu of data caps should be provided and it entails a higher (lower) loss of privacy for less (more) naive users, relative to the first best. We also show that compensating users for their data can efficiently incentivize effort, but might increase the privacy loss of more naive users.

Keywords: Data extraction, incentives, users' naivety, privacy.

JEL codes: D82, D83, D86, L12, L51

*We thank Maria Rosa Battaggion, Marc Bourreau, Steffen Hoernig, Byung-Cheol Kim, David E. Sappington, Yossi Spiegel and seminar participants at SIEPI 2020 and NERI 2020, as well as the editor and two anonymous referees, for useful comments and suggestions. Declarations of interest: none.

[†]Corresponding author. Politecnico di Torino, Department of Management, Corso Duca degli Abruzzi 24, 10129 Torino. Email: laura.abrardi@polito.it.

[‡]Politecnico di Torino, Department of Management. Email: carlo.cambini@polito.it.

1 Introduction

The massive and unprecedented scale of data collected by big digital platforms is creating serious concerns to policymakers for the large loss in terms of privacy as well as for its effects on competition, consumers, and society (Cremer et al., 2019; ACCC, 2019; EC, 2020). The policymakers' burden is aggravated by the fact that users often naively approach the privacy problem, and are nudged by questionable interface designs towards privacy intrusive options.¹ How to induce platforms to willingly choose privacy by design and frontally attack the issue of people's behavioral biases? In this paper we study the role of ex ante interventions towards a monopolistic digital platform that collects data from consumers' internet usage, and the incentives to adopt measures that reduce the privacy loss of users.

So far, the problem of excessive data collection by digital platforms has mainly been addressed by the privacy legislation, through the definition of the limits for the processing of personal information.² However, such measures require a substantial effort of enforcement to verify compliance, due to the large information asymmetries and rapid pace of innovation characterizing digital ecosystems (De-Streel et al., 2020). Regulating people's data is challenging also because few consumers are fully informed of, or fully understand, the scale and scope of data collected. While consumers claim to care about how their data will be used, in practice they often show little concern about it in their daily behavior (the so-called *privacy paradox*, Acquisti et al., 2015). Because of such inconsistency, consumers often relinquish considerable control over their uploaded content to digital platforms, thereby strengthening the platforms' potential of exploitation

¹Cookie consent prompts, for example, often present users with two options: to either immediately consent to the processing and disclosure of personal data to third parties, or to personalize the privacy settings by fine-tuning the disclosure after reading legalise and extremely long and detailed lists of trackers.

²For example, the 2018 European GDPR law states clearly that a platform should minimize the amount of data collected and processed for the purposes specified.

of the collected data (Solove, 2013).

In the face of the struggles of privacy law, a new idea is gaining consensus about the need for a more direct, regulatory-like approach to the problem of data and its use by internet giants (Stiglitz, 2019).³ In a report for the European Commission, Cremer et al. (2019) make the case for the creation of a new Authority for data regulation, to be enforced at a supranational level by an ad-hoc body (DeStreel et al., 2020). In January 2020, the UK government announced that companies sharing user-generated content will soon be subject to an independent regulator, and in June of the same year the European Commission launched a consultation to seek views on Digital Services Act package, arguing that Europe needs a modernized regulatory framework for digital services.⁴ Other countries are going along the same path: for example, the Australian Antitrust Authority recently argued that the regulatory framework governing digital platforms needs to be addressed (ACCC, 2019). If the problem of regulating digital services is at the forefront of public authorities, it is still unclear how the problem should be approached, missing a theoretical framework that could guide the policymakers' action.

In our model, the usage of the service of a digital platform –for example, a search engine or a social network– generates personal data about users, that can be monetized by the platform. The data extracted causes a privacy loss to users, which they naively underestimate, thus overdisclosing data. The platform could reduce the users' privacy loss by exerting a non-observable effort in data minimization. For example, the effort could entail developing a protocol that reduces the collection of data during the provision of the service. The socially

³The Economist, May 6th 2017, “Regulating the internet giants. The world’s most valuable resource is no longer oil, but data”.

⁴In the Commissions' words, “The current regulatory framework for digital services dates back twenty years. It helped the growth of European digital services but it does not give answers to many of today’s pressing questions on the role and responsibility of online platforms, especially the largest ones” (EC, 2020).

optimal provision of data trades-off the cost of effort and the cost of the privacy loss. However, the platform does not have an incentive to exert effort in the data-saving technology. We study two potential policy interventions: a cap on the data extracted, that might be applied to restore the platform incentives to enhance its effort, and redistribution of the platform’s revenues through compensations paid to users for their data. The cap on data is the maximum amount of data that the platform can extract, for example the information on the user’s name, location, contacts, etc. The optimal policy entails a cap on the quantity of data that the platform can collect. For example, the social planner could mandate the anonymization of some data before it is stored at the service provider, or limit the access to data such as the contact list or the geolocation information via a smartphone.⁵

Owing to the fact that the platform’s effort to protect users’ privacy might be difficult to observe by a regulator, the platform could gain a rent by blaming the users’ naivety for its large extraction of data, thus concealing its low effort. We show that the optimal second best intervention requires to offer to the platform a menu of contracts, characterized by a more stringent data cap (relative to the first best) when users have severe naivety, and a larger cap when users have mild naivety.

We also show that the possibility to compensate users for their data, jointly with data caps, improves welfare but worsens the privacy loss of more naive users relative to the case in which only data caps are imposed. In fact, the value extracted by the platform from data can be redistributed to users via the transfers. Given that users are reimbursed for the privacy loss when transfers are available, the effort to develop a data-saving architecture is lower and users suffer a larger privacy loss relative to the case where no compensations are feasible.

Our goal is to link the literature on data acquisition and its impact on privacy

⁵Data aggregation and masking are currently part of the set of proposals in the Digital Market Act to foster competition (Cabral et al., 2021).

with the literature pertaining to the inconsistency of preferences underlined by behavioral sciences (Kahneman and Tversky, 1979; Ainslie, 1992; Laibson, 1997). Our contribution to the theoretical literature is thus twofold. First, we model the effort exerted by the platform to protect users' privacy. In our model, the platform plays an active role in internalizing the network externalities via the prices, and needs incentives to reduce the privacy loss of users. Second, we introduce the element of the users' naivety, which hampers the effectiveness of initiatives such as consent policies and enhances the potential of data exploitation. Our approach is close in spirit to the literature on regulation of monopolistic firms under information asymmetries (e.g., Baron and Myerson, 1982), with one crucial difference. In our setup, the regulator needs to take into account not only the firm's lack of incentive to exert effort, but also the inefficiency generated by consumer's behavioral responses. Our results highlight that regulating a monopolist with behavioral users introduces an additional dimension to the usual efficiency vs. rent extraction trade-off, in terms of ability to mitigate users' (mis)behavior.

The next section of the paper presents our relation to the pertinent literature. Section 3 describes our model. Section 4 studies the optimal design of incentives when only caps on data can be imposed, while Section 5 analyzes optimal regulation when users can also be compensated for their data. In Section 6 we extend the model to the case in which the platform's effort increases its revenues. Finally, Section 7 concludes. All proofs and technical details are relegated into an Appendix.

2 Related literature

This paper lies at the intersection between two main strands of literature.

The first one is the literature that explores the role of users' naivety on privacy. Although excessive data sharing by some users might be consistent with rational behaviour when data externalities are present (Choi et al., 2019; Ace-

moglu et al., 2022), behavioral biases constitute an important channel of privacy loss. To the best of our knowledge, only three papers are dealing with this issue, despite its importance from a policymaking point of view (Tucker, 2018). Taylor (2004) provides a theoretical model in which privacy regulations limit the ability of individual merchants to sell customer transactions data to other merchants. A naive consumer does not anticipate that his present actions may affect the prices he faces in subsequent transactions, and faces excessive privacy loss. Acquisti and Varian (2005) apply a similar argument to the case of a monopolist with repeated transactions, assuming that tastes are intertemporally correlated. Finally, Baye and Sappington (2020) assume that the merchants operate inside a platform. In the aforementioned literature, not fully rational users are not able to act strategically, and fail to account for the exploitation of their data by firms during future purchases. We instead assume that users suffer from a self-restraint problem, that causes an inconsistency between their stated preference for privacy, and the preference they reveal through their behavior. As a consequence, they release too much personal information when using the service. Our setup is also different as it considers platforms offering their services free-of-charge, therefore it can be used to analyze situations such as people posting messages, personal information and likes on social media.

The second line of research to which we are connected is the recent and growing strand of literature on privacy protection, data acquisition, and public policy. Advances in information technologies and the spread of e-commerce and digital rights management has spurred firms' ability to adopt differentiated, behavior-based strategies (Fudenberg and Villas-Boas, 2007). Our work is close in spirit to the idea that our digital behavior can leave tracks that could be exploited by the firm to increase its profit (Acquisti and Varian, 2005). Indeed, the recent literature on the value of data has emphasized how the access to customer-level data may affect profits (and welfare) through several channels, e.g. by affecting firms' investment in quality (Lefouili and Toh, 2017), or by providing to firms

private information, which could be used for a competitive advantage (Montes et al., 2019; Casadesus-Masanell and Hervas-Drane, 2015). Recent papers have also analyzed the interplay between acquiring information to perform a better personalized price while accounting for the consumers’ reactions from the loss of privacy (Braulín and Valletti, 2016; Montes et al., 2019). Jullien et al. (2020) show that, when data sharing causes disutility from advertising, customer retention may motivate the firm to reduce data monetization. Moreover, the presence of consumers with different attitudes towards the loss of privacy influences the scope for price discrimination (Krähmer and Strausz, 2022). Several studies have focused on policy instruments that may affect the firms’ incentives for data collection. Ichihashi (2020) shows that, if firms can commit not to use consumer information for pricing, consumers have the incentive to share more data, thus increasing firms’ profits despite the inability to price discriminate. In addition, shifting the ownership of data to consumers might reduce the firm’s incentive to invest in data processing (Dosis and Sand-Zantman, 2022). We contribute to this literature by taking a regulatory perspective to study the efficiency and optimal design of data caps and revenue sharing mechanisms, in a setting where users are not fully rational.

3 A simple model

In our model, a platform provides a service to users, but the service entails a risk in terms of privacy loss. The platform can exert an effort to provide the service in a way that limits users’ privacy loss.

The platform. A monopolistic digital platform provides a service to a unit mass of users. Users receive from the service the utility $u(x)$, with $u(0) = 0$, $u_x > 0$, $u_{xx} < 0$, where $x \in [0, x^{max}]$ is the intensity of use of the service (e.g., the amount of time users spend on the platform). The use of the service provides to the platform an amount of data d about the users’ characteristics, such as their

identity, movements, interests or location. The platform can exert an effort a that reduces the extraction of data by developing a data-saving technology. For example, the platform's effort allows to design a protocol that anonymizes the data collected. Then, the total amount of data produced $d = d(x, a)$ depends positively on the intensity of use x , and negatively on the platform's effort a . For simplicity, let us assume that the data production function is linear: $d = x - a$.⁶

Developing a data-saving technology is costly for the platform. In particular, the platform incurs in the cost $\psi(a)$, with $\psi(0) = 0$, $\psi', \psi'' > 0$, for exerting effort a . The amount x of use generates a positive externality γx , with $\gamma > 0$ representing the commercial value of the service, which accommodates the presence of a second side of the market, for example advertisers, and it is completely internalized by the platform owing to its market power. Our focus is on the case where γ is sufficiently high, as in the case of big digital companies that earn large profits from the rich variety of data collected.

In order to build a flexible model that could be applied to a variety of platforms, we assume that users pay to the platform a fixed transfer F for accessing the service. The fixed transfer F can be positive (as in the case of subscription fees for accessing the platform's service, or the price for downloading the app), null (as in the case of free platforms) or negative, if users receive a compensation from the platform. In Section 4 we focus our attention on platforms that provides a service for free (i.e., $F = 0$), such as in the case of Google or Facebook. However, F might be the result of a regulatory choice. For example, users might receive a compensation in return of the data they provide while using the platform. In Section 5 we thus generalize the analysis by leaving F unconstrained and show how regulation should determine the optimal sign (and value) of F . The platform's profit is thus equal to $\pi(x, a, F) = F + \gamma x - \psi(a)$.

⁶Results can be extended to the case of a more generic function $d = d(x, a)$, as shown in Appendix A.2. Moreover, in Section 6 we consider the case where the data function is not separable in effort.

The users. Users incur in a constant marginal cost c for each unit of data extracted by the platform due to the loss of privacy. The total privacy cost cd can be due to the future exploitation of their personal data for price discrimination or to the negative feeling associated to the disclosure of personal information.

Users decide the amount x of usage of the service trading-off its benefit in terms of utility, and its cost in terms of loss of privacy. However, when taking this decision, users naively behave as if their privacy cost is only $\tilde{c} < c$, where the difference $c - \tilde{c}$ is a measure of users' naivety.⁷

There are two types $j = \{L, H\}$ of users, that differ for the severity of their naivety. Users of type j behave as if their privacy cost is \tilde{c}_j , with $\tilde{c}_L < \tilde{c}_H$. Therefore, users of type j maximize their utility, net of the perceived privacy cost: $\max_x u(x) - \tilde{c}_j(x - a) - F$. Their preferred level of usage solves the FOC

$$u_x(x) = \tilde{c}_j, \quad (1)$$

i.e., $\tilde{x}_j = u_x^{-1}(\tilde{c}_j)$ and $\tilde{d}_j(a) = u_x^{-1}(\tilde{c}_j) - a$, for $j \in \{L, H\}$.⁸ More effort in designing a data-saving protocol reduces the amount of personal data extracted by the platform for a given amount of usage. Moreover, we obtain $\tilde{x}_L > \tilde{x}_H$ and $\tilde{d}_L > \tilde{d}_H$: the trade-off between usage and loss of privacy is resolved by more naive users (L) in favor of usage, as they choose a more intensive use of the service than less naive (H) users, thereby suffering a larger privacy loss.

From a social planner's standpoint, the users' surplus is defined as the users' utility, net of the true privacy cost: $S(x, a, F) = u(x) - c(x - a) - F$. For a consistent treatment of the privacy problem, we focus on separating equilibria where both users choose positive amounts of usage. To this aim, we assume that

⁷The discrepancy between the users' attitude towards the privacy loss and their behavior reflects the phenomenon known as "privacy paradox" (Acquisti et al., 2015, 2016).

⁸Due to the monotonicity of the utility function $u(x)$, the user's problem is not bounded when $\tilde{c}_j = 0$. The assumption that x is upper bounded ($x \leq x^{max}$, e.g., the maximum amount of time users can spend on the platform) allows to ensure the existence of a solution to condition (1).

the privacy cost is sufficiently low, i.e., $c \leq u_x(0)$, implying that usage increases surplus, despite the loss of privacy.⁹

The regulator maximizes the weighted sum of the users' expected surplus plus the expected profit for the platform, where $\alpha \in [0, 1]$ is the weight of the profit component (Baron and Myerson, 1982).¹⁰ Specifically, the regulator's objective is to maximize the welfare function $W = S(\tilde{x}, a, F) + \alpha\pi(\tilde{x}, a, F)$.

The timing is as follows. First, the social planner chooses the effort a and the fixed fee F . Second, users choose the amount of usage x and pay the privacy cost. The platform obtains the profit π .

4 Platform regulation without transfers

Social networks and search engines, as well as many digital marketplaces, are currently free of any usage charge as well as any forms of user compensations. What are the implications in terms of platform's incentives for protecting users' privacy? How are these incentives affected by users' naivety? The absence of transfers has several well-known economic implications. First of all, from a redistributive point of view, the social planner cannot use transfers to convey some of the surplus from the platform to users. In addition, there are efficiency implications, as transfers cannot be exploited to efficiently provide to the platform the incentive to exert effort in a second best scenario. We now study how incentives can be provided in this framework, and the cost that they entail.

In this Section we assume that $F = 0$. We first characterize the perfect information benchmark, and then assume that the effort is not observable by the regulator or by users.

⁹If c is too high, the regulator would prefer no usage and could simply forbid the use of the platform.

¹⁰The parameter α allows to shift the regulator's focus from a privacy problem to the problem of internalizing the external value of data. When $\alpha = 0$, the regulator's objective coincides with that of a privacy authority, focused on minimizing the privacy cost.

4.1 Perfect information benchmark

The platform's effort that maximizes welfare, subject to the platform's participation constraint, solves the problem

$$\max_{a_j} W = u(\tilde{x}_j) - c(\tilde{x}_j - a_j) + \alpha [\gamma \tilde{x}_j - \psi(a_j)] \quad (2)$$

$$s.t. \pi(\tilde{x}_j, a_j) \geq 0. \quad (3)$$

If γ is high enough, the constraint (3) is not binding at the optimum level of effort a_j^* that maximizes welfare (2), i.e. we assume $\gamma \tilde{x}_j - \psi(a_j^*) > 0$.¹¹ Moreover, given the upper bound on the privacy cost, $c \leq u_x(0)$, the user surplus is positive for all $a_j \geq 0$, ensuring his participation. Using the supply of data (1), the welfare function can be rewritten as:

$$W = u(u_x^{-1}(\tilde{c}_j)) - c(u_x^{-1}(\tilde{c}_j) - a_j) + \alpha [\gamma u_x^{-1}(\tilde{c}_j) - \psi(a_j)]. \quad (4)$$

The FOC of (4) yields

$$a_j^* : \alpha \psi'(a_j) = c. \quad (5)$$

In the optimum, the marginal cost of effort $\psi'(a)$, weighted by α , is equal to the marginal benefit in terms of reduced privacy loss. The weight α is due to the fact that the social planner might give a lower weight to profits than to surplus in the welfare function. When the social planner focuses entirely on consumers' surplus (i.e., $\alpha = 0$), the optimal effort is the maximum that is feasible, given the platform's participation.¹² Conversely, when $\alpha > 0$, the social planner cares also about the commercial value of data and the cost of effort. Therefore, the optimal effort trades-off the cost of effort for implementing a data-saving technology, and the benefit in terms of lower privacy loss. In this case, the optimal effort is higher,

¹¹In the more general case in which γ is low, Problem (2) has a boundary solution with a_j such that $\pi(\tilde{x}_j, a_j) = 0$, with no loss of generality. Notably, such a situation is less likely to call for a regulatory intervention.

¹²This is, for example, the perspective adopted by the privacy law, whose focus is on the privacy cost. Indeed, data minimization is at the foundation of the European GDPR.

the higher is the cost of privacy c , or the lower is the weight α of profits in the welfare function. Moreover, by using (5), it is easy to find that $d_L^* > d_H^*$: very naive users are a richer source of data, as they underestimate to a higher extent the privacy cost and use more intensely the service (i.e., $\tilde{x}_L > \tilde{x}_H$). This fact has immediate implications on the platform's profit: we find that $\pi_L^* > \pi_H^*$, that is, profits are higher with more naive users.

We now show that, in the absence of regulation, a platform underexerts effort relative to the first best level a_j^* . In fact, for each type of user j , the platform chooses a_j^U (where the superscript U denotes the unregulated solution) so as to maximize the profit $\pi(x_j, a_j)$ subject to the user's supply of data (1):

$$\max_{a_j} \gamma \tilde{x}_j - \psi(a_j). \quad (6)$$

The solution of Problem (6) is straightforward and is expressed by the following Proposition:

Proposition 1 *In the absence of any regulatory intervention, the platform underexerts effort relative to the first best (i.e., $a_j^U = 0 < a_j^*$).*

The profit maximizing platform chooses not to devote any effort to protect users' data. In fact, the platform's effort is costly, but it does not increase its revenue. A regulatory intervention is thus needed to induce the platform to an efficient level of effort, in the form of limits to the quantity of data $d_L^* > d_H^*$, depending on the users' propensity to use the service. These data caps might for example entail limits to the information collected by the platform, such as the information on the user's name, gender, political interests, location, etc. They might also be enforced by imposing aggregation and masking obligations on the information stored.¹³ We will discuss its practical implementation in the final

¹³Such obligations are currently at the center of the policy discussion for data sharing purposes, to reduce gatekeepers' exclusive control over the data they collect (see, e.g., Cabral et al., 2021).

Section of the paper. In the analysis that follows, we study the situation in which the data-saving effort is the platform’s private information, and neither the regulator nor users can observe it.

4.2 Non-observable data saving effort

Users’ naivety on privacy decisions often takes the form of an inconsistency between users’ stated preferences about privacy, and the preferences they reveal through their actual behavior. In particular, users may state their concern about privacy and believe that the cost of a privacy loss is c , but then they naively behave as if the cost of a privacy loss is only $\tilde{c}_j < c$. Since in many cases users’ naivety emerges only during their actual interaction with the platform, the latter is endowed with first-hand information about the cost \tilde{c}_j that drives the users’ actual consumption of the service. The platform’s informative advantage about users’ actual behavior and its own effort has important implications. In fact, it might be difficult for the social planner to figure out whether a large volume of collected data $d = \tilde{x} - a$ is due to the user’s naivety, which causes an intense use of the service (i.e., a high \tilde{x}), or rather to the platform’s misbehavior (i.e., a low a). Conversely the platform, by observing the level \tilde{x}_j of consumption, can infer the degree of user’s naivety (i.e., the cost \tilde{c}_j revealed by their behavior), and exploit this private information to its advantage.¹⁴ In this Section we explain how the problem of optimal –second best– regulation could be approached.

Let us assume that \tilde{x}_j (or, equivalently, \tilde{c}_j) is the platform’s private information, although its probability distribution is common knowledge. The platform caters to a homogeneous mass of j users. Let us denote with λ the probability that users have low naivety (i.e., $j = H$), and with $1 - \lambda$ the probability that users

¹⁴Advances in information technologies and the spread of e-commerce and digital rights management has spurred firms’ ability to adopt differentiated, behavior-based strategies (Fudenberg and Villas-Boas, 2007), possibly relying on private information to gain a competitive advantage (Montes et al., 2019).

have $j = L$ type.¹⁵ We also assume that the platform's effort is not observable by the regulator or by users. While the platform's effort a and the users' preferred level of service \tilde{x}_j are unobservable by the regulator, the data extracted by the platform can be observed and verified. For example, under the current provisions of the European GDPR, companies have the legal obligation to specify which personal data they are collecting and for which purpose, how they mean to use it, how it may be disclosed and to whom.¹⁶

In this setup, the platform has the incentive to extract from less naive users (i.e., $j = H$) the same amount of data d_L that it would extract in the case of severely naive users, in order to save on effort. In particular, for a generic level of effort a_H , the data extracted from H users is $\tilde{d}_H(a_H) = \tilde{x}_H - a_H$. The effort \hat{a}_H , such that $\tilde{d}_H(\hat{a}_H)$ coincides with $\tilde{d}_L(a_L) = \tilde{x}_L - a_L$, is

$$\hat{a}_H(a_L) = a_L - (\tilde{x}_L - \tilde{x}_H),$$

which can also be expressed as

$$\hat{a}_H(a_L) = a_L - \Delta, \tag{7}$$

where $\Delta = \tilde{x}_L - \tilde{x}_H$. Given that $\tilde{x}_L > \tilde{x}_H$, then $\Delta > 0$ and $\hat{a}_H(a_L) < a_L$. In order to streamline the exposition, let us focus on internal solutions by assuming that $a_L^* - \Delta > 0$.

¹⁵Assuming an homogeneous population of users is consistent with the standard approach of the literature on regulation with asymmetric information (Baron and Myerson, 1982; Besanko and Sappington, 1987). In alternative, λ could be interpreted as the share of users in the population with type H , assuming heterogeneous users. We develop this case in Appendix A.1.

¹⁶Transparency requirements regarding data collection and processing are one of the pillars of the European GDPR and several provisions are meant to ensure it, including the right to be informed and the right to access to the information gathered, the need of users' unambiguous and specific consent, and the need to keep written documentation of data categories and the purpose of the processing. This documentation must be completely made available to authorities upon request.

Suppose that the regulator offers the first best menu $(d_L^*; d_H^*)$. Suppose also that the platform, after having privately observed that its users have low naivety (H), chooses a level of effort such that the level of usage \tilde{x}_H chosen by H users, produces the amount of data d_L^* . As $\hat{a}_H(a_L^*) < a_L^*$, a platform with H users can save on effort in the case it extracts the same amount of data d_L^* obtained when users are L . Therefore, under asymmetric information, the platform earns a rent when users are H by choosing a lower effort, thereby overextracting data from H users. The complete information outcome with the menu of contracts $(d_L^*; d_H^*)$ can no longer be implemented under asymmetric information.

The size of the informational rent obtained by the platform depends on the effort a_L . In particular, the profit obtained by a platform with H users, given the effort $\hat{a}_H(a_L) = a_L - \Delta$, is $\hat{\pi}_H = \gamma\tilde{x}_H - \psi(a_L - \Delta)$: a higher effort a_L reduces the deviation profit obtained by the platform when its users are H , because it increases the cost of effort. The optimal level of effort of the platform when users are L thus entails a trade-off. On the one hand, the increase of the effort a_L^{SB} reduces the rent that the platform earns when users are H . On the other hand, a higher a_L^{SB} entails an inefficiently high costs for the platform when users are L . Hence, the optimal second best menu entails the well-known trade-off between rent extraction and efficiency (Baron and Myerson, 1982).

When the platform has an informative advantage on its data-saving effort, the regulator must offer a menu of incentive compatible contracts, which we denote with $(d_L^{SB}; d_H^{SB})$, where the superscript SB denotes the second best solution. Each contract includes the specification of the quantity of data that could be extracted d_j^{SB} from each type j of user. In the optimal solution, the platform will properly self-select within this menu and choose the data cap d_H^{SB} if $j = H$, or d_L^{SB} if $j = L$. The correct self-selection of the platform is ensured by the incentive compatibility of the contracts. We now study the features of the menu of contracts in this second best scenario.

The regulator solves the following problem:

$$\max_{a_H, a_L} \lambda [u(\tilde{x}_H) - c\tilde{d}_H + \alpha(\gamma\tilde{x}_H - \psi(a_H))] + \quad (8)$$

$$+ (1 - \lambda)[u(\tilde{x}_L) - c\tilde{d}_L + \alpha(\gamma\tilde{x}_L - \psi(a_L))] \quad (9)$$

$$s.t. \gamma\tilde{x}_j - \psi(a_j) \geq 0 \quad \forall j, \quad (9)$$

$$\gamma\tilde{x}_H - \psi(a_H) \geq \gamma\tilde{x}_L - \psi(\hat{a}_H(a_L)), \quad (10)$$

$$\gamma\tilde{x}_L - \psi(a_L) \geq \gamma\tilde{x}_H - \psi(\hat{a}_L(a_H)). \quad (11)$$

Constraint (9) constitutes the participation constraint of the platform j , while constraints (10) and (11) ensure the incentive compatibility of the contract.

The internal solution of the regulator's problem is implicitly defined by the following conditions:¹⁷

$$\alpha\psi'(a_L^{SB}) = c + \frac{\lambda}{1 - \lambda}\psi'(a_L^{SB} - \Delta) \left(\frac{c}{\psi'(a_H^{SB})} - \alpha \right) \quad (12)$$

$$\psi(a_H^{SB}) = \psi(a_L^{SB} - \Delta) - \gamma\Delta \quad (13)$$

Proposition 2 compares conditions (12) and (13) with the first best solution described in Subsection 4.1.

Proposition 2 *Optimal regulation under imperfect information is characterized by $d_L^{SB} < d_L^*$ and $d_H^{SB} > d_H^*$. Moreover, the platform's profits are higher (lower) than the first best when users have low (high) naivety (i.e., $\pi_H^{SB} > \pi_H^*$ and $\pi_L^{SB} < \pi_L^*$).*

Proof. See Appendix A. ■

¹⁷As in the benchmark, we assume that the platform's participation constraints are not binding in the second best optimum, i.e. γ is sufficiently high. In the case this condition is not met, the incentive compatibility could be ensured by reducing the distortion and/or by allowing transfers from users to the platform, as we examine in Section 5.1. Fixed payments in the form of subscriptions are actually implemented by some types of platforms (e.g., Netflix), although not by others (e.g., Google or Facebook, where data generates a very high value).

Proposition 2 highlights that in a second best scenario, the amount of data that can be extracted by the platform must be distorted, by imposing limits that we could interpret as data caps. In particular, a pair of incentive compatible contracts entails a distortion for all type of users, although the direction of the distortion depends on the level of users' naivety. The data cap is reduced (relative to the first best) when users are highly naive (L), whereas it is raised above the first best level in the case of less naive (H) users. As a consequence, the platform is induced to overexert effort when users have high naivety, and underexert effort when users are less naive. This result follows from the fact that both efforts can be used as an instrument to ensure the incentive compatibility of the platform. The decrease of effort a_H raises the platform's profit in the case of H users (as it decreases its cost), while the increase of the effort a_L reduces the platform's profit in the case of L users. Overall, both effects induce the platform to a truthful revelation of information when users are H . In fact, the overexertion of effort required to the platform in case of users with high naivety reduces the platform's profit when users are L (i.e., $\pi_L^{SB} < \pi_L^*$), and decreases the incentive of a platform with less naive users to deviate.

Finally, note that the distortion of a_L^{SB} is higher, the higher is the probability λ that users have a mild naivety. Intuitively, when users with a mild naivety are more likely, the distortion of the effort of highly naive users has a negligible impact on welfare.

Figure 1 graphically illustrates the results of Proposition 2, and highlights how the second best menu of data caps reduces the gap $d_L - d_H$ relative to the first best menu.

The necessity to introduce a distortion on both efforts entails a significant welfare loss, which could be reduced if the social planner had an additional instrument in his hands. In the next section, we analyze the case where transfers are allowed between the platform and users.




Figure 1: Second best menu of data caps

5 Compensating users for their data

One set of proposals for curbing platforms’ potential of exploitation of user data is to allow individuals to accrue the value of the data itself.¹⁸ The data provided by users constitutes a fundamental input for platforms, yet in most cases it yields little financial value to users. This fact has raised some discussion about treating the data like property and compensating people for it (Stiglitz, 2019; Gans, 2020).¹⁹ ²⁰ According to this view, users could receive compensations or rewards for providing the data while using the service.

In this Section we analyze the extent to which revenues should be shared between the platform and users, from a socially optimal point of view. This form of regulation is substantially analogous to a taxation system, as in Bloch and Demange (2018) and Bourreau et al. (2018). However, while these works explore the performance of different taxation systems, our focus is rather on the effects of revenue sharing on the platform’s incentives for efficiency, in the spirit of Laffont

¹⁸“We need to own our data as a human right - and be compensated for it”, The Economist, January 21st 2019.

¹⁹“If people were paid for their data”, The Economist, July 7th 2018.

²⁰Stiglitz (2019) states that “Government could go further, assigning a minimal price as compensation for a firm using personal data” [p. 135].

(1994), Armstrong and Sappington (2006) and Iossa and Stroffolini (2005).

Allowing for compensations, in addition to limits on data, provides to the regulator an additional instrument through which he can channel incentives to the platform. As a consequence, this section analyzes the optimal design of such rewards to users, and how they could influence the incentives for the platforms to invest in data-saving technologies.

The level of the platform's effort that maximizes welfare, subject to its participation constraint, solves the problem

$$\begin{aligned} \max_{a_j, F_j} W &= u(\tilde{x}_j) - c(\tilde{x}_j - a_j) - F_j + \alpha [F_j + \gamma \tilde{x}_j - \psi(a_j)] \\ & \text{s.t. } \pi(\tilde{x}_j, a_j, F_j) \geq 0. \end{aligned} \quad (14)$$

The solution of Problem (14) is provided by the following Proposition:

Proposition 3 *Optimal regulation entails a_j^{**} such that $c = \psi'(a_j^{**})$ and $F_j^{**} < 0$ if γ is sufficiently high. Moreover, $d_L^{**} > d_H^{**}$ and $F_L^{**} < F_H^{**}$.*

Proof. See Appendix A. ■

If data externalities are sufficiently large (i.e., γ is large enough), the platform's revenue from the service, net of the effort cost, is positive: $\gamma \tilde{x}_j - \psi(a_j^{**}) > 0$. As profits have a lower weight than surplus in the welfare function, they are extracted via the fixed transfer, hence $F_j < 0$.²¹ Then, users receive a compensation that is larger, the larger the profit. In particular, the transfer $F_j^{**} = -\gamma \tilde{x}_j + \psi(a_j^{**})$ fully reimburses the platform for the cost of effort, and extracts the value it obtains from data. Since the optimal effort depends on the true privacy cost, rather than the perceived one, it is the same regardless of the users' naivety,

²¹Although our focus is on quasi-monopolistic digital companies, whose high γ generates severe privacy issues and might require a regulatory intervention, our results could be generalized to the case where γ is low. In this case, we would have $\gamma \tilde{x}_j - \psi(a_j^{**}) \leq 0$. Then, the platform must be compensated for the loss through a fixed transfer $F_j \geq 0$. This case, however, is less likely to call for a regulatory intervention.

i.e. $a_H^{**} = a_L^{**}$. Then, more naive users exhibit a higher usage than less naive ones ($\tilde{x}_H < \tilde{x}_L$), and should thus receive a higher transfer than users with a low naivety (i.e., $F_L^{**} < F_H^{**}$).

When the commercial value of the service is large, the transfer is a means to return to users the value of their data. Users with a severe naivety, whose use entails the extraction of large amount of data, should receive a larger transfer than users with a milder naivety. The possibility to redistribute profits to users via transfers also affects the optimal level of effort. From Proposition 3, in the optimum the marginal cost of effort $\psi'(a)$ is equal to its marginal benefit c in terms of reduced privacy loss, thus reflecting the trade-off between the cost of effort and the loss of privacy. Note that the optimal effort a_j^{**} does no longer depend on the weight α of the profit function on welfare, as was the case for effort a_j^* when no transfers are available. This is due to the fact that transfers allow to shift the profits to users, hence the regulator fully accounts for the cost of effort, even if profits have a low (or even null) weight.

It is also interesting to compare the optimal level of effort with the results described by condition (5) in the previous Section. We can observe that $a_j^* > a_j^{**}$: the platform is induced to exert more effort in a data-saving technology when transfers are not allowed. This is due to the fact that compensating people for data crowds out the need to develop the data-saving technology.

A graphical representation of the optimal levels of data and effort when transfers are feasible is provided by Figure 2, which summarizes the above discussion and compares the first best effort a_j^{**} with the effort a_j^* obtained in the absence of transfers.

The prohibition to compensate people for their data stimulates the platform to exert effort, and leads to the reduction of the privacy loss, relative to the case when transfers are feasible: $d_j^* < d_j^{**}$ for all j . Naturally, allowing transfers improves welfare because it induces an optimal level of effort. However, it might be worthwhile to note that welfare is higher when transfers are allowed only




Figure 2: Optimal regulation when compensations are feasible. First and second best levels of data (panel a) and effort (panel b)

because the user is compensated for his data, even though his privacy loss is larger than in the case transfers are not allowed. A fixed transfer allows to extract the platform's profit and compensate the user for his provision of data, whereas in the absence of transfers the users' data provision can be exploited and raises the platform's rents.

For a useful comparison, let us now describe the optimal choices of an unregulated platform, given the behavior of a naive user. The platform can exploit the user's naivety to over-extract value through the fixed fee. Then, for each type of user j , the platform chooses a_j so as to maximize the profit $\pi(x_j, a_j, F_j)$ subject to the user's supply of data (1) and their participation constraint:

$$\max_{a_j, F_j} F_j + \gamma \tilde{x}_j - \psi(a_j) \quad (15)$$

$$s.t. u(\tilde{x}_j) - \tilde{c}_j (\tilde{x}_j - a) - F_j \geq 0. \quad (16)$$

The solution of Problem (15) allows to obtain the profit maximizing level of effort

$$a_j^{U'} : \psi'(a) = \tilde{c}_j < c, \quad (17)$$

where the superscript U' denotes the unregulated solution when transfers are

feasible. By comparing (17) with the result of Proposition 3, we find that the platform under-exerts effort relative to the first best, as also illustrated in Figure 2. Developing a data-saving technology is costly, whereas its benefit in terms of lower privacy loss (which can be extracted by the platform through F_j), is reduced because of the user's naivety. On the contrary, the social planner accounts for the true benefit in terms of reduction of privacy loss.

The results in this Section are summarized by the following Proposition.

Proposition 4 *In the unregulated scenario, the platform underexerts effort and overextracts data relative to the social optimum (i.e., $a_j^{U'} < a_j^{**}$ and $d_j^{U'} > d_j^{**}$ for all j).*

In the absence of regulation, the platform curtails the costly (but data-saving) effort. Given the users' desired level of utilization of the service, the platform extracts a larger amount of data, and users incur in an excessive privacy loss. The insufficient level of effort thus reduces welfare. Quite importantly, the user's naivety plays a crucial role on this welfare loss. Indeed, if users were rational (i.e., $\tilde{c}_j = c$), the platform's preferred level of effort would coincide with the first best one ($a_j^{U'} = a_j^{**}$).

How does the presence of user compensations affect the optimal design of platform incentives in a second best scenario? We answer this question in the next section.

5.1 User compensations and incentives

The analysis in the previous section has highlighted how the possibility to compensate users for their data, jointly with constraints on the amount of data collected by platforms, allows to efficiently provide incentives to induce the platform's effort. In this Section we study how these incentives are affected if the platform is endowed with private information about its effort.

The social planner solves the following problem, which is analogous to Problem (8), but for the presence of transfers:

$$\max_{a_H, a_L, F_H, F_L} \lambda[u(\tilde{x}_H) - c\tilde{d}_H - F_H + \alpha(F_H + \gamma\tilde{x}_H - \psi(a_H))] + \quad (18)$$

$$+(1 - \lambda)[u(\tilde{x}_L) - c\tilde{d}_L - F_L + \alpha(F_L + \gamma\tilde{x}_L - \psi(a_L))] \quad (19)$$

$$s.t. F_j + \gamma\tilde{x}_j - \psi(a_j) \geq 0 \quad \forall j \quad (20)$$

$$F_H + \gamma\tilde{x}_H - \psi(a_H) \geq F_L + \gamma\tilde{x}_L - \psi(\hat{a}_H(a_L)) \quad (21)$$

$$F_L + \gamma\tilde{x}_L - \psi(a_L) \geq F_H + \gamma\tilde{x}_H - \psi(\hat{a}_L(a_H)) \quad (22)$$

The solution of the regulator's problem is defined in Proposition 5, where the superscript SB' denotes the second best outcome in the case transfers are feasible.

Proposition 5 *The optimal revenue sharing rule under imperfect information is characterized by $d_H^{SB'} = d_H^{**}$ and $d_L^{SB'} > d_L^{**}$. Moreover, $F_H^{SB'} > F_H^{**}$ and $F_L^{SB'} < F_L^{**}$. Finally the platform's profit is $\pi_L^{SB'} = \pi_L^{**} = 0$ and $\pi_H^{SB'} > \pi_H^{**} = 0$.*

Proof. See Appendix A. ■

Intuitively, the shift of some profits from the platform to users when they have a high naivety reduces the incentive for the platform to curtail its effort when users have a low naivety. Hence, the social planner leaves the platform with L users with minimum profits, such that its revenues merely cover the cost of effort. At the same time, the social planner provides a rent to the platform when users are H . In practice, this implies that the platform should be offered to choose between two contracts. One is characterized by a more relaxed cap on data, but a higher transfer to users, to compensate them for the larger privacy loss. The other contract entails a more stringent (and socially optimal) data cap, but a lower compensation.

Proposition 5 establishes two standard results in the literature of regulation with asymmetric information: the “no distortion at the top” for the platform

with users having a low naivety, and no informative rents for the platform with users having a high naivety. When users have a low naivety, the platform must be given an informative rent to induce its effort. Such rent is increasing in the effort exerted in the case of L users. Then, in order to reduce the informative rent, the platform's effort a_L is distorted downward, thus implying the necessity to relax the constraint on data in the case of L users. The second best solution is illustrated in Figure 2.

Interestingly, by comparing the results of Proposition 5 with those of Proposition 2, we find that the possibility to compensate users for data affects not only the extent, but also the direction of the distortion of the data caps (and consequently of the platform's effort). In fact, $d_L^{SB'} > d_L^{**}$ if user compensations are possible (i.e., the data cap should be raised), whereas $d_L^{SB} < d_L^*$ when compensations are ruled out (i.e., the data cap is more stringent). Recall also that $d_L^{**} > d_L^*$, i.e. compensations entail a higher first best level of data than no compensation. Then, when users have a severe naivety, second best regulation entails the further separation of the optimal level of data with and without compensations, relative to the first best. Conversely, for H users, second best regulation produces opposite effects. Specifically, $d_H^{SB'} = d_H^{**}$ if compensations are possible, while $d_H^{SB} > d_H^*$ without compensations.

The opposite results for H and L users depend on the fact that, if compensations are feasible, the distortion of the effort is employed to *reduce* the informative rent in a second best scenario, while the rent itself is provided by leaving part of the transfer within the platform. Then, L users are subject to a downward distortion of effort, in order to reduce the informative rent. Conversely, if compensations are ruled out, the distortion of the effort is needed to *provide* the information rent in the second best. Then, L users are subject to an upward distortion of effort, in order to provide the necessary informative rent to the truthful revelation of information by the platform in the presence of H users.

6 Extension: revenue-increasing effort

In our baseline specification, the effort is only a cost for the platform, as it does not increase its revenue. In practice, however, the platform's effort to ensure a safer online environment might also have a positive effect on profits. For example, the platform's effort to protect consumers' data from breaches or unauthorized use by third parties might spur usage, because users anticipate a lower risk of privacy.

In this Section we assume that the platform's effort decreases the amount of data released on the platform, for a given level of usage, but –differently from our baseline model–, its effect has decreasing returns. In particular, we assume that the platform can exert effort $a \geq 0$. The use x of the platform releases data $d(a) = xe^{-a}$, so that $d(0) = x$, $d_a < 0$ and $d_{aa} > 0$ for all $a \geq 0$. To maximize tractability, we assume a specific functional form for the users' utility: $u(x) = 1 - e^{-x}$, and we focus on the case where transfers are absent, i.e., $F = 0$.

Users choose usage to maximize their surplus $S(x, a) = u(x) - \tilde{c}_j x e^{-a}$, whose f.o.c. with respect to x provides the level of usage $\tilde{x}_j(a) = -\ln \tilde{c}_j + a$. The increase in the platform's effort reduces the risk to the user's privacy, making him more willing to use the platform. We assume that $\tilde{c}_j < e^{-1}$ to ensure that the data produced $\tilde{x}_j(a)e^{-a}$ are decreasing in effort, despite the increase of usage.

The platform chooses the effort to maximize the profit $\pi(\tilde{x}_j, a) = \gamma \tilde{x}_j(a) - \psi(a)$. Straightforward analysis allows to obtain the platform's preferred level of effort

$$a_j^{U''} : \psi'(a_j) = \gamma \quad (23)$$

for all j . Effort has a direct cost, but also a benefit for the platform, as it increases the users' willingness to use it and thus the platform's revenues. In the optimum, the platform chooses the effort such that its marginal cost equals to the marginal revenues obtained through the additional effort .

The welfare maximizing level of effort $a_j^{*''}$ solves problem (2), given the usage

$$\tilde{x}_j(a) = -\ln \tilde{c}_j + a.$$

Proposition 6 *If the platform's effort decreases the production of data at decreasing rates, the profit-maximizing level of effort is positive but lower than the first best level (i.e., $a_j^{*''} > a_j^{U''} > 0$).*

Proof. See Appendix A. ■

From a welfare perspective, effort has a direct cost, but three benefits. By inducing higher usage, effort increases the platform's revenues, and the users' utility. Moreover, a higher effort reduces the data provided, and thus the total privacy cost. The platform does not internalize the two benefits of effort related to user utility and privacy costs, therefore the platform chooses an inefficiently low level of effort.

Notably, since the amount of data produced is decreasing in effort but increasing in users' naivety, then a platform with H users can save on effort if it collects the same amount of data of the case of L users. Thus, if effort is not observable, results are qualitatively analogous to those analyzed in Section 4.2.

7 Discussion and conclusions

In a world where data is a cheap but valuable resource, consumers are most vulnerable to the risk of exploitation of their personal data by platforms, which are also endowed by significant market power owing to strong network externalities. Users' naivety further aggravates this market power and the potential of exploitation, by making it easier for platforms to collect large amount of data.

The time has come for the question of whether competition law or regulation is better placed to deal with the challenges arising from the digitalization of the economy. Some forms of ex-ante regulation, so far applied only to public utilities, might in principle be adopted for big digital companies as well. Indeed, many countries – e.g., the European Union, UK, Australia – are already rethinking

their policy towards digital giants and making a case for the creation of a new Authority for data regulation (Cremer et al., 2019; Stiglitz, 2019; DeStreel et al., 2020).

Still, there is currently no definite framework for regulating users' data, therefore it is not clear how the problem should be addressed. The current privacy legislation already puts constraints on the collection and processing of data. However, the flexibility that necessarily accompanies general provisions leaves to platforms some degree of discretion on what "legitimate" and "absolutely necessary" are. Platforms need not only rules, but also incentives to actively invest in data-protection initiatives, else they will take the minimum actions to formally comply to privacy law. We study how these incentives should be designed, and we focus on the case of big digital platforms, where the potential for the exploitation of users' data is larger.

This paper is a first attempt to analyze the role of a prospective data regulator in addressing the privacy problem, with complementary tools to those currently adopted through the provisions of the privacy legislation. We focus on the case of a platform that collects sensible, personal data about its users. This data constitutes a major and extensive source of revenues for the platform, which therefore lacks the incentive to try to reduce the users' privacy loss. Limits on the amount of data collected has the potential to discipline platforms on privacy concerns. For example, a social planner could mandate the anonymization of data, or impose limits to the access to specific data, such as the contact list or the geolocation information while using a smartphone. However, this work also shows that these limits may present two drawbacks. First, if they are combined with mechanisms that reward users for the data collected, users' privacy loss might be worsened. Second, in many situations it might be difficult to directly monitor the effort of the platform, and set limits on a case-by-case basis. When this happens, a menu of different caps, rather than a single constraint on data, can provide to the platform the incentive to invest in data-saving effort.

Our analysis is a first attempt to provide useful direction to policymakers regarding the instruments for an ex ante intervention in digital markets. We consider a type of intervention that is not particularly demanding, as it is based on the amount (and eventually type) of data collected by platforms. We explore how platforms can be induced to provide the desired service to users, while exerting an effort to protect their privacy. Notably, in practice caps exist in the form of limits on downloading and uploading of data, but are implemented only by some Internet Service Providers, and they rarely impact users other than the highest users (FCC, 2013). Other policies also consider the role of data access or data sharing while accounting for the impact on privacy. For example, to reduce the platforms' exclusive control over the data they collect, the recent Digital Market Act imposes data sharing obligations, possibly after some degree of aggregation and masking. While such measures are primarily intended to foster competition in digital markets, in this paper we show how limits to data collection could also be used to redistribute surplus between platforms and users.

Acknowledgement This work has been partially supported by Ministero dell'Istruzione, dell'Universita' e della Ricerca Award TESUN- 83486178370409 finanziamento dipartimenti di eccellenza CAP. 1694 TIT. 232 ART. 6.

References

- ACCC (2019). Digital platforms inquiry, final report. Report 06/19 1545, Australian Competition and Consumer Commission.
- Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, forthcoming.
- Acquisti, A., L. Brandimarte, and G. Loewenstein (2015). Privacy and human behavior in the age of information. *Science* 347(6221), 509–514.
- Acquisti, A., C. Taylor, and L. Wagman (2016, June). The economics of privacy. *Journal of Economic Literature* 54(2), 442–92.
- Acquisti, A. and H. R. Varian (2005). Conditioning prices on purchase history. *Marketing Science* 24(3), 367–381.
- Ainslie, G. W. (1992). *Picoeconomics: the strategic interaction of successive motivational states within the person*. Cambridge University Press.
- Armstrong, M. and D. E. Sappington (2006, June). Regulation, competition and liberalization. *Journal of Economic Literature* 44(2), 325–366.
- Baron, D. P. and R. B. Myerson (1982). Regulating a monopolist with unknown costs. *Econometrica* 50(4), 911–930.
- Baye, M. R. and D. E. M. Sappington (2020). Revealing transactions data to third parties: Implications of privacy regimes for welfare in online markets. *Journal of Economics & Management Strategy* 29(2), 260–275.
- Besanko, D. and D. Sappington (1987). *Designing Regulatory Policy with Limited Information*. New York: Harwood Academic Publishers.
- Bloch, F. and G. Demange (2018). Taxation and privacy protection on internet platforms. *Journal of Public Economic Theory* 20(1), 52–66.
- Bourreau, M., B. Caillaud, and R. De Nijs (2018). Taxation of a digital monopoly platform. *Journal of Public Economic Theory* 20(1), 40–51.

- Braulin, F. C. and T. Valletti (2016). Selling customer information to competing firms. *Economics Letters* 149, 10 – 14.
- Cabral, L., J. Haucap, G. Parker, G. Petropoulos, T. Valletti, and M. Van Alstyne (2021). The eu digital markets act: a report from a panel of economic experts. Report, European Commission - Joint Research Centre.
- Casadesus-Masanell, R. and A. Hervas-Drane (2015). Competing with privacy. *Management Science* 61(1), 229–246.
- Choi, J. P., D.-S. Jeon, and B.-C. Kim (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics* 173, 113 – 124.
- Cremer, J., Y.-A. De Montjoye, and H. Schweitzer (2019). Competition policy for the digital era. Report B-1049, Directorate-General for Competition - European Commission.
- DeStreel, A., M. Cave, R. Feasey, J. Kramer, and G. Monti (2020). Digital markets acts: making economic regulation of platforms fit for the digital era. Report, Centre on Regulation in Europe.
- Dosis, A. and W. Sand-Zantman (2022). The ownership of data. Available at SSRN: <https://ssrn.com/abstract=3420680> or <http://dx.doi.org/10.2139/ssrn.3420680>.
- EC (2020). Commission launches consultation to seek views on digital services act package. Press release IP/20/962, European Commission.
- FCC (2013). Policy issues in data caps and usage-based pricing. Annual report, Federal Communications Commission - Open Internet Advisory Committee.
- Fudenberg, D. and J. Villas-Boas (2007). *Behavior-based price discrimination and customer recognition*. Elsevier, Amsterdam.
- Gans, J. (September 24, 2020). The second wave is here. have you installed the covid alert app? *The Conversation*.
- Ichihashi, S. (2020, February). Online privacy and information disclosure by consumers. *American Economic Review* 110(2), 569–95.
- Iossa, E. and F. Stroffolini (2005). Price cap regulation, revenue sharing and information acquisition. *Information Economics and Policy* 17(2), 217 – 230.

- Jullien, B., Y. Lefouili, and M. Riordan (2020). Privacy protection, security, and consumer retention. Mimeo.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–291.
- Krähmer, D. and R. Strausz (2022). Optimal non-linear pricing with data-sensitive consumers. *American Economic Journal: Microeconomics* (forthcoming).
- Laffont, J.-J. (1994). The new economics of regulation ten years after. *Econometrica* 62(3), 507–537.
- Laibson, D. (1997, 05). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics* 112(2), 443–478.
- Lefouili, Y. and Y. L. Toh (2017). Privacy regulation and quality investment. TSE Working Paper, n. 17-795, Toulouse 26296.
- Montes, R., W. Sand-Zantman, and T. Valletti (2019). The value of personal information in online markets with endogenous privacy. *Management Science* 65(3), 1342–1362.
- Solove, D. J. (2013). Introduction: Privacy self-management and the consent dilemma. *Harvard Law Review* 126(7), 1880–1903.
- Stiglitz, J. (2019). *People, Power, and Profits. Progressive Capitalism for an Age of Discontent*. W.W. Norton and Company.
- Taylor, C. R. (2004). Consumer privacy and the market for customer information. *The RAND Journal of Economics* 35(4), 631–650.
- Tucker, C. (2018). Privacy, algorithms, and artificial intelligence. In A. Gans and Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

A Proofs

Proof of Proposition 2 Let us write the Lagrangean function L for Problem (8) subject to the IC constraint of the platform with H users (we verify the IC constraint for the platform with L users afterwards).

$$\begin{aligned} L = & \lambda[u(u_x^{-1}(\tilde{c}_H)) - c(u_x^{-1}(\tilde{c}_H) - a_H) + \alpha(\gamma u_x^{-1}(\tilde{c}_H) - \psi(a_H))] + \\ & + (1 - \lambda)[u(u_x^{-1}(\tilde{c}_L)) - c(u_x^{-1}(\tilde{c}_L) - a_L) + \alpha(\gamma u_x^{-1}(\tilde{c}_L) - \psi(a_L))] + \\ & + \mathcal{L}[\gamma u_x^{-1}(\tilde{c}_H) - \psi(a_H) - \gamma u_x^{-1}(\tilde{c}_L) + \psi(a_L - \Delta)], \end{aligned}$$

where $\mathcal{L} \geq 0$ is the lagrangean multiplier of the IC constraint for the platform with H users.

The FOC w.r.t. a_H and a_L are:

$$\lambda(c - \alpha\psi'(a_H)) - \mathcal{L}\psi'(a_H) = 0 \quad (1)$$

$$(1 - \lambda)(c - \alpha\psi'(a_L)) + \mathcal{L}\psi'(a_L - \Delta) = 0. \quad (2)$$

In $a_H = a_H^*$, it holds $\alpha\psi'(a_H) = c$. Hence, the l.h.s. of condition (27) in the first best solution is equal to $-\mathcal{L}\frac{c}{\alpha} < 0$. It follows that $a_H^{SB} < a_H^*$. Moreover, it holds $\alpha\psi'(a_L^*) = c$. Hence, condition (28) in the first best solution ($a_L = a_L^*$) becomes $\mathcal{L}\psi'(c - \Delta) > 0$. This implies that $a_L^{SB} > a_L^*$. The data disclosed is $d_j^{SB} = u_x^{-1}(\tilde{c}_j) - a_j$. Given that $a_H^{SB} < a_H^*$, we obtain $d_H^{SB} > d_H^*$. Moreover, given $a_L^{SB} > a_L^*$, there holds $d_L^{SB} < d_L^*$. Profits are $\pi_j = \gamma\tilde{x}_j - \psi(a_j)$. In the case of a H platform, $a_H^{SB} < a_H^*$ imply $\pi_H^{SB} > \pi_H^*$. Conversely, in the case of a L platform, $a_L^{SB} > a_L^*$ imply $\pi_L^{SB} < \pi_L^*$.

The IC constraint of the platform with L users, $\gamma\Delta \geq \psi(a_L^{SB}) - \psi(a_H^{SB} + \Delta)$, is also verified as $\gamma\Delta = \psi(a_L^{SB} - \Delta) - \psi(a_H^{SB})$ and $\psi(a_L^{SB}) - \psi(a_L^{SB} - \Delta) < \psi(a_H^{SB} + \Delta) - \psi(a_H^{SB})$.

Proof of Proposition 3 Using the users' supply of data (1), let us write the

lagrangean function L associated to the constrained maximization Problem (14):

$$L = u(u_x^{-1}(\tilde{c}_j)) - c(u_x^{-1}(\tilde{c}_j) - a_j) - F_j + \alpha [F_j + \gamma u_x^{-1}(\tilde{c}_j) - \psi(a_j)] + \mathcal{L}[F_j + \gamma u_x^{-1}(\tilde{c}_j) - \psi(a_j)], \quad (3)$$

where $\mathcal{L} \geq 0$ is the lagrangean multiplier associated to the platform's participation constraint. The Kuhn-Tucker conditions are:

$$\frac{\partial L}{\partial a_j} = c - \alpha \psi'(a_j) - \mathcal{L} \psi'(a_j) = 0 \quad (4)$$

$$\frac{\partial L}{\partial F_j} = -1 + \alpha + \mathcal{L} = 0 \quad (5)$$

$$\mathcal{L} [F_j + \gamma u_x^{-1}(\tilde{c}_j) - \psi(a_j)] = 0 \quad (6)$$

$$\mathcal{L} \geq 0, a_j \geq 0. \quad (7)$$

From condition (5), we obtain $\mathcal{L} = 1 - \alpha$. Then, condition (4) can be rewritten as $c - \alpha \psi'(a_j) - (1 - \alpha) \psi'(a_j) = 0$. Hence, $c = \psi'(a_j^{**})$.

Two cases may arise.

Case 1: $\alpha < 1$. Given that $\mathcal{L} = 1 - \alpha$, then for any $\alpha < 1$, it must be: $\mathcal{L} > 0$, implying that $F_j + \gamma u_x^{-1}(\tilde{c}_j) - \psi(a_j) = 0$ from condition (6). Then, $F_j^{**} = -\gamma u_x^{-1}(\tilde{c}_j) + \psi(a_j^*) = 0$ for any $\alpha < 1$, i.e. $\pi(\tilde{x}_j, a_j^{**}, F_j^{**}) \geq 0$.

Case 2: $\alpha = 1$. In this case, we have that $\mathcal{L} = 0$ from (5). Then, from (6), we have that any F_j^{**} such that $\pi(\tilde{x}_j, a_j^{**}, F_j^{**}) \geq 0$ is a solution of the problem.

Proof of Proposition 5. Let us analyze the problem subject to the IC constraint of the platform with H users and the individual rationality (IR) constraint for the platform with L users (we verify the IC for L and the IR for H afterwards). The lagrangean function is:

$$\begin{aligned} L = & \lambda [u(\tilde{x}_H) - c(\tilde{x}_H - a_H) - F_H + \alpha (F_H + \gamma \tilde{x}_H - \psi(a_H))] + \\ & + (1 - \lambda) [u(\tilde{x}_L) - c(\tilde{x}_L - a_L) - F_L + \alpha (F_L + \gamma \tilde{x}_L - \psi(a_L))] + \\ & + \mathcal{L}_1 [F_H + \gamma \tilde{x}_H - \psi(a_H) - F_L - \gamma \tilde{x}_L + \psi(a_L - \Delta)] + \\ & + \mathcal{L}_2 [F_L + \gamma \tilde{x}_L - \psi(a_L)]. \end{aligned}$$

where $\mathcal{L}_1, \mathcal{L}_2 \geq 0$ are the lagrangean multipliers.

The Kuhn-Tucker conditions are:

$$\frac{\partial L}{\partial a_H} = \lambda(c - \alpha\psi'(a_H)) - \mathcal{L}_1\psi'(a_H) = 0 \quad (8)$$

$$\frac{\partial L}{\partial a_L} = (1 - \lambda)(c - \alpha\psi'(a_L)) + \mathcal{L}_1\psi'(a_L - \Delta) - \mathcal{L}_2\psi'(a_L) = 0 \quad (9)$$

$$\frac{\partial L}{\partial F_H} = -\lambda(1 - \alpha) + \mathcal{L}_1 = 0 \quad (10)$$

$$\frac{\partial L}{\partial F_L} = -(1 - \lambda)(1 - \alpha) - \mathcal{L}_1 + \mathcal{L}_2 = 0 \quad (11)$$

$$\mathcal{L}_1 [F_H + \gamma\tilde{x}_H - \psi(a_H) - F_L - \gamma\tilde{x}_L + \psi(a_L - \Delta)] = 0 \quad (12)$$

$$\mathcal{L}_2 [F_L + \gamma\tilde{x}_L - \psi(a_L)] = 0 \quad (13)$$

$$\mathcal{L}_1, \mathcal{L}_2 \geq 0, a_j \geq 0$$

From (11), $\mathcal{L}_2 = (1 - \lambda)(1 - \alpha) + \mathcal{L}_1 > 0$. Then, from (13), the IR constraint for L is binding, i.e. $\pi_L^{SB'} = 0$. Moreover, from (10), $\mathcal{L}_1 = \lambda(1 - \alpha) > 0$. Then, from (12), also the IC constraint for H is binding, i.e. $\pi_H^{SB'} = \psi(a_L) - \psi(a_L - \Delta)$. By substituting $\mathcal{L}_1 = \lambda(1 - \alpha)$ in (8), we obtain $c = \psi'(a_H)$, i.e. $a_H^{SB'} = a_H^{**}$. Moreover, by substituting $\mathcal{L}_1 = \lambda(1 - \alpha)$ in (11), we obtain $\mathcal{L}_2 = 1 - \alpha$. We can thus rewrite condition (9) as:

$$(1 - \lambda)(c - \alpha\psi'(a_L)) + \lambda(1 - \alpha)\psi'(a_L - \Delta) - (1 - \alpha)\psi'(a_L) = 0. \quad (14)$$

After straightforward simplifications, condition (14) can be expressed as:

$$\psi'(a_L) = c - \frac{\lambda}{1 - \lambda}(1 - \alpha)[\psi'(a_L) - \psi'(a_L - \Delta)], \quad (15)$$

which implies $a_L^{SB'} < a_L^{**}$. Then, $F_L^{SB'} = -\gamma\tilde{x}_L + \psi(a_L^{SB'} - \Delta) < F_L^{**}$.

It is easy to verify that the IR constraint for the platform with H users is not binding. In fact, from the binding IC constraint of the H platform, $\pi_H^{SB'} = \psi(a_L) - \psi(a_L - \Delta) > 0$. This implies that $F_H^{SB'} > F_H^{**}$. Moreover, the IC constraint of the platform with L users can be expressed as:

$$F_L + \gamma\tilde{x}_L - \psi(a_L^{SB'}) \geq F_H + \gamma\tilde{x}_H - \psi(\hat{a}_L(a_H^{SB'})). \quad (16)$$

The l.h.s. is $\pi_L^{SB'} = 0$. The r.h.s. can be rewritten as $\pi_H^{SB'} + \psi(a_H^{SB'}) - \psi(\hat{a}_L(a_H^{SB'}))$. By substituting $\pi_H^{SB'} = \psi(a_L^{SB'}) - \psi(a_L^{SB'} - \Delta)$, condition (16) becomes:

$$0 \geq \psi(a_L^{SB'}) - \psi(a_L^{SB'} - \Delta) + \psi(a_H^{SB'}) - \psi(a_H^{SB'} + \Delta). \quad (17)$$

As $a_H^{SB'} + \Delta > a_H^{SB'} > a_L^{SB'} > a_L^{SB'} - \Delta$, the r.h.s. of (17) is negative. Hence, the IC constraint of the platform with L users is satisfied.

Proof of Proposition 6. Let us express welfare in explicit terms given the production function of data $d(x, a) = xe^{-a}$:

$$W(a) = 1 - e^{-\tilde{x}_j(a)} - c\tilde{x}_j(a)e^{-a} + \alpha(\gamma\tilde{x}_j(a) - \psi(a)), \quad (18)$$

where $\tilde{x}_j(a) = -\ln \tilde{c}_j + a$.

The total cost of privacy is $P(a) = c\tilde{x}_j(a)e^{-a}$. In explicit terms, it can be expressed as $P(a) = c(-\ln \tilde{c}_j + a)e^{-a}$. Note that $P(0) > 0$, $\lim_{a \rightarrow \infty} P(a) = 0$, and $P_a = ce^{-a}(1 - a + \ln \tilde{c}_j)$. In particular, a sufficient condition for having $P_a < 0$ is $\ln \tilde{c}_j < -1$, i.e., $\tilde{c}_j < e^{-1}$.

Under the assumption $\tilde{c}_j < e^{-1}$, the welfare function in (18) is concave in a and its maximum satisfies the FOC

$$e^{-\tilde{x}_j(a)} - ce^{-a} + c\tilde{x}_j(a)e^{-a} + \alpha(\gamma - \psi_a(a)) = 0,$$

which can be rewritten as

$$\tilde{c}_j e^{-a} - ce^{-a} + c(-\ln \tilde{c}_j + a)e^{-a} + \alpha(\gamma - \psi_a(a)) = 0. \quad (19)$$

Let us denote with \bar{a}_j the solution of condition (19). We distinguish between two cases.

i) $\pi(\tilde{x}_j, \bar{a}_j) \geq 0$ (i.e., the solution of condition (19) satisfies the platform's participation constraint). In this case, the solution of the constrained welfare maximization problem is $a_j^{*''} = \bar{a}_j$. In $a_j^{U''}$, we have that $\psi_a(a) = \gamma$. Then, the left handside of condition (19) becomes $e^{-a}(\tilde{c}_j - c + c(-\ln \tilde{c}_j + a))$, which is

strictly positive given the assumption $\ln \tilde{c}_j < -1$. Since the FOC (19) in (19) is positive, and the welfare function is concave, we conclude that $a_j^{*''} > a_j^{U''}$.

ii) $\pi(\tilde{x}_j, \bar{a}_j) \geq 0$ (i.e., the solution of condition (19) does not satisfy the platform's participation constraint). If $\pi(\tilde{x}_j, a_j) < 0$, then $a_j^{*''}$ solves the condition $\pi(\tilde{x}_j, a_j) = 0$. Given the concavity of the platform's profit function, we have that $a_j^{*''} > a_j^{U''}$.

A.1 Heterogeneous users

In this Appendix we consider the case of a population of heterogeneous users in a second best scenario without transfers. Let us denote with λ the share of users in the population with type $j = H$, and with $1 - \lambda$ the share of $j = L$ users. As in our main analysis, the effort a and the amount of usage \tilde{x}_j of each user are unobservable by the regulator, whereas the data d_j provided by each user can be observed and verified.

The platform might prefer to induce H users to provide the same amount of data d_L of L users. In particular, the H users' supply of data is $\tilde{d}_H(\hat{a}_H) = \tilde{x}_H - \hat{a}_H$. The effort \hat{a}_H , such that $\tilde{d}_H(\hat{a}_H)$ coincides with $\tilde{d}_L(a_L) = \tilde{x}_L - a_L$, is $\hat{a}_H(a_L) = a_L - \tilde{x}_L + \tilde{x}_H$, which can also be expressed as

$$\hat{a}_H(a_L) = a_L - \Delta, \quad (20)$$

where $\Delta = \tilde{x}_L - \tilde{x}_H$. Given that $\tilde{x}_L > \tilde{x}_H$, then $\Delta > 0$ and $\hat{a}_H(a_L) < a_L$. In order to streamline the exposition, let us focus on internal solution by assuming that $a_L^* - \Delta > 0$.

The regulator offers a menu of incentive compatible contracts, which solves a similar problem to that analyzed in Section 4.2, the only difference being that the platform's participation here is given by the ex-ante expected profit $\pi(a_H, a_L) \geq 0$, instead of constraint (9), where $\pi(a_H, a_L) = \lambda [\gamma \tilde{x}_H - \psi(a_H)] + (1 - \lambda) [\gamma \tilde{x}_L - \psi(a_L)]$.

Given that constraints (9) are not binding in the optimal solution (a_L^{SB}, a_H^{SB}) , then ex-ante participation is not binding as well, so that menu (a_L^{SB}, a_H^{SB}) is also the solution of this problem with heterogeneous users.

A.2 Generic data function

Let us consider a generic function of data, $d = d(x, a)$, with $d_x > 0$, $d_a < 0$, $d_{xx} \geq 0$, $d_{xa} = 0$, i.e. the function is separable in x and a . Users of type j

maximize their utility, net of the perceived privacy cost: $\max_x u(x) - \tilde{c}_j d(x, a) - F$.

Their preferred level of usage solves the FOC

$$u_x(x) = \tilde{c}_j d_x(x). \quad (21)$$

Let us define the function $g(x) = \frac{u_x(x)}{d_x(x, a)}$, and rewrite condition (21) as $g(x) = \tilde{c}_j$.

Given that $g_x = \frac{u_{xx}d_x - u_x d_{xx}}{(d_x)^2} < 0$, we obtain $\tilde{x}_L > \tilde{x}_H$ and $\tilde{d}_L > \tilde{d}_H$, for a given level of effort a . The welfare function can be rewritten as:

$$W = u(\tilde{x}_j) - cd(\tilde{x}_j, a) + \alpha [\gamma \tilde{x}_j - \psi(a_j)]. \quad (22)$$

The FOC of (22) yields the optimality condition $\alpha \psi'(a) = -cd_a(a)$, which corresponds to condition (5) when $d(x, a) = x - a$.

In a second best scenario, for a generic level of effort a , the data extracted from H users is $d(\tilde{x}_H, a)$. The effort \hat{a}_H , such that $d(\tilde{x}_H, \hat{a}_H) = d(\tilde{x}_L, a_L)$, is $\hat{a}_H = f(a_L)$, such that $f' > 0$. The regulator solves the following problem:

$$\max_{a_H, a_L} \lambda [u(\tilde{x}_H) - c\tilde{d}_H + \alpha (\gamma \tilde{x}_H - \psi(a_H))] + \quad (23)$$

$$+ (1 - \lambda) [u(\tilde{x}_L) - c\tilde{d}_L + \alpha (\gamma \tilde{x}_L - \psi(a_L))]$$

$$s.t. \gamma \tilde{x}_j - \psi(a_j) \geq 0 \quad \forall j, \quad (24)$$

$$\gamma \tilde{x}_H - \psi(a_H) \geq \gamma \tilde{x}_L - \psi(\hat{a}_H), \quad (25)$$

$$\gamma \tilde{x}_L - \psi(a_L) \geq \gamma \tilde{x}_H - \psi(\hat{a}_L). \quad (26)$$

Constraint (24) constitutes the participation constraint of the platform j , while constraints (25) and (26) ensure the incentive compatibility of the contract. The Lagrangean function for Problem (23) is:

$$\begin{aligned} L = & \lambda [u(\tilde{x}_H) - cd(\tilde{x}_H, a_H) + \alpha (\gamma \tilde{x}_H - \psi(a_H))] + \\ & + (1 - \lambda) [u(\tilde{x}_L) - cd(\tilde{x}_L, a_L) + \alpha (\gamma \tilde{x}_L - \psi(a_L))] + \\ & + \mathcal{L} [\gamma \tilde{x}_H - \psi(a_H) - \gamma \tilde{x}_L + \psi(f(a_L))], \end{aligned}$$

where $\mathcal{L} \geq 0$ is the lagrangean multiplier of the IC constraint for the platform with H users. The FOC w.r.t. a_H and a_L are:

$$\lambda(-cd_a(a_H) - \alpha\psi'(a_H)) - \mathcal{L}\psi'(a_H) = 0 \quad (27)$$

$$(1 - \lambda)(-cd_a(a_L) - \alpha\psi'(a_L)) + \mathcal{L}\psi'(f(a_L))f' = 0. \quad (28)$$

The solution of the regulator's problem is implicitly defined by the following conditions:

$$\alpha\psi'(a_L^{SB}) = -cd_a(a_L^{SB}) + \frac{\lambda}{1 - \lambda}\psi'(f(a_L^{SB}))f' \left(\frac{-cd_a(a_H^{SB})}{\psi'(a_H^{SB})} - \alpha \right) \quad (29)$$

$$\psi(a_H^{SB}) = \psi(f(a_L)) - \gamma(\tilde{x}_L - \tilde{x}_H) \quad (30)$$

that generalize conditions (12) and (13).