Taxing digital advertising revenue

Carmen Marchiori [1] Enrico Minelli [2]

September 2, 2025

[1]Carmen Marchiori, Università degli Studi di Brescia, Dipartimento di Economia e Management, via San Faustino 74b, 25122, Brescia, Italy. Email: carmen.marchiori@unibs.it

[2]Enrico Minelli, Università degli Studi di Brescia, Dipartimento di Economia e Management, via San Faustino 74b, 25122, Brescia, Italy. Email: enrico.minelli@unibs.it

Abstract

We propose a model of a social media platform in which platform's profit, consumer surplus and users' participation all depend on the platform's choice of data extraction intensity. The equilibrium of the model depends on two parameters: a parameter of 'data squeezing' which measures the platform's effectiveness in turning users' personal data into commercial revenue, and a parameter of 'privacy concern', which describes the distribution of the awareness of privacy loss in the population. We use this framework to analyze the implications of alternative business models (Pure Advertising; Pure Subscription; Freemium), and to provide insights into recent proposals to introduce a tax on digital advertising revenue.

JEL classification: D42. D43, D62

Keywords: Two-sided platforms, Digital tax, Privacy, Business models

1 Introduction

With approximately 4.9 billion users worldwide spending an average of three hours a day on digital platforms[1], social media have undeniably become an integral part of our daily life, fundamentally reshaping how we communicate, interact, and consume information.

While this profound transformation has certainly brought about benefits such as enhanced accessibility and a greater variety of content, there are growing concerns regarding the potential dangers and negative implications of social media.

Platforms like Facebook, Instagram, TikTok, X and YouTube provide personalized content to users, extract data from them, and finance themselves, partially or fully, by selling targeted ads space to firms. Within this business model, platform profits are tightly linked to user engagement: increased engagement facilitates more precise data extraction, which in turn enhances the value of targeted ads.

Although the specifics of algorithmic design remain largely opaque, recent empirical evidence suggests that engagement optimization is central to platform design. In particular, Guess et al. (2023), in collaboration with Meta, document that switching from algorithmically curated content to chronologically ordered feeds results in a 26% reduction in time spent on Facebook, suggesting that current algorithms are intentionally structured to maximize user attention.

This is concerning from a welfare perspective. As highlighted by McLoughlin et al. (2024), user engagement is positively correlated with outrage, and outrage with misinformation - both factors potentially degrading the quality of social and political discourse.[2] Beyond informational effects, large-scale incentivized experiments by Bursztyn et al. (2024) provide evidence of significant negative spillovers on individual well-being from excessive use of Instagram and TikTok; while Braghieri et al. (2022) document rising anxiety and depression linked to Facebook access.

The accumulation of evidence on the negative externalities of a business model based on engagement-maximization has prompted leading economists like Romer (2021), and Acemoglu and Johnson (2024) to advocate for the introduction of a tax on digital advertising revenues, with the explicit goal of inducing social media companies to switch to less socially detrimental

[1] See, for example, https://www.forbes.com/advisor/business/social-media-statistics/, and https://whatagraph.com/blog/articles/how-much-time-do-people-spend-on-social-media

[2] For a comprehensive survey of the societal and political impacts of social media, see also Zhuravskaya et al. (2020).

monetization strategies, such as subscription-based models, where revenues depend on the sustained quality of content and user experience rather than user engagement and clicks.

This idea has already gained traction in policy circles. Several U.S. States and European countries have proposed or enacted digital advertising taxes, albeit at rates significantly lower than those advocated by Romer, and Acemoglu and Johnson.[3] In parallel, the European Union enacted the General Data Protection Regulation (GDPR) in 2018, introducing significant constraints on personal data usage.[4] Unsurprisingly, these interventions have triggered strong reactions from digital platforms. Some companies have initiated legal proceedings to suspend the implementation of digital advertising taxes; while others, most notably Meta (Facebook and Instagram) and Alphabet (YouTube), have preemptively modified their business models by introducing 'pay or consent' schemes. Under this structure, which is often referred to as Freemium, users can either pay for an ad-free experience or consent to personalized ads and data extraction in exchange for free access.

In this paper, we develop a model of a social media platform to assess the implications of alternative monetization strategies (Pure Advertising; Pure Subscription; Freemium), and to provide insights into the design and effectiveness of taxation and other policy interventions. Through the lens of our framework, we can address the following questions: What is the minimum level of taxation on digital advertising revenues required to induce a switch to Pure Subscription? How does a platform's adoption of Freemium alter the design and effectiveness of such a tax? How does taxation of digital advertising revenues interact with other policy instruments, such as privacy protection policies?

Framework. The model features a single platform providing content/services to users while extracting data on their characteristics. The platform chooses the intensity of data extraction, which determines the granularity with which users are segmented in the product market according to their willingness to pay. This information is then transferred to the commercial sector (a monopolistic seller), which leverages it to discriminate among different segments of users. The platform proposes a take it or leave it offer to the commercial

[3]Maryland was the first U.S. State to adopt such a tax, with tax rates ranging from 2.5% to 10% depending on a company's global annual gross revenue. The proposals under discussion in other countries involve tax rates in the order of 3%, in contrast with the 50% rate suggested by Acemoglu for companies with annual digital ad revenue over $500 million.

[4]Goldberg et al. (2024) and Aridor et al. (2020) find evidence of online-traffic and revenue reductions of around $12 - 13\%$ following the implementation of the GDPR.

sector, and fully appropriates the commercial profit generated through targeted transactions with its users; we refer to this profit as Digital Advertising Revenue.

For a given mass of users, the platform's profit increases with the intensity of information extraction. This relationship is mediated by a single parameter that links the degree of data extraction to the number of identifiable market segments. We call this parameter *data squeezing*, as it captures the platform's ability to exploit fine-grained information from its users for commercial targeting purposes.

However, the size of the user base is itself endogenous to the platform's choice of data extraction intensity: each individual decides whether to access the platform by comparing the utility derived from its services with the disutility arising from the loss of privacy. The distribution of privacy costs across the population is governed by a second parameter, which is an index of the saliency of *privacy concern*. This parameter captures heterogeneity in users' sensitivity to privacy loss and determines how data extraction affects overall platform participation.

Outline of results. We first study the benchmark case of a platform characterized by a Pure Advertising business model. We show that the minimum level of taxation required to induce a switch to Pure Subscription is decreasing in *data squeezing* and increasing in *privacy concern*. This result suggests that complementary policies aimed at limiting the platform's capacity to exploit user data and at raising awareness about the risks of privacy loss may play a critical role in making the required level of taxation politically feasible.

We then turn to the case of a platform operating under a Freemium business model. As noted earlier, this monetization strategy has gained increasing prominence in practice, with major technology firms such as Meta and Alphabet adopting it in response to growing regulatory pressure and the prospect of digital advertising taxation. We characterize the platform's optimal choices as a function of the tax levied on Digital Advertising Revenue. We obtain three main results. First, the level of information extraction under Freemium is always higher than under Pure Advertising. Second, the level of taxation required to induce a switch to Pure Subscription is higher in the Freemium case. Third, for levels of taxation not high enough to induce the switch, increasing the tax rate has the perverse effect of inducing *higher* levels of information extraction.

Given the increasingly widespread adoption of the Freemium model, these findings raise important questions regarding the desirability of a digital advertising tax. In the final part of the paper, we employ our framework to

explore this issue. We derive analytical expressions for both users' net surplus and the platform's surplus, and argue that - for plausible values of the key parameters - the introduction of a tax on digital advertising revenues *is* justified on welfare grounds.

Related Literature. Anderson and Jullien (2015) review the literature on media platforms, a special class of two-sided platforms[5] delivering content to users while selling ad space to advertisers. A seminal paper in this literature is Anderson and Coate (2005), which studies a market where competing platforms choose both the type of content offered and the intensity of advertising. Our model differs from Anderson and Coate (2005), and the literature that followed, in two main respects: 1) we provide a micro-foundation of the advertisers' (and platform's) revenues which links these revenues, as well as consumer surplus, directly to the level of information extraction exerted by the platform; 2) we explicitly account for heterogeneous user preferences over privacy. Furthermore, our focus differs: we are interested in the design of taxation and other regulatory policies, which we study in the simplest possible setting of a monopoly platform and a monopoly seller; the literature surveyed by Anderson and Jullien (2015), on the other hand, does not discuss policy interventions, but carefully considers the effects of competition between platforms and/or sellers on market outcomes.

Closer to ours is the paper by Bloch and Demange (2018). They consider a monopolistic platform, and allow for buyers' heterogeneity in the cost of data extraction (point 2) above). Like us, they are interested in the design of tax policies. Where we differ from them is in the micro-foundation of platform's revenue (point 1) above). They assume separate, specific functional forms for the users' benefit from advertising and for the platform's revenue from the sale of advertising space; we derive both elements from a stylized model of information extraction and market segmentation. Our approach leads us to isolate two basic parameters affecting equilibrium choices, 'data squeezing' and 'privacy concern', which we can also use to describe policy interventions complementing the choice of taxation.

The modeling of platform's information extraction and market segmentation, our main innovation with respect to the literature discussed so far, is the object of an important literature reviewed by Bergemann and Bonatti (2019), with recent contributions by Galperti and Perego (2023) and Bergemann and Bonatti (2024). In this literature, the focus is on the role of platforms as 'information brokers': the platform designs mechanisms to elicit information from users in an incentive compatible way, and at the same

[5]See the seminal contributions by Rochet and Tirole (2003) and Armstrong (2006), and the recent textbook treatment by Belleflamme and Peitz (2021).

time controls the information transmitted to the sellers (and sometimes also to the users themselves). Subtle issues arise with respect to the possible correlation of the users' information, or the effect of information dissemination on the intensity of sellers' competition and the platform's bargaining power. In our model, the complexities of information elicitation and transmission are sidestepped by assuming that the platform chooses directly the intensity of data extraction, and that the way in which this choice translates into information used by the seller is controlled by a single parameter ('data squeezing'). We can think of this parameter as a reduced form representation of more fundamental conditions affecting the platform ability to extract information and monetize it through the design of information mechanisms. It is comforting that this literature obtains results supporting a high degree of bargaining power of the platform *vis-à-vis* the sellers (see e.g. Bergemann and Bonatti (2024), Proposition 4), which we assume. On the other hand, it is clear that modeling the effect of an increase in privacy protection simply in terms of variations of our two parameters does not allow to capture the subtle effects highlighted by Galperti and Perego (2023) and Bergemann and Bonatti (2024). With respect to this literature, the simplification in the modeling of platform as 'information broker' allows us a) to consider the interaction between information extraction and users' participation, b) to focus on the platform's choice of business model, and on the interaction between this choice and taxation and regulation policies.

Finally, in a recent discussion paper, Acemoglu, Huttenlocher, Ozgdaglar and Siderius (2024) provide, like us, a comparison of different platform business models and their effects on welfare. Their focus is on the manipulative role of advertising in shaping consumer behavior. In their model, consumers are uncertain about their own willingness to pay and a subset - referred to as *naive* - becomes overly optimistic about product quality after being exposed to ads. A key condition for the Pure Ad business model to yield higher profits than Pure Subscription (which always holds in our framework) is that naive consumers are sufficiently susceptible to manipulation. Under this condition, they show that welfare is lower under Pure Ad than under Pure Subscription for both naive and sophisticated consumers, as advertising-driven optimism raises demand and leads to higher prices. Our framework differs from theirs in several key respects. First, their model does not incorporate any form of information extraction by the platform, whereas in our setting this is a central feature. We explicitly model how user data is commercially exploited, providing a microfoundation for the link between information extraction, market segmentation, and revenue generation. Second, our analysis does not rely on advertising being manipulative. Advertising in our framework can enhance consumer surplus by expanding market reach. However, the ability

to segment the market and target users means that a substantial share of the resulting surplus is appropriated by the seller and the platform. We also differ in our emphasis on the analysis of taxation of digital ad revenues, and of the complementary role of other policies, such as privacy protection policies.

In the next section we introduce our model of information extraction and market segmentation. In section 3 we study the benchmark case of the Pure Ad business model, and in section 4 we address the new issues related to the introduction of Freemium. Section 5 provides insights into the desirability of a tax on digital ad revenues from a welfare perspective. Section 6 summarizes the main findings and concludes.

2 Profiting from information

In our model, a platform offers content/services to its users and decides the intensity with which to extract data from them (subsection 2.2). This information is valuable because it enables market segmentation, which in turn can be used by the production sector to increase its profits (subsection 2.1).

2.1 Market segmentation

The production sector is represented by a monopolist, with production cost normalized to zero, selling a homogeneous good to a mass $S \leq 1$ of platform users. Each user i has a unit demand for the good, and is characterized by a willingness to pay (WTP), w_i, uniformly distributed over the interval $[0, A]$.

As previously mentioned, the platform can leverage the data extracted from its users to partition them into target groups, thereby enabling the monopolist to increase earnings. Specifically, we model the process of market segmentation as a sequence of n bisections of the users' WTP. With $n = 1$, the platform can distinguish consumers with WTP in $[0, A/2]$ from those in $[A/2, A]$; with $n = 2$, it can distinguish consumers with WTP in $[0, A/4]$ from those in $[A/4, A/2]$, $[A/2, 3A/4]$, and $[3A/4, A]$, respectively. More generally, with n bisections the market can be divided into $N = 2^n$ cells $[0, A/N]$, $[A/N, 2A/N]$, $\ldots$ $[(N-1)A/N, A]$. Note that, $n = 0$ implies that there is no additional information for the monopolist, who will simply face a linear demand $Q(p) = S(1 - \frac{1}{A}p)$.

In the next proposition we characterize the optimal pricing and profit of the monopolist for any given n.

Proposition 1. *With $n \in \mathbb{N}_+$ bisections, the monopolist profit is:*

$$\Pi(n) = \frac{AS}{4}\left(2 - \frac{2^{n+1}-1}{2^{2n}}\right) \tag{1}$$

Proof. Let $N = 2^n$ be the number of cells in the partition of consumers. Order cells form 1 to N by increasing average WTP. In each cell, there is a mass S/N of consumers. Consumers in cells $J = 2, \dots N$ buy a quantity S/N at the lowest possible price $p(J) = (J-1)A/N$. Those in the cell with the lowest average WTP, i.e. cell $J = 1$, buy a quantity $S/2N$ at the "monopoly price" $p(1) = A/2N$.

Hence, for a given N, the total profit is:

$$\begin{aligned}\Pi(N) &= \frac{S}{N}\left[\frac{A}{N} + 2\frac{A}{N} + 3\frac{A}{N} + \dots (N-1)\frac{A}{N}\right] + \frac{S}{2N}\frac{A}{2N} = \\ &\frac{S}{N}\left[\frac{A}{N}\left(\sum_{a=1}^{N-1} a\right)\right] + \frac{S}{2N}\frac{A}{2N}\end{aligned}$$

Recalling that $N = 2^n$, this reduces to:

$$\Pi(n) = \frac{AS}{2^n}\left[\frac{1}{2^n}((2^n-1)2^{n-1})\right] + \frac{AS}{2^{2(n+1)}}$$

which simplifies to (1). □

Similarly, we can derive an expression for total consumer surplus as a function of n.

Corollary 1. *With $n \in \mathbb{N}_+$ bisections, consumer surplus is:*

$$CS(n) = \frac{AS}{4}\frac{2^{n+2}-3}{2^{2n+1}} \tag{2}$$

Note that, for $n = 0$, we obtain the monopoly profit and consumer surplus $\Pi(0) = \Pi^M = \frac{AS}{4}$, and $CS(0) = CS^M = \frac{AS}{8}$; while, as n goes to ∞, we tend to the solution of a perfectly discriminating monopolist, with $\Pi(n) \to \frac{AS}{2}$, and full surplus extraction, $CS(n) \to 0$ (see Figure 1).

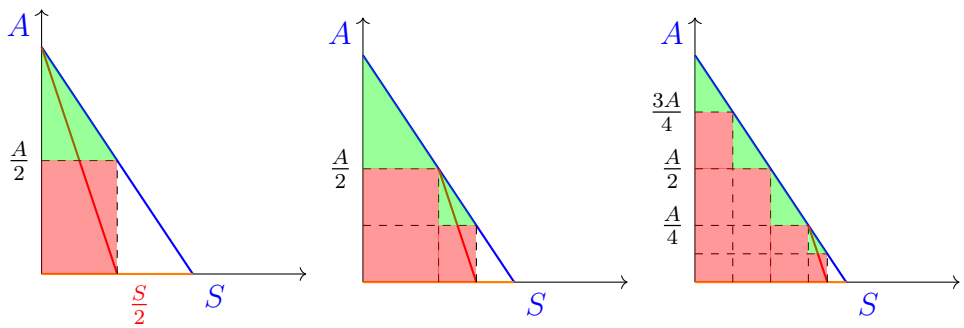

Figure 1: Market segmentation and pricing. The Seller's profit is in red, and consumer surplus in green, for $n = 0$ (no bisection, left), $n = 1$ (one bisection, center) and $n = 2$ (two bisections, right).

REMARK 1: Consumer surplus is non-monotonic in n. If $N = 2^n$ is the number of cells in the partition, we can write $CS(N) = \frac{AS}{4}\frac{4N-3}{2N^2}$. This expression is increasing in N for $N \leq 3$, and decreasing thereafter. This non-monotonic relationship reflects the trade-off inherent in market segmentation: on the one hand, it expands consumer coverage by allowing the firm to reach more buyers; on the other hand, it increases the firm's ability to extract surplus, potentially reducing consumer welfare. In our model of segmentation by bisection with uniformly distributed WTP, this tension resolves in favor of the first effect when going from no segmentation to one bisection, ($n = 1$, $N = 2$), but turns against consumers as soon as $n = 2$ and $N = 4$. This non-monotonicity will play a role when discussing Welfare (see section 5).

2.2 Information extraction

The number n of bisections of users' WTP (i.e., the degree of precision with which the platform is able to segment the market) is an increasing function of the platform's *data-extraction intensity*, which we denote by $x \in [0, 1]$. For simplicity, we assume that n is a continuous variable, linearly increasing in x:

$$n(x) = kx \tag{3}$$

with $k \geq 0$. For a given x, the number of cells in the partition (i.e. the number of target groups) is $N(x) = 2^{n(x)} = 2^{kx}$. The parameter k can be interpreted as a measure of the efficacy of data extraction. For example, with full data extraction intensity ($x = 1$), the platform is able to classify the users in: only $N = 2$ target groups if $k = 1$; and $N = 1024$ target groups if $k = 10$.

The variable x is a summary measure of how invasive the platform is in extracting information from users, for example by choosing personalized content to foster engagement, or by tracking users' activity across different apps. The parameter k, on the other hand, represents the efficacy with which the collected data can be transformed into *commercially valuable* information; crucially, as we will discuss later, k may be affected by regulatory policies that limit or shape the use of data.

In section 2.1, we took the mass S of platform users as given; yet, it is reasonable to think that a higher intensity of data extraction by the platform may negatively affect users' experience, thereby influencing their decision to

access the platform. To capture this, we model the utility an individual derives from accessing the platform as comprising two components. The first component, $v > 0$, represents the benefits from consuming the platform's content, and is for simplicity assumed to be identical across users. The second component captures the privacy cost associated with data extraction, and is parametrized by a user-specific parameter $\beta^i > 0$. With a data-extraction intensity x, and an access price $r \geq 0$, the utility of user i can then be written as:

$$v - r - \beta^i x.$$

User i stays on the platform if and only if this expression is positive. Therefore, if we define:

$$\hat{\beta}^i(x) = \frac{v - r}{x}, \tag{4}$$

all users with $\beta^i \leq \hat{\beta}^i(x)$ choose to stay on the platform.

If β has a distribution with CDF F, the mass of users on the platform, as a function of x, is:

$$S(x) = F\left(\frac{v - r}{x}\right)$$

Consistent with the empirical evidence collected by Lin (2022), we assume that the distribution of β in the population follows a Pareto distribution[6], with minimum value $\underline{\beta}$, and shape parameter α:

$$S(x) = F\left(\frac{v - r}{x}\right) = 1 - \left(\frac{\underline{\beta} x}{v - r}\right)^{\alpha} \tag{5}$$

A Pareto distribution for β represents a situation in which the privacy cost for most users is close to the minimum value $\underline{\beta}$, while a fraction of the population incurs substantially higher costs. The shape of the distribution is governed by the parameter $\alpha > 2$. For instance, the proportion of individuals with privacy costs exceeding twice the minimum is $1/2^{\alpha}$, which decreases with α. Therefore, higher values of α correspond to a population that is *less* sensitive to privacy concerns.

Remark 2: Users are described by a two dimensional type (w^i, β^i). We assume that the two characteristics, the willingness to pay for the commercial good and the privacy cost, are independently distributed in the population,

[6]Lin (2022) empirically estimates the intrinsic value of privacy and finds high heterogeneity across consumers, with a distribution skewed to the right.

so that, for any given x, the mass of users $S(x)$ is characterized by a uniform distribution of w on $[0, A]$, and the analysis in section 2.1 continues to apply.[7]

REMARK 3: In defining $S(x)$ we assumed that each user i takes her decision based on the utility for content, v, and the privacy cost $\beta^i x$. We did not consider the effect on this choice of the commercial surplus she might obtain by being on the platform. Indeed, as we discussed in section 2.1, for any given x, users are partitioned in $N(x) = 2^{kx}$ groups and some user in each group has positive surplus $(w^i - p(J_{i,x})) > 0$ where $p(J_{i,x})$ is the price for the target group she belongs to. The *marginal* user in each group always has zero consumption surplus (see Proposition 1 and Figure 1), so *all* users with a β^i satisfying 4 will choose to be on the platform. Notice, however, that *some* users with $\beta^i > \hat{\beta}^i(x)$ may decide to participate, if they were to consider also the positive consumption surplus they could obtain. Our expression for $S(x)$ is thus a conservative estimate of users' participation when the level of information extraction is x. We chose this modelling strategy for two reasons: i) as noticed in Remark 1, in our model, as soon as $n = kx$ is higher than 2 the gain in consumer surplus is eroded by the more precise price discrimination made possible by data extraction, ii) we aim at a minimal model to discuss 'the fundamental economic tension for policymakers: balancing platform profits via high-quality targeting and the value consumers place on their data' (Aridor et al.(2024), p.26). In the last section of the paper, when discussing Welfare, we will consider, as a robustness check, the polar case in which all individuals, independently of their participation choice, enjoy the commercial surplus generated by data extraction.

2.3 Commercial revenue

From our discussion in subsection 2.2 and 2.1, we conclude that, in our model, the commercial revenue generated by a given level of data extraction intensity x is:

$$R(x)S(x)$$

where, using equations 1 and 3, $R(x) \equiv \frac{A}{4}(2 - \frac{2^{kx+1}-1}{2^{2kx}})$, while $S(x)$ is defined as in 5.

[7] One could argue that the two dimensions may be correlated. If for example β^i was influenced by the 'cost of time', this would probably be positively correlated with income and willingness to pay for the good w^i. On the other hand, the empirical research is not conclusive on the magnitude and sign of the correlation; for example, Aridor et al. (2023)'s study of the effects of GDPR regulation finds a negative correlation: consumers deciding to allow for data extraction (thus revealing a low privacy cost β) turned out to have *higher* value to advertisers than those that opt-out.

In practice, as mentioned in the introduction, the platform monetizes the information it extracts from users by selling ad slots to firms. Firms are ready to pay a higher price for more precise information as this enables them to better tailor personalized advertising messages, and possibly to price discriminate among categories of users. We thus interpret $S(x)$ as the mass of users from which the platform extracts data used for targeted digital ads, and $R(x)S(x)$ as 'Digital Advertising Revenue', i.e. the revenue that the commercial sector obtains from accessing these users, and over which it negotiates with the digital platform.

In the next two sections we analyze two alternative business models that the platform may adopt: the *Pure Ad business model*, where revenues derive exclusively from the sale of digital ad space; and the *Freemium business model*, in which users are offered a choice: free access accompanied by targeted ads and data extraction, or a paid subscription at fee r that grants access to the platform's services without data extraction and targeted ads.

The two models have different implications for the effects of taxation and other possible policy interventions.

3 Pure Advertising

Under the *Pure-Ad business model*, the platform's profit as a function of x is:

$$\Pi^{PA}(x) = \gamma[R(x)S(x)] \tag{6}$$

where γ denotes the share of commercial profits appropriated by the platform, which we assume equal to one for simplicity.

Given that, under such model the access price is zero ($r = 0$), equation (5) implies that the mass of users on the platform is:

$$S(x) = 1 - \left(\frac{\beta x}{v}\right)^{\alpha}$$

We fix $\underline{\beta} = v$, so that for all $x < 1$ some users remain on the platform, while at $x = 1$ all users leave. This can be seen as a definition of $x = 1$ as 'maximal level of data extraction intensity'. Then:

$$S(x) = 1 - x^{\alpha} \tag{7}$$

Using Proposition 1, together with equations (3) and (7), we can re-write the platform's profit as:

$$\Pi^{PA}(x) = \frac{A}{4}\left(2 - \frac{2^{kx+1} - 1}{2^{2kx}}\right)(1 - x^{\alpha})$$

This expression depends on three key parameters: A, k, and α.

The parameter A captures the potential profitability of the commercial sector. An increase in A scales the seller's revenues linearly. Recalling that the maximum commecial profit in the absence of data extraction is $\Pi^M = \frac{A}{4}$, it is convenient to normalize $A = 4$, so that $\Pi^M = 1$, the profit of a non-discriminating monopolist, becomes the benchmark against which to compare the profit resulting from different data extraction choices.

Our final expression for the platform's profit under the pure advertising business model thus depends only on two parameters, k and α:

$$\Pi^{PA}(x) = [2 - \frac{2^{kx+1} - 1}{2^{2kx}}](1 - x^{\alpha}), \tag{8}$$

As previously discussed, the parameter k reflects the platform's efficacy in turning users' personal data into commercially valuable information. For any given level of x, a higher k enables more refined segmentation of users.

The parameter α, on the other hand, characterizes the distribution of privacy concern across the user base. Higher values of α correspond to a population which is *less* sensitive to the private cost of information extraction; a reduction in α thus captures an increase in the users' privacy concern.

Figures 2 and 3 illustrate how maximal commercial revenue $R(x)$ and users' participation $S(x)$ vary when k and α take different values.

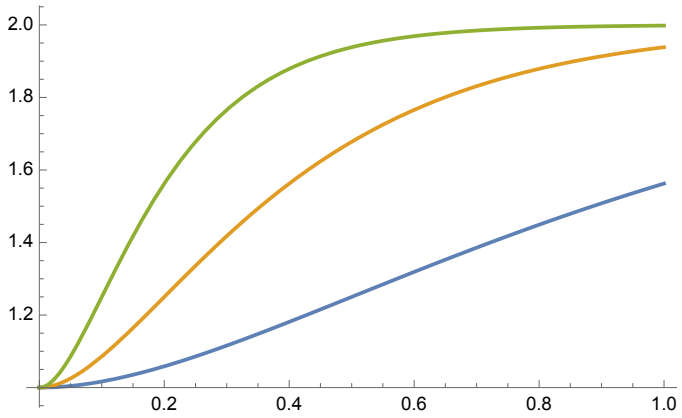

Figure 2: Maximal commercial revenue, $R(x)$, as a function of the level of information extraction $x \in [0, 1]$, for different values of k: $k = 2$ (in blue), $k = 5$ (in orange) and $k = 10$ (in green).

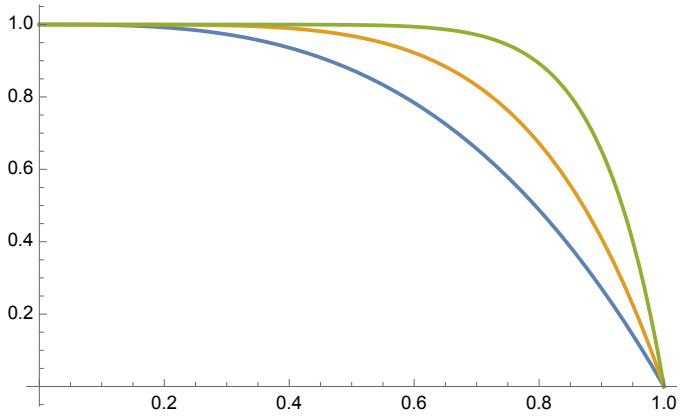

Figure 3: Platform participation, $S(x)$, as a function of the level of information extraction $x \in [0, 1]$, for different values of α: $\alpha = 3$ (in blue), $\alpha = 5$ (in orange) and $\alpha = 10$ (in green).

Both k and α may be shaped by policy interventions. For instance, regulations that require platforms to obtain explicit user consent for data collection and use are likely to reduce both k and α: they may limit the extent to which personal data can be processed and monetized (lowering k), while also increasing users' awareness of privacy risks, thereby amplifying their sensitivity to data extraction (lowering α).

We make the following assumption on the magnitude of k and α:

Assumption 1. $k \geq 2$, $\alpha \geq 3$.

A value of α larger than 2 is necessary for the Pareto distribution to have finite mean and variance; assuming a minimal value of $\alpha = 3$ means that the proportion of the population with a cost of privacy higher than twice the value of being on the platform does not exceed 12.5%. A value of $k = 2$, on the other hand, means that, even with maximal information extraction, users can only be partitioned into four target groups. We do not consider Assumption 1 to be very restrictive.

In the following Lemma we prove existence, interiority and uniqueness of the platform's optimal choice of data extraction intensity under the Pure-Ad business model, as a function of our two parameters. The proof requires some care because the function $\Pi^{PA}(x)$ is not concave in x, and can be found in the Appendix.

Lemma 1. *Under Assumption 1, the problem* $\max_{x\in[0,1]} \Pi^{PA}(x)$ *has a unique solution* $\hat{x}^{PA}(\alpha, k) \in (0, 1)$.

Proof. See Appendix. □

Taxing digital advertising revenue

Suppose now a tax $\tau \in (0, 1)$ is imposed on the platform's revenue $\Pi^{PA}(x)$. Under the Pure-Ad business model, such a tax does not alter the platform choice of data extraction intensity; yet, as suggested by Romer (2021) and Acemoglu and Robinson (2024), it may induce the platform to switch to a *Pure Subscription* business model.

Under Pure Subscription, users pay an access fee $r > 0$ and are neither exposed to personalized ads nor subject to data extraction. In our simple setting, where v is the same for all users, we can normalize $v = 1$. Then, under Pure Subscription: the platform sets the access price equal to the user value, $r = v = 1$; all users access the platform; and profit is $\hat{\Pi}^{PS} = v = 1$.

Denoting the platform equilibrium profit in the Pure-Ad model as $\hat{\Pi}^{PA} \equiv \Pi^{PA}(\hat{x}^{PA})$, a switch to Pure Subscription occurs if and only if the after-tax

profit under the Pure-Ad model falls below the subscription profit, that is: $(1-\tau)\hat{\Pi}^{PA} < \hat{\Pi}^{PS}$. Rearranging, the threshold tax rate that triggers the switch is:

$$\tau > \underline{\tau}^{PA}(\alpha, k) \equiv 1 - \frac{1}{\hat{\Pi}^{PA}} \tag{9}$$

The threshold $\underline{\tau}^{PA}$ varies in the same direction as $\hat{\Pi}^{PA}$ and, by direct calculations we can prove:

Proposition 2. *The minimum tax rate $\underline{\tau}^{PA}$ required to induce a switch from PA to PS is lower when data squeezing is lower (k is lower) and/or privacy concern is higher (α is lower).*

Proof. At the optimal choice $\hat{x}$, using the envelope theorem:

1. $\frac{\partial \hat{\Pi}^{PA}}{\partial k} = 2^{-k\hat{x}+1}(1-2^{-k\hat{x}})\log[2](1-\hat{x}^{\alpha})\hat{x}$, which is positive for any $\hat{x} \in (0,1)$; and
2. $\frac{\partial \hat{\Pi}^{PA}}{\partial \alpha} = -[2-2^{1-k\hat{x}}+2^{-2k\hat{x}}]\hat{x}^{\alpha}\log[\hat{x}]$, which is positive for any $\hat{x} \in (0,1)$.

□

The complementarity between taxation and privacy regulation implied by Proposition 2 may play a crucial role in enhancing the political feasibility of introducing a tax on digital ads revenues.[8] To give a sense of the magnitudes involved, consider that in 2023 Alphabet generated approximately 60 billions from the sale of digital ads space in the US. With 0.4 billions users in the US, and an estimated average willingness to pay for access to YouTube of 15$ per user per year, inducing a switch to a subscription-based model would require a tax rate on advertising profits of approximately 90%.

4 Freemium

In the previous section we considered two polar cases of platform business models: Pure Advertising and Pure Subscription. As discussed in the introduction, some of the largest platforms - possibly in response to public discussion on policy interventions - have recently moved in the direction of offering users a 'pay or consent' option. In this section we examine how the

[8]The impact of privacy protection policies on advertising profits can be substantial. Aridor et al. (2024b) quantify the effect of Apple's App Tracking Transparency (ATT) policy, which requires apps distributed via the Apple Store to obtain user consent before enabling device-level identification. They find that ATT reduced Meta's advertising profits by approximately 30%.

adoption of this more nuanced business model affects the analysis of taxation of Digital Advertising Revenue.

Let the platform offer its users a choice between *two* alternative options (a Freemium menu):

a) a subscription contract, whereby a user pays a subscription fee r for access and is exempt from any data extraction;

b) a free-access contract, under which the user incurs no monetary cost to access the platform, but is subject to data extraction at intensity x.

The marginal user choosing the free-access contract is characterized by a privacy sensitivity β^i that satisfies:

$$v - \beta^i x = v - r$$

which implies

$$\hat{\beta}^i(x) = \frac{r}{x}$$

Accordingly, all users with $\beta^i > \hat{\beta}^i(x)$ prefer the subscription contract. Given our assumption that β follows a Pareto distribution with parameters $(\underline{\beta} = 1, \alpha)$, the fraction of users selecting the free, ad-sponsored contract is given by: $S(x, r) = 1 - (\frac{x}{r})^\alpha$.

As in Sato (2019), one can show that the platform optimally sets $r = v = 1$. The profit of a platform offering two contracts, $(r_1, x_1) = (1, 0)$ and $(r_2, x_2) = (0, x)$, is then:

$$\Pi^F(x) = (1 - S(x)) + S(x)R(x) = x^\alpha + (1 - x^\alpha)\left(2 - \frac{2^{kx+1} - 1}{2^{2kx}}\right) \quad (10)$$

and we can prove:

Lemma 2. *Under Assumption 1, the problem* $\max_{x\in[0,1]} \Pi^F(x)$ *has a unique solution* $\hat{x}^F(\alpha, k) \in (0, 1)$.

Proof. Under Freemium, the first order condition for the maximization of profit at x is:

$$\Pi^{F'}(x) = R'(x)S(x) + S'(x)R(x) - S'(x) \equiv N(x) - S'(x) = 0.$$

The function $S'(x) = -\alpha x^{\alpha-1}$ is a concave decreasing function with $S'(0) = 0$ and $S'(1) = -\alpha$.

From Lemma 1, we know that the function $N(x) = \Pi'^{PA}(x)$ crosses the x axis from above only once, at $x = \hat{x}^{PA} > \underline{x}_k$, and it is decreasing thereafter, with $N(1) = -\alpha R(1) < S'(1)$, because $R(1) = [2 - \frac{2^{k+1}-1}{2^{2k}}] > 1$ for $k > 0$.

So there exists a unique $\hat{x}^F > \hat{x}^{PA}$ such that $N(\hat{x}^F) = S'(\hat{x}^F)$, i.e. $\Pi^{F\prime}(\hat{x}^F) = 0$. Also $N'(\hat{x}^F) = \Pi''^{PA}(\hat{x}^F) < S''(\hat{x}^F) < 0$, so the second order condition for a maximum is also verified. □

The introduction of a Freemium menu changes the platform's incentives in a fundamental way. In particular, the availability of a subscription option softens the trade-off between data extraction and user participation. This leads to a strictly *higher* optimal level of data extraction compared to the Pure Advertising model:

Proposition 3. *For all values of the parameters,* $\hat{x}^F(\alpha, k) > \hat{x}^{PA}(\alpha, k)$.

Proof. The first order condition for an interior solution in the PA model is:

$$R(\hat{x}^{PA})S'(\hat{x}^{PA}) + S(\hat{x}^{PA})R'(\hat{x}^{PA}) = 0 \tag{11}$$

while, from 10, in the Freemium model, $\hat{x}^F$ must satisfy:

$$R(\hat{x}^F)S'(\hat{x}^F) + S(\hat{x}^F)R'(\hat{x}^F) - S'(\hat{x}^F) = 0 \tag{12}$$

Comparing equations 11 and 12, we see that, at $\hat{x}^{PA}$, 12 is not satisfied. At $\hat{x}^{PA}$, the first derivative of Π^F is positive:

$$R(\hat{x}^{PA})S'(\hat{x}^{PA}) + S(\hat{x}^{PA})R'(\hat{x}^{PA}) - S'(\hat{x}^{PA}) = 0 - S'(\hat{x}^{PA}) > 0$$

so, if an internal solution $\hat{x}^F$ exists, it must be that $\hat{x}^F > \hat{x}^{PA}$. □

Taxing digital advertising revenue under Freemium

Suppose now a tax is imposed on digital advertising revenue. Under this policy, the platform's profit becomes:

$$\Pi^F_\tau(x) = (1 - S(x)) + (1 - \tau)R(x)S(x)$$

In the next Proposition, we prove that the Freemium model remains optimal, with a unique internal solution $\hat{x}^F_\tau$, when the tax rate is below a threshold $\underline{\tau}^F(k)$:

$$\underline{\tau}^F(k) \equiv 1 - \frac{1}{R(1)} \tag{13}$$

where $R(1) = [2 - \frac{2^{k+1}-1}{2^{2k}}]$; the threshold tax rate $\underline{\tau}^F(k)$ thus increases with the *data squeezing* parameter k. For tax rates above $\underline{\tau}^F(k)$, the platform switches to Pure Subscription:

Lemma 3. *Under Assumption 1, for all $\tau < \underline{\tau}^F(k)$, the problem $\max_{x\in[0,1]} \Pi_\tau^F(x)$ has a unique solution $\hat{x}_\tau^F(\alpha, k) \in (0,1)$. If $\tau \geq \underline{\tau}^F(k)$, the platform switches to Pure Subscription.*

Proof. Under Freemium, when the tax rate on Advertising revenue is $\tau \in [0,1]$, the first order condition for the maximization of profit at x is:

$$\begin{aligned} \Pi_\tau^{F\prime}(x) &= (1-\tau)[R'(x)S(x) + S'(x)R(x)] - S'(x) \\ &\equiv (1-\tau)N(x) - S'(x) = 0. \end{aligned} \tag{14}$$

By the same argument used in the proof of Lemma 2, if $(1-\tau)N(1) < S'(1)$, that is $-(1-\tau)\alpha R(1) < -\alpha$ (i.e. $\tau < \underline{\tau}^F$), there is a unique $\hat{x}_\tau^F > \hat{x}^F$ such that $(1-\tau)N(\hat{x}_\tau^F) - S'(\hat{x}_\tau^F) = 0$ and 14 is satisfied. Also as before, the second order conditions for a maximum are verified.

If, on the other hand, $\tau \geq \underline{\tau}^F$, i.e. $(1-\tau)N(1) = -(1-\tau)\alpha R(1) \geq S'(1) = -\alpha$, then the optimal choice is to set $\hat{x}_\tau^F = 1$ (implying $S(\hat{x}_\tau^F) = 0$), or, equivalently, switch to Pure Subscription. □

Lemma 2 and Lemma 3 are illustrated in Figure 4, for the case $k = 2$, $\alpha = 3$, $\tau = 0.2$.

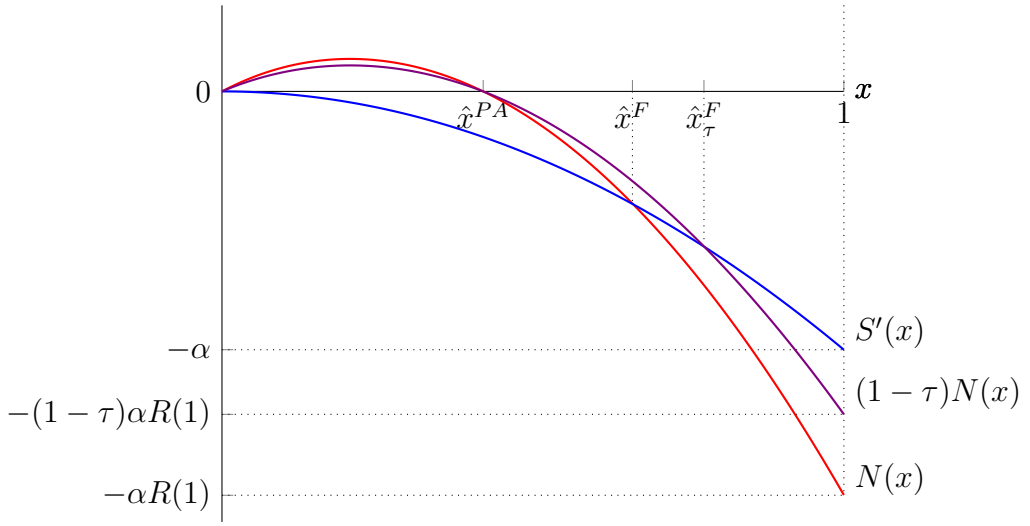

Figure 4: $N(x)$ (in red), $S'(x)$ (in blue) and $(1-\tau)N(x)$ (in purple), as functions of x. Parameters are: $k = 2$, $\alpha = 3$, $\tau = 0.2$, leading to choices $\hat{x}^{PA} = 0.42$, $\hat{x}^F = 0.66$ and $\hat{x}_\tau^F = 0.78$.

How does the adoption of Freemium alters the effectiveness of a tax on Digital Ads Revenue?

The first observation is straightforward but important. For all parameter values, platform profits under the Freemium model are strictly higher than in the Pure Advertising case; as a result, the tax rate required to induce a switch to Pure Subscription is also *higher*:

Proposition 4. *For all k and α, $\underline{\tau}^F(k) > \underline{\tau}^{PA}(k, \alpha)$.*

Proof. This follows immediately from the definitions of $\underline{\tau}^F$ and $\underline{\tau}^{PA}$, in equations 9 and 13, and the fact that $\hat{\Pi}^{PA} = R(\hat{x}^{PA})S(\hat{x}^{PA}) < R(1)$. □

Second, unlike in the Pure Advertising case, the level of data extraction now reacts to τ. For levels of taxation which are *not* sufficiently high as to induce a switch to Pure Subscription, an increase in τ actually leads to an *increase* in $\hat{x}^F_\tau$. Intuitively, this follows from the fact that, under Freemium, any increase in x has a *negative marginal cost* for the platform, because it adds users to the subscription base. At the optimal choice $\hat{x}^F_\tau$, the level of information extraction is thus pushed *beyond* the level at which the marginal advertisement revenue is zero.

Proposition 5. *If $\tau < \underline{\tau}^F$, the level of data extraction $\hat{x}^F_\tau$ is increasing in τ.*

Proof. By the implicit function theorem (using the SOC for maximization):

$$sign\frac{\partial \hat{x}^F_\tau}{\partial \tau} = sign\frac{\partial(\frac{\partial \Pi^F_\tau}{\partial x})}{\partial \tau}$$

Now, using the envelope theorem:

$$\frac{\partial \Pi^F_\tau}{\partial x} = (1-\tau)[R(\hat{x}^F_\tau)S'(\hat{x}^F_\tau) + S(\hat{x}^F_\tau)R'(\hat{x}^F_\tau)] - S'(\hat{x}^F_\tau) = 0 \qquad (15)$$

so

$$\frac{\partial(\frac{\partial \Pi^F_\tau}{\partial x})}{\partial \tau} = -[R(\hat{x}^F_\tau)S'(\hat{x}^F_\tau) + S(\hat{x}^F_\tau)R'(\hat{x}^F_\tau)] > 0$$

where the inequality follows from 15 and the fact that $S'(\hat{x}^F_\tau) < 0$. □

Third, the privacy concern parameter α *does not* influence the platform's choice between Freemium and Pure Subscription: as equation (13) makes it clear, the threshold tax rate $\underline{\tau}^F(k)$ is independent of α. This can be explained by noticing that, both the after-tax marginal revenue from the ad-sponsored contract and the marginal revenue from subscription depend on α

(through the change in the number of subscribers). When τ approaches the threshold $\underline{\tau}^F$, the level of $\hat{x}_\tau^F$ approaches one, but when $\tau = \underline{\tau}^F$, both the after-tax marginal revenue from advertising and the marginal revenue from subscription are equal to $-\alpha$, nobody chooses the ad-free contract, and the platform switches to Pure Subscription (see the proof of Lemma 3 and Figure 4 for an illustration).

However, this does not imply that privacy concern is completely irrelevant under the Freemium model. Indeed, for levels of taxation below $\underline{\tau}^F$, a policy that increases users' awareness of the risks associated with information extraction (i.e. a policy which *reduces* α) can effectively induce the platform to choose a lower x:

Proposition 6. *If $\tau < \underline{\tau}^F$, the level of data extraction $\hat{x}_\tau^F$ decreases when privacy concern increases (i.e. α is lower).*

Proof. We need to show that: $\frac{\partial \hat{x}\tau^F}{\partial \alpha} > 0$. By the implicit function theorem (using the SOC for maximization), for any parameter ν influencing the profit:

$$sign\frac{\partial \hat{x}_\tau^F}{\partial \nu} = sign\frac{\partial \Pi_\tau^F}{\partial x \partial \nu}$$

Now, rewriting 15 as:

$$\frac{\partial \Pi_\tau^F}{\partial x} = [(1-\tau)R(\hat{x}_\tau^F) - 1]S'(\hat{x}_\tau^F) + (1-\tau)S(\hat{x}_\tau^F)R'(\hat{x}_\tau^F)$$

we see that:

$$\frac{\partial \Pi_\tau^F}{\partial x \partial \alpha} = [(1-\tau)R(\hat{x}_\tau^F) - 1]\frac{\partial S'}{\partial \alpha} + (1-\tau)R'(\hat{x}_\tau^F)\frac{\partial S}{\partial \alpha}$$

which is positive because

$$\frac{\partial S}{\partial \alpha} = -x^\alpha \log x,$$

and

$$\frac{\partial S'}{\partial \alpha} = -\alpha x^{\alpha-1} \log x$$

are both positive for all $x \in (0,1)$; the same holds for $R(x)$ and $R'(x)$; while, using 15, at $\hat{x}_\tau^F$:

$$[(1-\tau)R(\hat{x}_\tau^F) - 1] = -\frac{S(\hat{x}_\tau^F))R'(\hat{x}_\tau^F))}{S'(\hat{x}_\tau^F))} > 0.$$

□

5 Welfare

Our analysis in the previous section shows that the platform's adoption of a Freemium model introduces additional challenges for the policy proposal advocating a tax on digital advertising revenue aimed at inducing a switch to Pure Subscription. Specifically: (i) the tax threshold required to induce such a switch is higher than in the Pure-Ad benchmark; (ii) prior to reaching that threshold, a higher tax has the perverse effect of inducing higher levels of data extraction; and (iii) policies aiming at increasing users' awareness of privacy costs do not help in lowering the necessary threshold.

Given the widespread diffusion of Freemium, these concerns are far from theoretical curiosities. Should we then conclude that the Romer and Acemoglu-Johnson proposal is either politically unfeasible or undesirable from a welfare standpoint in a context where the platform operates under a Freemium business model?

While any definitive conclusion must ultimately rely on empirical evidence and political judgment, the theoretical framework we have proposed can nonetheless illuminate key aspects of the issue by capturing some of the salient features of the real-world environment in which digital platforms operate. In what follows, we illustrate this point through a simple numerical example, using plausible values for the two central parameters in our model: the intensity of data extraction (k) and the degree of privacy concern among users (α).

We begin by expressing the social welfare generated on platform as the sum of users' surplus and platform's profit:

$$W(x) = CS(x) + \int_1^{\hat{\beta}(x)} (1 - \beta x) d\beta + \Pi^F(x) \tag{16}$$

In this expression, the first two terms constitute total users' surplus, $V(x)$: the sum of consumer surplus, as defined in equation (2), and the surplus derived from access to the platform, net of privacy costs. Recalling that $S(x) = 1 - x^\alpha$, and noting that, under a Pareto distribution, the aggregate privacy cost is:

$$x \int_1^{\hat{\beta}(x)} \beta d\beta = \frac{\alpha}{(\alpha - 1)}(x - x^\alpha)$$

we have:

$$V(x) = (1 - x^\alpha)\left(\frac{2^{kx+2} - 3}{2^{2kx+1}}\right) + (1 - x^\alpha) - \frac{\alpha}{(\alpha - 1)}(x - x^\alpha)$$

while the platform's profit under Freemium (as in equation 10) is:

$$\Pi^F(x) = (1 - S(x)) + R(x)S(x) = x^\alpha + (2 - \frac{2^{kx+1} - 1}{2^{kx}})(1 - x^\alpha)$$

Substituting these expressions into (16), total welfare simplifies to:

$$W(x) = 1 + (1 - x^\alpha)\left(\frac{2^{2kx+2} - 1}{2^{2kx+1}}\right) - \frac{\alpha}{(\alpha - 1)}(x - x^\alpha) \tag{17}$$

The first term in (17) captures the *platform access value*, which is always fully realized under the Freemium model, since all users access the platform. The second term represents *total commercial surplus*, the product of the share of users choosing the ad-sponsored contract times the sum of consumer surplus and platform profit. The third term is the *aggregate privacy cost*.

For given values of k and α, the platform's equilibrium choice of data extraction intensity, $\hat{x}^F$, determines its profit, $\Pi^F(\hat{x}^F)$, and total users' surplus, $V(\hat{x}^F)$, hence overall welfare, $W(\hat{x}^F)$.

In Figure 5, we plot the total welfare function and its components, for the parameter values $\alpha = 10$ and $k = 10$. Specifically, $W(x)$ is represented in red, while $V(x)$ and $\Pi^F(x)$ are represented in blue and purple, respectively.

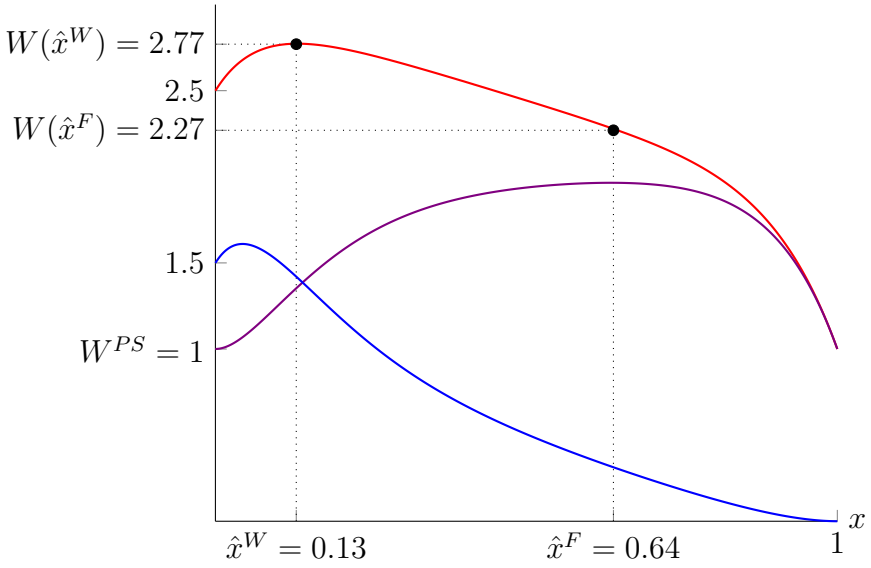

Figure 5: Welfare W (in red), users' surplus V (in blue) and profit Π^F (in purple), as functions of x. Parameters are: $k = 10$, $\alpha = 10$, leading to choices $\hat{x}^F = 0.64$ and $\hat{x}^W = 0.13$.

A value of $\alpha = 10$ captures a relatively low degree of privacy concern in the population: in equilibrium, the platform chooses a level of data extraction $\hat{x}^F = 0.64$, and yet nearly all users opt for the ad-sponsored contract, $S(\hat{x}^F) = 0.99$. Under these conditions, the firm's profit, fully appropriated by the platform, is equal to $\Pi^F(\hat{x}^F) = 1.97$, while user surplus amounts to $V(\hat{x}^F) = 0.30$, yielding a total welfare of $W(\hat{x}^F) = 2.27$. This parametrization corresponds to a situation characterized by low privacy concerns, a high intensity of data extraction and exploitation, and very large platform profits - indeed, $\Pi^F(\hat{x}^F) = 1.97$ is almost twice the profit that would accrue to a standard, non-discriminating monopolist. We take this configuration to be a reasonable representation of current conditions and use it as our baseline market equilibrium.

As a benchmark, note that a benevolent planner, capable of *imposing* the level of information extraction that maximizes equation (17), would select a substantially lower data extraction level, $x^W = 0.13$. This choice would result in a reduction in profits of more than 30%, lowering them to $\Pi^F(x^W) = 1.35$, while users' surplus would increase almost fivefold, reaching $V(x^W) = 1.42$. Consequently, total welfare would rise to $W(x^W) = 2.77$, a 22% improvement relative to the initial scenario.

Remark 4: In section 2.2, we observed that our definition of $S(x)$, the share of the population choosing the free, ad-sponsored contract, is a conservative one. Given that in our model commercial consumer surplus is generated *only* through the interaction between the seller and these users, this may lead to an underestimation of the welfare created by this interaction. As a robustness check, let us momentarily assume that *all* users, including those choosing subscription, enjoy the surplus generated by information extraction and targeted ads:

$$V(x) = \left(\frac{2^{kx+2} - 3}{2^{2kx+1}}\right) + (1 - x^\alpha) - \frac{\alpha}{(\alpha - 1)}(x - x^\alpha)$$

In our numerical example, with very low privacy concern ($\alpha = 10$), at equilibrium almost all users choose the ad-sponsored contract anyway; hence, nothing of substance changes. For lower values of α, that is for higher levels of privacy concern, a larger share of users chooses subscription. However, for these same users, the privacy cost is also higher. In numerical simulations it is never the case that the planner chooses a higher x^W than in the baseline model.

Consider now the case in which, possibly as a reaction to a high enough tax on Digital Advertising Revenue, the platform chooses to offer only a pure subscription contract. In our model, this implies the complete elimination of

commercial interactions via the platform, so that only the first component of welfare (the value derived from platform access) remains, resulting in a drastic reduction of welfare to $W^{PS} = 1$.

As discussed in the Introduction, the rationale for a tax on digital advertising revenues, designed to induce a switch towards Pure Subscription, does not hinge on its effects on commercial surplus. Instead, it is grounded in the accumulating evidence of a negative *social* externality inherent in business models where profitability is driven by the maximization of user engagement. Taking this externality into account, a benevolent planner should weigh the gain in total commercial surplus (the second term in (17)) not only against the privacy cost (the third term), but also against the negative external effect generated by the information extraction process, a component that is not explicitly modeled in (17).[9] Under our parametrization, a shift towards a Pure Subscription model would be justified on welfare grounds if the monetary value of the negative externality were on the order of $W(\hat{x}^F) - W^{PS} = 2.27 - 1 = 1.27$, which corresponds to roughly 56% of the total market-equilibrium profits generated through digital advertising.

Recent estimates from Statista indicate that, in 2025, the five largest digital platforms generated a combined total of 715 bilion dollars in profits from digital advertising.[10] With nearly 5 billion social platforms users worldwide, this implies that an externality cost as low as 80 dollars per user per year would suffice to justify the proposed policy intervention. Given the growing empirical evidence documenting the harmful effects associated with engagement-driven business models, it is highly plausible that the actual value of the externality exceeds this figure. Our numerical exercise therefore suggests that the policy proposal advanced by Romer and Acemoglu-Johnson does not hinge on implausibly large valuations of these negative effects to be considered a serious and credible option.

Alternative Policies

Our framework also lends itself to comparing the taxation of digital advertising revenue with alternative policy instruments aimed at reducing the extraction of users' data. A relevant example is the European Union's General Data Protection Regulation (GDPR), introduced in 2018, which mandates the informed consent of users as a prerequisite for data collection.[11] In the

[9] See e.g. Acemoglu, Ozdaglar and Siderius (2024) for an interesting first attempt to formalize one dimension of this negative externality, namely the potential for engagement-driven algorithms to induce belief polarization.

[10] https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/

[11] See Goldberg et al., 2024.

context of our model, such regulatory interventions can be interpreted as a reduction in one or both of the key parameters governing data extraction.

Consider, for instance, reducing α from 10 to 3 - a change that induces 30% of the population to select the ad-free subscription option - and simultaneously limiting the platform's data extraction capacity from $k = 10$ to $k = 2$. This combination represents a substantial tightening of users' privacy protection. Nevertheless, the equilibrium level of data extraction *increases* under this policy (see Figures 4 and 5). This somewhat counterintuitive result mirrors the mechanism observed in the case of taxation with tax rates below the switching threshold: the flexibility embedded in the Freemium contract enables the platform to strategically adapt in response to privacy regulations. As a result, policies that merely restrict the parameters of data extraction, while allowing personalized advertising to persist, may fail to achieve their intended outcomes.

In line with our earlier analysis of taxation, a more effective intervention may consist in a *complete ban* on the use of personalized data for advertising purposes. In our model, this corresponds to setting $k = 0$, which leads the platform to choose $x = 0$.[12] Importantly, such a policy does not preclude all forms of digital advertising. It allows non-personalized (e.g., contextual) ads, which - in the context of our example - would enable the platform to generate a total commercial surplus (the sum of consumer surplus and the platform's profit) of $CS(0) + \Pi(0) = 1.5$, and a total welfare of $W(0) = 2.5$, close to the maximum of $W(x)$ (refer again to Figure 5). The main issue with this drastic form of regulation lies in its implementability, as evidenced by the difficulty of drawing a clear boundary between 'behavioral' and 'contextual' advertising.[13]Hence, also from the perspective of relative simplicity in implementation, the proposal to tax digital advertising revenue seems particularly worthy of policy makers' attention.

6 Conclusion

In this paper, we proposed a model of how a platform's choice of data extraction intensity affects its profit, as well as users' decisions and well-being. The model is stylized, but the two parameters we identify - data squeezing (k) and privacy concern (α) - succinctly capture two crucial forces shaping

[12]A recent proposal in the US Congress, the Banning Surveillance Advertising Act of 2023, moves in this direction. See https://www.congress.gov/bill/118th-congress/house-bill/5534.

[13]See Bondi et al. (2025). Also, for a policy discussion: https://www.techpolicy.press/is-so-called-contextual-advertising-the-cure-to-surveillance-based-behavioral-advertising/

the intensity of data extraction: its usefulness for market segmentation and pricing, and the potential reduction in user participation due to increasing awareness of privacy costs.

Our findings indicate that designing an effective digital advertising tax to induce a switch in business model is more challenging when platforms offer Freemium contracts to their users. Nevertheless, when using the model to estimate orders of magnitude for the various components of social welfare, we find that the policy proposed by Romer, and Acemoglu and Johnson does not require implausibly high values of the social negative externality generated by the current business model to be justified. This policy also benefits from relative simplicity of implementation. Of course, political feasibility remains an open question; but analytical models can provide useful insights in to the ongoing debate.

References

Acemoglu, D., D. Huttenlocher, A. Ozdaglar and J. Siderius, 2024. Online business models, digital ads, and user welfare. NBER WP 33017, October 2024.

Acemoglu, D. and Johnson, S., 2024. The urgent need to tax digital advertising. https://shapingwork.mit.edu/research/the-urgent-need-to-tax-digital-advertising/

Acemoglu D., A. Ozdaglar and J. Siderius, 2024. A model of online misinformation, *The Review of Economic Studies*, 91, pp.3117- 3150.

Anderson, S.P. and Coate, S., 2005. Market provision of broadcasting: A welfare analysis. *The Review of Economic Studies*, 72(4), pp.947-972.

Anderson, S.P. and Gabszewicz, J.J., 2006. The media and advertising: a tale of two-sided markets. *Handbook of the Economics of Art and Culture*, 1, pp.567-614.

Anderson, S.P. and Jullien, B., 2015. The advertising-financed business model in two-sided media markets. In *Handbook of media economics* (Vol. 1, pp. 41-90). North-Holland.

Aridor, G., Jimnez-Durn, R., Levy, R.E. and Song, L., 2024a. The Economics of Social Media. *Journal of Economic Literature* forthcoming.

Aridor, G., Che, Y.K and T. Salz, 2023. The effects of privacy regulation on the data industry: empirical evidence from GDPR. *The RAND journal of economics*, 54(4), pp.695-730.

Aridor, G., Che, Y.K., Hollenbeck, B., Kaiser, M. and McCarthy, D., 2024b. Evaluating The Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apples App Tracking Transparency. Mimeo

Armstrong, M., 2006. Competition in two-sided markets. *The RAND journal of economics*, 37(3), pp.668-691.

Belleflamme, P. and Peitz, M., 2021. The economics of platforms. Cambridge University Press.

Bergemann, D. and Bonatti, A., 2019. Markets for information: An introduction. *Annual Review of Economics*, 11(1), pp.85-107.

Bergemann, D. and Bonatti, A., 2024. Data, competition, and digital platforms. *American Economic Review*, 114(8), pp.2553-2595.

Bloch, F. and Demange, G., 2018. Taxation and privacy protection on Internet platforms. *Journal of Public Economic Theory*, 20(1), pp.52-66.

Bondi, T. , O.Rafieian and Y. Yao, 2025. Privacy and Polarization: An Inference-Based Framework. *Management Science*, forthcoming.

Braghieri, L., R. Levy and A. Makarin, 2022. Social media and mental health. *American Economic Review*, 112(11), pp.3660-3693.

Bursztyn L., B.R. Handel and J.H. Oh, 2023. When product markets become collective traps: The case of social media. NBER WP 311771, October.

Galperti, S. and Perego, J., 2023. Privacy and the Value of Data. *AEA Papers and Proceedings* (Vol. 113, pp. 197-203).

Goldberg, S.G., Johnson, G.A. and Shriver, S.K., 2024. Regulating privacy online: An economic evaluation of the GDPR. *American Economic Journal: Economic Policy,* 16(1), pp.325-358.

Guess A.M. et al. 2023. How do social media feed algoriths affect attitudes and behavior in an election campaign. *Science*, 381, pp. 398-404.

Lin, T. 2022. Valuing Intrinsic and Instrumental Preferences for Privacy. *Marketing Science* 41(4), pp.663-681.

McLoughlin K.L. et al. 2024. Misinformation exploits outrage to spread online. *Science*, 386, pp. 991-996.

Rochet, J.C. and Tirole, J., 2006. Two-sided markets: a progress report. *The RAND journal of economics*, 37(3), pp.645-667.

Romer, P., 2021. Taxing digital advertising. https://adtax.paulromer.net

Sato, S., 2019. Freemium as optimal menu pricing. *International Journal of Industrial Organization*, 63, pp.480-510.

Wernerfelt, N., Tuchman, A., Shapiro, B.T. and Moakler, R., 2024. Estimating the value of offsite tracking data to advertisers: Evidence from meta. *Marketing Science*.

Zhuravskaya, E., Petrova, M. and Enikolopov, R., 2020. Political effects of the internet and social media. *Annual review of economics*, 12(1), pp.415-438.

Appendix

Proof of Lemma 1. With $\gamma = 1$, a necessary condition for an interior solution $\hat{x} \in (0,1)$ to the maximization of the platform's profit under the Pure Ad model, equation 6, is:

$$\frac{R'(\hat{x})}{R(\hat{x})} = -\frac{S'(\hat{x})}{S(\hat{x})} \tag{18}$$

Now:

$$S(x) = 1 - x^{\alpha}$$

and

$$S'(x) = -\alpha x^{\alpha-1}$$

so that RHS of 18 is:

$$h(x) \equiv -\frac{S'(x)}{S(x)} = \frac{\alpha x^{\alpha-1}}{1-x^{\alpha}}$$

which, for $\alpha > 1$ is increasing and convex function of x, with $h(0) = 0\lim_{x\to 1} h(x) = \infty$, while:

$$R(x) = [2 - \frac{2^{kx+1}-1}{2^{2kx}}]$$

and

$$R'(x) = 2\log[2]k\frac{2^{kx}-1}{2^{2kx}}$$

so that the LHS of 18 is:

$$f(x) \equiv \frac{R'(x)}{R(x)} = 2\log[2]k\frac{2^{kx}-1}{2(2^{2kx}-2^{kx})+1}$$

with $f(0) = 0$ and $f(1) = 2\log[2]k\frac{2^{k}-1}{2(2^{2k}-2^{k})+1} > 0$ for all $k > 0$.

To show that there exists a unique $\hat{x} \in (0,1)$ such that $f(\hat{x}) = h(\hat{x})$, we need to study the shape of the function $f(\cdot)$ on the domain $[0,1]$.

Notice first thats the first derivative of $f(x)$ ix:

$$f'(x) = \frac{R''(x)R(x) - (R'(x))^2}{R^2}$$

which is negative if and only if $R''(x)R(x) < (R'(x))^2$. Substituting the expressions for R, R' and R'', one can see that, for $k > 0$ and $x > 0$, this is true if and only if $2^{kx} > 1.7$, that is if:

$$kx\log[2] < \log[1.7]$$

i.e:

$$x > 0.77/k \equiv \underline{x}_k$$

Now, at the value $x = \underline{x}_k$,

$$f(\underline{x}_k) = 0.29k,$$

while:

$$h(\underline{x}_k) = \frac{\alpha(0.77/k)^\alpha}{1-(0.77/k)^\alpha}\frac{k}{0.77}$$

so that $f(\underline{x}_k) > h(\underline{x}_k)$ if and only if:

$$k^\alpha > \frac{\alpha + 0.22}{0.22}(0.77)^\alpha$$

which is always true for $k \geq 2$ and $\alpha \geq 3$.

We can conclude that, for $x \geq \underline{x}_k$, there exists a unique $\hat{x}^{PA} \in (0,1)$ such that $g(\hat{x}^{PA}) = h(\hat{x}^{PA})$, and 18 is satisfied.

To see that there is no other solution to the first order condition 18 for $x \in (0, \underline{x}_k]$, notice that:

$$f'(x) = \frac{R''(x)}{R(x)} - (\frac{R'(x)}{R(x)})^2 = \frac{R''(x)}{R(x)} - (f(x))^2$$

so:

$$f''(x) = \frac{R'''(x)R(x) - R''(x)R'(x)}{(R(x))^2} - 2f(x)f'(x)$$

which is always negative for $x \in (0, \underline{x}_k]$. Indeed, for all x, $R(x) > 0$ and $R'(x) \geq 0$, so $f(x) \geq 0$; also $R'''(x) \leq 0$, while , for $x \in (0, \underline{x}_k)$, $R''(x) > 0$.

So, for $x \in [0, \underline{x}_k]$, $f(\cdot)$ is an increasing concave function starting at $f(0) = 0$ with $f'(0) = R''(0) = 2[k \log[2]]^2 > 0$, while $h(\cdot)$ is an increasing convex function with $h(0) = 0$ and $h'(0) = 0$. As we showed above, for $k \geq 2$ and $\alpha \geq 3$, $f(\underline{x}_k) > h(\underline{x}_k)$, so, for $x \in [0, \underline{x}_k]$, the two functions only cross at $x = 0$.

Let us now check the second order conditions. Using a simplified notation we can write:

$$\Pi^{PA'} = R'S + S'R$$

so

$$\Pi^{PA''} = R''S + S'R' + S''R + S'R' = R''S + 2S'R' + S''R$$

Given that for all $x > 0$, $S''(x) < 0$ and $R(x) > 0$, a sufficient conditions for $\Pi^{PA''}(x) < 0$ is that:

$$R''S + 2S'R' < 0$$

At $\hat{x}^{PA}$, the first order condition $R'S = -S'R$ is verified, so that this expression can be written as:

$$-S'[\frac{R''R - 2(R')^2}{R'}] < 0$$

which is true for all $x \geq \underline{x}_k$. □