Targeted Advertising in Social Media Platforms with Heterogeneous Participants

Claudio Panico, Bocconi University, Milan

claudio.panico@unibocconi.it

Carmelo Cennamo, CBS, Copenaghen

cce.si@cbs.dk

Pietro Morino, Bocconi University, Milan

pietro.morino@phd.unibocconi.it

Abstract. Social media platforms face a strategic tradeoff when addressing the needs of different participants: while content providers would like to reach a large number of users, users are not always interested in the content and may suffer a disutility from advertising. How to create value for and thus attract participants from both sides? Using a formal model, we examine a platform that offers targeted advertising to increase the "click-through rate" while minimizing users' disutility from advertising. Providers offer content of heterogenous quality (vertical differentiation) that differs also for the relatedness to users' heterogeneous preferences (horizontal differentiation). Contrary to common wisdom, platform configurations can have an unbalanced number of participants on the two sides, pointing to unexplored mechanisms in the functioning of digital platforms.

Keywords: Social media platforms; targeted advertising; advertising intensity; heterogeneous participants; horizontal differentiation.

INTRODUCTION

Digital platforms, particularly social media platforms, are increasingly becoming "attention brokers" (Wu 2019), catalyzing the significant amount of time that individuals spend online. According to Statista (2022), the average user worldwide spends 145 minutes on social networks every day. Given that users' attention is a limited and scarce resource (Boik, Greenstein & Prince, 2016), there is intense competition from content providers and advertisers vying for a share of this attention. A market for user attention can emerge around a digital platform that engages users on one side also informing them and stimulating their interest toward the content of providers on the other side (Arrate Galán et al. 2019; Prat & Valletti, 2021). Digital platforms that adopt this business model usually offer a free service to the end users and then charge the providers in their ecosystems that want to advertise their content. For instance, Facebook offers a few services for free while generating about 98% of its revenue ($84.17 billion) through online advertising (Facebook Inc. Annual Report, 2020). Alphabet also operates in a similar manner, with about 80% of its revenue ($146.92 billion) generated by online advertising (Alphabet Inc. Annual Report, 2020). Yet, a platform that aims to leverage (direct and indirect) network effects to attract participants on both sides must also consider that users might dislike too much advertising, especially if the ads content is irrelevant (i.e., promoting products or services that do not match users' preferences) or of low quality.

Target advertising supported by deep knowledge of users' preferences can solve in part the problem of irrelevant content (Marotta et al., 2021; Zhang & Katona, 2012). Yet platforms are hardly fully informed about the individual characteristics of the participants populating their ecosystem. For example, users may differ in their tolerance (or disutility) for being the target of advertising, and providers may differ in the quality of the content they can offer. Assuming that the participants on the two sides have private information about these characteristics, the relatedness between users and providers becomes a key factor for understanding a platform's ability to attract participants and to generate value. Mainstream theory on multi-sided platform markets (Armstrong, 2006; Caillaud &

Jullien, 2003; Parker & Van Alstyne, 2005; Rochet & Tirole, 2003, 2006) has it that platforms efficiently coordinate the multiple sides of the market. Through the proper pricing strategy, a platform internalizes the externalities, positive and negative, generated by the different groups of participants. However, this theory is largely based on transaction platforms, whereby there is symmetry of incentives for both groups on the distinct sides to interact to transact and exchange (economic) value. This leads to the instantiation of the positive indirect network effects across the sides of the market that characterize much of the value-enhancing results of these platform-based market models. However, search engines and social media platforms share features that depart from this canonical model. What gets exchanged is information across the sides; information can be more or less relevant for different users and can generate asymmetric network effects, even of opposite signs (Cennamo 2021).

This peculiarity of social media platforms has led some scholars to conclude that they have the incentives to "exploit" their user base and favor the side of the market generating profits, i.e., content providers. Specifically, users are exposed to a greater level of advertising than it would otherwise be optimal (if the platform were to internalize the cost of user exposure to advertisement) (e.g., Anderson & Gabszewicz, 2006). However, given that in some case advertisement can generate costs that are greater than the benefit of becoming informed, it can lead to platform usage frictions, decreased user engagement, lower attention, and ultimately lower benefits also for the providers. Thus, a platform must still care for both sides of the market, and indeed digital platforms have proved successful in managing these network externalities by means of coordination strategies (e.g., Hagiu, 2006; Panico & Cennamo, 2022). The setting we develop closely resembles existing social media platforms (e.g., Facebook), and might help to shed light on the ongoing debate concerning the extent to which these platforms profit to the expense of their users by ignoring negative ad externalities.

We model a digital platform that enables the interaction among a population of users, functioning like a social network, and between users and content providers, informing users about providers' content. The platform makes profits by charging providers based on the mass of users that they reach through the platform. On one side, providers choose the advertising intensity on the platform based on the pricing and on the type of content they can offer. On the other side, users may "click" on a specific content (e.g., to read it or to buy it) if they find it interesting, experiencing a utility; if instead they are reached by content deemed uninteresting, they are annoyed by the advertising and experience a disutility. The two sides connected by the platform enjoy direct as well as indirect network effects, with intensity and sign, we posit, that depend on the interrelation of users' and providers' characteristics. In line with recent developments in the literature (e.g., Panico & Cennamo, 2022, ADD MORE), we aim to offer a more nuanced view of the role of indirect network effects, going beyond the much-studied effect of quantity and size. Our approach and our findings depart from the common wisdom that more participants on one side attract more participants on the other side, offering new insights into the functioning of digital platforms.

ADD BRIEF OVERVIEW OF THE MAIN RESULTS

THEORETICAL BACKGROUND

The concept of "attention economy" (Wu, 2016, 2019) is increasingly used inside and outside[1] of academia to characterize some of the existing digital markets (e.g., search engines or social networks). Attention commonly refers to the time users spend being actively involved on the internet, deciding which digital actors deserve a share of their time and what kind of interactions they elicit. As an example, when internet users spend time on a social network platform, they are implicitly making the social network a (temporary) monopolist of their attention. In such a situation, the value that can be extracted through proper targeted advertising is tremendous (Choi et al., 2020) and indeed, there is a growing body of literature studying the evolution and the fierce competition that characterizes the attention economy (Evans, 2013, 2017). Many digital platforms act as brokers between users and content providers; after having captured users' attention, they extrapolate content preferences (from users' actions on the platform) and sell targeted ad slots to content providers. The more time users

[1] See for example this interview that appeared on the Financial Times in 2019, https://www.ft.com/content/f9f849c1-29cb-4838-a25f-0746d3836038.

spend on a platform's premises, the more it will learn who they are, their interests, and their needs, perfecting ad targeting techniques (Hagiu & Wright, 2020; Jones & Tonetti, 2019). Thus, digital platforms thrive on users' attention, a particularly scarce and contested resource due to the large number of platforms and the finite nature of users' time (Boik, Greenstein & Prince, 2016). It is increasingly recognized that digital platforms are not satisfied with small fractions of users' attention (paralleling the idea of "tipping point" in traditional multi-sided platforms – e.g., Parker, Petropoulos & Van Alstyne, 2021); A platform's profits depend on the quality of the matches between users and advertisers; users need to manifest interest in the content and eventually click on it for the platform to realize gains. Our work positions within the growing body of research on attention markets, and studies the role of digital platforms as attention brokers.

Users join digital platforms for different reasons but rarely with the main goal of seeing ads (Goldstein et al., 2014; Jenkins et al., 2016), hence a platform needs to manage the trade-off between profiting from advertisers and showing users too many ads (Todri, Ghose & Sing, 2019). For example, users are on social media platforms (e.g., Facebook, Instagram, or Twitter) to socialize and create content; when they allocate attention to these platforms, the latter learn about them and sell their attention and preferences to matched providers which ads are placed within content fruition. Ads interruptions might generate positive utility when providers' content is related to users' preferences but might also generate disutility when content is of low quality or unrelated to users' preferences. However, there is evidence that social networks or search engines (among the others) tend to overwhelm users with ads, sometimes with little concerns over quality or relatedness leading to users experiencing ad disutility (Choi et al., 2020).

Disutility caused by advertisement has been largely studied in management (e.g., DeValve & Pekeč, 2022; Marotta et al., 2021). Much of these studies focus on users' attempts to avoid ads, not much because users dislike ads per se but rather because ad-intensive websites are often populated by low-quality and unrelated ads. Many companies provide users with tools that can be used to elude ads; for example, ad blockers filter ads and only let a small fraction pass through (Gritckevich, Katona

& Sarvary, 2021). The widely accepted underlying reason according to which users want to avoid ads is that ads generally harm consumers and are a net nuisance (Anderson & Gans, 2011; Hann et al., 2008). Nevertheless, some authors adopt a view on ads that expands their spectrum and is more in line with our conceptualization of ads. For example, according to Johnson (2013) improved ad targeting may benefit users because they "may gain by witnessing more-relevant ads" (p. 129). As anticipated, attention platforms need to properly manage the trade-off between profiting from advertisers and offering users qualitative and related ads. Thus, there is a fundamental difference between traditional platforms and attention platforms: when attention is exchanged, indirect network effects might be positive as well as negative (Cennamo, 2021).

Classic two-sided markets theory hinges on the assumption that indirect network effects are positive across the sides (Rochet & Tirole, 2003, 2006). This is true if goods or services are exchanged on the platform, whence the latter facilitates transactions between buyers on one side and sellers on the other interested in completing mutually beneficial economic transactions. In this case, buyers and sellers experience positive externalities whereby buyers are better off the more sellers there are and vice versa. Digital platforms that operate according to this business model devise the pricing strategy to internalize indirect network effects (Parker & Van Alstyne, 2005); the side generating stronger externalities (i.e., the "loss" side) is subsidized to foster its participation, while the other side (i.e., the "money" side) pays to participate. This theoretical framework has been used to study multiple instantiations of two-sided platforms. For example, Caillaud & Jullien (2003) study alternative pricing schemes available to platforms that face the "chicken-and-egg" problem, whence a young platform needs to acquire its first participants. Armstrong (2006) focuses on the impact that different types of fees have on the strength of network externalities across the sides. Hagiu (2006) looks at asynchronous participation dynamics on the two sides to infer the optimal pricing strategy. Finally, Panico & Cennamo (2022) study the videogame industry with a focus on the pricing structures that platforms can devise to orchestrate their ecosystem. Attention platforms differ from these more "traditional" two-sided platforms oriented towards facilitating transactions between users, pointing at new venues to study the efficacy of pricing in managing ecosystems with asymmetric network effects.

MODELING TARGETED ADVERTISING

We examine how a platform's choice of (a linear) price for targeted advertising affects the composition of users and providers that populate the two sides, and the interaction between them. For greater concreteness, we model two pillars of online advertising. First, the number of users who see an ad, generally referred to as "impressions" (e.g., Balseiro et al., 2014; Danaher, Lee & Kerbache, 2010). We refer to this metric as the "reach" or the "advertising intensity" of a provider. Second, we model the number of users who eventually click on the ad, or the "click-through rate", that is usually computed for a given reach and dependson the relatedness of an ad to user preferences (e.g., Chatterjee, Hoffman & Novak, 2003; Ghose & Yang, 2009). These two metrics are used and evaluated complementarily by advertisers. For example, Lee, Hosanagar & Nair (2018) show how the most successful ads on Facebook optimize both impressions and click-through-rate. The characterization of the reach and the click-through-rate enables us to relate them to the platform's pricing, and our model then helps to understand the equilibrium configuration of the platform by establishing which users and providers decide to join.

To begin, we assume that the platform knows users' preferences about the content as well as the type of content offered by providers. These aspects correspond to the location of each user (x) and each provider (y) on a two-dimensional metric space, with $(x, y) \in \Re^2$. Geometrically, our model is a generalization of the well-known models of product differentiation à la Hotelling. We refer to the (Euclidean) distance $d(y, y')$ as the *horizontal* differentiation between providers y and y' ("Ad providers" in Figure 1); instead, the distance $d(x, y)$ measures the mismatch between user x's most preferred content and provider y's type of content ("Users and providers" in Figure 1). We also refer to the distance $d(x, x')$ between two users when modeling the direct network effects that depend both on the mass of users on the platforms and on the distance between users. We further assume that there is a unitary mass of users and providers at each point on the plane, though only a fraction of users and providers decide to join the platform.

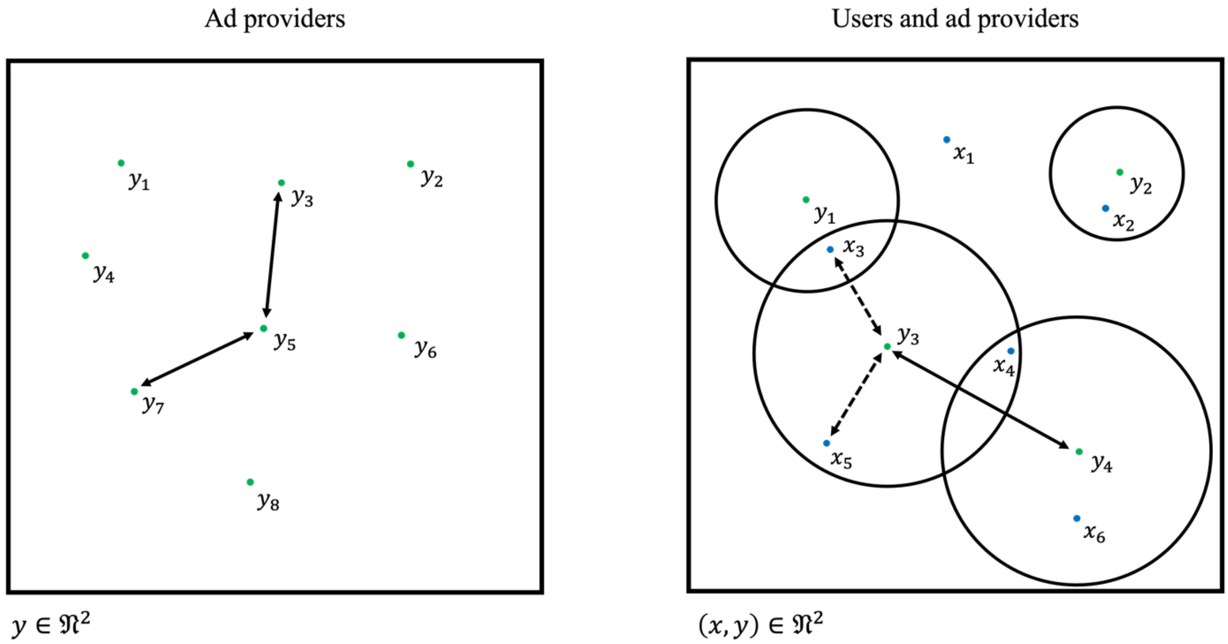

Figure 1 – Geometric (or social) graphs of providers and users. On the left, providers differ in the content they offer: double-edged solid arrows represent horizontal differentiation ($d(y, y')$); on the right, users differ in their preferences for content: double-edged dashed arrows reflect relatedness between providers and users ($d(x, y)$). The circles are the reach of ad providers, characterized by the radius; users can be reached by more than one advertiser simultaneously (e.g., x_3 or x_4).

Besides their location, users and providers differ in their individual characteristics, or *types*. Users differ in how much they get annoyed by the ad they see on the platform, according to a parameter $c \geq 0$. Providers differ in the quality of their content, according to a parameter $\theta \geq 0$. Thus, θ captures the *vertical* differentiation between content providers. Although the platform knows the exact location of each consumer (x) and each provider (y), it does not know the participants' types (c and θ), that are private information. We assume that types are i.i.d. with $c \sim F[\underline{c}, \overline{c}]$ and $\theta \sim G[\underline{\theta}, \overline{\theta}]$. Let $B \subseteq [\underline{c}, \overline{c}]$ and $S \subseteq [\underline{\theta}, \overline{\theta}]$ be the subsets (to be characterized) of the types of users and providers (at each point on the plane) that decide in equilibrium to join the platform, and $R = \frac{B}{S}$ be the ratio between users and providers. As for the platform, it charges a linear price p for the mass of users reached by a provider to advertise its content. Thus, we limit the analysis to the case of non-discriminatory pricing schemes.

Users located at x that are reached by the content of quality θ from a provider located at y obtain a (dis-)utility:

$$\theta - c(1 + d(x, y)).$$

Besides the (dis-)utility from seeing the content, users enjoy direct network effects that decay exponentially with the distance from the (mass of) other users on the platform, according to the function $e^{-d(x,x')^2}$. Thus, we assume that network effects are stronger the closer the other users are to the focal user. It can be immediately proved that when seeing the content of providers in a set $A_x \subset \Re^2$, a user x's expected payoff corresponds to:

$$u = 2\pi B \int_0^\infty x e^{-z^2} dz + SE\left[\int_{y \in A_x} [\theta - c(1 + d(x, y))] dy\right]. \qquad \text{(Eq.1)}$$

By noting that a user decides to click on a specific content if $c \leq \frac{\theta}{1+d(x,y)}$, a provider's payoff when reaching users within a distance r can be written as:

$$U = 2\pi \frac{B}{S} \int_0^r zF\left[\frac{\theta}{1+z}\right] dz - pB\pi r^2. \qquad \text{(Eq.2)}$$

The ratio $R = \frac{B}{S}$ captures both the indirect and direct network effects for the providers, according to the fact that their payoff is greater when there are more users but less providers competing for users' attention.

To provide closed-form solution, we consider the case when $c \sim U[0,1]$ and $\theta \sim U[0,1]$. With this specification, we can characterize users' and providers' equilibrium choices for a given price p. This allows to investigate the relationship amongst i) the linear pricing chosen by the platform; ii) the subset of users and providers active on the platform; iii) the advertising intensity (or reach); and iv) the average click-through-rate.

To begin, note that $F\left[\frac{\theta}{1+z}\right] = \frac{\theta}{1+z}$ and a provider solves:

$$max_r 2\pi \frac{B}{S} \int_0^r z \frac{\theta}{1+z} dz - pB\, \pi r^2. \qquad \text{(Eq.3)}$$

We can immediately prove that the optimal reach is $r^*(p) = \frac{\theta}{pS} - 1$, and therefore we can conclude that only the types $\theta \geq \hat{\theta}(p) = \frac{p}{1+p}$ join the platform. Therefore, the mass of providers that join the platform at each point in the plane is $S = 1 - \hat{\theta}(p) = \frac{1}{1+p}$. The payoff of a provider with type θ at the optimum can be written as:

$$U^*(p,\theta) = \pi \frac{B}{S}\left[\frac{\theta^2-(pS)^2}{pS} - 2\theta log\frac{\theta}{pS}\right]. \qquad \text{(Eq.4)}$$

Turning now to the other side, users' utility given the providers' optimal reach and given the (mass of) providers that populate the platform is:

$$u = \pi B + \frac{2\pi}{1+p}\int_0^{\frac{1}{p}} zE[\theta - c(1+z)|\theta \geq pS(1+z)]dz. \qquad \text{(Eq.5)}$$

We can show that among users only those with type $c \leq \hat{c}(p) = \frac{p}{(1+p)^2[8+12p(1-p^2(1+p))]}$ will join the platform, determining a mass of users at each point on the plane $B = \hat{c}(p)$. The payoff of a user at the optimum can then be written as:

$$u^*(p,c) = \pi\hat{c}(p) + \frac{\pi}{(1+p)p^4}\left[\frac{1}{12}\left(\frac{p}{1+p}\right)^2 - cp\left(p + \frac{2}{3}\right)\right] \qquad \text{(Eq.6)}$$

As a starting point for our analysis, we can offer the following results:

Proposition 1. For a given price p chosen by the platform, content of greater quality has greater reach, with $r^*(p,\theta) = \theta\frac{1+p}{p} - 1.$

Proposition 2. For a given price p chosen by the platform, the two sides are populated by users with a cost of seeing content $c \leq \hat{c}(p) = \frac{p}{(1+p)^2[8+12p(1-p^2(1+p))]}$, *and by providers with a quality of content* $\theta \geq \hat{\theta}(p) = \frac{p}{1+p}$.

Propositions 1 and 2 inform us about the populations of participants and the level of activity on the platform. As depicted in Figure 2 below, our results highlight the trade-off faced by the platform when choosing the pricing for advertising, with the thresholds $\hat{c}(p)$ and $\hat{\theta}(p)$ being both increasing with p.

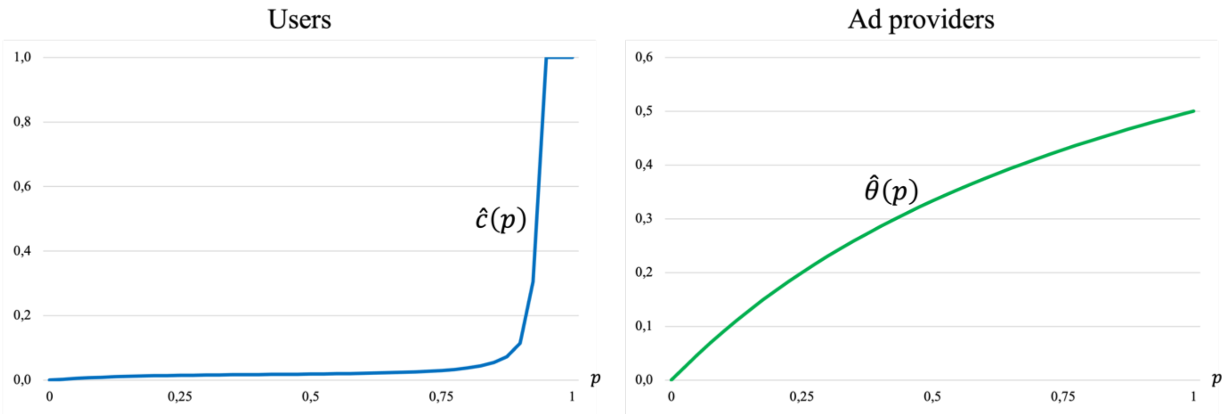

Figure 2 – Thresholds of users' participation ($\hat{c}(p)$) and providers' participation ($\hat{\theta}(p)$) as p increases.

The equilibrium participation thresholds reveal that the mass of users on the platform increases with the price whereas the mass of providers decreases with the price, as can be seen from Figure 3.

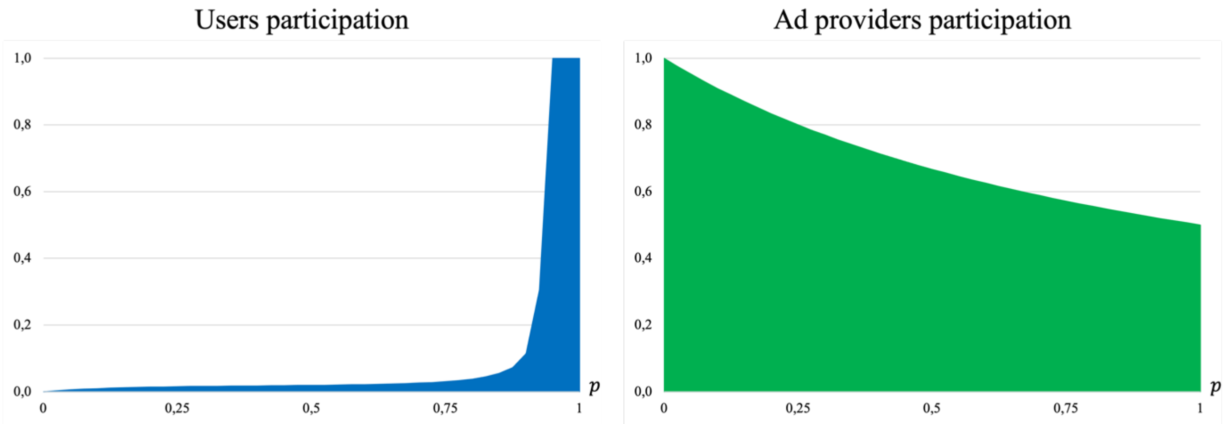

Figure 3 – Participating masses of users (left) and ad providers (right) as p increases.

On the one hand, the platform can attract more users by charging a high price and inducing less providers to join the platform (high values of p in Figure 4); in this way, the platform offers less content but of greater quality, and the larger number of users enjoy greater direct and indirect network effects. With this high-price configuration, there is a greater mass of users with a larger average cost of seeing content matched with a smaller mass of providers that offer a content of high average quality. On the other hand, the platform can decide to attract more providers by charging a low price, inducing less users to join the platform (low values of p in Figure 4). With this low-price configuration, there is a smaller mass of users with a smaller average cost of seeing ads, who are then exposed to more advertising on content of lower average quality. Given the population of users and providers on the two sides, Proposition 1 informs us about the advertising intensity for different levels of content quality. Based on these results, we could compute the average click-through-rate as to characterize the levels of activity on the platform corresponding to different configurations.

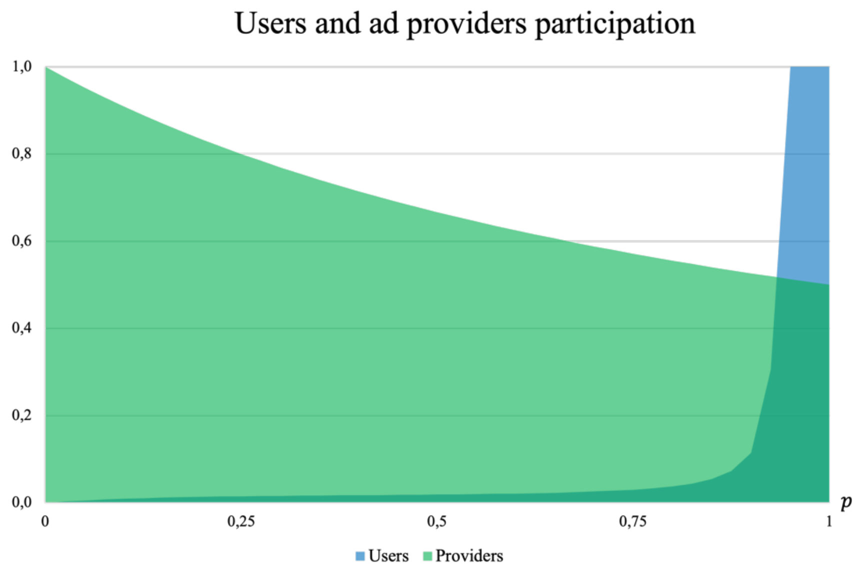

Figure 4 – Different configurations of users' and ad providers' participation as p increases. Users' maximum participation happens for p close to 1, while providers' maximum participation happens for $p = 0$.

The interesting finding behind Propositions 1 and 2 is that once we account for individual characteristics of the participants on the two sides of the platforms, the role of network effects is more nuanced. This is interesting because it departs from the usual setting where a larger mass on one side necessarily attracts a larger mass on the other side. When sheer numbers are not all that matter but also the quality of participants (on one side) plays a role, such that network externalities can be negative, then we might observe unbalanced participation in equilibrium. Moreover, the resulting dynamics are such that, as the price increases and user participation increases, on the other hand ad providers participation decreases, and vice versa. Participation on the two sides is specular, whence the platform cannot have many users and many ad providers but need to choose which side will be more populated. Equipped with these initial findings, we can proceed further in the analysis and develop a welfare analysis related to the different sides, in order to study to which extent the actors of a platform ecosystem vary in their preferences for the possible equilibrium configuration.

Next steps and extensions

We plan to investigate the platform's choice of the price based on alternative objective functions. For example, it can be shown that a platform which aims to maximize its profits Π when charging a price p to providers would maximize the following function:

$$\Pi = E\left[(p-C)\pi\left(r^*(p)\right)^2 | \theta \geq \hat{\theta}(p)\right]$$

$$= (p-C)\left(1-\hat{\theta}(p)\right)\pi \int_{\hat{\theta}(p)}^{1}\left(r^*(x)\right)^2 dx, \qquad \text{(Eq. 6)}$$

where C corresponds to the marginal cost of connecting users and providers. Instead, a platform that aims to maximize the total welfare of the participants on the two sides would maximize:

$$W = E\left(u^*(p,c)\right) + E\left(U^*(p,\theta)\right) = \int_0^{\hat{c}(p)} u^*(p,c)dc + \int_{\hat{\theta}(p)}^{1} U^*(p,\theta)d\theta \qquad \text{(Eq.7)}$$

The immediate next step is to compare the profit- and the welfare-optimizing choices, and to investigate which are the preferred configurations for the users, the providers, and the platform. This welfare analysis allows to investigate to what extent users are exposed to a greater level of advertising than it would otherwise be optimal for them or for the overall welfare of the platform participants.

A second step moves in the direction of considering alternative characterizations of users' content preferences. Assuming that the platform exactly knows this information is not too far from reality in some cases and is a useful simplification, but it also depicts an omniscient platform which omits many other real-world scenarios. Platforms are pervasive and skilled in their quest for user data (Krämer, Schnurr & Wohlfarth, 2019) and if it is true that in many cases these data may be exploited to offer a better service (Hagiu & Wright, 2020), there have also been illustrious misusages which eventually harmed users (e.g., the Facebook-Cambridge Analytica scandal in the 2010s). As a result, what users get from platform usage decreases in value if they need to forgo most of their personal information and are bombarded by ad in return. Debates on privacy are gaining prominence among academic and practitioners, with a growing share of users becoming increasingly sensible to the collection and handling of personal data by platforms (Biggar & Haimler, 2021; Cusumano, Gawer & Yoffie, 2021; Kira, Sinha & Srinivasan, 2021). Regulators are also pushing towards more transparency from platforms when it comes to user data, how they are collected, stored, shared, and used; for instance, the European General Data Protection Regulation (GDPR) or the California Consumer Privacy Act (CCPA) are prominent examples. We identified at least three relevant areas concerned with the issue of data and privacy, two of which we will extend in this second step. A first area is concerned with ad-blockers that enable users to curtail a platform's advertising (Despotakis, Ravi & Srinivasan, 2021; Gritckevich, Katona & Sarvary, 2021). Ad-blockers essentially stop most of ad providers' content to pass through to the final user; these tools offer a free service to the user and charge the platform to let some advertising pass through. A second area is focused on freemium pricing strategies, whence users can pay a price to access the service without seeing any ad instead of receiving the service for free but seeing ad (Belleflamme & Peitz, 2022). A third area considers privacy in a broader manner and investigates the instances under which platforms may be unable to collect (Acquisti, Taylor & Wagman, 2016; Casadesus-Masanell & Hervas-Drane, 2015; Goldfarb & Tucker, 2011) or unwilling to use (de Cornière & de Nijs, 2016) full information on users.

Our first extension will look at the case in which the platform does not have full information on user content preferences, that is to say it does not exactly know users' location x. This deficiency will impact the platform's ability to profit from advertisers since it impairs its capacity to match users with related ads, lowering the click-through-rate. In addition, it will also impact user participation by exposing them to less related ad, thus increasing the chances that users experience disutility and lower value during platform usage. The second extension will instead focus on the freemium pricing model, whence the platform gives users two versions of the service (see Belleflamme & Peitz, 2022, pages 173-175). Users who choose the first version can use the service for free but are exposed to advertising; instead, users who choose the second version pay a price but are not subject to advertising. This pricing formula is mainly used by digital platforms to price discriminate between users who have different willingness to be exposed to advertising; it enables them to extract more surplus from those users who suffer a greater disutility from ad. For example, both YouTube and Spotify offer a free version where ads interrupt the consumption of videos or songs and an ad-free version available under a monthly payment where consumption is never interrupted.

References

Acquisti, A., Taylor, C., & Wagman, L., 2016. The economics of privacy. *Journal of Economic Literature, Vol. 54: 442-492*.

Alphabet Inc., 2020. Annual Report 2020.

Anderson, S.P., & Gabszewicz, J.J., 2006. The media and advertising: A tale of two-sided markets. *In The Handbook of the Economics of Art and Culture, Ed. Ginsburgh V.A., & Throsby, D.: 567-614. Amsterdam, Elsevier*.

Anderson, S.P., & Gans, J., 2011. Platform siphoning: Ad-avoidance and media content. *American Economic Journal: Microeconomics, Vol. 3: 1-34*.

Armstrong, M., 2006. Competition in two-sided markets. *RAND Journal of Economics, Vol. 37: 668-691*.

Arrate Galán, A., Gonzales Cabanas, J., Cuevas, A., Calderón, M., & Cuevas Rumin, A., 2019. Large-scale analysis of user exposure to online advertising on Facebook. *IEEE Access, Vol. 7: 11959-11971*.

Balseiro, S.R., Feldman, J., Mirrokni, V., & Muthukrishnan, S., 2014. Yield optimization of display advertising with ad exchange. *Management Science, Vol. 60: 2886-2907*.

Belleflamme, P., & Peitz, M., 2022. The economics of platforms: Concepts and strategy. Cambridge University Press, New York.

Biggar, G., & Heimler, A., 2021. Digital platforms and the transactions cost approach to competition law. *Industrial and Corporate Change, Vol. 30: 1230-1258*.

Boik, A., Greenstein, S., & Prince, J., 2016. The empirical economics of online attention. *NBER Working Paper 22427*.

Caillaud, B., & Jullien, B., 2003. Chicken & egg: Competition among intermediation service providers. *RAND Journal of Economics, Vol. 34: 309-328*.

Casadesus-Masanell, R., & Hervas-Drane, A., 2015. Competing with privacy. *Management Science, Vol. 61: 229-246*.

Chatterjee, P., Hoffman, D.L., & Novak, P.T., 2003. Modeling the clickstream: Implications for web-based advertising efforts. *Marketing Science, Vol. 22: 520-541*.

Choi, H., Mela, C.F., Balseiro, S.R., & Leary, A., 2020. Online display advertising markets: A literature review and future directions. *Information Systems Research, Vol. 31: 556-575*.

Cusumano, M.A., Gawer, A., & Yoffie, D.B., 2021. Can self-regulation save digital platforms? Industrial and Corporate Change, Vol. 30: 1259-1285.

Danaher, P.J., Lee, J., & Kerbache, L., 2010. Optimal internet media selection. *Marketing Science, Vol. 29: 336-347*.

de Cornière, A., & de Nijs, R., 2016. Online advertising and privacy. *RAND Journal of Economics, Vol. 47: 48-72*.

Despotakis, S., Ravi, R., & Srinivasan, K., 2021. The beneficial effects of ad blockers. *Management Science Vol. 67: 2096-2125*.

DeValve, L., & Pekeč, S., 2022. Optimal price/advertising menus for two-sided media platforms. *Operations Research, Articles in advance*.

Evans, D.S., 2013. Attention rivalry among online platforms. *Journal of Competition Law & Economics, Vol. 9: 313-357*.

Evans, D.S., 2017. The economics of attention markets. *Working paper*.

Facebook Inc., 2020. Annual Report 2020.

Ghose, A., & Yang, S., 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science, Vol. 55: 1605-1622*.

Goldfarb, A., & Tucker, C.E., 2011. Privacy regulation and online advertising. Management Science, Vol. 57: 57-71.

Goldstein, D.G., Suri, S., McAfee, R.P., Ekstrand-Abueg, M., & Diaz, F., 2014. The economic and cognitive costs of annoying display advertisements. *Journal of Marketing Research, Vol. 51: 742-752*.

Gritckevich, A., Katona, Z., & Sarvary, M., 2021. Ad blocking. *Management Science, Articles in advance*.

Hagiu, A., 2006. Pricing and commitment by two-sided platforms. *RAND Journal of Economics, Vol. 37: 720-737*.

Hagiu, A., & Wright, J., 2020. Data-enabled learning, network effects and competitive advantage. *Working paper*.

Hann, I.H., Hui, K.L., Lee, S.Y.T., & Png, I.P.L., 2008. Consumer privacy and marketing avoidance: A static model. *Management Science, Vol. 54: 1094-1103*.

Jenkins, J.L., Anderson, B.B., Vance, A., Kirwan, C.B., & Eargle, D., 2016. More harm than good? How messages that interrupt can make us vulnerable. *Information Systems Research, Vol. 27: 880-896*.

Johnson, J.P., 2013. Targeted advertising and advertising avoidance. *RAND Journal of Economics Vol. 44: 128-144*.

Jones, C.I., & Tonetti, C., 2019. Nonrivalry and the economics of data. *Working paper*.

Lee, D., Hosanagar, K., & Nair, H.S., 2018. Advertising content and consumer engagement on social media: Evidence from Facebook. *Management Science, Vol. 64: 5105-5131*.

Kira, B., Sinha, V., & Srinivasan, S., 2021. Regulating digital ecosystems: Bridging the gap between competition policy and data protection. Industrial and Corporate Change, Vol. 30: 1337-1360.

Krämer, J., Schnurr, D., & Wohlfarth, M., 2019. Winners, losers, and Facebook: The role of social logins in the online advertising ecosystem. Management Science, Vol. 65: 1678-1699.

McIntyre, D.P., & Srinivasan, A., 2017. Networks, platforms, and strategy: Emerging views and next steps. *Strategic Management Journal, Vol. 38: 141-160*.

Marotta, V., Wu, Y., Zhang, K., & Acquisti, A., 2021. The welfare impact of targeted advertising technologies. *Information Systems Research, Articles in advance*.

Panico, C., & Cennamo, C., 2022. User preferences and strategic interactions in platform ecosystems. *Strategic Management Journal, Vol. 43: 507-529*.

Parker, G.G., & Van Alstyne, M., 2005. Two-sided network effects: A theory of information product design. *Management Science, Vol. 51: 1494-1504*.

Parker, G.G., Petropoulos, G., & Van Alstyne, M., 2020. Digital platforms and antitrust. *In The Oxford Handbook of Institutions of International Economic Governance and Market Regulation (Chapter), Ed. Brousseau, E., Glachant, J.M., & Sgard, J.: 1-33. Oxford University Press, Oxford*.

Prat, A., & Valletti, T., 2021. Attention oligopoly. *Working paper*.

Rochet, J.C., & Tirole, J., 2003. Platform competition in two-sided markets. *Journal of the European Economic Association, Vol. 1: 990-1029*.

Rochet, J.C., & Tirole, J., 2006. Two-sided markets: a progress report. *RAND Journal of Economics, Vol. 37: 645-667*.

Statista, 2022. Daily time spent on social networking by internet users worldwide (consulted on January 22, 2022).

Todri, V., Ghose, A., & Singh, P.V., 2019. Trade-offs in online advertising: Advertising effectiveness and annoyance dynamics across the purchase funnel. *Information Systems Research, Vol. 31: 102-125*.

Wu, T., 2016. The attention merchants: The epic struggle to get inside our heads. *Atlantic Books, London*.

Wu, T., 2019. Blind spot: The attention economy and the law. *Antitrust Law Journal, Vol. 82: 771-806*.