"In medio stat virtus": Targeted advertising in social media platforms with heterogeneous participants

Pietro Morino, Bocconi University, Milan

pietro.morino@phd.unibocconi.it

Claudio Panico, Bocconi University, Milan

claudio.panico@unibocconi.it

Carmelo Cennamo, CBS, Copenaghen

cce.si@cbs.dk

Abstract. Social media platforms face a strategic trade-off when addressing the needs of different user groups: while content providers would like to reach many users with their ads, users are not always interested in the advertised 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 how the platform manages this positive-negative externalities trade-off by resorting to targeted advertising to minimize users' disutility from advertising. We find that contrary to mainstream theory's prediction of positive, self-reinforcing indirect network effects, platform configurations can have an unbalanced number of participants on the two sides. We show that "in medio stat virtus" principle applies indeed to this governance trade-off: the optimal governance strategy for the platform lies in-between the pricing level maximizing user benefits and that maximizing provider benefits.

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

1. Introduction

Digital platforms, particularly social media platforms, are increasingly becoming "attention brokers", catalyzing the significant amount of time that individuals spend online (Evans, 2019; Wu 2019). 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 (Peitz, 2020; Van Zandt, 2004), content providers and advertisers must devise strategies to engage users and get a share of this attention (Rossi & Rubera, 2021). A market for user attention can thus emerge around a digital platform that engages users on one side, also informing them on the advertised content of providers on the other side (Jullien & Sand-Zantman, 2021). Compared to canonical transaction platforms like eBay, where indirect network effects are positive and the sides interact to exchange economic value (Rochet & Tirole, 2003), the same does not necessarily hold for platforms in the "attention economy" (Wu, 2016).

Information platforms like social media platforms, that leverage user attention and facilitate information flows within and across the sides (Cennamo, 2021), face a different matching problem when governing platform's distinct market sides. Users join Facebook to interact usually with other users sharing similar interests and access content of interest to them. Content providers and advertisers join Facebook to intercept part of these users' attention with their own content's value proposition. Thus, compared to transaction platforms wherein users join the platform to search for and interact with providers, in information platforms such as Facebook, content providers' activity may be a nuisance to users and in fact, demote the core value proposition of the platform to them to the extent that providers' content is irrelevant to them or of low quality (Cennamo 2021).

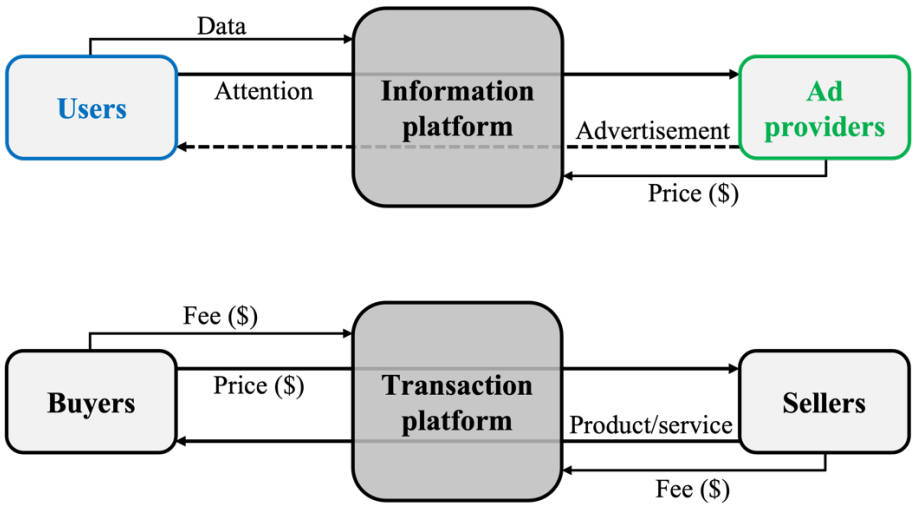

Figure 1 – On an information platform, indirect network effects from users to content providers are positive (bold and solid line), while indirect network effects from content providers to users might be negative (bold and dashed line). On a transaction platform.

Indirect network effects on information platforms can thus go from positive to negative, giving rise to an "asymmetry" between the two sides which is depicted in Figure 1 (top graph). This asymmetry creates a management problem, for which current theorizing on platform dynamics is of little help. There is a voluminous literature exploring the optimal pricing and governance structure for platforms with positive indirect network effects (bottom graph in Figure 1) (e.g., Armstrong, 2006; Caillaud & Jullien, 2003; Hagiu, 2006). However, with few exceptions (see Anderson & Gans, 2011; Lin, 2020), we know little about how platforms manage negative indirect network effects. Also, some scholars contend that, because of this asymmetry, social media platforms have the incentives to "exploit" their user base and favor the side of the market generating profits, i.e., content providers (Monti, 2020). Specifically, users might be exposed to more 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). In the cases in which advertisement might 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 for the providers as well. Standing the asymmetry, and these

possible consequential negative effects, a natural question is whether and to what extent platforms can leverage and extend their user base to attract advertisers, and how the different pricing governance design affects user-advertiser interactions and the value of the platform to each user group.

Platforms might collect data about users' activity on the platform and through that develop deep knowledge of users' preferences. This, in principle, can solve the problem of irrelevant content, and thus management of the asymmetry of indirect network effects (Marotta et al., 2022; Zhang & Katona, 2012). However, while platforms may know users' preferences for content and the nature of content offered by providers, they have but very imperfect information about users' tolerance (or disutility) for the focal advertisement and the quality of the focal content being sponsored by the provider (e.g., Despotakis et al., 2021). Consider social media platforms like Facebook. They usually offer a free service to the end users and then charge the content providers who want to advertise their content[1]. Users derive value from Facebook to the extent that they can exchange relevant information with other users. Facebook designs and manages the platform to enable valuable interactions among users and to compile users' preferences profiles for targeted advertising. Intensive monitoring enables the social media platform to have accurate accounts of what users like or dislike (Srinivasan, 2019). Yet, it is much harder for Facebook to know whether users are annoyed by advertising and, if so, to what extent. For example, users with similar (or even identical) content preferences may vary widely in their perception of advertising, some being annoyed by any ads, others enjoying related ads and some others being indifferent. Facebook knows that a click on the ad is the result of successful targeting, while no click is an ambiguous outcome that might be due to unsuccessful targeting,

[1] 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).

inattention, or annoyance. In addition, Facebook does not know the quality of providers' content, which may influence users' perception of an ad. Indeed, Facebook does not conduct specific screening of the providers who decide to advertise on the social media as long as they adhere to few basic principles (e.g., no violence, no abuse, etc.). Facebook enables providers to decide how many users will be reached by their content, but the platform ultimately chooses which users exactly are displayed the content through content relatedness (probabilistic) models.

In our analysis we start from these departing features of information platforms to study the market (for user attention) dynamics and the platform's optimal governance of it. The setting we investigate closely resembles existing social media platforms like Facebook, and might help to shed light on the ongoing debate concerning the extent to which such platforms profit at the expense of their users by ignoring negative ads' externalities. We assume that users and providers have private information about their cost of seeing unrelated ads and their quality respectively, and model the interactions dynamics between users and content providers in relation to such heterogeneity characteristics[2].

We find that, in equilibrium, user participation on the platform increases while content provider participation decreases when the platform increases the price for targeted advertising. As a result, information platforms may be characterized by unbalanced participation: for low prices

[2] The platform makes profits by charging content providers based on the number of users that they reach through the platform. On one side, providers choose how much advertising they will conduct on the platform based on the pricing and on the quality 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 and experience a disutility. The two sides connected by the platform enjoy direct as well as indirect network effects, whose intensity and sign depend on the interrelation of users' and providers' characteristics.

(charged for advertisement), there will be many content providers but few users; for high prices, the opposite is true. Also, we show that platform's value to users (or "user welfare") is higher for high prices, while platform's value to providers (or "provider welfare") is higher for low (but not too low) prices. This creates a governance challenge for the platform: users are better off with a platform configuration whereby there are many other users and few providers sponsoring high-quality content. Providers instead are better off with a platform configuration with many users but low pricing for ads, granting them the choice to advertise more. Increasing (or decreasing) the price charged to content providers for their ads is thus a critical mechanism for the platform to govern these asymmetric platform participation benefits for users and content providers, and thus minimize the negative network externalities while enabling both sides' participation. We show that the optimal pricing strategy lies in the middle of the two extremes (as per Aristoteles's maxim "in medio stat virtus"), i.e., in-between the pricing maximizing user welfare and that maximizing provider welfare. In this optimal platform configuration, there will be many users participating but only high-quality providers. Thus, this way, the platform solves the trade-off standing in the benefits from platform participation between the users and content providers by fostering user participation while pushing low-quality providers out of the platform.

We contribute to platform theory and the recent developments in the platform literature (e.g., Lin, 2020; Panico & Cennamo, 2022), offering a more nuanced view of the role of indirect network effects and the governance trade-offs and challenges faced by the platform. We show and elaborate on the importance and consequences of negative indirect network effects for information platforms managing attention markets. Our findings, which depart from mainstream theory's positive indirect network effects and self-reinforcing loops (e.g., Rochet & Tirole, 2003), enhance the general understanding of information platforms where the sides

trade information instead of goods, of which social media platforms are the most excellent representative.

Our findings can also contribute to the debate on user exploitation and provide important policy implications. On the one hand, we show that some degrees of user exploitation may be in place in such contexts since the platform's optimal price is lower than the price that would maximize user welfare, in line with the arguments put forth by some scholars (e.g., Monti, 2020). In equilibrium, there will be more providers sponsoring content of relatively lower average quality than would be best for users. On the other hand, we also show that the platform's optimal price is higher than the price that would maximize provider welfare. This implies that the platform is willing to lose participation from possibly a significant number of providers (with a lower reservation price for platform participation than the one being set by the platform) to attract more users and maintain a strong value proposition for them. This conclusion is consistent with the argument advanced by other scholars that information platforms cannot grant themselves the luxury of ignoring their user base without consequences (see Peitz, 2020).

The paper is organized as follows: In Section 2 we draw on the literature to develop more on where this work situates and how it contributes to the debate on platforms. In Section 3, we introduce the analytical model employed to study attention platforms and targeted advertising. In Section 4, we develop a welfare analysis and introduce our main results. In Section 5, we discuss our findings and why they are relevant to advancing our understanding of platforms. Finally, in Session 6, we conclude and trace a roadmap for the future.

2. Theoretical background

Mainstream literature on multi-sided markets (Armstrong, 2006; Caillaud & Jullien, 2003; Eisenmann et al., 2006; Hagiu, 2006; Katz & Shapiro, 1994; McIntyre & Srinivasan, 2017; Parker & Van Alstyne, 2005) maintains that digital 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, which leads to the instantiation of positive indirect network effects across the sides of the market. The notion of positive INEs has characterized much of the value-enhancing results of platform-based market models (Rochet & Tirole, 2003, 2006). Positive INEs (often) materialize when goods or services are exchanged, wherein the platform 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 (Cennamo & Santaló, 2013; Evans, 2003).

Information platforms differ from more traditional transaction platforms in that information platforms are concerned with facilitating the exchange of information among users, generating value through social interactions and engagement (Cennamo, 2021; de Coernière & Sarvary, 2022; Toubia & Stephen, 2013). For instance, users of social media platforms value gaining access to relevant information that gets exchanged within and across the sides, but information can be of relevance to different users to different degrees (Jullien & Sand-Zantman, 2021). Users enjoy creating and sharing content with a selected circle of users but may not enjoy (potentially irrelevant) content that they receive from advertisers (Anderson & Gans, 2011; Hann et al., 2008). That is, information platforms are characterized by asymmetric indirect network effects across the sides, which deeply undermine the accepted notions of platform coordination (see for example Rochet & Tirole, 2003). For example, Lin (2020) investigates second-degree price discrimination in media platforms (e.g., online newspapers) and finds that depending on users' annoyance from seeing ads, different platform configurations may emerge in equilibrium.

The relevance of the information and content available on a platform, which influences the time that users are willing to dedicate to it, is an element of paramount importance within the

governance of information platforms (Evans, 2019). The time that users dedicate to a platform reflects the attention that the platform is capable of eliciting (Wu, 2019); a subset of information platforms (e.g., search engines and social networks) trade users' attention as the main currency, giving rise to "attention markets" (Wu, 2016). During the time a user is on an "attention platform", the platform becomes a (temporary) monopolist of her attention and can profit by selling this attention to business providers (Prat & Valletti, 2022). For example, when a user is engaged on Facebook, her attention is completely absorbed by the social media and cannot be simultaneously captured by a competitor. The value that attention platforms can extract by exploiting the temporary monopoly on users' attention is tremendous, as the growing body of literature studying attention platforms points out (Evans, 2017; Iyer & Katona, 2016; Rossi & Rubera, 2021).

One reason why user attention is so valuable to attention platforms is that while users are hooked up, relevant information can be inferred from their behaviors and actions on the platform (Alaimo & Kallinikos, 2017; Iyer et al., 2005). On the one hand, platforms' knowledge of users has several dimensions which correspond to user preferences concerning different attributes and features of products and services (Long et al., 2022; Trusov et al., 2016). On the other hand, whether all of these dimensions are necessary to conduct targeted advertising (Despotakis & Yu, 2022; Liu et al., 2022) and whether attention platforms can infer the extent to which users dislike being the object of advertising remain ambiguous (Despotakis et al., 2021; Gritckevich et al., 2021). For example, Arrate Galán et al. (2019) show how Facebook collects more than 1,000 preferences per user but only uses a small fraction (in the order of 0.7% to 22.8%) to conduct targeted advertising. All in all, deep user knowledge tends to increase the effectiveness of advertising through ads that are more aligned to user preferences (e.g., Marotta et al., 2022).

Another reason why attention platforms vie for user attention is that it is a limited and scarce resource (Peitz, 2020; Van Zandt, 2004). Attention is limited because there is a cap to the time users can spend online in a day (i.e., the 24 hours), and essential activities like sleeping or working consume much of users' attention. Moreover, the growing number of platforms providing online offerings aimed at capturing users' attention make it a contested and, thus, scarce resource, easy to compete away (Anderson & de Palma, 2012; Evans, 2013). In the past decade of internet boom, users have been frequently switching their attention from one platform to another (Boik et al., 2016). One can think of those that were successful a decade ago and have (almost) disappeared today like MSN, Myspace, or Google Plus. Users have instead proved much less prone to changing the share of time they spend online and the number of platforms they are able to dedicate their attention to (Boik et al., 2016).

All in all, advertisers value reaching customers through targeted ads on attention platforms because the platform's knowledge of users guarantees higher chances that users are interested in the sponsored content compared to traditional advertising methods (Johnson, 2013). For this reason, some say that attention platforms tend to exploit users through intensive (data collection and) advertising aimed at extracting as much profit from advertisers as possible (Monti, 2020). User exploitation arises when intensive advertising, even if targeted, sacrifices relevance and quality for quantity and harms those users who suffer more from being the target of ads (Choi et al., 2020; Todri et al., 2019). This creates a tradeoff between the platform value to users and to providers. Attention platforms have been held up as unable to internalize the (potentially) negative externalities that excessive advertising imposes on some users (Anderson & Gabszewicz, 2006). Opposed to this view, some others say that platforms have no incentive to exploit users as this might disengage them from the platform (Peitz, 2020). For example, Anderson & de Palma (2009) show that the presence of a network coordinator (i.e., a "gatekeeper" platform) in an email network might reduce users' overcrowding with messages

(i.e., "spam" emails) from advertisers by charging a price per message sent[3]. On another note, Prat & Valletti (2022) show that a monopolistic attention platform might impose limitations to excessive advertising to create an "attention bottleneck" that reduces advertisers' access to users and enables the platform to charge higher prices to advertisers.

There is no consensus yet on how platforms act as coordinators to maximize profits from targeted advertising while at the same time internalizing the (potentially) negative externality of excessive targeting. In this paper we build theory through mathematical modelling that combines information, attention, and targeted advertising with typical elements of two-sided markets to understand how a social media platform uses (pricing in) targeted advertising to coordinate the two sides of its attention market given the standing asymmetry. We study the impact of the coordination strategy adopted by the platform on the configuration of the two sides, that is on user and provider participation, and on user and provider welfare as well as on the platform profits.

3. The model

We set up a model to examine how a social media platform's pricing strategy for targeted advertising affects the composition of users and providers that populate the two sides of the platform and the interaction between them. We distinguish between the number of users who see an ad, or impressions (Danaher et al., 2010), which we refer to as the "reach" or the "advertising intensity" of a provider, and the number of users who eventually click on an ad, which we refer to as the "click-through rate" (Chatterjee et al., 2003; Ghose & Yang, 2009). The click-through-rate is usually computed for a given number of impressions, and it is

[3] Here, Anderson & de Palma (2009) compare the case of an open-access network and that of a network with a gatekeeper platform that internalizes the user-attention-as-a-common-property-resource problem but, at the same time, maximizes profits.

interpreted as users' appreciation for an ad and a measure of success of targeted advertising (Balseiro et al., 2014). These two metrics are used and evaluated complementarily both by advertisers and digital platforms like social media platforms. For instance, Lee et al. (2018) show that the most successful ads on Facebook are the ones that optimize on both impressions and click-through-rate. The characterization of providers' reach and click-through-rate enables us to realistically model the pricing of targeted advertising as a coordination mechanism by a social media platform like Facebook.

We envision the users' preferences for a given content and the specific content being promoted by the advertisers distributed in a two-dimensional geometric space, such that they correspond to the location of each user (x) and each provider (y) on the geometric space, with $(x, y) \in \Re^2$. Geometrically, our model is a generalization of the well-known models of horizontal product differentiation à la Hotelling. We refer to the (Euclidean) distance $d(y, y')$ as the *horizontal* differentiation between providers y and y' ("Content providers" in Figure 2); 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 content providers" in Figure 2). 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 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 these users and providers decide to join the platform in equilibrium.

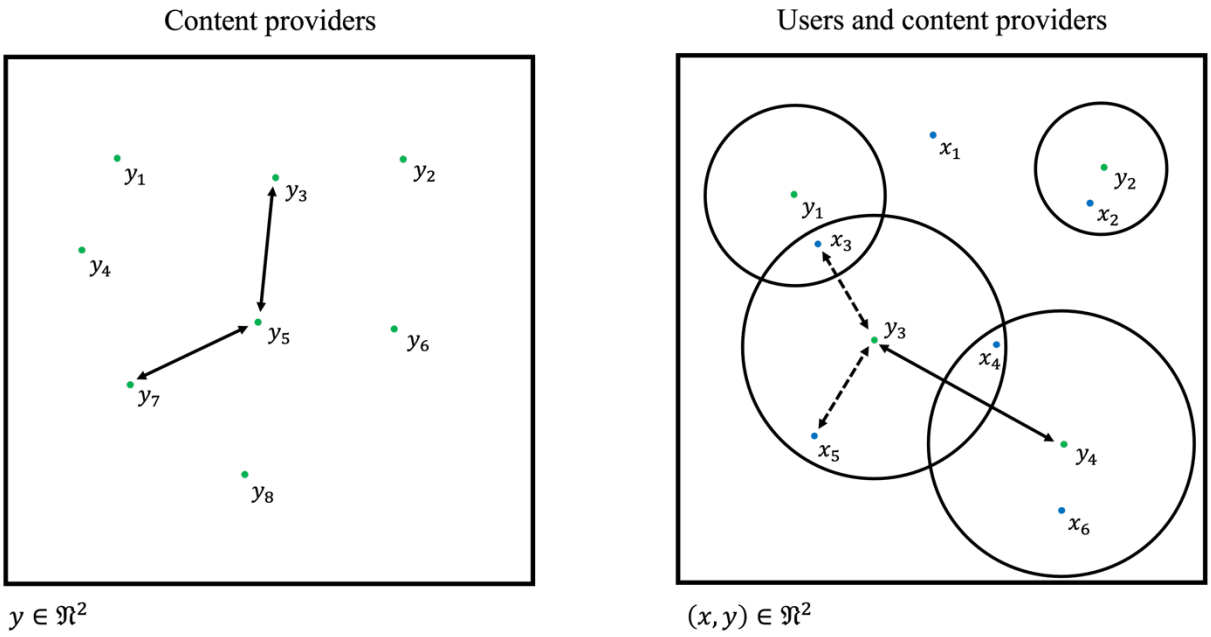

Figure 2 – Geometric (or social) graphs of content providers and users. On the left, providers differ in the content they offer: double-edged solid arrows represent horizontal differentiation ($d(y, y')$). On the right, users differ in their preferences for content: double-edged dashed arrows reflect relatedness between providers and users ($d(x, y)$). The circles are the reach of content providers, characterized by a certain 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. Given the amount of information that social media platforms collect on their users (Arrate Galán et al., 2019; Srinivasan, 2019), we assume that the platform knows users' preferences about the content as well as the nature of the content offered by providers. That is to say, the platform knows the exact location of each consumer (x) and each provider (y). On the other hand, the platform is not informed about how much users are annoyed by advertising and the quality of the content sponsored by providers. It is difficult for social media platforms to determine the extent to which users are bothered from advertisement (e.g., Gritckevich et al., 2021). Moreover, social media platforms are not interested in thoroughly classifying the quality of the content providers,

rather they use indirect mechanisms (such as pricing) to select out low-quality providers (e.g., Lin, 2020). This is modelled by assuming that types are i.i.d. with $c \sim F\left[\underline{c}, \overline{c}\right]$ and $\theta \sim G\left[\underline{\theta}, \overline{\theta}\right]$. The platform charges a linear price[4] 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. Let $B \subseteq [0,1]$ and $S \subseteq [0,1]$ be the mass of types of users and providers at each point on the plane that decide to join the platform in equilibrium, and let $R = \frac{B}{S}$ be the ratio between users and providers. Users located at x that are reached by the content of quality θ from a provider located at y obtain a (dis-)utility:

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

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^{-t^2}. Thus, we assume that network effects are stronger the closer the other users are in terms of preferences 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:

$$u = 2\pi B \int_0^{\infty} te^{-t^2} dt + 2\pi S \int_{y \in A_x} E\big[\theta - c\big(1 + d(x,y)\big)\big] dv, \qquad \text{(Eq. 1)}$$

with $d(x,y) = v$.

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:

[4] Many platforms use a probabilistic model and auction theory to conduct targeted advertising. We believe that linear pricing is able to capture the intuition behind these models because these platforms initially charge providers for the number of users they aim to reach and only after use probabilistic methods to understand which specific users will be shown an ad (as for example Facebook).

$$U = 2\pi R \int_0^r vF\left[\frac{\theta}{1+v}\right] dv - \pi pBr^2. \qquad \text{(Eq. 2)}$$

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 and less providers competing for users' attention.

To solve the model, we consider the case when $c \sim U[1,2]$ and $\theta \sim U[1,2]$. 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 content providers active on the platform; iii) the advertising intensity (or reach); and iv) the average click-through-rate.

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

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

We can immediately prove that the optimal reach is $r^*(p,\theta) = \frac{\theta}{1+pS} - 1$, and therefore we can conclude that given $S(p) = 2 - \hat{\theta}(p)$, only the types $\theta \geq \hat{\theta}(p) = \frac{1+2p}{1+p}$ join the platform. As a consequence, the mass of providers that join the platform at each point in the plane is $S(p) = \frac{1}{1+p}$. The payoff of a provider with type θ at the optimum can be written as:

$$U^*(p,\theta) = \pi B(1+p)\left\{2\theta\left[\frac{\theta}{\hat{\theta}(p)} - 1 - \log\left(\frac{\theta}{\hat{\theta}(p)}\right)\right] - \hat{\theta}(p)\left(\frac{\theta}{\hat{\theta}(p)} - 1\right)^2\right\}. \qquad \text{(Eq. 4)}$$

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

$$u = \pi B + 2\pi S(p) \int_0^{r^{max}(p)} E[\theta - c(1+v) | r^*(p,\theta) \geq v] dv. \qquad \text{(Eq. 5)}$$

A user is reached – and consequently experiences a (dis-)utility – by all providers whose reach is between 0 and $r^{max}(p) = \frac{1}{1+2p}$, conditional on the chosen reach (i.e., $r^*(p,\theta)$) being larger than the distance between the provider and the user (i.e., $v = d(x,y)$). We can show that only the users with type $c \leq \hat{c}(p) = \frac{(1+2p)\{6p(1+2p)[4+p(5+2p)]-5\}}{2(1+p)\{3p(1+2p)[5+4p(2+p)]-2\}}$ will join, such that the mass of users at each point on the plane is $B(p) = (\hat{c}(p) - 1)$. The payoff of a user at the optimum can then be written as:

$$u^*(p,c) = \pi B + \frac{\pi}{1+p}\left[\frac{11+34p+24p^2}{6(1+2p)^3(1+p)} - \frac{5+6p}{3(1+2p)^3}c\right]. \qquad \text{(Eq. 6)}$$

These initial results lead to a first proposition.

> *Proposition 1. For a given price p chosen by the platform, the two sides are populated by users with a cost of seeing advertised content $c \leq \hat{c}(p) = \frac{(1+2p)\{6p(1+2p)[4+p(5+2p)]-5\}}{2(1+p)\{3p(1+2p)[5+4p(2+p)]-2\}}$, and by providers with a quality of content $\theta \geq \hat{\theta}(p) = \frac{1+2p}{1+p}$.*

Proposition 1 reveals how user and provider participation depend on the price chosen by the platform, in such a way that higher prices will enhance user participation and dampen provider participation. This trade-off faced by the platform can be better understood by looking at Figure 3, where both participation thresholds, $\hat{c}(p)$ and $\hat{\theta}(p)$, increase with p. A price close to zero attracts all providers but elicits low participation on the user side, while higher prices attract all users but less and less providers.

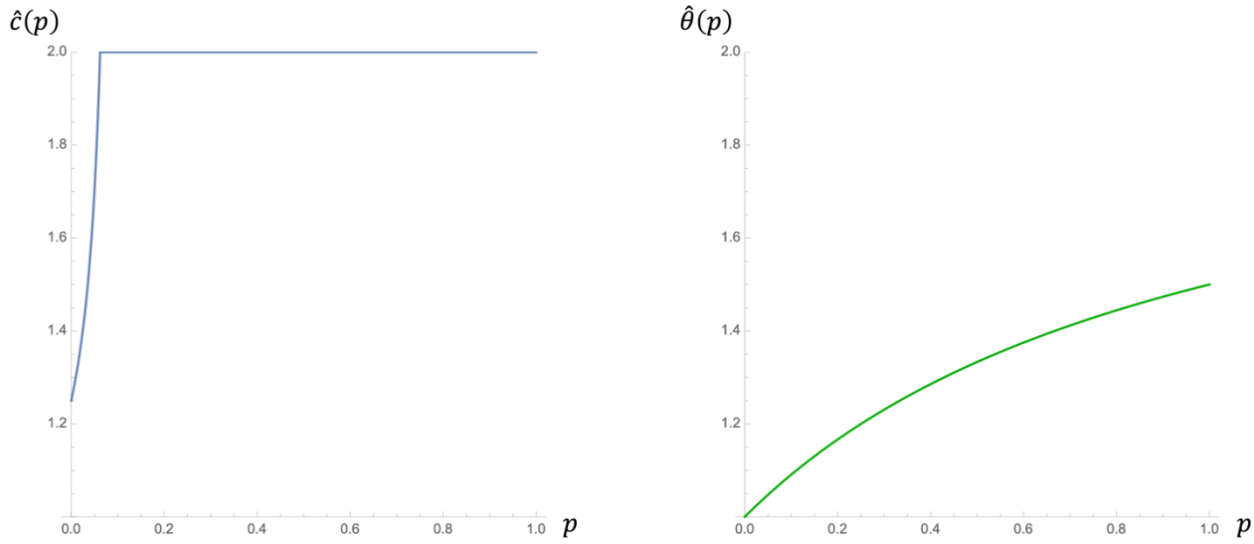

Figure 3 – User ($\hat{c}(p)$) and provider ($\hat{\theta}(p)$) participation threshold.

On the one hand, the platform can attract more users by charging medium to high prices and inducing less providers to join the platform (see Figure 4). In this way, providers on the platform offer less content but of greater quality, and the larger number of users enjoy greater direct and indirect network effects. With this medium/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 a greater mass of providers by charging a low price (see Figure 4), while inducing less users to join the platform. With this low-price configuration, there is a smaller mass of users with a lower average cost of seeing ads, who are then exposed to more advertising on content of lower average quality.

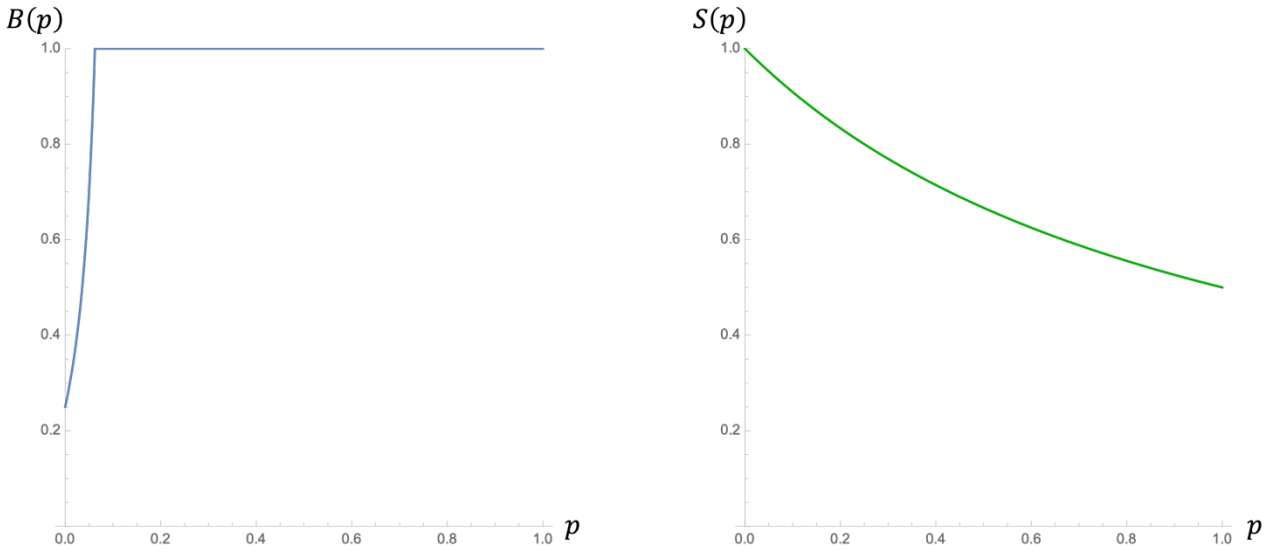

Figure 4 – Mass of users ($B(p)$) and providers ($S(p)$) on the platform.

The interesting finding behind Proposition 1 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 than in standard platform models. This is interesting because it departs from the usual wisdom that 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 the provider side) plays a role, network externalities can be negative and we might observe asymmetric participation in equilibrium.

Given the population of users and providers on the two sides, we obtain a second result.

> *Proposition 2. For a given price p chosen by the platform, content of greater quality has a greater reach, with* $r^*(p,\theta) = \frac{\theta}{1+pS(p)} - 1$.

Proposition 2 informs us about the advertising intensity for different levels of content quality θ. Providers whose content is of higher quality choose a higher reach and advertise more broadly for any price charged by the platform, as can be seen in Figure 5 whence $r^*(p,\theta)$ is higher for any p for those providers with a higher θ. Together with Proposition 1, this tells us

that when the platform charges high prices the network will be populated by fewer providers of higher average quality whose content reaches more users.

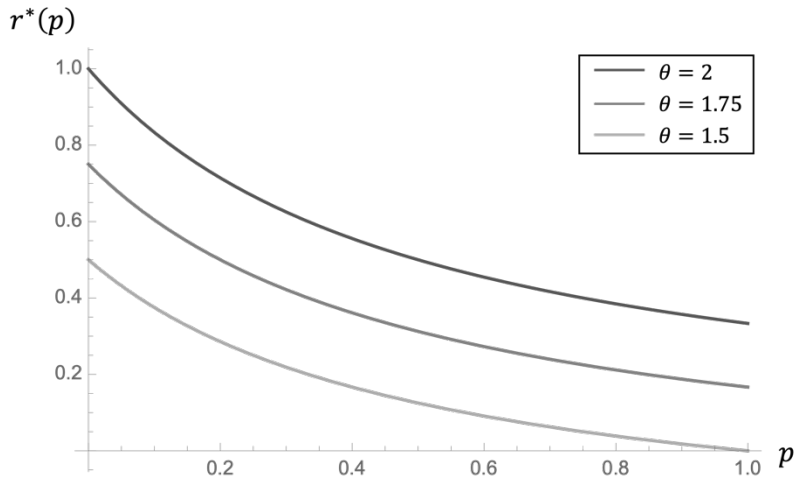

Figure 5 – Providers set different reaches (or advertising intensities) according to different θ.

4. Welfare analysis

In this section we analyze user and provider welfare given the endogenous participation thresholds $\hat{c}(p)$ and $\hat{\theta}(p)$, i.e., for a given price. We first aim to understand how the two sides fare under different price levels set by the platform, and then we investigate the levels of welfare corresponding to the optimal price chosen by the platform to maximize its own profits.

First, because the platform is populated only by users whose cost of seeing (unrelated or low-quality) ads is low enough, users' welfare correspond to:

$$W^U(p) = \int_1^{\hat{c}(p)} u^*(p,c)dc \,, \qquad \text{(Eq. 7)}$$

that after plugging the equilibrium values reduces to:

$$W^U(p) = \pi\,(\hat{c}(p)-1)^2 + \frac{\pi}{1+p}\left\{\frac{11+12p}{6(1+2p)^2(1+p)}(\hat{c}(p)-1) - \frac{5+6p}{6(1+2p)^3}\left[\left(\hat{c}(p)\right)^2 - 1\right]\right\}. \quad \text{(Eq. 8)}$$

User welfare encompasses both direct and indirect network effects. Direct network effects capture what users get from the interaction with other users on the platform, while indirect

network effects pertain to the interaction between users and providers. As can be seen in Figure 6, user welfare quickly grows as the platform increases the price and elicits more participation from users by decreasing the number and increasing the quality of providers on the platform. When all users join the platform ($\hat{c}(p) = 2$) and that side of the market becomes fully covered, indirect network effects become prominent. Every increase in the price means that users experience lower nuisance from unrelated or low-quality sponsored content. Yet, this benefit decreases at the margin, such that for high prices the platform is not able to significantly improve user welfare by reducing the number and increasing the quality of providers. Overall, users are best off with high prices, that is to say when the platform is populated by few and high-quality providers and there are many other users with whom to interact. Thus, our model suggests that users' welfare is maximal when users are exposed to few but high quality and relevant ads (i.e., $p = p^{max} = 1$). We can then prove the following results:

Proposition 3. Users' welfare is monotonic in the price and it first accelerates and then decelerates when the market is fully covered.

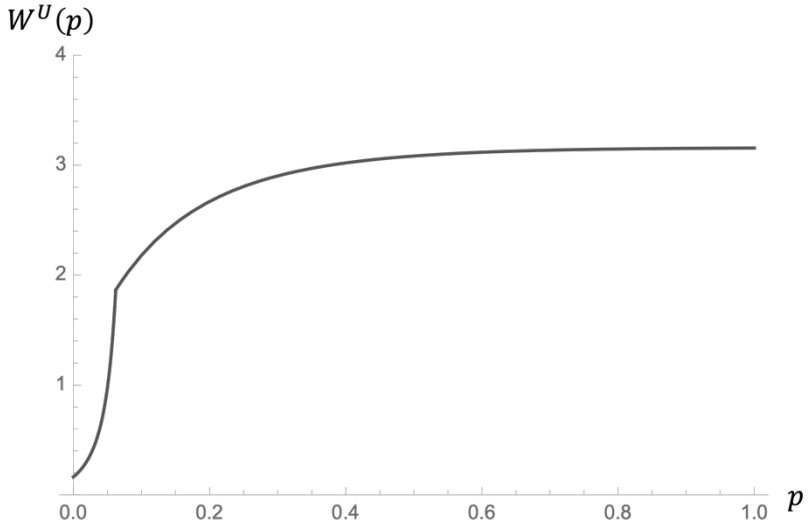

Figure 6 – User welfare.

Equivalently, because only those providers whose quality of the advertised content is high enough will join the platform, the welfare of providers is:

$$W^P(p) = \int_{\hat{\theta}(p)}^{2} U^*(p,\theta)d\theta. \qquad \text{(Eq. 9)}$$

that after plugging the equilibrium values reduces to:

$$W^P(p) = \frac{2\pi(\hat{c}(p)-1)}{3(1+p)}\left\{\frac{6p^2+14p+6}{1+2p} - 3(1+p)^2\left[log(4) + 2\,log\left(\frac{1+p}{1+2p}\right)\right]\right\}. \qquad \text{(Eq. 10)}$$

The overall provider welfare is driven by the click-through-rate, that is how many users click on the sponsored content, for a certain reach optimally set by the provider. Providers of higher quality choose a greater reach and remain on the platform even when the platform charges high prices, converting more of the reached users into clicks. On the other hand, providers of lower quality leave the platform as the price increases. This generates a non-monotonic welfare function, as can be seen in Figure 7. As long as more users join the platform when the price increases, the providers' welfare increases due to the greater indirect network effects. Provider welfare increases for low prices because of the low direct (negative) impact of the price and the rapid initial increase in user participation, which is the indirect (positive) impact of the price. When the user market becomes fully covered, the overall welfare of providers decreases with the price as there are less and less active providers and each pays a high price for the reach thus gaining less profits. When all users are onboard, for a price $p' : \hat{c}(p') = 2$ (and $B(p') = 1$), then the direct impact of the price becomes predominant and drives down the welfare of providers. We can then prove the following result:

Proposition 4. Providers' welfare increases with the price as long as more users keep on joining the platform, and then decreases when the market becomes fully covered.

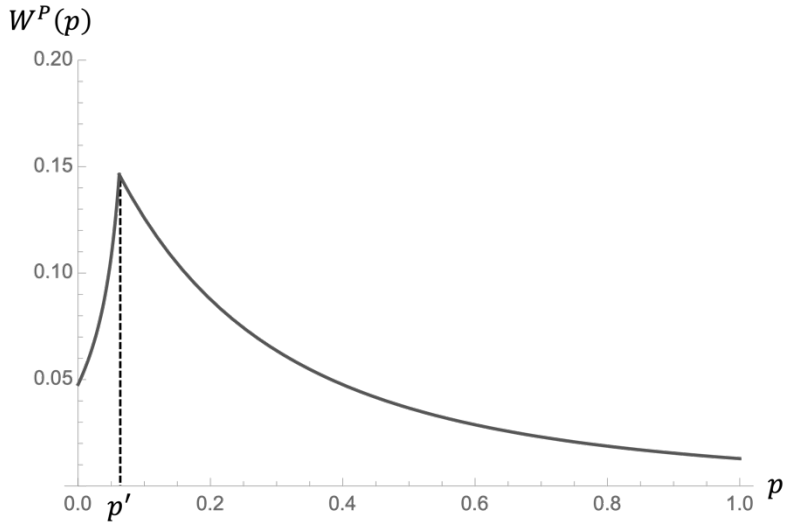

Figure 7 – Provider welfare.

We now aim to understand the optimal price chosen by the platform and how this impacts those users and providers who decide to join the platform. The platform's profits write as follows:

$$\Pi(p) = (p - C) \cdot S(p) \cdot E\left[\pi\left(r^*(p,\theta)\right)^2 | \theta \geq \hat{\theta}(p)\right]. \qquad \text{(Eq. 11)}$$

The platform incurs a marginal cost $C \in [0,1]$ for targeting and advertising a provider's content, with a margin $p - C$ for the mass of users reached to sponsor the providers' content when they join the platform, i.e., only those providers whose quality is higher than the threshold $\hat{\theta}(p)$. After plugging the equilibrium values, we obtain the following profit function

$$\Pi(p) = \frac{\pi}{3} \frac{p-C}{(1+p)(1+2p)^2}. \qquad \text{(Eq. 12)}$$

By taking the first order condition, we can then immediately obtain the platform's optimal price $p^* = \frac{\sqrt{5+14C+9C^2}+3C-1}{4}$. Interestingly, the platform chooses a price that lies in-between the level of prices that maximize users' and providers' welfare. Thus, the platform faces a trade-off between charging low prices and having many providers but few users onboard and charging high prices and having many users but few providers onboard. Such trade-off is resolved by balancing participation on the two sides, charging a price that elicits full participation by users

while attracting a considerable mass of providers. Interestingly, the choice of the platform seems to support the evidence that social media platforms (e.g., Facebook) tend to overcrowd their users with more ads than they would prefer, but it also shows that the platform is not completely inconsiderate in doing so.

Proposition 5. The monopolist platform charges a price $p^* = \frac{\sqrt{5+14C+9C^2}+3C-1}{4} > p'$.

Finally, we look at the total welfare in the case of a social planner maximizing the aggregate welfare of users, providers and the platform. This exercise enables us to understand how the price set by a social planner compares to the price that maximizes the platform's profits. The aggregate welfare is

$$W^T(p) = \Pi(p) + W^U(p) + W^P(p), \qquad \text{(Eq. 13)}$$

which is depicted in Figure 8. We can observe that users are better off in the presence of a social planner who would set a high price (namely $\tilde{p}$). This means that the socially optimal configuration is one with many users on the platform and only few providers who sponsor high-quality content. Although users would be better off with a configuration characterized by even less providers, the social planner takes also into account the welfare of providers and the profits of the platform. This result suggests the existence of a tendency for (targeted ads-driven) social media platforms to expose their users to more ads than they would otherwise prefer in order to achieve higher profits.

Proposition 6. The social planner adopts a price configuration that favors users compared to providers and the platform ($\tilde{p} > p^* > p'$)*, this configuration is characterized by many users and few providers who sponsor high-quality content.*

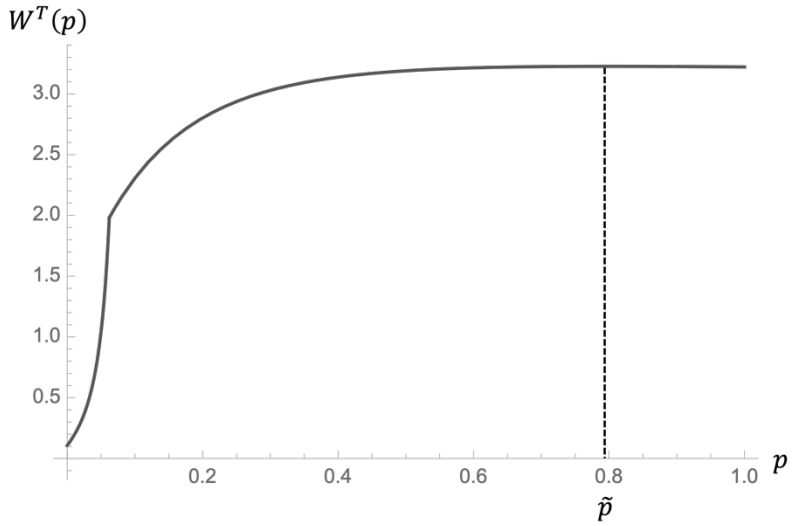

Figure 8 – Total welfare.

5. Discussion

Information platforms brokering user attention, like social media platforms, use pricing as a coordination mechanism to elicit participation of two sides with partially conflicting objectives. On the one hand, users join the platform to interact with other users, but they may benefit from ads when the content is related to their preferences and/or of high quality. Some authors argue that the excessive presence of low-quality or unrelated ads may hurt users (e.g., Choi et al., 2020; Monti, 2020), thus creating platform usage frictions and eventually damaging both providers and the platform. On the other hand, providers want to reach as many users as possible, both related and unrelated, with their content but are less willing to do so if they have to pay a high price. Users whose preferences are unrelated to the sponsored content do not click, and providers are less willing to reach them when the price increases due to lower chances of converting a view into a click. Throughout the paper, we characterize these (potentially) asymmetric indirect network effects to highlight the tension faced by social media platforms when setting the price of targeted advertising. Providers choose the reach based on their quality (θ) and the price (p), while users incur a (dis-)utility from the content that reaches them which is based on their cost of seeing ads (c) and the ad relatedness and quality.

By starting from the literature on targeted advertising and that on two-sided markets, we have built an analytical model with which to provide novel insights to inform the debate on excessive advertising on social media platforms. We show that an information platform charges a medium price, this means that it chooses to have many users onboard but not many providers. The platform attempts to solve the participation trade-off by laying off low-quality providers; this decision is based on the fact that a low price coupled with low-quality providers enables a broad reach for low-quality content, which hurts users with moderate to high nuisance cost. The optimal price (p^*) is higher than the price that maximizes provider welfare (p'); nonetheless, it is approximately half the price that maximizes user welfare, pointing towards the existence of a certain degree of user exploitation. Users may therefore be exposed to more low-quality ads on social media than they would otherwise prefer.

Our work has elements in support to both sides of the debate on user exploitation, which eventually enables a better understanding of the phenomenon. On one side, we provide novel evidence to scholars and regulators worried about user exploitation; an information platform tends to host many advertisers and show more ads than its users would prefer. This speaks to the ability of social media platforms like Facebook to profit from advertisers by leveraging their attention brokering position toward users. On the other side, we show that an information platform sets a price higher than that which would elicit full participation on the provider side, and it does so because admitting low-quality providers on the platform would drive off its user base. In doing so, we provide analytical support to those arguments advanced by scholars who do not think of information platforms as organizations that can disregard their demand and exploit it for profit-maximizing self-interests. We show that information platforms are neither reckless organizations aiming to extract all welfare from their user base, nor benevolent entities worried that their users will be annoyed by excessive ads. Rather, information platforms trade

some of their users' welfare for increased profits but are not willing to disrupt their user base with unbearable usage frictions.

6. Conclusions

In this paper we study how an information platform manages participation of users and providers on the two sides of its network by means of pricing its targeted advertising. We develop a model that encompasses providers' vertical and horizontal differentiation and users' heterogeneity in their preferences for content and cost of seeing (unrelated or low-quality) ads. First, we show that in equilibrium there will be unbalanced participation on the two sides of the platform as a consequence of (potentially negative) indirect network effects of ads. The platform uses pricing as a coordination mechanism to choose the configuration of users and providers (and the size of ads' externalities), whence for low prices there will be many providers and few users while for high prices there will be few providers and many users. Second, we demonstrate that content of high quality has a greater reach than content of low quality for any price charged by the platform due to higher conversion of views into clicks. Third, our core result is that we show the existence of some degrees of user exploitation by the platform. The platform's optimal price is indeed lower than the price that maximizes user welfare, meaning that the chosen configuration will present more providers who advertise content of lower average quality than what maximizes value for users. Nevertheless, we also demonstrate that the platform's optimal price is well above the price that would maximize value for providers, meaning that the platform does not disregard users' requests of less ads. The platform moves on a fine line between providers and users, pricing in-between the two to foster participation on both sides. Finally, we show that a social planner would opt for a price that lies between the platform's optimal price and the price that maximizes users' welfare, thus going in the direction of users and pointing to a further reduction in the number of providers on the platform.

We make several contributions to ongoing debates. First, we contribute to the emergent understanding of attention platforms (Wu, 2016, 2019) by shedding light on the user exploitation debate. We show that social media platforms tend to host more advertisers compared to what users or a social planner would choose, thus highlighting the existence of some degrees of exploitation. We also show that the platform cannot disregard the user base, thus setting a price that pushes off the platform low-quality providers in order to foster user participation. As often, the truth lies in the middle; user exploitation happens but the platform needs to curate its users and accommodate some of their requests to win their presence on the platform. Second, we contribute more broadly to platform theory by providing novel insights on information platforms (Cennamo, 2021) whence (potentially) negative indirect network effects may take place. In this regard, we propose is that asymmetric network effects will lead to asymmetric platform configurations, a finding that goes against the notions pushed forward in traditional (transaction) platform literature. Moreover, the pricing strategy seems once more to be a powerful coordination tool available to information platforms. Finally, we peripherally contribute to the advertising literature by developing a model in which advertising can bring a benefit to users and is not necessarily a net negative (in line with Johnson, 2013).

Our paper has some limitations that may also suggest avenues for future research. First, we consider a platform that has full information about providers' nature of content and users' preferences for content. Even if information platforms and in particular attention (and social media) platforms have incredibly large quantities of data on every single participant (Arrate Galán et al., 2019), we acknowledge that it might be relevant to understand what happens when the platform has incomplete knowledge of its participants. Second, we use linear pricing even though most social media platforms and similar employ probabilistic and auction-based pricing strategies. We believe that the model we presented captures the intuition behind the fact that many platforms charge for the number of reached users and employ probabilistic methods

mainly to pair providers and users. As an example, with one of its advertising packages, Facebook charges providers for a pre-determined reach, and then employs a probabilistic model to auction an ad slot to the best match.

We think that this paper sets the ground for interesting future ideas. First, scholars could look into what happens when privacy (or other events) limits the platform's knowledge of users or providers. This will make the platform's targeting less accurate with likely consequences on participation, profits, and social welfare. Second, many attention platforms (including social media platforms like YouTube) are introducing freemium models whence users can pay to be relieved from all or part of the advertisement (e.g., Rietveld, 2017; Shi et al., 2019). Freemium may not only be a source of additional revenue but may also foster participation on both sides by enabling the platform to recruit low-quality providers and shield users with a high cost of seeing ads by means of the ad-free subscription.

References

Alaimo, C., & Kallinikos, J., 2017. Computing the everyday: Social media as data platforms. *The Information Society, Vol. 33(4): 175-191*.

Alphabet Inc., 2020. Annual Report 2020.

Anderson, S.P., & de Palma, A., 2009. Information congestion. *RAND Journal of Economics, Vol. 40(4): 688-709*.

Anderson, S.P., & de Palma, A., 2012. Competition for attention in the Information (overload) Age. *RAND Journal of Economics, Vol. 43(1): 1-25*.

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(4): 1-34*.

Armstrong, M., 2006. Competition in two-sided markets. *RAND Journal of Economics, Vol. 37(3): 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(12): 2886-2907*.

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

Caillaud, B., & Jullien, B., 2003. Chicken & egg: Competition among intermediation service providers. *RAND Journal of Economics, Vol. 34(2): 309-328*.

Cennamo, C., 2021. Competing in digital markets: A platform-based perspective. *Academy of Management Perspectives, Vol. 35(2): 265-291*.

Cennamo, C., & Santaló, J., 2013. Platform competition: Strategic trade-offs in platform markets. *Strategic Management Journal, Vol. 34(11): 1331-1350*.

Chatterjee, P., Hoffman, D.L., & Novak, P.T., 2003. Modeling the clickstream: Implications for web- based advertising efforts. *Marketing Science, Vol. 22(4): 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(2): 556-575*.

Danaher, P.J., Lee, J., & Kerbache, L., 2010. Optimal internet media selection. *Marketing Science, Vol. 29(2): 336-347*.

de Cornière, A., & Sarvary, M., 2022. Social media and news: Content bundling and news quality. *Management Science, Articles in advance*.

Despotakis, S., Ravi, R., & Srinivasan, K., 2021. The beneficial effects of ad blockers. *Management Science, Vol. 67(4): 2096-2125*.

Despotakis, S., & Yu, J., 2022. Multidimensional targeting and consumer response. *Management Science, Articles in Advance*.

Eisenmann, T., Parker, G., & Van Alstyne, M.W., 2006. Strategies for two-sided markets. *Harvard Business Review, Vol. 84(10): 92-101*.

Evans, D.S., 2003. Some empirical aspects of multi-sided platform industries. *Review of Network Economics, Vol. 2(3): 191-209*.

Evans, D.S., 2013. Attention rivalry among online platforms. *Journal of Competition Law & Economics, Vol. 9(2): 313-357*.

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

Evans, D.S., 2019. Attention platforms, the value of content, and public policy. *Review of Industrial Organization, Vol. 54: 775-792*.

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(10): 1605-1622*.

Gritckevich, A., Katona, Z., & Sarvary, M., 2021. Ad blocking. *Management Science, Vol. 68(6): 4703-4724*.

Hagiu, A., 2006. Pricing and commitment by two-sided platforms. *RAND Journal of Economics, Vol. 37(3): 720-737*.

Hann, I., Hui, K., Lee, S.T., & Png, I.P.L., 2008. Consumer privacy and marketing avoidance: A static model. *Management Science, Vol. 54(6): 10941103*.

Iyer, G., & Katona, Z., 2016. Competing for attention in social communication markets. *Management Science, Vol. 62(8): 2304-2320*.

Iyer, G., Soberman, D., & Villas-Boas, J.M., 2005. The targeting of advertising. *Marketing Science, Vol. 24(3): 461-476*.

Jacobides, M.G., Cennamo, C., & Gawer, A., 2018. Towards a theory of ecosystems. *Strategic Management Journal, Vol. 39(8): 2255-2276*.

Johnson, J.P., 2013. Targeted advertising and advertising avoidance. *RAND Journal of Economics Vol. 44(1): 128-144*.

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

Jullien, B., & Sand-Zantman, W., 2021. The economics of platforms: A theory guide for competition policy. *Information Economics and Policy, Vol. 54: 1-19*.

Katz, M.L., & Shapiro, C., 1994. Systems competition and network effects. *Journal of Economic Perspectives, Vol. 8(2): 93-115*.

Lee, D., Hosanagar, K., & Nair, H.S., 2018. Advertising content and consumer engagement on social media: Evidence from Facebook. *Management Science, Vol. 64(11): 5105-5131*.

Lin, S., 2020. Two-sided price discrimination by media platforms. *Marketing Science, Vol. 39(2): 317-338*.

Liu, X., Li, X., & Sarkar, S., 2022. Cost-restricted feature selection for data acquisition. *Management Science, Articles in advance*.

Long, F., Jerath, K., & Sarvary, M., 2022. Designing an online retail marketplace: Leveraging information from sponsored advertising. *Marketing Science, Vol. 41(1): 115-138*.

Marotta, V., Wu, Y., Zhang, K., & Acquisti, A., 2022. The welfare impact of targeted advertising technologies. *Information Systems Research, Vol. 33(1): 131-151*.

McIntyre, D.P., & Srinivasan, A., 2017. Networks, platforms, and strategy: Emerging views and next steps. *Strategic Management Journal, Vol. 38(1): 141-160*.

Monti, G., 2020. Attention intermediaries: Regulatory options and their institutional implications. *TILEC discussion paper*.

Panico, C., & Cennamo, C., 2022. User preferences and strategic interactions in platform ecosystems. *Strategic Management Journal, Vol. 43(3): 507-529*.

Parker, G.G., & Van Alstyne, M., 2005. Two-sided network effects: A theory of information product design. *Management Science, Vol. 51(10): 1494-1504*.

Peitz, M., 2020. Economic policy for digital attention intermediaries. *ZEW discussion paper*.

Prat, A., & Valletti, T., 2022. Attention oligopoly. *American Economic Journal: Microeconomics, Vol. 14(3): 530-557*.

Rietveld, J., 2017. Creating and capturing value from freemium business models: A demand-side perspective. *Strategic Entrepreneurship Journal, Vol. 12: 171-193*.

Rochet, J.C., & Tirole, J., 2003. Platform competition in two-sided markets. *Journal of the European Economic Association, Vol. 1(4): 990-1029*.

Rochet, J.C., & Tirole, J., 2006. Two-sided markets: a progress report. *RAND Journal of Economics, Vol. 37(3): 645-667*.

Rossi, F., & Rubera, G., 2021. Measuring competition for attention in social media: National women's soccer league players on twitter. *Marketing Science, Vol. 40(6): 1147-1168*.

Shi, Z., Zhang, K., & Srinivasan, K., 2019. Freemium as an optimal strategy for market dominant firms. *Marketing Science, Vol. 38: 150-169*.

Srinivasan, D., 2019. The antitrust case against Facebook: A monopolist's journey towards pervasive surveillance in spite of consumers' preference for privacy. *Berkeley Business Law Journal, Vol. 16(1): 39-101*.

Statista, 2021. Number of active advertisers on Facebook from 1st quarter 2016 to 3rd quarter 2020 (consulted on October 5, 2022).

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

Toubia, O., & Stephen, A.T., 2013. Intrinsic vs. image-related utility in social media: Why do people contribute content to twitter? *Marketing Science, Vol. 32(3): 368-392*.

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(1): 102-125*.

Trusov, M., Ma, L., & Jamal, Z., 2016. Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting. *Marketing Science, Vol. 35(3): 405-426*.

Van Zandt, T., 2004. Information overload in a network of targeted communication. *RAND Journal of Economics, Vol. 35(3): 542-560*.

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(3): 771- 806*.

Appendix

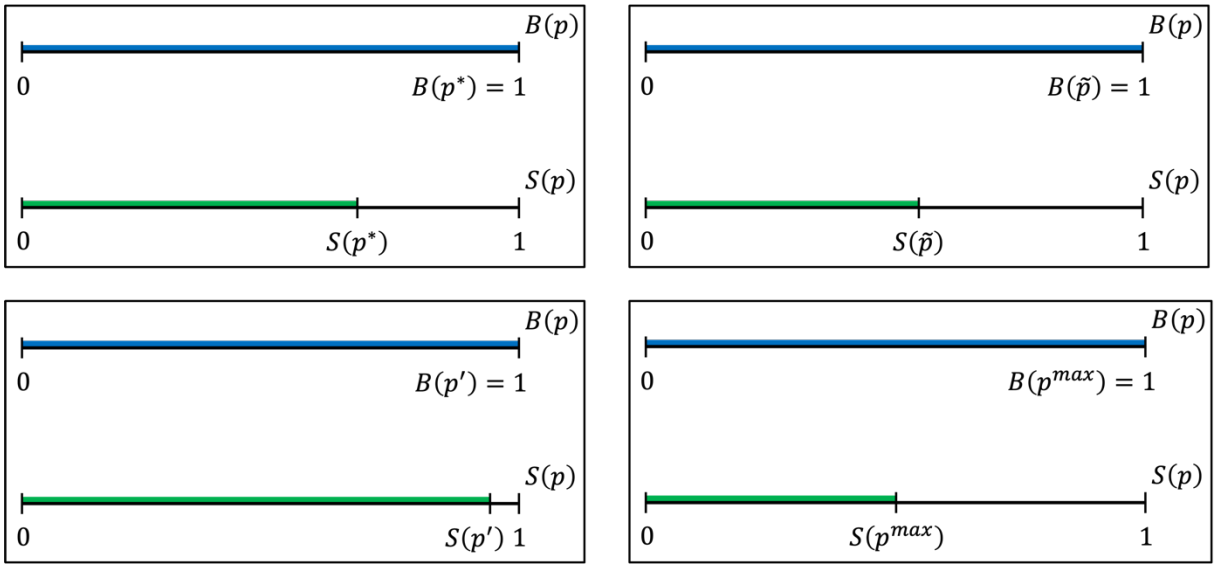

Figure 9 – Platform configurations ($B(p)$ is the mass of users and $S(p)$ is the mass of providers) for the different prices. The price configurations follow: p^* is the price that maximizes the platform profits; p' is the price that maximizes provider welfare; $\tilde{p}$ is the price that maximizes the social welfare; p^{max} is the price that maximizes user welfare.